

Louvain School of Management

Generative Artificial Intelligence in Consulting Organizations: Implications for Human Expertise, Trust and Transparency

Author(s): Hamuli Ymane
Supervisor(s): Decaux Loïc
Academic year 2025-2026
Dissertation for the master of Business Engineering with major in
Business Analytics
Daytime Schedule

Declaration Regarding AI Tool Usage in Master's Thesis

Although I used generative AI platforms like Gemini, ChatGPT-4, and Claude 3.5 Sonnet to assist with this work, they served purely as a secondary support system. Their application remained strictly within the authorized institutional guidelines at all times. Deepl was also used to translate the French interviews in English.

I confirm the following statements:

Statements on responsible use of generative AI	Confirmation:
	Yes
The content of my/our master's thesis is my/our original input, even if I/we used generative AI to help prepare it.	☒
I/we master the theories, tools and methods that I/we use for the thesis.	☒
I/we verified that the output presented in the thesis is accurate and valid, and that it contains no hallucination, no false reference, etc.	☒
I/we acknowledge that I/we master the final output and that I am/we are able to explain it and answer questions about it.	☒

Author last name, first name	Date	Signature
1 Ymane Hamuli	17th, May 2026	

Acknowledgements

First and foremost, I would like to thank my supervisor. M. Loic Decaux for his guidance, valuable advice, and support throughout this project.

I am also very grateful to all the interview participants who generously gave up their time to share their insights and experiences with me. This thesis simply could not have been done without them.

A special thanks goes to Aymeline Brabant, whom I met at an LSM event. Her help in opening up her network was incredibly valuable and made a real difference in helping me find my participants.

Finally, thank you to my family and friends for always being there and keeping me motivated until the very end.

Abstract :

The rapid rise of generative artificial intelligence (GenAI) is reshaping knowledge-intensive industries. It fundamentally challenges the traditional value proposition of management consulting. While these tools offer unprecedented efficiency, they disrupt the historical asymmetry of information between consultants and clients. This disruption introduces critical tensions around human expertise, trust, and transparency. This research investigates how the progressive integration of GenAI influences these three foundational pillars within the consulting-client relationship.

This qualitative study is based on sixteen semi-structured interviews with consulting professionals across Belgium. Using thematic analysis across specializations like ERP, risk advisory, and AI implementation, it captures the changing reality on the ground. The findings reveal a structural transformation. As GenAI automates raw knowledge production, traditional contributory expertise is steadily losing ground to contextual, interactional expertise. Furthermore, algorithmic imperfections actively weaken traditional ability-based trust. To counter this, consultants increasingly rely on benevolence and integrity, demonstrating trustworthiness through strict data governance. Regarding transparency, explicit disclosure strategies remain rare. This creates a fragmented reality marked by a "double opacity." However, when strong interpersonal trust is established, or when AI use is formalized upfront in contracts, clients focus entirely on final deliverables. This outcome-driven approach makes the algorithmic black box acceptable, significantly reducing the pressure for full technical transparency.

Ultimately, this study concludes that GenAI is not displacing the consultant. Instead, it reshapes the profession from within and redefines the professional's true added value. As knowledge becomes democratised, the economic value of consulting shifts from intellectual production to accountability. The core product being sold today is the professional signature. This represents the capacity to filter algorithmic noise, exercise contextual judgment, and assume responsibility for strategic decisions. In conclusion, clients no longer pay consultants merely to generate information. They pay them for risk mitigation, human perspective, and the ultimate accountability needed to navigate uncertainty.

UNIVERSITÉ CATHOLIQUE DE LOUVAIN
Louvain School of Management

Place des Doyens, 1 bte L2.01.01, 1348 Louvain-la-Neuve
Boulevard Emile Devreux 6, 6000 Charleroi, Belgique
Chaussée de Binche 151, 7000 Mons, Belgique
www.uclouvain.be/lsm

Table of content

1	Introduction	1
2	Literature Review.....	3
2.1	Management Consulting as a Knowledge-Intensive Industry.....	3
2.1.1	<i>Defining Management Consulting</i>	3
2.1.2	<i>Management Consulting : A Knowledge Industry</i>	4
2.1.3	<i>The Evolution of the Consulting Industry.....</i>	5
2.2	The Three Pillars of the Client-Consultant Relationship	6
2.2.1	<i>Human Expertise in Client Interactions.....</i>	6
2.2.2	<i>Trust in Client-Consultant Relationships.....</i>	6
2.2.3	<i>Transparency in Consulting Engagements.....</i>	7
2.3	Generative AI as Disruptive Knowledge infrastructure	9
2.3.1	<i>Definition and Conceptual Positioning of Generative AI</i>	9
2.3.2	<i>Implications for Human Expertise</i>	9
2.3.3	<i>Implications for Trust.....</i>	11
2.3.4	<i>Implication for Transparency.....</i>	12
2.4	Research Gap.....	13
3	Methodology	14
3.1	Introduction	14
3.1.1	<i>Research propositions.....</i>	14
3.1.2	<i>Sample Composition.....</i>	15
3.1.3	<i>Data Collection</i>	16
3.1.4	<i>Data Analysis</i>	16
3.1.5	<i>Limits of Methodology.....</i>	17
4	Analysis and Findings	17
4.1	Introduction	17
4.2	What Clients Valued Before GenAI.....	18
4.3	P1: How Does GenAI Transform Consultant Expertise?.....	19
4.3.1	<i>The reduction of Manual Production</i>	19
4.3.2	<i>Relational Expertise as The Primary Differentiator</i>	20
4.3.3	<i>The Impact on Juniors.....</i>	21
4.3.4	<i>The Impact on Jurisdiction.....</i>	22
4.4	P2: How Does GenAI Reconfigure Client-Consultant Trust?	23
4.4.1	<i>Introduction.....</i>	23
4.4.2	<i>The Commodification of Technical Ability.....</i>	23

4.4.3	<i>The Over-Investment in Benevolence</i>	23
4.4.4	<i>The Over-Investment in Integrity</i>	24
4.4.5	<i>Cybersecurity as a Concrete Expression of Integrity</i>	25
4.5	P3: How Do Client Expectations Change?	25
4.5.1	<i>Reliance : Clients Delegate Speed to GenAI</i>	25
4.5.2	<i>Hallucinations: The Mechanism That Fragments Trust</i>	26
4.6	P4: The fewest implications for transparency	27
4.6.1	<i>Transparency on Mechanical Tasks</i>	27
4.6.2	<i>Strategic Opacity for High Value Tasks</i>	28
5	Discussion	29
5.1	<i>Redefining Human Expertise</i>	29
5.2	<i>Adapting the ABI Framework</i>	31
5.3	<i>The Real Drivers of Transparency</i>	33
5.4	<i>What Remains of the Consultant's Role in the GenAI Era</i>	35
5.5	<i>The limitations of the Research</i>	38
5.6	<i>Recommendations for Future Research</i>	38
6	Conclusion	39
7	Bibliography	41
8	Appendices	45
	Appendix 1 : Definitions	45
	Appendix 2 : Participants Data.....	47
	Appendix 3 : Interview Guide	49
	Appendix 4 : Codebook.....	50
	<i>Research proposition 1 : Shift in Expertise</i>	50
	<i>Research proposition 2 : ABI model Compensation</i>	51
	<i>Research proposition 3 : Reliance in GenAI , Trust in Consultant</i>	52
	<i>Research proposition 4 : Selective Transparency</i>	53
	Appendix 5 : Interviews	54
	<i>Interview ERP01</i>	54
	<i>Interview ERP02</i>	65
	<i>Interview ERP03</i>	73
	<i>Interview ERP04</i>	82
	<i>Interview ERP05</i>	92
	<i>Interview ERP06</i>	103
	<i>Interview ERP07</i>	111
	<i>Interview RISK01</i>	119

<i>Interview RISK02</i>	131
<i>Interview RISK03</i>	138
<i>Interview TECH01</i>	146
<i>Interview TECH02</i>	156
<i>Interview TECH03</i>	166
<i>Interview TECH04</i>	174
<i>Interview TECH05</i>	182
<i>Interview FUNC01</i>	191

1 Introduction

In most sectors, generative AI is mainly associated with efficiency. A radiologist can read more scans in less time, and a financial analyst can process massive amounts of data almost instantly. In these contexts, the value of GenAI is primarily linked to productivity.

Consulting, however, is different. Consultants do not simply sell faster analysis, but something far more intangible: expertise, trust, and transparency. These factors are closely related to human judgment and influence the value of the consultant-client relationship. For decades, this value was based on a significant asymmetry of knowledge. Clients couldn't simply access or replicate the frameworks, benchmarks, and analytical capabilities that consultants had. This disparity justified both their authority and the faith that clients had in them.

Generative AI upsets this equilibrium in a way that prior technologies did not. Today, AI systems can create strategic frameworks, market evaluations, and document syntheses in seconds. As a result, the debate is no longer only about productivity, but also about what makes consulting work distinctively human. This dilemma exists in three dimensions, which generative AI is changing.

The first dimension is expertise. Traditionally, expertise was closely associated with possessing specialised knowledge. However, as access to knowledge becomes increasingly democratised through AI, expertise can no longer simply mean “knowing more.” It increasingly raises questions about judgment, contextualisation, and the ability to critically assess AI-generated outputs. A second dimension is trust. Consulting relationships have always involved a degree of uncertainty because clients cannot fully assess the quality of recommendations before implementation. Traditionally, firms addressed this uncertainty through reputation, recognised methodologies, professional credentials, and brand prestige. The third dimension is transparency. The growing use of AI in consulting raises unresolved questions. Should consultants disclose the use of AI? Who becomes responsible when AI-generated outputs contain errors? And how transparent should AI-assisted processes actually be? These issues increasingly touch broader questions of ethics, governance, and professional legitimacy.

The Deloitte Australia case illustrates these tensions particularly well. In 2025, an AI-assisted audit produced significant factual errors, damaging the credibility of the engagement (Fortune, 2025). The case simultaneously raised questions about expertise, trust, and transparency. Clients could no longer clearly distinguish between human judgment and machine-generated

outputs, while responsibility became increasingly ambiguous. Yet despite these risks, Deloitte expanded its partnership with Anthropic to deploy Claude at scale (Fortune, 2025).

These tensions lie at the heart of contemporary consulting practice. If generative AI reshapes expertise, trust, and transparency simultaneously, then it potentially reshapes the profession itself. Beyond productivity gains, this raises broader questions about professional legitimacy, human judgment, and the future value of consultants in a knowledge-based economy.

Consequently, this research investigates the following question: *From the consultant's perspective, how does the progressive use of generative AI in consulting organisations influence the preservation of human expertise, trust, and transparency in the client-consultant relationship?*

To answer this question, the study draws on three key theoretical frameworks. For expertise, it distinguishes between contributory expertise, the ability to perform analytical work, and interactional expertise, which involves interpreting and communicating knowledge across contexts (Collins & Evans, 2007). This is supported by the SECI model of knowledge creation, explaining how tacit and explicit knowledge circulate within organisations (Nonaka & Takeuchi, 1995), and Abbott's (1988) concept of professional jurisdiction, which addresses how professions maintain authority. For trust, the analysis uses Mayer et al.'s (1995) ABI model, defining trust through ability, benevolence, and integrity, alongside the idea of a "leap of faith" that acknowledges inherent uncertainty (Nikolova et al., 2014). For transparency, the study applies Schnackenberg and Tomlinson's (2016) framework, which breaks transparency into disclosure, clarity, and accuracy. This is combined with Pasquale's (2015) concept of the "black box" and recent insights on the transparency paradox, showing that more transparency does not always increase trust (Schilke & Reimann, 2025). Finally, Strobel et al.'s (2024) taxonomy helps classify AI use based on whether systems generate, synthesise, or reimagine content.

Methodologically, the research is based on semi-structured interviews with 16 consulting professionals working in Belgium. The sample includes consultants at different seniority levels and across various specialisations such as ERP, risk, data analytics, and AI. This diversity allows for a more nuanced understanding of how GenAI is used in practice across different roles and contexts. The data were analysed using thematic analysis (Braun & Clarke, 2006). This technique compared the codes driven by the theory and others codes more inductive enabled to capture unexpected patterns emerging from the field.

The thesis is structured as follows. The first chapter reviews the literature on expertise, trust, and transparency in consulting, and explores how GenAI is reshaping each of these dimensions. The second chapter presents the methodology. The third chapter reports the empirical findings structured around the four research propositions. The fourth chapter discusses these findings in relation to the literature and introduces emerging concepts such as prompt engineering and double opacity. The final chapter reflects on the limitations of the study and proposes directions for future research.

2 Literature Review

The following section reviews the existing academic literature relevant to this research. The literature for this study was found by searching Google Scholar and exploring the academic databases of both UCLouvain and Università Bocconi. This review builds its analytic framework around three dimensions that are central to management consulting: expertise, trust, and transparency. It then explores how generative AI is reshaping each of these pillars. These three dimensions are rarely examined together in the existing literature. Bringing them into dialogue allows for a more grounded understanding of what actually changes when AI enters the consulting relationship. Additionally, a table summarizing the key definitions used in this section can be found in Appendix 1.

2.1 Management Consulting as a Knowledge-Intensive Industry

2.1.1 *Defining Management Consulting*

Defining management consulting remains a complex and debated task in academic literature (Nissen, 2019). Nevertheless, Appelbaum and Steed (2005, p.68) offer a widely used working definition: “Management consulting is an advisory service provided by trained persons who help organizations identify and analyse management problems and assist in the implementation of solutions”. Despite this definition, consulting remains an ambiguous profession with what has often been described as a chameleon-like nature. Consultants frequently adopt the language, dress, and cultural norms of their clients in order to gain acceptance and legitimacy (Cerruti et al., 2018). There are several broad domains of consulting practice including human resources, finance, marketing and IT, and what counts as legitimate expertise varies considerably across them (Kubr, 2002).

Building on prior research, consulting can be decomposed into three key dimensions. First, consulting involves the construction and framing of organizational problems. It is a

collaborative process: the actual problem only comes to light through continuous dialogue and co-production with the client organization (Schein, 1990, as cited in Appelbaum & Steed, 2005; Cerruti et al., 2018). Second, because of this shared process, even the best recommendations will fail if the relationship between both parties is weak. Since the outputs of management consulting are largely intangible, clients often have to rely on what Nikolova et al. (2014) describes as a leap of faith when evaluating consultant's credibility. Third, consultants can be understood as knowledge brokers. They transfer and recombine knowledge across organizations and industries to foster innovation (Hargadon, 1998). As a result, value often comes from recombining what already exists elsewhere which make it usable in a context where it was not available before.

2.1.2 Management Consulting : A Knowledge Industry

Understanding what makes consulting distinctive as a profession, however, requires going beyond its definition. The nature of the knowledge it produces and transfers is what truly sets it apart from other advisory services. One of the key aspects that gives consultants their authority is the ability to structure, mobilise and transfer knowledge in a usable way (McKenna, 2006).

A classic distinction in the knowledge management literature helps clarify this. On the one hand, explicit knowledge is knowledge that "can be expressed in words and numbers and easily shared" (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998, p.673). On the other hand, tacit knowledge is harder to capture. This type of knowledge is developed through experience, intuition and practice. It is often difficult to express it explicitly or communicate it to others because it is closely tied to personal contexts and actions (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998, see Appendix 1). This distinction becomes very tangible in consulting practice. Consultants can be understood as improvising bricoleurs who rarely rely on standardized tools, choosing instead to combine and adapt whatever resources are currently at hand (Lévi-Strauss, 1966 as cited in Visscher et al., 2017). Much of this happens through trial and error, feedback from the field, and iterative refinement, rather than through abstract reasoning alone. Because of this dynamic, consulting knowledge is rarely a static asset; it is something that is constantly renegotiated every time a consultant steps into a new client's organization.

A useful way to capture this dynamic is the SECI model developed by Nonaka and Takeuchi (1995). The model describes knowledge creation as a continuous movement between four

processes. It starts with socialisation, where tacit is shared through observation and shared experience. Then comes externalisation, when these intuitive insights are turned into more explicit forms such as models or metaphors. Combination follows, where different pieces of explicit knowledge are organised into structured frameworks. Finally, internalisation takes place when individuals learn through practice, turning explicit knowledge back into tacit understanding. Knowledge moves continuously between tacit and explicit forms, and it is in this movement that learning really happens.

2.1.3 The Evolution of the Consulting Industry

Knowledge has been at the core of the consulting profession for years. What has changed, however, is the technological context in which this knowledge is produced and delivered. The consulting profession has not evolved in a linear way. It has moved through successive waves, each shaped by a major shift.

The first wave, in the late nineteenth and early twentieth centuries was centred on workplace efficiency and the rise of scientific management (Kipping & Clark, 2012). At that stage, consulting was largely about making work more rational, more measurable, and more productive. Then came a second wave, between the 1930s and 1970s. The focus shifted. It was no longer only about making work more efficient on the ground, but also about how companies are structured and managed at a higher level. This period also saw the emergence of firms that would later define the industry, such as McKinsey & Company and Boston Consulting Group (Curnow et al., 2003). A third wave followed in the late 1970s, which was driven by the rise of information technology. Large accounting firms such as Deloitte and PwC expanded into IT consulting. They also started to integrate systems and digital infrastructure into their service offerings (Curnow et al., 2003). As a result, consulting gradually became more technical and more data-driven.

Today, we are in the fourth wave of consulting. Historically, consulting firms have tended to adopt new technologies cautiously within their own operating models (Crisan & Marincean, 2023). Previous technologies mainly expanded what consultants could offer. But this current wave is characterized by the usage of more advanced tools, and more recently, generative AI. We will see that GenAI goes further : it changes how knowledge is created , how it is shared with clients, and even how trust is established.

2.2 The Three Pillars of the Client-Consultant Relationship

2.2.1 *Human Expertise in Client Interactions*

Expertise is the fundamental element of consulting profession. To understand how professionals navigate their fields, it helps to look at Collins and Evans' (2007) distinction between two specific types of expertise. The first, contributory expertise, is deeply practical. It involves the hands-on ability to actually execute tasks within a specialized area. This kind of competence comes from deep immersion and experiential learning, relying heavily on tacit knowledge (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) that is notoriously hard to codify. On the flip side, interactional expertise is entirely different. It involves mastering the specific language and concepts of a field well enough to discuss, debate, and interpret them convincingly, even if the person cannot perform the core technical tasks themselves. For consultants, success actually depends on blending both. They need enough contributory expertise to design solutions that actually work. But interactional expertise is what allows them to connect with clients, handle complex organisational challenges, and make their analysis understandable. Additionally, as emphasized by Clark (1995), consulting goes beyond applied knowledge: it is something that is performed and demonstrated in interaction with clients.

As argued by Abbott (1988), professional expertise acquires meaning only when it is socially recognized and institutionally protected through formal mechanisms and stakeholder acknowledgement. Abbott introduces the concept of jurisdiction which describes the exclusive territory a profession claims over a specific set of problems. When a profession holds jurisdiction, society and institutions recognize it as the sole authority qualified to diagnose issues and offer solutions in that space. Management consulting positions itself as the profession best equipped to handle ambiguity. It turns unclear situations into structured problems that can be analysed and addressed by someone that is external to the client's company. This often involves strategy, digital transformation or change management. Unlike medicine or law, the consulting profession does not require official certification. In reality, everything depends on the client and whether they believe the service is worth paying for. This is precisely why the profession is so vulnerable today, especially when technological tools are able to deliver nearly identical results.

2.2.2 *Trust in Client-Consultant Relationships*

Beyond expertise, trust is another key pillar of the client-consultant relationship. Without it, even strong technical recommendations may fail to convince or to be implemented. Trust is a

psychological state in which a person is willing to be vulnerable because they expect the other party to act with good intentions or behave positively (Rousseau et al. ,1998; as cited in Nikolova et al., 2014).

In their ABI model, Mayer et al. (1995) show that trust is actually based on three key elements. First, competence, which simply refers to the consultant's level of expertise. Second, benevolence. What the client should feel is the impression that the consultant is working in the client's best interest rather than focusing on revenue. The final element is integrity. This refers to the consultant's ability to act with honesty and take responsibility when the situation calls for it. In practice, however, clients rarely have enough information to evaluate these dimensions with certainty. They often make decisions under conditions of incomplete knowledge, ambiguity, and information asymmetry. Risk is therefore unavoidable. Therefore, trust becomes a decision made under uncertainty. This trust can be described as a leap of faith: "A confident suspension of uncertainty within the context of a newly forming relationship" (Nikolova et al.,2014, p. 234). In other words, clients choose to rely on consultants even when the available evidence remains incomplete. As Nikolova et al. (2014) similarly argue, credentials and reputation may provide good reasons to trust, but they do not remove uncertainty. Trust still has to be built through interaction.

Historically, we can see that consulting has never enjoyed automatic trust anyway. Mckenna (2006) notes that since the early twentieth century, consultants have often been met with scepticism, sometimes seen less as neutral experts and more as corporate enablers. In response, consulting firms have long relied on professionalism and impression management to signal credibility (Mckenna, 2006). Consequently, aesthetic presentation and impression management are integral to producing legitimacy. As a result, structured presentations, polished frameworks, and persuasive storytelling becomes more than communication tools: they are part of how legitimacy is produced. This links back to Section 2.1.1, where consultants' adaptation to client norms was presented as central for their legitimacy (Cerruti et al., 2018).

2.2.3 Transparency in Consulting Engagements

As highlighted by the ISO 20700 standard, clients are increasingly demanding greater transparency in consulting engagements. This expectation is not new, but it is becoming harder to ignore. Transparency can be defined as "the perceived quality of information that is intentionally shared by a sender" (Schnackenberg & Tomlinson ,2016, p. 1788). In other words, organisations actively shape how transparent they appear by deciding what to communicate,

how to communicate it, and when (Dholakia,2023). Schnackenberg and Tomlinson (2016) conceptualise organisational transparency as a multidimensional construct built around three main components. The first is disclosure, which refers to the act of openly sharing relevant information to the stakeholders exactly when it is needed (Schnackenberg & Tomlinson, 2016). The second dimension is clarity, meaning “the perceived level of lucidity and comprehensibility of information received from a sender” (Schnackenberg & Tomlinson, 2016, p. 13). In other words, simply providing information is not enough. Even when disclosure is high, communication can still fail if the message is filled with jargon, overly technical, or expressed in vague terms. The third dimension is accuracy, meaning how correct the information is between the sender and the receiver (Schnackenberg & Tomlinson , 2016). In short, whether stakeholders perceive a piece of information as trustworthy enough to act upon.

These challenges connect closely to the idea of the black box: “a system whose workings are mysterious; we can observe its inputs and its outputs, but we cannot tell how one becomes the other” (Pasquale, 2015, p. 3). Although this term technically refers to the decision-making process of algorithms, we can also try to apply it to the traditional consulting model. It is fair to say that there is a certain information asymmetry between the solutions the consultant proposes and what the client knows about those solutions and how they were designed. Sometimes, it creates uncertainty and often scepticism (Appelbaum & Steed, 2005). This scepticism is visible in empirical work. Appelbaum and Steed’s case study found that only 36% of respondents agreed that consultants effectively adapted their recommendations to the specific needs of the organisation. Ultimately, these findings show that clients are quite suspicious of recycled solutions.

Dholakia (2023) captures this perfectly by framing transparency as a matter of perception. When organizations make their processes visible, they manage how stakeholders interpret what they are shown. In consulting, this means that transparency is often as much about signalling and reassurance as it is about information itself.

These three dimensions do not exist in isolation from the technological context in which consulting operates. As generative AI becomes embedded in consulting workflows, each of them is being reshaped: sometimes in ways the existing literature anticipated and sometimes in ways it did not.

2.3 Generative AI as Disruptive Knowledge infrastructure

2.3.1 *Definition and Conceptual Positioning of Generative AI*

The rapid diffusion of digital technologies has made the concept of artificial intelligence increasingly difficult to pin down. “AI” is now used to describe a wide range of tools, capabilities, and systems, often with little precision. One of the most recent and visible stages of this evolution is generative artificial intelligence. Generative artificial intelligence refers to a set of models and algorithms capable of creating original content, whether it be text, images, or strategic decisions. (He et al. ,2025). For instance, popular AI models like OpenAI’s ChatGPT (Belcic & Stryker, 2026) and Google’s Gemini (Koster, 2025) are highly versatile and can generate content for almost every purpose. In contrast, to protect confidential client data, major consulting firms have deployed secure and internal systems with GenAI capabilities such as Deloitte’s Ai called Solaria (Deloitte, 2024) or McKinsey’s Lilli (McKinsey & Company, 2023).

Large Language Models, in particular, are designed to generate natural language by predicting the most probable sequence of words in a given context (Larson, 2023). Their outputs can appear coherent, persuasive, and even insightful. Yet, as Bender and Gebru (2021) emphasise, these models do not “understand” language in the human sense. They generate text through statistical prediction rather than genuine reasoning or contextual comprehension.

While Strobel et al. (2024) propose a very broad classification of AI functions, we will focus primarily on two of these roles, as they are truly the most relevant to the subject of this thesis. On one hand, we have the “data processors.” Their role is to do the heavy lifting: they crunch through massive amounts of data to extract truly useful insights, such as benchmarks. On the other hand, the “re-inventors” take over. They take this raw material and shape it into polished presentations, ready for executives. Looking at things from this perspective is very useful: it helps us understand exactly which tasks are likely to be automated and which still require genuine human judgment.

2.3.2 *Implications for Human Expertise*

Empirical evidence increasingly suggests that generative IA can dramatically expand the analytical capacity of non-specialists. Indeed, a controlled study conducted at Boston Consulting Group illustrates this clearly. In the experiment, generalist consultants equipped with GPT-4 were able to perform advanced data-science tasks such as writing code, building

predictive models, and interpreting statistical outputs, despite having no prior expertise in these areas (Dell'Acqua et al., 2023). This supports what Benbya et al. (2024) describe as a democratization of knowledge application: generative AI reduces traditional skill gaps between consulting profiles and increases the productivity of workers who are not formally specialised.

Despite these capabilities, the exact nature of the expertise provided by AI remains ambiguous. From the perspective of Collins and Evans (2007), generative AI can mimic elements of interactional expertise. It produces outputs that sound competent and appear contextually appropriate. Yet it does so without possessing the tacit understanding (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) that characterises contributory expertise, the form of expertise grounded in lived practice. Crucially, as noted earlier, this process relies purely on statistical prediction rather than comprehension (Bender and Gebru, 2021; see Section 2.3.1).

Consequently, while AI can boost performance, it introduces new vulnerabilities. Benbya et al. (2024) warn, for instance, that generative AI may produce plausible but false content, also known as hallucinations (Larson, 2023). If we rely on it too heavily, we simply risk losing the perspective and nuance that are characteristic of human judgment. The BCG experiment reflects this duality. Consultants equipped with AI perform about 40% better on purely creative tasks. However, they sometimes struggle with complex problems that require a thorough investigation to identify the root cause (Dell'Acqua et al., 2023). In other words, AI can amplify productivity, but it may also weaken performance precisely where complex analytical thinking is required.

If we look closely at how these systems actually function, they operate almost entirely within the combination phase. This is where explicit knowledge (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) is gathered, reorganized, and scaled up. They are highly effective at recombining information. But this is also where their capability ends. Because they lack the capacity for shared, lived experiences, they cannot truly participate in socialization. Similarly, they cannot engage in internalization simply because they cannot learn by doing within the complex, relational reality of an organization. It reflects the fact that AI systems do not possess embodied experience in the first place (Nonaka & Takeuchi, 1995).

Such a shift carries deep implications for the consultant's role and professional identity. As Mohan (2024) observed, the early stages of consulting still rely heavily on contextual judgment and human interaction. Consequently, AI tools cannot replicate crucial tasks like understanding implicit client needs, framing ambiguous problems or building trust. Yet at the delivery stage, the situation is shifting rapidly. Generative AI can now produce polished slide decks,

summaries, and content ready-made for presentations within seconds (Korzynski, 2023). As a result, the material artefacts that consulting has historically relied on to signal expertise are becoming less exclusive (Clark, 1995; McKenna, 2006).

Tronnier et al. (2025) and Mohan (2024) both point to the same shift: consultants are spending less time producing analysis and more time checking, reframing, and taking responsibility for outputs, they did not generate themselves. This might look like a minor shift on the surface, but it has massive implications for why professions exist in the first place. Susskind and Susskind (2015) remind us that technical superiority is not everything; humans remain essential because accountability is strictly human. An AI does not have a reputation to ruin or legal liability to shoulder. Taking this logic a step further, Abbott (1988) offers an even sharper interpretation. If analytical production becomes partly automated, then the consulting profession's jurisdiction over problem framing and strategic analysis may become increasingly contested. This authority could be challenged either by new actors entering the market, or by the tools themselves as they take on the cognitive work that once defined the profession.

2.3.3 Implications for Trust

When examined through the ABI framework of Mayer et al. (1995), generative AI presents a mixed picture. On the one hand, its ability is difficult to deny. Large language models can process vast amounts of information, generate fluent text, and produce outputs that often appear highly competent. Yet this competence is not grounded in experience (Larson, 2023). As established in Section 2.3.1, AI performance is fundamentally computational (Bender and Gebru, 2021). It resembles what Collins and Evans (2007) would describe as simulated or interactional expertise, rather than genuine contributory expertise based on lived practice.

The limitations become even clearer when looking at the other two dimensions. Benevolence presupposes intentions, empathy, and relational engagement. AI systems do not possess these qualities. They cannot genuinely care about a client's success, nor can they adapt their behaviour through emotional understanding or social sensitivity (Larson, 2023). Integrity raises an even deeper issue. Generative AI has no moral agency, can generate hallucinated information, and cannot be held responsible for the consequences of its recommendations (Larson, 2023; Susskind & Susskind, 2015). It may sound confident, but confidence is not accountability.

As a result, the leap of faith described by Nikolova et al. (2014) might become more complicated. Clients are no longer placing trust solely in a human advisor. Part of the advisory

process is now mediated by an algorithm, whose logic may remain unclear to both the client and, at times, the consultant.

Finally, Durán and Pozzi (2025) highlight a crucial distinction: clients may come to rely on AI because it is fast, consistent, and efficient, without this reliance turning into genuine human trust. This difference becomes particularly important in situations where the client cannot easily judge whether an AI-generated output is accurate, well-founded, or simply plausible.

2.3.4 Implication for Transparency

The integration of generative AI fundamentally reshapes transparency challenges in consulting. Rather than simply making consulting practices more or less transparent, the literature identifies two primary impacts that complicate this dynamic, creating a genuine transparency paradox.

The first major challenge concerns the issue of transparency: what are the risks of admitting to using AI? Schilke and Reimann (2025) highlight a real dilemma. In fact, a professional who openly admits to using it often inspires less trust than one who does not mention it. To explain this loss of trust, the authors show that it all comes down to “perceived legitimacy.” In concrete terms, legitimacy is simply the fact that a practice is deemed “appropriate” or normal within a specific context (Suchman, 1995, cited in Schilke & Reimann, 2025). In the area of consulting, where high value analytical work is closely tied to exclusive human judgment; clients are often surprised when they realize that their consultant's report was produced with the help of AI. As a result, even if the report's content is accurate or helpful, it may not seem as legitimate to a client as they would normally expect from “human” consulting work because it is produced with the assistance of AI. When legitimacy falls short, trust will most likely go along with it (Bitektine & Haack, 2015; Zucker, 1977, cited in Schilke & Reimann, 2025). In other words, it is not only the quality of the work that matters, but also whether the way it is produced fits with what people consider appropriate in a professional context.

Yet, actively hiding AI use carries an even severer risk, as post-hoc discovery by clients can result in a devastating loss of trust. Furthermore, disclosure does not yield a simple, linear outcome. Yu and Li (2022) show that transparency can produce dual, contradictory effects. On the one hand, explaining how AI processes information can reinforce perceptions of the consultant's technical competence. On the other hand, excessive transparency, especially when AI closely mimics human reasoning, may generate discomfort or threaten the clients' sense of human uniqueness, ultimately reducing trust rather than strengthening it.

The second impact concerns the relationship between transparency, model accuracy, and the inherent "black box" nature of AI. The problem is that these models can easily make mistakes and incorporate biases, and they do so with absolute certainty (Larson, 2023). Given this lack of reliability, Tronnier et al. (2025) emphasize that transparency is not a magic solution. On the contrary, if the result turns out to be false, having been transparent about the use of AI will actually accelerate the loss of trust. Ultimately, this completely debunks the common misconception that "the more transparent we are, the more people trust us" (Huynh & Aichner, 2025).

2.4 Research Gap

Despite the growing interest in generative AI in professional services, the existing literature remains fragmented. Research has developed along the three main topics mentioned earlier: expertise, trust, and transparency, but these streams rarely speak to one another. As a result, important blind spots remain, especially when the focus shifts to management consulting.

Current research tells us how AI handles analytical work and alters project stages (Mohan, 2024), but it tends to treat expertise as a static trait. In reality, consultants construct their authority daily through highly specific tasks interactions. What remains missing from the conversation is how AI affects this day-to-day reality, especially since the experience varies drastically depending on a consultant's specialization.

Second, trust research continues to rely heavily on inter-human frameworks such as the ABI model (Mayer et al., 1995). While foundational models offer a solid starting point, they struggle to account for the new reality of hybrid advisory work involving generative AI. Today, the landscape has completely changed. Customer trust is no longer limited to human experts; it must also incorporate the results of machines. And yet, we still know very little about how professionals handle this situation, especially when they are the ones who must take ultimate responsibility for work generated by AI.

But the real problem with current literature is that it tends to treat expertise, trust, and transparency as three completely separate elements. This perspective somewhat misses the reality on the ground, where everything is interconnected. In practice, changes in expertise directly impact client trust; this trust then dictates choices regarding transparency, which inevitably ends up altering how the consultant's expertise is perceived. Because previous studies have ignored this cyclical dynamic, there is a severe lack of research on what truly constitutes a consultant's added value in the era of generative AI. This thesis aims precisely to

fill this gap by taking a closer look at how professionals navigate these three dimensions on a daily basis.

3 Methodology

3.1 Introduction

To address the following research question: “From the consultant’s perspective, how does the progressive use of generative AI in consulting organisations influence the preservation of human expertise, trust, and transparency in the client-consultant relationship?”, a literature review was conducted. Based on the insights drawn from this review, four research propositions were formulated.

These four research propositions follow a sequential logic. P1 diagnoses how GenAI transforms the consultant's expertise; P2 and P3 examine what this transformation does to trust, first from the consultant's side, then from the client's; P4 describes how consultants behaviourally respond to this reconfiguration through differentiated disclosure strategies.

While presented sequentially for analytical clarity, this ordering does not imply a strict causal chain. In practice, the three dimensions pull on each other in both directions. How a consultant manages disclosure (P4) shapes how the client perceives their expertise (P1). How the client responds to AI involvement (P3) conditions what the consultant feels they need to protect through integrity and benevolence (P2). The sequence helps structure the analysis, but the underlying phenomenon is better understood as a knot than a chain. For the sake of conciseness, these research propositions will hereafter be referred to as P1, P2, P3, and P4.

3.1.1 Research propositions

The first research proposition posits that the integration of GenAI shifts the consultant’s added value from contributory expertise, the ability to directly produce analysis or deliverable (Collins and Evans, 2007) toward interactional expertise, focused on interpretation, translating needs, and problem framing.

As a result, this change makes it necessary to re-examine the issues of trust and transparency addressed in the following research propositions.

The second research proposition (P2) posits that if P1 is correct, then the ability dimension of trust (Mayer et al., 1995) will not be a differentiating factor anymore due to commoditization of technology. When high-quality content becomes easily accessible, consultants can no longer

rely on production skills alone to signal competence. Instead, they must increasingly build trust through benevolence and integrity (Mayer et al., 1995).

The third research proposition (P3) shifts the perspective from the consultant's response to the client's expectations. It posits that *reliance* (dependence on an efficient tool), and *interpersonal trust* (trust placed in a human) coexist. Specifically, clients rely on GenAI for operational efficiency, but they maintain a strong need for human, interpersonal trust when making strategic decisions. However, this assumes the client is aware that the consultant is using GenAI, which leads to the final research proposition.

The fourth research proposition (P4) posits that consultants, faced with the transparency paradox (Schilke & Reimann, 2025), develop a differentiated disclosure strategy based on the symbolic value of the task. Drawing on the taxonomy of Strobel et al. (2024), consultants are comfortable disclosing the use of tools acting as *reimaginators* for simple, repetitive tasks, but they tend to minimize or hide the use of *synthesizers* when performing tasks that are at the core of their recognized expertise.

To address these research propositions, this research employs a qualitative approach using semi-structured interviews (DiCicco-Bloom & Crabtree, 2006). A qualitative approach is particularly suitable for this research because the core phenomena cannot be reduced to measurable variables. Unlike closed questionnaires, semi-structured interviews provide the necessary depth to understand how consultants make sense of their practices (Braun & Clarke, 2006).

3.1.2 Sample Composition

The sample consists of 16 consulting professionals working in various consulting firms across Belgium. To ensure confidentiality, the names of the specific firms are not disclosed. However, the samples encompass a diverse range of organizations, including “Big Four” firms, mid-sized players and more specialized structures. This diversity aims to provide a representative view of the local consulting market and to avoid biases tied to any single firm’s corporate culture. To give an idea of the types of firms involved without revealing their identities, indicative data regarding their size range can be found in the Appendix 2.

The interviewees cover various hierarchical levels, from junior to senior consultants, as well as diverse specializations: management consulting (ERP implementation), risk management, artificial intelligence and sales force. This mix of profiles and specialisations was intentional. Indeed, it allows for the comparison of different perspectives on the integration of GenAI.

Participants were identified and contacted through several complementary channels: professional networks, LinkedIn connections, and email.

3.1.3 Data Collection

The 16 interviews were conducted online via Microsoft Teams, a method that offers flexibility and accessibility. Fifteen interviews were conducted in French and one in English, in accordance with the participants' linguistic preferences. However, the interviews conducted in French were all translated in English for better cohesiveness. The interviews lasted between 25 and 50 minutes.

The interviews were guided by a structured interview guide (Appendix 3). This guide served as a baseline but was not a rigid script. Indeed, additional questions were used to delve deeper into specific points and allow participants to freely elaborate on their answers.

All participants are anonymised. Besides, to facilitate the reading of the results, interview excerpts are not attributed to participants' names but to anonymized codes. Each interview is identified by a category corresponding to its field of specialization ("ERP" for ERP system implementation, "RISK" for risk management, "TECH" for AI&Data and "FUNC" for functional lead), followed by a number indicating the chronological order of the interview. Additional demographic and professional details, such as their job position, the exact sub-domain, the Generative AI used, and the anonymized profile of their firm, are summarized in a comprehensive table in Appendix 2. Besides, the transcripts of the interviews can be found in Appendix 5.

3.1.4 Data Analysis

Data analysis was conducted using thematic analysis (Braun & Clarke, 2006), combining deductive coding based on the research propositions with inductive coding to allow themes not anticipated by the theoretical framework to emerge.

Initially, a codebook was constructed based on the four research propositions. Each code includes a concise definition, a detailed definition, and precise inclusion criteria to ensure consistent application across all interviews (DeCuir-Gunby et al., 2011) (Appendix 4).

The coding process quickly brought to light several emerging themes, concepts that were missing from the initial theoretical framework but surfaced repeatedly across the data. Rather than filtering these out, these themes were kept as distinct analytical categories because they capture real-world dynamics that enrich the original research propositions.

Finally, two levels of reading were employed. A horizontal analysis, focusing on each participant, provides an understanding of their profile and how they perceive GenAI, trust, and transparency. A vertical analysis, focusing on each research proposition, allows for the comparison of responses across participants. The combination of these two approaches strengthened the coherence and robustness of the analysis.

3.1.5 Limits of Methodology

This study has several limitations that call for caution when interpreting its conclusions.

First of all, the absence of director- or partner-level profiles in the sample represents a significant analytical gap. Partners and directors operate at precisely the level where transparency decisions are made which makes their absence from the sample a real limitation for P4 in particular. Recruitment efforts were made through consulting firms at this management level. However, these candidates proved difficult to find within the given timeframe and through the available recruitment channels.

Besides, this study is grounded exclusively in the Belgian context. That alone already sets clear boundaries. What is observed here cannot be automatically extended to other national or cultural settings, where institutions, norms, and professional practices may differ in subtle but important ways. What holds true in one specific environment may simply not apply elsewhere.

Finally, there is the question of time. Generative AI is evolving at a very fast pace. As a result, the practices, perceptions, and norms described in this study should be understood as part of a transitional moment rather than a stable configuration of the profession.

4 Analysis and Findings

4.1 Introduction

Building on the methodological framework established in the previous chapter, this section presents the detailed findings of our thematic analysis. The results follow a clear hierarchical structure. We begin with a foundational finding about what clients value, then examine the four research propositions. Each of the four research propositions constitutes a main section, which is then further divided into specific subsections dedicated to analysing its associated codes. To capture the full richness of the data, this structure continuously integrates a horizontal perspective that accounts for the individual profiles and specific contexts of our participants.

Consequently, the different analytical points are systematically demonstrated and illustrated using direct verbatims from the consultants.

4.2 What Clients Valued Before GenAI

Before examining how GenAI transforms the consulting profession, it is essential to establish the foundational baseline of the consultant's value proposition as perceived by the participants themselves. Although the interview guide did not directly ask participants to compare consulting “before” and “after” the arrival of AI, their responses still reveal an implicit understanding of what traditionally defines the profession. This is partly because several interviewees entered consulting at a time when generative AI tools were already beginning to spread. Nevertheless, when participants explained why clients continue to rely on consultants, they indirectly described the core sources of value in their work. Across the interviews, five recurring elements emerged.

The first is technical expertise. ERP01 confirms that clients primarily seek advanced mastery of systems and tools they do not possess internally. The second driver is the precision of contextual judgment. As FUNC01 emphasizes, the objectivity of an external perspective is crucial. Companies do not expect mere technical theory. They demand an analytical framework capable of adapting these solutions to their own organizational complexity. The third source of value is the capture of tacit knowledge. Many participants emphasized that their mission is to identify informal company practices and transform them into formalized, replicable processes. The fourth dimension is human support. According to ERP02, technical expertise alone is never enough. The consultant must guide teams through the change process and ensure the adoption of new tools. Finally, all these skills revolve around a non-negotiable condition: trust. All participants stressed this point. Reliability and integrity remain the only true guarantees of a lasting client relationship.

Taken together, these elements help explain the traditional value attributed to consultants. Their role is not limited to providing technical knowledge alone. Clients also rely on consultants to interpret complex situations, structure internal knowledge, support organisational change, and provide advice they consider trustworthy. This broader role provides an important point of comparison for evaluating the impact that generative AI may have on the profession.

4.3 P1: How Does GenAI Transform Consultant Expertise?

This first research proposition posits that the integration of GenAI shifts the consultant's added value from the capacity to produce analysis (contributory expertise) toward the capacity to interpret (interactive expertise) and make sense of it. Analysis of the interviews shows that this research proposition is broadly confirmed, however, with several important nuances and implications: while the reduction in manual production is very widely observed, it does not necessarily mean the disappearance of human expertise, but rather a redistribution of key competencies.

4.3.1 *The reduction of Manual Production*

Most participants describe AI as a tool enabling the generation of initial drafts at the beginning of projects, tasks that were previously produced manually. For instance, the reading and synthesis of documentation, particularly in ERP implementation contexts. The use of GenAI is notably pronounced in the middle of projects and among the most intensive consultant profiles such as technical or data environments.

ERP06 formulates this shift when he estimates that *"the pure creation phase (writing or coding) has already been largely reduced, [...] to about 90% today. If we go back five years, to produce code, you really had to be very technically competent."* While this figure is an individual estimate, the underlying trend is consistent with what many participants described in the interviews. For example, TECH02 confirms the pattern from AI engineering, saying that a significant portion of the code is generated or assisted by AI. His role consists of defining requirements, providing context, and assigning tasks. Once this foundation is configured, AI can generate code or execute certain actions autonomously, which he then adjusts. This logic shows a clear shift away from the traditional way consultants work, where they build a deliverable step by step from scratch. Now, the initial draft is automated, moving the human effort toward reviewing and adjusting what the tool generates. The same goes for hybrid tasks that mix functional requirements with technical logic. For example, ERP05 noted using GenAI regularly to rewrite certain paragraphs, create quick summaries, or generate Excel formulas.

However, several respondents offer a more nuanced view. ERP01, but also all the risk consultants, freely admit that GenAI has taken over part of their daily production, yet they insist the essence of their role has not changed. Some participants are indeed opposed to the idea of a potential reduction in the role of human expertise. This is understandable because, in risk management, the relevant technical skill lies in analytical judgment, not in production capacity.

