UCLouvain

LOUVAIN
School of Management

Louvain School of Management

From Market Rates to Deposit Outflows: A Threshold Model of the Cyclical Component of Belgian Regulated Savings

Author: Théo Delbouvry
Supervisor: Frédéric Vrins
Academic year 2025-2026
Dissertation for the master of Business Engineering
Master subject and focus Financial Engineering
Daytime schedule

Declaration Regarding AI Tool Usage

During the preparation of this master's thesis, the author used **Claude (Anthropic)** as a support tool. In accordance with the principles of responsible AI use, the author critically assessed every output, verified all quantitative results independently and retained full authority over the analytical and methodological choices of the work.

The tool was used for the following purposes:

1. **Writing and structure.** Reviewing, restructuring and copy editing drafts for clarity, coherence and consistency of notation; the arguments, interpretations and conclusions are the author's own.

2. **Code support.** Drafting and reviewing Python scripts used to run diagnostic and verification tests (robustness checks, alternative-model comparisons, filter-bias audits) and to generate figures. All scripts were executed by the author on the author's own data and every reported figure was reproduced and checked by the author.

3. **Methodological discussion.** Discussing modelling alternatives, their trade-offs and their justification (for example, the choice of functional form, the treatment of the trend–cycle decomposition and the decision not to pursue a machine-learning approach). Each decision was made, and is defended, by the author.

4. **Source identification.** Suggesting and locating references, which the author then consulted, verified and cited directly from the original sources.

After using this tool, the author diligently reviewed and edited all content it contributed, and independently re-estimated and verified every empirical result presented. The author takes full responsibility for the final content of this thesis. By signing this declaration, the author affirms that the content of this master's thesis reflects the author's original work, augmented by the responsible use of AI, and that the tool contributed to the work but did not replace human judgment.

Théo Delbouvry

Signature Date: 20/07/2026

T.D.

Abstract

The 2023 Belgian government-bond issue drew €17.8 billion out of regulated savings accounts in a single month, the sharpest decline on record, after the slowest deposit-rate pass-through of the past two decades left savers with a widening incentive to leave. Behavioural deposit models calibrated on the long preceding decline of interest rates had little to say about such an episode. The gap is acknowledged at the supervisory level: banks know that a rising spread between market alternatives and the rate they pay will eventually push deposits out, but the standard calibration lacks an empirically grounded answer to *how much* leaves and *from what threshold*.

This thesis provides that answer for the Belgian regulated savings account and couples it to the deposit rate. It proceeds in two connected stages. First, the client rate is modelled with an error-correction specification anchored on a funding-value pillar: a medium-term swap rate averaged over the preceding five years and a lagged short rate, with the regulatory floor imposed explicitly. The model tracks two decades of administered pricing across three structurally distinct regimes. Second, the cyclical component of deposit volumes is shown to obey a threshold rule: the deposit base is inert while the spread between the retail alternative and the client rate stays below about 1.35 percentage points. Beyond that threshold, it loses roughly €15 billion per further point of spread. The threshold is not calibrated to the episode it explains: the underlying decomposition is pinned down by an independent frequency rule and validated economically against an independently constructed measurement of where household savings actually went.

Chained together and driven by market rates alone, never shown the rate the banks actually paid, the two models reproduce the 2023 outflow to within €2.1 billion of the realised figure, with an error that is systematically conservative. Benchmarked against the regulatory standardised approach, the model is silent in the months where nothing is at stake and accurate in the months where something is, whereas the supervisory method signals a fixed outflow capacity in every month alike.

The result is best read as a building block. A market-wide aggregate identifies a single behavioural threshold; a bank's own account-level data would resolve it into the distribution of thresholds that its heterogeneous depositors hold. This thesis establishes, on public data, that the mechanism is real, that its parameters can be measured and that the chain from rates to outflows holds against realised data. It leaves a structure that granular data would extend rather than replace.

Keywords: non-maturity deposits; deposit-rate pass-through; interest-rate risk in the banking book (IRRBB); threshold model; error-correction model; replicating portfolio; deposit outflows; Belgian regulated savings.

Acknowledgements

I would first like to thank my supervisor, Professor Frédéric Vrins, for his guidance throughout this work. He was consistently available and constructive in his feedback, helped me structure the thesis and suggested the credit-risk channel explored through CDS spreads, which became one of its more instructive threads.

I am grateful to the Balance Sheet Management team at Crelan, and in particular to Olivier Dewell, for taking the time to read the thesis and to respond in detail. Their practitioner perspective sharpened several of the arguments and helped me frame the contribution more accurately.

My thanks also go to the ALM Treasury team at BNP Paribas Fortis, whose members took the time to teach me a great deal during my internship. That experience is what gave me the desire to take up this subject in the first place.

Finally, I would like to thank my family for their unfailing support throughout my studies.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Research Context

For nearly a decade, the modelling of bank deposits was a problem one could afford to get approximately right. Between 2015 and 2022, euro-area policy rates sat at or below zero and the rate paid on retail savings accounts in Belgium lay flat against its regulatory floor. With market rates pinned low and client rates immobile, the behavioural questions at the heart of asset–liability management (ALM): "how fast does the deposit rate follow the market?" and "how much volume leaves when better-paying alternatives appear?" had little practical bite. The answers did not move the balance sheet, because nothing moved at all.

The monetary tightening that began in mid-2022 ended that tranquillity abruptly. Between July 2022 and September 2023, the European Central Bank raised its deposit facility rate from −0.50% to 4.00%, in what the ECB itself describes as an exceptionally rapid tightening campaign by the standards of the monetary union.[1] For the first time since the prevailing deposit models were built, banks observed how their liabilities behave when rates rise sharply. Consistent with a long-standing regularity, the transmission was sluggish: pass-through from policy rates to deposit rates has historically been slow and incomplete, and slower still during the 2022–2023 episode (Greenwald, Schulhofer-Wohl, & Younger, 2023). In Belgium, the gap that opened between market rates and the rate paid on savings accounts was closed not by the banks but by the State: in September 2023 the Treasury issued a one-year retail government bond (*bon d'État*) whose tax treatment lifted its net yield far above what savings accounts offered. More than 500,000 households subscribed, for €21.9 billion. The counterpart is visible directly on the liability being studied here: regulated savings balances fell by €17.8 billion in August 2023 alone and by €28 billion between July and September, a fall the National Bank attributes jointly to the bond issue and to a parallel shift of households into term deposits.[2] Models calibrated exclusively on the preceding long decline of rates were, suddenly, navigating in fog. This thesis is about that fog, about building one instrument to see through part of it.

The object of study is the Belgian regulated savings deposit (*dépôt d'épargne réglementé*, hereafter the regulated account), the country's dominant retail savings vehicle, whose institutional features: a legally floored remuneration, a favourable tax treatment, a highly concentrated market are set out in Chapter 3. It belongs to the class of **non-maturity deposits** (NMDs): liabilities without a contractual maturity, which the depositor may withdraw at any time and which the bank may reprice at any time.

[1] Source: European Central Bank (2024); ECB key interest rates. The deposit facility rate moved from −0.50% (July 2022) to 4.00% (September 2023).

[2] National Bank of Belgium (2023a) and National Bank of Belgium (2023b) for these figures and their attribution

Neither their effective maturity nor their interest sensitivity is therefore observable from the contract; both must be *modelled* from behaviour. That is what makes NMDs at once the cheapest and the most analytically awkward item on a bank's balance sheet.

Two behavioural relationships govern NMDs and both run through this thesis. The first is the **deposit rate pass-through** (or *deposit beta*): the degree to which a movement in market rates is transmitted to the rate paid to clients. The second is the sensitivity of **deposit volumes** to the gap between the returns available elsewhere and the rate the account pays, the incentive depositors have to move their money. Both run through this thesis, but they are not equally apparent in the raw data: the first is visible directly (Figure 1.1), while the second, muted in the aggregate series, is what this thesis sets out to extract. Figure 1.1 plots the regulated client rate against the ECB policy rate over 2003–2026: the client rate stays near its regulatory floor for most of the sample and follows the sharp 2022–2023 tightening only reluctantly and with delay (panel a). Its month-to-month adjustment also *appears* asymmetric. Visibly stronger when market rates fall than when they rise (panel b). It is a pattern the literature characterises as *downward-flexible but upward-sticky* (Driscoll & Judson, 2013). Whether that appearance survives estimation, or dissolves into the composition of a market-wide average, is itself one of the questions this thesis settles (Chapter 4). What is not in doubt is the stickiness of panel (a), and the deposit outflows it ultimately drives.

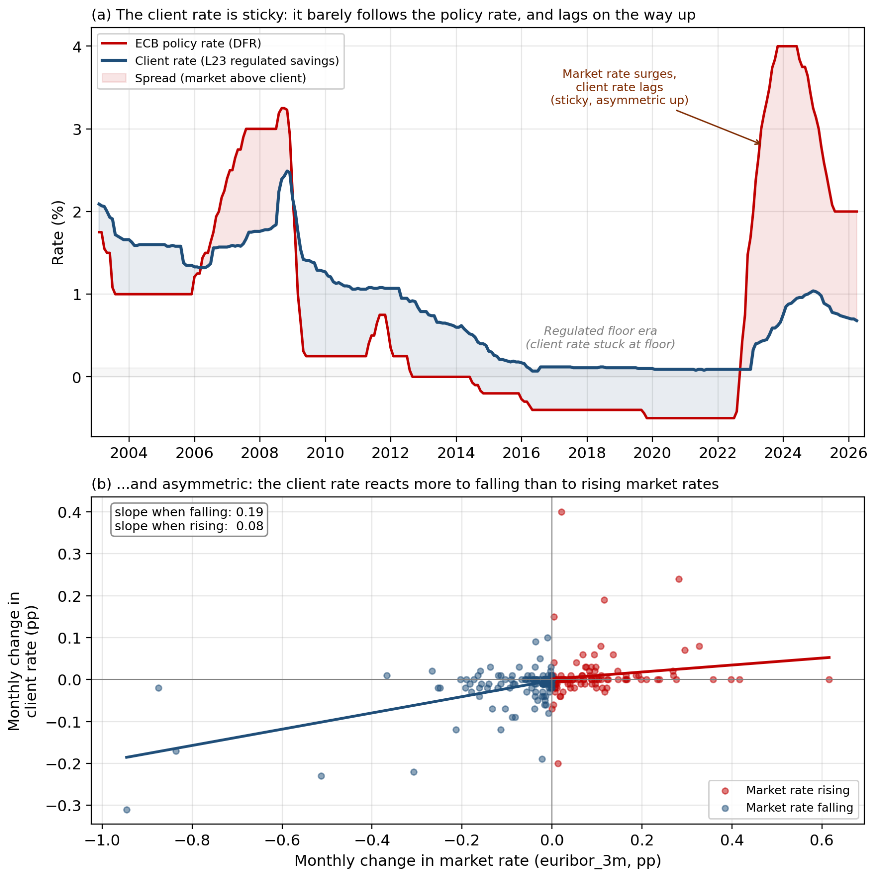

Figure 1.1: Two salient features of the regulated client rate, Belgian data (2003–2026). (a) Client rate (L23) versus ECB policy rate; the shaded area is the market-above-client spread. (b) Monthly change in the client rate against the monthly change in the market rate, split by direction: the fitted slopes are 0.19 when rates fall and 0.08 when they rise. Whether this reflects an asymmetric mechanism is examined in Chapter 4. Sources: Chapters 3 and 4.

1.2 Problem Statement

The behavioural modelling of NMDs is not an optional refinement but a regulatory requirement: under the Basel IRRBB standards and their transposition by the European Banking Authority, banks must model the repricing and volume behaviour of these deposits, the key assumptions must be "conceptually sound, documented and rigorously tested" (Basel Committee on Banking Supervision, 2016; European Banking Authority, 2022b). That framework is reviewed in Chapter 2; what matters here is how it is met in practice. Deposit rates and deposit volumes are typically modelled separately, the link between them, the fact that outflows respond to the gap between the returns available elsewhere and the rate the account pays, is supplied largely through expert judgement: withdrawal rates and attractiveness thresholds set from historical experience rather than estimated from it (Section 2.4).

The limitation of this practice is not one of diligence but of data. Until 2022, the historical record contained no sharp upward rate cycle against which an outflow response could be calibrated; behavioural models were, in the ECB's own assessment, calibrated only on falling rates, which adds model risk to the resulting interest-rate-risk measures (Molyneux, Pancotto, Reghezza, & Rodriguez d'Acri, 2020). The consequence became visible during the tightening: banks whose deposits were the most interest-rate-sensitive did not update their NMDs maturity assumptions as rates rose. A finding the ECB closes by underscoring how critical the timely and accurate calibration of those assumptions has become (Coulier, Pancaro, Pancotto, & Reghezza, 2025).

The gap is therefore specific and acknowledged at the supervisory level. Banks know that a widening spread between market alternatives and the rate they pay will eventually push deposits out; what the standard calibration lacks is an empirically grounded answer to *how much* leaves and *from what threshold*. The pre-2022 sample could not provide one. The 2022–2023 episode can and the stake it exposed is not hypothetical. Reviewing the Belgian episode a year later, the National Bank records that the sudden loss of some €28 billion of deposits over the 2023 summer compressed banks' interest margin sharply and forced them to refinance on the market, at market rates, the funding that had walked out of the savings accounts (Cappoen, 2024). Quantifying that displacement is the gap this thesis addresses for the Belgian regulated account.

1.3 Research Question

This thesis addresses the following question:

> How can the client rate and the cyclical volume dynamics of the Belgian regulated savings account be modelled and coupled, so that deposit outflows can be deduced from the interest-rate environment ?

Neither model is taken as given. The specification of each (Drivers of the client rate and threshold of response) is itself an outcome of the thesis, selected from competing configurations and validated empirically (Chapters 4 and 5). The alternatives considered and the evidence behind each choice are reported in Appendices B and C. The central contribution, however, is not either model taken alone but their *coupling*: outflows respond not to market rates directly but to the *spread* between the returns available to depositors and the rate their regulated account pays: a spread the bank itself sets, with delay. The thesis therefore proceeds in two connected stages: first a model of the client rate, then a model of cyclical volumes driven by the resulting spread. The two are then chained and the chain is confronted with the realised outflows of the sample.

1.4 Contribution

The contribution of this thesis is threefold.

First, the thesis specifies and estimates a spread-to-outflow relationship for the Belgian regulated account over a sample that, uniquely, now spans a full easing and tightening cycle. It documents that this relationship is **non-linear and threshold-shaped**: below a spread threshold the deposit base is inert; beyond it, outflows accelerate sharply. This gives an empirical, Belgian-specific content to the generic spread term of existing frameworks. Benchmarked against the standardised approach prescribed by the regulator (European Banking Authority, 2022b), the estimated relationship proves *conditional*, silent when no arbitrage signal is present, accurate when one is, where the supervisory method posits a fixed, unconditional outflow capacity; this echoes recent evidence that standard models overstate depositors' rate sensitivity (Cirelli & Ólafsson, 2025).