AI still struggles with this kind of reasoning. These differences suggest that the reduction of human expertise depends on the complexity of the tasks. Routine and repetitive work is clearly becoming commonplace, while in-depth analytical work may retain its value for longer. Therefore, the technical capacity for production might not be a differentiating factor anymore. If GenAI can generate code, documentation and analyses, a consultant's added value can no longer rest solely on their production capacity.

4.3.2 *Relational Expertise as The Primary Differentiator*

This is the point on which consensus is most unanimous: only TECH03 does not explicitly address this dimension. Moreover, all other participants confirm that human relationships, communication, and understanding of client needs become central competencies in the face of AI, more than it was ever before. This interactive expertise unfolds along two dimensions.

The first dimension of this expertise is relational. What GenAI does not capture is the nuanced human dynamics: the hesitations, the unspoken elements, the power dynamics and the emotions within a meeting room. ERP consultants have clearly highlighted this evolution. ERP02 emphasized that *“honesty, communication, patience, pedagogy, and even humour”* are elements that allow clients to feel comfortable and wanting to stay. Ultimately, trust depends more on interpersonal skills and human connections than on technical skills or GenAI. As RISK03 points out, even when AI processes a meeting transcript, it remains completely unable to *“interpret the emotions, reactions, or concerns”* of the clients. Interactive expertise requires the ability to interpret human situations. This involves being present (often physically) and understanding what is left unsaid. This is reinforced by RISK01, which emphasizes that a report is never merely a document, but is based on *“exchanges, clarifications, adjustments, and recommendations grounded in interaction.”*

The second dimension, more cognitive in nature, is the framing. It naturally extends the first dimension: understanding the client also means knowing how to translate what they express, often in a vague manner, into a precise and operational problem. Building on this, ERP05 stresses the importance of “active listening”. Unlike machines, the consultant must be able to *“interpret what the client says”* in order to correctly diagnose the *“underlying problem”*. This interpretive dimension, understanding not only what clients say, but also what they mean, represents tacit knowledge (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) that resists codification. This has a direct impact on the requests made to GenAI.

Indeed, the interviews highlighted another dimension, aside from the theoretical one, mentioned by nine participants: prompt engineering the ability to : “find the most appropriate prompt to allow a LM [*Language Model*] to solve the task at hand” (Liu et al. ,2023, p. 195:2) , as a new valuable competency. It is essential to note that prompt engineering is not synonymous with interactive expertise, but rather its continuation. This distinction is fundamental to understanding the future of consultants' added value. Among those who confirm this, ERP05 makes it a fundamental competency formulated explicitly: “ *Knowing how to prompt well is essential, because it is based on the question that the AI will give an answer.* ” Similarly, ERP04 describes a progressive, quasi-methodological approach, where the prompt becomes a sequential approach: “*[...] if it [ChatGPT] gives me 12 points, I take point 1 and ask a bunch of questions about it. Once I understand it well, I move on to point 2, etc. I really work in steps, a bit like sprints.*” This observation illustrates that the use of GenAI does not necessarily reduce intellectual work: it modifies the way it is carried out. Therefore, the framing addresses previously makes sense: if consultants do not understand the client’s need, how can they prompt correctly? This dimension directly joins the logic of framing and translating needs, but it also adds to the consultant's skillset by broadening their expertise in a certain way.

However, prompt engineering is only one facet of this transformation. The other side concerns how the consultant values their work in this new environment shaped by GenAI. As AI takes over the routine work, the consultant's value is no longer about creating from scratch, but about quality control. FUNC01 compares this new reality to “*the role of a manager who checks the work produced by juniors.*” The consultant becomes a “*guarantor of quality*” rather than a “simple producer” (TECH02). This shift implies a strong reliance on the consultant's ability to spot AI errors. Interestingly, participants feel increasingly confident in doing so. As TECH01 notes, while hallucinations were a major issue in the early days of GenAI, the tools have significantly improved: “*Today, it is less and less the case. There is therefore greater trust in artificial intelligence. That said, you must always keep a critical eye.*” This critical eye is precisely what their pre-existing expertise provides. Still, this means the consultant carries all the risk. As ERP06 warns, failing to review the AI's output can lead to major errors and lost client trust.

4.3.3 *The Impact on Juniors*

However, this transformation of the consultant’s role has consequences that go beyond the individual. It affects the very structure of firms and access to the profession for new generations. Eight participants mentioned significant repercussions on junior positions. As stated by

TECH05: “[...] it has become harder to sell junior consultants. [...]When I graduated from university, I had many companies to choose from. If I were graduating today, I honestly think it would be much more difficult to enter the job market.” While this statement could be applied to several others industries, this is particularly relevant. This is because certain tasks, formerly performed by juniors like note-taking, creating slides or repetitive tasks can all be replaced as suggested previously. Nevertheless, it also wipes out different ways of thinking that used to make team collaboration so strong and valuable.

TECH04 also raises a concerning issue: “Juniors are already mainly using AI: they ask ChatGPT, then they check and adjust. There is almost no completely manual production of code or documentation anymore.” If juniors no longer produce from scratch, they do not acquire the judgment necessary to validate. To properly check the AI’s work, consultants need solid prior expertise. The problem is that today’s junior consultant risk never developing this expertise, since the AI now does the work for them. Again, TECH04 explicitly warns against this danger: the risk of cognitive dependency. He stresses that if GenAI replaces human reflection instead of just clearing away redundant tasks, the learning process simply collapses. The threat is real. If we train an entire generation to rely on machines to perform basic analyses, we simply run the risk of large-scale intellectual stagnation.

4.3.4 *The Impact on Jurisdiction*

Finally, the last consequence resulting from this first research proposition is a threat to the very jurisdiction of the consultant profession. Even if the participants never named it explicitly, this threat deserves in-depth analysis. ERP06 formulates the stakes most clearly: "If this expertise becomes much more accessible thanks to AI, clients will logically have more room to maneuver and might think they can do more in-house with their own teams." If the client can produce themselves what they used to commission from the consultant, jurisdiction collapses not in favor of another profession, but in favor of the tool itself. However, we must not generalize this risk to all of consulting’s field. What the corpus shows is a selective threat: it bears on the most codifiable and reproducible tasks: code production, documentation, writing standardized deliverables. Because AI cannot provide a solution that is “genuinely adapted and specific” as said by RISK03, missions involving internal politics or change management remain secure. Consequently, the profession’s jurisdiction is not disappearing but fragmenting. This points toward the trajectory envisioned by ERP02: the end of the traditional generalist consultants, and the rise of the “*expert consultants*”. Even if consultants already operate within specific fields, this evolution means they must deepen their existing expertise even further.

4.4 P2: How Does GenAI Reconfigure Client-Consultant Trust?

4.4.1 Introduction

This second research proposition states that if GenAI commodifies the technical dimension, the capacity to produce, the consultant must compensate through the two other dimensions of trust, benevolence and integrity, to remain differentiated in the client's eyes. The data reveal a more complicated story than P2 initially suggests: most participants do not spontaneously report degradation of consultant-client trust due to GenAI use. Yet deeper analysis using the ABI framework reveals a reconfiguration of trust's foundations. This research proposition suggests that trust does not disappear but transforms.

4.4.2 The Commodification of Technical Ability

As established in P1 (see Section 4.3.1), the capacity to produce analysis is no longer a reliable differentiator because GenAI has made it broadly accessible. What P2 adds is the trust implication of this shift: if technical ability can be replicated by a tool, on what basis do clients decide to trust a consultant?

ERP06 captures this directly: *"Today, even a business graduate, whether at bachelor's or master's level, can produce something functional thanks to AI."* TECH05 draws the implication for professional hierarchy: *"In a way, it can turn a junior into someone who performs more like a senior."* As the ability to produce work is no longer an exclusive skill, technical seniority alone is no longer sufficient to guarantee client trust.

This dynamic does not apply uniformly, however. Risk consultants largely pushed back, in their context, the relevant technical ability is analytical judgment, not production capacity, and AI does not credibly replicate that (see Section 4.3.1). Where commodification does apply, benevolence and integrity (Mayer et al., 1995) become the primary differentiators. Where it does not, traditional trust dynamics largely persist.

4.4.3 The Over-Investment in Benevolence

If technical capacity ceases to be a differentiator, what remains, and what is confirmed by 15 out of 16 participants, is the demonstration that the consultant acts in the client's genuine interest, not merely in the firm's interest.

Many participants identify communication as fundamental to benevolence, positioning it as the foundation from which all relational trust emerges. As RISK02 emphasized, while some

individuals naturally possess good communication skills, it is primarily a competence developed through experience with various clients. This observation connects directly to the interactive expertise identified in P1. ERP05 identifies how critical benevolence is: *“There is also benevolence: showing genuine care towards the client, not just trying to sell them features at all costs simply because they are a client and you want to maximize revenue. [...] This can even go as far as advising them against certain options if they don't truly match their needs or goals.”* Therefore, benevolence is not passive but requires the consultant to prioritize client welfare over firm revenue.

Additionally, a concept mentioned in many interviews was that genuine care and human commitment emerge as irreplaceable value. This shift is critical because it cannot be replicated by GenAI. As ERP01 argues: *“[...] the human component and the change management on the human side are important, and I think AI won't be able to replace that.”* ERP03 articulates benevolence around availability, a dimension often overlooked: *“For me, I think it's availability. [...] Because clients want everything, right away, all the time. And as a consultant, you must find the right balance: help them understand that you are there as support, that you are available.”* Being available is how consultants actually show they care. In a world where AI can instantly answer any question, the fact that a human takes the time to listen and engage becomes much more meaningful.

4.4.4 The Over-Investment in Integrity

The integrity dimension has also been confirmed by all participants without exception: the consultant becomes the guarantor of the reliability of what is produced, even when it is GenAI that produced it.

The question of responsibility was asked to all participants: in case of incorrect deliverables involving GenAI use, who was responsible? ERP03 explains: *“But regarding the use of AI specifically, for me, the blame rests solely on the consultant's shoulders. Because they are the one using it.”* Many other participants confirm this position. This matters because it closes any possible escape route: if the tool fails, the consultant owns the failure. TECH05 arrives at the same conclusion: *“The responsibility lies with the people using the tool, because they are the ones who must verify the information.”* The consultant is responsible for the deliverable they produce, regardless of whether AI is involved. As ERP05 states: *“If we provide a deliverable, it means we have checked and validated it.”* While the first proposition (P1) treated verification as a practical step, P2 looks at its deeper impact on trust. The core of the argument is that a

consultant who delivers AI-generated work without checking it first risks breaking the client relationship. The issue here is not dishonesty, it is simply that they failed to apply the professional judgment expected of them. Responsibility is also manifested in concrete operational practices. One such practice, the management of data security and confidentiality, has emerged as a critical expression of integrity but also another key added value for the consultant in the context of GenAI.

4.4.5 Cybersecurity as a Concrete Expression of Integrity

Eleven participants explicitly mentioned concerns related to cybersecurity and confidentiality as central to their AI use. ERP03 illustrates this through tool selection: *“I prefer using Copilot over ChatGPT: with ChatGPT, it's not written in black and white that the data you send it respects confidentiality rules, whereas with Microsoft, it is.”* This choice reveals that confidentiality management has become an important competency, a form of operational integrity that complements moral integrity. Nevertheless, large consulting firms using internal AI tools that do not disclose clients’ data face fewer of these dilemmas. Technology consultants described client specific restrictions that further constraints GenAI use. As TECH05 notes, highly regulated industries, like banking, *“strictly forbid the use of any external generative AI tools”* because the data are too sensitive. Therefore, using public GenAI in those cases is risky because those data will be stored in global databases. Today, a consultant’s added value also resides in carefully protecting sensitive data while using GenAI tools. Failing to do so risks to destroy permanently the client trust and consequently, losing the business relationship.

4.5 P3: How Do Client Expectations Change?

This third research proposition marks a significant shift in analytical focus: the unit of analysis moves from the consultant to the client. The research proposition posits that when confronted with AI, clients develop a certain reliance toward GenAI but still trust human judgment.

4.5.1 Reliance : Clients Delegate Speed to GenAI

A third of participants confirmed that clients increasingly rely on AI tools to gain speed and efficiency. This reliance becomes increasingly institutionalized. As RISK03 observes: *“[...] some clients even clearly tell us that they use AI themselves. So they naturally expect to receive a deliverable faster since they know we have additional tools at our disposal.”* For clients who are not resistant to GenAI use in projects, they expect faster results and view GenAI as a competitive advantage in terms of productivity. Indeed, TECH05 has observed that clients increasingly ask whether consultants use AI, perceiving it as a factor of competitive

differentiation. However, this reliance is fragile because it rests on speed rather than on an understanding of GenAI's outputs and limitations. Clients do not ask whether the results are correct; sometimes they do not have the knowledge to affirm it which is problematic. ERP03 experienced several time clients saying "ChatGPT told me this," expecting it to be true, "but in 99% of cases, the result is incorrect." Clients then turn to the consultant to remedy the situation, revealing that this reliance was not genuine trust but rather conditional dependence on performance. This observation reveals the fragility of reliance: it depends entirely on GenAI functioning correctly.

4.5.2 *Hallucinations: The Mechanism That Fragments Trust*

The gap between reliance and trust comes down to one thing: GenAI fails in ways that are hard to detect. Hallucinations are mentioned by fourteen of sixteen participants, and what they describe converges: GenAI produces plausible but false outputs, often without signalling the error itself. While hallucinations were mentioned briefly under P1 and P2, their actual impact on the client relationship only becomes clear at this stage.

In P1, we observed that AI standardizes production capacity, which in turn forces consultants to rely more heavily on their interactive expertise. Hallucinations are the exact reason why this shift must happen: a consultant without deep domain knowledge simply might not be able catch a plausible AI error. Domain expertise is no longer just a simple advantage. It has become an absolute necessity, primarily because AI-generated errors frequently disguise themselves as perfect answers. Spotting these mistakes is nearly impossible without deep industry knowledge. This reality brings us back to our second research proposition. As basic technical skills become increasingly standardized across the industry, consultants must rely much more heavily on their professional integrity. Hallucinations illustrate this dynamic perfectly. When a machine produces an error, it is always the human consultant who assumes full responsibility before the client, regardless of whether that specific output was rigorously verified or not. The consultant's integrity is tested not by the quality of their original work but by their willingness to take responsibility for AI-generated errors. This is why hallucinations matter beyond the technical level, they are what keeps human judgment necessary.

ERP02 illustrates GenAI limitation from the angle of domain specificity: *"Right now, when you ask Copilot or ChatGPT something specific about Business Central, they start talking about all the ERPs in the world. And right now, they're not very precise, so they invent features."* For instance, what GenAI does not master yet is contextual and domain-specific knowledge which

exactly constitutes the expertise of a consultant. Hallucinations are therefore more frequent where domain expertise is most necessary. This means that AI does not eliminate the need for expert consultants; rather, it makes their expertise more visible and more necessary.

4.6 P4: The fewest implications for transparency

This fourth research proposition posits that consultants are not uniformly transparent about their GenAI use, they modulate this transparency based on the nature of the task and its perceived symbolic value. P4 is the most contested of the four research propositions. Unlike P1, P2, and P3, which were broadly confirmed across the sample, the picture here is more divided. Roughly a third of participants describe deliberate opacity on high-value tasks, while the majority do not report any systematic disclosure strategy at all.

4.6.1 *Transparency on Mechanical Tasks*

Firstly, transparency regarding GenAI use is relatively common, though far from systematic. About one-third of the participants indicate they do not hide their use of the tool when it comes to routine activities such as rephrasing or drafting emails.

Transparency regarding the use of GenAI is widely accepted by clients when it is justified. For example, ERP01 openly states that it uses AI for rewriting, while ERP04 spontaneously mentions ChatGPT for quick checks. Consultants agree on one point: acknowledging this use poses no social issues, as AI is perceived as a simple time-saver, not as the core of their unique value proposition.

In large firms, transparency regarding the use of AI is addressed upfront through contracts: it becomes an institutional rule rather than an individual choice. Once the agreement in principle is approved, the client no longer seeks to know whether a specific task required the use of AI, as RISK02 observes in its infrastructure projects. This contractual framework relieves consultants of the need to justify their use of AI on a case-by-case basis. This provides a level of security that is often lacking in smaller firms.

Some participants explicitly state that they do not see the point of informing the client about this type of use, or that they simply do not do so. ERP05, for example, states that it is not necessary to inform clients, while ERP06 answers unequivocally, “*no, not at all*” when asked about defining the terms of AI use with the client in advance. In these cases, the lack of transparency does not necessarily stem from a strategy of concealment, but rather from an implicit logic where the use of digital tools is considered part of the consultant’s internal

operations. Beyond institutional guidelines, transparency is highly context-dependent.: TECH02 explains that he had to inform the client only when the project explicitly involved significant automation and reduced human intervention. In short, transparency on routine tasks happens, but almost never by design.

4.6.2 *Strategic Opacity for High Value Tasks*

When it comes to high-value tasks such as problem framing and analysis, one-third of the participants explicitly confirm that they deliberately maintain opacity about GenAI use. These are the tasks that carry high symbolic value that clients perceive as the core of consulting expertise. Therefore, those consultants articulate a clear mechanism: transparency about GenAI use for strategic work risks triggering in clients a questioning of whether the consultant is necessary at all. When asked if they revealed the use of GenAI during projects, ERP03 responds: *“No, I don't specify it. I don't do it because I feel like the client might say to themselves: “If he can do it with AI, then I can do it too.” And besides, I use it differently than the client. Because I know the tool I'm implementing.”* Similarly, ERP01 expresses caution and argues that the absence of contractual clauses regarding AI, as it is usually in smaller firms, is problematic. If the practice becomes too widespread, it risks devaluating the consultant's contribution in the eyes of the client.

However, this view is far from being shared by the entire sample. Two thirds of the participants do not describe any deliberate strategy of opacity and assert that the use of GenAI does not pose a risk of undermining credibility as long as the final result is of high quality. TECH01 explicitly advocate state idea that one should “not hide”, as the use of GenAI is becoming an industry standard as supposed earlier (see Section 4.5.1). Similarly, several respondents note that the client is primarily results-oriented, especially in ERP implementation, and does not question the production process, making the issue of transparency secondary.

A consultant's position in the hierarchy strongly influences their attitude toward transparency. Drawing on their technical expertise to establish credibility, much like TECH02, senior consultants have less hesitation about embracing the use of AI. As a result, senior consultants find it easier to admit that they have used AI. For them, using AI does not mean a loss of competence; in fact, it is quite the opposite. They draw on their expertise and experience to maintain their critical judgment. For junior or generalist roles, the situation is much more delicate. Their value to the client still relies heavily on basic output, that is, the specific task that AI is now automating.

But it also depends entirely on the type of tool used. There is a major difference between relying on public platforms and using secure, internal systems. As RISK02 points out, having access to a proprietary AI actually allows consultants to turn transparency into a reassuring selling point. By speaking openly about their internal tools, consultants directly reassure the client. They prove to the client that their sensitive data will not be leaked.

5 Discussion

The results section examined each research proposition individually. This discussion steps back to ask a broader question: what do these findings mean together ? The goal here is not to repeat what has already been shown, but confront it with the literature, identify what is confirmed , what is challenged , and what is genuinely new.

5.1 Redefining Human Expertise

Collins and Evans (2007) distinguished two types of expertise: contributory, which is the ability to produce work directly, and interactional, which is the ability to make sense of it, frame it, and communicate it. The literature predicted that GenAI would shift consulting value from the first toward the second. This shift is broadly confirmed by the data. Besides, the SECI model offers a useful way to explain why this shift occurs. GenAI primarily operates in the *combination* phase, where explicit knowledge (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) is organised, and reformulated at scale. However, it remains structurally limited when it comes to socialisation and internalisation : the stages through which consultants acquire tacit (Nonaka & Takeuchi, 1995, as cited in Nonaka et al., 1998) , embodied knowledge through experience and interaction.

This limitation is also reflected in the data. As some participants point out, even if AI has access to a full transcript of a client call, it cannot fully capture what happens beneath the words: emotions, subtle reactions, or unspoken concerns. Yet these relational cues are precisely what the socialisation process makes visible, and what often shapes consultants' deepest understanding of a client situation.

The data introduce an empirical dimension that the literature had not yet documented in the consulting context: the emergence of *prompt engineering* (Liu et al. ,2023) as a new competency, mentioned spontaneously by nine participants. This finding is theoretically significant because it shows that interactional expertise, in an environment shaped by GenAI , does not only manifest in the relationship with the client : it also manifest in the relationship

with the tool. The ability to translate a complex client need (Abbott,1988 ; Schein,1990 , as cited in Appelbaum & Steed, 2005) into precise instructions for the AI and to critically evaluate the outputs produced is more than a standalone technical skill , it represents the operational extension of problem framing. As mentioned in the interviews, if consultants do not understand the client's need , they cannot prompt correctly. Prompt engineering becomes therefore contemporary expression of interactional expertise proposed by Collins and Evans' framework. In this sense, it also extends the literature by highlighting a concrete and observable way in which interactional expertise is enacted in AI-assisted work.

Nevertheless, the data introduce an important nuance that the literature did not sufficiently anticipated : not all consultants experience this shift equally. While ERP and technology consultants describe a clear reduction in contributory work, risk consultants do not. This is because in risk advisory, the core expertise is analytical judgment, not production which AI cannot replace. This matters because it means the framework of Collins and Evans applies differently across contexts. The commoditization of contributory expertise is real, but it is selective.

Furthermore, this shift toward a validation role adds a critical nuance to the literature's warnings regarding AI hallucinations and the erosion of human judgment (Benbya et al., 2024; Dell'Acqua et al., 2023). At first glance, the data seems to contradict Dell'Acqua's findings, which suggested that AI lowers performance on complex problems because users over-rely on the machine. Indeed, the interviewed consultants do not experience this vulnerability today. They express strong confidence in their ability to detect errors, applying a critical eye grounded in deep, pre-existing contributory expertise. Because they learned their craft before the GenAI era, these senior professionals do not currently suffer from the automation bias predicted by the literature.

In the long run, however, data from the field confirm all the concerns raised in the literature. Just as Dell'Acqua warned about the loss of analytical rigor, participants also point to a real risk of cognitive dependency among newcomers. The danger is clear. If junior staff use AI to avoid thinking, rather than using it solely to automate tedious tasks, their entire learning process collapses. We are thus faced with a fascinating paradox. Today, the industry's confidence in AI remains strong because it is firmly protected by the expertise of current senior professionals. But this shield is only temporary, and it may simply cease to exist for the next generation.

This problem has another consequence that was not necessarily addressed in the literature review. There is a real risk to jobs and training for junior staff, particularly in consulting, as confirmed by eight participants. It is by performing these basic tasks that juniors build their tacit knowledge. If AI automates them, the SECI cycle is broken from within. The profession thus faces a severe paradox. The new paradigm, centred on validation, interpretation, and assembly, presupposes robust prior contributory expertise to exercise proper judgment (Collins & Evans, 2007). If young consultants rely on AI instead of learning by doing, they miss out on the basic practice they need to really learn the job (Nonaka & Takeuchi, 1995). Without that experience, it becomes very hard for them to develop real professional judgment.

This is where Abbott's (1988) idea of "professional territory" comes in. The data shows that AI is a threat, but not everywhere. Instead of taking over the whole consulting job, AI is only replacing people on simple, repetitive tasks. By contrast, missions involving stakeholder management, organizational politics, and contextual judgment remain largely beyond current AI capabilities. Consulting jurisdiction is therefore not disappearing. It is being redrawn. Several participants suggest that this evolution may push the profession toward a more specialised model, in which consultants increasingly differentiate themselves through domain expertise (Kubr, 2002).

5.2 Adapting the ABI Framework

The literature, drawing on the ABI model developed by Mayer, Davis, and Schoorman (1995), suggested that if generative AI increasingly commoditises the *ability* dimension of trust, consultants would need to compensate by strengthening the other two dimensions, namely benevolence and integrity. This research proposition is broadly confirmed by the data. Fifteen out of sixteen participants identify benevolence, understood as genuinely prioritising client interests over short-term firm revenue, as a key differentiating factor in AI-assisted consulting work. Several respondents describe benevolence as a real value that they apply consistently at work. For example, some of them actively tell clients not to buy a service if it does not fit their needs, even when it means losing money for the firm. This choice goes straight against short-term business profits. This behaviour actually expands the ABI model, which only talks about benevolence as a general feeling of goodwill.

Beyond confirming the ABI model, the data introduce an enrichment that the original framework did not anticipate. In environments shaped with the help of AI, integrity increasingly includes the management of cybersecurity and data confidentiality. This dimension was

mentioned spontaneously by eleven participants, suggesting that it is an emerging core element of professional legitimacy.

Integrity therefore takes on a more operational meaning. In the context of GenAI, integrity revolves around security aspects and protecting clients' data. We see this plainly in how consultants choose between ChatGPT and Microsoft Copilot, opting for one over the other purely based on confidentiality guarantees. This means technical setups are now actively redefining what professional integrity looks like, and future research using the ABI framework will need to account for this operational side. Inevitably, this evolution is disrupting how consultants manage their professional image, what McKenna (2006) already called "impression management". Historically, reliability was largely based on delivering flawless presentations. But those days are over. Generative AI can now produce a superficially perfect rendering in an instant; the aesthetics of a deliverable are no longer enough to gain trust. The real added value has shifted. Clients now value a professional's legitimacy based on how they guide the project and the quality of their human interactions.

Consequently, the rules of trust now change depending on the consultant's specific field. The need to lean heavily on this relational and ethical engagement fluctuates significantly across different consulting disciplines. For example, in IT and ERP consulting, AI handles the data and does the work much faster. Because of this, standard tech skills are no longer enough to stand out. Consultants in these areas feel a much stronger pressure to make a difference through their human relationships and good behaviour. By contrast, in risk advisory, where analytical judgment remains difficult to automate and where decision-making often depends on contextual interpretation, traditional trust dynamics appear more stable. These results therefore support a domain-sensitive application of the ABI model, rather than a universal one.

Finally, the data reveal a key mechanism that explains why trust does not shift fully toward AI systems. Hallucinations, mentioned by fourteen out of sixteen participants, are not only described as a technical risk. They act as the underlying factor that prevents reliance on AI from turning into full trust.

This finding empirically strengthens the conceptual distinction between reliance and trust proposed by Durán and Pozzi (2025), by identifying a concrete mechanism through which this distinction becomes relevant in consulting practice. More importantly, it creates an unexpected convergence with the findings on expertise. The sectoral and interactional expertise highlighted in P1 is precisely what enables consultants to detect hallucinations that appear plausible within

a given domain. Without sufficient domain expertise, a consultant may not be able to identify an error presented with confidence.

The existence of remaining hallucinations therefore make human judgment, and especially sector-specific expertise, irreplaceable. This challenges the assumption, present in parts of the literature, that AI simply commoditises consulting work because AI simultaneously accelerates production and reinforces the need for the domain expertise required to validate what it produces. The consultants most at risk are therefore not those whose production AI replaces, but those who lack the expertise to catch what AI gets wrong.

5.3 The Real Drivers of Transparency

Of our four research propositions, this one is the most debated. We initially proposed that consultants would systematically adjust their transparency based on the task, hiding AI use for strategic work while being open about routine tasks. However, the data tells a different story. The type of task matters a little, but the seniority of the consultant and the contractual context matter much more.

When consultants use GenAI for routine work, such as drafting emails or summarizing notes, they rarely hide it. In these cases, transparency is almost never a calculated move. It happens naturally or is simply guided by the firm's standard policies.

Things get more complicated with high-value strategic work. Here, AI acts as an assistant to help structure ideas and synthesize data. About a third of the participants reported keeping their AI use quite in these situations. However, this is not really about a deliberate culture of secrecy. It is simply that senior and junior staff experience the risk differently. Because a junior's value is still closely tied to their production speed, admitting that an algorithm did the heavy lifting feels like a severe professional risk.

Senior consultants, along with those who possess deep technical expertise, have much less trouble being open. Because their foundational knowledge is already solid, transparency does not discredit them. Their legitimacy comes from their architectural judgment rather than their execution speed. Confident that their core value cannot be replaced by a machine, they feel virtually no need to conceal their tools.

This practical reality directly challenges the dual-effect model proposed by Yu and Li (2022). The authors originally theorized that transparency could actively reinforce an image of professionalism. Our data shows that this positive effect is not truly confirmed. Today, even

when consultants openly discuss using secure, internal AI platforms, they are not doing it to boost their credibility. Reassuring a client about data security is a mandatory procedural requirement, not a strategic move to look innovative.

This procedural logic is deeply connected to how projects are framed. If the client approves the AI from day one, the issue of transparency is automatically resolved. The initial agreement firmly establishes the framework. As soon as the tools are approved, the client completely detaches themselves from their consultant's production process. Their sole evaluation criterion becomes adherence to KPIs and the results achieved. In this ideal scenario, imposing continuous transparency becomes pointless. But the reality on the ground is often quite different. In the absence of shared rules, there is total ambiguity. Each professional is then forced to improvise their own approach, adapting on the fly according to the limitations of their contract and the specific requirements of the task at hand.

Because of this upfront framing, the core dilemma highlighted by Schilke and Reimann (2025) is only partially confirmed. They warned that professionals who openly admit to using AI are trusted less. Our data shows that this threat does not materialize in formal contexts where internal tools are used or where agreements are made in advance. On the other hand, the second dimension of Yu and Li's paradox, along with Schilke and Reimann's legitimacy risk, becomes highly accurate in a different scenario. The dynamic changes entirely when consultants rely on everyday public tools like ChatGPT. In these cases, being fully transparent can actually backfire. A client might easily reason that if everyone has access to the exact same software, the consultant's personal added value becomes much less obvious. This explains why some professionals remain very careful about what they disclose when using public platforms.

All of these dynamics culminate in what this study calls "double opacity." Consulting has always been a bit of a black box, as clients can never fully observe the expertise they are buying. Generative AI adds a second "black box" into the equation. The reality is that even the consultants do not fully understand how these tools operate behind the scenes. As long as the work is correct, this lack of visibility goes almost unnoticed. Clients primarily judge the project based on the final results anyway. However, the situation becomes highly critical at the very first mistake. If a consultant misses an AI hallucination and delivers a false analysis to the client, the trap immediately closes. That is precisely when the issue of transparency, which was previously ignored, suddenly becomes the centre of attention.

As the interviewed participants made clear, clients will not blame the machine. Any unverified AI error becomes the direct fault of the consultant. This is exactly where the choice between transparency and opacity takes on its full meaning. Taking responsibility for a mistake is always difficult. It becomes much harder when the consultant has to answer for work they did not actually produce themselves. Assuming this new kind of professional liability has to start with being transparent about how the work was done. If a consultant hides their use of AI and the tool hallucinates, the loss of trust is total.

However, simply disclosing that AI was used is never enough to save the relationship. Transparency must always be backed by rigorous human validation (Susskind & Susskind, 2015). Ultimately, the only real safeguard against this double opacity is the consultant's own domain-specific expertise. This knowledge is what allows them to catch the errors, correct them, and confidently stand behind the final outcome.

5.4 What Remains of the Consultant's Role in the GenAI Era

Taken together, both the established literature and the insights from our participants confirm that the traditional value of consulting has long relied on a stable balance of five elements: technical expertise, contextual judgement, tacit knowledge translation, human guidance, and trust under uncertainty (Appelbaum and Steed, 2005; Abbott, 1988; Nonaka and Takeuchi, 1995; Mayer et al., 1995). For many years, this combination created a clear division of labour. Consultants were more than just knowledge providers. They were interpreters, translators, and guarantors of decisions in complex environments.

What was only mentioned by the literature, however, is how transparency worked in consulting before the emergence of GenAI. Historically, the relationship between expertise, trust, and transparency heavily relied on an asymmetry of information. Consulting has actually always functioned as a sort of black box (Pasquale, 2015). Clients paid for expertise they could not fully evaluate themselves, which naturally created uncertainty and scepticism. As Appelbaum and Steed (2005) highlighted, clients often feared receiving generic, recycled solutions rather than tailored advice. To counter this doubt, transparency acted primarily as a tool of perception (Dholakia, 2023). It was less about revealing every single technical step and much more about signalling competence to reassure the client. Once the client felt reassured by these careful signals, they stopped demanding to see the inner workings of the process. Ultimately, because clients could never fully eliminate the asymmetry of information, evaluating the consultant's actual work always involved a degree of uncertainty. Trusting the professional therefore

required a leap of faith (Nikolova et al., 2015). Clients were able to make this leap simply because of the consultant's proven expertise. This expertise justified their authority, and this authority generated the trust necessary for the client to take a leap of faith. This nurtured the ability dimension of the ABI model (Mayer et al., 1995). This is exactly why the traditional black box of consulting remained perfectly acceptable. Trust was built on the authority of the expert, rather than on the visibility of the process itself.

Generative AI disrupts this balance by democratising access to knowledge. What the interviews confirmed is that when a client can instantly generate the same analytical frameworks or summaries as the consultant, the traditional asymmetry of information disappears. This creates a potential crisis of legitimacy for the industry. This technological shift forces the sector to redefine its core value. Today, production tasks are increasingly automated, while judgement and responsibility remain central. But what is easily forgotten is how this transition shakes the very foundations of why people trusted consultants in the first place.

This is where the connection between expertise, trust, and transparency becomes critical. These three dimensions do not exist in isolation from the technological context. As generative AI becomes embedded in consulting workflows, they are being reshaped and reconfigured together. They are being transformed, sometimes exactly as the literature anticipated, and sometimes in completely unexpected ways. As GenAI augments human capacity, simply possessing knowledge loses its exclusivity. The issue of trust follows exactly the same trajectory. Previously, the equation was simple: the client trusted the consultant because the latter possessed exclusive knowledge. But those days are over. Now, the real added value has shifted. The client is primarily looking for someone capable of making decisions and providing guidance in the face of uncertainty. At the same time, the demand for transparency is becoming a real headache for the profession. In practice, total transparency is not always essential on a daily basis. However, it becomes a central issue in two very specific cases: when the customer relationship is still being built, or when it is necessary to take responsibility for an error.

In other words, this transformation reflects a profound shift. We are moving from a world where clients trusted consultants because of who they were, to a world where trust depends on what consultants are able to stand behind. It is no longer enough to simply appear competent. Consultants must now be able to justify, explain, and take full responsibility for outputs that are co-produced with artificial systems.

In reality, the consultant loses very few of their traditional qualities. They only lose the mechanical credit of having produced the analysis entirely from scratch. Empathy, emotional security, and deep human understanding remain irreplaceable, especially during times of organisational change. More importantly, the very definition of domain expertise is transforming. The modern expert is no longer just someone who knows. They are the person who contextualizes, prioritizes, arbitrates, and understands the deep human and business consequences of a given strategy. This deep expertise provides the vital foundation needed to critically engage with AI. It allows consultants to effectively prompt systems, interrogate outputs, and detect the limits or hallucinations of the machine.

In an environment where AI guarantees unprecedented efficiency, the true economic value of a consultant shifts toward risk mitigation. Their main job is reducing the cost of a bad decision. As a result, the core product being sold is the professional signature: a commitment to stand behind decisions, defend their rationale in front of stakeholders, and absorb the reputational consequences if outcomes fail. Companies continue to rely on consultants for the presence of a responsible actor who can validate decisions within their unique organisational context.

This shift fundamentally redefines the jurisdiction of the consulting profession (Abbott, 1988). The traditional monopoly over knowledge production is eroding, while a new jurisdiction emerges entirely around accountability. Generative AI cannot structurally occupy this new domain. In reality, this limitation has little to do with technology. It is deeply rooted in legal constraints and corporate practices that cannot be automated. AI will never understand the silent power dynamics of a boardroom. Similarly, it will never be able to shoulder the burden of failure before shareholders. Responsibility inherently requires a physical actor. Companies will always need a legitimate professional, capable of answering for their actions and absorbing the shock in a crisis. The arrival of generative AI is not just a simple technological evolution. It completely redefines the social contract that traditionally binds the client to their consulting firm.

The profession is moving from a stable system of implicit trust toward a more fragile and continuously negotiated reality. Legitimacy must now be constantly demonstrated, justified, and assumed. The deep interdependence of expertise, trust, and transparency is therefore not incidental. It is precisely what allows us to understand how consulting retains its immense value in a world where knowledge production is now permanently shared between humans and machines.

5.5 The limitations of the Research

Beyond the methodological limitations already discussed, the findings themselves should be interpreted with caution. They come with clear boundaries that need to be acknowledged.

The first limitation concerns the generalisability of the concept of double opacity. Although it is one of the most original contributions of this research, it is grounded in a relatively small sample of sixteen professionals operating within a single national context. At this stage, we are still dealing more with a conceptual proposal than with a well-established reality.

The second limitation relates to the temporal context. This research is taking place during a period of transition. GenAI is still being integrated into consulting practices, tools are evolving very rapidly, and firms' internal rules have not yet fully stabilized. The results therefore primarily reflect a snapshot at a given moment of an environment in transition.

Finally, the third limitation stems from the fact that the trends described in P1 vary significantly across different types of activities. For example, they are very pronounced in IT or ERP, but much more limited in risk consulting. That said, the sample does not cover all consulting sectors. Areas such as marketing or organizational transformation have been less explored and may follow different patterns.

5.6 Recommendations for Future Research

Naturally, the limitations and tensions raised in this study open the door to further reflection.

The first directly concerns the customer. In practice, companies' levels of maturity regarding artificial intelligence vary considerably. This significant difference in technological culture necessarily influences two key elements: the dynamics of trust (P3) and the requirements for transparency (P4). This is a blind spot that absolutely must be explored. For future research, it would be interesting to interview these customers directly, through in-depth interviews or surveys. The objective would be to understand how they, on the other side of the mirror, accept and experience this so-called "double opacity" in concrete terms.

A second avenue involves adopting dynamic and extended timeframes. To address the immediate risks posed to junior consultants, future research should empirically investigate how AI impacts skill acquisition over time. Longitudinal studies (Ployhart & Vandenberg, 2010) tracking the career progression and competency development of new consultants working in AI-intensive environments are needed. Academic research must explore alternative

pedagogical models to understand how interactional expertise and professional judgment can be cultivated when the traditional "learning-by-doing" phase of contributory work is automated.

Finally, a third avenue opens up around the creation of genuine sector-specific standards. Indeed, this is one of the most striking findings of this study: no common rules currently exist. To fill this gap, it would be particularly interesting to observe how transparency requirements will be structured in the future, notably under the impetus of professional federations and regulators. But it is also necessary to broaden the scope. Conducting comparative studies between different countries, or between more branches of consulting (such as marketing or supply chain), would help answer a central question.

6 Conclusion

This research set out to explore a simple but increasingly important question: *From the consultant's perspective, how does the progressive use of generative AI in consulting organizations influence the preservation of human expertise, trust, and transparency in the consulting-client relationship?*

After conducting sixteen semi-structured interviews with professionals across Belgium, the reality on the ground proved to be highly nuanced. Ultimately, when we examine the four propositions of this research, only three were confirmed by the interview data. The first (P1) is verified: the profession is indeed evolving from production expertise to interaction expertise. However, this transition is observed much more rapidly in IT integration (ERP) than in risk management. Our second hypothesis (P2) is also validated. As the tool trivializes basic technical skills, consultants now emphasize integrity and empathy. Today, this ethical requirement also extends to data security. The observation is the same for the third hypothesis (P3). While companies seek the speed of AI, the risk of algorithmic error maintains an absolute need for interpersonal trust. In contrast, P4 is only weakly confirmed, if at all. Since there is still no common practice in the market, many consultants disclose their use of AI based on their perspective on the situation, the need to share it, and their ethical values. The only standard to date remains when the issue of transparency is addressed in the contract.

These findings are important because expertise, trust, and transparency are closely linked in a self-reinforcing cycle. The use of generative AI creates a "double opacity" problem, meaning consultants can no longer fully explain the inner workings of their tools. As a result, the entire burden of transparency shifts to guaranteeing the final output. To achieve this, consultants must act as a critical safety net, applying their deep domain knowledge to filter out hallucinations

and secure client data. This active validation is precisely how professionals demonstrate their integrity and benevolence. Once clients realize the consultant is actively protecting their interests, strong interpersonal trust is established. This closes the relational loop: because the client trusts the consultant's judgment and receives an accurate outcome, the demand for algorithmic transparency simply disappears.

For decades, the value of a consultant was tied directly to technical skill, massive data processing, and analytical outputs. Now that generative AI handles this type of raw production at scale, the human practitioner's role is undergoing a profound shift. Clients are redirecting their expectations toward dimensions that machines cannot mirror. They actively seek out an authentic relational presence, the emotional intelligence needed to navigate complex corporate environments, and a deep sense of benevolence. In this environment, neutralizing the technological "black box" highlights the true worth of the professional, where value is measured through human judgment and trusted partnerships.

Taken together, these results point to a broader conclusion: generative AI is not disrupting consulting from the outside. It is reshaping it from within: redefining what expertise looks like, how trust is built, and what transparency can realistically deliver. Yet this transformation is unfolding without clear rules, without shared professional standards, and without consensus on who bears responsibility when AI-generated outputs turn out to be wrong. For now, consultants are often left to improvise individual responses to challenges that arguably require collective answers.

Ultimately, the real question is not whether human expertise will survive generative AI. This research suggests that it will and that it must. As long as AI remains imperfect, and hallucinations continue to illustrate this imperfection in concrete ways, someone will be needed to validate outputs with experience, judgment, and domain knowledge. What remains uncertain is whether the profession can develop the shared norms, safeguards, and institutional frameworks needed to make this redefinition explicit, coherent, and credible. The answer to that question will shape not only the future of consulting, but also the trustworthiness of the advice on which organisations increasingly depend.

7 Bibliography

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor*. University of Chicago Press. <https://tmie.hypotheses.org/files/2011/11/Abbott-Andrew-The-System-of-Professions.-Introduction.pdf>
- Appelbaum, S. H., & Steed, A. (2005). The critical success factors in the client-consulting relationship. *Journal of Management Development*, 24(1), 68-93. <https://doi.org/10.1108/02621710510572362>
- Belcic, I., & Stryker, C. (2026, 19 mars). ChatGPT. IBM. <https://www.ibm.com/fr-fr/think/topics/chatgpt>
- Benbya, H., Strich, F., & Tamm, T. (2024). Navigating Generative Artificial Intelligence Promises and Perils for Knowledge and Creative Work. *Journal of the Association for Information Systems*, 25(1), 23-36. <https://doi.org/10.17705/1jais.00861>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Mitchell, M. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610-623. <https://doi.org/10.1145/3442188.3445922>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77-101. <https://www2.uwe.ac.uk/services/Marketing/students/Newstudents/HAS/Using%20thematic%20analysis%20in%20psychology.pdf>
- Collins, H., & Evans, R. (2007). *Rethinking expertise*. University of Chicago Press.
- Crişan, E. L., & Marincean, A. (2023). The digital transformation of management consulting companies : a review. *Information Systems And e-Business Management*, 21(2), 415-436. <https://doi.org/10.1007/s10257-023-00624-4>
- Curnow, B., & Reuvid, J. (Eds.). (2003). *The international guide to management consultancy: The evolution of management consultancy and its role in changing the world of business* (2nd ed.). Kogan Page.
- DeCuir-Gunby, J. T., Marshall, P. L., & McCulloch, A. W. (2011). Developing and using a codebook for the analysis of interview data: An example from a professional development research project. *Field Methods*, 23(2), 136-155. <https://doi.org/10.1177/1525822X10388468>
- Dell'Acqua, F., McFowland, E., III, Mollick, E., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Kraye, L., Candelon, F., & Lakhani, K. R. (2023). *Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality* (Harvard Business School Working Paper No. 24-013). Harvard Business School. <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>

- Deloitte. (s. d.). *Beyond the code: Meet our people powering Gen AI*. <https://www.deloitte.com/uk/en/about/story/impact/beyond-the-code-meet-our-people-powering-gen-ai.html>
- Dholakia, U. (2023). *Transparency in Business*. <https://doi.org/10.1007/978-3-031-12145-6>
- DiCicco-Bloom, B., & Crabtree, B. F. (2006). The qualitative research interview. *Medical Education*, 40(4), 314-321. <https://doi.org/10.1111/j.1365-2929.2006.02418.x>
- Durán, J. M., & Pozzi, G. (2025). Trust and trustworthiness in AI. *Philosophy & Technology*, 38(1), Article 16. <https://doi.org/10.1007/s13347-025-00843-2>
- Fortune. (2025, 7 octobre). *Deloitte AI Australia government report hallucinations technology 290000 refund*. <https://fortune.com/2025/10/07/deloitte-ai-australia-government-report-hallucinations-technology-290000-refund/>
- Hargadon, A. B. (1998). Firms as knowledge brokers: Lessons in pursuing continuous innovation. *California Management Review*, 40(3), 209-227. <https://doi.org/10.2307/41165951>
- He, R., Cao, J., & Tan, T. (2025). Generative artificial intelligence: A historical perspective. *National Science Review*, 12(5), nwaf050. <https://doi.org/10.1093/nsr/nwaf050>
- Huynh, M., & Aichner, T. (2025). In generative artificial intelligence we trust: Unpacking determinants and outcomes for cognitive trust. *AI & Society*, 40(8), 5849-5869. <https://doi.org/10.1007/s00146-025-02378-8>
- Kipping, M., & Clark, T. (2012). Researching management consulting: An introduction to the handbook. In *Oxford University Press eBooks*. <https://doi.org/10.1093/oxfordhb/9780199235049.013.0001>
- Korzynski, P., Mazurek, G., Krzyrkowska, P., & Kurasinski, A. (2023). Artificial intelligence prompt engineering as a new digital competence: Analysis of generative AI technologies such as ChatGPT. *Entrepreneurial Business and Economics Review*, 11(3), 25-37. <https://doi.org/10.15678/eber.2023.110302>
- Koster, A. (2025, 1 octobre). Définition : Gemini. *Sfeir*. <https://www.sfeir.dev/ia/definition-gemini/>
- Kubr, M. (Ed.). (2002). *Management consulting: A guide to the profession* (4th ed.). International Labour Office. <https://imc-armenia.org/wp-content/uploads/2016/01/Management-Consulting.pdf>
- Larson, A. (2023). Large language model (LLM). *Research Starters*. EBSCOhost. <https://www.ebsco.com/research-starters/language-and-linguistics/large-language-model-llm>
- Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2023). Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9), 1–35. <https://doi.org/10.1145/3560815>

- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709. <https://doi.org/10.2307/258792>
- McKENNA, C. D. (2001). The World's Newest Profession: Management Consulting in the Twentieth Century. *Enterprise & Society*, 2(4), 673–679. <http://www.jstor.org/stable/23699625>
- McKinsey & Company. (2023, 14 août). *Meet Lilli, our generative AI tool*. <https://www.mckinsey.com/about-us/new-at-mckinsey-blog/meet-lilli-our-generative-ai-tool>
- Mohan, S. K. (2024). Management consulting in the artificial intelligence -LLM era. *Management Consulting Journal*, 7(1), 9-24. <https://doi.org/10.2478/mcj-2024-0002>
- Nikolova, N., Möllering, G., & Reihlen, M. (2014). Trusting as a “leap of faith”: Trust-building practices in client–consultant relationships. *Scandinavian Journal of Management*, 31(2), 232-245. <https://doi.org/10.1016/j.scaman.2014.09.007>
- Nissen, V. (Ed.). (2019). *Advances in consulting research: Recent findings and practical cases*. Springer. <https://doi.org/10.1007/978-3-319-95999-3>
- Nonaka, I., Reinmoeller, P., & Senoo, D. (1998). The 'ART' of knowledge: *European Management Journal*, 16(6), 673-684. [https://doi.org/10.1016/s0263-2373\(98\)00044-9](https://doi.org/10.1016/s0263-2373(98)00044-9)
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press. https://tetrazolelover.at.ua/Frank_Pasquale-The_Black_Box_Society-The_Secret_AI.pdf
- Ployhart, R. E., & Vandenberg, R. J. (2010). Longitudinal research: The theory, design, and analysis of change. *Journal of Management*, 36(1), 94-120. <https://doi.org/10.1177/0149206309352110>
- Schilke, O., & Reimann, M. (2025). The transparency dilemma: How AI disclosure erodes trust. *Organizational Behavior and Human Decision Processes*, 188, 104405. <https://doi.org/10.1016/j.obhdp.2025.104405>
- Schnackenberg, A. K., & Tomlinson, E. C. (2016). Organizational transparency: A new perspective on managing trust in organization-stakeholder relationships. *Journal of Management*, 42(7), 17841810. <https://doi.org/10.1177/0149206314525202>
- Strobel, G., Banh, L., Möller, F., & Schoorman, T. (2024). Exploring generative artificial intelligence: A taxonomy and types. *Proceedings of the Annual Hawaii International Conference on System Sciences*. <https://doi.org/10.24251/hicss.2024.546>
- Susskind, R., & Susskind, D. (2015). The Future of the Professions. In *Oxford University Press eBooks*. <https://doi.org/10.1093/oso/9780198713395.001.0001>
- Tronnier, F., Bernet, R., Löbner, S., & Rannenber, K. (2025). Consulting in the Age of AI: A Qualitative Study on the Impact of Generative AI on Management Consultancy Services. *Proceedings Of The . . . Annual Hawaii International Conference on System*

Sciences/Proceedings of The Annual Hawaii International Conference on System Sciences. <https://doi.org/10.24251/hicss.2025.017>

Visscher, K., Heusinkveld, S., & O'Mahoney, J. (2017). Bricolage and Identity Work. *British Journal of Management*, 29(2), 356-372. <https://doi.org/10.1111/1467-8551.12220>

Yu, L., & Li, Y. (2022). Artificial intelligence decision-making transparency and employees' trust: The parallel multiple mediating effect of effectiveness and discomfort. *Behavioral Sciences*, 12(5), 127. <https://doi.org/10.3390/bs12050127>

8 Appendices

Appendix 1: Definitions

Concept	Academic Definition (English)	Source(s) (APA 7th)
Management Consulting	An advisory service provided by trained persons who help organizations identify and analyse management problems and assist in the implementation of solutions	Appelbaum & Steed (2005, p. 68)
Explicit Knowledge	Knowledge that "can be expressed in words and numbers and easily shared."	Nonaka et al. (1998, p. 673)
Tacit Knowledge	Knowledge, "such as subjective insights or emotions, [that] is non-articulated and embedded in contexts and actions. It is highly personal and hard to verbalize or communicate."	Nonaka et al. (1998, p. 673)
Improvising Bricoleurs	Situational tinkering, 'making do' and recombining available materials in a creative manner	Lévi-Strauss (1966 as cited in Visscher et al., 2017, pp. 1-2)
Contributory Expertise	What you need to do an activity with competence (p.14). Contributory expertise, as its name suggests, enables those who have acquired it to contribute to the domain to which the expertise pertains: contributory experts have the ability to do things within the domain of expertise (p.24)	Collins & Evans (2007, pp.14,24)
Interactional Expertise	The ability to master the language of a specialist domain in the absence of practical competence	Collins & Evans (2007, p.14)
Jurisdiction	The link between a profession and its work (p. 20), which expresses the body of exclusive claims, recognized rights, and socially established control that a profession enforces over a specific set of tasks (p. 33)	Abbott (1988, pp. 20, 33)
Trust	A psychological state comprising the intention to accept vulnerability based upon positive expectations of the intentions or behavior of another	Rousseau et al. (1998, p. 395 ; as cited in Nikolova et al., 2014, p. 233)
Leap of Faith (Trust)	A confident suspension of uncertainty within the context of a newly forming relationship, achieved through socio-cognitive and emotional interactions that create a shared social space	Nikolova et al. (2014, p. 234)
Transparency	The perceived quality of information that is intentionally shared by a sender	Schnackenberg & Tomlinson (2016, p. 1788)
Disclosure (Transparency)	The extent to which relevant information is openly shared with stakeholders, where timing and relevance matter just as much as the quantity of information.	Schnackenberg & Tomlinson, 2016, p. 12)
Clarity (Transparency)	The perceived level of lucidity and comprehensibility of information received from a sender	Schnackenberg & Tomlinson, 2016, p. 13)
Accuracy (Transparency)	The perception that information is correct to the extent possible given the relationship between sender and receiver	Schnackenberg & Tomlinson (2016, p. 14)
Black Box	The term "black box" is a useful metaphor for doing so, given its own dual meaning. It can refer to a recording device, like the data- monitoring systems in planes, trains, and cars. Or it can mean a system whose workings are mysterious; we can observe its inputs and outputs, but we cannot tell how one becomes the other.	(Pasquale, 2015, p. 3)
Generative AI (GenAI)	A group of AI algorithms and models that are capable of producing new content, including texts, images, videos and	He et al. (2025, p.1)


	problem-solving strategies, with human-like creativity and adaptability	
Large Language Models (LLMs)	Advanced artificial intelligence systems pre-trained on massive datasets (or corpora) that use tens or hundreds of billions of parameters to "predict the most appropriate sequence of words" in order to understand and communicate in natural human language	Larson (2023)
AI Generators <i>(Taxonomy)</i>	GAI applications that "offer capabilities to create novel and innovative content based on only few user inputs," focusing around creating entirely new artifacts (e.g., text-to-image or text-to-video)	Strobel et al. (2024, p. 4550)
AI Synthesizers <i>(Taxonomy)</i>	GAI applications that "provide the capability to generate entirely synthetic data" strongly based on original inputs, used to train AI models, mitigate bias, or establish test landscapes without violating data privacy	Strobel et al. (2024, p. 4550)
AI Reimaginators <i>(Taxonomy)</i>	GAI applications that "pursue the goal of reinterpreting data in a novel way," innovating existing material (such as style transfers or image expansions) while keeping the underlying core semantics stable	Strobel et al. (2024, p. 4550)
Hallucinations (AI)	Information that conflicts with the source data (intrinsic hallucination) or cannot be verified by the source data (extrinsic hallucination)	Larson (2023)
Legitimacy (Micro-institutional)	Whether an action, tool, or process feels "appropriate" or socially acceptable within a given professional context, heavily dictating whether the disclosure of AI use reinforces or undermines professional trust.	Suchman (1995, as cited in Schilke & Reimann, 2025)
Prompt Engineering	Finding the most appropriate prompt to allow a LM [Language Model] to solve the task at hand	Liu et al. (2023, p. 195:2)
Longitudinal Research	Research emphasizing the study of change and containing at minimum three repeated observations (although more than three is better) on at least one of the substantive constructs of interest" (p. 97). It requires collecting data over time from the same units of observation in a manner that enables these units to be linked across time waves	Ployhart & Vandenberg (2010, p. 97)

Appendix 2 : Participants Data

Interview	Size Range	Company's Category	Job Title	Job Position	GenAI used
ERP01	250-4999	Mid-size	Business Consultant (Atlassian)	Senior	Copilot, ChatGPT
ERP02	250-4999	Mid-size	Business Consultant (Business Central)	Senior	Copilot, ChatGPT
ERP03	250-4999	Mid-size	Business Consultant (Dynamics 365)	Junior	ChatGPT, Copilot, Claude
ERP04	<50	Small Entreprise	Junior Consultant (Business Analysis/Processes)	Junior	ChatGPT
ERP05	> 5000	Large Entreprise	ERP Consultant (NetSuite)	Junior	IA interne BDO, Copilot
ERP06	> 5000	Large Entreprise	Consultant Tax & Technology (SAP GRC)	Junior	Internal LLM, Agent ABAP, ChatGPT, Claude, Gemini
ERP07	> 5000	Large Entreprise	Consultante Oracle TNT (NetSuite)	Junior	Internal LLM
RISK01	> 5000	Large Entreprise	Junior Risk Control Consultant	Senior	Internal LLM
RISK02	> 5000	Large Entreprise	Finance Transformation Consultant (Internal Audit)	Senior	Internal LLM

RISK03	> 5000	Large Entreprise	Consultant Risk Advisory	Junior	Internal LLM, Gemini, Copilot
TECH01	> 5000	Large Entreprise	Consultant AI and Data	Junior	Internal LLM, Microsoft Copilot
TECH02	> 5000	Large Entreprise	Consultant Digital Solutions	Senior	Claude (Opus), GPT, Codex, Mistral
TECH03	<250	Medium-sized Enterprise	AI Consultant / Engineer	Junior	Copilot, Internal Secure GPT
TECH04	<250	Medium-sized Enterprise	Generative AI Consultant & Business Analyst	Senior	Copilot, Copilot Studio, Azure
TECH05	> 5000	Large Entreprise	AI Consultant	Senior	Copilot, GitHub Copilot, Claude, Gemini, Maya, Heidi, Satric (internes)
FUNC01	> 5000	Large Entreprise	Consultant Functional Lead (Salesforce)	Senior	Internal LLM , Internal Coding Assistant

Appendix 3 : Interview Guide



Appendix 4 : Codebook

Research proposition 1 : Shift in Expertise

P1a REDUCTION CONTRIBUTIVE EXPERTISE			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	Reduction in the ability to directly produce analysis or deliverables	Contributive expertise (the ability to produce outputs "from scratch") is diminished. Consultants no longer create entire outputs but use AI to generate initial drafts. This reduction is measurable and quantifiable	Code when the participant mentions: • Reduction in manual production • Shift from "scratch" to "draft" • Reduction in production capacity
P1b INCREASE INTERACTIONAL EXPERTISE			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	Increase in expertise regarding interpretation, needs translation, and problem framing.	Interactional expertise (the ability to interpret, translate needs, and frame problems) becomes more critical. Consultants must better understand the client's needs to interpret AI outputs and frame problems strategically. This expertise becomes the new source of added value.	Code when the participant mentions: • Importance of interpretation • Translation of needs • Problem framing
P1c JURISDICTIONAL REDEFINITION			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	Redefinition of professional jurisdiction and boundaries	Based on Abbott's theory, AI challenges the "jurisdiction" of consultants by automating the abstract knowledge they once controlled. This leads to a professional struggle to reclaim jurisdiction over AI-mediated tasks and maintain cognitive control over the profession's core boundaries	Code when the participant mentions: • Control over knowledge • Professional boundaries • AI as a competitor for tasks • Reclaiming value
D1 PROMPT ENGINEERING MASTERY			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Data-Driven	Emerging technical skill in optimizing AI interactions	Mastery of prompt engineering as a specific, iterative skill. It involves the ability to refine instructions (iterative prompting) to obtain high-quality professional outputs. This skill bridges the gap between general expertise and technical AI execution	Code when the participant mentions: • Iterative process; • Refining instructions • Prompting techniques • "Trial and error" with the tool

D2 JUNIOR REDUCTION AND ROLE SHIFT			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Data-Driven	Generative AI reduces the need for junior consultants and transforms their roles	Generative AI replaces tasks traditionally assigned to juniors (note-taking, documentation, data filtering). This leads to a decrease in junior recruitment and shifts their remaining roles toward validation and supervision. Juniors must master AI quickly to remain relevant	Code when the participant mentions: • Reduced junior recruitment • Disappearance of junior tasks • Evolution of junior roles • Necessity for juniors to master AI • Impact on the learning curve

Research proposition 2 : ABI model Compensation

P2a TRIVIALIZATION TECHNICAL ABILITY			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	AI trivializes technical ability (The 'A' in the ABI model)	Technical ability is no longer a differentiator. AI can perform tasks previously reserved for consultants. This means technical skill alone is insufficient for building trust. Consultants must compensate through other dimensions (Benevolence and Integrity)	Code when the participant mentions: • Trivialization of technical skills • "Anyone can do what I do" • Loss of technical differentiation
P2b OVERINVESTMENT BENEVOLENCE			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	The consultant must over-invest in Benevolence (The 'B' in ABI).	Given the trivialization of technical skill, the consultant must demonstrate strong intentions, genuine care for the client's success, and altruism. Benevolence becomes a key differentiator	Code when the participant mentions: • Importance of showing good intentions • Demonstrating interest in the client • Empathy or understanding the client
P2c OVERINVESTMENT INTEGRITY			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	The consultant must over-invest in Integrity (The 'I' in ABI).	Integrity becomes a key differentiator by being honest, transparent, reliable, and responsible. The consultant becomes the "shield" against AI errors or data leaks	Code when the participant mentions: • Consultant's responsibility • Protection against AI errors • Protection against data leaks

D3 CYBERSECURITY AND DATA PRIVACY			
Type	Brief Definition	Detailed Defintion	Inclusion Criteria
Data-Driven	Critical management of data protection as a prerequisite for professional trust	Professional practice is constrained by strict cybersecurity rules. The consultant acts as a "gatekeeper" to ensure sensitive data remains outside public LLMs. This responsibility creates a new "ethical boundary" where the consultant's value lies in guaranteeing a safe environment for AI use	Code when the participant mentions: • Data leakage risks • Client confidentiality • Secure environments (VPN/Private instances) • GDPR compliance • Ethical data handling

Research proposition 3 : Reliance in GenAI , Trust in Consultant

P3a TECHNICAL RELIANCE DELEGATED TO AI			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	AI trivializes technical ability (The 'A' in the ABI model)	The client develops a form of "reliance" (technical dependency) toward the GenAI tools used by consultants. This reliance is based on an appreciation for the speed and efficiency that AI brings to the outputs	Code when the participant mentions: • Confidence in AI's capacity • AI reliability • -AI performance
P3b DECISIONAL TRUST RETAINED BY HUMAN			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	The consultant must over-invest in Benevolence (The 'B' in ABI).	The consultant remains responsible for the final decision. They validate, decide, and assume moral responsibility for the outputs. This confidence is based on interpersonal relationships and moral responsibility, not just on technical performance	Code when the participant mentions: • Lacks of critical thinking (GenAI) • Increasing Use • Verification • Validation
D4_HALLUCINATIONS_AND_OUTPUT_VERIFICATION			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Data-Driven	Generative AI produces hallucinations requiring systematic human verification	Generative AI regularly produces inaccurate, inconsistent, or fabricated responses (hallucinations). This creates a critical need for human verification before use. Consultants must invest time in checking, which reduces expected productivity gains	Code when the participant mentions: • Necessity to verify outputs • AI errors or inaccuracies • Time spent correcting outputs • Limited reliability

Research proposition 4 : Selective Transparency

P4a_TRANSPARENCY_MECHANICAL_TASKS			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	Transparency regarding AI use for mechanical tasks	The consultant is transparent with the client about using AI for mechanical tasks (formatting, layout, generating simple summaries). These tasks are perceived as time-savers rather than core added value. Transparency here does not threaten the perception of professional value	Code when the participant mentions: • Transparency for mechanical tasks • "I tell the client I use AI for..." • Mechanical tasks as time-saving.
P4b_OPACITY_STRATEGIC_TASKS			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	Opacity regarding AI use for strategic tasks	The consultant remains opaque about AI use for strategic tasks (analysis, framing, recommendations). These tasks are viewed as the "magic" of consulting and the primary source of added value. Opacity protects the client's perception of value and the consultant's professional legitimacy	Code when the participant mentions: • Opacity for strategic tasks • "I don't tell the client I use AI for..." • Strategic tasks as the source of value
P4c_VISIBLE_LABOR_REDUCTION			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Theory-Driven	AI disclosure reduces the perception of visible labor	When the consultant discloses AI use for outputs, the client perceives less "visible labor" (apparent effort, man-hours, visible complexity). If the client sees that AI generated the deliverable, they perceive less effort from the consultant, which threatens the perception of added value	Code when the participant mentions: • Reduction of visible labor after disclosure • Perception of reduced effort • AI "doing everything" or "generating automatically"
D4c_ECONOMIC_DILEMMA_AND_FEES_PROTECTION			
Type	Brief Definition	Detailed Definition	Inclusion Criteria
Data-Driven	Structural barrier to transparency due to fee-based business models.	The reduction in visible labor creates a structural barrier to transparency. Since clients are less willing to pay high fees when they perceive less effort, consultants maintain opacity to protect their fee-based economic model. Total transparency is viewed as a threat to revenue	Code when the participant mentions: • Structural barrier to transparency • The dilemma between transparency and revenue • Client wanting to pay less

Appendix 5 : Interviews

Interview ERP01

It's very kind of you to accept. So, I'll get straight to the point. Can you introduce yourself? Where do you work? What is your role in the company you work for? All of that, yes.

So I think I'm the one who will have to start the recording though, otherwise I'm not sure you'll get the copy, and then I'll send it by email once the recording is done. Since I'm the one who sent the link, I don't think you have control over all that. So there, I'll start the recording, I think we're good. At least, it says so. Anyway, regarding the introduction. My name is [Consultant Name], I am a business consultant at [Company Name] in the Atlassian team. So that is a very specific brand of products. They are products, project management software, collaboration software in general. And also ticketing. I don't know if you're familiar with the term. In a company, to manage internal or even external requests, we set up what we call ticketing, so it's the creation of a ticket. This actually represents a work item; based on a client portal, people enter their request. It actually reaches the operational teams, and they can pass on the necessary information to the people... well, the people who created the ticket, and actually follow up on this work item with all the necessary information. The most common example is when a new employee joins a company; they don't have a computer, they need a computer to work. They can create a ticket with the IT infrastructure team to request a computer, that's just one example among many. That's a hardware request, for example. So that's the ticketing part.

My role, primarily, is to support clients in implementing the tools of this particular brand. Whether it's through workshops, which we call "workshops" in our jargon, since we use a lot of English words, where we help them express their needs regarding problems they might have encountered internally, regarding their current situation, or a situation they would like to evolve towards. And then afterwards, since we work specifically with the brand and all its affiliated products, we analyze the needs gathered through the workshops. So there is the functional analysis part where we describe the proposed solution in non-functional terms, that is to say, quite simply, imagine if there was someone who had acne and was looking for a solution to get rid of it. The functional analysis would be to say, "Here, it will be a cream that, once applied, will clear up the pimples." And then after that, we write up everything related to the technical analysis, meaning we explain in technical terms how we are going to configure the solution. In our case, it's configuration since it's software implementation. But if we go back to the cream

example, the day cream... There, we would detail, for example, the entire composition of the cream, what elements actually make up the cream, the packaging, and all those things. That is my job. Generally speaking, the job of a business consultant applies to a huge number of topics; here, it's in the context of IT.

OK, so I have two questions regarding what you said. I'm not sure I heard what type... what type of clients you have. Sorry, what type of clients you deal with?

All types of clients, all company sizes, though slightly fewer small businesses. So it's mostly medium and large enterprises, and from all sectors. It ranges from management... management... we have a lot of hospitals, universities, really all types of clients. Actually, we target clients who have a need for project management or document collaboration in general. Collaboration and knowledge sharing on documents that need to be a bit more dynamic. It goes a little further than just storing documents in a SharePoint. But essentially, we cater to all types of companies, of all sizes, and across different sectors.

You mentioned that most of the time, you produce two types of analysis: functional and non-functional, or technical and non-technical. When you produce these kinds of analyses, do you share with the client how you obtained your results, or not at all?

The analyses are subject to client validation, since they will have to review them to make sure that what we put in the section describing their needs actually corresponds to what they explained. To validate that they have correctly understood the functional analysis and that it will indeed meet their needs. The technical part is for their information. Since we are consultants, they generally don't have the in-house expertise, so very few are familiar with the products when we implement them for the first time. So, it's subject to validation. Beyond that, regarding how we obtain information, that comes down to our business expertise, the fact that we know the product and can advise them as best as possible on the solution to consider. To guide our thinking, besides the certifications we took right at the beginning when we joined the company to become experts on the product and the brand, there is also the various user documentation available for the products. It's like with Microsoft: when you have a question, there are always FAQs explaining how to use the product, how to configure certain things. That also exists for the products we implement and for all the marketplace applications that exist to enhance the product's features. So, it's a combination of what we know about the product because we learned to use it, were trained to use it, and passed exams to upskill on it, alongside the various resources we find that are written by the publisher itself. And then, and I think this

is where your thesis topic comes in, AI helps too. OK, when we have a question, just asking a simple question to the AI, at least to guide us to the right sources, and so on.

Well, I was actually just about to follow up with my main question: which AI tool, specifically generative AI? We're talking about ChatGPT, Gemini, Copilot. What do you use or what are you authorized to use at [Company Name]?

Well, internally, we mainly have the Microsoft suite, so it's mostly Copilot. ChatGPT would be more for private use. And as for permissions, it's always about being careful with data confidentiality. So, never any confidential information related to the internal information of our clients, nor our own. So no client names, no information we aren't supposed to share. In fact, any information that would be part of a contract or a confidentiality clause must never end up in the chat with Copilot in this case, nor any information that would allow the creation of information, like creating new elements.

Ok, and so Copilot is more authorized at your company because of that confidentiality aspect, or does that have nothing to do with it compared to ChatGPT?

No, I would say it's more about integration with all the tools we already have internally. It's because it's an AI available with the Microsoft suite, which is obviously configurable and can be controlled at the enterprise level, since the administrators aren't the individual users themselves. A regular employee in the company generally doesn't have admin rights. So, they [the IT department] can define rights at different levels and restrict the end user's rights. Copilot integrates into this whole Microsoft suite, and I think that's actually what justified using that chat over another one.

Alright. And so for you personally, when you are on assignments, in what context do you use the generative AI from Copilot?

I would say in several contexts. The first is simply when I'm writing an analysis: for rephrasing, writing assistance, proofreading, synonyms, and sometimes translations, even though I admit that tools like DeepL are sometimes still more accurate. The thing is, I find AI to be very good at contextual translation. Yes. There you go. Whereas DeepL provides pure, literal translation but doesn't take into account the context of the sentence and all the explanation I might have provided alongside it. Besides that, as I mentioned, for finding information, because user documentation can be hard to navigate since there is so much of it. There are also blog sites, which we call community sites, where all the users of the products can write blog posts that

people can reply to. It's good for finding a specific topic that might not be in the official user documentation. However, and this is the problem with AI, you have to go and verify all the sources because it tends to present results as absolute truth, when it might actually just be a user's research proposition that hasn't been verified by the brand or by the actual functioning of the tool.

I was actually just getting to that, it was my next question. So, do you use it more like an Assistant, let's put it that way. Do you use it more as an assistant, or is it something that really takes up a large part of your work? Meaning, if you use it as an assistant, do you validate the analysis and the inputs that the AI produces, or not at all?

Personally, I think it's a tool that supports my work, because I am the one who has all the product knowledge. So, it's to go a bit further or to help me with research, that's generally when I call on the AI, and then everything must be verified. Because precisely, it is not capable of interpreting the way a person would, nor can it propose the advanced configurations that we would need.

So exactly on that note, regarding expertise and skill sets: what are you able to do that the AI cannot do at that level, in terms of interpretation, etc.?

I think I would say... well, it has access to all the documentation. The thing is, it cannot identify on its own: "Hey, this configuration works for this specific case, so we need to set up this configuration in this way." It will be able to produce a list of all possible configurations. It will be able to produce a list of everything it finds on the internet recommending one configuration over another. But it will not be able to produce the configuration itself. When we document the configuration, we go through each parameter and indicate: "We configured this parameter this way, for this reason." The AI won't be able to do that as precisely as a consultant. Because it will rely on everything it finds in its databases, in all the sources it has access to, but it doesn't interpret it perfectly with regard to the client's context or what we know about the product. So it might propose things that aren't suitable, either because it didn't quite understand the question and suggests things that have nothing to do with the type of instance, like whether it's cloud-based or if we are on a private on-premise data center server. So it will suggest things that are not applicable to the client's case, or things that people do but are not applicable given the client's size or the products the client wants to implement. And besides that, all the interpretation, it doesn't know how to do that. It can proofread and rephrase, but it remains very generic actually. It's not precise enough that if we asked it, for example, to provide a complete

configuration for the Jira or Jira Service Management product, it wouldn't go point by point giving all the parameters to configure in this or that way based on the client's needs.

Ok, so it really lacks that interpretation side, the human side, that is needed to properly understand the client's needs.

Exactly, there is a limit to what it can understand and what it can extrapolate, let's say. Meaning, it can only reinterpret things written on websites, in documentation, etc., and it cannot apply them to the client's case precisely by saying, "Here, we will configure this, this, and this parameter, but then we will have to change the name of this because it's not the client's name." It will base itself on very generic things, it will propose results. At least, that's my experience with it. It won't say, "Careful, in menu 3 named 'implementation of such configuration,' there is another drop-down menu. You should choose this option instead because of this or that." It will only know how to say that if someone has already mentioned it on a forum for a very specific case.

OK, so in your case, you use GenAI mostly for rephrasing, reports, not for the performative tasks that are normally included in your skill sets. If I understand correctly?

For me anyway, since it's very limited in what it can propose on its own and since it doesn't know how to do an analysis from A to Z, it's more... well, I have a bullet-point outline that I'll have it rephrase more comprehensively, or for meeting notes, I put in a few elements and say, "Here are the meeting notes, without confidential information obviously. Compare the meeting notes for me and see if I haven't forgotten anything in the analysis for this point."

But you have to be very precise because otherwise it mixes everything up, because it brings in elements that have nothing to do with the way we structure our analyses. We work based on user stories. So the way we break down user stories, I don't know if you've heard the term before. Ok, that's fine then. But so, you always have the description of the feature, the role it is intended for, and the objective of that feature. There are times when it struggles a bit, at least Copilot does, to properly distinguish each of these elements and to focus solely on that theme when you give it raw meeting notes.

Ok, and so, when you use AI for the few tasks you do, do you inform the clients about it or not at all?

There is no formal disclosure as such. Now, I don't particularly hide it either. During my initial contact with them, I usually mention that it's generally for rephrasing, but there is nothing written as such in the analyses or in the contracts.

And when you specify this kind of thing directly, in the eyes of the client, doesn't that call you into question as a consultant?

As long as it is well explained and I explain exactly what I used it for or not, I've never had any issues so far. But I think since we are in the IT world, many people are quite up to date with what is being done nowadays, and I believe most of them already use AI a lot in general. So I don't notice any pushback.

That said, if we were to move into much more extensive use, I think it would be worth mentioning the use of AI more clearly, and especially *how* it was used. Because if it does my job for me, it's clear that my added value in all this is called into question. Here, since I create everything myself, since it's just the final touch, let's say the final polish, I've never had any issues on that level.

Ok, and precisely on that note, this transparency you have with your clients is fine for you because it's not like you're doing your whole job with GenAI, but for those who use it more extensively, do you think transparency is a good thing or not? Specifying things, I mean.

It is important, but it's not easy to implement because the line is very thin regarding what value the consultant adds. Generally, consulting exists because you call upon someone for skills you don't have; so from the moment the consultant makes extensive use of AI, it raises the question for the client: "Wait, couldn't I just do it myself then, if it's AI?" So the issue is delicate.

Yes. Have you ever observed this with your colleagues, then? Maybe you don't necessarily want to answer the question, but specifically cases where their added value is called into question because they use AI a bit "too much," quote-unquote, compared to what they should basically perform as a consultant.

That problem is heavily mitigated within the team in the sense that there is always a peer review of everything we produce. So in principle, at least for the analyses, since they are reviewed by a colleague, the colleague will have gone over it and corrected it or pointed out if there is any issue. So we never release things that are borderline or that don't meet the quality standard we

expect for our analyses vis-à-vis the client. So I think we generally manage to catch this problem before it reaches the client.

I don't have an example of someone using it to such an extent that we say they aren't doing their job anymore. Not in my team anyway, but I can't speak for other teams. I don't think so, in any case. The problem with AI is that it shows when you use it too much. It relies on very generic phrasing... The coherence from one sentence to the next isn't always maintained, the tone of voice changes. And when you know the person a little, it becomes obvious very quickly when their writing is vastly different from what they usually produce. So there's nothing that jumps out at us like that.

This is a bit of a silly question, but I'll follow up on it right after. Is trust important in client-consultant relationships?

It is. Always, of course. First of all, because the client must trust that we are going to produce quality work for them. Besides that, they need to know that we have the skills to do what we do. And so, it's also a guarantee of quality to be able to prove that we have the skills to do the work, either based on certifications or client experience. Because sometimes there are no certifications for the implementation of a specific product, but through repeatedly implementing it for clients and being in contact with the publisher and support, you acquire mastery of that product and thus become the expert in the field.

And then, on the other side, we need to have a point of contact on their end who is available, who is open-minded, because we generally question their internal processes. With the implementation of a tool, there is a whole upheaval that takes place, there is a lot of change management to handle. So that kind of client-consultant relationship is important, as is transparency in the relationship. And then, expectations and the act of communicating them and meeting them, that's super important.

So precisely, how could the abusive use, or non-use, of GenAI disrupt this relationship of trust between the consultant and their... [client]?

[Their] client? It calls into question the quality of the consultant and their work in general if it's the consultant who is misusing AI. It calls into question the reputation of the company, of the consultant, the quality of the work, and the competence of the consultant. Because, is the consultant really competent if all the work is done by AI? The line between using AI too much and using it *in* your work can be crossed very quickly if you don't do it right, I think. And so,

that damages the client relationship because the client is bringing in someone external. Consulting comes at a certain price because it's a service, you pay for a service. From the moment you pay for a service and feel you aren't getting what's fair, it causes a lot of issues in maintaining a lasting relationship with the client.

When, for example, there have been cases of companies getting caught submitting reports full of hallucinations because they used AI, in those cases, who is held responsible? The individual consultant, their manager, the entire company?

That depends on the company. The thing is, vis-à-vis the client, it's both the consultant and the company. The credibility of the company and the consultant working internally for the client are at stake. I admit that since we haven't had this happen, it's not a question I've really asked myself, and I think it comes down to everyone facing their own values, whether they feel responsible or not. For us, responsibility is actually one of our company's core values. It's among the values listed on our website that we try to implement daily with our clients and colleagues. So, taking responsibility for the situation is important. I would say the consultant, first. But alongside that, the managers or the people who have to review the consultant's work as such. And then at the company level, because it remains a client of the company. The responsibility is multi-level.

OK, alright. So, we went faster than I thought, but at the same time, you were very clear. So much the better. Yes, so you answered a few of my questions in one go. We're going to jump straight into the future. Since you were saying earlier that clients pay for a service you provide, but now most of them have access to GenAI. In the end, aren't there times when they are paying for nothing, since they could do it themselves? I would say no, because...

A huge part of our job is guidance. We sit down next to the client, we teach them how to use the tool, we explain why we chose one element over another, we review what they currently have in place, we advise them on the best course of action. And those are things, to a certain extent, that AI cannot do. At least not right now; in a few years, I don't know where it will be, but the human component and the change management on the human side are important, and I think AI won't be able to replace that.

And so that is a part of the expertise as a consultant that AI cannot yet replace for now.

It's really the whole package of having someone external come in to cast an objective eye on ways of doing things, internal processes, and tools that are perhaps not being used well or to their full potential. AI doesn't have that perspective because there is a massive amount of information in a company, in its processes, and in how it operates that is informal meaning it's not written down anywhere. It's just someone who knows how things work in their head but doesn't write it down.

And so? The bulk of our work is being able to formalize these ways of working to turn them into a process that can be shared by many people, things we put into Confluence (the tool that allows for this collaboration and knowledge management), so, putting processes, operating methods, and organizational charts in writing. And when it's in people's heads, AI doesn't have access to it.

So it cannot organize meetings, which sometimes involve debates and long discussions on how to align internally, how so-and-so operates, how to reorganize when the person in charge of a specific activity isn't there. Those are things AI can't do, because it's purely human contact. Yes, once we've taken the notes and made contact with those people and shed some light on all the informal elements happening in a company, inevitably, by feeding the meeting notes to an AI, we can reorganize them by themes and identify certain points. But the AI doesn't know how to go and extract the information from the people themselves. Because that is purely about human relationships.

I'm going to make a comparison between now and the future. For now, do you find that AI is an assistant or a threat, and vice versa for the future?

The thing is, like with everything in IT and in general, it remains a tool to me. And the thing with tools is that it depends on what you do with them. If you take a fork and use it to stab someone, it's a threat. But your fork, if you use it to eat... If you start stabbing it into your laundry, well, maybe you just don't know how to use your fork properly.

It's the same with AI: if you use it well, sparingly, and in accordance with your company's values, and your own values too, because ethical and ecological values come into play here. Some people don't want to use AI, or at least prefer AIs that are a bit more ecological or sustainable. So, there are many things that make this question complicated. To me, it remains a tool, it remains a support, it remains an assistant; it shouldn't become the thing doing the actual work. Not in my job, anyway.

And then in the future, maybe they will evolve enough to be able to threaten certain parts of the job, but I think there is a deep expertise that remains profoundly human. So I don't see it as a threat. What is threatening is the *misuse* we can make of it. So sharing confidential information, that is dangerous, because information that ends up on the internet can be used by anyone. And since we usually deal with sensitive internal corporate governance topics, it can be used for malicious purposes. For me, I think the threat comes from how we use AI, not from the AI itself or what it is capable of doing.

My next question would be: does GenAI create barriers to entry for employment, for people who want to become consultants in the future?

I don't think it creates them; I think it removes a lot of them. Because when you are lost on a topic, you can ask it to help, there are very typical prompts for this, they don't come to mind exactly right now, but saying "identify the 20 most important topics to know about this product, or the 20 most important features," and I think that can clear up a lot of the ambiguity.

It removes a lot of friction in how you draft an email, an analysis, or anything else. I no longer waste 20 minutes wondering how to phrase a specific email in a professional and polite way. Once my draft is written, I ask it to correct the tone or make sure it aligns with a specific writing style. And the email is written in a fraction of the time it might have taken. The same goes for analyses. So I don't think it puts up barriers.

On the other hand, I think we have to remain vigilant. We must continue to have people review the work to ensure we don't leave AI artifacts behind. You know, things like: *"Here is a list of this and that. Do you also want me to write this in a more polite form?"* You definitely can't leave that in a report. So if we use AI, we must remain vigilant about how we used it and ensure the professional aspect is fully maintained. But again, that comes down to responsibility. It is up to each user to use the tool correctly and guarantee to their clients and other stakeholders that the job was done well. For me, that falls under the quality of the work we deliver, making sure we don't hand the client something sloppy.

Ok. And so, we won't lie to ourselves, GenAI is evolving and will continue to evolve. In your opinion, what are the qualities and skills that a consultant in the 2030s and beyond should have to remain relevant in the face of GenAI?

I think it's a bit of the same kind of answer as before. I think we will need to focus on human values and human interaction. We will have to develop soft skills: everything related to

communication, to pedagogy. What is certain, in my field anyway, which is change management, ITSM (Incident and Request Management), etc., for ticketing, is really having a deep knowledge of management subjects, leadership, and management in general to be better equipped to handle human relationships and navigate them. Because that is the core component of our job.

Even when we recommend something that is perfectly suited to the client's context, sometimes for purely internal, political, or other reasons, they just don't want to implement the solution. I think AI will hit a wall by just proposing things "as is," whereas a human will be able to make concessions and compromises. I don't think AI is capable of doing that. It might propose alternative solutions, but it lacks that concept of compromise or finding a win-win situation. So the key for me is soft skills, and especially interpersonal and purely human skills, truly interacting with people and guiding them.

Ok. And in the next five, even ten years, maybe five is a bit too soon, but do you see the consulting profession evolving or changing due to this GenAI context? Are we going to transition from being an "expert" to someone who just validates AI outputs, or will the job not necessarily change that much?

I think it will evolve, yes, it will. Until a few years ago, there was no AI, and we couldn't ask it to rephrase an email or dig through documentation to provide a product comparison, for example. But to go from there to becoming just simple verifiers of what the AI produces... I don't think it will come to that, at least not in 5 to 10 years. But I could be wrong, I don't have a crystal ball for that. I think I know enough about AI to anticipate the future fairly well. But I believe the job relies so heavily on social interactions that a huge part of the work will remain our own doing.

Interview ERP02

As I mentioned in my email, this interview is for my thesis topic, which is the use of generative AI in the consulting field, specifically regarding how much it will influence client-consultant relationships in terms of human expertise, trust, and transparency. I have a bunch of little questions for you, and you just need to answer as spontaneously as possible.

Perfect, let's go.

Hello, can you introduce yourself? Where do you work, how long have you been there? What kind of projects do you manage?

My name is [Consultant Name], I am 26 years old, I have been working for 3 years, almost 4, 3 and a half years actually, at [Company Name]. I am a consultant, a business consultant for an ERP called Business Central. An ERP is software that allows you to manage all the operations of a company, but electronically. Managing purchasing, sales, production, and generally everything related to accounting to automate everything a bit. I am on different projects and different themes. I have a client in pharmaceuticals, one in construction, one in window frames, and another one in gas detection and those kinds of trades, those kinds of functions. Generally, in these kinds of projects, we have two types: implementation projects and support projects. For an implementation project, we will do workshops and some analyses to set up the solution. At the end of this implementation, inevitably, there are always little requests, questions here and there. The implementation project becomes a support project. Then, once in a while, a client asks questions because something isn't working, or they would like to improve a process or things like that. We are there to help them and continue the process when needed. As a very good summary, that's it.

As a business consultant, what are the skill sets, the hard skills, and the soft skills that shape your expertise?

In terms of soft skills, you need to have quite a bit of autonomy because it's an environment where you are asked to be independent very quickly. We do get quite a bit of training, but you have to be ready quickly. In other words, you need to know how to communicate and how to be very patient because sometimes we have clients who are resistant to change, and we have to keep training them since we have to teach them the program. Since they don't want to change programs, it is sometimes complicated. You have to stay very patient and re-explain the different concepts of the new program several times. You need a lot of pedagogy since we give

training sessions, we are a bit like corporate teachers in the end. Quite a bit of rigor, and you also need to know how to say no, because otherwise they put too much workload on you, but that's in any job I would say. Knowing how to say no and being able to step back from the situation, not taking things to heart. You need to maintain a certain distance from the situation. Otherwise, you end up in burnout, which isn't very practical. For hard skills, it's a good mastery of the programs, and therefore the Business Central ERP program, knowing how to make PowerPoints, knowing how to manage Excel because we do data imports, you need to be able to work with Excel. Frankly, working with Outlook, who doesn't know how to do that now. Knowing how to do all the reporting, documentation, we have lots of various documentation, procedures to write, that is also something you need to master.

Sorry to interrupt. When you talk about implementation and the tools you need to master, could you clarify, please?

Could you repeat the question because it cut out for a moment?

Sorry. I was saying, you mention implementation, are you talking about mastering the tools that allow you to build the ERP, or is it something else entirely?

No, more like knowing what is inside the ERP. Yes, and in our case, we don't create the ERP, it's really just an existing ERP from Microsoft. We install that ERP and then we train the clients on this ERP and help them understand the different features inside it. We have a developer side that allows for adding features within the program, but that's a part that isn't directly tied... I mean, we work in collaboration, but we aren't the ones doing the development.

I would add one last soft skill: you have to continuously adapt because the program is constantly evolving. There are always updates and always new things. You have to constantly retrain and re-learn a bit of everything.

In your opinion, what are the right ingredients to establish a good relationship between the consultant and their client, especially in terms of trust?

A consultant who plays by the rules would say to keep a distance and not commit too quickly to certain things or that sort of thing. My personal opinion is to remain honest, to admit when you've messed up, you've messed up and it's not a big deal because you remain human, and if the client understands that, it's normal. Really taking your time and sometimes putting yourself in the client's shoes to understand their situation and ensure they feel understood and their needs are understood. As long as you are honest, kind, and have good manners, like with anyone, it

generally goes well. A touch of humor is always nice and always useful for building a trusting relationship with clients.

Do you find that this human aspect is fundamental or rather secondary in the client-consultant relationship?

I think it is quite fundamental. Up to now, that is what makes clients feel good and keeps them with us.

When you implement ERPs, is it a standardized solution in the sense that it will be the same process for your clients, or is it completely different?

We have a baseline that is common to all clients. The way we implement the program is common to almost all clients. We go through a needs analysis phase, then a phase where we draft a document to make the various needs factual, and that becomes the contract between them and us.

Then we have a phase where we develop various additional features because the standard, the bare ERP, does not necessarily suit all clients. For example... did it glitch or not?

No, no, we're good. Perfect. For example, I have a client who works in spirits. They need excise duties, obviously, since it is a different tax from the normal tax. Excise duties are not managed in Business Central, our Microsoft ERP. Given that it is not managed, it still needs to be developed since it is part of their daily routine. There, we have the developer team that develops the various additional needs that are not in the standard. If we already have pre-made solutions, what we call add-ons, ISVs, extensions, it depends on the name you give them, we install those add-ons and it adds a whole range of features without us having to do any development since they are already made. We have an add-on called Continia, for example, which manages OCR. Basically, it's the automatic scanning of invoices to avoid manual data entry. Users scan and everything is already encoded in Business Central. For example.

After that, we have the whole training part. That is also common to all clients. We continue to give training sessions to teach the program, and then we transition to support. The area where it changes a lot is really everything related to development. Depending on that, the training adapts a bit. Some clients take all the modules in Business Central, and others do not take everything. For instance, my client in pharmaceuticals produces machines, they need the production module. A module is a part of Business Central. There are other clients who do not produce anything, this production module is useless for them and we don't implement it for

them. In the standard, we have some modules that are sometimes implemented, others not. Then we have the development that allows us to add features or not depending on the client's needs.

Last question for this more general part. During the training sessions you give to your clients, do you explain the whole process you went through to implement the ERPs, or is it not necessary and not at all what you do in the training sessions?