Second, and as its central contribution, the thesis **chains the two models and shows that the chain holds**. An interest-rate environment feeds a specified client-rate model; the modelled rate determines the spread; the spread drives outflows through the estimated threshold. This closes a loop that practice typically leaves open, in which rates and volumes are modelled side by side but never joined. The claim is not merely that the loop can be closed but that it survives the closing: driven by market rates alone, and never shown the observed client rate, the chain reproduces the 2023 episode to within €2.1 billion of the realised outflow, with an error that is systematic in direction and conservative in sign (Chapter 5). What is demonstrated is therefore not a projection but a property, that the outflows of this deposit base are, to a useful approximation, a function of the rate environment.

Third, the thesis is explicit about its **methodological itinerary**: rather than presenting a finished model, it documents the specifications tested and the evidence that led certain choices to be retained and others set aside. That itinerary is substantial. A trend anchored in macroeconomic fundamentals and the parametric trends standard in the deposit-modelling literature are both rejected on these data; a below-threshold slope, statistically present, is shown to rest on the least reliable observations of the sample and is discarded; a bank-credit-risk channel and a policy-uncertainty channel are tested, found correctly signed, and still not retained: the first because it cannot be identified on aggregate data, the second because it cannot be conditioned on a rate scenario. It also includes the deliberate decision not to pursue a machine-learning approach: because the client rate is in part an *administered*, semi-discretionary decision, bounded by a legal floor and shaped by commercial and competitive considerations, an interpretable structural model, in which each term carries an economic meaning, is more appropriate to the object than a black-box predictor optimised for fit alone.

The methodological tools these contributions rest on are established and the thesis draws on them as such: the error-correction framework follows the standard pass-through literature (de Bondt, 2005; Driscoll & Judson, 2013); the funding-value anchor of the client rate adapts the replicating-portfolio logic applied to Belgian regulated deposits by Maes and Timmermans (2005); and the threshold form of the volume response draws on the least-squares threshold estimator of Hansen (2000) and on the non-linear deposit-behaviour literature (Blöchlinger, 2015; You & Noh, 2025). What is new is not the instruments but their specification, selection and coupling.

Taken together, these contributions are best read as a **building block**. The relationship estimated here is the reaction function of a single representative depositor with a single exercise threshold because a single threshold is what a market-wide aggregate containing two arbitrage episodes can identify. Yet the deposit base is not homogeneous: depositors differ in attentiveness, in switching costs and in the size of the balance at stake, and each holds, in effect, an option to move funds with its own strike. On account-level data, the same specification estimated segment by segment would yield a portfolio of thresholds: a multi-strike outflow model of direct operational use, of which the form retained here is the one-bucket case. This thesis does not deliver that model and does not claim to. It establishes, on the data available, that the mechanism such a model would be built on is real, that its parameters can be measured and that the chain from rates to outflows holds against realised data; it leaves a structure that richer data would extend rather than replace. That trajectory is developed in Chapter 6.

1.5 Structure of the Thesis

The remainder of the thesis is organised as follows. Chapter 2 reviews the literature on NMDs modelling, deposit-rate pass-through and the replicating-portfolio approach. It also positions the present work within it. Chapter 3 describes the institutional setting of the Belgian regulated account, the data and their construction. Chapter 4 develops and validates the client-rate model. Chapter 5 isolates the cyclical component of deposit volumes, develops the threshold model that links it to the market-to-deposit spread, chains that model to the client rate of Chapter 4 and sets the result against the regulatory standardised approach. Chapter 6 concludes: it sets out what the model establishes and what it does not, discusses its limits and traces the path from the single-threshold specification retained here to the multi-strike model that account-level data would support.

Chapter 2

Literature Review

This chapter situates the thesis within four strands of literature. It covers, in turn, non-maturity deposits and the regulatory framework governing their modelling (Section 2.1), deposit-rate pass-through (Section 2.2), the replicating-portfolio approach underpinning the funding-value view of the deposit rate (Section 2.3) and the modelling of deposit volumes (Section 2.4). Section 2.5 draws the threads together and states the gap the thesis addresses.

2.1 Non-Maturity Deposits and Behavioural Maturity

Non-maturity deposits occupy a peculiar position on the bank balance sheet. Contractually, they are floating-rate liabilities of zero maturity: the depositor may withdraw at any time and the bank may reset the rate at any time. Behaviourally, neither option is exercised at every movement of rates. A large and stable fraction of balances, the *core*, remains with the bank across years and through rate cycles. The rate paid adjusts only partially and with delay to market conditions. This divergence between contractual and behavioural characteristics is the central analytical difficulty: the effective maturity and the interest-rate sensitivity of NMDs are not observable from the contract but must be inferred from observed behaviour.

The economic source of this stability is what the recent literature terms the *deposit franchise*. Drechsler, Savov, and Schnabl (2017) show that banks' market power in deposit markets is the foundation of a stable, low-cost funding base: depositors are imperfectly mobile and value the liquidity and convenience of their accounts, so banks can pay well below market rates without triggering large outflows. In a follow-up, Drechsler, Savov, and Schnabl (2021) argue that this market power lets banks fund long-duration assets with short-term deposits while bearing little interest-rate risk: the deposit rate moves so little that the franchise behaves like a fixed-rate liability. The deposit franchise thus simultaneously explains the stability of volumes and the stickiness of rates, the two behavioural relationships at the heart of this thesis.

This behavioural character is not merely an academic observation; it is the object of explicit prudential regulation. Under the Basel standards on interest-rate risk in the banking book (Basel Committee on Banking Supervision, 2016), banks must model the repricing and volume behaviour of NMDs. The European Banking Authority operationalises these standards for EU institutions (European Banking Authority, 2022b), notably capping the assumed repricing maturity of NMDs at 5 years, an acknowledgement that, unconstrained, behavioural models can assign these instantaneous liabilities a long effective life. The modelling choices this thesis examines are therefore made within, and partly shaped by, a binding regulatory frame.

2.2 Deposit-Rate Pass-Through and Its Asymmetry

The transmission of market rates to deposit rates: the pass-through or deposit beta, has been studied for three decades and two robust regularities emerge: pass-through is *incomplete* and it is *asymmetric*.

That deposit rates respond only partially and sluggishly is among the most durable findings in banking. Hannan and Berger (1991) document that deposit rates are rigid, adjusting slowly to market movements and Neumark and Sharpe (1992) link this rigidity to market structure: deposit rates are stickier in more concentrated markets, where banks face less competition. For the euro area, de Bondt (2005) confirm that this pass-through is slow, incomplete and variable across countries and products.

The asymmetry is equally well established and its *level* of operation matters for what follows. Driscoll and Judson (2013), examining a panel of over 2,500 branches, sharpen the regularity: deposit rates are *downward-flexible but upward-sticky*. Banks cut it promptly when rates fall and raise it reluctantly when they rise. It is a pattern they rationalise with a *menu-cost model*, in which each bank faces a fixed cost of changing its posted rate and adjusts only when the benefit exceeds that cost. This is, by construction, a theory of *individual* pricing decisions and its central empirical signature is heterogeneity: the degree of asymmetric rigidity varies across deposit types, bank sizes and even across branches of the *same* institution. Stickiness, the slow and partial adjustment, survives aggregation and is detectable in market-wide average rates; the asymmetry, rooted in institution-level menu costs, is a finer phenomenon that averaging across heterogeneous rate-setters may wash out. The implication, to which we return in Chapter 4, is that an aggregate, market-level series is the level at which stickiness is expected to be identifiable but at which the individual-level asymmetry need not be.

Market structure conditions both. Kho (2025), studying the euro area, shows that deposit-rate pass-through is shaped by deposit-market concentration: in more concentrated markets, rate increases are passed through more slowly and decreases more quickly, so that market structure governs the speed at which ECB policy reaches depositors. This is directly relevant to Belgium, whose retail deposit market is highly concentrated (Chapter 3), and it bridges the institutional structure of the Belgian market and the sluggish pass-through this thesis estimates. The mechanism also has a contemporary edge: a Federal Reserve analysis of the post-2022 euro-area tightening finds pass-through to have been even slower than in previous hiking cycles, attributing the sluggishness partly to abundant liquidity on top of imperfect competition (Greenwald et al., 2023).

A complementary literature stresses that the deposit beta is not constant but *state-dependent*. Work on the post-2022 period documents a *dynamic* beta that rises with the level of rates: as market rates climb, banks are forced to pass through a growing share, so that the sensitivity of deposit rates increases precisely when rates are high (Greenwald et al., 2023). This non-constancy is the empirical motivation for modelling the client rate with a regime-aware specification rather than a single fixed coefficient, the approach taken in Chapter 4.

2.3 Replicating Portfolios and the Funding-Value Interpretation

If the deposit rate is sticky and the balance is stable, what is the "right" rate against which to price and hedge the deposit? The replicating-portfolio approach answers this by mapping the non-maturity deposit onto a portfolio of market instruments that reproduces its margin behaviour. The deposit is represented by a rolling ladder of fixed-income tranches, typically swap rates of staggered maturities, whose blended yield approximates the return a bank earns by investing the stable deposit base. The average yield of this ladder defines the funding value of the deposit and, in transfer pricing, the rate at which the treasury remunerates the business line for the deposits it gathers.

The foundational contributions formalise this idea as an optimisation. Kalkbrener and Willing (2004) develop a general quantitative framework for managing the liquidity and interest-rate risk of non-maturing liabilities, representing the deposit by a replicating portfolio chosen to match its risk characteristics. Frauendorfer and Schürle (2007) and the related work of Paraschiv and Schürle (2010) cast the problem as a dynamic optimisation, modelling jointly the client rate and the volume so as to construct a replicating portfolio that is robust across interest-rate scenarios. Blöchlinger (2015) provides a theory of deposit pricing that identifies and values the embedded options in non-maturity deposits, linking the replicating-portfolio margin to the optionality the depositor holds, a reading developed in Section 2.4.

For the Belgian case specifically, Maes and Timmermans (2005) measure the interest-rate risk of regulated savings deposits using a replicating-portfolio methodology; their work is the direct antecedent of the funding-value pillar used in this thesis. It is useful, at this point, to make explicit a decomposition that organises both the literature and the present work. The deposit balance can be split into a *structural* (or stable) component: the core that remains with the bank across rate cycles and a *cyclical* (or non-structural) component: the portion that expands and contracts with the rate environment as depositors arbitrage between their account and outside options. These two components call for different treatments. The structural component is the natural domain of the replicating portfolio in its full sense: a long-maturity ladder, optimised so that its blended yield tracks the slow-moving funding value of a base that is, by construction, stable. It is for this stable core that the optimisation problems of Kalkbrener and Willing (2004), Frauendorfer and Schürle (2007) and Maes and Timmermans (2005) are designed.

The present thesis works on the other component: the behaviour of the *cyclical* part of deposit volumes. This part is governed not by long-term funding value but by short-term arbitrage. We therefore do not construct and optimise a full replicating portfolio: that machinery answers a question about the stable base, whereas the question of this thesis concerns the movements of the cyclical layer. What we retain from the replicating-portfolio literature is its funding-value insight, in the form of a single constructed *pillar*: a trailing moving average of a medium-term swap rate, proxying the yield of a rolling replication ladder, which serves as the structural anchor of the client-rate equation (Chapter 4). The replicating-portfolio logic is thus not incidental to this thesis but the source of the client rate's principal driver, which is why it is set out here at length; the term "pillar" is reserved throughout for that single series, to avoid conflation with the full optimisation problem.

This division of labour also locates the contribution of the thesis. The structural component and its replicating portfolio are well understood; the cyclical component, by contrast, is precisely where the post-2022 environment exposed the limits of existing calibration. It is the dynamics of that cyclical movement, left generic by the structural machinery, that the coupling of rate and volume in Chapter 5 sets out to quantify.

2.4 Modelling Deposit Volumes

The literature reviewed so far concerns the rate. The volume side has its own, smaller tradition, and it converges on a common architecture: the balance is split into a slow-moving component and a residual that responds to the rate environment. Kalkbrener and Willing (2004) model deposit volumes as an Ornstein–Uhlenbeck process around a linear trend, correlated with market and deposit rates. Bardenhewer (2007) splits the volume explicitly into a trend component: linear, quadratic or exponential and an unexpected component, alongside moving averages of market rates. The regulatory standardised approach adopts the same architecture under different names, separating a *core* share, capped by supervisory convention, from a *non-core* share deemed susceptible to outflow (European Banking Authority, 2022b). Whatever the vocabulary: trend and unexpected component, core and non-core, structural and cyclical, the frame is the one set out in Section 2.3. This thesis adopts it.

Two features of this literature bear directly on what follows. First, the trend is typically *parametric* and estimated *jointly* with the rate response, in a single equation. That choice has real advantages: no smoothing parameter to defend, no pre-filtered series carried into a second stage. Chapter 5 tests it on Belgian data before adopting a filter-based decomposition instead; the comparison is reported in Appendix C. Second, and more consequentially, the term that links volumes to the rate environment is left *generic*. The frameworks posit that outflows respond to the gap between market alternatives and the rate paid, but the shape of that response, from what gap it begins and how steep it becomes, is supplied by assumption or by expert judgement rather than estimated. It is this term that the present thesis calibrates.

The reason the response should not be linear is best seen through the option view of the deposit. Blöchlinger (2015) formalises the non-maturity deposit as a bundle of embedded options: the bank holds an option on the rate it pays, the depositor an option to withdraw. Read this way, the depositor's decision has the structure of an exercise decision. Leaving the account is not free, it costs attention, time, the loss of a fidelity premium, the friction of opening an account elsewhere. The option is exercised only once the return gap exceeds those costs. The gap at which this happens is, in the language of options, a *strike*: below it the option is out of the money and the balance is inert; beyond it, exercise becomes rational and money moves. This is why a threshold, rather than a slope, is the natural form of the volume response and it is the form Chapter 5 estimates.

The option view also states, immediately, what a single threshold leaves out. Depositors are not identical: they differ in attentiveness, in switching costs, in financial sophistication and in the size of the balance at stake, so their strikes differ too. The response of an entire deposit base is therefore the aggregate of many individual exercise decisions with dispersed strikes, a smooth, curved function of the spread, of which a single kink is the parsimonious approximation, a structure whose extension is taken up in Chapter 6.