In the training sessions, we explain what is inside the ERP itself, how to create an invoice, how to create an order, how to create a delivery note, or how to enter a new client or a new supplier. We explain that to them. We also explain the different developments that were made. For the company with the alcohol, we tell them: here are the excise duties, how it works, it's like this, you have to fill it out like this to get that. We explain more of the functional side, as we call it. Really what allows them to work every day. Personally, I sometimes explain the technical side when I see that the client is advanced and mature enough to understand what is happening inside. Because there are some, they are 60 years old, they don't have the patience to understand the technical side, they are not people who have a feel for computers, it's useless to explain it to them. The only thing they want is to know what they have to do and what they have to click on. Then we have the young people, they want to know everything, how it works, the smallest detail. There I take the time to explain so they can understand because we have to work with all generations, you have to know how to adapt.

Now, I will get to the heart of the matter, namely the use of generative AI. The first question is: what tools are you authorized to use within [Company Name]? Do you use any when you do your job as a business consultant?

Yes, I use Copilot since it is part of Microsoft and we are in the Microsoft ecosystem. We can also use ChatGPT or other programs, but Copilot ensures confidentiality because it would be silly to disclose information from various clients, since we have a guarantee that we must not disclose all that. Copilot, theoretically, ensures this security and confidentiality, and we work a lot with Copilot, especially to draft emails, reports, those kinds of things.

You don't necessarily use Copilot to help you implement the ERP or analyze client needs or other things?

It has happened sometimes that I ask it a question to know how it would do certain things, how it would organize a training session, how a workshop could be held, things that are rather functional outside the ERP. For now, when you ask Copilot and ChatGPT a question to know exactly something in Business Central, they will start talking to you about all the ERPs in the

world. At the moment, they are not very precise, they invent features. You try to look for the features and you can't find them. Everything related to Business Central is not yet optimal.

There are a lot of hallucinations that don't necessarily allow you to use it for those tasks?

Exactly. It is really more for administrative tasks. There, they help us a lot, well I would say, for me at least, they help me a lot. I don't know about the others. There are some at our company who never use it because they are used to working their way, they don't see the point of using Copilot or ChatGPT.

When you use Copilot or ChatGPT, do you have to specify it to the client, or is it not necessary?

For now, it is not necessary and we don't care, for the moment.

They are not necessarily interested in that part of the ERP development, they don't care if AI contributed or not?

No, for now. There is one who is very interested in artificial intelligence and AI. But the fact that I use it, he couldn't care less. On the other hand, in Business Central, our ERP, we have Copilot assistants that allow you to create different functions. But again, it remains a bit sub-optimal, it's messy, but you can see these are the early stages, it is not often used by the client. Generally, AI in Business Central is used little or not at all.

And then, in your opinion, the fact of not telling clients to what extent... well, having or not having transparency on that, is it a good thing or not? In the sense that some people in the literature say it might not be a good thing to be too transparent, because it can call into question your value as a consultant. Do you agree with that or not?

I think that if I tell them whether I write my emails with Copilot or not, they won't necessarily... I mean, it won't change anything about the situation. For my knowledge of Business Central and my way of working with Business Central, I don't use ChatGPT. Generally, for everything that is truly the knowledge of my profession, Copilot is not involved. Whether I say it or not won't necessarily change anything. In my opinion, they probably suspect that when my email is extremely well-structured with lots of bullet points, ChatGPT plays a part. They have never said anything about it. When we don't have time, when there are plenty of other tasks, and we have to meet various deadlines, we use it to make the most of the time we have. I don't think it would be dramatic if I say it, as long as it stays at the level of email management, report

management, etc. If, on the other hand, I start telling them, "Yes, every time you ask me a question about Business Central, I ask ChatGPT," then... well.

Exactly, since you are paid as consultants to provide a service, to do something for the client, could clients stop paying you and do what you do themselves using AI?

For now, no. I was actually confronted by one of my clients who started telling me, "Yes, but with ChatGPT, it estimated 10 minutes of development for me, I don't see why it takes you 2 hours." The thing is that ChatGPT can create whatever you want, but it isn't necessarily true, it isn't necessarily correct yet for the moment. For everything related to development, right now ChatGPT doesn't know how to take a piece of code and put it into Business Central, that is impossible. Up until now, we don't have the tools to do it. Their code is only half correct because, once again, it takes the knowledge it has, but sometimes it mixes different codes, making it a bit complicated. It can give you advice, but developers remain very important for the moment as well. Consultants remain very important for the moment because ChatGPT does not have the ability to step back from the situation, focus solely on Business Central, and provide useful information. It can help with a lot of little things, but for the moment we are protected in that regard. I am not saying... perhaps in several years, when it is truly high-performing and knows Business Central perfectly by heart, then we can start worrying, but right now we are good.

So you really see it more as a tool than as a competitor?

Yes, exactly, it is really not a competitor for the moment.

Which dimension of the consulting profession, in your opinion, is not yet replaceable by generative AI? What are the aspects of your job that cannot be replaced?

First of all, not in my job directly, but everything related to development, that is not replaceable for the moment. Even if it is not part of my job as such, it is part of the process. I would say... in the future, or even now... what is not replaceable... it is true that it is a complicated question. Let's think about it. For example, if we take the beginning of a project... just being there in person, being in front of the client, that is not replaceable. We will say that is the first thing. Otherwise, for the beginning of a project, analyzing the needs, overall we could be replaced if ChatGPT or Copilot ask relevant questions, they could replace us on that. Doing reports, they can replace us on that. Doing development, for the moment it is complicated. Giving training sessions and giving advice, taking the time to re-explain... technically they could replace us, but it still remains sometimes more pleasant to have human contact rather than just a program

that says step 1, click on this, step 2, click on this, step 3, click on this. I think we are replaceable, but I think losing the human side would be very sad. Because there isn't much that is complicated in what we do.

Regarding all this context, do clients ask you for more transparency on how you reached a conclusion, an output, considering that you might have used AI? Do they ask for more explanations or not at all?

That still depends heavily on the different clients. There are some who just want the solution and don't want to have the explanation, and in that case, there is no need to explain everything. There are others who are very meticulous, since we bill them for the hours we spend, and they want to know exactly what we did, the different tests we carried out, and understand what we did in our resolution of the problem. In that case, we have to explain everything. In this kind of situation where there is a blocking issue and we have to solve something within Business Central, personally I never use Copilot, so even if they ask me to justify, I would tell them that the only time I use Copilot is to write the email. Once again, here it doesn't cause anything negative regarding the client. Yes, there are times when clients ask questions to know what was done, why, how, and to have a few details to understand why they are going to pay €2000 for an answer. Because consulting is expensive.

Do you have the impression that the barriers of the consulting profession are collapsing with this generative AI context? Is there a risk that consultants will become more like verifiers of what AI produces rather than creators of solutions?

For now, no. In five, ten years, I think the question could be raised again. I think that at some point, all IT professions will go through this, meaning becoming AI verifiers rather than... I don't have the right word... executors.

And then, last question, well, you've already answered a lot at the same time. In your opinion, in the future, maybe in five, ten years, what would be the added value of a consultant to remain relevant against generative AI?

Knowing how to manage all the ChatGPT or Copilot prompts, and therefore really mastering Copilot as a true tool, as an assistant, as something with major added value. Just knowing how to manage that will be a big point. And showing that human contact remains important. I'm thinking... could you repeat the question to be sure?

Yes, of course. In the future, what would be the added value of a consultant to remain relevant against generative AI?

Knowing how to master it and keeping the human side. Overall, there isn't much more. I admit that at some point, in ten years, we might not need consultants anymore. It's a bit like the accounting profession in the end. Before, accountants recorded invoices; now we have OCR and ERPs that do it, and becoming an accountant today is becoming a bit useless. It's the chartered accountants who remain to handle financial management, account auditing, those kinds of things. I think that if we continue down this path, a consultant won't be useful as such anymore, but "expert-consultant" could become a new profession. That would be verifying that Copilot is used correctly, that the tools implemented in the company are used correctly, etc. Verifying that procedures are done well, properly implemented, and perhaps optimizing all that, the management of Business Central, those kinds of things. In my opinion, it's going to evolve like that.

I have one last question, sorry. If consultants become AI verifiers, won't that drastically reduce the size of the teams? For example, if today there are 5 or 6 of you, would it go down to just one person?

We have 25 consultants and 5 developers in our team. [Company Name] has 2000 people, which is another thing entirely. And [Company Name] is part of the [Parent Company Name] group, so even more people. But if I just take my team, and we move from a consultant role to an expert-consultant role, I think we could reduce by about twenty people. Developers, in my opinion, will stay as long as Business Central doesn't perfectly know the code and have access to everything. But at the consultant level, it will depend on the number of clients and projects, because one person cannot manage 20 clients in a week. There will always be a need for people, but we will probably be able to manage more clients than today. Today, we manage about 6-7 projects each, and in the future, we could manage more.

I have everything I need. Thank you very much for your time, you were very insightful in your answers.

I'm sorry.

No, on the contrary, it means the questions were well formulated. Thank you very much again.

Interview ERP03

Right, so as I mentioned in my email, my thesis focuses on the interaction between the consulting world and GenAI, specifically regarding client-consultant relationships, particularly focusing on human expertise, trust, and transparency. So I'm going to ask several questions. Can we keep it informal?

Yes, of course, we can keep it informal, no problem.

I'm going to ask several questions and you just need to answer as spontaneously as possible.

Sounds good, sounds good. But just to clarify things a bit: this is about the daily use of AI, in my life as a consultant, right? In relation to the different projects I have depending on my clients? OK, makes sense, no problem.

That's it. So first question: can you describe your role? Where do you work, in which consulting firm, and what is the nature of your projects?

OK. So, I am an ERP business consultant. I work at [Company Name]. It's a company 100% in the IT sector. We have several units: we do cybersecurity, we do ERP implementation, we really do everything related to IT. And so, my specific role, as an ERP business consultant... You see, for example Odoo, Microsoft... they have management software that we implement for clients to make their lives easier. Meaning facilitating their accounting management, inventory management, warehouse management, production, etc. So for us, at [Company Name], we have three teams: an Odoo team, a Business Central team (which is a Microsoft ERP mostly dedicated to SMEs), and then we have Finance and Operations, so that's the team I'm in, linked to large enterprises. In my team, we are divided into three sectors: we have the warehouse/production part, which is the one I'm in (warehouse management, inventory management, production management, etc.); we have the supply chain part (purchasing-sales), and we have the finance part. So, what is my daily job going to be? We sign projects with large enterprises, and the projects are divided into several phases. The first phase is the needs analysis. We meet the clients, who we call Key Users. They explain their way of working to us. And based on that, we design a solution while respecting the tool's standard, so that it matches their way of working. Once this analysis is done, we start showing them how the system works and how we tried to configure it according to their needs. And there, there's a lot of back-and-forth: "no, that doesn't work like that," "that's good," "that's not as good." And actually, step by

step, we build the solution so that it meets 100% of the client's needs. The goal is that in the end, they can use the tool autonomously in their daily routine. That's the general idea of my job. And so there you go: these are very long projects that generally last between one and two years.

I noted that among your clients, you have SMEs and large enterprises. What is the typical profile of your clients? Is there a specific industry or not at all?

No, exactly. As I was saying, there are SMEs and large enterprises. But in my team, it's exclusively large enterprises. And there is no specific sector. I could easily have a client... I'll give you examples: I have a client that makes waffles, so more in the food industry. And I have another client that makes large electronic cabinets for clients like Proximus. It has absolutely nothing to do with each other. And that's why this job is very varied: the way of producing, managing stock, etc., is going to be extremely different between a client in food and another in electronics, for example. So there you go, we don't have a specific sector: we take all types of clients and all types of companies.

In your opinion, precisely regarding the clients, what are the best ingredients for a good relationship between the consultant, the consulting firm, and the client?

For me, I think it's availability. Why do I say that? Because clients want everything, right away, all the time. And as a consultant, you have to find the right balance: making them understand that we are there as support, that we are available, etc., but not too much, because otherwise they will exaggerate a bit. So for me, the first thing to have is to make yourself available. You absolutely must put client satisfaction first, in the foreground. For me, that's really the most important thing. And if the client sees that we make ourselves available and do everything for them to be satisfied, a relationship of trust will automatically be established. After that, obviously, there is also the human side. Meaning that when I arrive at the client's site, I am always smiling, I try to be in a good mood, because it's important. The client shouldn't feel like I can't be bothered to be there. It needs to be an exchange, not just client-consultant, but rather two people trying to create something for the good of both companies in the end.

So now I will get to the heart of the matter. What GenAI tools, like ChatGPT, Copilot, Claude, or even internal AIs, are you authorized to use?

Listen, I use two AIs and I'm starting to use a third, and it's exactly the three you mentioned. So the one I used to use a lot before was ChatGPT. I still use it a lot now, but a bit less, and I think we'll talk about that later. So ChatGPT and Copilot are the two I use the most. Why do I

use Copilot a lot? Because it's a Microsoft tool. And since it's a Microsoft tool, and the tool I implement for clients is also a Microsoft tool, it's centralized. And I realized that Copilot was much more precise in its answers than ChatGPT. ChatGPT always gives fairly coherent answers, but Copilot is a bit better. And now, I'm starting to use Claude, but very lightly.

Is there also a question of confidentiality when you use GenAI?

Yes, exactly. And that's the second reason why I prefer using Copilot over ChatGPT: with ChatGPT, it's not written in black and white that the data you send it respects confidentiality rules, whereas with Microsoft, it is. And since I use these intelligences to do, for example, data synthesis from notes I've taken with clients, I prefer going through Copilot because I know there is respect for data. And since it's somewhat private data between the client and us, I really prefer using Copilot for that, because ChatGPT isn't 100% legit for those kinds of things. That's why I prefer using Copilot now.

You just mentioned one of your uses for AI, namely data synthesis. What else do you use it for on a daily basis?

Yes. I use a tool that is heavily documented on the internet. And so, when I can't find certain information, I ask the AI by sending it, for example, URL links of that documentation, so it can explain the practical functioning to me. Because in the use of an ERP, there is the theoretical way, so the way we are going to explain it to the client, but then there is the practical way it applies. And so what I like to do is ask the AI to create real and concrete examples of using certain features of the tool. Then, as I told you, I do a huge amount of data synthesis, especially from Excel files. We sometimes take notes in Excel files, and so I ask it to draft my meeting minutes, meaning a summary of the meeting we had with the client, based on the Excel file where I took my notes. And then I always review it. And sometimes also, simply to send long emails, and to make it more professional, because I'm not a good writer, I ask the artificial intelligence to rewrite what I wrote in a slightly more professional way.

Now for a more general question about the consulting profession. In your branch of consulting, what characterizes a consultant in terms of skills, both hard skills and soft skills?

In terms of soft skills, we've already talked about it a bit. For me, what is essential is to be autonomous and flexible. I even think that flexibility is probably the most important skill for a consultant. You also need to be somewhat extroverted, and above all, have self-confidence. Because sometimes, when you arrive at a client's site, they will challenge you a lot. Clients are

spending money on a project, so they want to ensure that the person in front of them is competent. They ask a lot of questions, they test you. And in that context, you have to be able to speak up, not hesitate, and explain things clearly. So self-confidence is really important. I think in terms of soft skills, that's already a good foundation. Well, in itself, the D365 tool, I didn't know it at all initially. And honestly, if you've never worked with this tool, it's impossible to really know it. It's something you learn on the job. But in terms of hard skills that I've needed here, I would say it's mostly the Microsoft suite: Excel, Word, and PowerPoint. And what I also needed is Power BI. Power BI is a data visualization tool. And so, we also do that type of work. For me, having already seen it at university helped me a lot. So, in summary, regarding hard skills, the most useful ones for me here are really the Office suite and Power BI. What is also interesting is having some coding knowledge, even very basic. For example, I never code anymore today, but I took Python classes. It's a fairly simple language and it mostly helps me understand the developers. Because in reality, we work with developers. We propose a standard solution to our clients, but obviously, depending on their business, there are sometimes things that don't work exactly as they want. So they request specific developments. And in that case, developers step in. I am sort of the link between the technical side (development) and the functional side, meaning the client's business process. Sometimes, the developer shows me their code. Even if I don't code myself, it makes sense to me: I can understand if what they are doing matches what we expect. So it is useful to have a minimum of coding knowledge, especially to be able to communicate properly with them.

Alright. Among all the hard skills you mentioned, like coding or Power BI, does generative AI help you in one way or another?

Yes, clearly. For coding, I know it can help immensely. I don't do it daily, but I'm starting to take more of an interest in Claude, and they have a tool called Claude Code. I don't know if you're familiar with it, but it's really impressive. I've already looked at a few little things with it and it's quite exceptional. Then, for Power BI, yes, I really use it. In Power BI, you can create what we call measures. That is also code, called DAX. And when I try to create a measure and it doesn't work, I ask Copilot. For example, I explain to it: "I need a measure to display this data based on that other data, and I want it to be dynamic in this way." I describe what I want to do, and I ask it to turn that into a measure. It then generates a DAX measure for me, I copy and paste it into Power BI. If it works, perfect. If it doesn't work, I send back the error message and continue working with the AI until it works. And several times, it has generated measures for me that I could use directly in my Power BI reports. And regarding D365, it's a bit like what I

was explaining before: I mostly ask it to explain concretely and practically what a specific configuration will produce in the system in reality.

So I have two questions regarding what you said earlier. The first is that you mentioned clients pay a lot for your skills and expertise. But with the rise of generative AI, which continues to develop very quickly, couldn't they try to do this work themselves?

Yes, but exactly, imagine that I have a lot of clients telling me that now. Actually, when a client calls a consultant, it's because there is a problem: something is wrong in the system, or there is a blockage. And it has happened to me several times, not just once or twice, but really often, that a client tells me: "I asked Copilot, and it told me this." And often, what Copilot tells them is incorrect. Clients think they can be replaced by AI or that AI can replace the consultant, but for now, it's clearly not possible. Even I, when I work with Copilot or ChatGPT to ask questions about the tool, they don't always answer correctly. So on that level, honestly, I think we are still far from being replaced. You also have to keep in mind that there is a human relationship in consulting. An AI will never be able to perfectly understand a client's real need. So yes, a lot of clients try to use AI before contacting me, but in 99% of cases, they come back to me saying: "It told me this," and I reply: "It said nonsense, you shouldn't listen to it." But indeed, a lot of clients do that. So this whole interaction and human relationship part is something that AI, for now, cannot replace.

Alright. In the literature, many authors say that consultants could transition from being experts to employees who simply validate GenAI outputs. Do you agree with that?

That depends on the timeline they are talking about. Are they talking about a year from now, or more like ten years from now?

More like now, in the near future.

Then no, I don't really agree for the moment. However, I would say that in three or four years, yes, we will see a huge improvement. I started using ChatGPT when I was in my second year of my Master's, so that was three years ago. And the gap between what we could do back then and what we can do today is enormous. So I tell myself that in three years, it will be even more impressive. But in the short term, no. In the very near future, I don't think so. But in three or four years, I think AI will start to gradually replace certain tasks.

When you or your colleagues use generative AI, do you specify it to the client or not?

No, I don't specify it. I don't do it because I feel like the client might say to themselves: "If he can do it with AI, then I can do it too." And besides, I use it differently than the client. Because I know the tool I'm implementing. The client doesn't master it yet. So when I ask the AI a question, I am much more precise, my request is more logical and structured than theirs. So no, I never tell the client that I used AI, whether it's for a Power BI report or to find a solution in the system. I don't mention it.

And in your opinion, the fact of saying or not saying that you use generative AI, is it a good thing? Is it a bad thing?

The fact of not saying it?

The fact of specifying it: is it a bad thing to say it or not?

No, for me it is not a bad thing to not say it. Because in the end, what the client expects is an answer. So whether I found it alone or with the help of AI, the client doesn't really care. As long as they get their answer and are unblocked, they are happy. So I don't think it's a problem not to tell them. On the contrary, I think telling them could sometimes discredit you. The client might say: "He is a consultant, but he uses artificial intelligence to help me." And then it comes back to the idea: "Couldn't I just use artificial intelligence instead of him?" So I think not saying it is perfectly fine.

Alright, and when you propose a solution, do clients never question your expertise? Do they never think: "Wasn't this done by..."?

No, never. Why? Because we have a huge amount of interaction with the clients. We communicate with them often, whether by email or on Teams. We regularly have calls, meetings, and so, since a bond of trust is already established, they do not question our expertise, in the sense that we constantly reassure them. So no, in my case anyway, I have never had a situation where a client asked me if something was done by artificial intelligence. Never. And it also has to be said that the software I implement, today, AI cannot manage it on its own. It cannot do the configuration the way I do it so that the system works in a certain way. I am the one, and only me, who sets everything up in the software. That being said, Copilot is starting to be quite powerful, I must admit. It manages to do more and more things. But in terms of system configuration, what we do on a daily basis, AI is incapable of doing it for now.

In the event that... normally this doesn't happen, and I think it will never happen, but let's imagine an extreme case: you propose a solution that is not at all what the client expected,

maybe even completely off the mark. And if you generated part of this solution with AI, who is responsible? The consultant individually? The entire firm? The consultant's manager?

For me, it's the team. Not necessarily the whole company, but really the team in charge of implementing the software. Why? Because AI, in the end, is a tool. And the one who uses the tool is the consultant. So it is their responsibility to use it correctly and to ensure that the answers obtained via AI truly meet the client's expectations. So for me, the responsibility falls first on the consultant in charge of the project, and then possibly on the manager or the team if the project really isn't working. But regarding the use of AI specifically, for me, the blame rests solely on the consultant's shoulders. Because they are the one using it. And for all we know, they might use it without informing their manager or even their team, so they might not even be aware that artificial intelligence was used. So no, for me, it is solely the consultant's responsibility.

OK, alright. Actually, since you answered several of my questions very quickly, I'm almost done. We are going to try to project ourselves 10 or 20 years into the future: what impact do you think generative AI will have on the consulting profession? We already talked about it a bit.

Yes. For me, even within three years, there will be a big change. And in 10 to 20 years, we will lean more towards what you were telling me earlier, regarding your literature review. And on that, I quite agree. I think AI will be able to do incredible things. At that point, it might be able to do what I do today, especially configuration, which it is currently unable to do. So I might be there more as a backup, in the sense that I will use AI extensively to do everything. That would save me time, and projects would move much faster. So yes, in 10 or 20 years, I think AI will be capable of doing much more than it does today. But I still think a human touch will always be needed in my job, because as I told you, it's a relationship of trust. Now, if AI becomes too powerful, then my job might no longer exist, because clients will no longer hire a consulting firm to implement the software. But despite everything, I think there will always be a need for people who are experts in the product. Because if there is no one at the clients who is a product expert, then they will have no anchor point. And what could also happen is that fewer consultants will be needed on a project. Because a single consultant could almost manage the project alone, since they would have all the necessary tools thanks to AI. So yes, in 10 to 20 years, I think that's where we will be heading.

I forgot to ask, but today, do you see generative AI more as a support, a tool, or already as a competitor?

No, 100%, it's a support.

Alright.

100% a support.

What advice would you give to people who want to become consultants today, specifically facing generative AI? What added value should they have?

Actually, for me, ChatGPT arrived when I was in my second year of my Master's. So I started using it then. But if we are talking strictly about artificial intelligence, I think future generations, and even consultants in general, need to train themselves in artificial intelligence: learn how to use it, learn how to write prompts, understand how it works. And even then, I say "write prompts," but sometimes I even ask the artificial intelligence to create my prompts for me. You see what I mean? But I think you really have to train yourself in artificial intelligence. Why? Because later on, there might even be consultants specializing in artificial intelligence. We will have artificial intelligences capable of doing an enormous amount of things, and these tools are going to be implemented for clients. So experts will be needed to master them. It's a bit like me today: I am considered an expert in the product I implement for my clients, so I can sell that expertise. And I think it will be the same with certain AI tools. There will be AI tools capable of doing exceptional things in a few years. So anyone who wants to become a consultant later should really train themselves at least on artificial intelligence. And I already see that, for example, since I started using Claude, Anthropic (the company behind Claude) offers completely free training on how to use artificial intelligence. And they are not the only company doing it. There are other artificial intelligences that also offer this kind of training. So for me, being able to use artificial intelligence in the future is a real added value.

My last question: in this context of generative AI, how can we continue to maintain a relationship of trust between the consultant and the client?

Listen, as I told you, it's a bit tied to what I was saying earlier. The client isn't really informed that I use artificial intelligence. Maybe they suspect it, but their main point of reference, with us, is the consultant. And so the relationship of trust is built on everything I explained to you: the calls, the meetings, the way I reassure them, and the fact that when they challenge or test me, I am able to answer their questions. All of that allows me to prove my value to the client.

And once you have proved your value, the client doesn't care whether you use artificial intelligence to deliver a solution. They just want it to work. If it works, great. If it doesn't work, then it's a problem. But whether it works because I did it myself or whether it works because it was 90% me and 10% AI, they don't care. So for me, this relationship of trust is established gradually over the course of meetings, calls, and exchanged messages. And the client eventually realizes that they can just simply trust you. And artificial intelligence, in my opinion, doesn't really play a big part in this relationship of trust yet. We'll see maybe later, but today, since I mostly use it as a support, there isn't really an impact between the client and me on this relationship of trust.

Interview ERP04

Let's start: could you introduce yourself? Who are you? What kind of company do you work for? And what is your role?

I finished my studies in August 2025 and started working in October, two months later. It's been exactly six months now. I work in a company called [Company Name]. It's a small structure: about 15 to 20 people internally, with 15 to 20 consultants and three partners (three bosses). It's a digitalization consulting firm. Then again, digitalization can mean everything and nothing. For example, some consultants work on ERPs like Odoo: they will mostly develop applications or implement an ERP. In our case, we work more on processes, developing Power BI reports, etc. I completed a first assignment with a client to help them produce a CSRD report. The goal was to digitize and automate the data collection needed for this report. We remain in a digitalization context, even if the ultimate goal was mainly a sustainability report, and not the development of an application per se. Currently, I am working with another client in a more finance-oriented context, because the company's main clients are in the financial markets. There is a new regulation dictating that from October 2027, certain trading operations will have to be executed faster. Today, there is a two-day delay between certain transactions, and this will have to be reduced to one day. Therefore, their processes need to be improved, testing must be done, and we need to ensure everything works correctly so that all of Europe can adapt to this acceleration. That is the general context of digitalization in the company.

It is clear. If I understand correctly, your assignments vary a lot from one project or client to another. You have quite varied assignments, and you don't really have a fixed procedure: everything depends on the client.

When you say "modus operandi", do you mean a procedure to follow? Not really. Since it's a small company, the organization is much less operational and structured than in large firms like the Big Four, which have very well-established methods. Our approach is much more client-centered. We adapt more according to their needs and their way of working. And that is exactly what some clients appreciate: a more flexible approach that is different from the Big Four.

In your opinion, what are the soft skills and hard skills that define your profession?

For soft skills, I would say it's easier: adaptability and flexibility. The client is never easy to manage. They have their habits, and it's not easy to change a process overnight.

I would also say that the human side is extremely important. Personally, I like it a lot, so I don't mind spending time communicating and explaining things, even for hours if necessary. In fact, I was talking about this with a colleague who was my manager on my first project: she realized how essential it was, especially because it was a public sector client. In the public sector, people are often less responsive than in the private sector, change is harder, and departments are very siloed. Departments communicate little with each other. Therefore, the consultant really has to ensure coordination, bridge the gap between teams, and facilitate communication. And for that, you need to have very good interpersonal skills. Ultimately, it is sometimes more of a human challenge than a purely technical one. For soft skills: adaptability, flexibility, and human/interpersonal skills. For hard skills, I would say analytical skills. We recruit a lot of profiles oriented towards business analysis. It's important to understand what the different departments do for the project to work. We can't stay on the surface, we really have to go into detail. Personally, to be able to tell people what they need to do and how to move forward, I need to understand their way of working, what they understand or don't understand, etc. You need to have a good level of analysis and an ability to quickly understand sometimes complex processes. As for other hard skills... I'm thinking, because I don't really know what examples to give. What would be an example of a hard skill?

Most of the time, I've met a lot of people working in ERP, and it was often more oriented towards coding and an interest in technology. But if you can name just one, that's already very good.

Coding is very important, yes. We have several profiles: they recruit quite a few people who know how to code, develop, etc. I don't really fit that profile, but I still have my place in the company because my profile is more oriented towards project management. I am not the person who is going to do the direct development. But I would say that mastering basic tools, like Excel, is essential. And even if you don't know how to write a formula, you have to know how to look it up, especially by using AI.

When you have to deliver an output to the client, what do you generally base it on? On your knowledge, your intuition, or your experience acquired over various projects?

I wouldn't say I base it solely on my knowledge, because ultimately, in consulting, especially at the beginning, you don't have a huge amount of knowledge yet. I am just starting out and this is only my second or third project, so I am clearly learning on the job.

I think what helps me the most is my perspective, my analytical skills, and my ability to understand situations, which allows me to figure things out. To deliver an output, I would say the main thing is to always base it on what the client wants. You have to constantly refocus on their needs, repeatedly ask yourself what they expect, even if the client sometimes changes their mind. To produce deliverables or solutions for the client, what do you concretely base them on? On your intuition, your experience, or your personal knowledge? I rely mostly on what the client wants, but also on logic and common sense. Sometimes, you just have to move forward logically. And sometimes, certain things are not possible: in that case, we think about the why, the how, and what we can propose as an alternative. Then, we do research to find a solution and reach the goal.

I'm going to jump back for a second: what type of clients do you have? You mentioned the public sector, but are they mostly large companies, startups, SMEs?

In the public sector, we only have one or two clients. They are still major players, but it was mostly a strategic decision to have a few of them.

Otherwise, our core clients are mainly private clients, and they are generally very large players.

In your opinion, what are the essential elements to build a good relationship of trust between a consultant and their client?

I think you always have to show a positive mindset and avoid presenting things negatively in front of the client. You have to show that you are motivated and, above all, anticipate as much as possible. The client then sees that we are invested, that we support them, and that we are involved in the project.

The human side is also very important. Then again, it depends on the profiles: some people are less relationship-oriented but are very efficient, and the client can still trust them.

You also have to deliver the deliverables on time, be available if the client has questions, and provide constant support. Most importantly, you need to know what to do: even if the client is sometimes lost, it is up to the consultant to ask the right questions and guide the project. That is what builds trust.

The two projects you told me about are quite different. At the end, do you have to do a presentation for the client to explain what you are proposing and detail what it entails?

I don't think you should wait until the end to do a presentation. In my two projects, it was more about presentations done progressively, at each step and for each deliverable.

Because there isn't just one deliverable. There is a final deliverable, but also many elements around it. For example, for the CSRD report, the main deliverable was the report, but we also had to provide a detailed roadmap of the steps to follow month by month. And that roadmap was almost as important as the report itself.

For my current assignment, I work more directly with the client's team, so I don't know yet if I will have to do as many presentations as an external consultant. But yes, I have already done presentations and slide decks to explain my understanding of the project, especially at the beginning. So I would say presentations are mostly done progressively.

And generally, do you have to go into details? For example, for the CSRD, did you have to explain the technical aspects in depth?

It really depends on the audience. For example, in steering committees, which are meetings scheduled every one, two, or three weeks depending on the project's urgency, with the main stakeholders. When talking to directors, you have to stay very "high level": they want to know what is progressing, what is blocked, and where the project stands, without too many details.

But recently, I presented the roadmap of steps to follow to the data teams. There, I had to go much more into detail and clearly explain each step, because they are the ones who will then have to do the work.

I practiced a lot beforehand to explain things simply, and I adapt to the audience. The goal is for the information to be clear, because at some point, they have to take over.

You were talking earlier about the follow-up with the client and the human side. When you have doubts about the project, do you talk directly to the client about them?

Not directly in the sense that I won't just say "I have a doubt." First, I write down the problem on my end. Then, I ask myself if the doubt comes from me, if I lack knowledge, if the client wasn't clear, or if the baseline contract is vague.

I then develop these questions and work collaboratively with a colleague. I go to her with all my questions. Internally, I am completely transparent.

Then, when we go to the client, we frame the question in a positive way. But above all, we arrive with one or two possible solutions already, because we anticipated.

If we arrive at the client's site expressing a doubt without proposing a solution, it might stress them out and make them lose trust.

Now, I will move on to the topic of generative AI with my first question: which generative AIs are you authorized to use?

It's always a bit complicated, because our clients don't necessarily have the same rules and policies as we do.

In my company, we are allowed to use any AI tools we want, as long as we do not put sensitive client data into these tools. The company encourages the use of AI: they trust us, and they consider it a time and efficiency saver. They also train their own bots and AI tools to get fast and concrete results, so we can move faster.

For example, I started my assignment with my new client two weeks ago. My manager had already worked with a lot of information from the internet, which he integrated into his AI tool (ChatGPT, which he trained). He then provided me with a very comprehensive document containing the basic information needed to understand the financial aspect related to the mission. That was very efficient and really helped me, because as soon as I arrived, I could already show that I understood certain concepts.

But on the client side, we can't do everything. At my current client, which is in finance, the rules are much stricter. I have access to almost no tools, except for the internal bots that the company developed itself. There might also be a Copilot, but within very strict limits, notably because I cannot import documents into it.

On my first project, for the CSRD reports, there were fewer restrictions. Since it was a public sector client, the policies were less strict. That being said, I was using the client's PC, so some access might have been blocked. I can't check anymore because I returned the PC. But in any case, we had agreed with the client that using AI was authorized, since the report data was going to be made public anyway. So it was acceptable to use AI to rephrase or rewrite certain content, especially since the client themselves was already using AI to generate their report.

So, in the end, it really depends on whether the client authorizes the use of AI or not.

Yes, exactly.

And for you, what tasks do you use generative AI for? You mentioned your manager who trained ChatGPT, but what do you use it for?

I use it mostly to understand everything and anything. I use AI a lot, I love it. Of course, I am careful: I never put client data into it. I use it mostly to understand the basic principles I need. I've only been working for six months and consulting is not an easy environment. There is a huge amount of things to learn, manage, and deliver. For example, to understand financial markets, I used AI, because I first needed to master the basics. I didn't need the client's data: I mainly needed to understand the context to ask the right questions.

I then create documents for myself with the information provided by the AI, and I always check it, because I know there can be errors.

For example, if I need to understand how financial markets work, I rely on several documents. But since these documents are often 100 pages each, I ask the AI to explain certain parts to me in small chunks. That works better than asking for a global summary.

Then, I verify in the official documents. But since I have already understood the basics thanks to the AI, I find the information faster and I can confirm if it is correct. It saves me time.

Alright, sorry, my apologies.

No worries. I can expand on that: I also use it to write emails. For example, to rephrase a message in a more diplomatic way. I always try to be very polite and professional with clients, so AI helps me a lot with that. I also use it for Excel formulas, or to understand how to do certain operations in Word, how to export files, etc.

You were just talking about the fact that you have to verify the answers, in case there are hallucinations or unreliable sources. Do you think you spend more time verifying the answers than doing the actual analysis yourself?

No, because I find that the tool I use is already quite reliable. I use the paid version, and I find it really high-performing. In general, the answers are well-structured and fairly accurate, especially if I phrase my requests correctly. I take the time to write my prompts, that's true. I can sometimes spend a lot of time on it, because I write fairly long prompts. But then, I get a really useful answer. Afterwards, once I have the information, I just look for the keywords in my documents, and I can quickly verify whether it's correct or not. Usually, it's quite accurate.

So in your opinion, is it a good tool to deepen your knowledge of a subject, or does it remain rather superficial?

No, it really allows me to go deeper, because I dig a lot. I ask a question, the AI gives me a very comprehensive answer with a lot of details, and then I go through it point by point.

For example, if it gives me 12 points, I take point 1 and ask a bunch of questions about it. Once I understand it well, I move on to point 2, etc. I really work in steps, a bit like sprints.

Alright. And so, we were saying that the client validates the use of generative AI or not.

Anyway, on their PCs, it is blocked.

For example, for your first project, the CSRD report, you said they agreed to let you use AI, at least in certain cases, even if you don't remember exactly which ones.

Yes, it was OK.

At the end, do they ask to what extent you used generative AI?

No, not at all. Since they had developed their own Copilot to write their report with AI, they completely didn't care. They obviously checked everything, but they needed a content base. So no, they didn't ask at all.

Alright.

But maybe that's specific to that client. Other clients might be more attentive to that.

In your opinion, does the fact that your company is open to generative AI make you seem more modern, or is it more a question of transparency with clients? Or both?

I don't think it's really to seem more modern. I think their main goal is to be efficient.

But AI is mostly used in the way we work, in internal processes, and not necessarily directly to produce the client's final deliverables. So, I don't necessarily see the direct link with transparency towards the client.

I mean: do you tell the client from the start "we use AI," and then they accept it or not? Is it more to be honest from the start or just to be more efficient and faster?

Honestly, I don't know how the partners, so the three bosses, handle this topic with the clients, because they are the ones who do all the upfront work. We come into the project afterward.

But personally, I don't hide it. For example, on the CSRD project, we had a question with the client, and I said that I had quickly looked it up on ChatGPT and that it suggested this answer. Since regulations are sometimes ambiguous, it helped to clarify. And on their side, they replied that they had done the same and gotten a similar answer. So I never hid it. But it was mostly

for small day-to-day questions, to understand certain things. On the other hand, with my current client in finance, I showed them my understanding of the financial markets, but I didn't explicitly tell them that I had used AI. I used the official documents, but I understood them faster thanks to AI. It was just to save time in my understanding. And in certain tasks, like testing in one of their applications, I didn't use AI at all. It wouldn't have helped me. If I had a question, I had to ask the colleagues, because it's their system and their application.

In your opinion, are clients going to ask more and more whether you used AI or not? I mean generally speaking in the consulting world.

Yes and no. I think yes, because AI is becoming more and more central, and everyone is asking more questions about data protection. I actually had some training on this recently: it has become a key topic, and we have to be very careful. So I think there will be more and more demands for transparency. But on the other hand, I think companies have already taken measures. For example, on my current client's PC, I cannot access ChatGPT, Copilot, or DeepL. It's just not possible. They also monitor sent emails: if we send too much information to a private address, cybersecurity steps in. It's particularly strict in the financial markets, because it's highly regulated. On the other hand, with my first client for the CSRD report, if I had wanted to send myself information and use AI with the data, I probably could have done it without anyone noticing. So I think there will be more awareness and more controls, but also that many companies are already putting systems in place to prevent what they don't want.

Once again, this is your opinion: let's take the case where a consultant produces a completely bad deliverable because they used AI irresponsibly, without checking its results. Who is responsible in your opinion: them individually, their team, their manager, or the whole consulting firm?

That's a good question. Personally, I think it's mostly the responsibility of the person, of the consultant who messed up. Today, companies are increasingly raising their employees' awareness of these risks. Everyone is at least informed. After that, it's mostly a matter of common sense. In my company, one of the main values is precisely common sense and trust. So if I do something bad because I didn't check my sources and did everything with AI, it is clearly my responsibility.

Does being transparent and open about the use of AI in a project run the risk of diminishing the respect clients have for your expertise?

Yes, there could be a form of discrediting, and I think that's going to happen in the consulting world. But in my opinion, if we are transparent about using AI, it is mainly to explain that we use it in the processes or to understand certain things, not to completely replace the actual concrete work. Since I started, I have always thought that if you don't think with your human brain at some point, you can't produce something of quality. So yes, there could be a loss of credibility, but it's up to us to explain that AI is an assistance tool: it serves to understand, to dig into certain topics, and to structure ideas. But in the end, there is always a human brain that must verify, analyze, apply logic, and finalize the deliverable. AI helps on the surface, but to produce a detailed, concrete, and high-quality deliverable, the human element remains indispensable.

And on that note, have you ever had clients who say: "Ultimately, I could do the project myself thanks to generative AI"?

No, because for the CSRD report, they used AI, but they really needed us to coordinate everything: collecting the data, ensuring communication between departments and with our clients. Here, it is about improving their processes and internal system, so it is not something AI can do on its own.

Ok, I'll move on to my last questions, which are more about the future of the consulting profession. For now, at this stage, do you see AI as a competitor or rather as a tool?

In digitalization consulting, I think it is clearly a tool. Digitalization is booming, companies need to automate their processes, so with or without AI, they will need us. We use AI to be more efficient in certain research, but it doesn't replace our role. That said, in other areas of consulting, AI will indeed replace many tasks, and I am aware of that. I think some consulting sectors might decline in the coming years.

Ok, so do you think that ultimately, consultants will end up mostly verifying what the AI proposes rather than creating the solutions themselves?

Yes, a little, but AI cannot analyze a company's specific situation. I saw this when testing AI for my thesis: it gave very superficial, unadapted solutions. The real solution came from me, by discussing it with my mom. AI remains limited to the surface; it does not replace human reflection.

Perfect. And how do you see the consulting profession evolving in 5 to 10 years?

I see myself continuing in the same role over the next ten years, because the human aspect will remain essential. AI can progress, it can impress, but it will not replace the human role of coordination and communication between IT and business. Perhaps the job will be more focused on the human side than the technical side, and that doesn't bother me; I love that aspect.

And in your opinion, how can a consultant maintain their added value in the face of generative AI in 5 to 10 years? What skills will be indispensable?

You have to know how to use AI correctly, be transparent, and not put sensitive data into the tool. You must keep a critical mind, common sense, and human judgment. Even if AI articulates certain things very well, it is always necessary to verify the information. Human reflection remains what makes the deliverable truly qualitative.

So, remaining relevant means combining AI and human work?

Exactly. The final deliverables must reflect our own reflection and judgment. For example, to understand the financial markets, I asked the AI for explanations, but I verified the information and structured it myself in a PowerPoint. AI cannot organize or prioritize information optimally; that is our role. Then, we can delve deeper into certain points through our own research.

Interview ERP05

As I specified in my email, my thesis focuses on the use of generative artificial intelligence, like ChatGPT or Gemini, in the consulting world. More specifically, I am analyzing how these tools influence the relationship between the consultant and the client, particularly in terms of transparency, trust, and human expertise. That is just to give you some context. I am going to ask you several questions, and the goal is simply for you to answer as honestly as possible.

No problem.

Let's start with an introduction. Can you introduce yourself? What company do you work for? What is your role and how long have you been there?

Yes, of course. My name is [Consultant Name], I am 24 years old. Currently, I work at [Company Name] as an ERP consultant, specializing in a particular ERP software called NetSuite, which is developed by Oracle.

I started working at [Company Name] in September 2025. Today, if my calculations are correct, it's been about 6 to 8 months that I've been with the company, probably around 7 or 8 months.

What is the nature of the projects or assignments you are given?

Concretely, I implement ERP solutions within the various companies we support. The assignments include, for example, data migration, the implementation of the software itself, as well as everything related to setting up the solution.

This also includes training the different stakeholders and users who will work with the tool.

There is also a part that is not directly linked to client assignments, but corresponds to my learning and developing my expertise on the software. The goal is to become increasingly proficient in NetSuite to be able to support clients more and more effectively over time.

Within the assignments, there is also a significant support component. Once the solution is implemented at the client's, we support them throughout the usage phase. This support can last indefinitely, or be time-limited if the project is planned for a specific duration.

Overall, my work therefore consists of implementing the ERP solution, and then supporting the client throughout its lifecycle.

The type of ERP you implement is generally used for what kind of tasks at the client's end? For example, is it to manage billing or accounting?

NetSuite is a software primarily suited for financial management. It allows for managing billing, as you mentioned.

For example, there is a specific module in NetSuite called SuiteBilling, which manages billing, including for complex products. This can involve phone subscriptions, Windows licenses, or other similar types of products that require specific billing rules.

The software also allows for managing accounting, as well as the management of fixed assets within a company. This includes, for example, buildings, land, or other fixed assets. These assets sometimes need to be depreciated each year and follow specific accounting rules. NetSuite therefore offers modules that manage these types of elements.

There is also multi-currency management, as well as the management of companies with multiple subsidiaries in different countries, with different accounting contexts to consider.

There are many possible features, but generally speaking, NetSuite is mostly geared towards financial management.

What types of clients generally need this type of ERP? What types of clients do you encounter in your assignments and projects?

They are generally clients who have already reached a certain level of maturity. If we were to classify them into small, medium, and large enterprises, I would say they are more medium-sized enterprises.

They are structures that already have a certain organization and are looking, for example, to expand internationally, or to acquire new markets abroad. NetSuite is therefore a solution well suited to this type of client.

I wouldn't say they are small companies, nor necessarily very large groups either, although it can also be suitable. But overall, they are more medium to large-sized companies.

In your opinion, what are the essential skills in your role, whether in terms of soft skills or hard skills?

I will start with soft skills. First of all, I would say that listening skills are essential, particularly active listening, in order to fully understand the client's needs. You have to know how to interpret what the client says and understand the underlying problem.

You have to be able to listen, but also to ask the right questions and focus on the real issues.