2.5 The Gap in the Literature

The four strands reviewed above converge on a coherent picture, but they leave a specific question open. The deposit-franchise literature (Drechsler et al., 2017, 2021) explains *why* non-maturity deposits are stable and their rate sticky, but treats these as broadly structural properties rather than quantifying how the stable base erodes when arbitrage becomes attractive. The pass-through literature (de Bondt, 2005; Driscoll & Judson, 2013; Hannan & Berger, 1991; Neumark & Sharpe, 1992) characterises the *rate* side in fine detail: its incompleteness, its asymmetry, its dependence on market structure (Kho, 2025), but stops at the rate, without carrying the analysis through to the deposit *volumes* that the rate, via the spread to outside options, ultimately drives. The replicating-portfolio literature (Blöchlinger, 2015; Frauendorfer & Schürle, 2007; Kalkbrener & Willing, 2004; Maes & Timmermans, 2005) provides the machinery for the *structural* component, optimising a funding-value ladder for the stable core, but by construction says little about the *cyclical* layer that expands and contracts with the rate cycle. The volume literature (Bardenhewer, 2007; Kalkbrener & Willing, 2004) supplies the right decomposition but leaves its arbitrage term generic. The regulatory frameworks (Basel Committee on Banking Supervision, 2016; European Banking Authority, 2022b) do impose a treatment of that cyclical layer but a deliberately generic and prudential one: a fixed, unconditional outflow capacity which, as the standardised approach itself concedes, is a prudential simplification rather than a calibrated model, which the most recent evidence suggests overstates depositors' true rate sensitivity (Cirelli & Ólafsson, 2025).

The gap thus lies at the intersection of these strands: the *empirical calibration* of the cyclical component of deposit volumes and its *coupling* with the client rate, so that the two sides of the deposit, modelled side by side in the literature, are joined. The component-based decomposition that structures current practice provides the frame; what it leaves open is the cyclical term, particularly for a national market observed over a genuine tightening episode. This is the space the thesis occupies.

Chapter 3

Data and Institutional Context

This chapter describes the institutional setting of Belgian regulated savings deposits, the data sources and sample, the construction of the key series and the data-quality checks performed prior to estimation. Particular attention is paid to the legal and fiscal specificities of the Belgian regulated savings account, which shape the deposit-rate dynamics modelled in Chapter 4.

3.1 The Belgian Regulated Savings Account

The empirical analysis focuses on the Belgian regulated savings deposit, the dominant retail savings vehicle in Belgium, with an aggregate outstanding amount of approximately €311 billion.[1] Three institutional features distinguish this product from a standard non-maturity deposit and motivate the modelling choices made in this thesis.

First, the remuneration is **regulated and floored**. The total return combines a base rate and a fidelity premium, both subject to legal minima set by Royal Decree.[2] Over the sample, the effective floor on the observed client rate is approximately 0.11 %, a level that becomes binding throughout the zero-lower-bound period and which we treat explicitly in the rate model.

Second, the product benefits from a **favourable tax treatment**: interest income up to a legal threshold is exempt from the withholding tax that applies to most other savings instruments.[3] This exemption is a structural source of stickiness: it raises the after-tax return of the regulated account relative to nominally higher-yielding alternatives and partly explains the low price elasticity documented in Chapter 5.

Third, the Belgian retail banking market is highly **concentrated**. On the ECB's harmonised measure, the five largest credit institutions hold approximately 74% of total banking-sector assets, placing Belgium among the more concentrated euro-area systems, well above Germany (about 35%) and France (about 46%).[4] The regulated-savings market in particular is dominated by four institutions: BNP Paribas Fortis, KBC, Belfius and ING, whose combined regulated-savings balances, of the order of €220 billion, represent the large majority of the €311 billion market.[5]

[1] Source: European Central Bank (2026b), series `BSI.M.BE.N.A.L23.A.1.U2.2250.Z01.E` (deposits redeemable at notice up to three months, households, Belgium; outstanding amounts), value for March 2026.

[2] See Article 2(4) of the Royal Decree implementing the Belgian Income Tax Code (AR/CIR 92), as amended by the Royal Decree of 21 September 2013, which defines the remuneration structure of regulated savings deposits (base interest rate and loyalty premium).

[3] Source: SPF Finances, « Revenus de l'épargne et des placements (revenus mobiliers) », https://fin.belgium.be/fr/particuliers/declaration-impot/revenus/epargne-placements, accessed on 18/06/2026. Montants applicables aux revenus 2025–2026.

[4] Source: European Central Bank (2026a), share of total banking-sector assets held by the five largest credit institutions (CR5), series `SSI.A.BE.122C.S10.X.A1.Z0Z.Z` for Belgium and the corresponding series for Germany and France, [2026].

[5] Order of magnitude from the four banks' published balances; the precise figures are not central to the argument.

This oligopolistic structure is directly relevant to the deposit-rate dynamics modelled in this thesis. Deposit-market concentration slows pass-through (Kho, 2025), so Belgian concentration predicts a *sluggish* transmission, and that is what Chapter 4 estimates.

3.2 Sources and Sample

The dataset is monthly and spans January 2003 to March 2026 (279 observations). It combines data from the European Central Bank, the National Bank of Belgium, the Belgian Federal Debt Agency, Bloomberg, Eurostat and Statbel. Table 3.1 summarises the variables, their definitions and notations.

Table 3.1: Variables, notations and definitions

Variable	**Notation**	**Definition**
Core series		
Client rate (L23)	r_t	Regulated savings rate (base + loyalty premium)
Deposit volumes (L23)	Vol_t	Outstanding regulated savings balances
Market rates and drivers		
€STR / EONIA		Overnight risk-free rate (spliced)
Euribor 3M/6M/1Y		Interbank reference rates
5Y swap €STR, 60M MA	p_t	Funding-value pillar
OLO 1Y, 2Y		Belgian sovereign yields (tested)
Bon de caisse 3M, 12M	c_t	Retail term-deposit rate (competing instrument)
Substitute instruments (reallocation, §5.1.4)		
Term deposits (L22)		Household deposits with agreed maturity
Retail government bonds		*Bons d'État*, retail holdings
Auxiliary channels (tested, App. C)		
CDS 5Y, iTraxx		Bank credit-risk measures
Macro controls		GDP, confidence, inflation, unemployment, EPU, real estate

Note. Series are monthly, spanning 2003M1–2026M3; several cover shorter windows. Exact start dates, end dates and sources are in Appendix A.1.

Two groups of series support arguments made outside the main model. The **substitute instruments**: household term deposits (ECB balance-sheet statistics) and retail government bonds (reconstructed from Belgian Federal Debt Agency records) are the statistical counterparts of the destinations the National Bank identifies for the July–September 2023 outflow (National Bank of Belgium, 2023b), and validate the cyclical component in Section 5.1.4 against an independent, flow-based measurement. The **bank-credit-risk** series, five-year senior CDS spreads for the major Belgian deposit banks and the iTraxx Europe Senior Financials index, are used in Appendix C.6 to test, and ultimately reject, a credit-risk channel. This panel includes Dexia SA, whose spreads carry the Belgian banking stress of 2008–2011 that its successor does not; CDS quotation begins only in 2006, a shorter window than the core sample.

3.3 Construction of Key Series

Three constructed series play a central role in the analysis and warrant a detailed description.

Spliced overnight rate. The risk-free overnight rate combines EONIA (until its discontinuation) and the €STR. To ensure continuity, EONIA values are adjusted by the official spread of 8.5 basis points fixed by the ECB upon the transition,[1] yielding a single consistent series over the full sample.

Replication pillar. Following the static replicating-portfolio logic of Maes and Timmermans (2005), we construct a *pillar* as a trailing moving average of a market rate, computed over a **rolling window of fixed size** (here 60 months): at each month t, the pillar is the arithmetic average of the market rate over the preceding 60 months, so that the window advances one month at a time while keeping a constant length. The maturity of the underlying rate and the length of the window are not fixed a priori: they are selected empirically in Section 4.3.

Reallocation series. To measure, independently of the L23 series, the money that left the regulated account for rate reasons, we construct a *reallocation* series from the two substitute instruments (household term deposits, detrended; retail government bonds, cumulated from net subscriptions), sign-inverted. It serves as the external validation of Section 5.1.4; its full construction is detailed in Appendix A.2.

3.4 Data Quality

Prior to estimation, the series were screened for coding inconsistencies. A non-trivial issue affected the *bons de caisse* series: several missing observations were coded as exact zeros rather than as missing values, notably in the 12-month series during four months of the September 2023 government-bond episode (Chapter 5). Left uncorrected, these spurious zeros collapse the market-to-deposit spread that drives the volume model, precisely over the window that identifies its threshold. The affected values were re-coded as missing and linearly interpolated between adjacent observations of the order of 3.5 %, a safe correction applied throughout. The episode underscores the need to distinguish encoded missing values from genuine zeros before using a rate series.

[1] ECB recalibration of EONIA as €STR + 8.5 basis points, effective 2 October 2019.

Chapter 4

Modelling the Client Rate

The coupling at the heart of this thesis runs from interest rates to deposit volumes: cyclical outflows are driven by the *spread* between the alternatives available to depositors and the rate paid on the regulated account (Chapter 5). Quantifying that spread requires, as a first step, a model of the client rate itself. This chapter specifies and validates one.

The chapter is organised in six sections. Section 4.1 sets out the general modelling framework and the choices it leaves open; Section 4.2 establishes an empirical constraint any specification must respect: the pass-through is regime-dependent; Section 4.3 resolves the open choices through a systematic out-of-sample search; Section 4.4 states the retained model; Section 4.5 interprets its estimates and Section 4.6 validates it.

Throughout, reported t-statistics are computed with heteroskedasticity and autocorrelation consistent (HAC, Newey–West) standard errors. This choice is not incidental: as Section 4.6 shows, the residuals of an administered, step-wise rate are autocorrelated by construction and HAC inference is precisely the tool that remains valid under these conditions. OLS standard errors, which would overstate significance, are not used.

4.1 Modelling Framework

We model the client rate within the error-correction framework standard for retail-rate pass-through, following the European Central Bank's own methodological guidance (European Central Bank, 2006). This framework fixes a general structure but deliberately leaves several choices open; the purpose of this chapter is to resolve those choices empirically rather than by assumption.

The general structure has three elements. First, a **long-run equilibrium** relates the observed client rate r_t (the L23 series of Chapter 3) to one or more market rates that proxy the value the bank derives from its deposit base:

$$r_t = f(\text{market rates}_t) + \varepsilon_t, \tag{4.1}$$

where ε_t is a disturbance. The fitted value of this relationship defines the **equilibrium rate** r_t^*: the level towards which the client rate gravitates given market conditions. Second, a **short-run dynamic** governs the speed at which the observed rate adjusts towards that equilibrium. Third, the **regulatory floor** of the Belgian product truncates the rate from below at 0.11%.

Within this structure, three choices are left free and must be settled: (i) *which* market rates enter the equilibrium, how the funding value of the deposit is proxied and whether the policy stance enters directly; (ii) the *transmission lag* with which market movements reach the administered rate; and (iii) the *functional form* of the short-run adjustment, in particular whether it is symmetric. The alternatives considered for each choice and the evidence behind each decision, are reported in Appendix B.

4.2 Stylised Facts and Regime Identification

4.2.1 A Regime-Dependent Relationship

The pass-through from market rates to the regulated savings rate is not stable over the sample. Figure 4.1 plots the client rate against a representative market rate, revealing three phases. Rather than impose these breaks a priori, we test for them formally and anchor the retained dates on the associated policy events.

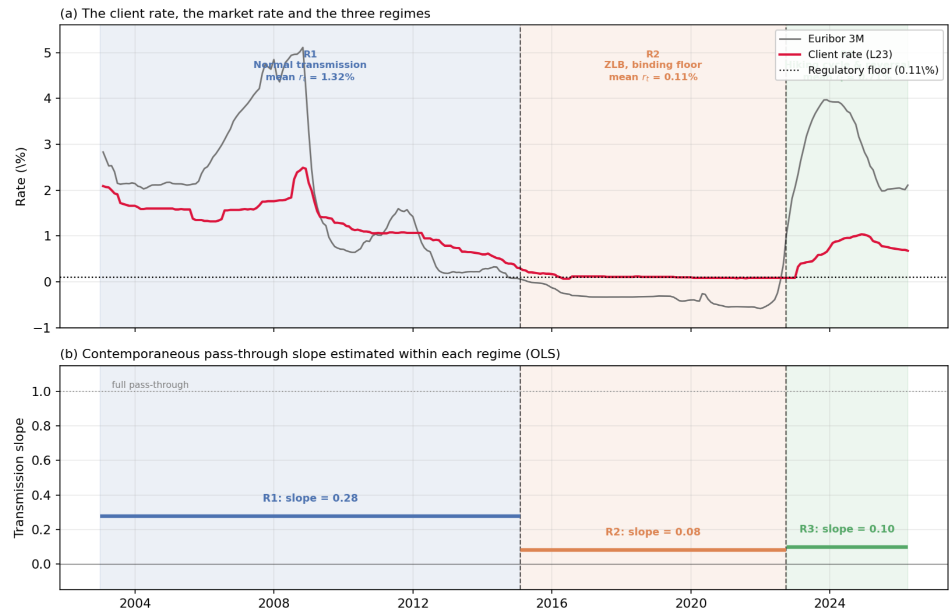

Figure 4.1: The three pass-through regimes. The regulated client rate (L23) plotted against a representative market rate, 2003–2026, with the two retained break dates marked and the three regimes shaded. Source: author's computations.

We assess the stability of the market-to-client relationship with three complementary procedures. A CUSUM and CUSUM-of-squares analysis of the recursive residuals signals parameter instability from the mid-2010s onward. The Andrews *SupF* test for a joint break in the intercept and slope strongly rejects stability (*SupF* = 438.8, against a 1% critical value of 13.15). Finally, the sequential Bai and Perron (2003) procedure, applied to the level of the client rate and to the market-to-deposit spread, rejects the null of no break at the 1% level, with statistics an order of magnitude above their critical values (e.g. $SupF(0 \text{ vs } 1) = 152.3$ on the spread, against a 1% critical value of 14.5).

The tests establish unambiguously that a single linear specification cannot describe the full sample. They locate the instability in the early-to-mid 2010s and again in 2022. As the precise break date varies by a few quarters across procedures and variables, we retain two dates with a clear institutional interpretation rather than the point estimate of any single test: **January 2015**, when the ECB announced its expanded asset-purchase programme, anchoring the curve deep in negative territory and making the regulatory floor binding, and **September 2022**, when the deposit facility rate turned decisively positive. Table 4.1 reports the three regimes and, for each, the simple arithmetic mean of the monthly client rate over the regime's months.

Table 4.1: Regimes, dates and characteristics

Regime	Period	n	Mean client rate	Characteristic
R1	2003M1–2015M1	145	1.32%	Unconstrained transmission
R2	2015M2–2022M9	92	0.11%	ZLB, binding floor
R3	2022M10–2026M3	42	0.71%	Hiking cycle and reversal

Note. The mean client rate is the simple arithmetic mean of the observed monthly client rate over the months of each regime. Break dates are selected on monetary-policy events and are consistent with the structural-break tests discussed in the text. Source: author's computations.

4.2.2 Economic Reading of the Three Regimes

The three regimes pose distinct modelling challenges (Table 4.1, Figure 4.1). **R1** (2003–2015) has positive market rates and an unconstrained pass-through. **R2** (2015–2022) is the zero-lower-bound plateau: with the policy rate negative from 2014, the market-implied equilibrium falls below the legal minimum, the observed rate is pinned at the 0.11% floor, and any pass-through estimated over this window is mechanically contaminated by censoring. **R3** (2022–2026) is the most demanding: the first genuine hiking cycle, followed by an incipient reversal, over which the market-to-deposit spread, defined throughout as the competing market rate minus the client rate, so that it is positive when the outside option pays more, averages +2.19% as market rates outrun the sluggish client rate. These are the constraints the specification search must accommodate.

4.3 Specification Search

We resolve the three open choices of Section 4.1: which market anchor, how many, at what lag by out-of-sample performance across expanding-window origins. The literature suggests three families of specification (von Feilitzen, 2011), which are *nested*: a pillar is itself a smoothed combination of market rates, so a single-pillar model (Family 1) is the parsimonious special case of a two-pillar model (Family 2) and of a free multi-rate combination (Family 3). We estimated all three across a grid of maturities, windows and lags, 54 specifications in all, evaluating each on the out-of-sample RMSE of the recursively simulated rate beyond 2019 and its median across eleven expanding-window origins. The three families, the grid and the full comparison are set out in Appendix B; here we state the outcome.

The verdict is clear-cut and it favours the most constrained family. The **two-pillar family** fits well but is not identified: its anchors measure overlapping funding values, so no specification achieves both cointegration and significance (the smallest market t-statistic never exceeds 1.5), its good fit is the signature of collinearity, not of a second economic force. The **multi-rate family** is dominated throughout: without a pillar's smoothing, it inherits the spot-rate volatility the administered client rate lacks. The **single-pillar family**, one funding anchor plus a lagged short rate, wins decisively. Its margin is largest on the demanding 2022–2026 window. That the richer, less constrained families lose is the point: the parsimonious funding-anchor-plus-transmission structure is not a default but an empirically established optimum. The retained specification is

$$r_t = \alpha + \beta_p\, p_t^{(M,W)} + \beta_s\, m_{t-L} + \varepsilon_t, \tag{4.2}$$

where m_{t-L} is a money-market rate (Euribor) at transmission lag L and $p_t^{(M,W)}$ the W-month moving average of the maturity-M swap rate, the proxy for a replicating ladder's yield. Its parameters M, W and L are fixed next.

4.3.1 The Transmission Lag

The lag L is selected by the same out-of-sample criteria. Figure 4.2 sweeps lags of one to twelve months: the RMSE profile is sharply U-shaped, falling to a minimum at **four months** (0.063) and deteriorating steeply on either side: 0.11 at lag 3, 0.24 at lag 5. The discrimination is strong: the retained lag cuts its nearest competitor's error by more than 40% and the asymmetry is informative: overshooting the true delay is punished far more than undershooting it, consistent with a rate set on a fixed institutional calendar that a longer lag simply misses. The expanding-window criterion confirms the ordering (median RMSE 0.072 at lag 4 against 0.078 at lag 3).

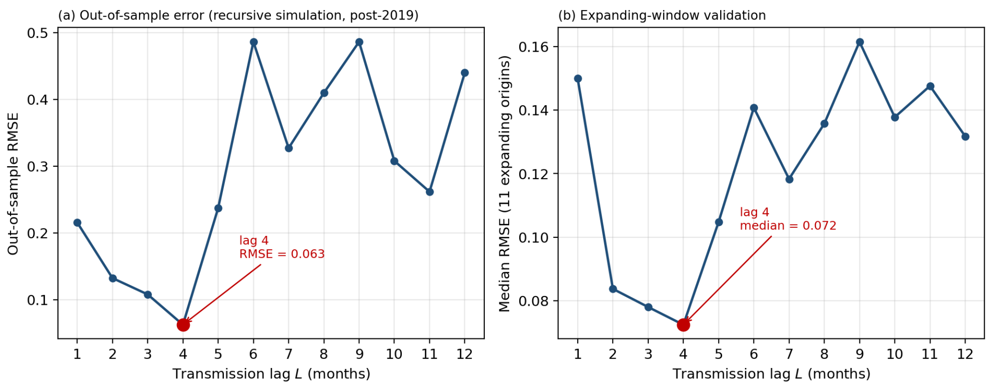

Figure 4.2: Selection of the transmission lag. **(a)** Out-of-sample RMSE of the recursively simulated rate (post-2019) for lags of one to twelve months. **(b)** Median RMSE across eleven expanding-window origins, confirming the four-month minimum. Source: author's computations.

The empirical optimum has a direct **institutional interpretation**: the asset–liability committee that sets the regulated rate meets quarterly and implements its decision the following month, three months of governance plus one of execution, a four-month delay. The lag retained by the data is the lag implied by the governance of the rate. Two further refinements were tested and discarded on the same grounds; the most instructive, a curve-slope factor is documented in Appendix B.3.

4.4 The Retained Model

We now state this retained model in full and report its estimates.

4.4.1 The Long-Run Equilibrium

The equilibrium relationship is Equation (4.2), now with its parameters fixed by the search of the previous section: the pillar p_t is the 60-month arithmetic moving average of the 5-year €STR swap rate and the short rate enters as the 3-month Euribor lagged four months. Writing $p_t \equiv p_t^{(5Y,60)}$ and $m_{t-4} \equiv \text{Euribor3M}_{t-4}$, the estimated equilibrium rate is

$$\hat{r}_t^* = \hat{\alpha} + \hat{\beta}_p\, p_t + \hat{\beta}_s\, \text{Euribor3M}_{t-4}.$$

The **replication pillar** p_t proxies the *funding value* of the deposit. Following the replicating-portfolio logic of Maes and Timmermans (2005), the blended yield of a rolling fixed-income ladder is the return a bank earns against its stable deposit base and hence the anchor of what it is willing to pay. The pillar is the slow-moving, structural component of the client rate. Its construction, a medium-maturity swap averaged over a window of comparable length, reflects a behavioural duration of the stable deposit base of roughly four to five years, consistent with both Belgian ALM practice and the five-year cap on assumed NMD maturity imposed by European Banking Authority (2022b).

The **lagged short rate** (3-month Euribor, lagged four months) captures the *monetary-transmission* channel: the current policy stance reaches the client rate with a delay that reflects the time the bank takes to reprice. The four-month lag is not a free parameter but is grounded both institutionally and empirically (Section 4.3.1). This term is what allows the model to track the 2023–2024 repricing and the subsequent 2025–2026 turning point, which a pillar-only specification cannot reproduce.

Estimated over the in-sample period (through 2019) with HAC standard errors, both market forces are strongly significant and carry the expected positive sign; the coefficients are reported in Table 4.2.

Taken together, the two terms give the specification an economic reading that a purely statistical fit would lack. Every element carries a meaning rather than a curve-fitting role. The model is therefore *structural* rather than descriptive: it does not reproduce the client rate by flexible parametrisation but reconstructs it from the two forces that economically determine it. This is what underpins the out-of-sample performance documented in Section 4.6: the model is expected to hold wherever its mechanisms operate, outside pathological episodes such as the 2008 confidence crisis, where the rate responds to other forces.

4.4.2 Short-Run Dynamics

Let $\text{dev}_{t-1} = r_{t-1} - \hat{r}_{t-1}^*$ denote the deviation of the observed rate from its estimated equilibrium in the previous period. The short-run dynamic combines the immediate transmission of equilibrium movements with a gradual correction of past deviations:

$$\Delta r_t = c + \rho\, \text{dev}_{t-1} + \gamma\, \Delta \hat{r}_t^* + \eta_t, \tag{4.3}$$

where γ measures the contemporaneous pass-through of a change in equilibrium, ρ the error-correction coefficient pulling the rate back towards target and η_t a disturbance. In estimation, all quantities in Equation (4.3) are computed from the observed series r_t. The estimates are reported in Table 4.2.

On symmetry. A natural question, raised by the menu-cost literature on individual deposit rates, is whether the correction speed differs between deviations above and below equilibrium: the celebrated *downward-flexible, upward-sticky* asymmetry of Driscoll and Judson (2013). We tested an asymmetric variant in which the error-correction term is split by the sign of the deviation,

$$\Delta r_t = c + \rho^+ \max(\mathrm{dev}_{t-1}, 0) + \rho^- \min(\mathrm{dev}_{t-1}, 0) + \gamma\, \Delta \hat{r}^*_t + \eta_t, \tag{4.4}$$

so that ρ^+ governs the correction of rates *above* equilibrium (pressure to cut) and ρ^- the correction of rates *below* equilibrium (pressure to raise). A Wald test of $H_0 : \rho^+ = \rho^-$ does not reject symmetry on the in-sample window ($p = 0.35$), nor on the full sample ($p = 0.82$). Because the binding floor in R2 could mechanically mask an asymmetry, we re-estimated the dynamics excluding that period under three distinct criteria (by dates, by restricting to observations above the floor and on the in-sample window); symmetry is rejected in none (p between 0.78 and 0.83). This result is consistent with what aggregate data can identify: as Section 2.2 discusses, the upward-stickiness asymmetry is a *micro*, menu-cost phenomenon that a market-wide average washes out. We therefore retain a single, symmetric ρ, a choice that is both parsimonious and faithful to the level at which the mechanism operates.

4.4.3 The Regulatory Floor

The regulated product cannot legally pay below 0.11%. Estimation needs no special treatment, the observed r_t already incorporates the floor, but simulation generates the rate recursively, so the constraint must be applied at each step. Denote by $\underline{r}_t$ the simulated (floored) rate and by $\tilde{r}_t$ the latent rate the dynamics would produce absent the constraint:

$$\tilde{r}_t = \underline{r}_{t-1} + c + \rho\left(\underline{r}_{t-1} - \hat{r}^*_{t-1}\right) + \gamma\, \Delta \hat{r}^*_t, \qquad \underline{r}_t = \max\left(0.11\%,\ \tilde{r}_t\right). \tag{4.5}$$

The error-correction term is computed on the *floored* rate $\underline{r}_{t-1}$, the rate that actually prevails, so the recursion is well defined: the floor sets each step's starting point, the max truncates only its outcome. The constraint binds throughout R2, reproducing the long 2015–2022 plateau. This places the model among *cashflow-based* NMD models, which, unlike pure replicating-portfolio models, accommodate non-linear pricing features such as rate floors (Zanders, 2025): the appropriate framework for a regulated product with a binding floor.

Table 4.2: Estimated coefficients of the client-rate model

	Parameter	Estimate	t-stat (HAC)
Long-run (Eq. 4.2)	α (constant)	0.104	2.2
	β_p (pillar)	0.239	9.7
	β_s (Euribor3M$_{t-4}$)	0.215	9.5
Short-run (Eq. 4.3)	γ (pass-through)	+0.645	3.71
	ρ (error correction)	−0.0667	−1.69
	R^2 (long-run)	0.964	

Note. Long-run relationship estimated over the in-sample period (through 2019) with HAC (Newey–West) standard errors. On the full sample the long-run coefficients are stable.

4.5 Results and Interpretation

4.5.1 Fit

Figure 4.3 overlaps the modelled rate on the observed series over the full sample, with the residual in the lower panel: the model tracks the secular decline of R1, the long floor of R2 and, most demanding, the 2023–2024 repricing and its 2025–2026 turning point. The fit is close and free of systematic deviation outside the 2008 confidence episode. The correlation between modelled and observed rates is 0.985 over the full sample and 0.988 out-of-sample, for an R^2 of about 0.97.

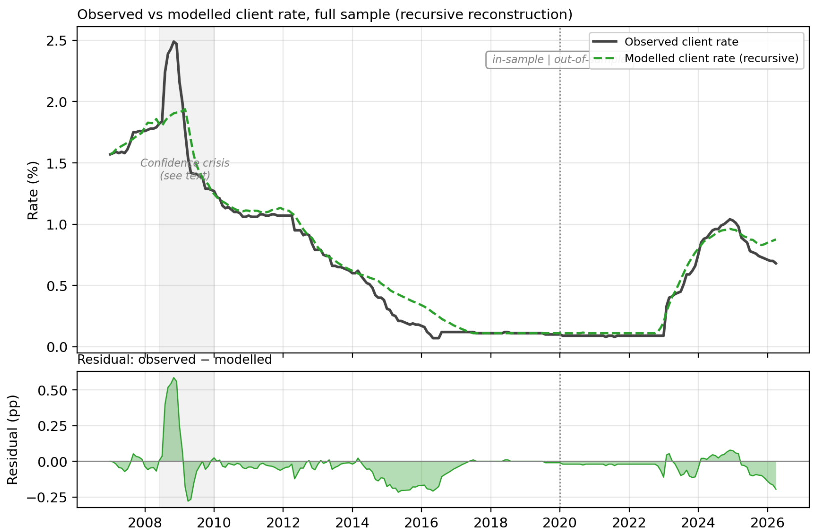

Figure 4.3: Observed versus modelled client rate, 2003–2026. The modelled rate is reconstructed recursively from Equations (4.2),(4.3),(4.5). Source: author's computations.

4.5.2 The Dynamics of Adjustment: Where the Stylised Facts Live

The estimates give quantitative content to the two behavioural features documented in the introduction (Figure 1.1). The **stickiness** of the client rate arises at two levels. It is partly structural: the funding anchor is a five-year moving average, so the equilibrium moves slowly by construction and the policy rate enters with a four-month delay. The estimates add a behavioural layer on top: only about two-thirds of any equilibrium change reaches the client rate within the month ($\gamma = 0.645$) and the pull back towards equilibrium is slow ($\rho = -0.0667$, a half-life of roughly ten months). The rate *follows but does not chase* and not merely because of the smoothed anchor: even relative to its own equilibrium, it adjusts with measured delay.

The **asymmetry** visible in the raw data (Figure 1.1, panel b) is accounted for rather than contradicted. At the aggregate level, once the binding floor and the error-correction structure are controlled for, no residual asymmetry in the adjustment speed survives (the Wald tests of Section 4.4.2); the descriptive asymmetry of the introduction reflects the floor episodes and the composition of the sample rather than an asymmetric correction mechanism. The micro-level asymmetry documented in the literature dissolves in the market-wide aggregate, a first instance of a theme that recurs throughout this thesis.