I would also say it is important to be proactive, meaning proposing solutions and anticipating needs. You shouldn't necessarily wait for the client to come back to you, because the client often has many other matters to handle. Their core business remains their priority, and the ERP is just a tool designed to help them in that activity. Sometimes, the client might forget certain aspects of the project or not give it their full attention. In this context, being proactive is very useful.

Another important soft skill is communication. As in many jobs, you need to know how to convey the right information and, above all, be able to make yourself understood. It is not just about saying what you want to say, it is also about making sure the person opposite has understood and correctly received the message.

Finally, I think you also have to be someone who works seriously and is dedicated. Those are the main skills that come to mind.

Regarding hard skills, in my case, NetSuite is software that I had to learn upon joining the company. I wouldn't say that it necessarily requires precise knowledge of this software before joining the firm.

However, certain common skills in this sector are clearly useful, like Excel, Outlook, or PowerPoint. Excel is particularly important, because we use it regularly in our assignments, especially to process and analyze data.

The SQL language is also very useful, as it is often used in data management. I would also say that basic programming skills are an advantage, because you might sometimes have to learn new languages depending on the project's needs. Having basics makes it easier to adapt.

Finally, in the current context, knowing how to use artificial intelligence tools is also an important skill. That happens to be the theme of this interview, and I think these tools can help speed up work and make you more efficient.

I will come back to artificial intelligence a bit later. For now, I am first trying to fully understand the challenges of your job before delving further into my topic. You were saying that you implement ERPs and that you also train clients on how to use them.

Exactly. During the implementation phase, there is a step dedicated to training the client. The goal is to teach them how to use the tool: how to navigate inside it, how to use it to perform certain recurring tasks they will have to do daily, and how to actually go about it. It is not a

training that lasts several months, but rather a phase where we show the client the basics and how to operate the tool.

Do you have to simplify the use of the ERP, or do you stay quite technical in the way you teach them how to use the tool?

I would say that for navigating the tool, there isn't necessarily a need to simplify things much. Why? Because, as I said, NetSuite is software geared towards financial management. The people who will use it are therefore often accountants or people working on financial tasks. They generally already know the jargon and vocabulary specific to this field. It is not always necessary to simplify heavily. Furthermore, the tool works like a web application with menus and buttons: you mainly just need to know where to go to perform the desired tasks. The main thing is therefore to show them the right paths within the software.

In your opinion, what are the essential ingredients to build a trusting relationship between you, as a consultant, and your client?

I think the first element is reliability. Being someone reliable, trustworthy. This involves, notably, respecting what you say, keeping your word, and being consistent in your commitments. It also involves paying attention to what the client expresses, so they feel heard and understood. All of this contributes to building a solid relationship and ensuring the client can trust us. For me, reliability is really point number one. Then, there is also transparency. To build a trusting relationship, the more transparent you are with the client, the more it strengthens that trust. Even though I think you don't necessarily have to say everything in every case. There is also benevolence: showing genuine care towards the client, not just trying to sell them features at all costs simply because they are a client and you want to maximize revenue. The goal must be to advise them as best as possible.

This can even go as far as advising them against certain options if they don't truly match their needs or goals. In my view, all of these elements help build a long-term relationship of trust.

I wanted to come back to what you said regarding transparency. Do you sometimes have doubts about what you do, the recommendations you give, or the tool you are implementing? And when you say that you shouldn't necessarily say everything, what do you mean exactly by that?

When I say that you shouldn't say everything, I meant it in the sense that, of course, we talk about transparency, but we aren't going to expose the entire business model of the company to the client, for example. It's more in that sense.

It's not a matter of hiding important things from the client. It's more about being transparent regarding the solutions we implement: what the impacts of these solutions are, what the advantages are, what the disadvantages are. You must not lie to the client about what you are implementing for them, or about what you are concretely proposing to them.

I understand better. Thank you for the clarification. In that case, I will now get to the heart of the matter regarding generative AI. The first question, quite obvious: what AI are you authorized to use within [Company Name]? Are they AIs like Gemini or Copilot, or rather internal AIs?

We have an artificial intelligence that was developed internally, and that is primarily the one authorized.

In addition to that, it is possible to obtain other licenses through a formal request. You have to make the request to the people in charge of managing these access rights.

For example, you can request a license for Copilot, which is a tool found in many companies, notably through Microsoft solutions, which are very widespread and widely used today.

In our case, it is possible to request this type of license.

But other than that, the company has mainly developed an internal AI, and we are strongly encouraged to use it. The reason is that the more we use it, the more data we feed it, and the more performant it can become over time. So it is an AI that we currently use on a daily basis.

And since it is an internal AI, I am not familiar with it. Concretely, what does it do?

It basically works like ChatGPT in everyday life. You could say it is an enterprise version of the same type of tool.

If [Company Name] implemented this artificial intelligence and gave us access to this solution, it is mainly to avoid sharing client data with other external artificial intelligences. The goal is to maintain total control over client data.

Since the AI is internal, we can enter whatever data we want without any problem, knowing that it will remain confidential.

Personally, for what purpose do you use this AI? For what tasks do you use it, whether it's this internal AI or possibly Copilot?

I mostly use it for fairly basic tasks. For example, when we have to write documentation, I can ask it to rewrite certain paragraphs so that they are clearer, more professional, or simply better phrased.

I also use it a lot to make summaries. For instance, when there is a large document to read, or when we have written a lot of text and want to summarize it, the AI allows for a faster and more structured summary, while remaining professional.

I also sometimes use it for anything related to Excel formulas, or even to write formulas linked to different programming languages. This could be SQL, Java, or other languages. The AI helps optimize these formulas and save time.

Sometimes too, even for the software we use, NetSuite, I have certain specific questions that I can ask the AI, and it helps me find solutions.

Those are the main uses that come to mind for now.

In your opinion, what aspects of this internal AI will never be able to replace the consulting profession? What will it not be able to do, or what makes it unable to replace you?

I think AI cannot be as flexible and as precise as a consultant in understanding the client's expectations. We, as humans, are capable of understanding the client's real need, what is expected, what needs to be implemented, and above all, adapting to changes. In a project, there are often several stakeholders involved. At one point, some people might add constraints, then others might remove some, then add more again. And all this while keeping the history of what was initially requested. This ability to manage daily unforeseen events is essential, and unforeseen events happen all the time and always will. I think that today, AI is not yet capable of managing these kinds of situations as well as a human being. There is also the dimension of trust. With a human being, you can establish a relationship of trust. AI, in the end, remains a robot. You can choose to rely on it, but the day there is a problem, you can't really turn against a robot. It only executed what it was asked to do, according to what it understood. You can't really discuss with it like with a person, nor ask it to adapt its way of doing things or to think differently. Furthermore, in certain situations, you have to find a creative solution that has perhaps never been imagined before. However, AI works based on a set of existing data: it proposes answers based on what it already knows. But when you have to step outside the box, outside of what is known, it is much more complicated for it.

Finally, there is also the question of the regulatory context. For example, when you want to implement something for a client or propose a more innovative solution than what is usually done, you have to make sure that what you are doing complies with regulations. In our context, we might ask ourselves: is this compliant in Belgium? And if we are in France, can we do the same thing? Are the rules identical or not? Being able to adapt to this type of context, depending on the countries and regulations, is essential. And I am not sure that an AI today is 100% reliable for these kinds of issues.

Precisely, when you use AI and write prompts to get a result, are there often hallucinations in the answers? Do you spend a lot of time checking that there aren't any? For example, you mentioned Java earlier: does it hallucinate a lot in that type of case, or not?

Personally, I haven't really had to use Java yet, or very little, so I can't really base it on that. But if I take the example of Excel, I have asked the AI for a formula, and the formula it gave me didn't work. Then, asking again, it gave me another formula that still didn't work. At that point, by testing other things, trying to phrase the request differently, or by testing other solutions myself based on what it already suggested, I sometimes end up finding something that works. I can also do research on the Internet and see what has already been done. Yes, it still happens that the AI hallucinates.

You were just saying earlier that sometimes AI recommendations don't work, and that you can't really turn against a robot. However, many companies have already had cases where AI gives outputs with hallucinations. In that case, in your opinion, who is responsible? Is it the individual consultant, or is it the entire structure, the company?

I would say that, from an internal point of view, it depends on the freedom given to the consultant and their level of responsibility. For example, if it's a junior consultant, then for me, the person directly above them should be responsible. But that is only from an internal point of view. From the client's point of view, however, the company is liable, because we act as representatives of the company. Regardless of the decisions we make, if a problem occurs, it is the company we represent. The company can therefore be held responsible, in my opinion.

I found my question. Do you find that when you use generative AI, you are more productive or not? Does it have an impact on your productivity? For me, yes.

Yes, I find that I am more productive.

And exactly, the time you save, for example thanks to AI for documentation or other tasks, do you then dedicate it to doing other things? Does it allow you to be more productive in other aspects of your consulting job?

Yes, clearly. The time saved allows me to focus on another task more quickly. It also allows me to respond faster to client needs. And if we respond more quickly when a client has a problem or asks a question, it improves their satisfaction. Yes, I think the impact is positive.

When you use AI like that, are you obligated to inform the clients? Is it not necessary? How does that work?

It is not necessary to inform the clients. I don't even think we necessarily do it. As soon as we deliver something to the client, it means we consider it ready to be delivered. We aren't going to tell the client: "We are handing this over to you, but watch out, AI did it." If we provide a deliverable, it means we have checked and validated it. In my opinion, there isn't necessarily any specific communication to make to the client on that level.

And on their end, won't they ask the question? For example, asking if AI helped you on certain parts? Won't they ask the question themselves?

No, I have never encountered this case yet, and I haven't heard of it happening either.

Do you think that saying you used AI for an Excel formula or something else could break the client's trust? Or in your case, does that not really apply?

I think it depends on the client, and especially on how they perceive AI. But in any case, if a client has a problem with the fact that we used AI, they could also try not going through a consulting firm. What I tell myself is that if the client has a problem with consultants using AI, they could try doing without consultants, and see if they manage to achieve their goals alone. Then, by making the comparison, they could realize whether they should continue to hire consultants or not. But... can you remind me of the question, please? I think I forgot it a bit.

Sorry. I was asking if that wouldn't risk harming the client's trust in the long term. If we talk about the use of AI, could that change this perspective of trust, or does that not apply here?

I think that in the long term, it won't necessarily create a loss of trust. However, it could push clients to think more before calling a consultant, and to ask themselves if they wouldn't be capable of doing what they want to achieve themselves before asking for help. But regarding trust, I don't think it will decrease. What builds trust in the work of consultants are the elements

I mentioned earlier, which an AI cannot necessarily reproduce. AI remains a robot. Entrusting great responsibilities to a robot is a risk that people can choose to take or not. I don't think the fact that consultants use AI reduces the trust clients have in them. On the contrary, it is more of a tool that helps consultants do their jobs better. I don't think it reduces trust.

You have answered so many questions so efficiently that we are almost at the end. For you, right now, AI is more of a tool, a support, rather than a direct competitor to the consulting profession.

Today, it is clearly a tool. Sometimes, clients ask questions that go beyond the strictly technical scope of the implementation. In those cases, it is more about expertise and experience. For example, a client might say: "Regarding this module, we are hesitating because we currently have this solution in place, and we don't know if we want to switch to yours. What do you recommend, based on your experience? What have other clients with a similar profile to ours already done with you? What do you think?" In this type of situation, we step outside the purely technical framework where it would be enough to give a simple answer, like a logic of $1 + 1 = 2$. Here, it is experience that allows for a relevant answer. These situations happen very often, at least from what I have seen. In those contexts, AI cannot do the consultant's job as efficiently or with as much relevance.

I will ask the question again: in 5 to 10 years, will you still have the same opinion? Do you think that in 5 or 10 years, AI will still simply be a tool?

In 5 to 10 years, given the current advances in AI and all the new technologies being put in place, I know that the goal of many world leaders is to develop intelligent agents. Agents capable of doing things like humans, capable of learning, and which no longer rely solely on what they know or on an existing database, but would be capable of thinking and reasoning like humans. I know that is the goal. Will they actually get there? I don't know. Even if they do, it would still mean placing trust, as well as the strategic pillars of your company, in robots that can be fallible. For example, if there is a hack, a computer crash, or malware that manages to infiltrate, it could influence the answers these robots will give. That could have dramatic consequences. I think the question goes beyond pure competence. If we only talk about technical competence, AI could become more competent than human beings. But there is always this dimension of trust: being sure that what the AI says is not a hallucination, but the truth. And that is more complicated. For me, leaving 100% of your trust to a robot remains difficult. Within ten years, I think there will still be consultants.

In my literature research, many authors say that consultants could transition to another role. This role would consist of verifying AI outputs rather than creating the solutions themselves. It would be the AI producing the solutions, and the consultants would verify if these solutions are correct or not. What do you think about this?

I don't think AI will create solutions on its own. I think it will rather help to create them. In the end, it will always be the consultants who have to ask the AI the right questions so that it provides the right answers. It will also be necessary to give it the right information so that it can produce relevant results. AI, in itself, does nothing alone. For example, if I open ChatGPT, ChatGPT does nothing as long as I don't ask it anything. If I ask a question a certain way, then ask the same question with a different phrasing, I can get two different answers. AI helps create a solution, but ultimately, there is always a human being behind it who must ask the right questions and provide the right input.

Right. So, my last question: if we look 5 to 10 years into the future, what do you think would be the added value a consultant must have to remain relevant in the face of generative AI?

I think the added value will lie mostly in the trust and reliability of the proposed solutions. Yes, definitely the reliability of solutions. There is also trust... Wait, I had something in mind that I am starting to forget. You were asking what skills would be needed in ten years, right? Well, I would say that we will also need to be able to bring a strategic dimension to companies. Being able to guide them on a strategic level, and helping them better navigate the ambiguity that all companies face on a daily basis. In any case, for most companies that are starting up or wanting to move to another level, there is often a form of uncertainty, because we don't know what the future holds. So being able to guide them, and being a trusted partner on the strategic decisions they must make based on their objectives, will be essential. I think this is a skill that consultants should develop: the ability to anticipate what might happen tomorrow and to support clients as a reliable partner, based on their goals.

One last very quick question, because you mentioned it at the beginning: wouldn't it also be relevant for a consultant to know how to use AI well? For example, knowing how to prompt effectively and knowing how to detect hallucinations. Would that be an important skill for the future?

Yes, I think it is important. Knowing how to prompt well is essential, because it is based on the question that the AI will give an answer. So if you ask bad questions, you will get bad answers.

And in that case, there is no point in blaming the tool. Often, these tools are highly powerful, but we don't give them the right information, or not in the right way. So yes, even today, you need to know how to use AI. Otherwise, it will take much longer to do what you have to do. Today, it is a time saver, so not using it makes you lose efficiency. I, for example, know people who do not use AI, and I see that certain tasks take them much longer than other people who do use it. So yes, I think knowing how to use AI will be an important skill to have.

Right, I am done with my questions. Thank you very much, you have been really thorough and highly relevant. Thank you so much.

Thank you.

Interview ERP06

I am [Consultant Name], I am 26 years old and I work in Technology Consulting at [Company Name]. I joined the firm almost two years ago, I hope to be promoted to Senior in two weeks. I did my master's in business engineering at [University Name]. I started it at [University Name] and not in Louvain-la-Neuve, but I still did a minor in Louvain. Most of my studies took place at [University Name]. I majored in Business Analytics and minored in International Finance.

What kind of projects do you work on in Tax and Technology?

I work mainly on ERP implementation projects related to tax, as well as data reconciliation projects between accounting and VAT reports.

What ERP tool do you implement?

We implement different tools, but we work mostly with SAP, which is [Company Name]'s main partner in this field. We also use Oracle, although that is much less frequent in our case. In my personal case, I am working on an SAP project called SAP GRC, for Governance, Risk and Compliance. It is a platform intended to perform compliance checks, verifications, and data monitoring within companies. In this project, I am mainly focused on the tax aspects.

How does a project with a client generally unfold?

It varies a lot, it is not always identical. Generally speaking, a project lasts about six months on average. First, there is a design and kick-off phase. Then, a build phase during which we implement and build the solution at the client's site. Then comes the testing phase, where we verify the implemented solution. After testing, there is the sign-off, then the hypercare phase, during which we remain on support until the client is completely operational with the implemented solution.

In your opinion, what are the essential soft skills and hard skills for your job?

For soft skills, it is quite common to all consulting jobs. Communication is point number one. You also need an analytical mind, organization, and the ability to structure your way of working to manage large projects over fairly long timelines. Personally, I would put particular emphasis on communication, because we are very often in contact with the client, and very early in our careers. From the first year, we are already in contact with important clients holding high positions in large companies. You must therefore know how to communicate both with clients and internally, and be able to manage teams and collaborative work. This requires a skill set

close to what you learn in management. Regarding hard skills, there are no strict prerequisites, but certain elements are an advantage. Accounting skills are useful, because we work in an audit and finance-oriented environment. Basics in computer logic or programming are also a plus, especially understanding the logic behind the systems. Today, with tools like AI, it is possible to generate code without knowing a language perfectly, but it remains important to understand the underlying logic. Finally, knowledge in tax is obviously an advantage if it was acquired during your studies.

What, in your opinion, builds a good relationship of trust between the client and the consultant?

I would say first clear communication, with transparency on the situation and the progress. The client needs to feel reassured and confident. Then, meeting deadlines is essential. To build a relationship of trust, you have to deliver what was promised within the announced deadlines. Finally, there is also an important human aspect: interacting with the client informally, making small talk, checking in on them, and creating a more relaxed bond. This helps build a closer relationship and shows your commitment to the collaboration.

Which generative AI do you use in your department?

At [Company Name], we have our own platform to guarantee data confidentiality. In this platform, we use many different AIs, practically all the ones available on the market. We have also developed internal AI agents that can, depending on the use case, combine several models. For example, the same agent can use ChatGPT, Claude, and Gemini in parallel. Regarding the basic chatbot, I don't know exactly which model is used in the backend, but overall we have access to a wide variety of AIs, without any real limitations.

Which ones do you use and in what contexts or for what tasks?

I am not going to go into too much detail about [Company Name]'s internal tools, but if you are referring to tools like ChatGPT or Claude, I can answer on both dimensions. At [Company Name], I primarily use [Internal LLM]. I am part of the Agent Tech AI Builder team, I am supposed to develop AI agents, even though I haven't really had the time to work on it yet.

[Internal LLM] is the main tool I use professionally. Before, we also had Perdy, which I really liked because it gave very concise answers, but it will soon be replaced. Today, [Internal LLM] is mostly used for tasks like creating flowcharts, designing design structures, or defining business requirements. For these types of tasks, I find it more performant.

I also work a lot on SAP, and there is notably an ABAP agent that I use for technical questions. For example, I can ask it how to join tables together to create a specific table, or ask technical questions related to SAP, ABAP, or CDS. However, I don't really use certain agents like the ones for PowerPoint, because I find their results are still too unreliable or irrelevant for now. Outside of work, I sometimes use ChatGPT or Claude for personal projects, but that remains marginal compared to my professional use.

Does generative AI help you more to better understand your projects or mostly to save time?

It depends on the case. Sometimes, it helps me understand better. For example, if I have a meeting or a call and I didn't understand everything, I take notes and then ask the AI to explain it to me, rephrase it, and give me context. It helps me clarify certain things. In other cases, especially when I already master the subject, I use AI to save time, for example to create a deliverable or speed up a task. But in about 80% of cases, the main goal remains saving time. Either because I already know what I want to do and simply want to go faster, or because it's more efficient in a work context where speed is important.

You mentioned that AI answers are not always optimal and that there can be hallucinations. Do you spend a lot of time checking if the AI outputs are correct or if they are hallucinations?

Yes, absolutely. Especially when it comes to technical tasks or even design. Not checking can have a significant cost, especially if you present an output from [Internal LLM] to the client and it contains errors or incorrect information. That can cause problems. I've had cases where the AI gave me an answer that seemed very good, but when I presented it to a developer, they told me it wasn't possible at all, even though [Internal LLM] had confirmed it. You must always remain cautious. In my field, in technology, I think vigilance must be even higher, because as soon as you get into technical topics, AI can say a lot of incorrect things. It happens quite often, especially in SAP. There are good uses, as I mentioned with the agents, but you must always double-check several times before considering an answer valid. On the other hand, for tasks like drafting or listing business requirements, AI is very useful. Since it is mainly text, language models are very good at structuring and formulating that type of content. But as soon as you get into technical questions, for example how to retrieve data in SAP by combining certain tables, which tables to use or which keys to employ, the AI can easily produce incorrect or invented answers.

Have you observed, since you have been at [Company Name], consultants who are less experienced on paper but capable of performing tasks of more experienced consultants thanks to generative AI?

Yes, absolutely. I myself am not very senior; I have almost two years of corporate experience. I clearly think that if someone had been in my place before the arrival of AI, they would not have achieved the same results. Today, even as a junior, regardless of the task, we are much more willing to take it on, even without prior knowledge. We figure we can always ask the AI for help and get something decent. For example, I also see this with my brother, who is interning at [Another Company Name]. Juniors are much more inclined to accept new tasks, even when they have no initial expertise in them. With ChatGPT or other tools, we sort of feel like we have superpowers. We think we can do anything, because AI can help us produce something functional quite quickly. I therefore think it makes juniors much more efficient and faster in their delivery. However, this mainly concerns the technical and operational aspects.

During a project, do you define upfront with the client whether generative AI can be used or not?

No, not at all.

How does that work then? Does the client ask at the end of the project if certain parts were done with AI?

I haven't encountered a case where that was asked yet, no. To give you a little context, I don't produce a report. My output in my field is not a written document. I am more likely to create a workflow or a type of code, to simplify. Whether it's with or without AI, the client generally cares very little. What matters to them is the final result: an interface where they can enter their parameters, run a process, and get a functional output. They are satisfied as long as it works. So, in this context, the use of generative AI isn't really visible to the client. They don't see how the solution was built in detail. They only see the final result. I understand that in some cases, notably for reports, there can be controversies over the use of AI, with examples of poorly executed documents. But in our case, at least in my work, it's not problematic at all.

Do you think that in the future clients will demand more transparency on the use of generative AI, even in ERP implementation projects and not just for reports?

No, well, that's my personal opinion, so it's subjective, but I don't think AI will go in that direction. On the contrary, I think AI will mostly increase client expectations in terms of speed

and productivity. They will expect faster and more efficient results. I don't think we will move towards more regulation or restrictions on the use of AI, because today, among major clients, it is mostly perceived as an advantage rather than a risk. We are currently in a mindset where the goal is to use AI in as many contexts as possible. So it's hard for me to see an evolution towards limitations or strict conditions on its use.

So if I understand correctly, the clients themselves see the use of AI as an advantage to save time and be more productive?

Yes.

Last question, a bit more oriented towards the future of consulting. Do you think that some clients could eventually tell themselves that they can do the work themselves thanks to generative AI, without going through consulting firms?

Yes, clearly. It's a real issue. Generally speaking, the consulting profession consists of bringing expertise to the client. If this expertise becomes much more accessible thanks to AI, clients will logically have more leeway and might think they can do more internally with their own teams. So yes, I think AI represents a significant risk for the consulting business model, especially in the short term. That's where the pressure will be the strongest. In the longer term, I think the profession will evolve. Consultants will have to offer other forms of value and advice. The model won't disappear, but it will change in nature. The most impacted areas aren't necessarily pure consulting, but rather activities like compliance, accounting, or certain legal and tax areas, where part of the work can be automated. In consulting and advisory, there could be a reduction in the number of juniors, because seniors can produce more thanks to AI and delegate certain tasks to tools like ChatGPT or Claude rather than to entire teams. My personal opinion is therefore that, in the short term, productivity will strongly increase and structures will reorganize. But in the long term, the sector will stabilize and return to growth, because there will always be a demand for advice, even if its nature has evolved. I think it will also depend on the fields. For example, Tax Technology is still a fairly recent field, growing strongly, which recruits a lot. So depending on the sectors, the impact will not be the same, but in any case, firms will continue to find value to sell to clients.

Today, do you see generative AI as a support or a competitor? And is it the same in the future?

Both. It's a support, because it helps us produce more efficiently. But it's also a competitor, because clients can easily ask ChatGPT directly instead of going through us. And that is already the case.

In the literature, some authors suggest that the consulting profession will evolve toward less direct production and more verification of generative AI outputs. Do you agree with that?

Totally. That is already the case today. Juniors are already mainly using AI: they ask ChatGPT, then they check and adjust. There is almost no completely manual production of code or documentation anymore. At most, we write a draft, then ask the AI to improve it, and then produce an output that can eventually be slightly refined. The pure creation phase (writing or coding) has already been largely reduced, I would say to about 90% today. If we go back five years, to produce code, you really had to be very technically competent. Today, even a business graduate, with a bachelor's or a master's level, can produce something functional thanks to AI.

How do you see the consulting profession evolving? And how can one remain indispensable face-to-face with generative AI?

As I mentioned, in the short term, I think headcounts will continue to decrease, notably with fewer recruitments of juniors. Seniors and management will be less impacted, because they already possess the know-how and can amplify it with AI to produce what juniors used to do before. The greatest impact therefore concerns juniors, particularly access to their first job and the ability to produce concrete value as soon as they enter the profession. Seniors are more involved in decision-making, and that remains difficult for AI to replace. Juniors, on the other hand, make few decisions, so their value relies more and more on communication, client relations, and the ability to present results. Pure operations are increasingly outsourced to AI. Thus, in the short term, I see a reduction in headcount, and this is already observable. For example, this year, we recruited about six people in Tax and Technology Consulting, compared to ten the previous year, even though the workload remains very high. This does not come from a lack of work, because everyone is overloaded, but rather from uncertainties regarding the commercial pipeline and future demand. There is also a risk linked to AI itself. For example, a client could compare a solution proposed by [Company] with a solution generated by an AI like Claude. Even if the quality of the AI is inferior, it becomes an option in the decision-making process. A CEO could ask if the problem has already been tested via an AI before signing a contract with a consulting firm. Finally, we must add the current macroeconomic context,

marked by a lot of uncertainties (wars, economic instability, etc.). AI adds to this environment and constitutes an important piece of the puzzle of the transformations currently underway in the sector.

Ok. I have one last question. Let's assume we reduce the number of juniors, as you said, and as a result we only keep seniors, not uniquely, but a large majority of seniors. Then, how can the client be persuaded? How will the client maintain trust precisely toward the people who remain for the projects, etc.? Why would they want to pay for those people?

Why wouldn't they want to? Again, because precisely, there will be generative AI. That means they might not have managed to do what they wanted using generative AI, and so they will call upon seniors, is that it? In any case, they never really called upon juniors. I mean, that wasn't even conceivable from the start. When a client signs a contract, they talk with the managers, and it is the managers who propose or create a team to provide the solution. The junior is never directly in contact with the client.

No, no. Well sorry, that might be a poor formulation on my part, but at what point, then, will we no longer call upon juniors? That's what I mean.

That managers will no longer call upon juniors? Internal managers at [Company]? Well, what is the problem? I still don't understand the question.

It's because actually, in your vision of the future of the profession, you said there would perhaps be fewer graduates and juniors, and that we would mostly keep experienced seniors. But then why? My question remains the following: why would [Company] managers call upon seniors for projects, knowing that there is still this generative AI and that the client could also try to do the project themselves?

Yes, but it is not that categorical. The client will try to do things, that might reduce the number of projects, marginally or considerably, I don't know. In any case, there will be contracts. We won't reach a stage where there are no projects at all, that's not possible. From there, managers will delegate to their seniors, and the seniors will solve the problems with the help of AI, quite simply. But again, in my opinion, that is short-term. I think that in the long term, we will find a balance. Or perhaps the profession will go back to how it originally was: experts answering specific problems. Today, we are no longer at all in the configuration of when the profession was created. At the time, the profession was composed of small teams of highly experienced consultants capable of answering questions that clients could not solve themselves. Today, fresh graduates out of school who are called "consultants," sometimes that doesn't really make sense

anymore. Even being a "senior" after two years, to me, doesn't really make sense either. We are not yet experienced enough. We are mostly supports for the managers, etc. So the real expertise, in my view, is truly located at the highest levels. Everything below that is more about operations. Except that now, seniors are being asked to take on more responsibilities, to have more client contact, to step into roles that managers occupied a few years ago. The process is currently accelerating.

Interview ERP07

So, I have several questions. You just need to answer them as spontaneously as possible. First, can you introduce yourself, tell me what position you hold, and how long you've been there?

Yes, my name is [Consultant]. I have been with [Company] as an employee for 2 and a half years now. I am a consultant in the Oracle TNT department, and before that, I did a 4-month GPC internship in the same team.

Sorry, this might seem like a silly question because I've already asked you these before, but I need them. So, what kind of projects do you work on?

So I work on ERP implementation projects, but also on support. On one hand, there is everything related to implementation: these are clients who are starting on NetSuite for the first time, and we help them through this transition. And otherwise, there are also clients who are already on the system, and there we help them with new projects or daily questions.

Ok. In your opinion, what are the skills that build your expertise as a consultant, whether in terms of soft skills or hard skills?

I think, for soft skills, that you need to know how to ask questions and be able to gather a company's needs, and then challenge them. Because often, the company has its own way of seeing things, but with the other projects we've had, it's interesting to be able to challenge those visions and see if we can't redefine certain processes. For hard skills, I mainly think of critical thinking. You have to be able to challenge the system and ensure that the proposed solution is the most suitable one in the long term. We are in an environment where companies grow a lot, so the idea is that the solution works for one entity or for 35, and that it can evolve with the system. These are the main points I see.

Ok. And in technical terms regarding implementation, do you sometimes need to code? Not necessarily.

Yes, there are some technical aspects. Sometimes, there is code, whether it's HTML or SQL, but we also have more technical people on the team, so it's also a choice. We aren't always obliged to do it ourselves, we can. With artificial intelligence, we can also rely on tools to help us. There are also workflows in NetSuite, where it's essentially a process to build. And in that case, there are sometimes technical aspects, for example knowing at what moment an action must be triggered to avoid crashing the whole process.

Ok. When you work on an implementation project, in your opinion, what are the factors that guarantee the quality of the result?

A point I have already mentioned is ensuring that the solution is designed for the long term, to prevent companies from coming back in 3 years saying that nothing works anymore. There are also comparisons with other tools: for example, in the team, some work on Odoo, and we might say that NetSuite is slower or make those kinds of remarks. The idea is for the solution to be durable, but also as high-performing as possible. Let's imagine that a process takes a lot of time: we can consider launching it at 12:01 AM, at night or on the weekend, when no one is on the system, in order to make the experience as smooth as possible for the user.

Ok. So when you were doing an implementation project, what guaranteed the quality of the results and of the solution?

Yes, I think there is a lot of expertise in the team. And there are also times when the client wants to do something, but we are able to tell them that it's really not a good idea, for reasons A or B. And that is tied to expertise and to what has been done in the past. I also think that what makes the difference with the competition is all the surrounding material: the documentation, explaining why we chose one option rather than another, and properly training the users. In a system, it is easy to make a mistake that can sometimes cost a lot, so it is important to properly guide and train the users. I find that to be quite important.

Ok. In your opinion, what builds a good relationship of trust between a consultant and the client?

I think transparency is essential, so being honest, especially regarding the budget or deadlines. At my level, for example, if I tell a client that a task will be done by today and it's not the case, I have to explain: "actually, it's taking more time than expected, sorry". If we ever make a mistake, we must say it clearly: "my mistake", and explain the implications that it might have. You always have to verify when you aren't sure and stay transparent about those aspects. I also think that having worked with other clients plays a role. A question we often get is: "How did you do that with other clients?" And there, we can bring experience by saying: "With them, they did this." It shows that we don't just follow what the client asks, but that we challenge them a bit. Sometimes, it can be a bit frustrating for the client, because they might say "no, I want to do it like this," but by being challenged, they understand there is a reflection behind it. And that also creates a relationship of trust.

Ok. So, my next question is: which generative AIs are you allowed to use in your department, and which ones do you personally use?

Well, I just use [Internal LLM], which is [Company]'s AI. We can also use copilots, but I don't use them. That's more for coding, so since I'm more on the functional side and have less coding to do, [Internal LLM] is often enough. And actually, NetSuite is currently developing its own artificial intelligence, but it is still at an immature stage and mostly available in the United States for the moment. So we don't have it yet, but I think we will have it on the system in the future.

Ok. And so, in your opinion, what tasks is [Internal LLM] useful for, and what tasks is it useless for?

I think that, from a PowerPoint, slides, Excel, or Word perspective, it is a bit of a disaster. It completely fails to generate what we expect, so I don't use it for that at all. What I use it a lot for is, for example, when we have new countries. Recently, I had an entity in Portugal, and I had never worked on that country. There are specific requirements for each country. I can search the internet, of course, but it's always useful to have a quick overview generated by [Internal LLM]: for example, to understand a specific tax, or how a system works in a given country. In the United States, for example, they don't have VAT but another system. We can ask why they don't have VAT and how it works. So it's very useful for getting a global view of a new process. Otherwise, I also use it for code. We can always ask a technical person, but sometimes you just need to modify one line. So I ask [Internal LLM] how to change it. Or, for instance, if a script causes an error, I can ask it to explain what the script does or generate documentation. It's also useful for PDF or HTML. For example, for syntax errors in a script that is hundreds of lines long, it can help quickly spot the problem. Or to compare two documents and identify the differences. These are simple tasks, but they save a lot of time. And also, when there is an error, I can ask it what the possible causes are, and it often suggests several leads.

Ok. Do you find that generative AI changes the way consultants are perceived in terms of expertise? In the sense that it's often said you have to learn in the field by doing things yourself. Do you think it changes that aspect?

I feel like there are two dimensions. Before, for example, anything involving scripts or technical aspects was hard to tackle without help, and you could easily make mistakes. Today, it has allowed me to gain technical skills and dare to dive into more technical topics, because I can rely on [Internal LLM] instead of bothering a developer with very basic questions. And

ultimately, that is their job, so it's better not to disturb them for that. It allows me to do in 15 minutes what would have taken an expert 5 minutes, but the next time, it will take me 5 minutes too. So it speeds up learning. It also broadens expertise, because we better understand what is behind the systems. Sometimes, we only have the functional aspect, and everything technical remains blurry. I had an example two or three weeks ago: there was a discussion with a client and an integrator. If I have to participate in this type of exchange, I need to understand what is happening in the background. Otherwise, I would always need a developer with me, which isn't always necessary. The second point is that some very junior tasks are disappearing, like taking notes or certain repetitive tasks. AI is partially replacing them.

Ok. So, when you use generative AI, even minimally in a project, do you tell the client about it?

Not necessarily. I know some clients explicitly tell us they don't use artificial intelligence, while working with Excel files full of errors. They then ask us to clean or structure them with [Internal LLM]. We are quite transparent: if I use something, I have no problem saying it's [Internal LLM]. I simply clarify that it's an internal tool at [Company], with secure data, and there is generally no problem on that front. An interesting point is that more and more clients are using ChatGPT or other tools themselves to check things in NetSuite. But we observe a lack of critical thinking: they take the answers at face value. And sometimes they cite Claude or ChatGPT as a reference, which makes it hard to explain to them that the tool can be wrong, whereas we are not. It's sometimes complicated.

Ok. And have there ever been clients who were against the use of artificial intelligence in projects?

No, I have never encountered any, and I haven't heard of any from my colleagues either. We work in highly technological industries, so companies are quite open and even try to look into these subjects closely. I think that in our industry, it's less of an issue than in other sectors.

My next question is: do you think that being transparent about having used generative AI in a project can diminish the client's perception of your expertise?

No, I don't think so. For example, if I take a simple example, like generating a PDF, sometimes you have to write some code behind it. Saying that I used artificial intelligence to help write that code saves them money in the end, because before it could take 4 hours and now it takes 40 minutes. So it really makes a difference. And I think the added value that [Company] can bring, and us as consultants, is to challenge certain things. For example, we've seen this with

workflows: now NetSuite is starting to create workflows itself, which are things we do for clients. I'll give a concrete example I have in mind: when a user creates a vendor and attaches banking information, that's how payments are made. Often, what we see is that an approval is needed; otherwise, someone could create a vendor in the system and make a 3 million euro payment to themselves without any control. In this case, [Company] sets up an approval process to validate these payments. The issue, for example, is making sure that as long as it isn't approved, the payment cannot be made. And we've seen that NetSuite sometimes simply offered an "approve/reject" option, but without real consequences on what the user could do in the system.

This means that even if the bank details are not approved, the user can still use that bank account, which doesn't make sense. What we do is set these accounts to inactive so they cannot be used until they are validated by the right person. There are also other safeguards, for instance: even if I am the CEO of the company, I cannot approve something I created in the system myself. It's rules like that. Yes, artificial intelligence can provide good solutions, but this aspect of critical thinking and business validation is still complicated for it. And that is what we need to show clients: that they cannot replace us on this aspect yet.

Ok. Exactly, earlier you said that transparency was important to guarantee trust. In your opinion, what is the impact of using Generative AI on this transparency? Does it change anything?

How to put it... does being transparent about using AI create more trust? I think so. Sometimes it's important for clients to realize we use these tools, because it saves them time and money, as I said. The ultimate goal is to be as efficient as possible. So if we can use these tools, which are actually internal tools at [Company], even if clients don't always realize it, we use them to be more efficient and improve the quality of what we deliver.

Ok. And do you think that in the future, clients will be more demanding about AI outputs? Will they be more demanding, thinking that AI is supposed to be "perfect"?

Yes, I think they sometimes have unrealistic expectations. For instance, with Excel or simple tasks. And I think we also have this bias ourselves. Sometimes I ask [Internal LLM] for something that seems very simple to me. For example, I've had cases where I wanted an Excel formula combining several columns with specific conditions. And [Internal LLM] gave me something horrible. At some point, it was faster to do it manually than to try to fix the formula. So I think there can be a form of naivety, on both the client and consultant sides, regarding what

AI can actually do. And when it doesn't know how to do something, it can be very frustrating, especially when these are tasks supposed to save time, like cleaning a 300-line Excel sheet. Because if in the end you have to do everything manually, you go from 1 hour to 6 hours. So you also have to be honest with the client: no, it's not always instantaneous, and no, AI cannot do everything. It is very useful for certain tasks, but completely ineffective for others.

Who should take responsibility in case of failure? For example, in an implementation project, if certain things were done poorly with the help of generative AI, who is responsible? The consultant, their manager, or the entire firm?

At my level, on a project, the responsibility would likely be shared between me and potentially my manager (N+1), because a manager's role is often quality assurance, ensuring that what is delivered is correct. I also work with more junior profiles, and when they do something in the system, my role is to verify that it's done correctly, that the idea behind it is the right one, and to coach them. So I think the responsibility is shared between the two: the person who creates must ask themselves if what they are doing is correct, and then there is someone who does a review and must ensure everything is done well.

Ok. But my last question is more about the future of the profession. You mentioned earlier that clients are using ChatGPT and questioning what you propose. In the future, will clients turn to consultants less and less, and try to build the solutions themselves using generative AI?

I think so. Especially in our case, for instance with NetSuite, we already see them implementing this type of feature. And as I mentioned with workflows, we have already discussed this internally. We knew clients would show up with these types of tools, so we try to align ourselves and see how we can best "embrace" this. Because there is clearly a part of what we currently do that AI will be able to do instead of us. For example, with the vendor case I gave you, we need to show that our added value isn't simply implementing in NetSuite, but especially everything around it: challenging decisions, anticipating impacts on other processes. An AI will generate a workflow, but it won't necessarily think about the fact that there are six different types of vendors and that different processing is needed depending on the case. For example: if it's an internal vendor, is it really necessary to verify the bank details? A company has so many implications and interconnections that, when a question is asked, we can think about the impacts on multiple streams, whereas AI doesn't always have this global overview. So I think our added value should be there: not in technical execution, but in critical thinking and experience.

Ok. So today, do you see artificial intelligence as a support or a tool? And the same question for the future?

I see it as a support today. I think support and tool are close, but I distinguish them a bit. Today, it's more of a support that helps me occasionally. In the future, and even for the last 6 months, I see that changing. It will be something as natural as Teams or PowerPoint: a tool that you have to master and use on a daily basis. Today, for example, I haven't even opened it yet in my day. But in the future, I think it will be constantly open and used every hour. On the other hand, I don't think it's a competitor. Maybe in some very junior roles, but even then, we will probably redefine the tasks. For example, when I started as a B, I used to do meeting minutes. Today, that makes almost no sense, because AI can do it. But mostly, it frees up time to do something else and bring more value to different tasks.

Ok. So my second-to-last question: how do you see the profession evolving in 5 to 10 years?

I think we will have to become more specialized and more expert in certain areas, and be able to bounce back on that. It's a difficult question. But clearly, we are going to have to develop things like artificial intelligences specific to [Company] for NetSuite, Odoo, etc. Today, we already have [Internal LLM], but tomorrow there will probably be more specific versions depending on the tools. In our team, we have what we call "assets", which NetSuite doesn't have. These are solutions developed internally, sometimes by entire teams over months. Today, we can reuse them in a few clicks, but creating them took an enormous amount of time and effort. For example, for a 150-slide PowerPoint, certain parts often come back. What was created the first time over several days can today be reused and adapted in two hours for a client. And I think it will be the same with artificial intelligence: we will be able to sell "packaged" solutions, for example an AI developed by [Company], installed at the client's, and adapted to their environment. It could cost, for example, 5,000 euros, but it would be trained and personalized differently from what NetSuite offers.

OK, and my last question, even though you kind of answered it earlier, what will make clients still turn to consultants in 5 to 10 years? What will be the added value of a consultant compared to generative AI?

Well yes, as I said, I think critical thinking, even though I imagine at some point artificial intelligence will also train itself a bit more on errors and all that. Because now we sometimes have system hallucinations, but I think in 5 or 10 years, we will have almost none left, not at

all in the same way as now. I would say, yes, critical thinking, and I am a bit skeptical about how... Now we really have to see all of people's requirements.

How can AI help in discussions with clients?

Imagine, we are in a conversation, there is a third person, you say, "Look, I want the process to be this." The third person says, "No, I think it's this." Our role too, we've already found ourselves in conversations where clients do not agree with each other. And the idea is also that with AI, they will both say, "I want this, I want this." Ok, for both, we can do that. There is this as a positive point, there is this as a negative point for each idea, but who is going to lead a discussion with the two people? I think that's a role we also have sometimes as an intermediary, and sometimes too, I think clients are happy to say, "Ah, it's [Company] who was going to do it like that," because then when there are impacts in the company, it's a bit like "well, we were advised to do this and we followed it, and there you go." It's harder to say that with artificial intelligence. Especially since you formulate your question. By formulating your question, there is a name for that, but you somewhat submit what you want as an answer sometimes. There is a psychology term for it. I saw it in school, I remember. If you look for an answer in a certain way, you will often get the first answer the way you want, but what you put in isn't necessarily always true. That, I think, is something we can still bring. Yes.

Ok, OK, perfect. Well, that was my last question

Interview RISK01

I have prepared a series of questions, some a bit hidden. You can answer in the most honest and spontaneous way possible. Don't hesitate to take your time to think if a question surprises you. Alright, let's start.

Okay, sounds good.

Can you introduce yourself? Where do you work, what is your role, and how long have you been with this company?

I have been a consultant at [Company] since last August. I joined as a risk controller. We work on risks related to internal control, not financial risk. Concretely, this includes, for example, checking revenue workflows, security procedures, and so on. Our role is to perform internal control, which is called internal audit, to be distinguished from external audit which concerns accounts, which we do not do. At the same time, I am pursuing my university activities. My academic background includes a master's degree in economics. I started in the public sector, notably with an internship at the OECD and a stint at the Council of the European Union, before moving to the private sector.

Alright. And concretely, what do your missions at [Company] consist of? You've already mentioned it, but can you elaborate a bit more?

Of course. We work with public or private organizations. There are two teams: one specializing in the public sector and one in the private sector. Over time, everyone specializes more. For my part, I have only worked in the private sector, but I could have worked on public institutions; some colleagues are assigned to them. Each mission is different. For example, in a revenue control mission, we analyze the financial functioning of the company to identify potential revenue leaks or risks that can be avoided. Conversely, some missions are more focused on security audits, evaluating management transparency, personnel compliance with rules, the use of personal protective equipment (PPE), and its availability. These are two typical examples, but it illustrates the diversity of our missions.

For this type of mission, how do you proceed concretely? Do you go on-site, analyze documents, interview the client...?