Two properties of this weak correction deserve to be stated plainly. The first is statistical: at $t = -1.69$, the error-correction coefficient is significant at the 10% but not the 5% level. Three considerations support retaining it: its sign is the one theory requires, the Johansen test rejects the absence of a cointegrating relation decisively (Section 4.6) and the weakness of ρ is itself the finding rather than a defect of the estimate, an administered rate that barely corrects towards its equilibrium is precisely what a sticky, regulated product should look like. The second property is practical. A strong ρ would make simulated paths highly sensitive to their starting point, producing a sharp mechanical catch-up whenever a simulation is launched far from equilibrium, an artefact of timing rather than economics. The weak ρ instead yields smooth trajectories, robust to the starting position, at the modest cost of a bounded tendency to under- or over-state the margin when a gap persists. For the use made of the model here, feeding a spread over a long sample, and, in prospective applications, over a projection horizon, this is a desirable trade-off (see Figure 5.4).

4.6 Validation

This section subjects the model to a full battery of validation, organised in two parts: the standard diagnostics required of an error-correction model and a set of robustness checks in the spirit of those applied to the volume model (Chapter 5). Table 4.3 summarises the battery; the subsections discuss what each result means and, where a test signals a departure from textbook conditions, why that departure is expected for an administered rate and how the estimation handles it.

Table 4.3: Validation battery of the client-rate model

Dimension	**Test**	**Result**	**Reading**
Cointegration (Eq. 4.2)	Engle–Granger	$p = 0.08$	borderline (conservative test)
	PSS bounds (F)	3.85	above the lower I(0) bound
	Johansen (trace)	39.8 > 29.8	**cointegration, 5%**
Residuals (η_t, Eq. 4.3)	Ljung–Box (12)	$p < 0.001$	autocorrelated (step-wise repricing)
	ARCH LM	$p < 0.001$	regime-dependent volatility
	Breusch–Pagan	$p = 0.14$	not linked to regressors
	Jarque–Bera	$p < 0.001$	non-normal (discrete jumps)
Stability	Expanding β_p	0.21–0.32	stable
	Expanding β_s	0.18–0.22	stable
	HAC t vs. truncation	$11.1 \to 9.7$	inference not lag-dependent
Robustness	Macro controls (6)	$\Delta R^2 \leq 0.008$	no omitted macro driver
	Coupling (Ch. 5)	τ: $1.35 \to 1.40$	fit for purpose

Note. Cointegration tests on the long-run relationship (4.2); residual diagnostics on the short-run disturbances of Equation (4.3); expanding-window estimates over 2015–2025. Symmetry of the short-run adjustment is tested in Section 4.4.2. Source: author's computations.

4.6.1 Order of Integration and Cointegration

The ECM framework requires the variables to be integrated of order one. Augmented Dickey–Fuller and KPSS tests confirm that the client rate and the lagged short rate are I(1). The pillar, being a 60-month moving average, is by construction highly persistent and its integration order is ambiguous, a feature we report transparently rather than obscure; it bears on the choice of cointegration test below.

Cointegration is assessed with three procedures, which we read jointly rather than in isolation. The single-equation Engle–Granger test does not reject the null of no cointegration ($p = 0.08$); this test is known to be conservative and is ill-suited to a regression mixing a smoothed moving-average regressor with a market rate. The Pesaran, Shin, and Smith (2001) bounds test yields $F = 3.85$, above the lower I(0) bound. The Johansen trace test, which treats the system multivariately and does not rest on a single residual series, rejects the null of no cointegrating relation decisively (trace statistic 39.8 against a 5% critical value of 29.8). The weight of evidence thus favours cointegration: Johansen rejects firmly, the bounds test leans the same way and theory demands a long-run tie between a deposit rate and its funding value. The conservative Engle–Granger result reflects the smoothed pillar and the low power of residual-based tests when ρ is small, precisely our case.

4.6.2 Residual Diagnostics

The residuals depart from the classical ideal: autocorrelated, heteroskedastic and non-normal (all $p < 0.001$), though with no heteroskedasticity linked to the regressors ($p = 0.14$), in an economically meaningful way. The source is the step-wise nature of an administered rate: in 44% of months the client rate does not move at all (30% even outside the floor period), a process of discrete jumps punctuating long plateaux that mechanically generates such residuals, independently of the explanatory variables. Attempts to whiten them fail and, worse, cannibalise the structural coefficients (Appendix B). We therefore retain the parsimonious specification and rely on HAC inference, consistent under exactly this combination of autocorrelation and heteroskedasticity; the non-normality affects only efficiency, not consistency. These residual properties are not a flaw but a manifestation of the very stickiness the model sets out to capture.

4.6.3 Robustness

The long-run coefficients are stable across expanding estimation windows: between 2015 and 2025 the pillar coefficient remains in the range 0.21–0.32 and the short-rate coefficient in 0.18–0.22, with the R^2 constant at 0.95–0.96. The coefficients estimated in-sample (through 2019) are confirmed on the full sample, including the unseen hiking cycle (Table 4.2, note). The HAC t-statistics are stable across the choice of lag truncation (from 11.1 at six lags to 9.7 at twelve for the pillar), so the inference does not hinge on an arbitrary Newey–West parameter.

Finally, we test whether the client rate is driven by macroeconomic conditions beyond the two market forces. We add each of a battery of controls: GDP, consumer confidence, the price index, inflation, unemployment, economic-policy uncertainty one at a time to the equilibrium relationship. None adds materially to the fit ($\Delta R^2 \leq 0.008$) and none destabilises the pillar or the short rate. The single exception, the **unemployment rate**, is statistically significant ($t = -3.7$) but economically negligible: it adds almost nothing to explanatory power ($\Delta R^2 = 0.008$) and lacks a direct channel to an administered pricing decision. We interpret it as a spurious trend correlation between two slowly declining series rather than a determinant of the rate and do not retain it. This confirms the economic premise of the model: the regulated rate is a pricing decision driven by the bank's funding costs and competitive position, not a response to the business cycle.

4.6.4 Coupling Coherence

A final validation, specific to this thesis, verifies that the modelled rate is fit for its purpose: feeding the spread that drives volume outflows. Substituting the modelled rate for the observed one in the outflow model of Chapter 5 leaves the mechanism essentially unchanged. The rate model can therefore substitute for the observed series without distorting the downstream mechanism; the full coupling validation, including the propagation of rate-model error into the predicted outflows, is conducted in Section 5.3.2.

Chapter 5

Modelling Cyclical Deposit Volumes

The client-rate model of Chapter 4 supplies one half of the coupling at the centre of this thesis; this chapter supplies the other. Its object is the *cyclical* component of deposit volumes, the portion of the balance that expands and contracts as the rate environment makes outside options more or less attractive. Its aim is to calibrate, empirically and for the Belgian regulated account, how much of that balance leaves and from what point.

The chapter proceeds from the isolation of the object to its model and validation. Section 5.1 separates the structural trend from the cyclical component, a statistical decomposition whose smoothing parameter is pinned down by two independent criteria and whose extracted cycle is then validated economically against the observed reallocation of household savings. Section 5.2 specifies the retained threshold model of outflows, flat below an attractiveness threshold, decoupling beyond it, including the evidence behind that functional form. Section 5.2.2 establishes the competing instrument that drives the spread. Section 5.3 validates the model, including the coupling with the modelled client rate of Chapter 4 and Section 5.4 benchmarks it against the regulatory standardised approach. Throughout, the retained choices are stated in the body; the alternatives, specifications tested and set aside, are documented in Appendix C.

5.1 Trend–Cycle Decomposition

5.1.1 The Structural and Cyclical Components

Recall from Section 2.4 that the deposit balance splits into a *structural* component and a *cyclical* component reflecting short-term arbitrage. For the observed volume series Vol_t (the L23 balances of Chapter 3),

$$\mathrm{Vol}_t = \mathrm{Trend}_t + \mathrm{Cycle}_t + \varepsilon_t, \tag{5.1}$$

where Trend_t is the slow-moving structural level, Cycle_t the cyclical component modelled here, and ε_t a residual. The structural component is the domain of the long-maturity replicating portfolio; this thesis models the cyclical component, which that machinery leaves generic. The trend is not itself the replicable core: because the cyclical component can fall *below* it, as in 2023, part of what the trend represents is itself liable to leave, and sizing the stable base is the distinct problem of the structural machinery (Section 2.3).

By *arbitrage* we mean, in a deliberately narrow and rate-driven sense, the reallocation of funds between the regulated account and a concretely available outside option, mainly the retail term deposit or a government note, when the return gap grows large enough to overcome the convenience and inertia that otherwise keep the balance in place. This is distinct from movements driven by liquidity needs or by confidence shocks; the latter, relevant to 2008, lie outside this scope (Section 5.3).

Neither component is observed; any split rests on assumptions. The most natural candidate, a trend tied to macroeconomic fundamentals, fails empirically: deposits and GDP are not cointegrated and the cyclical component a GDP anchor implies overstates the documented 2023 outflow threefold (Appendix C.1). The trend must therefore be extracted statistically, under two disciplines: the filter's free parameter must be anchored outside the model and the resulting cycle validated against an independent economic measurement, the safeguards of the next subsections.

5.1.2 *The Filter and Its Setting*

We extract the trend with the Hodrick–Prescott filter (Hodrick & Prescott, 1997), with a smoothing $\lambda = 100{,}000$. Since λ is the method's one degree of freedom, it is worth being precise about how little work that choice does. The resulting decomposition is shown in Figure C.1 (Appendix C.2): a smooth structural trend and a cyclical component whose arbitrage episodes, the 2008 crisis and the 2023 *bon d'État*, are what the threshold model of Section 5.2 sets out to explain.

Two independent criteria select the same region of the parameter space, as Table 5.1 summarises. The first is statistical and owes nothing to the Belgian data: for monthly observations, the frequency-adjustment rule of Ravn and Uhlig (2002) prescribes $\lambda = 1600 \times 3^4 = 129{,}600$, the default for monthly series in standard statistical software.[1] The second is economic: $\lambda = 100{,}000$ is the value at which the amplitude of the 2023 cyclical episode matches, in order of magnitude, the deposit outflow the National Bank documented for that summer, some €22 billion (Appendix C.2).

The convergence settles the question of circularity. Imposing the rule-based $\lambda = 129{,}600$ outright, a value blind to the Belgian episode, moves the 2023 trough by only €0.4 billion (Table 5.1). Had the episode never been observed, the frequency rule alone would have produced essentially the same decomposition: the economic anchoring *corroborates* the setting, it does not manufacture the result. We retain $\lambda = 100{,}000$ and report the rule-based value throughout as a robustness check.

Table 5.1: The smoothing parameter: two independent criteria

Criterion	λ	**2023 trough (€bn)**
Economic (episode-matched, retained)	100,000	−23.9
Statistical (Ravn–Uhlig frequency rule)	129,600	−24.3

Note. Both criteria place the trough in November 2023, €0.4 billion apart. The documented episode is described in Section 1.1. Source: author's computations.

[1] Notably `statsmodels` (Python) and Stata. The value 14,400 often used for monthly data corresponds to a squared, not a fourth-power, adjustment and is not the rule's recommendation.

The extracted trough sits within the documented range and at the right moment.[1] This is the first pillar of the decomposition's credibility: the cyclical component is not a statistical artefact vulnerable to the critique that filters manufacture spurious cycles; it is anchored to an independently documented event, by a setting that two independent criteria agree on. The second pillar, validation of the entire cyclical *path*, not just its trough, follows in Section 5.1.4.

5.1.3 Robustness and the End-of-Sample Caveat

The decomposition is robust to its assumptions, with one exception we quantify and then track. Varying λ across two orders of magnitude keeps the 2023 trough within €−19.0 to €−24.3 billion. Two rigid alternatives fail instructively: the filter of Hamilton (2018) overstates the trough by two-thirds (€−39.6 billion) and dates it late, while the parametric trends standard in the deposit literature (Bardenhewer, 2007; Kalkbrener & Willing, 2004) are rejected outright on these data: non-stationary residuals, implied troughs of €−41 to €−51 billion (Appendix C.3). Since the HP filter is a penalised spline whose trend tends to a line as $\lambda \to \infty$, these parametric trends are its rigid limiting cases: what the data reject is not the approach but its rigid limit, which is why a finite λ is required. The band-pass filter of Christiano and Fitzgerald (2003), by contrast, agrees with the HP cycle almost everywhere (correlation 0.94) and on the trough magnitude (€−24.0 billion) but dates it five months late; the HP filter, which dates the episode correctly, is therefore the primary decomposition and the band-pass its edge-robust control (Appendix C).

The exception is the end of the sample, where the HP cycle is measurably biased upward: over 2024–2026 it sits several billion euros above the edge-robust control, and a real-time re-estimation revises the final values downward as later data arrive (quantified in Appendix C.4). This has three consequences, which we flag now. First, the apparent cyclical rebound of 2024–2026 is overstated; we return to it in Section 5.3. Second, any estimate that leans on these boundary observations must be treated with caution, an argument that proves decisive for the functional-form choice of Section 5.2.3. Third, because the inflated points sit *above* the decoupling line, they flatten rather than steepen the outflow slope estimated below: the estimate therefore errs on the conservative side. Reassuringly, the threshold and slope hold regardless ($\tau = 1.35\%$, $\delta = -14.5$ on the band-pass cycle, against 1.35% and −14.7 on the HP cycle): the mechanism is a structural property of the data, and the filter governs only the dating and the boundary.

[1] The comparison is one of order of magnitude: the cyclical component measures a *deviation from trend*, whereas the National Bank's figures measure the *decline in outstanding balances* over a window. Because the trend was still rising through 2023, the swing of the cyclical component from its mid-2023 peak to its November trough, not the trough level itself, is the quantity comparable to the observed fall. Nothing below rests on the exact correspondence: the threshold and slope are stable across λ (Section 5.1.3), and it is the *path*, not the trough level, that Section 5.1.4 validates.

5.1.4 Economic Validation: The Cyclical Component as Observed Reallocation

A statistical filter identifies the cyclical component only up to its own assumptions. The decisive question is whether the extracted series corresponds to an economic *object*: money households actually moved. The 2023 episode anchors the trough; this subsection anchors the *path*, by measuring the destination of the funds directly. If the cyclical component of regulated deposits is genuinely the arbitrage layer of household savings, its mirror should be visible in the substitute instruments those savings flowed into. We construct that mirror from two independently observed families: household *term deposits* (the ECB BSI series for deposits with agreed maturity) and *retail government bonds* (the *bons d'État*, reconstructed from Belgian Debt Agency records).[1] Their sum, sign-inverted, is an *independent, flow-based measurement* of the money that left the regulated account for rate reasons, built without any reference to the HP filter, its λ, or the L23 series. The two series, one extracted statistically from L23 balances and the other measured from competing-instrument balances, move together over the whole sample, with a correlation of 0.914 (Figure 5.1). Over the region that matters, the cyclical component the filter extracts *is* the reallocation the substitutes record. The end of the sample confirms the *direction* of the 2024–2025 reflux. The first *bons d'État* matured and term deposits began to unwind, so money demonstrably flowed back while its HP-measured *amplitude* is inflated by the boundary bias of Section 5.1.3; we read the rebound qualitatively and lean on its magnitude nowhere.

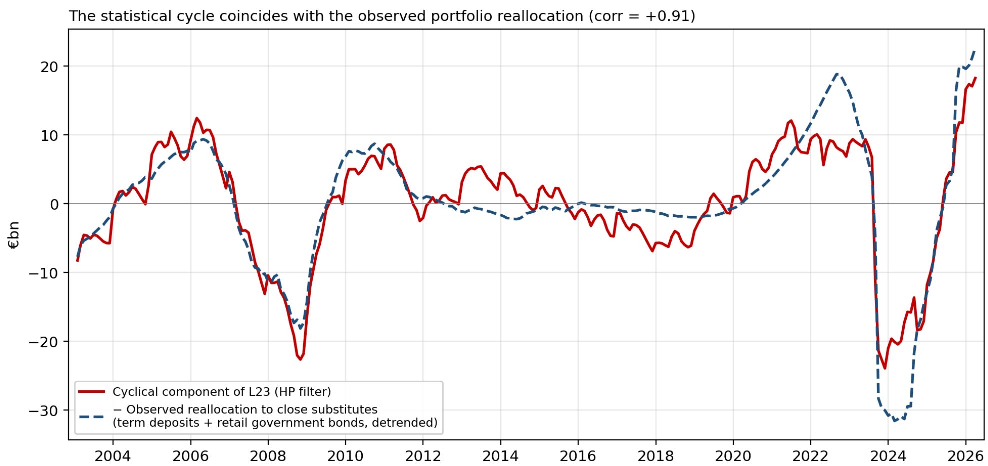

Figure 5.1: Economic validation of the cyclical component: the HP cyclical component of L23 volumes (solid) against the sign-inverted cyclical balance of substitute instruments (dashed), measured independently. Source: author's computations on ECB BSI and Belgian Debt Agency data.

Two checks sharpen the validation. First, re-estimating the retained threshold model on the reallocation series in place of the HP cycle recovers the same mechanism, and recovers it more sharply: a threshold of the same order (τ = 1.50% against 1.35%) and a decoupling that is, if anything, cleaner ($t = -12.0$ against -9.5; $R^2 = 0.64$ against 0.56). The differences are those the two measurements should exhibit: the reallocation series tracks the outstanding balances of the substitutes, which persist after the flows that created them, so its trough is

[1] Construction and sources in Appendix A.2. Term deposits enter as deviations from their own slow-moving trend, so both sides are cyclical; retail bond holdings are cumulated from gross subscriptions net of redemptions.

deeper and later; part of it, moreover, originates from liquid holdings other than the regulated account. Second, the check operates at the level of the *mechanism*: the same spread, through the same threshold, explains a measurement of the outflow constructed entirely outside the decomposition.

One temptation must be resisted: *substituting* the measured reallocation for the filtered cycle as the dependent variable. The reallocation series measures where money demonstrably *arrived*, within a perimeter of substitutes we chose; the object of this thesis is what *leaves* the regulated account, whatever its destination, including destinations outside any perimeter (current accounts pending investment, securities, cross-bank purchases). Conditioning the model on the destination measurement would redefine the modelled object by an ex-post classification choice, the very misidentification this thesis criticises in the standardised benchmark (Section 5.4). The reallocation series therefore serves as *validation*, not as calibration target: the HP cycle, externally corroborated, remains the primary series, and that the same spread, through the same threshold, explains both it and an independent measurement of the outflow is the strongest evidence that the mechanism is real rather than filtered into existence.

5.2 The Threshold Model of Outflows

5.2.1 The Threshold Form and Its Open Choices

The central empirical relationship of this thesis links the cyclical component of volumes to the spread between a competing instrument and the client rate. The relationship is not linear. Below a spread threshold, the deposit base is *inert*: depositors tolerate a moderate rate disadvantage, held by the convenience and inattention that underpin the deposit franchise (Section 2.1). Beyond the threshold, outflows accelerate sharply as the rate gap becomes large enough to overcome that inertia. The retained model is the direct translation of this behaviour, flat below the threshold, linear decoupling beyond it:

$$\mathrm{Cycle}_t = \alpha + \delta_{\mathrm{excess}} \max(0,\ \mathrm{spread}_{t-L} - \tau) + u_t, \tag{5.2}$$

where $\mathrm{spread}_{t-L} = c_{t-L} - r_{t-L}$ is a competing rate c_t minus the client rate r_t (the L23 series modelled in Chapter 4), entering with a reaction lag L; τ is the attractiveness threshold, estimated by grid search to minimise the residual sum of squares, the least-squares threshold estimator of Hansen (2000); α absorbs the mildly positive mean of the cyclical component outside arbitrage episodes; and u_t is a disturbance.

This piecewise formulation is the empirical counterpart of the behavioural component-based frameworks that structure modern NMDs modelling, which posit precisely such a non-linear outflow response to the market-to-deposit spread but typically leave its parameters, the threshold and the slope, generic and uncalibrated. Our contribution is to give these parameters empirical, Belgian-specific content over a sample that, uniquely, now spans a full hiking cycle. The threshold form also echoes a recent and independent strand documenting threshold effects in deposit behaviour: You and Noh (2025), on Korean institution-level data, find that deposit-pricing responses reverse beyond identifiable stress thresholds. It is a different variable from our volume threshold, but it reflects the same structural intuition: that depositor behaviour is regime-dependent around a tipping point.

The form has, finally, a structural reading that matters for where this research leads. Equation (5.2) is the reaction function of a depositor with a single *strike*; as anticipated in Section 2.1, a heterogeneous deposit base is a portfolio of such reactions, and the flat-below-threshold form retained here is precisely the one that nests into their sum, where a below-threshold slope would not. Beyond the linear decoupling segment, the response is moreover unlikely to remain linear indefinitely: as in mortgage-prepayment modelling, it should follow an S-curve that eventually saturates (*burnout*) as readily movable balances are depleted and only precautionary money remains. The single-kink model retained here is the segment of that richer structure identifiable on aggregate data with two arbitrage episodes; the bucketed, S-shaped extension is developed among the perspectives of Chapter 6.

Three choices are left open by this form and are settled empirically in what follows: the *competing instrument* c_t against which the spread is measured (Section 5.2.2), the *functional form* below the threshold (Section 5.2.3), and the *reaction lag* L (Section 5.2.4). Section 5.2.5 then states the retained model in full and reports its estimates.

5.2.2 The Competing Instrument

The first choice is the rate c_t representing the depositor's outside option, and it is an economic decision we make explicitly. Six candidates are compared, each over the full grid of spread lags from one to nine months with τ re-estimated in every cell, since a rate that reaches the depositor slowly need not share the reaction delay of one that reaches him fast. Table 5.2 reports each at its own optimum.

Table 5.2: Candidate competing instruments, each at its own optimal spread lag

Candidate	**Lag**	τ	δ_{excess}	t	R^2
Interior optimum					
Euribor 3M	5	1.70	−15.71	−15.3	0.62
Euribor 6M	6	1.75	−14.81	−14.8	0.60
Bon de caisse 3M	6	1.30	−14.17	−8.5	0.57
Boundary optimum					
Swap 1Y	9	1.30	−15.67	−10.5	0.57
OLO 1Y	9	1.40	−14.56	−10.4	0.57
Bon de caisse 12M	9	1.40	−15.03	−10.9	0.56

Note. Grid of spread lags 1–9 months, τ re-optimised in every cell; the reported row is each candidate's best lag. HAC (Newey–West) standard errors; δ_{excess} significant at the 1% level throughout. Source: author's computations.

Two features matter. First, only the short-dated instruments reach a maximum inside the grid, a genuine reaction delay of five to six months. The three longer-dated candidates peak at nine months simply because their fit is still rising there: phase alignment between two slow-moving series, not a behavioural response, and they are set aside on that ground rather than on their R^2. Second, and more consequentially, **the decoupling coefficient is essentially invariant across all six candidates** (-14.2 to -15.7 billion euros per point of spread). The threshold moves, the lag moves, the fit moves; the force of the reaction does not. The choice of instrument governs where arbitrage *begins*, not how hard it bites.

The retained driver is the three-month *bon de caisse*. The instrument that fits better, the three-month Euribor (0.62 against 0.57), is set aside on economic grounds: the Euribor is an *interbank* rate, not the one the representative depositor watches. The *bon de caisse* is a concrete product, distributed by the depositor's own bank and subscribed directly as the salient alternative to the savings account. The reallocation evidence of Section 5.1.4 confirms it, since the destinations actually observed are term deposits and retail government paper. We accept a modestly lower fit in exchange for fidelity to the mechanism.

The proximity of the two fits is itself telling. Banks price their retail term-deposit books against the swap curve and the three-month Euribor, so the *bon de caisse* rate is, to a first approximation, the Euribor transmitted to the retail counter after margin. This is why the wholesale rate tracks the arbitrage signal almost as well, and why the threshold estimated against it is mechanically higher: 1.70–1.75% for the two Euribor maturities against 1.30–1.40% for every retail instrument. That systematic gap of some forty basis points *is* the retail margin, which the depositor never sees but which the wholesale spread embeds. As with the asymmetry of the client rate (Section 4.4.2), the aggregate wholesale reference displaces a mechanism that operates, at root, through the concrete product; on individual-bank data, the product-level driver would sharpen its dominance.

5.2.3 The Functional Form: Why Flat Below the Threshold

A natural enrichment of Equation (5.2) is a slope below the threshold, a term $\beta_{\text{base}}\,\text{spread}_{t-L}$ letting the base respond, mildly, before arbitrage sets in. We estimated this variant, Figure 5.2 shows both fits against the monthly observations. The variant does produce a gently rising segment below τ ($\beta_{\text{base}} = +2.7$, $t = 1.8$, significant at the 10% but not the 5% level) and a marginally better in-sample fit. We nonetheless retain the flat specification, for a reason the figure makes visible.

The below-threshold slope is carried by the points highlighted in Figure 5.2: the end-of-sample observations (≥2024), precisely the region where Section 5.1.3 measured the HP cycle to be inflated upward. These boundary points sit high on both sides of the threshold: below τ, they manufacture the rising slope; beyond it, they sit above the decoupling line and attenuate its steepness. Excluding them, the slope loses its significance ($t = 1.1$) and the fit advantage of the richer form collapses; and no behavioural interpretation of the slope survives scrutiny:

the candidate readings (a remuneration-level effect on saving, a rate-momentum effect, a market-level effect) were each tested against the spread and rejected, a winnowing documented in full in Appendix C. A coefficient whose significance rests on the least reliable observations of the sample and which resists economic interpretation is not a mechanism; it is an artefact of the boundary. The flat form, by contrast, states exactly what the data robustly support and is the form that nests into the multi-strike extension outlined above. One implication bears repeating: since the same inflated boundary points also flatten the decoupling segment, the slope δ_{excess} retained below is conservative, the true decoupling is, if anything, steeper.

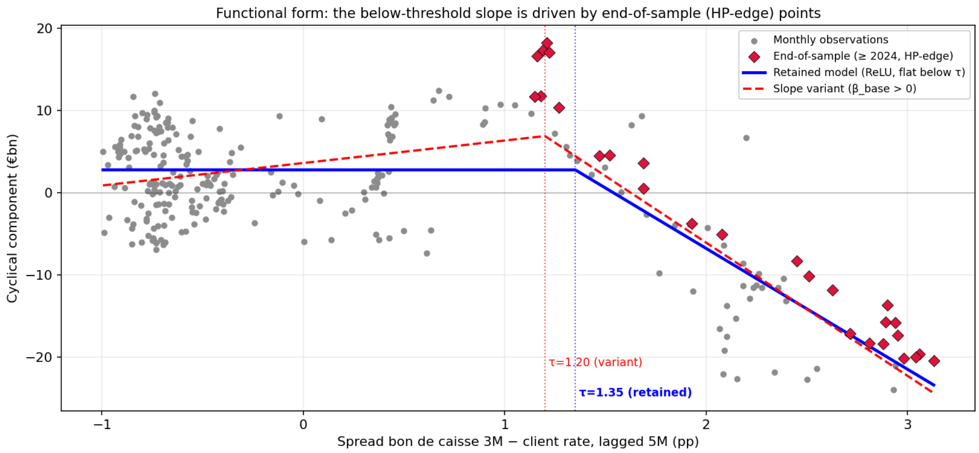

Figure 5.2: The functional form of the outflow relationship: cyclical component against the lagged spread, with the retained flat-below-threshold model (solid) and the below-threshold-slope variant (dashed). Highlighted points are the end-of-sample observations (≥2024). Source: author's computations.

5.2.4 *The Reaction Lag*

The spread enters with a five-month lag, representing the time depositors take to perceive a widening gap, decide, and act [1]. This delay is consistent with the documented inertia and inattention of depositors (Lu & Wu, 2026) and is where the data place the optimum: the explanatory power of Equation (5.2) rises steadily with the lag to a plateau over four to six months (R^2 of 0.54–0.57) and deteriorates markedly beyond it, with the threshold and slope stable across the plateau (τ = 1.30–1.40%, δ between −14 and −15). Differences within the plateau are immaterial; we retain five months, the centre of the plateau and the optimum under the alternative functional form as well, so that the selection is robust to the choice of form. The lag profile is shown in Figure 5.3. The choice also respects an internal-consistency condition with the rate model: the depositor reaction lag (five months) exceeds the bank's repricing lag (four months, Section 4.3.1). This ordering is economically necessary, were depositors quicker to flee than the bank to reprice, the bank would face outflows before its own adjustment and would be forced to react faster. That the bank can afford a four-month delay reveals that depositors react no faster than that.

[1] Table 5.2 reports the *bon de caisse* 3M at its grid optimum of six months; we retain five, the centre of the plateau (Section 5.2.4).

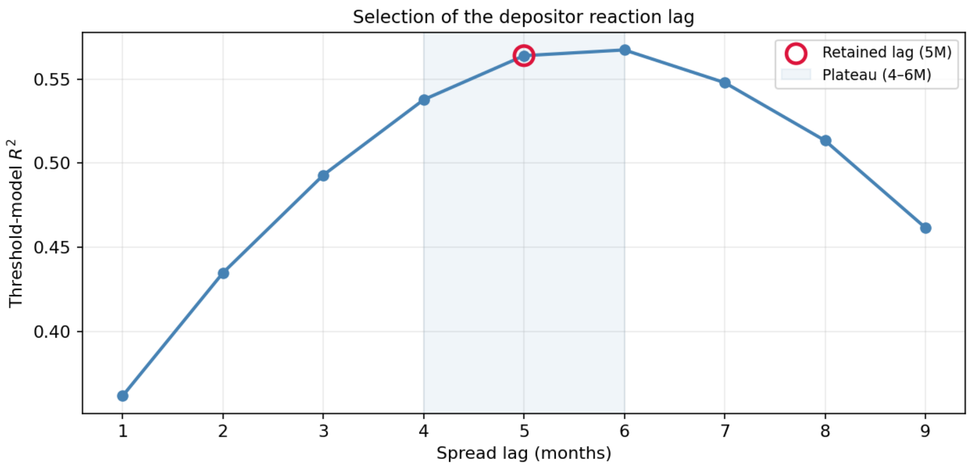

Figure 5.3: Selection of the depositor reaction lag: threshold-model R^2 as a function of the spread lag. Source: author's computations.