Generally, missions unfold as follows. Going to the client's site is, in my opinion, essential. To understand how a company works, it is not enough to just have reports or numbers in front of you. Being on-site allows you to ask questions immediately, get additional information, and

talk directly with the people involved. It also allows you to get to know the client you are working with better, which is always an advantage. However, some missions require less physical presence. For example, in a liquidated revenue control mission, it is not really necessary to meet the client multiple times. We usually conduct an initial interview to gather information and break the ice, possibly a mid-term interview, and a final interview at the end of the mission. These meetings mainly serve to formalize the follow-up rather than provide critical information. On the other hand, for missions like security audits, being on-site is indispensable. You have to concretely observe compliance with the rules, verify that personal protective equipment (PPE) is available and used correctly, and ensure that employees are informed of the procedures and new rules in place. This involves looking at how employees work on a daily basis, how management communicates with them, and conducting in-depth interviews with management and employees. For this type of mission, the on-site presence is therefore much more intensive and structured. What is interesting about this job is that every mission is different. Depending on your profile and the way you work, you can be assigned to varied missions that match your preferences. If you don't like routine, this also allows you to regularly change mission types and gain new experiences each time.

Alright. In your opinion, what are the essential skills for your job as a consultant, especially in the risk domain, in terms of hard skills and soft skills?

For hard skills, it is not necessary to know how to code or master Python. Personally, I haven't done any since university, so it is not essential. The required computer skills are rather basic, but you need to know how to fully exploit certain software. For example, Excel is widely used, and you have to be able to leverage it fully, whether to save time in analysis or to produce accurate reports. Knowing shortcuts and advanced functions allows you to work much more efficiently. For report presentations, it is also important to know how to use Power BI. What is good about [Company] is that, despite it being a major firm, training is really taken care of. Even if someone is not fully trained on a software or doesn't know it perfectly, they will be trained without any problem. Regarding soft skills, some are essential in all companies. For example, being pleasant and sociable is important. A company doesn't just look at technical skills; it also evaluates if the person will integrate well into the team and if the team will be comfortable working with them daily. This doesn't mean you have to be constantly in a good mood or participate in all social activities, but simply not be closed off or bitter, not remain isolated, and not refuse to interact with colleagues. Other important soft skills, particularly at [Company], include the ability to recognize one's limits. This means knowing how to say, "I

don't know how to do this," and accepting to listen to or follow someone who masters the task. It's not about taking on responsibilities that you are not capable of fully assuming. In my team, I was lucky to have exceptional supervisors: seniors, managers, directors, who are always available to guide and train. Thus, if you don't know how to complete a task, you can ask for help without a problem. To summarize, hard skills include the advanced use of basic software like Excel or Power BI, while soft skills include sociability, humility, and the ability to ask for help and recognize your limits.

Ok. And when it comes to the internal audit or the security procedures you mentioned earlier, what knowledge do you have to apply? Do you receive training for that?

Yes, absolutely. When you arrive at [Company], we have two weeks of mandatory training, a sort of intensive full-time program. Then, once this period is over, we become mission leads, but we are never left alone. Field experience is essential to truly master the job. For each mission, we are always accompanied by a senior, a manager, and sometimes a director. In practice, the majority of the work is supervised by the senior and the manager. The director, on the other hand, manages the overall organization and the follow-up of the team. For example, in the risk domain, we use the Risk Control Matrix a lot. If it's the first time you have to build one, it is difficult to do it alone. The best approach is to ask a senior for help, who guides us and explains how to avoid mistakes and what to pay attention to. At [Company], there is also the [Company] Academy. Depending on your department, you have access to specific training that allows you to familiarize yourself with the concepts and tools specific to your field. For example, if you want to learn how to build a risk control matrix, you will find detailed training modules to support you. Finally, we also use an artificial intelligence interface specific to [Company] for certain internal practices. We do not use ChatGPT or other external AIs like DeepSea for confidentiality reasons. Everything remains on [Company]'s servers and nothing is shared. For training, we do not yet use client data; it is only used to familiarize ourselves with the concepts. The use of client data is still more sensitive, but the internal AI is planned to be adapted for it eventually.

Last question, more about your expertise. When you prepare a deliverable for a client, what do you base it on? Is it on your theoretical knowledge, your accumulated experience, your intuition?

When I create a deliverable, several elements come into play. There is of course what I have learned, my intuition, but also my past experience. This is why the intervention of seniors and

managers is important during the mission. Generally, once a week, we have a meeting with the senior, the manager, and sometimes the director to present what we have done. Depending on their experience, they indicate what is correct, what needs to be modified or improved. Thus, before transmitting a deliverable to the client, it is reworked several times a week based on the advice and indications of the supervisors. In summary, the deliverable is based on a combination of knowledge, experience, and constant feedback from the team.

Ok. We were talking about the client and their data. In your opinion, what creates a relationship of trust between the consultant and the client?

It is essential to keep the client informed. For example, before using an artificial intelligence tool, whether with or without client data, we explain how our interface works, how the data is processed, and why security is maximal, so that there is no risk of a leak. Of course, no system is totally risk-free, but our internal [Company] interface is almost completely secure and has never caused a problem. We also present the benefits: these tools allow us to speed up the processing of certain time-consuming tasks, like data filtering, which could take a week manually but only an hour with the tool, if the client agrees. So, it's beneficial for the client and for us. The important thing is to clarify the functioning and the advantages of the tools, and to use them only if the client gives their consent.

So, this is clearly established with the client beforehand?

Yes, it's mandatory. At each client meeting, either the director or the manager presents the tool with slides to clearly explain how it works and ask for the client's agreement. Even for economic reasons, it is advantageous to obtain this agreement, because certain tasks can then be carried out more efficiently.

And do you also present technical details to the client?

Not in depth, we are not IT specialists. We rather explain the benefits, the risks, the general functioning, and the servers used, insisting on security. The idea is to reassure the client and clarify that we do not use ChatGPT or other public AIs to process client data. All our tools and interfaces are developed and controlled by [Company].

When you mention interfaces in the plural, do you use other generative AI tools besides those created by [Company]?

No, only those created by [Company]. We have a ChatGPT-type tool, a sort of chatbox to ask questions and get answers. We also have other tools to automate certain tasks, but I cannot go

into details for confidentiality reasons. All these tools remain internal, created and controlled by [Company].

Ok. And in what cases are you allowed to use these tools, and why?

So, last question, rather about your expertise. When you prepare a deliverable for a client, what do you generally base it on? Is it on your theoretical knowledge, your accumulated experience, your intuition?

When I create a deliverable, several elements come into play. First, what I have learned, obviously, then my intuition, but especially my past experience. This is why the intervention of seniors and managers is crucial during the mission. For example, during a mission, we generally have, at least once a week, a meeting with the senior, the manager and sometimes the director to present what we have accomplished.

During these meetings, the seniors and managers, thanks to their experience, point out what is correct, what needs to be modified, what can be improved, and sometimes give advice on points we hadn't thought of. So, before handing the deliverable over to the client, it is reworked several times a week, following these advice and indications. In summary, the preparation of a deliverable relies on a combination: the theoretical knowledge acquired, the experience accumulated on previous missions, intuition to identify key points, and regular feedback from seniors and managers to ensure the final quality.

Ok. And precisely, we were talking about the client and their data. In your opinion, what contributes to establishing a solid relationship of trust between the consultant and the client?

The main thing is to keep the client informed. For example, before using an artificial intelligence tool, whether we work with client data or only with internal data, we always explain to the client how our AI interface works. We detail how the data is processed, the security levels in place, and why there is no risk of a leak. Of course, no IT system is completely risk-free, but our internal [Company] interface is extremely secure and, so far, no incident has occurred. We also explain the concrete benefits: these tools allow us to speed up very heavy and time-consuming tasks. For example, for data filtering that could take a week manually, the tool can do it in an hour if the client agrees. This is advantageous for us and for the client. The idea is therefore to clearly present how the tool works, its advantages and its limits, and to only use it with the explicit agreement of the client.

So, this agreement is established upstream with the client?

Yes, absolutely. Personally, at every client meeting, either the director or the manager presents the tool with detailed slides to explain how it works, the benefits and the risks, and to ask if the client accepts the use of the tool or not. It is mandatory, because we want to make sure the client understands and validates the process.

Even from an economic point of view, it makes sense. Our missions are generally billed based on working days. If the tool reduces the time needed for a task, it is beneficial for the client but can reduce billable time. Nevertheless, the priority is to get an agreement and guarantee trust, rather than saving time without validation.

And when you present the tool to the client, do you go into technical details?

Not in depth. We are not IT specialists. Instead, we explain the general functioning, the benefits, the risks, and the security measures, such as the fact that the servers are highly secure. The objective is to reassure the client and clarify that we do not use ChatGPT or other public AIs to process sensitive data. All tools are created and controlled by [Company], guaranteeing confidentiality and security.

When you mention interfaces in the plural, do you use other generative AI tools besides those created by [Company]?

No, only those created and controlled by [Company]. We have a ChatGPT-type tool, a sort of chatbox to ask questions and get structured answers. But we also have other tools to automate different specific tasks, which I cannot detail for confidentiality reasons. All these tools remain internal, created and supervised by [Company], guaranteeing that no data leaves the secure system.

Ok. And in what context are you allowed to use them and why?

We use them only when it has been validated by the client, for specific and time-consuming tasks, like the automation of repetitive or tedious data processing. The objective is to save time on these tasks while maintaining quality and security. These tools are not used to directly generate consultancy reports intended for the client; they are mainly used to facilitate internal work and make the process more efficient, always with the client's agreement.

Ok, and so, in what cases are you allowed to use these artificial intelligence tools and why do you use them?

Well, the use depends on the context. If it's for purely personal learning, we are free to use it whenever we want. For example, if my question is "How do you build a risk control matrix?", I can use the tool without restriction. On the other hand, as soon as we work on real data from a client for a specific mission, there are two indispensable conditions: The agreement of my supervisor, which is generally clear from the beginning of the mission, because we propose to the client whether the tool will be used or not. The explicit agreement of the client. If these two conditions are met (supervisor and client), then I can use the tool for the mission and I have total freedom within this framework.

And what is this tool used for in practice? For example, to produce risk matrices or reports?

Not for the direct production of reports or risk matrices, unlike some other firms that use AI to entirely write a consultancy report from data. At [Company], this is absolutely not the case. Superiors constantly remind us that AI must never replace the consultant's added value. AI is used to automate repetitive and time-consuming tasks that do not require a consultant's expertise. For example, data filtering: previously, this was done manually by consultants, but it is not work requiring a specific skill. If you take a student and show them the method, they could do it in a week. On the other hand, everything related to the consultant's expertise, such as report writing, matrix building, and in-depth analysis, remains the consultant's responsibility. AI is simply a tool to save time on mechanical tasks, but never to replace our expertise or judgment in a report or an analysis.

Ok. In your opinion, what aspects of the generative AI implemented by [Company] cannot yet replace the consultant's job?

For me, it's everything that requires specialized expertise. AI relies on all available information, but it makes a sort of "general mix" of all this information coming from different sectors or types of consultancy. The result is a report that is too global, which can include irrelevant or generalized information. The consultant, on the other hand, brings sectoral expertise, acquired through experience. Even a junior consultant has specific training and methods, but more experienced colleagues possess fine expertise in their sector, know past mistakes, and know how to avoid them. AI cannot reproduce this nuance. AI is very effective for repetitive and time-consuming tasks, like data filtering, but for anything requiring critical analysis, judgment, or contextualization, it is not up to par. It will never be able to replace intuition, human expertise, and direct contact with the client, which are essential.

Does generative AI help, in your opinion, to better understand the client's needs?

Yes, but mainly for someone discovering the subject. For example, on a first mission, if a client sends a very technical email, AI can help decipher the vocabulary and understand what it means. Once you have gained experience over several missions, AI becomes less necessary for this type of understanding.

And when you do not use generative AI, how do you present your conclusions to the client? Are they interested in the process or only in the final result?

We have to present and justify everything. Even for the final report, every point must be explainable. If a client asks a question about a specific detail, we must be able to demonstrate how we reached that conclusion, what data was used, and the logic applied to get there.

And when there are moments of uncertainty, difficulties, or unclear points in the analysis, do you talk to the client about it?

Yes, the client is informed during the mission. For example, if obtaining a specific type of information is difficult, we discuss it with them: "This point is a problem, can we find a solution?" If the information is missing or poorly distributed, we identify the cause and propose recommendations. For example, for safety rules, if the client only sent an initial email and did not distribute the information afterward, it creates a real risk. We explain this to the client and make recommendations to improve the situation. The goal is for the client to understand the difficulties encountered and how we addressed them, and for corrective measures to be proposed if necessary.

Ok. When presenting solutions to the client, and they know you used generative AI, do they ask you which parts were done with AI and how it arrived at a particular solution?

Yes, it happens, but it depends on the type of solution we present. At [Company], anything that falls under the consultant's real expertise, such as report writing or in-depth analysis, does not use generative AI. This is a clear directive from our supervisors. So yes, clients sometimes ask the question, but generally, it's not on a specific or complex point. They simply ask, during the report, for which parts we used generative AI. In that case, we mention it openly. But since this mainly concerns mechanical tasks, like filtering or data collection, there isn't much to justify: it's simply to explain what was automated.

And this desire to specify to clients beforehand whether generative AI will be used, is it to build trust with them or rather to show that you are technologically advanced?

No, it is not about showing that we are technologically advanced. [Company] is not trying to "appear" cutting-edge in technology. The objective is primarily to save time on repetitive and time-consuming tasks. However, when we present the use of AI to the client, it is for reasons of total transparency. We want to be perfectly clear about what we do and how we do it. Of course, this contributes to the relationship of trust, but the main point is that we must, as one of the largest firms, demonstrate complete transparency to the client regarding our methods and tools.

And do you think that in the future, clients will increasingly demand this transparency regarding the use of generative AI?

Regarding transparency in general, I think they already are. We work with companies that are generally of a significant size, where the question of AI has already been raised. So they are already aware and ask questions about AI if necessary. For our part, we anticipate this. For example, for presentations or slides, we prepare them even before the first interview, in order to explain the AI solutions we might use. Regarding the use of generative AI, clients are increasingly demanding it because they realize it reduces mission costs. Previously, a large part of the preparatory work was very time-consuming and had to be billed. Thanks to AI, this work can be accomplished much faster, which reduces the final bill for the client, sometimes, for example, from around 300,000 euros to 200,000 or 250,000 euros. Clients understand this gain and readily accept it.

Ok. Let's take a hypothetical case where a consulting firm provided reports made with generative AI, with errors or hallucinations, without informing the clients. In your opinion, who is responsible? The consultant, their manager, or the firm as a whole?

For me, it is the firm as a whole. If a report containing such errors is delivered to the client, it means there were no clear rules established by the firm and that the supervisors did not pass these rules on to the juniors. It would be too simple to say the fault lies only with a junior. Even if a junior completely writes a report, it must pass through the hands of a senior, a manager, and a director before being sent to the client. If errors persist upon delivery, it is therefore the responsibility of the whole firm, because each hierarchical level has a share of control and responsibility. At [Company], this kind of situation does not happen, because the directives are clear: AI must never replace the consultant's added value. A junior would never dare to disobey this directive. Even if they tried, the director or the manager would detect it and ask them to correct the work. In summary, responsibility is collective; it is distributed from management

down to managers and juniors: a fault in the delivery of a report or a deliverable is therefore attributable to the entire firm.

OK, now we will move on to my last questions. They concern the future of the consulting profession. You mentioned several times that clients pay a lot for missions and that the use of AI helps reduce costs. But in the long term, with the evolution of generative AI, don't you think clients might end up doing what you do themselves?

Well, I don't think so, but there are two aspects to consider. First aspect: for clients to be able to totally replace a consultant's work with AI, this AI would have to reach a human level in terms of expertise, analysis, and the capacity to produce quality work. I personally think that will never be the case. Second aspect: even if, hypothetically, AI reached this human level, clients would still demand human contact. Why? Because behind a consultancy report, there are always strategic decisions to be made. For example, in a security report or a report on revenue losses, the recommendations will guide important decisions. If the AI made a mistake, the client would suffer the consequences directly, when in reality, it's the AI that would have failed. Moreover, even if some clients tried a transition phase where they only used AI, they would quickly return to the human consultant. The discussion, experience, and judgment of the consultant remain essential: a report is never just a document; there is exchange, clarification, adjustment, and recommendations based on the interaction with the client. To illustrate, I'll take a concrete example: cashiers. AI could very well replace certain repetitive tasks, like scanning items. But even in this case, customers would end up preferring human contact rather than a computer for everything. So for jobs requiring expertise or human judgment, AI will never be able to completely replace the consultant.

Ok. So for you, AI is really a tool or a support, but not a competitor for the consulting profession?

Exactly.

In scientific literature, some authors believe that eventually, consultants would mainly become people who verify the outputs of generative AI, and that the majority of the work would be done by AI. Do you agree with this perspective?

No, I disagree. In my opinion, this scenario would be the death knell for the largest firms, like the Big Four. If consultants merely had to verify AI outputs, clients might wonder why they should pay ten times more at [Company], [Competitor 1], or [Competitor 2] when they could hire an independent consultant or internalize the verification. Today, clients pay for the

expertise, experience, and real added value of consultants. At [Company], for example, I have worked with highly competent and specialized colleagues. Their experience is valuable and is part of what clients pay for. If the job came down to simple verification of AI-generated outputs, the point of using a firm would disappear. Even the verification could be entrusted to another AI. So, I don't think we will reach that scenario. AI will remain a powerful tool to save time and automate certain tasks, but it will never replace human value in consulting.

OK. So in that case, how do you see the profession evolving in 5 to 10 years? If it evolves.

It is clear that there will be an evolution, and AI will be at the center of this evolution, that is for sure. I think there will be two types of firms. Some will say: "We use AI for everything." Why? Because they will be able to massively reduce costs and work much faster than their competitors. They will therefore be cheaper and faster. Except that, ultimately, the work will be of lower quality. We could end up with situations like the case of a major firm that was sanctioned because some reports contained inconsistencies or nonexistent sources. Conversely, other firms, and in my opinion this is the right way to proceed, will say: "We use AI, but not to replace the consultant's expertise," which is typically what [Company] does. So, in my opinion, we will have two possible paths: one path where AI is used massively to reduce human work, but at the risk of producing lower quality deliverables, and a path where AI is a tool serving human expertise, allowing time to be saved on certain repetitive tasks without sacrificing added value. In the long term, I think it is the firms that follow the second path that will survive.

OK. And my last question: in this future, in 5 to 10 years, what would still be the added value of a consultant? How could they remain relevant in the face of generative AI?

For me, the first added value is experience. I will take a concrete example: in the construction sector, there are general rules or recommendations that are applied during safety inspections. These recommendations generally work, but during a mission, a consultant may encounter a specific case where this general recommendation does not work for a particular reason. Then, several years later, if the consultant encounters another company in the same sector, they will be able to remember this specific exception and adjust their recommendation accordingly. AI, on the other hand, relies on general rules: it might propose the same recommendation, and the problem could happen again. So, for me, what distinguishes the consultant is the accumulated experience, which allows anticipating exceptions and adapting solutions, as well as the human relationship, because interactions with clients are essential and often neglected. Human contact allows for understanding subtleties that AI cannot perceive. In summary, even in 5 to 10 years,

AI will be a tool, but the consultant's relevance will always come from their experience and their ability to manage human relations.

Interview RISK02

Can you introduce yourself, specify your position, your company, and how long you have been working there?

I am a consultant at a professional services firm, where I have been working for two years now. I work within the Finance Transformation department. In this department, we intervene on transformation projects, whether digital or organizational, to help companies rethink and optimize their financial strategy. This department is structured into several sub-teams, each having a specific role in this global transformation. For my part, I work more specifically on internal audit missions within Finance Transformation.

Concretely, what do your internal audit projects consist of?

In internal audit, we work with different clients and profiles. The main objective is related to risk management. We analyze the clients' operational processes to identify whether they have appropriate internal controls. These controls allow for the reduction of operational risks, their mitigation, or even their elimination when properly implemented. We intervene on different types of processes, such as purchasing processes, supplier payments, or IT processes. The goal is to verify if controls exist and if they ensure the efficient and secure functioning of the organization.

In your opinion, what are the essential soft skills and hard skills for the consulting profession?

I would say that one of the main skills is the ability to maintain a good client relationship. This primarily involves communication, which can be innate for some people but is also learned extensively through experience and different clients. It is also important to develop knowledge of the sectors in which we operate to better communicate and understand the client's challenges. Improving English, or other languages, is also an important skill in this profession. Regarding hard skills, I would say mastery of tools like Excel and PowerPoint, as well as other analysis tools like Power BI. In some cases, more technical tools can be used or implemented depending on the project. At my firm, we are well supported in these aspects with training that allows us to build skills progressively. Overall, I think everything can be learned over time.

When producing a deliverable, what takes precedence in your opinion: expertise, experience, theory, or intuition?

It depends a lot on each person's role in the project. For example, a junior does not yet have all the expertise, but they are expected to learn and progress. The client is generally aware of this because the teams are composed of varied profiles: juniors, consultants, seniors, and managers. Everyone has a well-defined role in the project. A junior will be more in a support role, while a senior consultant will already have expertise built from past experiences on various projects and clients. The manager plays a role of coordination and supervision, ensuring the overall quality of the deliverables. During commercial proposals or projects, these elements are clearly presented, notably through the CVs and experiences of the mobilized teams.

You mentioned the importance of the client relationship. In your opinion, what allows for building a relationship of trust between a consultant and their client?

I think the relationship of trust is built above all through direct contact with the client. Initial face-to-face meetings are important to establish a human relationship and explain everyone's role. Then, this relationship develops progressively through regular exchanges and meetings throughout the project. It is essential to maintain continuous communication with the client and not wait until the end of the project to share progress. Regularly showing what is being done, even before finalizing the deliverables, allows for involving the client and gathering their feedback. This enables continuous work improvement and strengthens collaboration. Finally, it is also important to close the project properly with an end-of-mission meeting to conduct a review and potentially mark the end of the collaboration in a positive way.

When you have doubts about an aspect of a project, should you communicate them to the client or does it remain internal?

It all depends on the type of doubt. If it is a lack of information or a need for clarification, we can certainly organize a meeting with the client or call them directly to obtain the missing elements. This allows us to quickly clear up areas of uncertainty. In cases where the client is unable to answer, we also have an "expertise hub" at the firm with more experienced people who can step in to help us. In our proposals, the different experts involved are identified beforehand. This ensures there is always someone competent to solve problems or answer technical questions. So overall, doubt is not a problem in itself because there is always a resource or an experienced person to support the resolution.

When delivering a project, is it necessary to detail the technical aspects and how you arrived at the results? Yes, throughout the mission, there are follow-up meetings with the client and the team. During these status meetings, we explain precisely the progress: who

worked on what, at what time, and what actions were carried out. This allows us to document and follow the entire project in a transparent manner. That is why regular communication with the client is essential: everything is detailed progressively over the course of the exchanges, and not solely at the end of the project.

What generative artificial intelligence tools do you use in your department and do you have the right to use them freely?

At our firm, we mainly use internal tools such as our firm's proprietary AI tool. It is a secure tool that functions similarly to ChatGPT, but also allows us to process client data within a controlled and compliant framework, with their agreement. It enables searching, sorting, and analyzing information.

At what stages of a project do you use generative AI?

I use it mostly as a support tool at different stages of the project. For example, when we receive a lot of documents, AI can help to get a quick initial analysis and an overview of the information. It is also useful for improving deliverables, identifying what is missing, or suggesting areas for improvement.

Does AI also allow you to acquire new knowledge on projects?

Yes, absolutely. The tool is very useful in that sense. It can be seen as a support comparable to a senior consultant guiding us through the analysis. However, it is essential to maintain a critical mindset and never trust it completely. AI remains a support tool. Personally, I prefer to work on my own deliverable first, then use AI to proofread it, improve it, or identify missing elements. It can also provide inspiration, but the information is not always 100% accurate, so it must always be checked and validated.

Do you think the arrival of generative AI has changed the skills required to be a consultant?

For the moment, no. We are still at the beginning of this evolution. At our firm, we are still in the learning and adoption phase of these tools. Many people do not use them fully yet, and training is regularly set up to support this transition. Therefore, the profiles sought after have not fundamentally changed for now. Companies continue to recruit the same types of consultants, while progressively introducing AI tools into daily practices, notably for project management or analysis. We are thus still in a phase of progressive evolution rather than a complete transformation of skills.

When a project is launched, do you have to define beforehand with the client whether you can use generative AI? Yes, of course, it is mandatory because we work with confidential data. We explain to the client that we have our own internal artificial intelligence, so the data is secure. It is not like using a public AI such as ChatGPT or others. Doing that is dangerous and forbidden. In any case, even with our internal AI, we must always ask for the client's authorization. They say yes or no. Generally, this is done at the time the proposal is signed, with the managers framing the contract and the terms of use.

Can the use of generative AI devalue your expertise in the eyes of the client?

I think the client relationship is very important in this case. We often need to hold status meetings with the client to show them the project's evolution. We do not always know what their reaction will be, but as we show them the progress, they understand how the work is done. In reality, they quickly see whether the work is solely generated by AI or if there is genuine human expertise behind it, involving proofreading, reflection, and experience. Even if AI is used, the result remains of high quality. Therefore, I believe the difference lies in transparency and project follow-up, not in whether AI is used or not. The client is not shocked at the end to discover a deliverable "made by AI" because they have been involved throughout the entire process. It is precisely by showing the evolution of the work step by step that we can integrate their feedback and adjust the deliverable. Personally, I have never encountered a situation where a client told me that the work had no value because AI was involved. Our managers also provide a lot of guidance to verify that everything is thoroughly reviewed and validated.

In follow-up meetings, do you specify which parts were done with AI or not?

No, not really. The authorization is given beforehand, so once it is validated, the client knows that we can use AI in the project. After that, they do not necessarily ask whether a specific part was done with or without AI. If the question is asked, we can answer, but it is not something we systematically detail. In any case, the client knows that the work is challenged, proofread, and validated by us. If there is something we do not understand or cannot explain, it shows immediately and we address it.

Are there clients who are resistant to the use of generative AI?

Yes, absolutely. If a client does not want AI to be used, it will not be used. That is why the client's agreement is essential right from the start.

Do you think clients will demand more and more transparency regarding the use of AI in missions?

That is a very good question. Today, I would say that everyone is already using AI, so this question depends heavily on the client's profile. In my projects, no client has asked me not to use AI; on the contrary, they see the added value: time savings, better quality, and greater efficiency. They know it helps consultants work better. So for the moment, there is not really a strong demand for transparency on every single use of AI. However, this could evolve in the future, particularly with cybersecurity challenges and risks related to digital transformation. If the risks increase, clients might ask for more control and more transparency regarding the use of these tools. Since we request prior authorization and use our own secure internal AI, clients do not necessarily demand more transparency for now. However, if risks intensify in the future, I think we will need to be even more explicit.

Let's imagine the case of a junior consultant who completes 80% of their project with generative AI, but in a completely botched manner, leaving the client dissatisfied. In your opinion, who would be responsible? The junior, their team, or the firm? In our team, a junior never works alone. They are always supervised by a senior consultant and a manager. Responsibility is therefore shared: it falls on both the junior and the manager. The manager has a supervisory role and must check the work produced while framing the use of AI. That said, if AI is used for 80% of the work but the junior has rigorously verified the results and the deliverable is excellent, it will not cause any problem. We only notice poor usage when the result lacks logic or does not meet expectations. Honestly, I have never encountered this situation because the levels of review are very strict. I cannot tell you what sanction would be applied, but the manager is there to ensure we do not get to that point.

In your opinion, does generative AI already have an impact on the relationship of trust between the client and their consultant?

No, because trust relies primarily on final client satisfaction and the quality of the work delivered. As long as we work in close collaboration with them throughout the project, there are no bad surprises at the end. Human contact and progressive validation of the work by the client protect this relationship. A negative impact would only appear if the work were botched due to a lack of control over the AI. Everything depends, once again, on the quality of the final deliverable.

Alright. So, my last questions are rather about the future of the consulting profession. They might be a bit more complicated. With the evolution of generative AI, will clients eventually think they could do the work all by themselves with the help of AI rather than hiring consultants and paying for it?

Well, I don't know if you've seen this, but I've already seen in the press that some consulting firms have started offering AI consultants. So these are not real people; it's the AI doing everything, delivering the report to the client, and doing all the work completely on its own. So it has already begun, and it's moving very fast. But I would ask, will it be possible in the future? I don't know, because people used to say we would never see artificial intelligence or robots, and here we are today. So I would never say it's impossible. However, currently, I think we are not yet ready for the client to use AI for everything. Today, we still need expertise. I'm really telling you this from experience: we use artificial intelligence, but if we don't have meetings with the right people to get the right information, the feedback, and the feelings, the AI will never be able to know them. People are not going to respond the same way as they would face-to-face in an interview. And that is truly where the importance of the consultant lies: going out to meet people, conducting interviews, and gathering information. The information we gather is incredibly important for the report. These are data points that nobody can find any other way. AI can help with documentation, but it cannot do that specific work. So today, I find that AI is still just a support tool for humans and cannot do everything. But in the future, I don't know how far it can go: will it be able to read emotions, or conduct interviews with humans? Maybe. We might end up in a world where AI can hold meetings directly with people. But for the moment, that is not possible. Today, it's mostly about data collection and report production, but it's not yet complete. A human is still required for this work.

During my research for my thesis, some authors state that the consultant's profession will transition from creating solutions and deliverables toward a role of verifying AI outputs. Do you agree with this? Can you repeat the question for me? So the transition is going from creating solutions and having the expertise, to just becoming a consultant who verifies what the AI has produced?

Yes, exactly: verifying AI outputs, meaning the AI would be the one generating the solutions. I feel like saying that in certain cases, we are already a bit into that. There are people who use AI to create a deliverable, and then they just proofread it to see if it's okay or not. But it depends on the tasks and missions. I wouldn't say that for a final report, because at that level, I haven't yet seen a 100% correct result where you would just need to proofread. You still have to modify

a lot: the output, the layout, etc. So for the moment, on final deliverables, we are not there yet. In the future, maybe yes. But for that to happen, the AI would need to be able to hold meetings directly with the various people on the client's side, gather information, and clarify doubts. So we are not there yet, but if that becomes possible, then yes.

How do you stay relevant? What will be the added value of the consultant in 5 to 10 years faced with generative AI? Well, the first piece of advice is that our firm prepares us well for this, notably by providing us with training to use artificial intelligence. I think there will be jobs that disappear, but also jobs that will be created, so it will balance out. We can already see it evolving year after year. Before, there was a lot of fear, but humans adapt very easily. The important thing is to always stay up to date, to train yourself, even on your own, and to look at the new opportunities that appear. Even if AI evolves in 5 to 10 years, humans will always find something else to do. New positions will be created, especially in technological sectors, where there will be plenty of opportunities. Therefore, you must always keep learning, be agile, and adapt to new opportunities and sectors where human input will remain indispensable.

Okay, that's it, that was my last question. I hope it didn't take too much of your time. No, no, not at all. I hope I was able to provide some useful insights.

Interview RISK03

I have several questions, notably about generative AI and its impact on consulting, particularly on the client relationship. You just need to answer as spontaneously as possible. Can you introduce yourself: where do you work, what is your role, and how long have you been there? My name is [Consultant's name]. I am currently a consultant in Risk Advisory at [Company's name]. I haven't been here very long, about a month. Before that, I did consulting at another firm in digital strategy, marketing, and commerce. I worked there for a little over a year and a half, right after my studies.

What did your missions consist of at your previous firm, and today at your current firm?

At my previous firm, our missions focused mainly on digital optimization. We looked for processes that could be automated using CRMs or other technologies. So it was very much oriented toward digital transformation. Part of our work also consisted of training clients on certain tools, notably on CRMs. It was truly support and consulting, with a strong technological dimension. We also analyzed the economic impact of these automations, particularly in terms of time and financial gains: what the company could concretely get out of it. We also accompanied certain companies at a European level to analyze how the implementation of a new technology could modify their way of working. Regarding [Company's name], it's a bit more recent so I have less hindsight, but I am already on two missions. There, it is pure Risk Advisory: we analyze risks within a company's processes. Risks can exist at several levels, whether in decision-making, operations, or financial management. We perform an assessment, like an internal audit, to understand their way of working and identify areas of weakness. For example, recently, we worked on an audit related to the insurance of a real estate company. The objective was to verify if all their buildings and assets were well insured, what types of policies they had, and what risks existed if certain assets were not covered. Then, based on this inventory, we formulate recommendations and support them in optimizing their risk management.

Okay. So if I understood correctly, whether in Risk Advisory or in your old position in digital, you had to produce reports or presentations at the end of the missions? Yes, all the time. By the way, are you at your firm on a permanent contract or doing an internship?

I am doing an internship.

I figured as much. In any case, in all consulting firms, we work extensively with reports. You have surely already noticed this during your internship. We do presentations, documents, working papers... We always document all the steps of the project.

My next question: in your opinion, what is your expertise based on as a consultant, in terms of skills?

I would say that at a junior level, you aren't necessarily an expert right away. You are more in a support role for more experienced consultants. On the other hand, I think a junior consultant must above all be very flexible and capable of adapting quickly. They must be able to integrate into any environment. They must also have a quick capacity for analysis and a concrete way of thinking, in order to quickly understand a situation and propose solutions rapidly. In short, a junior must be a bit like a chameleon: capable of working in different industries, adapting, and having a quick understanding of very varied contexts before gradually specializing.

And when you were doing digital optimization, did you need to code or implement something?

No, the implementation was not on our side. We had colleagues who took care of development and implementation. We intervened more on the post-implementation part.

Okay. When you realize a deliverable in a project, what do you base yourself on the most: your accumulated experience over projects, your theoretical knowledge, your intuition... or a bit of all three? I would say a bit of all three. Intuition is useful, but we always avoid starting from scratch. When you have to produce a report, you try to reuse what you have learned on other projects, feedback from colleagues, or similar situations that you have already encountered. You look for similarities between problems already solved and the one you are facing. So yes, we often start from an existing base, which we adapt to the context.

In your opinion, what are the essential elements to build a relationship of trust between a consultant and their client?

It's a bit of a cliché in consulting because people often say there is a lot of "bullshit" and that solutions are not always concrete. But in my opinion, to ensure a good relationship with a client, you first need to clarify right from the start what their expectations are: what they are looking for, what they want to achieve, and why they are calling on us. In Risk Advisory in particular, the biggest job at the beginning is truly to understand the client: knowing their business, analyzing their activity, understanding their issues, and identifying their pain points. You need

to have a clear vision of what you are getting into. Then, you must not "oversell," meaning promising more than what you are capable of delivering. You have to be realistic about what you can produce and the expected results. Finally, you must maintain regular and transparent communication. The missions that go the best are those where there is a real collaboration with the client. For example, if we need documents or information, we need to be able to call the client easily, explain what we are looking for, why it is necessary, and clarify certain points with them. This cooperation is very important.

When you do your presentations or documentation, do you explain all the technical details of the project, or not necessarily?

We try to be as concrete as possible, but it depends on the level of detail required. You also need to know how to adapt the content to the client. That is precisely why you need to know them well: to understand what is relevant to them and what is not. For example, in a working paper, we won't necessarily detail the time spent on each activity if it's not relevant. If we did testing or sampling, or if certain elements were still in a draft state, it's not always useful to include everything. On the other hand, in the documentation we provide to the client, the essential information must appear: what we did, how we did it, and why we did it.

Right. My next questions concern the use of generative AI. Did you already use generative AI at [Company's name]? If so, which one? And the same question for your current firm.

To be completely honest, in terms of AI, we use it in all consulting firms. But since we work with clients, its use is very limited. We cannot use open AIs like ChatGPT or other public models because there is a risk of information leakage, and certain data is confidential. At my previous firm, I remember we had an [Internal LLM]. And at [Company's name], it's the same. These are generally tools developed in-house, so "in-house" AIs. We can use public models like Gemini, but not for direct client work. It's more for doing general research on public information, as if we were doing university work. It's frequent that we need this type of information to understand a sector or get context on a client.

Precisely, you said you used them to better understand the client's context. For what other tasks would you use AI in consulting?

First, for anything related to rephrasing or translation. In Belgium, for example, in certain regional contexts, we can come across documentation in a language that is not our mother tongue. So AI is very useful for quick understanding. It can also be used to do a quick market scan. This isn't necessarily for deliverables directly intended for the client, but rather for internal

missions, like market research. For example, when we look for new potential clients, AI can help to quickly analyze a sector or an industry. It can also be useful for identifying people: knowing who the CEO or CFO of a company is, following public information like recent acquisitions, or strategic changes. Finally, with internal tools like Copilot, which operate within a secure environment, we can analyze large documents. If we have a 1,000-page document, we don't have time to read everything in detail. In this case, AI is very useful for synthesizing or quickly finding specific information we are looking for.

I'll ask the question the other way around: for what tasks do you not use AI? In other words, what are the tasks for which generative AI cannot replace you?

First, everything concerning the relationship with the client. Even if you give the AI the transcript of a call, it cannot interpret the emotions, reactions, or concerns of the interlocutors. It cannot truly understand what is important or a priority in a discussion. So, for that, you need to be present during the exchange, whether on a call or in a meeting. Ideally, we also try to go on-site, to the client's premises. And that, AI cannot replace. I think AI is a good assistant. I know it can replace certain tasks, notably those carried out by junior profiles. But the experience a consultant acquires after several missions, AI cannot have it the same way. For example, making connections between different companies within the same industry, or comparing situations experienced in different contexts, is something that comes with personal experience. AI can help to draft texts or give a working base, but a final deliverable cannot be produced solely by AI. It absolutely requires human proofreading, verification, and validation, otherwise it's not possible.

Right. Exactly, you were talking about junior profiles. It is often said that we learn on the job, by doing things ourselves. Do you think generative AI has an impact on juniors at that level?

Yes. From what I've been able to observe, juniors use AI a lot on a daily basis. The problem is that they might stop there and not go further in their reflection. They won't necessarily seek to understand more or deepen things on their own. So it has an impact on the development of their skills because they rely too much on AI. In addition, the main impact is that we need fewer people to do certain tasks. And therefore, juniors are going to have to produce more, and especially produce more qualitative work. We also see it in client expectations: they demand more quality because they know that now we have AI assistants that can help us constantly. The level of requirement is therefore higher, and that is the main impact on junior profiles.

When you do a presentation or documentation for your clients, do you explicitly mention if generative AI was used for certain aspects of the project?

It depends. If, in a report or a presentation, we give a figure, we are strictly required to indicate where it comes from. If this figure comes from an AI, then it must be specified. On the other hand, do we explicitly mention, like in a thesis, that we used AI for this or that task? No, in general, we don't do that. In the corporate world, since clients use AI themselves, they assume that everyone uses it. It's not surprising. And some clients even clearly tell us that they use AI themselves. So they naturally expect to receive a deliverable faster since they know we have additional tools at our disposal. Overall, clients prefer that we have access to these tools because it makes the work more efficient.

Right. So if I understand correctly, it doesn't necessarily impact the client's perception of your expertise.

No, because in reality, AI cannot provide a solution that is genuinely adapted and specific to the client's problem. It mostly gives general answers that might sometimes work, but rarely in a completely satisfactory way. AI can primarily help execute certain tasks faster, but it won't build a customized solution. Therefore, clients do not necessarily see us as less experienced. However, I think they are becoming much more suspicious. They challenge our solutions, our proposals, and even our data a lot more. Today, you can no longer present a benchmark or a data point without specifying where it comes from, how it was calculated, and what sources it is based on.

So they are really pushing for an optimal, or even perfect, result? Yes, exactly. I would say they are more demanding and more rigorous. They pay much closer attention to the information sources we use to propose a solution to them.

Do you think generative AI impacts the relationship of trust between the client and the consultant?

No, not directly. Trust does not deteriorate simply because we use AI. On the other hand, if we sell a poor-quality deliverable and it is obvious that it was generated by AI, then yes, trust can deteriorate. For example, if you start sending reports or emails where you can clearly see the content comes from an AI, the client is naturally going to be more suspicious. It's not that they stop trusting you immediately, but they will challenge you much more because they feel like you didn't produce the work yourself. And in the end, the client pays for a consulting service; they do not pay for an AI. If they only wanted an AI, they could buy a subscription and do the

same thing themselves. So the impact on trust depends mostly on how the AI is used. If it is abused, then yes, it can harm the relationship.

Do you think consultants will have to adapt and be more transparent about whether or not they used AI during a project? Yes, clearly. But at the same time, AI also opens up a new market and new opportunities. There are more and more projects where companies tell us: "We don't use AI yet, but we know it exists, we know it can be useful, how can we use it to improve our processes?". So it creates new consulting missions around AI. As I said, an essential skill of a consultant is knowing how to adapt: adapting to the client, to their industry, but also to the tools we use. If a new tool allows us to improve quality or efficiency, we will inevitably adopt it. It's as if flying cars existed tomorrow: we would prefer to use them to arrive faster at our clients' locations. In consulting, we are never going to ignore a tool that allows us to get a better, clearer, more precise, or higher-quality result. We constantly adapt, and we are already doing it.

Right. You mentioned earlier that care must be taken so that the use of AI isn't visible, and that the content must be thoroughly verified. If despite this there is a problem, we cannot "fire" the algorithm or the AI. In that case, who must take responsibility for the failure: the consultant, the team, the manager, or the entire company? You mean if a report was generated by AI?

Yes. For example, if a report is generated by AI and it is completely wrong, opposite to the client's expectations, and does not match at all what was requested. In that case, there are two levels of responsibility. First, toward the client, the responsibility lies with the project manager. They are the visible person, the one carrying the project, who sold it, and who is responsible for it. Then, internally, the responsibility goes to the person who produced the deliverable, potentially a junior. The first problem would be that this junior generated entire content using AI without thinking, without checking if it was correct or not. The second problem would then be the N+1, meaning the manager or the senior who did not review the content. This means several people did not do their verification work. In reality, this kind of situation shouldn't normally happen. First, because internally we are quite open about using AI: personally, if I use AI for a specific task, I will tell my senior so they know exactly what I did and how. We are quite transparent about our use of these tools. But if despite everything a project was totally wrong because of AI, then the responsibility would be shared between the manager and the junior, and more broadly among the entire project team.

Okay. Last question about the future of the profession: you said earlier that clients can pay for an AI subscription and try to do the projects themselves. Why, in the future, would they still pay for consultants? Because what we sell is expertise. AI, in the end, remains a tool based on public content. In a way, it mostly replaces Google: instead of searching for information by typing a question on Google, you ask the question to the AI and you get a more direct answer. It saves time, but it remains based on public information. However, consultants do not only sell data. We do not sell raw information. We use this information to produce a deliverable tailored to the client's context. AI can give sources or information, but it doesn't necessarily know how to interpret them correctly, nor how to apply them to the right context. And above all, AI does not have the experience that a consultant develops after several years or several missions in the same sector. A consultant knows what works, what has already worked, what failed, and why. It is this knowledge that AI cannot replace.

During my literature review, many authors state that even now, and even more so in the future, consultants will transition: they will no longer necessarily create the solutions themselves, but rather become verifiers of the outputs produced by AI. Do you agree or disagree with this? In the future, it's possible, yes. I think I can agree with this idea.

Okay. And how do you see the profession evolving in five to ten years? In five to ten years, I think it won't solely be about verification. It will mostly be the fact that we will be able to complete many projects in much less time. We will be able to produce the same deliverables faster, so the volume of work will increase. We will be able to take on more projects in parallel and manage a larger volume. Many small tasks, such as verifications, investigations, calculations, or controls, will be able to be performed much faster and in a more automated way. Today, teams are often limited. When your whole team is staffed full-time, you are forced either to recruit (which is not always easy) or to refuse missions. So we are sometimes limited by the availability of resources. I think AI will precisely allow us to overcome this limit and increase our capacity to deliver more.

Okay. My last question: in the future, how do you remain indispensable as a consultant faced with generative AI?

What should be the added value of a consultant? For me, the main added value will be having a critical mindset, knowing how to give an opinion, and above all, taking ownership. Humans need someone to take responsibility. An AI cannot be held responsible. If something goes wrong, especially at a board of directors' level, you cannot shift the blame onto the AI: a person

has to assume it. And in today's society, many people seek to transfer responsibility to someone else, and notably to a consultant. So that is an essential point. Next, I would say that to remain indispensable, there is also everything related to networking and human relationships. This trend is not going to change: people prefer talking to a human rather than an AI. We already see it with call centers: most people prefer to talk with someone real because it is more reassuring than a robot. So I would say that the consultant's added value will be the ability to communicate, to create a bond, to have a network, and to maintain a critical mindset rather than simply following what the AI suggests. very

Okay, that was my last question. Thank you much.

Interview TECH01

Well, as I explained in my message, the thesis I am writing focuses on the use of generative AI in the consulting profession, and on how it influences relations between consultants and clients, notably in terms of trust, transparency, and human expertise. To start, can you introduce yourself? Where do you work? What role do you occupy in the company? And how long have you been there?