5.2.5 *The Retained Model and Its Estimates*

The three choices are now settled: the spread is measured against the three-month *bon de caisse*, the response is flat below the threshold, and the lag is five months. Estimated on the observed client rate with HAC standard errors, the model yields[1]

$$\tau = 1.35\%, \qquad \alpha = +2.79, \qquad \delta_{\text{excess}} = -14.71^{***},$$

with $R^2 = 0.564$[2]. The parameters tell the behavioural story in three numbers. The threshold $\tau = 1.35\%$ is the empirical measure of the Belgian depositor's tolerance: the rate disadvantage that can be sustained before arbitrage sets in. Below it, the base is flat at a mildly positive level ($\alpha = +2.8$ billion), the counterpart of the negative arbitrage episodes in a zero-mean cyclical series. Beyond it, each additional point of spread is associated with roughly €15 billion of cyclical outflow. At the 2023 *bon d'État*, the spread reached well above the threshold, placing the episode squarely in the decoupling regime and accounting for the observed outflow.

5.3 Validation

5.3.1 *Specification and Residual Diagnostics*

The volume model is subjected to the same scrutiny as the client-rate model of Chapter 4. Table 5.3 reports the battery, organised along the same four dimensions as Table 4.3: specification, residual diagnostics, parameter stability and robustness.

The residuals are well behaved. Stationarity is not rejected by the KPSS test ($p = 0.09$), and the augmented Dickey–Fuller test confirms it once the end-of-sample observations, inflated by the boundary bias of Section 5.1.3, are removed ($p = 0.03$). Unlike those of the

[1] Significance: $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$, on HAC (Newey–West) standard errors; δ_{excess} has $t = -9.5$.

[2] At the retained five-month lag.

rate model, the residuals are also normal ($p = 0.61$), the arbitrage episodes being few and large rather than heavy-tailed. As in the rate model they are autocorrelated and episodically heteroskedastic, which is why all inference is conducted with HAC standard errors. The threshold and the decoupling coefficient are stable across sample truncations that contain the identifying episode ($\tau = 1.30$–1.45, δ between -14.1 and -14.7) and across the choice of decomposition. Truncating before 2023 removes the very episode that identifies δ, and is therefore a test of identification rather than of stability, a distinction the model makes visible rather than hides.

Table 5.3: Validation battery of the volume threshold model

Dimension	**Test**	**Result**	**Reading**
Specification (Eq. 5.2)	Observations	$n = 274$	56 above threshold, 218 below
	KPSS on u_t	$p = 0.09$	**stationarity not rejected**
	ADF on u_t (≤2024)	$p = 0.03$	stationary once edge removed
Residuals (u_t)	Ljung–Box (12)	$p < 0.001$	autocorrelated (HAC used)
	ARCH LM	$p < 0.001$	episodic volatility
	Jarque–Bera	$p = 0.61$	**normal**
Stability	τ vs. truncation	1.30–1.45	threshold stable
	δ vs. truncation	-14.1–-14.7	stable once episode included
Robustness	HP vs. band-pass	$\tau = 1.35/1.35$	decomposition-invariant
	Coupling (Ch. 4)	τ: $1.35 \rightarrow 1.40$	fit for purpose

Note. Threshold model of Equation (5.2), $n = 274$. Residual stationarity: KPSS does not reject stationarity on the full sample; the augmented Dickey–Fuller test, weak against the near-unit-root drift induced by the end-of-sample HP bias (Section 5.1.3), does not reject a unit root on the full sample ($p = 0.13$) but does once the inflated 2025–2026 observations are removed ($p = 0.03$). Truncations report the model re-estimated on samples ending in 2024 and on the full sample; ending before 2023 removes the arbitrage episode that identifies δ. HAC (Newey–West) standard errors. Source: author's computations.

5.3.2 Model Fit and Coupling Coherence

Figure 5.4 presents the central evidence of this section: the observed cyclical component against the model's prediction, fed first with the observed client rate and then with the modelled one, over the full sample (a) and the zoomed 2022–2026 episode(b).

Three readings follow. First, the *fit*: the model reproduces the observed cyclical pattern, capturing the 2023 decoupling in both timing and depth, a predicted trough of €−23.4 billion against the cyclical trough of €−23.9, while remaining flat and close to the observed series through the calmer regimes. Second, the *coupling*: re-calibrating the model on the spread computed from the modelled rate yields a near-identical mechanism. The threshold shifts from 1.35% to 1.40%, one step of the estimation grid, and the decoupling slope from -14.7 to -14.9 billion euros per point of spread. The modelled rate triggers the outflow at the right point and with the right force, the two predicted series are visually near-indistinguishable (average discrepancy €0.55 billion). Third, the *propagation of rate-model error*, visible on Figure 5.4 at the peak of the episode: the chain fed with the modelled rate predicts a trough of €−21.8 billion against €−23.4 through the observed rate, a gap of €1.6 billion, about 7% of the episode. The direction is systematic: at the height of the tightening the modelled rate repriced marginally faster than the observed one, compressing the spread and hence the predicted outflow. Rate-model error thus attenuates, and never amplifies, the predicted outflow; a conservative propagation. This is the central result of the chapter: the outflows of this deposit base can be recovered from the rate environment alone, without observing what the bank actually paid.

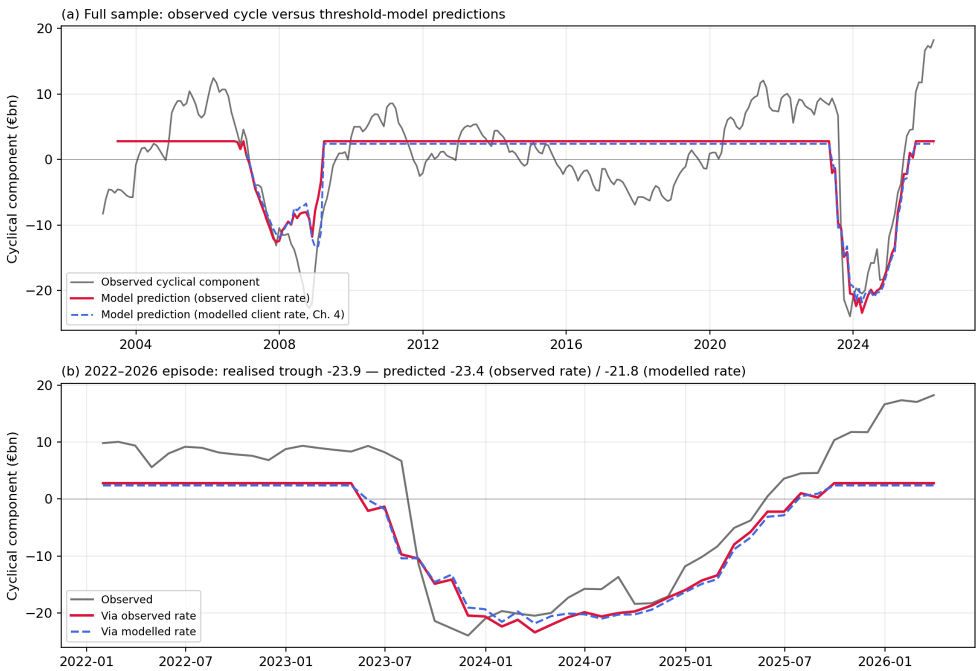

Figure 5.4: Model fit and coupling coherence, (a) full sample and (b) 2022–2026 episode. The observed cyclical component of L23 volumes against the threshold-model prediction, fed with the observed client rate and with the modelled client rate of Chapter 4. Source: author's computations.

5.3.3 *Fidelity Below the Threshold*

A specific concern deserves its own verification: is the model faithful in the calm regime? By construction, the retained form predicts a constant ($\alpha = +2.8$ billion) whenever the spread is below the threshold; the question is whether the observed cyclical component in that regime is indeed flat and mildly positive, or whether the flatness suppresses genuine dynamics. Over the 218 months concerned, the observed cyclical component averages €+2.8 billion (median +2.0), matching the model's constant almost exactly, with a mean absolute deviation of €4.5 billion around it. What the flat form foregoes is the slow drift of the below-threshold mean across sub-periods (mildly higher in the 2000s, mildly negative in the floor years); as established in Section 5.2.3, no tested economic mechanism accounts for that drift, and its recent portion is a boundary artefact. Below the threshold the model is, as it should be, quiet and approximately right. This is the model's declared scope. It is built to detect and size the arbitrage episodes that move the deposit base by tens of billions, not to track the residual variation of a calm regime which, as Section 5.2.3 establishes, no rate-driven mechanism accounts for.

5.3.4 A Limit of Scope: Confidence Shocks and the Credit-Risk Channel

The only material discrepancies between the two volume predictions concentrate in the 2008–2009 financial crisis, where they reach up to €8 billion against a full-sample mean of €0.6 billion. This is not a failure of the coupling but a limit of the model's scope, and the distinction is diagnostic: at the very months when the modelled and observed rates diverge most (October–November 2008), the volume discrepancy is negligible; the large 2008 gap arises where the *outflow model itself*, calibrated on rate arbitrage, cannot capture a systemic confidence crisis. In 2008, deposits responded to a flight to safety driven by fear and not by the spread.

That fear is, in principle, measurable, so we tested it rather than merely acknowledging it. With the deposit guarantee capped at €100,000, holders of larger balances had an incentive to reduce exposure to any single stressed institution regardless of its rate, a credit-risk channel. Using five-year senior CDS spreads of the major Belgian banks, we examined both directions in which bank credit risk could enter the framework, on the rate side and on the volume side; the full battery is in Appendix C.6. The verdict is consistent: the channel carries the expected sign on the volume side but cannot be robustly identified on aggregate data, and the arbitrage mechanism is untouched in every variant. The reason is by now familiar: the guarantee-threshold mechanism is a *large-balance, bank-specific* phenomenon, operating in the cross-section of depositors and institutions, precisely the dimension a market-wide aggregate averages away, as with the asymmetry of the client rate (Section 4.4.2) and the choice of competing instrument (Section 5.2.2). We document the channel without retaining it: 2008–2009 lies outside the model's declared scope, an acknowledged limit without bearing on the 2022–2024 period of interest, where the coupling is faithful to within €2.1 billion.

5.4 What the Model Adds to the Standardised Approach

The reference against which a bank's own model is judged is the Basel/EBA standardised treatment, the prudential fallback the regulator prescribes and the benchmark any internal model must improve upon. It splits the balance into a *core* and a *non-core* component. For retail non-transactional deposits, the category to which the regulated savings account belongs under Basel Committee on Banking Supervision (2016, §111), the core share is capped and the non-core portion scaled by a supervisory factor under a rising-rate scenario (European Banking Authority, 2022a).[1] Applied to the deposit base of Section 3.1, with the non-core share at the standardised 30% scaled by 1.3, this prescribes an outflow capacity of roughly €121 billion.

[1]We follow the standardised-approach caps of Basel Committee on Banking Supervision (2016, §110–115 and Table 2) as transposed by European Banking Authority (2022a); the precise treatment of regulated deposits may be refined in national implementation, so the capacity figure below should be read as an order of magnitude rather than an exact supervisory number.

The decomposition is ours under another name: core and non-core are structural and cyclical. But the standardised approach holds the boundary between them *fixed*, a non-core share set once for every month and every state of the world; the mechanism estimated here shows it is not. Beyond the threshold, part of what a fixed rule counts as core in fact leaves, so the frontier moves with the spread. This thesis does not, for that reason, deliver the split itself: as Section 5.1 noted, the core/non-core partition must be fixed *a priori* to define what a replicating portfolio hedges, and calibrating it is the distinct problem of the structural machinery. What the model adds is a reading of the state that partition is set in: when the split is next revised, it locates the deposit base relative to its threshold, dormant, approaching, or already beyond it, so that the *a priori* choice can reflect the arbitrage pressure of the moment rather than be made blind to it. The role is one of monitoring, not recalibration: the model does not reset the boundary, it tells a bank where that boundary currently stands.

What the estimated model then shows is that realised outflows sit far inside the prescribed bound. On the September 2023 *bon d'État*, the sharpest deposit episode in the history of the series, the model predicts €−23.4 billion against an extracted cyclical trough of €−23.9 billion, where the standardised approach signals a capacity several times larger and does so in every month alike. Over the 56 months in which the spread exceeds τ, the model tracks realised outflows with a mean absolute error of €4.2 billion; over the remaining 218 months, in which the deposit base was in fact mildly growing, it stays flat and neutral. Even the 2008 confidence crisis, outside the model's rate-arbitrage scope, moved the cyclical component by a fraction of the prescribed capacity. A permanent non-core provision, whatever its level, is not a conservative reading of this deposit base; it is one the data never approach, in twenty-three years that include two arbitrage episodes and a global financial crisis. This does not argue for a bank holding less than the regulator requires: the cap is a prudential floor, and a treasurer is free to plan against a figure more conservative than the model's. What the model adds is not a lower number but a **dated one**. The spread is observable, and deducible from market rates alone through the client-rate model of Chapter 4; the estimated threshold marks the level at which outflows historically begin; and the five-month reaction lag turns a crossing into a warning with a horizon. Beyond the threshold, the slope gives the order of magnitude, some €15 billion per additional point of spread. The standardised capacity carries none of this: it is the same number when the spread is deeply negative and when it is about to breach the threshold.

The 2023 episode is the illustration. The prescribed capacity was identical in January and in December; the spread crossed τ in the spring, and the outflow arrived in the autumn. A fixed capacity tells a bank how much of its base is conventionally at risk; it cannot tell it that the risk is about to materialise. That interval, between the signal and the outflow, is time enough to hedge, to arrange replacement funding, or to reprice before rather than after. In 2023 it was not used: the deposits left, the interest margin compressed, and the funding was replaced on the market at market prices (Cappoen, 2024). The model does not tell a bank to hold less; it tells it when the buffer is about to be drawn.

Chapter 6

Conclusion

6.1 What Was Asked, and What Was Found

This thesis asked how the client rate and the cyclical volume dynamics of the Belgian regulated savings account can be modelled and coupled, so that deposit outflows can be deduced from the interest-rate environment. The question was posed by an episode: in 2023, the gap between market rates and the rate paid on regulated savings, left open by the slowest pass-through over two decades, was closed by the State rather than by the banks; €17.8 billion left the accounts in a single month. Models calibrated over two decades of falling rates had nothing to say about it.