Yes, of course. My name is [Consultant's name]. I have a background in business analytics. I have been working since September 18th, to be precise, as an AI and Data consultant at [Company's name]. So it has been nearly 7 months that I've been working with them, and I have been on a mission at a client's site since the beginning of October.

Alright. And what do your projects or missions consist of?

Currently, for my client, I work mainly on test automation. Concretely, this consists of automating everything related to API flows and non-API flows tied to websites and applications. For example, if a client needs to transfer money from one account to another, the API flow corresponds to the transaction taking place correctly. The non-API flow, on the other hand, corresponds to the case where the client cannot manage to make this transfer. So our objective is to automate these sets of tests.

Alright. And to do that, do you have to code? Do you need to install a module? How does it happen concretely?

It's a bit complicated. In my case, I am currently working for a client that is a young bank, which needs to establish itself in the banking sector. For this specific project, we use a no-code tool instead, in order to carry out these sets of tests. However, I know that with other clients, the approach is more oriented toward coding for this type of test. And for me, as a consultant, the objective would be to implement artificial intelligence into this testing process. But it is not a simple task. Sometimes, the client is reluctant regarding these tests, and especially regarding this transition toward artificial intelligence.

This might sound like a silly question, but if they are reluctant, why do they contact you then?

It's a bit complicated, because they contact us primarily to get advice. But sometimes, they are not necessarily ready to hear the advice we have to give them. Moreover, as a junior, I have been with the firm for about seven months, as I explained. So I was staffed on a mission that is

not necessarily directly linked to artificial intelligence. But overall, the objective of a company, and notably of a Big Four firm, is not only to help the company with the tasks it must accomplish, but also to advise it on its future and on how to implement artificial intelligence into its operations. Now, some clients are reluctant, others are less so. You have to deal with it.

And is it part of your mission to then explain to your client how to use the solutions you provide, or not necessarily?

Yes, that is indeed part of the job. In consulting, there is an implementation side, so we are the ones putting the solution in place. There is also a pricing aspect, so you have to be capable of determining the price of the services. Obviously, when you are a junior, you don't necessarily master this part. In my case at [Company's name], it's rather the role of the managers, who will help with pricing. But in any case, you have to know how to pitch the solution and teach people how to use it. Now, to put this into context, I work with people who do not touch a single line of code. So if we implement a solution, it can be complicated to transfer this knowledge to the client. You have to compose with the environment in which you work.

Alright. So your current client is a bank, if I understood correctly. But in general, do you have other types of clients or do you work solely with banks?

So, I work in what we call FS, meaning Financial Services. Therefore, we work with banks, insurance companies, and wealth managers. I could give specific examples, but I cannot, so I won't do it. But there you go, those are the types of financial institutions. Let's say that banks, insurance companies, asset managers, and wealth managers are generally large companies. These are institutions that we often see in advertisements in Belgium, and that many people use on a daily basis.

That makes sense. In your opinion, what are the essential skills one must master to practice your profession?

I am talking about both soft skills and hard skills. That's a very good question. I would say that regarding skills, we aren't necessarily looking solely for profiles that are highly competent on a technical level. We are looking above all for passionate people who are ready to invest themselves and put in the necessary hours to train and learn. To give you an example, I have colleagues who did a PhD in physics, others a PhD in chemistry, and who today work in financial services consulting. Yet, fundamentally, they aren't necessarily capable of writing code. What we look for most are people who have a good capacity for reflection and who are able to combine this reflection with a business understanding. This then allows us to propose

solutions to clients, or even to develop internal solutions at [Company's name] that will enable the company to grow.

And in your opinion, what are the skills that allow for maintaining a good relationship of trust with the client? First of all, I would say it's honesty toward the client. It is possible that after a certain amount of time, when working on a mission where we wanted to establish ourselves, we sometimes have to take a step back and say to ourselves: "Okay, this specific task, we are ultimately not going to be able to carry it out," because we were too optimistic compared to what was requested. So honesty is an essential element. Next, there is also the confidentiality aspect, which is very important. We receive an entire training session on this topic: for example, we must not post on social networks, on LinkedIn or elsewhere, which client we are working for. There is a logic of anonymization that plays a major role in our work. That is actually why I almost made a mistake earlier when talking about potential clients we have at [Company's name]. When you propose solutions to your clients, what do you base yourself on?

Is it your personal experience, what you have learned at [Company's name], your intuition, your skills? What truly guides your project and your mission? I don't have the word in French, but... what's driving your project and mission?

Yes, I think I understand what you mean. You want to know more about how we deal with our knowledge and our experience to propose something to the client. In fact, what you need to know is that sometimes it's the client who comes directly to a consulting firm saying: "Here, I have this task to perform. Can you propose something to me?" In that case, we call it an RFP, meaning a Request for Proposal. The client asks: "I have this task to do, do you have a solution to propose?" And at that point, the client will go around the consulting firms to see who offers the best proposal, not only in terms of price but also in terms of the ability to correctly answer the request. That is the first part. On the other hand, internally, we also develop solutions that can then be proposed to clients. So we are going to look into current technologies related to artificial intelligence in order to identify solutions that could correspond to our clients' needs. For example, currently, clients do not necessarily think about the security aspect linked to Large Language Models. Yet, this is a topic that all consulting firms are working on today in order to be able to offer something to sell to the client.

So in fact, you also work on LLMs in part, is that right?

Well, we don't design them ourselves, but we use them. Those who directly design LLMs earn several hundreds of thousands of euros a year. We are not in that part of the business. On the

other hand, we conclude contracts on this basis. We use what we call cloud services, which give us access to these LLMs. This allows us to build an architecture that can then be put into production and correspond to what our clients are looking for. These cloud providers offer different language models. Typically, we find AWS, Azure... We don't necessarily use Google, but these types of providers offer several different LLMs. And then it's up to us to choose which one to use, depending on the tasks requested by the client.

We already talked about this a little bit earlier, but do clients sometimes ask you to explain how you arrived at a solution or a recommendation? Or do they prefer just to get the result without necessarily knowing how it went?

In my personal case, it's not necessarily something I see a lot. But from what some colleagues report to me, there is really a documentation role that must be carried out from end to end. This means that we have to explain how we thought of the base idea, then how we arrived at the proof of concept, or the minimal viable product, and then how we move through all the steps up to production and the final use of the tool. All these stages are documented and explained to the client.

Alright. This is a question purely related to the consulting aspect: does it happen that you have doubts when you produce solutions?

And if you have doubts, do you talk about them to the clients? For example, if an aspect of the project is not feasible or is too complicated. Do you discuss it, or not at all? Yes, absolutely, we talk about it. In companies, we often apply what we call the Agile or Scrum method. So we move forward step by step: we start with a large project, which we cut into several small projects, then we implement these small projects one by one. And if, for example, right from the first small project we encounter a problem, then we consult the client, explain the situation, and look for a solution together. Then, once we have found a solution for project 1, we move on to project 2, and so on. And in the end, the combination of these projects constitutes the large project that will be delivered at the end. So there is always a consultation step between the expectations expressed by the client and what we have the capacity to offer.

Alright. Now, I will move on to the main topic of my thesis: generative AI. First question: which generative AIs are you authorized to use at [Company's name]?

That's a very good question, because if you had asked me this question about 4 or 5 months ago, there wasn't really a problem regarding this topic. But 5 or 6 months ago, our firm, at a global level, took the decision to restrict the use of these services. Global management decided

to limit employees' access to AI tools. For example, ChatGPT is restricted, Perplexity is restricted, Gemini is restricted, and there are others as well. I don't have them all in mind, but the majority of providers are restricted. Today, in the company, what we mainly use is either Microsoft Copilot, which is widely used, or, as we see more and more in companies, solutions developed internally. At [Company's name], we have a solution called [Internal LLM]. However, we have to be honest: Microsoft Copilot and [Internal LLM] are two totally different tools. They do not have the same role and are not made for the same uses.

Can you elaborate? Because I don't know [Internal LLM], I am not at [Company's name], as you have noticed.

Yes, of course. We all know Microsoft Copilot: it allows you to ask simple and general questions. It can use Internet searches, or help edit Word documents, PowerPoints, etc. [Internal LLM], on the other hand, is more of an LLM type model developed by our firm. It is not usable in a client context. Its role is mostly to help consultants on different aspects. For example, it can help with the formatting of Requests for Proposal. It is also used for internal documentation. Many companies, whether at our firm, at clients' sites, or in other consulting firms, are today developing their own chatbot based on their internal documentation. The objective is to help employees find their way more easily through the often vast documentation available on internal platforms.

And you personally, for what purpose do you use [Internal LLM] or Microsoft Copilot? For what tasks?

In my case, since I am a junior, I have less need to use [Internal LLM] because I don't have to do RFPs. So [Internal LLM] is not really a tool I use much. On the other hand, Copilot is used much more. For example, we can use it to generate slides that are specific to our work. But you have to be careful: even if you use it in a client context, you must anonymize the data. These past months and years, there has been a lot of use of Microsoft Copilot and other tools like ChatGPT without anonymization. And today, we find ourselves with annual reports, as we have seen in the press, that are generated entirely by artificial intelligence tools, without anyone even wondering if the work was correctly done. This poses a real problem. And the measures taken 5 or 6 months ago are linked to this situation.

Alright. So you use it mostly for slides, but not necessarily to help you with the implementation of tools for your clients?

Yes, for that too. But in that case, there isn't really the anonymization problem. A piece of code remains a piece of code, and it could be used for just about any client. So yes, for implementation, AI helps a lot. We also use it for code implementation.

At what level exactly, if you can specify? Because you mentioned several tools, notably a no-code tool. Does AI also help with this type of tool?

In that case, it's rather if we need to perform a search on this tool that we can use artificial intelligence. But since it is a very specific no-code tool, there is very little documentation available. And artificial intelligence models are not trained on that specific documentation. So for the no-code tool, no, AI doesn't really help: it is mostly personal learning. I followed about a forty-hour training course on this tool, so now I am trying to master it. On the other hand, as soon as you have to implement things in code, it is ten times simpler to use artificial intelligence because it can advise effectively on the implementation.

By using artificial intelligence, does it allow you to have a better understanding of what you are supposed to do? Or do you already have a clear idea in mind and it's mostly a matter of productivity and saving time?

Sometimes it's a mix of both. There are times when we have a task to do, but we have what we call the blank page syndrome: we don't know where to start. In that case, personally, I use generative AI step by step. The first thing is to clearly define the problem. Then, piece by piece, we structure our thinking to develop a solution. So there is that aspect. There are also subjects where going directly to consult the documentation linked to the code or the libraries used would have taken me a huge amount of time. In that case, going through artificial intelligence allows me to save time. Its ability to summarize information is very useful, and the time saving is enormous.

I am asking the question, but I already more or less know the answer: do you find it positive to use it for these purposes?

Yes, it is positive, but you must maintain a critical mindset and not blindly copy what is written. You must not be what we call a "script," meaning someone who takes the provided code, implements it directly, and tries to run it without thinking. It doesn't work that way. You have to read and understand the code that the AI provides. And that's where it connects to what I was saying about my colleagues coming out with a PhD, for example in chemistry: these are people capable of using this tool with a critical mindset. They don't necessarily have the technical level in coding at the start, but they have this capacity to understand what the AI proposes and

especially this ability to learn quickly, to be fast learners, and to see very quickly whether what is proposed is compliant or not.

So in the end, is it also a way to learn and extend one's skill set?

Yes. Afterward, personally, I do not base all my knowledge solely on that. But indeed, it allows you to learn. For example, if I ask what the complete stack of an end-to-end project should be, and I ask: "Explain to me what I should put in the front-end and in the back-end," the AI will explain to me step by step what needs to be integrated, as well as the structure and architecture of the project. So yes, it helps a lot. A year ago, I didn't have all the skills and all the knowledge that I have today, and it is partly thanks to AI.

According to your experience with generative AI, what aspects of the consultant's profession, and more specifically of your field, cannot be replaced by generative AI?

What is it unable to do? For the moment, what is mainly developed are LLMs. And an LLM is essentially a string of words where each word is selected based on the probability of the next word: at each step, you choose the most probable word, which ultimately produces a text. But what cannot be replaced is, first of all, the relationship with the client, the client relations. Next, there is always this essential role of verification: verifying what the artificial intelligence proposes. And we don't necessarily think about this, but AI also opens up new professions. For example, there are all the security aspects linked to artificial intelligence that need to be developed. Because it's fine to use it, but guardrails must be put in place to prevent the AI from providing inappropriate answers. There is also the topic of artificial intelligence governance. You also have to understand the client's vision regarding artificial intelligence. And therefore, all this structuring around AI governance, security, and strategy is not at all replaceable by artificial intelligence today. Now, given the speed at which it is evolving, we will have to see if in the coming years we won't reach a point where we could be 100% replaced. But for now, we are not there yet. There is still this critical mindset that AI cannot replace, in addition to client relations.

Do you find that you, or other consultants you know, spend more time verifying AI outputs to detect hallucinations rather than producing the analysis directly yourselves?

I think that was more the case 4 or 5 years ago, at the beginning of artificial intelligence when there were indeed a lot of hallucinations. But today, that is less and less the case. There is therefore a greater trust toward artificial intelligence. That said, you must always maintain a critical eye.

Alright. Now, when you present solutions to your clients, do you tell them that you used artificial intelligence for certain aspects of the project? Or are they not interested in that, and only the result counts? I think clients are no longer naive. There is a certain transparency around the fact that we use artificial intelligence. Now, I personally wouldn't know how to answer this question in detail because I use AI, but probably less than some colleagues who have to implement more things.

Right, but then do you personally talk about it to your clients or not? Is it an obligation at [Company's name]? Or is it not necessary?

One of the obligations at [Company's name] is transparency. Now, in the deliverables I have proposed, the client is more attentive to the final result than to whether or not we used artificial intelligence tools. But there is still this notion of trust: the client trusts the consultant, notably because they assume the consultant maintains a critical eye on what they produce. Normally, that is.

Precisely, if you tell the client: "I used artificial intelligence to do A, B, C," do you think that's a good thing? Or could it discredit you as a consultant? Is saying that you use AI positive or negative?

I don't think it discredits you. If we don't use it, we are falling behind compared to other companies and other consulting firms. We shouldn't hide away from that.

At the very beginning of the interview, you mentioned certain clients who were reluctant to change and to technological advances linked to generative AI. How do you try to make them understand that it is in their interest?

Sometimes it's difficult. Yes, it's in their interest, but there is often a difference between the person you are going to pitch the solution to and the final implementation. The person you present the solution to must then go through various people internally, and it happens that there is a blocking element that prevents going all the way to the solution we would truly want to propose. It's not always the case, I am dramatizing a bit, but it happens. For example, at my current client, my manager is not really comfortable and has no knowledge of the world we live in today with all the artificial intelligence, so it scares him a bit. But he is not the only blocking element. There is also the security context specific to financial services, and therefore to banks: all access is controlled. For example, I can only access the systems from home, and I cannot work in a café near my place. So there is this whole security aspect that makes the implementation of artificial intelligence solutions more difficult.

Alright. In that context, with the evolution of artificial intelligence and GenAI, do you think clients are going to demand more and more transparency regarding deliverables? Are they going to demand more transparency? There is already transparency in consulting firms. Now, do all consulting firms truly practice transparency? I am unable to answer. But I think as long as the final product and the solution that were pitched at the start are delivered, the client does not necessarily demand more transparency. We work with banks, and I think most of them, even if I cannot say all, also work with people based in India. Are those people completely transparent about using artificial intelligence? I don't think so. In the end, as long as the deliverable is there, we don't necessarily question the way it was produced.

Alright. We already talked about this a bit earlier, but several corporate executives have produced reports, and in the press, it is said that these reports contain a huge number of hallucinations, notably because they prompted the AI poorly or because they solely used AI. In that case, in your opinion, who is responsible? The individual consultant, the N+1, or the consulting firm as a whole?

I would say that everyone is responsible. First, the consulting firm because it has a role of information, prohibition, or regulation regarding the use of artificial intelligence. Next, the N+1 can also be questioned because there is an education role to play at the management level toward the consultants. But above all, it is also up to the consultant to be aware that the use of artificial intelligence can be harmful if it is not questioned. And globally today, in all consulting firms, we are becoming increasingly strict regarding the use of these artificial intelligences. I think those famous reports containing hallucinations sounded the alarm on the necessity of framing the use of these tools.

Alright. We are going to move on to my last questions, which focus more on the future of the consulting industry. Before, in the past, it was simpler: a client asked a consultant for a solution because they did not have the expertise internally. Now, with generative AI, couldn't the client stop paying a consultant and try to produce a solution all by themselves?

In the near future, I am not sure. On the other hand, in 20 or 30 years, it is probable. Today, why don't I think this is the case? Because there is a whole research aspect that banks do not carry out. Notably, all the research linked to new technologies is something that banks do not really do themselves. A large part of the consultants' work consists precisely of doing this research and putting solutions in place, sometimes even before the client explicitly comes to

ask for something. But will this always be the case in 20 or 30 years? I think we are heading toward increasingly specialized AIs by domain: in the banking sector, in medicine, etc. Each sector will probably have its own specialized LLM.

Alright. And for now, is generative AI mostly a tool or already a potential competitor to the consultant profession?

It is a tool, and it will remain a tool. I think the profession of consultant is not going to stop from one day to the next. There will always be a need to understand what is being done behind it. Today, we can generate code, but we cannot generate the entire infrastructure of a project. That is still impossible. For example, the DevOps side, meaning setup, development, and operations, involves all the infrastructure management on servers, and the concrete transition between a proof of concept and a solution in production. That, artificial intelligence cannot do today.

Alright. And how do you see the consultant profession evolving in the next 5 to 10 years?

That's a very good question. I have the impression that as a consultant, we almost live day by day. With the evolution of artificial intelligence, we don't really know what tomorrow will be made of. I am also speaking as a person who entered the working world only 6 or 7 months ago. But if we go back 4 or 5 years and had asked ourselves this question, could we have predicted that we would move from generative AIs to the agentic side? It is not predictable. I would be unable to say what the consultant profession will be in 5 to 10 years. But if we look at artificial intelligence from a more global perspective, a consultant could help a company implement artificial intelligence into a robot. For example, tomorrow, instead of cooking yourself, a robot could, following our prompts, go cook, take food out of the fridge, start the machines, and automate this whole process. I think this type of automation could arrive in the near future. So the consultant profession will always exist. But will the consultant profession specifically for financial institutions always exist? That is the real question.

Would there be a risk that we transition from the consultant profession as we know it today toward a role of a simple verifier of AI outputs?

For certain tasks, yes. I wouldn't know how to answer with absolute certainty because indeed, we are getting closer and closer to a verification role. But there are also tasks for which we still need to build the solution ourselves, and potentially get advice from artificial intelligence, but without relying 100% on it.

Interview TECH02

Regarding client-consultant relations, notably in terms of trust, transparency, and human expertise, I have prepared several questions that you just need to answer. With pleasure.

To start, can you introduce yourself, specify your role, and indicate how long you have been in this position?

My name is [Consultant's name], I am 27 years old. I started in this field at another firm. By the way, I see that you are currently at that firm: are you doing an internship or something else?

Yes, I am doing an internship.

I started as an intern at that firm, then I stayed there for three years in the digital team, working on consulting missions in technology, cloud, and generative artificial intelligence in particular. Today, I have about three and a half years of experience. I am currently working at [Company's name]. In between, I also worked in a small startup, still as a consultant, within a small team. At [Company's name], I have been on a project for almost two months in the Digital Solutions team. It's a model close to consulting, but more oriented toward engineering. I am mainly an AI engineer, with a focus on robotics and process automation.

What types of projects or missions do you generally work on?

I have a strong background in the public sector, notably at my first firm for about three years, on missions related to the European Commission in Brussels. I worked with local governments and different offices of the Commission on digital topics. Then, I worked in the banking and financial sector for about a year, on AI and data science topics. Today, I work in the private sector. I can intervene in different fields depending on the projects, notably defense, banking, and other traditional sectors. My current project is mainly in the private sector.

Concretely, do you implement AIs or chatbots?

Yes, I mainly develop conversational artificial intelligence agents, a bit like ChatGPT-type systems for companies or specific use cases. I have also worked on data science algorithms to process and analyze data, notably in the banking sector. Currently, on my project, I work in a rather particular configuration: I am largely replacing humans. It's strange to say, but I am developing systems composed entirely of AI agents. I am the only technical person on the project. Usually, we work with developers, designers, and various traditional roles. Here, there is only a project manager and a frontend developer colleague, who hasn't contributed yet because he is on vacation. Everything else has been built via AI. My role is therefore to design

these agents, initialize them, give them a form of "personality," and then make them work on a daily basis to execute tasks, code, answer questions, and replace certain human functions like those of an IT architect or a UX/UI designer.

So you code and then you do the implementation?

Exactly. I code, analyze, and also do research. I cover both the business and technical aspects, from understanding the problem to the solution. I also do a lot of architecture, systems engineering, and documentation.

In your opinion, what are the essential soft skills and hard skills for your profession

Regarding hard skills, I would first say mastery of Python and more broadly programming languages. Next, cloud technologies like AWS or Azure are essential because they allow for the deployment of AI, machine learning, and data science services. I think the future of AI is strongly linked to the cloud, so these skills are fundamental. You also need solid foundations in logic, mathematics, statistics, and algorithmics, even if it's not always highly advanced theory on a daily basis. Understanding the mechanisms of systems and algorithms remains important. Finally, there is the whole ecosystem of data science and AI tools, such as notebooks, machine learning, model fine-tuning, or AI engineering services. There are also the APIs and protocols used in technological projects related to AI. Regarding soft skills, I would say communication, which is essential, as well as a certain political intelligence, especially in consulting environments with large clients. You need to know how to protect yourself, anticipate the challenges, and prevent certain decisions from turning against you. Working in a team also remains very important, even if AI is automating more and more tasks. The human side and collaboration will not disappear, notably in client relations and stakeholder management, which is central to all projects. Finally, you must be able to organize, structure your work, and apply effective methods.

In your opinion, in implementation projects, what takes precedence among theoretical knowledge, practical experience accumulated over projects, and intuition?

I would say all three: theoretical knowledge, practical competence, and intuition. Regarding theory, I think it is the foundation of everything. We go from our studies to concrete application in professional life. Having a solid theoretical background in technology is always essential. In my case, I did not initially come from an AI specialization. I started working in this field directly at the beginning of my career because I had this inclination, and my managers and colleagues encouraged me. But a technical and technological base was still necessary to understand the

concepts and evolve in AI. The practical side is just as important because it is on projects that you truly learn. My first project at my first firm was a real school. I acquired skills there that I still retain today and that are highly valued in my current work. At the beginning of my career, I was fortunate to always be assigned to projects and not to the "bench." This allowed me to learn quickly, notably through my junior mistakes. For example, regarding soft skills, I had to progress in communication, writing emails, and interacting professionally. I also learned to manage sometimes stressful situations, which made me evolve a lot. Practice also allows you to develop your skills in real situations. In technology, you can learn many concepts at school, on the internet, or via YouTube, and today even with AI. But it is truly by being confronted with concrete cases, in a project, that learning is strongest. In my work, I am constantly led to do research, use technical resources, and adapt to varied technological environments, today even more so with AI. Then, there is intuition. In this type of project, it is essential. In technology, you often have to show creativity and propose solutions that are not necessarily traditional. You also have to be capable of reacting quickly when a problem occurs in a system or an AI solution. In those moments, you often become the main point of contact, and you must be able to analyze the situation, understand the origin of the problem, and quickly propose a short-term solution, then a more sustainable solution. Intuition is also important because projects evolve constantly. We often move from a MVP to a broader and more complex solution with the client. You must therefore be able to anticipate this evolution, understand the direction the project might take, and adopt the right mindset to guide the technical and functional choices. It is this aspect of anticipation and projection that makes intuition particularly important in my profession.

You mentioned different types of clients, whether governments or private actors. How do you maintain a relationship of trust with your clients?

I think it relies mainly on what I call a "political" posture, in the professional sense of the term. In these environments, we evolve with many interlocutors from varied organizations, sometimes even from governments or large corporations, with very different expectations. Not all clients have the same relationship with technology: some are more business or legally oriented, others are very technical, like developers or engineers, and some have a hybrid profile. You must therefore constantly adapt to these different levels of understanding and expectation. This ability to adapt is essential to establish a relationship of trust. To this is added the structuring of communication: it is important to clearly express your ideas, formalize them, organize meetings when necessary, and involve the right people at the right time to move

subjects forward rather than letting them stagnate. Stakeholder management is central: you need to know how to identify the right stakeholders, include them in exchanges, keep them informed, solicit them at the right time, and not hesitate to use the right communication channels, such as emails with the right people in copy or targeted meetings. Over time, this rigor and transparency contribute to building a good professional reputation, clients and colleagues speak positively about you, which naturally reinforces the relationship of trust. Finally, there is also a fundamental human aspect: staying professional, respectful, and avoiding toxic behaviors.

How do you manage moments of doubt or hesitation on a project?

At the beginning of a project or a career, I think it is essential to ask a lot of questions. This learning phase must be fully accepted. It allows you to understand the context, expectations, and best practices. It is important to go and find information from the right people, those who are truly in a position to bring relevant answers. This allows you to clear up doubts faster and progress effectively. With experience, you naturally ask fewer questions because you acquire reflexes and a better understanding of projects. But at the start, this curiosity and active approach are indispensable.

When you implement AI agents, do you explain how they function to the clients or are they solely interested in the results?

That depends on the client's profile. If I have an architect, an engineer, or a developer in front of me, there is a strong chance that I will go into technical detail and explain the functioning. On the other hand, if the interlocutors are rather business-oriented, they are mainly interested in the results, the KPIs, and the solution's capacity to solve their concrete problem. In this case, the technical explanation is much less central.

Which generative AIs are you authorized to use in your company at [Company's name]?

We have access to several types of language models, notably Claude (including Opus), GPT by OpenAI, as well as Codex, which is often used for AI agents in development. We can also use other models like Mistral, depending on the use cases.

Within the framework of your projects, for what tasks do you use these AIs?

I use AI at several levels. First, in the analysis and design phase: it helps me conduct research, conceptualize systems, structure ideas, produce diagrams, and write functional and technical documentation. I also use it to challenge my ideas and improve the clarity of my notes so that they are usable by all stakeholders. Next, in the development phase: today, a significant part of

the code is generated or assisted by AI. My role consists of defining the requirements, structuring the agents, and giving them a role, a context, and tasks. Once this base is configured, the AI can generate code or execute certain actions autonomously, which I then adjust. Finally, I also use AI for communication and documentation tasks, for example to proofread, correct, or structure texts, notably during internal training or for deliverables.

In your profession as a consultant, on which tasks do generative AI or LLMs seem useless or difficult to replace?

I wouldn't say "useless," but rather that there are aspects of the profession that AI will not be able to replace easily. The main one is everything that falls under stakeholder management. This is a fundamental aspect of the consultant's job. AI can help structure, analyze, or automate certain tasks, but it cannot replace human relationships, direct communication, and the political dimension of the role. In this profession, there is a real necessity to interact with clients, report back in real time, adapt one's speech, and manage sometimes complex relationships between different stakeholders. It is not solely a question of production or algorithms, it is also a question of human presence and trust. Clients also need direct contact. In certain projects, they regularly ask to see the consultants, to exchange ideas in person, and to validate decisions together. This dimension of proximity and human interaction remains essential. Finally, there is everything concerning the control, validation, and orchestration of deliverables. Even if AI can produce many things, it remains indispensable to verify, validate, and pilot what is put into production. This role of supervision and responsibility cannot be fully automated today.

Using AI allows you to be more productive. Does it also give you a better global understanding of your projects?

Yes, clearly. Today, I have a much more global vision of the projects. It is almost as if I were managing the work of a significant set of agents, in my case, that represents about twenty AI agents functioning like a virtual team. Even if the work is highly accelerated, I keep a complete vision of what is happening. I am not confined to a single role, such as only development or only business. I have a 360-degree understanding of the project. I also maintain control over what the AI produces. I do not take its results for granted: I challenge them, I verify them, and I do debugging when necessary. This allows me to thoroughly understand each step and to always know precisely what is happening in the system.

Do you think that less experienced consultants could, thanks to generative AI, reach a level equivalent to yours?

I think it is possible, but it depends a lot on the individuals. I wouldn't say that someone is incapable of it by nature. What makes the difference is above all curiosity, the capacity to learn, and the ability to use AI without apprehension. Today, in companies, we already see a difference between generations and profiles. Some seniors, sometimes with 20 or 30 years of experience, are still reluctant toward AI or certain new practices, while more junior profiles are much more comfortable. In certain cases, the youngest can even surpass more experienced profiles simply because they adopt these tools and these new ways of working more naturally. But in the end, everything depends on personal investment, curiosity, and the willingness to learn and adapt.

At [Company's name], you can use different LLMs. Does this have to be validated with the client beforehand?

That heavily depends on the type of client and the compliance level of the sector. In highly regulated environments, such as banking, the public sector, or certain sensitive fields, there are strong constraints on the tools used. In my current project, it is not a blocking issue. We have great freedom in the choice of LLMs used for development. The models mainly serve to assist creation and are not directly integrated into the final product. On the other hand, in other contexts, notably in banking, the constraints are much stricter. For example, some clients refuse non-European models like GPT or Claude, and privilege more sovereign or secure solutions. There are also cases where internal solutions are developed, with models hosted on private servers, such as internal chatbots using models like [X], in order to guarantee a maximum level of security and compliance. So, everything really depends on the sector, the level of data sensitivity, and the compliance requirements of the client.

When your deliverables are completed, do you specify yourself which parts were created or assisted by AI, or is that not necessary?

It depends on the projects. On certain projects, it was not necessary to specify it. On the other hand, on the current project, I had to do it because the client was informed right from the start that there would be little to no traditional human intervention on development, and that a large part of the work would be carried out with the help of AI. The client was actually quite curious about this way of functioning. In the estimation of tasks, we therefore integrated the fact that certain activities would be accelerated thanks to AI, notably after defining and configuring the agents.

Can the use of AI in your deliverables discredit your expertise?

No, I don't think so. On my end, it doesn't discredit me because I possess a solid foundation in Python and backend development. For example, on my project, there are also frontend and design aspects, which are not my core business. I therefore have no problem recognizing that these parts were generated or assisted by AI. In any case, I am the one who designs and structures the agents in Python, defining their rules and logic. So the role of architecture and design remains entirely human. During demonstrations, the feedback is even rather positive. People are often impressed by the ability to orchestrate a set of AI agents that interact with each other, exchange documents, and function almost like a virtual team in agile mode. Thus, the perceived value is linked more to the design and orchestration of the system than to the fact of coding each element manually.

In your opinion, does generative AI have an impact on the relationship of trust between the client and the consultant?

Yes, it can have one. AI must be used with discernment. The risk is falling into a superficial use, for example simple copying and pasting without a real understanding of the content or without mastering the project. There are already examples where errors related to AI in professional documents had a significant negative impact and generated controversies. This shows that it is essential to always proofread, verify, and understand what is produced, whether it has been generated by AI or not. It is an individual responsibility of the consultant to ensure the quality and consistency of what is delivered.

In this type of situation, when there are hallucinations or errors in a deliverable generated with AI, who is responsible: the consultant, the team, or the consulting firm?

In practice, the responsibility rests first on the consultant, but more broadly, the consulting firm is also impacted. It is generally the name of the firm that is exposed first, notably in the event of a problem or scandal. Internally, the consequences can directly concern individuals. There are cases where major professional errors led to sanctions or even dismissals, even if the contexts are often more complex.

Do clients today demand more transparency regarding the use of generative AI in projects?

Not necessarily. In most cases, clients use AI themselves internally, often for simple uses like writing emails or translation. They are not particularly focused on how AI is used in detail, as long as the results are effective and bring concrete value. There is even sometimes a form of

curiosity or a "magic" effect around these tools, which means they ask relatively few questions about their precise use.

In the long term, do you think clients could use AI to replace consultants and carry out these projects themselves?

In the long term, it is possible that some clients will try to internalize more of these capabilities thanks to AI, notably for simple or standardized cases. However, the role of the consultant is not limited to producing code or deliverables. It also includes understanding the challenges, the architecture of the solutions, stakeholder management, and strategic accompaniment. Even with AI, these dimensions remain difficult to replace entirely, especially in complex environments where human expertise and field experience remain decisive.

You mentioned that the labor market has evolved a lot with AI, notably for juniors. Can you explain your point of view on this evolution?

This is a question I ask myself very often. In three or four years, the labor market has changed enormously. Today, companies and clients are much more reluctant to integrate junior profiles, meaning people who are fresh out of school or who have very little experience. Part of this comes from the fact that many tasks traditionally assigned to juniors are now carried out by AI. For example, tasks like translating documents, extracting information, or doing simple analyses are automated today. A senior no longer necessarily needs to delegate these tasks to a junior since they can do them directly with AI. This phenomenon is not solely linked to AI. There are also economic and structural factors: a large number of graduates, pressure on costs, and an increased requirement for productivity. But AI clearly accentuates this trend by reducing the share of "entry-level" tasks. That said, I think the real question is: how to reposition juniors? Perhaps precisely by entrusting them with more complex tasks or by training them directly in the use of AI. Today, companies wonder what added value a junior can bring faced with a more senior consultant who is already operational with AI. And this perception varies greatly depending on the organization. We also observe a strong difference in maturity depending on the client. In certain sectors, notably the public sector, interlocutors are sometimes less familiar with these tools, and their adoption of AI remains limited or superficial. Conversely, in the private sector, some are already using these technologies in a more advanced way. In the end, I think juniors who know how to master AI quickly will remain indispensable. The subject is not only generational, but above all linked to the capacity for adaptation and the speed of adoption of these tools.

Today, in your opinion, is generative AI a tool or a competitor for the consulting profession?

I think it is both. It is a real question that I ask myself too: how to remain irreplaceable in this context. In my team, I am among the people most advanced in AI, notably on the creation of agents, so people often ask me to understand how I proceed. But at the same time, I am aware that certain roles are going to disappear or be significantly reduced. For example, tasks that were once outsourced to freelancers, like copywriting, are today increasingly automated or integrated directly into AI tools. In my own project, I have already replaced a significant part of human work with AI agents, which clearly shows this transformation. So yes, AI is both an extremely powerful tool and a potential competitor, depending on how it is integrated into organizations and depending on the consultants' skills.

You mentioned the necessity to verify AI productions to avoid errors or hallucinations. Does this mean that the consultant's profession is evolving toward a validation role rather than a production role?

Yes, I think there is clearly an evolution in this direction. The role of the consultant tends more and more toward the supervision, validation, and control of outputs generated by AI. AI can produce a huge amount of content or solutions, but it is not infallible. It is therefore essential to verify, audit, and ensure that the results are consistent, relevant, and adapted to the client's context. In this sense, the consultant becomes more of a guarantor of the quality and relevance of deliverables than a simple producer of content or code. This does not mean that production completely disappears, but value is progressively shifting toward the capacity to orchestrate, challenge, and validate what AI produces.

In your opinion, will the consultant's role evolve toward a role of validation and orchestration of AI productions?

Yes, I think that is already the case and that it will continue to strengthen. The role of the consultant is evolving toward more validation and orchestration of outputs produced by AI, regardless of the type of mission or profile. That said, we must not reduce the profession to simple copying and pasting or to a total dependency on AI. The idea remains that the work produced remains the intellectual and professional property of the consultant. AI helps structure, improve, and accelerate production, but it is always the person who carries the concepts, explains them, and takes responsibility for them. In this context, everyone is going to be confronted with these tools, including in other professions like law or finance. It is therefore

not an isolated case for the consulting profession. On the other hand, this creates a strong need for profiles capable of orchestrating AI, controlling it, avoiding hallucinations, and challenging its answers. This implies a certain amount of training, whether through internal training, e-learning, or skills in prompt engineering. You also have to be capable of detecting the biases and limits of the models.

How do you think the consultant profession will evolve in the long term with AI?

I think the consultant profession will continue to exist, and even in a significant way, because companies increasingly resort to consultants and freelancers rather than internal recruitment, notably for reasons of flexibility and costs. This allows them to adjust resources more easily depending on the projects. On the other hand, the consultants who will stand out will be those who know how to use AI effectively: those who know how to exploit it, challenge its outputs, defend their ideas, and explain their choices clearly. AI allows for a considerable gain in time, which frees up time for higher value-added tasks. For example, in my job as an engineer, the interest is not solely to code. What is truly interesting is designing systems, structuring them, representing them in a global manner, almost like a living system. Coding becomes a part that is more automatable thanks to AI, which allows us to focus on a more global and strategic vision of projects. This also changes the nature of work: previously, developers were very centered on code and business profiles on their domain. Today, AI allows for a much more transversal and 360-degree approach. Finally, this can also improve the quality of life at work, with more varied and less repetitive tasks. Work becomes faster, which can leave more time to reflect, design, and step back from projects.

Interview TECH03

We are going to begin. Hello, can you briefly introduce yourself, indicating where you work, your role in the company, as well as your seniority?

My name is [Consultant's name]. I finished my civil engineering studies in 2023 with a specialization in artificial intelligence. I then worked for two years at a telecommunications company, and I recently joined a consulting firm active in the energy sector. Within the framework of this mission, I work for a client organization created jointly by several energy distributors. This company was set up to facilitate the exchange of contracts and client information, as well as the management and resolution of network related problems. Concretely, they developed software to centralize these operations, and my role consists of intervening in case of problems with this system. For example, when a client has no electricity, it involves analyzing the situation, understanding the origin of the problem, and identifying potential malfunctions in the processes in order to restore the situation. It is therefore mainly IT work, without field intervention, focused on the analysis and resolution of problems at the level of the software and associated processes.

Very well. I am now going to enter the core of the subject. Can you further detail the missions you generally carry out?

My main activity consists of solving problems related to the software I just mentioned. As I just started in January, a large part of my work currently relies on understanding the processes and reading the associated documentation. I notably had to assimilate a significant volume of documentation, about 800 pages during the first weeks. The work is therefore quite technical and requires a lot of understanding and memorization, or at least assimilation. Subsequently, learning is mainly done on the job, as cases are encountered.

Alright. To return to the subject, which focuses on GenAI in consulting, and more precisely on its influence on consultant client relations in terms of transparency, trust, and human expertise, can you specify which GenAI tools you use in your missions or on a daily basis in your work?

It's interesting that you approach this point because there are two main axes I can evoke. The first concerns upskilling in prompting. In my current firm, I recently initiated an approach aiming to help collaborators better integrate GenAI into their professional daily life, notably by developing basic skills in prompting. The objective is to raise awareness across the entire organization regarding best practices and to allow them to identify relevant use cases. This

initiative relies notably on my past experience. The second axis concerns the question of client data confidentiality. This is a point that particularly struck me when I started consulting the internal documentation. Intuitively, I would have been tempted to use tools like ChatGPT to process or summarize the documents, for example to make the information more digestible. However, it was clearly indicated to me that this type of use was not authorized for confidentiality reasons. It is possible to use certain solutions like Copilot within a framework defined by the company, but not other external tools or unapproved environments. For example, tools like NotebookLM are not accessible in this context, and even models like Gemini are not authorized for this type of use with client data. This constraint was already present in my previous company, but it is even more framed here. In my past position, it seemed more logical to the extent that partnerships existed with certain technological suppliers, notably in projects linked to the cloud and collaboration with sector actors. In my current firm, similar agreements also exist, which explains the authorization of certain tools like Copilot, but the prohibition of other solutions.

Alright. So, at your client, if I understand correctly, they solely use Copilot and do not have other GenAI tools developed internally that you could use. Is that correct?

No, that's not exactly it. These are not tools developed internally. In reality, the authorized tools are strictly framed, and only certain environments are available for compliance and confidentiality reasons.

Alright, very well.

You caught me a bit off guard with that question.

So, if I understand correctly, in your consulting firm, you can use ChatGPT, but with certain restrictions?

Yes, that's it. The use of ChatGPT is authorized, but it is forbidden to use it with documents belonging to the client.

Alright. In which cases could you use ChatGPT then?

The client has two types of documentation. On one hand, there are public documents accessible online via the client's website. These documents notably concern the APIs used by the different actors in the energy sector. They allow energy suppliers and distributors to connect to the systems. In this context, different actors intervene: energy distributors, who manage the physical infrastructures, and suppliers, who commercialize the energy. These actors must

interact with the client's systems via these APIs, which requires consulting voluminous technical documents, sometimes of several hundred pages. These public documents can be used with GenAI tools. In my team, for example, a custom GPT was configured based on these documents, which allows us to interact with them in a question and answer format. This makes it possible to interrogate the documentation in a faster and more efficient way, rather than manually browsing through numerous files. On the other hand, internal and confidential documents of the client can under no circumstances be used in external tools. The rules regarding confidentiality are very strict, and the consequences in case of non-compliance can be significant.

Alright. With this custom GPT that you set up as a team, what can you do concretely?

It is mainly a tool for assistance in searching for information. Given that the processes are very technical, it allows us to ask questions when we do not understand the origin of a problem or when we are looking to identify specific information in the documentation. For example, when a problem occurs, it is possible to ask a question directly to the GPT to get indications on the relevant sections of the documentation, rather than having to manually browse through a large number of documents. This represents a significant time saving. However, it is not a tool allowing the processing of confidential information, but rather a support to explore public documentation and facilitate the understanding of the systems.

Alright. Outside of that, what other uses of GenAI do you make?

We also use Copilot, as well as assistance tools for writing emails. For example, I frequently rewrite my emails using these tools to improve their wording. I also use them in certain presentation contexts or when I have to answer complex questions, in order to structure my answers. On the other hand, certain features like automatic recording and transcription of meetings are not yet set up in my current company. This is a tool that I have been able to use in other contexts and which brings significant time savings, notably for the automatic generation of meeting minutes and action lists.

Alright. So if I understand correctly, you use GenAI more for administrative or support tasks, and not for tasks directly linked to the production of your consulting work, like development for example?

Yes, actually, it depends on the case. For example, we also use GenAI in a development context. We notably work with the SPL language to interact with large databases. In certain cases, it is necessary to perform complex queries involving several tables containing very large volumes

of data. The performance of these queries is therefore essential. In this context, I use Copilot as a support to generate or optimize code. On the other hand, I do not use ChatGPT for this type of task due to constraints linked to data confidentiality.

Alright. In your team, are you developers?

No, not specifically. We are not a team of developers strictly speaking.

Alright. In that case, given the use of AI and tools like Copilot, can we say that these tools mimic languages like SQL or Python?

And in this context, does the role of the consultant limit itself to validating AI outputs, or does it remain indispensable? No, clearly, we are in a logic of support. AI does not replace the consultant. It is used as a tool to improve efficiency and speed of execution. Our role does not consist of solely validating outputs generated by AI. We use AI as a support, but reflection, analysis, and validation remain human. At this stage, AI does not replace the consultant.

In that case, do you communicate to the client when you use these tools, or is that not necessary since AI intervenes solely in support?

In practice, we use AI indirectly to assist us in certain tasks. We do not systematically present outputs generated directly by AI to the client. AI is used upstream to facilitate the work, but the final deliverables are produced and validated by ourselves. If we use AI to generate code or suggestions, we then provide the necessary information to achieve an adapted result, but always with a human intervention in the validation and adaptation. In this context, it is important to specify that the tool does not have access to internal databases or their structure, notably to columns or schemas. Consequently, I do not necessarily know in advance what information I need to look for. There is therefore no direct integration with these systems. It is also for this reason that, for now, these tools cannot replace consultants. An integration of this type with such complex platforms would still require a lot of time. Setting up communication interfaces, like APIs, remains technically complex, particularly within the framework of GenAI, and current solutions are still relatively limited from this point of view.

Alright. And in this case, when you use these tools for prompting, do you talk about it to the client or not?

No, we do not talk about it. We know that Copilot can be used, but other tools are not authorized within the framework of client data. If this were to be discovered, the consequences could be significant. However, in practice, the client does not have direct visibility on the use of these

tools. There is no monitoring mechanism allowing them to know whether or not a consultant uses an AI on their workstation. There is therefore no specific supervision, except in case of voluntary disclosure. The situation could eventually be escalated to a manager, but in the absence of communication, there is no direct control. The consequences therefore essentially fall under individual responsibility.

Alright. The firm remains quite clear nonetheless on authorized tools and those that are not.

Yes, exactly. The rules are very explicit regarding authorized tools and forbidden uses. The company is particularly strict on these aspects.

Another question I was asking myself concerns cases where GenAI tools produce erroneous results or hallucinations, for example in reports. In this type of situation, who is responsible?

I have not been directly confronted with this type of case. In my team, which is rather technically oriented and not audit-oriented, this kind of situation does not present itself frequently. In theory, responsibility depends on the hierarchy. In a consulting context involving several participants, responsibility could go up to a high managerial level within the consulting firm, or even to the executive level, depending on the severity of the problem. It is therefore a shared responsibility, but one that tends to be assumed at a higher level in case of an incident.

Alright. I have a somewhat particular question. Even if the company is open to the use of GenAI, could the fact of explicitly mentioning its use, for example to generate a SQL query, devalue you in the eyes of the client?

Yes, I think that could be perceived as unprofessional. Clients are very attentive to these aspects. They notably organize regular training sessions aimed at everyone working with them, internal as well as external, to raise awareness about data confidentiality, GDPR, and best practices. The rules are therefore clearly established and strictly applied. Explicitly mentioning that one used an AI in a context where it is not authorized could be perceived negatively, or even as a lack of professionalism. It is not necessarily considered a betrayal, but it would go against the established guidelines and the expected trust.

Alright. One last question: have you already observed, without going into details, cases where consultants would use GenAI to reach a higher level of expertise, for example in development, when they do not initially possess the required technical skills?

To a certain extent, yes and no. In my company, knowing how to code is not an absolute prerequisite for entry. Some consultants arrive without technical skills in programming and receive initial training. Then, they use tools, including GenAI, to help them progress, notably in query optimization or the understanding of certain processes. In this framework, GenAI can contribute to a progressive upskilling, helping to better formulate queries or improve technical understanding. However, this remains framed by the company. Furthermore, if a mission exceeds the consultant's initial skills, it is generally accompanied by the consulting firm, which can offer appropriate guidance or upskilling. In case of need, a consultant can also escalate the situation to their company to be supported in acquiring the necessary skills.

In my opinion, entrusting a single person, without prior expertise, with a 100% usage of ChatGPT to perform a task for which they possess no skills does not seem relevant to me. To my mind, the use of these tools should fit more into a logic of partnership. Within this framework, this could also have organizational implications, notably in terms of remuneration or the recognition of responsibilities. For example, if a person is led to produce work involving more responsibilities thanks to the use of AI, this could justify an evolution in their position or remuneration. However, in practice, this type of situation does not really present itself. From my experience, it is very rare to be assigned to a mission about which you have absolutely no knowledge. On the other hand, it is always possible to use these tools autonomously within a framework of personal experimentation, a bit like in a laboratory environment. This remains, however, more relevant for personal projects than for paid professional missions.

I would like to bounce back on a point you mentioned. Certain profiles, who did not initially have skills in SQL, used ChatGPT or other GenAI tools to rephrase or improve their queries. Do you consider this practice as positive, notably in a logic of broadening skills?

Yes, I consider it to be globally positive. There are actually two possible approaches. The first is a short-term approach, where you seek to quickly obtain a solution by relying directly on AI. The second is a longer-term approach, oriented toward learning and upskilling. To develop sustainable skills, it is not enough to entirely delegate one's tasks to ChatGPT. In that case, real learning is limited. On the other hand, one can use AI as a pedagogical tool. For example, it is possible to ask it to explain a SQL query step by step in order to understand its functioning, its stages, and its logic. In this configuration, the tool becomes more of a trainer than a simple provider of answers. It is therefore a matter of two different uses: either a use oriented toward

an immediate solution, or a use oriented toward learning, where the AI accompanies understanding rather than replacing cognitive effort.

Alright. As previously mentioned, in your opinion, what parts of your current activity in consulting cannot be replaced in the short term by GenAI?

In the short term, I think everything concerning the integration of solutions will remain difficult to replace by AI. In the context of the platforms used for managing clients, notably in telecommunications, I do not currently observe any significant integration of GenAI into existing systems. The company in which I evolve is relatively young, with about ten years of existence, and it is currently developing a second version of its software. The objective is to stabilize this new version by 2030. In the current roadmaps, the integration of AI or GenAI remains very limited, if not non-existent. It is probable that these features will be integrated into a later version, but that will take time. If the development of the first version required 15 to 20 years, it is reasonable to think that version life cycles remain long. Thus, even if the integration of GenAI is called to develop in the future, the industry is not yet fully ready, notably in sectors like telecommunications or energy.

Do you think that, in this context, clients are going to become more demanding toward consulting firms, notably by asking for more transparency on the methods used to produce results?

That depends on the type of consulting. In technical consulting, the main stakes remain the quality of the final result. Clients expect above all concrete deliverables and operational solutions. As long as the results are achieved and justifiable before the client, the use of tools like GenAI in the background does not fundamentally change perception. On the other hand, in fields like audit or the production of reports, the situation is different. Clients, who invest significant amounts of money, are particularly attentive to the quality, reliability, and rigor of the deliverables. In this case, the use of GenAI could spark more questions, notably in terms of credibility and traceability. One can therefore think that the market self-regulates: firms capable of producing high-quality deliverables, with or without AI, will be privileged by clients.

In technical consulting, do you think clients could question the skills of the consultants due to the use of GenAI?

In the short term, I don't think so. In my current experience, clients mainly evaluate the results provided rather than the means used to achieve them. As long as the deliverables are relevant and meet the expressed needs, the tool used upstream matters little. In the longer term, it is

possible that this perception will evolve, but for now, it does not seem to constitute a major stake.

Another question: given the existence of GenAI, could clients be tempted to carry out certain tasks themselves that were previously entrusted to consulting firms?

It is a possibility. In certain cases, we can observe a reduction in staff rather than a complete replacement. AI allows for the automation of certain tasks, which reduces the need for human resources without necessarily eliminating functions entirely. For complex technical tasks, it remains difficult to replace an expert with AI alone. On the other hand, for more standardized tasks, like report production or documentary synthesis, GenAI tools are already very powerful. This could therefore lead to a progressive transformation of organizations, with fewer consultants for an equivalent volume of work. Some companies could also reallocate their resources to other activities. In my experience, notably in environments like telecommunications, we already observe staff reductions linked to these evolutions. Executives may sometimes try to minimize this impact in their communication, but in fact, structural adjustments take place, including unfilled departures or team reductions. Thus, rather than a total replacement of consultants, we are witnessing more of a progressive transformation of skill requirements and an optimization of staff. No, I was simply saying that you are not the first person to encounter this type of difficulty, and you won't be the last. We therefore perfectly understand the challenges you are facing.

Interview TECH04

As my mother explained, my thesis focuses on consultants and the consulting field, and more specifically on the impact of generative AI on client relations, notably in terms of trust, human expertise, and transparency. I have a series of questions to ask you: you just need to answer them. To start, can you introduce yourself? Where do you work? What is your role? And how long have you been with the company? I work at a professional services firm.

I am a generative AI consultant and functional analyst. That is my role.

What type of mission or project do you carry out there? What do your missions consist of?

I work as a consultant. We intervene at clients' sites to sell and implement Microsoft Dynamics solutions. We also set up the AI associated with these solutions, notably via Copilot and Copilot Studio. This includes managing artificial intelligence tools, as well as implementing agents and chatbots directly for the client.

Alright. So it is mainly to make chatbots, or are there other possible uses for the client?

In fact, Copilot Studio and Copilot in general serve mainly to implement agents. Today, we talk about agentic AI, meaning agents capable of automating tasks on behalf of the client: organizing meetings, executing certain actions, or even interacting as a chatbot by relying on knowledge sources. But today, the architecture of agents is much more advanced and especially oriented toward automating the client's daily routine. The objective is not only to create a chatbot that returns information, but to integrate useful automations into everyday processes. That said, chatbots remain very important, notably because a huge number of documents and information sources exist, and employees do not always have the time to consult everything. Thanks to generative AI, we can quickly provide answers based on these knowledge sources.

So, for example, this agentic AI could automate data filtering in Excel?

Exactly. That is also the objective of Copilot. We work specifically with Microsoft tools, so Copilot is also integrated into Excel. It allows you to target certain elements of the file and create customized agents capable of processing and analyzing Excel data.

How does a mission generally unfold? Does a client come to see you saying: "I need an AI, what do you recommend?" How does it happen? Generally, the client is curious about the technology but does not know exactly what to use, nor how, nor why. The first step is for them

to give us an overview of their data: what they possess as a database, which sources are available, and especially what they would like to accomplish. Our role, as consultants, is to explain to them what they can do with their data and to propose the best way to set up a solution that truly facilitates their work. We analyze which data is available: is it an internal source? Is it accessible? Is it only documents? Or databases that require specific connectors? Our mission is to guide the client throughout this process, taking into account their objectives, the most adapted solution, but also financial constraints, because each solution involves a cost.

And if it's not a secret, what type of clients do you have? Is it a particular sector? Is it mostly SMEs? We work with all types of clients, truly across all sectors.

In your opinion, in your job as an AI consultant, what are the indispensable soft skills and hard skills?

For soft skills, you must above all be comfortable in client relations. You need to present well, communicate well, and especially know how to explain things simply because the technology behind AI is complex. You must be capable of translating this complexity into understandable terms. Sometimes the audience is very technical, sometimes much less, so you have to know how to adapt. You also need to know how to read the room and understand who you are speaking to. For hard skills, you must know generative AI models and how they function: why a model answers in a certain way, why one model is more adapted than another, how to connect databases, how to configure tools, etc.

Alright. This is perhaps a bit more technical, but since I am not very familiar with AI consulting: do you code to implement AI agents, or not necessarily?

There are two ways to implement agents. Microsoft offers Copilot Studio, for example, which is a low-code tool allowing you to easily create an agent, notably by integrating it into Microsoft tools like Teams, Outlook, etc. But there is also a more technical approach where we code the agents, for example via Python. Certain connections or features are not possible with Copilot Studio but can be achieved with Python. You must therefore have a global vision of the solution's architecture.

Precisely, we were talking about that a bit before. When a client is curious and asks you to implement an AI, what do you base yourself on to choose the solution? Do you base yourself on your past experience, your intuition, your skills?

You must first do thorough research on what can be implemented. That is why we build a solution architecture. Once this architecture is defined, we identify what is possible or not, what must be done in a certain place, and what must be treated differently. Then, we evaluate the skills available in the team: who is the most competent on a certain topic, who masters a certain tool better, etc. Then we choose the right people to implement each part of the solution for the client.

In your opinion, what are the essential elements that allow for building a relationship of trust between the consultant and the client?

You must be transparent right from the start. If the client wants something that is not possible, whether because of budget, deadlines, or technological limitations, you must say so immediately. You have to explain that it is out of scope, that it is not feasible at the moment, or that it will need to be developed later. You must not say yes to everything just to win the client. Transparency is at the core of the relationship of trust. If a company promises something in two days when in reality it takes two months, the client will eventually find out. And in that case, the relationship of trust collapses. Transparency must therefore remain central: the client must understand that some things are possible, others are not, and that time and budget constraints are an integral part of the discussion in order to propose the most suitable solution.

Which generative AI tools are you authorized to use?

We work mainly with Microsoft, so we use Copilot Studio and the models available through Copilot Studio. We also use Azure. Globally, we are very centered on the Microsoft ecosystem.

And is it for confidentiality reasons?

We have trust in the security offered by Microsoft. We are a Microsoft partner, and we know the security mechanisms put in place. We know that the client's data is secure, which allows us to implement solutions with peace of mind.

And personally, in which cases do you use generative AI?

Mainly for building agents. Of course, today everyone uses tools like Copilot in daily life to get general information. It has become very common. But in our case, the objective is above all implementation: making sure that the client has a tool adapted to their own environment and specific needs.

In your opinion, what are the tasks that generative AI will never be able to perform, or at least those for which it cannot help you in your job as a consultant?

I think anything that touches the human side will remain very difficult to replace. The client must choose which consultant or which consulting company they want to work with. And it is not an artificial intelligence that is going to convince a client over another. It is truly human contact that puts the client at ease and creates this relationship. AI is very useful, but it has limits when it comes to human skills.

Alright. So, on one hand, it still allows you to be more productive.

Of course. If you use artificial intelligence daily, for example to summarize documents, it saves you a huge amount of time. If you have a 15-page document, instead of spending time reading everything, you can quickly get a precise answer to a question. This saves time for the team and allows them to focus on other tasks. But the goal is not to save time just to waste it pointlessly afterward: it is to save time to produce more value for the client.

Can generative AI also help you better understand a project, or learn new things?

Yes, of course. For example, when you do a search on an engine like Bing, you can get about fifteen or twenty pages on your topic. With generative AI, you ask the question directly and you get an answer faster, without having to browse all the sites and look for the relevant part. On the other hand, you must always check the sources. The mistake would be to take the answer as it is without checking it. But yes, it allows you to structure your ideas and better understand a subject. Then, you can go deeper by investigating further. You must not simply copy the information. Let's take a simple example: if you ask how to make a lemon cake, you can get a recipe and send it directly to someone without having tested it. But you cannot be sure it is a good recipe. You must always check and test before transmitting something.

Yes, I see. I am going to take an extreme case: imagine a junior who uses AI to do 80% of their project. It is often said that we learn by doing things ourselves. Don't they risk learning in a superficial way and gaining nothing from their experience?

That is a very good question. You have to understand the true purpose of artificial intelligence. If you use it to do everything for you, summarize all your emails, write your emails, or even write your thesis, the result might be well written, but you won't truly learn. AI should be used for redundant tasks, not to replace reflection. If you use it too much to think in your place, you risk becoming dependent. And that is precisely what we want to avoid: a generation unable to write an email without Copilot. You must keep your capabilities, otherwise other people will progress faster than you in their way of thinking, and you risk stagnating.

When you or your team use generative AI, do you talk about it to the client beforehand, at the end of the project, or not at all?

We use AI either in the implementation of solutions for the client, or directly in the solutions delivered to the client.

I mean: if during a project, for the implementation, you don't know how to do something and you use AI to help you, do you need to have the client's consent? Is it necessary to talk about it or is it not required?

Yes, it is important. Some clients clearly say right from the start that they do not want us to use AI in their projects. Others accept, but on the condition that it is not abusive. And some clients use AI themselves. The most important thing, as I said, is that the client receives a solid and functional solution. Now, whether you used AI or not, the essential thing is to test what you deliver. If you use AI to generate code or text, it must be clean, clear, reliable, and above all tested. We cannot put a piece of code into production without having checked it. But some clients explicitly ask: "I do not want any AI at all in the code, I want it to be written manually." So, it really depends on the client.

Does it happen that some clients then ask which parts were done by AI, or is it mostly the final result that counts?

No, in general, it is the result that takes precedence. Of course, as I said, the client must agree beforehand and everything produced must be tested. But if the solution is functional, the client is generally not a developer, or not necessarily a developer. There are therefore many technical aspects that they do not know and cannot necessarily know. They must therefore trust the developer and the whole ecosystem working on the project.

Do you think that being too transparent about having used generative AI, even in a minimal way, could make the client respect your work less?

As a client, I imagine they mostly want a well done, functional result that respects the budget and deadlines. AI helps facilitate certain tasks, notably those that are redundant or take a lot of time. So if in the end the client receives a tested, clean, and reliable solution, and the project respects the defined time limits, I don't think that reduces trust. Trust is built right from the start. The final objective is to have something that works. If you promise a solution in a week, then in a week the client must have that solution. That is what establishes trust, not the fact of having used AI or not.

There have been several cases where consulting firms were caught red handed giving recommendations full of hallucinations because the sources were not good. In those cases, in your opinion, who is responsible: the consulting firm, the individual consultant, or their team?

From my point of view, it is the individual consultant. You have to realize the importance of what you deliver, and it is essential to proofread, check, and validate. As I said, using AI is a good thing, but using it in a responsible, clear, and documented way is even better. The objective is to produce tested and approved work. Making life easier is good, but the most important thing remains that the client is satisfied and receives a truly functional solution.

I am going to move on to my last questions, rather on the future of the consultant profession. Do you think that, in the long term, clients will not try to implement AI agents themselves thanks to AI? Will consultants still be useful, knowing that missions are expensive and companies might try to do everything internally?

Developing an AI agent requires resources and a certain know-how. Of course, a company can have someone internally capable of creating agents, but that implies allocating time and resources to this person, or even hiring specialized profiles. That is why calling upon external consultants can be more interesting: if a company wants to develop an AI solution on a punctual basis, it is more profitable to bring in consultants to carry out the project and then end the mission, rather than paying salaries over the long term. You also have to find a balance: if an internal person works on an AI project, they will have to divide their time between this project and other priorities. And the role of the consultant is not limited solely to AI. It also includes the global integration of the tool: how the company will adopt it, how it will evolve with it, and how the tool integrates into existing processes. The consultant also brings a global vision, an expertise, and an experience acquired over several projects. Internal teams can develop this, but they will have to learn on the job, which requires time and energy. During this time, they are no longer focusing on the core business of the company. That is why consultants still have an important value today, which is human, technical, and strategic all at once.

Do you think that in the future, since you said you always have to check and double check, consultants will end up becoming mostly people who check the work of the AI, rather than truly producing the solution themselves?

So you think that, in the future, people internal to the company will directly produce the solution, and that consultants will become mostly verifiers. That is a possibility. But as in many

sectors, we cannot truly predict: it will be above all the human aspect that will determine this evolution. If a company is comfortable with this model, it will go in that direction. Conversely, some companies are more rigorous and methodical, and prefer a solution to be entirely handled by an external entity because ownership of a solution is important. So, in the end, it is really the human factor that will determine which is the best approach. We can draw a parallel with certain professions that are partially disappearing, like that of translator. If this profession is less demanded today, it is because people decided that the translation produced by AI was sufficient. But we must also recognize that translators remain important: they bring a quality, a style, and above all a clear responsibility. If a translation is bad, we are not going to blame an artificial intelligence because it is only a tool. On the other hand, when a human is responsible, we know to whom to attribute the error. It is the same for consulting: if an agent is bad, we must be able to identify who is responsible for it. And if it is very good, we also know who to thank.

In your opinion, how will your profession evolve in the next 5 to 10 years? Even if we cannot know the future, I found it important to ask the question.

I think that needs will change. And I also think that the human side will become increasingly rare. This means that my technical expertise could evolve toward a more human expertise, more oriented toward the management of all these tools. Rather than focusing solely on implementation, it will also be necessary to manage everything around it: organization, monitoring, adoption, and governance. The implementation itself will perhaps require fewer people, but global management will remain essential.

And in your opinion, in the next 5 to 10 years, what added value should a consultant have to remain relevant in the face of generative AI?

It will be necessary above all to remain attentive to the client. Today, prompting well or prompting poorly makes a big difference in the quality of an answer. But if a client explains something in a vague or not totally clear way, a human consultant will have this capacity to understand the real needs and clarify expectations. It is truly listening and understanding that will make the difference. I am sure that you too, Imane, have already asked a question to ChatGPT or another tool and the answer was not satisfactory, then you had to prompt again several times. That happens to everyone. With a human person in front of you, the exchange is more direct: they can guide you toward the right solution without you needing to rephrase several times. No need to "prompt well" to be understood.

Alright, that was my last question. Thank you.

Perfect, thank you.

Interview TECH05

I'm starting the recording. First, could you introduce yourself? Where are you currently working, what is your job position, and what kind of projects are you involved in?

I am currently one of the two AI leads within [Company's Name]. Concretely, this means I have an internal role focused on developing the AI practice: recruiting the right profiles, supporting sales activities, managing people, making sure consultants stay up to date, creating learning paths, and ensuring they can access relevant training. In parallel, I also work on client projects, mainly in project management, data strategy, AI strategy, and team leadership. At the moment, my role is more hands-off, meaning I do less coding than I used to. My background is in statistics and data science, but over the last three years I have been more involved in project management and team leadership.

What does your day-to-day work look like? What are the main stages of your projects?

Currently, I am working on two projects, so I can use them as concrete examples. The first project is with an energy supplier called DOS24. My involvement there focuses on defining the planning, structuring the work packages, ensuring the right resources are allocated, setting up the governance, organizing team meetings and sprint planning, and guiding the team throughout the project. I also monitor the budget, provide daily support, handle escalations when needed, and manage stakeholders. Overall, it is a very planning-oriented role, with a lot of day-to-day support and team leadership.

The second project is at SNCB/NMBS, the Belgian railway company. There, I also act as a team lead, but with a stronger focus on organizing the data team. This includes governance questions, defining key priorities, managing the roadmap, and translating business needs into data projects. For example, when there is a general business question, my role is to determine how it should be addressed: whether it requires a dashboard, an analysis, or the development of a data science model. So this project is more focused on governance, team leadership, and analytical translation.

In your opinion, what are the main soft skills and hard skills required to do your job?

In terms of soft skills, communication is clearly the most important one. I spend around half of my time in meetings, so I need to understand what is happening, communicate the right information, and ensure that actions are taken efficiently. The second key skill is stakeholder

management, which is closely linked to communication. These two are by far the most important skills in my role.

You mentioned that one of your clients is SNCB. What type of clients does [Company's Name] work with? Do you target specific sectors such as SMEs or startups?

It is quite diverse. Historically, we have had a lot of experience with marketing departments. In terms of industries, we have mainly worked in telecommunications, energy and utilities, and retail. These are the sectors where we have the strongest expertise. However, since we were acquired by [Company's Name], we have expanded into many other industries. We are organized with account managers dedicated to different sectors, such as finance, retail services, defense and space, and others. So our market coverage is quite broad. For example, we currently have a consultant working in defense. We also work more and more with the European Commission and public services. Overall, we have a broad range of interests, but of course we do not cover every possible sector.

In your opinion, what are the key elements needed to build a strong trust relationship between consultants and clients?

For me, clarity and honesty are essential. As a consultant, you are expected to provide expert advice, and sometimes that advice means saying no or advising against a certain direction. Clear communication, transparency, and honesty are critical. Trust is built when you are able to openly communicate when things go wrong or do not go as planned. If you are honest and transparent, clients will naturally trust you more.

You mentioned transparency. For example, if there are doubts or diverging opinions during a project, do you communicate that to the client?

Yes. In most projects, there are different possible approaches. The key is to clearly explain the pros and cons. Just because something is technically possible does not mean it should be done. Sometimes it is better to explain that a certain approach is not worth the time or resources. As a consultant, your role is to provide expert guidance, but in the end the client remains responsible for the final decision. If the discussion is based on transparency, honesty, and clear communication, then a strong trust relationship is created.

When you work on a project, how do you make decisions? Do you rely on expertise, knowledge, intuition, or something else?

It is a combination of all of those. As a consultant, you are expected to be an expert, meaning you have already encountered similar situations many times. Intuition is also useful, but what really differentiates a consultant is experience and lessons learned from previous projects. Expertise and experience are the main reasons why clients hire consultants instead of assigning someone internally. Of course, for new situations, logical thinking and intuition also play a role. But ultimately, we are hired because we have done it before and because we are expected to know what we are talking about.

My next questions will focus on generative AI, which is the main topic of my master's thesis. What types of generative AI tools can you use inside [Company's Name] ?

It depends heavily on the client, because when you work at a customer site you must follow their internal policies and rules. For example, some clients, especially in banking, strictly forbid the use of any external generative AI tools. They may develop their own internal chatbots, but employees are not allowed to use tools like ChatGPT on their professional laptops. For other projects, we do have subscriptions to Microsoft Copilot. We use Copilot as a conversational interface, similar to ChatGPT or Gemini. We also use GitHub Copilot, which is an extension integrated into coding environments such as VS Code. It allows developers to receive assistance directly while writing code. In addition, many people use Claude from Anthropic, although I do not know the exact model versions. There are also initiatives at the group level to develop internal AI products. For example, there are two assistants called Maya and Heidi. One is specifically designed for project management. The other is more like a wrapper around an LLM, tailored for a specific type of user. Finally, there is an internal chatbot called [X], designed to help with writing RFPs. So overall, we have several internal products to support daily work. However, the first reflex for any consultant should always be to check what is allowed by the client. For example, at NMBS there is no strict AI policy, so we can use most tools. At DOS24, we cannot use external tools. And at BNP, we cannot use generative AI tools at all. So it strongly depends on the client.

So before the project even starts, it is already specified whether AI tools can be used or not?

Yes.

For which tasks do you personally use generative AI in general?

For many tasks, honestly. One advantage of project management is that a lot of work can be done outside the customer's environment, which makes it easier to use AI tools. I try to use AI

as much as possible. For example, I use it to create project planning, to discuss risk management, governance, and strategy. Personally, I have a subscription to Google Gemini, and I use it frequently in my day-to-day work.

When you use generative AI, do you think it helps you gain a deeper understanding of the topic, or is it mainly about saving time and being more efficient?

Both. I would also add a third benefit: filling in gaps. Most of the time, I prepare something first, for example a project planning, and then I ask the AI if I missed anything. It becomes a kind of discussion, and sometimes it identifies something I forgot. So I start with my own preparation and then improve it with AI feedback. Sometimes, I use it differently: I ask the AI to give me a completely different perspective on a topic, even without any preparation, just to challenge my way of thinking. Recently, I have also started automating parts of my work by creating specific personas, called “Gems” in Gemini, to handle recurring tasks.

Do you learn from your experience and from experimenting with different prompts? Do you gain new knowledge that you didn’t have before?

Yes, I think so. Again, it helps me fill in gaps and identify things I missed. Also, in my free time, I created a persona focused on learning. It provides short learning sessions of about 15 minutes. I try to do that around three times a week, outside office hours. For example, it can teach me frameworks related to analytical translation or change management. So generative AI helps me in my day-to-day work by improving my thinking and filling gaps, but it also supports additional learning through specific personas.

It is often said that in consulting you learn by doing, through experience and repetition. When you use generative AI, do you think you develop a more superficial understanding compared to the traditional “learning by doing” approach?

Yes and no. If you do not know anything about a topic, then generative AI will usually give you a superficial, high-level answer. However, if you already have some knowledge, you can go deeper by asking better questions and prompting more effectively. In that case, you can reach more detailed insights. But overall, I would say that around 80% of what you get remains superficial, and only about 20% can be considered a real deep dive.

Do you think less skilled consultants could, thanks to generative AI, do the work that more experienced consultants would normally do?

Yes, definitely. If someone learns how to use generative AI properly, they can significantly improve their performance. In a way, it can turn a junior into someone who performs more like a senior.

In your opinion, what is generative AI not able to do yet? Which tasks do you think it will not replace you in?

I think creativity is one key area. If everyone uses AI, they will often end up with very similar outcomes because the models generate comparable answers to similar prompts. In Dutch, we have a phrase: if you ask ten doctors, you get eleven opinions. With generative AI, it feels like the opposite: many people will end up with the same answers. So creativity will become an important way to differentiate yourself as a human. Honestly, I am not sure what else AI will not be able to do. Creativity is probably the most important point. I think a large part of my work could potentially be done by AI. It is not a very reassuring future for me, but we will see how things evolve.

That's also why I am writing this master's thesis. You mentioned earlier that whether you can use AI depends on the client's policy. Why would your company allow its employees to use generative AI? Is it mainly to save time, to stay technologically competitive, or something else?

I would say there are three main reasons: saving time, increasing quality, and gaining a competitive advantage. For example, if we sell a project that takes 40 days, but another company can deliver it in 30 days thanks to generative AI, then we are no longer competitive. So we need to adapt our ways of working. Overall, the main reasons are efficiency, quality, and competitiveness.

When you are allowed to use generative AI in a project, do you tell your customers which parts were done using AI, or are they only interested in the final results?

I have never explicitly mentioned which parts were done using AI. I know some people do it in final presentations, for example by listing tools like ChatGPT or Gemini in the acknowledgements section. But in terms of specifying which tasks were supported by AI, no, I have never done that. I have also never been asked. I think most people already assume that everyone is using AI.

So they care more about the results than the process. When you propose a solution, do you explain in detail how you obtained it and how it works, or do they not necessarily care?

It depends on the type of project. For example, in a dashboarding project, clients are mostly interested in how the dashboard works and how to use it, not in the backend configuration. However, in traditional data science projects, we usually involve the client at every step. We make sure they understand what is being used in the model. Transparency is essential in those cases, because the client needs to understand the value, but also the limitations and risks. Often, the client also needs to justify the project to their manager or a steering committee. This level of transparency creates long-term partnerships, but it is not always the case for every type of project. When it comes to traditional AI projects, we definitely focus on transparency. With generative AI projects, it is slightly less systematic. In the end, every project is different.

You mentioned that you have never been asked whether you used AI. But do you think in the future clients will demand more transparency regarding whether insights were generated with generative AI or not?

I am not sure. I think the answer will be somewhere in between. At a high level, clients should probably know what tools were used. But I am not convinced they need to know every single step in detail. However, for anything that will be used externally, transparency becomes more important. For example, if you create a marketing campaign and generate an image using AI, then the client should definitely be informed. There are important reasons for that, such as intellectual property concerns and reputational risk if something is wrong with the content. So in those cases, I believe the client should know where AI was used and where it was not. But I do not think this level of transparency is necessary for every task.

Let's simulate a case. There have been situations where consulting firms delivered reports containing hallucinations because the sources or numbers were wrong. In that case, who do you think is responsible for this failure: the consultant, the manager, or the consulting firm?

All three. Definitely not the generative AI itself. The responsibility lies with the people using the tool, because they are the ones who must verify the information. You cannot blame ChatGPT if something goes wrong. So in my opinion, responsibility is shared between the consultant, the manager, and the company. That being said, it is not clearly defined anywhere, and I am not sure how it works legally. I don't know whether a company like [Company's

Name] could be held legally accountable. At the very least, the consultants involved should be responsible for the final output. But it also depends on the severity of the hallucination: if the mistake is major, then the whole company could be held accountable, whereas if it is minor, it might not be treated the same way. So I don't have a fully conclusive answer

In your opinion, does the use of generative AI impact the trust relationship between the consultant and the client? Does it strengthen it, undermine it, or have no impact at all?

At the moment, I don't think it has a significant impact. I believe everyone assumes that most professionals are already using AI tools, but still mostly in an experimental way. However, if we develop internal assets or products and implement them for clients, and the result is poor, then of course the relationship could be damaged. But for now, I would say it has no real impact.

Now, my final questions concern the future of the consulting industry. Do you think clients will be able to do the work themselves with generative AI that consultants currently perform?

Partly, yes. Things will definitely change. But I would compare it to what happened with IKEA. When IKEA appeared, people thought professional carpenters would disappear, because you could buy a closet for a fraction of the price. Yet carpenters still exist today. The same will happen with consulting. Some companies will prefer standardized, one-size-fits-all solutions, and they will be satisfied with that. But consulting is based on expertise and tailored solutions. So consulting will evolve, but it will not disappear. In fact, we already see a shift. Before, clients asked us mainly to build AI solutions. Now, they increasingly ask for advice: which AI product should they buy, how to manage change, how to ensure adoption, how to use AI correctly and compliantly, what the do's and don'ts are. So the type of work is changing, but consultants will still have a role.

In my literature review, many authors argue that consultants may shift from designing solutions to validating the outputs generated by AI, and then delivering that validated result. Do you agree with this idea?

Yes, partly. For easy and repetitive tasks, generative AI will likely do more of the work, and consultants will validate the output. But for complex tasks, and especially tasks that involve human aspects like change management, I don't think generative AI will replace consultants. You can already see this with self-driving cars. There are ethical discussions, and many people still prefer human drivers because they want to feel in control, even if self-driving cars could statistically reduce accidents. I think the same logic will apply to generative AI: people will use

it, but they will still want to speak with humans. They will still want an emotional connection and human interaction, rather than relying entirely on a cold chatbot. So yes, some tasks will become automated and validated, but not everything.

How do you see the consulting profession evolving over the next five to ten years?

That is a very good and very broad question, and I think the consulting profession will evolve on several levels. First, we have had internal discussions about the traditional consulting “people pyramid.” Usually, consulting firms have many juniors, fewer mid-level profiles, and even fewer seniors. It follows a classic pyramid structure. But now we are increasingly talking about an “obelisk” instead of a pyramid, because many junior-level tasks can be automated. We already see this on the market: it has become harder to sell junior consultants. And it makes sense, because when clients hire consultants, they pay for expertise, not for someone who has just graduated and needs to be supervised by a senior. For example, today you might have two data engineers and a junior assisting a senior. In the future, this junior role could increasingly be replaced by a generative AI product supporting the senior consultant. I am not saying that juniors will completely disappear, because that would not be sustainable. But the market is clearly changing. When I graduated from university, I had many companies to choose from. If I were graduating today, I honestly think it would be much more difficult to enter the job market. So I believe the “golden age” of consulting, where firms could easily recruit juniors and staff them on projects, will likely be over. Second, there will be major changes in technology. Traditionally, the consulting model was mainly about selling people: selling expertise through staffing. In the future, it will likely become a combination of consultants and AI agents, as I mentioned earlier. These agents and assets will probably be developed internally to ensure that we maintain quality, efficiency, and competitiveness. This also means that the model will shift. It will no longer only be about sending consultants to the client to produce slides and deliver work on-site. There will likely be more internal roles as well, with people working inside the consulting firm to create, improve, and maintain these AI assets, which will then be used in client engagements. Third, sales and business models will also change. If we invest money in building internal assets and automation, we might be able to deliver a 40-day project in 30 days. But we cannot simply reduce delivery time while charging the same daily rates without rethinking the model, because we would be spending money to earn less money, which would not be sustainable. So there are discussions about moving toward business models that include more intellectual property, where we are not only selling people but also selling products and reusable assets. There is also some discussion about outcome-based models,

although I don't think this will happen in the near future. Outcome-based means, for example, that if a solution generates two million euros of profit, [Company's Name] would receive a percentage of that value, such as 10%. So payment would be based on results rather than on days worked.

Finally, day-to-day consulting work will change significantly. Many tasks will become automated and rechecked through AI. Consultants will likely spend less time on repetitive work and more time focusing on complex topics that require expertise and human judgment.

What is the essential skill consultants should have in the near future to remain indispensable in the face of generative AI?

I would go back to what I mentioned earlier: communication skills will remain the most essential soft skill. The second key skill will be the ability to work with AI. If you cannot work with AI, you will be left behind. We have even stated internally that if someone does not learn to work with AI, they may no longer be relevant within three years. So the two essential skills are strong communication, which is a human skill, and the ability to adapt and collaborate effectively with AI tools.

That was my last question.

Thank you.

Interview FUNC01

As I mentioned on LinkedIn, my thesis focuses on the use of generative AI in the consulting world and on how this could influence relations between consultants and clients. I am going to ask you several questions, and you just need to answer them. Alright, no problem.

Can you introduce your role to us? Where do you work and for how long?

I work at a professional services firm in the Sales and Service department, where we mainly implement Salesforce and Microsoft Dynamics. I have been at the firm for three and a half years. I started right after my studies as an analyst, and today I am a consultant. On projects, I have a Functional Lead role, which means I manage workshops with clients, write user stories, draft test scripts, and also test solutions within Salesforce and Microsoft Dynamics. I therefore mainly take care of the functional part, and not the technical part.

How does a project generally unfold in your function as a Functional Lead?

That depends on the type of project: either an implementation project or a more strategic project. Most of the time, they are implementation projects. Generally, we start with a discovery phase. We organize workshops with the client to understand their needs, but also to identify the best way to respond to them. This notably includes the choice of tool: is Salesforce suitable, or would another tool fit better? During this phase, we also build the data model, meaning the way the objects will articulate within the CRM. Then, we move into agile mode, with sprints of three to four weeks. During these sprints, we refine the user stories with the business during several refinement sessions per week, generally two. We discuss with them, we show them how it materializes in Salesforce, we verify if it answers their needs, what data is necessary, etc. In parallel, the developers develop the user stories based on the acceptance criteria. After that, we move to the testing phases: first internally, then we write the test scripts, and then the solution is tested by the client. That is generally how a sprint unfolds. After several sprints, we finish with end-to-end tests to verify that everything corresponds well to the business processes. For more strategic projects, it is different: we produce more deliverables like PowerPoint slides. The objective is rather to recommend a solution, for example explaining why they should choose Salesforce, Microsoft, or another tool. We also create backlogs of user stories, roadmaps, and other planning documents. It is therefore a work more oriented toward analysis and presentation, with a lot of PowerPoint and Excel.

In your projects, which soft skills and hard skills do you use the most?

Concerning hard skills, I would say the main skill is knowledge of Salesforce. When I arrived at the firm, I only had basic foundations. Then, we are encouraged to pass certifications to better understand the solution, learn how to configure it, and be capable of proposing the right features. Even if we are functional, we must be capable of configuring Salesforce. You also need to be comfortable with IT in a general way. But in my opinion, soft skills are even more important. You need to know how to communicate well, both internally with your team and externally with clients. Exchanges can sometimes be complicated, notably because there is always a defined scope, and there can be misunderstandings. Certain terms can be interpreted differently by the client, which generates discussions, clarifications, and sometimes pushback. You must therefore be capable of communicating well, but also of adapting very quickly. With each project, you work in a new industry or at a new client, with completely different processes. You must manage to understand their business quickly, sometimes in a few weeks, whereas for them it is their daily life. I also think you need to be curious, not be afraid to ask questions, challenge the client when necessary, but also challenge yourself. Above all, you must keep a constant desire to learn, because if you rest on your laurels, you quickly fall behind.

If I understood correctly, you had training sessions to learn how to use the tool, is that right?

Yes, exactly. When you join the firm, you follow a two-week training session, a bootcamp. At the end, you must be capable of presenting something configured, which allows you to get a first glimpse. Then, you learn mostly on projects. We also pass Salesforce or Microsoft certifications. These are exams in the form of multiple-choice questions, but you have to study seriously and practice in the environment. This is often done in addition to work, so on your personal time. We also have training sessions more oriented toward soft skills: communication, assertiveness, running workshops, agile methodology, writing user stories, etc. Those training sessions are done in person.

In your opinion, what takes precedence most in a project: accumulated expertise, experience, personal knowledge, or a mixture of both?

I think expertise and especially experience play a very important role. Of course, it is useful to learn how to use Salesforce, but what truly brings value is the experience acquired on projects. With experience, we draw lessons on what to do or avoid, as much on the human level (how to manage the client, how to dialogue) as on the technical level. For example, we learn that a solution that is too complex can become difficult to maintain later on. I therefore think that the

more projects we carry out, the more we are exposed to different clients and situations, the more we develop experience, and the easier it becomes, especially from a functional point of view. On the other hand, from a technical point of view, when it comes to coding or developing complex features, if a person does not know how to code, it can obviously become more difficult.

In your department, what type of clients do you have? Mostly SMEs, large companies, or public institutions?

Honestly, we have a bit of everything. Personally, I work more with the public sector. My current project, for example, concerns the Brussels government, so the Brussels administrations. I have also worked with clients like Samsonite or Coca Cola. But that is linked to the fact that I am specialized in the public sector and consumer goods. We also have many clients in the financial sector, like Argenta, Arag, etc. Globally, it is very varied: there are Belgian clients, international clients, large companies, and also smaller clients. So the portfolio is truly diversified.

In your opinion, what are the ingredients to build a good relationship of trust between the consultant and the client in a project?

I think the first element is to show the client that we are listening to them. Even if we do not always know how to answer all their questions immediately, it is important to ensure a good follow up. Another essential element is to regularly show what is being developed and to make things concrete. In our field, it can be difficult for a client to project themselves. So the more we show them visible elements in the environment, the more they are reassured. They tell themselves that the project is moving forward, that it is concrete, and that it corresponds to their expectations. Finally, communication is essential. Being present with the client, organizing sessions together, and going on site when possible also allows for strengthening the link and building a relationship of trust.

Alright. So for you, is it important that the client understands the technical details during implementation?

Not necessarily the highly technical details, because that can quickly lose them. On the other hand, it is important that they have a global vision of what is being put in place and especially of the way it will translate concretely into their daily routine. The essential thing is that they can visualize what they are going to use.

And when you have moments of doubt, or you are not certain about an aspect of the project, do you talk about it directly to the client?

Yes, but in general, I do not immediately share my doubts with the client. I prefer to say that I am going to come back to them with a more precise answer. Then, I take the time to consult a colleague or to think it over calmly. I find that it is better than improvising an answer on the spot. In general, the client is understanding: we cannot know everything. So I prefer to check internally and take into account the expertise of my colleagues before answering.

Now, my next questions focus on the main subject of my thesis: generative AI. I am also at the firm internally in data analytics, but I am still going to ask you: which AIs do you use in your department?

In our department, we mainly use [Internal LLM]. We also have Codify, which is a tool oriented toward development. Codify is a tool linked to Salesforce. It notably allows for generating code. It is also an internal tool specific to the firm.

And you personally, in your missions, for which tasks do you use these tools? Personally, I use it in a fairly basic way, for example to correct emails. Codify can also be used to generate user stories. Of course, this must then be reworked and refined, but it allows at least to have a baseline to start from. I also use [Internal LLM] to get ideas for slides. I do not directly reuse the slides as they are proposed because often the result is more adapted to an HTML format than to a PowerPoint. But it gives me leads and ideas. I also use it to structure my ideas or do brainstorming. For example, if I have already prepared a list of questions on a precise theme, I can ask [Internal LLM] if I forgot certain points. It is therefore mainly a tool to complete my work.

Does using generative AI sometimes allow you to better understand your current project, or is it mainly a way to save time?

I would say that it is mostly a way to save time. Honestly, it truly makes me gain time. For example, at the beginning of a project, when a new person joins the team, AI can be very useful for onboarding. Instead of reading long documents, which are sometimes difficult to retain, we can get a clear and rapid summary. In a few hours, we can already understand the context well, whereas it would have perhaps taken two days otherwise. This also allows for having a better understanding of the business. On the other hand, once the refining sessions start and the client comes with their requests, there we work together in meetings. In those moments, we do not use AI: these discussions are not recorded, and the AI does not have access to these exchanges.

Yet it is precisely these conversations that allow for understanding the client's processes in depth.

And here it is purely your opinion: do you think that by dint of using generative AI, notably among juniors who have less experience, they risk learning less on the job and developing a more superficial understanding of their project?

Yes, I think that is possible. For juniors, AI can represent a risk, notably in terms of creativity. Before the arrival of AI, if we asked several juniors to create slides or to propose a first draft of user stories, each would have brought a different vision. It was actually interesting because juniors are not yet completely formatted and can bring a fresher or more original approach. Today, if we ask for the same thing, everyone is going to use AI, which is logical. Then, each person will rework the result a little, but in the end, we risk obtaining something very similar, very smooth, and sometimes very "AI". So yes, I think that can reduce creativity.

My question focused mainly on the way juniors are going to evolve with generative AI, knowing that we learn a lot "on the job". Does AI risk diminishing this learning experience?

Yes, I think that juniors indeed risk thinking less by themselves and therefore perhaps retaining certain information less. It is not the same thing to reflect on a problem for two hours as it is to ask AI directly to generate an answer, then to do a copy-paste. The intellectual work is not the same. That said, we cannot ask juniors not to use AI. That would be hypocritical and above all a waste of time. Today, AI is part of the tools that allow us to be more productive, and that is going to become increasingly common.

Still according to your opinion, do you think that consultants with fewer skills will be able, thanks to AI, to do the work of more experienced consultants, notably in implementation projects where it is sometimes necessary to code?

I have a concrete example on my project, on the technical implementation side. We have a developer who does not work at the firm, and he uses AI enormously. He does not use the firm's tools, but other external tools. The problem is that AI can generate code that works, but that does not necessarily correspond to the firm's standards or to the specific standards of the client. Each client sometimes has their own frameworks, their own methods, and their own rules. Result: we obtain code that works, but where best practices are not respected. Perhaps with a better prompt, that could be improved. But today, there are still quite a few limits and errors. An experienced developer is going to spot these problems quickly and correct them, which

allows them to be very efficient. On the other hand, if a person lacks technical perspective or simply seeks to go fast, and they trust AI too much, that can become problematic.

So I think that if we are very good technically, we remain relatively safe because we know how to analyze and correct what AI proposes. And then there is also a human aspect, especially in the functional part, that AI will not have.

Precisely, in your opinion, what aspect of the consultant profession does AI not replace and will not be able to replace?

I am thinking notably of the refining sessions where we are directly with the client. We discuss with them, we understand their needs, we clarify expectations, we manage exchanges live. I have a hard time imagining a client discussing with an AI, explaining their requests to it, and the AI adapting correctly in real time. I am curious to see if one day that will be possible, but today, I do not see how AI could replace this human interaction. It would have to be capable of adapting to different personalities, to different communication styles, and especially to more delicate situations, like when a client is unhappy or frustrated. Moreover, even on more technical tasks, AI still has limits. For example, when we ask it to generate acceptance criteria, there are often elements that do not correspond to the specific needs of the business. So for now, I would say that what will be replaced last are precisely the refining sessions and the whole human relationship with the client.

And when you make a prompt to generate user stories, are there many hallucinations ?

Yes, that happens, but it depends on the tools and the context. With Codify, what is good is that we can provide an enormous amount of context. And the more context we give, the fewer hallucinations there are. However, AI cannot take into account things that were discussed orally but never noted down. So if certain important pieces of information are not documented, the AI obviously cannot integrate them. But globally, I see that tools are improving more and more.