The answer developed here is a chain of two deliberately simple models. The client rate is described by an error-correction specification anchored on a funding-value pillar, a five-year swap rate averaged over its trailing five years, and a lagged short rate, with the regulatory floor imposed explicitly; it tracks two decades of administered pricing, through three structurally distinct regimes, with a mean error of a few basis points. Deposit volumes are split into a structural trend and a cyclical component, and the cyclical component obeys a threshold rule: inert while the spread between the retail alternative and the client rate remains below approximately 1.35 percentage points, and beyond that threshold losing roughly €15 billion per further point of spread, with a five-month delay. Chained together and fed with market rates alone, never shown the rate the banks actually paid, the two models reproduce the 2023 outflow to within €2.1 billion of the realised figure, with an error whose direction is systematically conservative. Benchmarked against the regulatory standardised approach, the contrast is stark: where the supervisory method signals a fixed outflow capacity of some €121 billion in every month of the sample, the threshold model is silent in the 80% of months where nothing is at stake and accurate in the months where something is.

The thesis is equally explicit about what it discarded, and why. A trend anchored in GDP, the parametric trends standard in the deposit-modelling literature, a below-threshold slope, a bank-credit-risk channel and a policy-uncertainty channel were each tested and set aside, on documented evidence rather than convenience (Appendices B and C). What remains is not the largest model the data could carry but the smallest one they robustly support.

6.2 What the Model Establishes—A Personal Assessment

Beyond the estimates themselves, I draw three conclusions from this work, and state them as my own reading of the evidence.

The first is that the **primary signals of Belgian deposit behaviour are now measured, and they are strong**. A tolerance threshold near 1.35 percentage points; a decoupling of the order of €15 billion per point of spread beyond it; a depositor reaction lag of five months, consistently longer than the banks' four-month repricing lag; a pass-through so weak that the error-correction pull is barely significant, each of these is a behavioural constant of the Belgian retail market that, to my knowledge, had not been estimated on a sample containing a genuine tightening cycle. They are not artefacts of a filtering choice: the threshold and the slope are stable across decomposition methods, smoothing parameters, sample truncations and functional forms, and the cyclical component they act on is corroborated by an independent, flow-based measurement of where the money actually went. For the big movements, the question that matters to a treasurer or a supervisor, the model does the job: it predicts the large volume swings of the Belgian aggregate from the rate environment alone. That these orders of magnitude, the tolerance threshold, the force of the reaction, the ordering of the lags, are recognised as plausible by balance-sheet practitioners reinforces the reading that they are genuine features of the market rather than artefacts of estimation.

The second is that **aggregate data have been taken to their ceiling**. This is, in my assessment, the honest reading of the recurring pattern documented throughout the thesis. The asymmetry of deposit-rate setting, well established at the institution level, dissolves in the market-wide average. The credit-risk channel of 2008, economically undeniable, carries the right sign but cannot be identified once balances above and below the guarantee ceiling are summed. The competing instrument that drives arbitrage is a concrete retail product, yet the aggregate barely distinguishes it from the wholesale rate it is priced from. The outflow response itself, which theory describes as a distribution of individual exercise decisions, is identifiable in the aggregate only as a single kink. Four times, the same lesson: the aggregate reveals the mechanism and absorbs its structure. I do not regard this as a failure of the modelling: it is a property of the data, one this thesis has, I believe, mapped rather than suffered.

The third follows from the first two: **the right way to read this model is as a building block**. It is not, in its current form, an internal model a bank could deploy: a single representative threshold cannot price the heterogeneity that risk management ultimately cares about. But that was never what aggregate data could deliver. What they could deliver, and what this thesis set out to extract, is the demonstration that the mechanism is real, that its parameters are measurable, that the chain from rates to outflows closes, and that the signals are strong enough to build on. The primary signals are there; the task ahead is to refine them on granular data and to recover the ones the aggregate absorbed.

6.3 Limits

Four limits bound these conclusions, and I state them plainly.

First, the decoupling coefficient is identified essentially by **one great episode**. The threshold is corroborated across methods and truncations, but the sample contains a single full *rate-driven* arbitrage cycle; the constancy of δ_{excess} across future episodes: under different fiscal treatments, competing products, or depositor attentiveness is an assumption and not a finding. Second, inference on the decoupling coefficient is **conditional on the estimated threshold**: the grid search that locates τ is not reflected in the reported standard errors, a known feature of threshold models, in which τ is unidentified under the null of no threshold effect and must be treated as an estimated nuisance parameter (Hansen, 1996, 2000). At $t = -9.5$ the conclusion would survive any reasonable correction, but the formally rigorous treatment, bootstrap inference in the spirit of Hansen (1996), is left to future work. Third, the model is a **conditional, reduced-form** description: in 2023 banks also responded to the outflows, raising fidelity premia to stem the outflow, so the client rate is not strictly exogenous to volumes; the five-month reaction lag mitigates, but does not eliminate, this simultaneity. A related caveat concerns what an L23 outflow represents. Part of what leaves the regulated account does not leave the bank: it migrates to the institution's own term deposits, whose rates banks raised precisely over the 2023 episode. The chain therefore measures the outflow from the L23, not a one-for-one loss of funding. For the purpose of this thesis this is less limiting than it appears: from an asset–liability standpoint the cost of the outflow is substantially the same whichever way the money moves, since it is the assumed duration of the L23 that breaks, forcing a rehedging of the position regardless of the destination. What the account gathers back in new term deposits is the only offsetting gain. Fourth, the linear decoupling segment is estimated on observed spreads of up to roughly three points; **beyond the observed range it should not be extrapolated**, for a reason the next section makes structural: the true response must eventually saturate.

6.4 The Trajectory: From One Strike to Many

The extension this thesis points to is not a list of refinements but a single programme, and the option view of Section 2.4 states it compactly. The depositor holds an option to move funds; the estimated threshold is its strike; and a deposit base is a portfolio of such options with dispersed strikes. On account-level data, the specification retained here, estimated segment by segment on buckets formed by reactivity, balance size or channel, becomes a multi-strike outflow model: a portfolio of threshold reactions whose aggregate reproduces, and refines, the single kink measured in this thesis. The flat form was chosen with that extension in mind, since it nests into such a sum where a below-threshold slope would not. The analogy with the asset side is exact: mortgage prepayment is modelled as a response to the refinancing incentive, aggregated over borrowers whose thresholds differ, and it carries a further implication for the shape of the response.

A portfolio of dispersed strikes implies, on its own, an *accelerating* aggregate response: as the spread widens, successive strikes are crossed and the marginal outflow steepens. Economic reasoning implies the opposite curvature far beyond the threshold: outflows draw first on excess, readily movable savings, and as the spread keeps widening the remaining balances are increasingly precautionary, so the marginal response must eventually *saturate*, the counterpart of *burnout* in prepayment modelling. The two forces together trace an S-shaped response, steepening past the threshold and flattening at high incentive, of which the linear decoupling segment estimated here is the central, identifiable portion. With account-level data, a logistic specification would locate both the acceleration at the front of the curve and the saturation at its back and would, incidentally, resolve the extrapolation limit stated above.

Granular data would also allow the recovery of the **signals the aggregate absorbed**, each documented in this thesis at the point where aggregation erased it: the asymmetric, menu-cost-driven rate setting of individual institutions; the guarantee-threshold credit-risk channel, which operates in the cross-section of balances above and below €100,000; and the product-level competing instrument, whose dominance over wholesale references should sharpen at the level of a single bank's shelf. Alongside this central programme, three further extensions are worth naming: probabilistic regime identification (Markov-switching) in place of fixed break dates; a smoothing parameter conditioned on the rate environment; and the prospective application of the validated chain, feeding simulated rate scenarios through it to project outflows and their net-interest-income cost, the natural operational use once the behavioural parameters are trusted.

6.5 Closing Remark

The 2023 episode showed what it costs to be surprised by depositors: a compressed margin and a forced return to market funding, at market prices, for money that had quietly walked out (Cappoen, 2024). The prudential answer to that risk is a bound that holds in every state of the world, and it did hold. What it could not say was when the state would change. This thesis has tried to show that, even from public, aggregate data, that second question has an answer: measure the tolerance of depositors, the force of their reaction and the delay of their attention, chain these to the rate environment, and the outflow becomes something a bank can see coming rather than merely provision against. The signals are real, they are Belgian, and they are now on the table. Building the granular model they call for is, in my view, the most valuable next step, for the institution that holds the accounts and for the supervisor who must decide how much prudence the fog still justifies.

References

Bai, J., & Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, *18*(1), 1–22.

Bardenhewer, M. M. (2007). Modeling non-maturing products. In L. Matz & P. Neu (Eds.), *Liquidity risk measurement and management: A practitioner's guide to global best practices* (pp. 220–256). Singapore: John Wiley & Sons.

Basel Committee on Banking Supervision. (2016). *Interest rate risk in the banking book* (Standards No. d368). Basel, Switzerland: Bank for International Settlements.

Blöchlinger, A. (2015). Identifying, valuing and hedging of embedded options in non-maturity deposits. *Journal of Banking and Finance*, *50*, 34–51.

Cappoen, S. (2024, November). *Où sont passés les 22 milliards d'euros qu'a rapportés le bon d'état?* NBB Blog. Retrieved from https://www.nbb.be/fr/blog/ou-sont-passes-les-22-milliards-deuros-qua-rapportes-le-bon-detat (Accessed July 2026)

Christiano, L. J., & Fitzgerald, T. J. (2003). The band pass filter. *International Economic Review*, *44*(2), 435–465.

Cirelli, F., & Ólafsson, A. (2025, October). *Households' inaction in the deposit market.* VoxEU column, CEPR. (https://cepr.org/voxeu/columns/households-inaction-deposit-market)

Coulier, L., Pancaro, C., Pancotto, L., & Reghezza, A. (2025). *Banking on assumptions? How banks model deposit maturities* (Working Paper Series No. 3140). Frankfurt am Main, Germany: European Central Bank.

de Bondt, G. J. (2005). Interest rate pass-through: Empirical results for the euro area. *German Economic Review*, *6*(1), 37–78.

Drechsler, I., Savov, A., & Schnabl, P. (2017). The deposits channel of monetary policy. *Quarterly Journal of Economics*, *132*(4), 1819–1876.

Drechsler, I., Savov, A., & Schnabl, P. (2021). Banking on deposits: Maturity transformation without interest rate risk. *Journal of Finance*, *76*(3), 1091–1143.

Driscoll, J. C., & Judson, R. A. (2013). *Sticky deposit rates* (Finance and Economics Discussion Series No. 2013-80). Washington, DC: Board of Governors of the Federal Reserve System.

European Banking Authority. (2022a). *Final draft regulatory technical standards on the standardised approach and simplified standardised approach for IRRBB* (Regulatory Technical Standards No. EBA/RTS/2022/09). Paris, France: European Banking Authority.

European Banking Authority. (2022b). *Guidelines on the management of interest rate risk arising from non-trading book activities* (Guidelines No. EBA/GL/2022/14). Paris, France: European Banking Authority.

European Central Bank. (2006). *The pass-through from market interest rates to bank lending rates in the euro area* (Working Paper Series No. 136). Frankfurt am Main, Germany: European Central Bank.

European Central Bank. (2024). *Annual report 2023* (Tech. Rep.). Frankfurt am Main, Germany: European Central Bank.

European Central Bank. (2026a). *Banking structural statistical indicators — share of total assets held by the five largest credit institutions (cr5).* ECB Data Portal, Structural Statistical Indicators (SSI). Retrieved from https://data.ecb.europa.eu/data/datasets/SSI/SSI.A.BE.122C.S10.X.A1.Z0Z.Z (Series SSI.A.BE.122C.S10.X.A1.Z0Z.Z (Belgium), and equivalent series for DE and FR. Accessed July 2026)

European Central Bank. (2026b). *Deposits redeemable at notice up to 3 months (l23), households, belgium — outstanding amounts.* ECB Data Portal, Balance Sheet Items (BSI). Retrieved from https://data.ecb.europa.eu/data/datasets/BSI/BSI.M.BE.N.A.L23.A.1.U2.2250.Z01.E (Series BSI.M.BE.N.A.L23.A.1.U2.2250.Z01.E. Accessed July 2026)

Frauendorfer, K., & Schürle, M. (2007). Dynamic modeling and optimization of non-maturing accounts. In L. Matz & P. Neu (Eds.), *Liquidity risk measurement and management: A practitioner's guide to global best practices* (pp. 327–359). Singapore: John Wiley & Sons.

Greenwald, E., Schulhofer-Wohl, S., & Younger, J. (2023). *Deposit convexity, monetary policy and financial stability* (Working Paper No. 2315). Federal Reserve Bank of Dallas. doi: 10.24149/wp2315

Hamilton, J. D. (2018). Why you should never use the Hodrick–Prescott filter. *Review of Economics and Statistics*, *100*(5), 831–843.

Hannan, T. H., & Berger, A. N. (1991). The rigidity of prices: Evidence from the banking industry. *American Economic Review*, *81*(4), 938–945.

Hansen, B. E. (1996). Inference when a nuisance parameter is not identified under the null hypothesis. *Econometrica*, *64*(2), 413–430. doi: 10.2307/2171789

Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, *68*(3), 575–603. doi: 10.1111/1468-0262.00124

Hodrick, R. J., & Prescott, E. C. (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking*, *29*(1), 1–16.

Kalkbrener, M., & Willing, J. (2004). Risk management of non-maturing liabilities. *Journal of Banking and Finance*, *28*(7), 1547–1568.

Kho, S. (2025). Deposit market concentration and monetary transmission: Evidence from the euro area. *European Economic Review*, *173*. (Article 104933)

Litterman, R., & Scheinkman, J. (1991). Common factors affecting bond returns. *Journal of

Fixed Income, *1*(1), 54–61.

Lu, X., & Wu, L. (2026, February). *Banking on inattention: When deposits hedge or amplify interest rate risk.* VoxEU column, CEPR. (https://cepr.org/voxeu/columns/banking-inattention-when-deposits-hedge-or-amplify-interest-rate-risk)

Maes, K., & Timmermans, T. (2005). *Measuring the interest rate risk of belgian regulated savings deposits* (Financial Stability Review). Brussels, Belgium: National Bank of Belgium.

Molyneux, P., Pancotto, L., Reghezza, A., & Rodriguez d'Acri, C. (2020). *Interest rate risk and monetary policy normalisation in the euro area* (Working Paper Series No. 2496). Frankfurt am Main, Germany: European Central Bank.

Moura, A. (2024). Why you should never use the Hodrick–Prescott filter: A comment on Hamilton. *Journal of Comments and Replications in Economics*, *3*(2024-1), 1–17. doi: 10.18718/81781.31

National Bank of Belgium. (2023a, October). *La ruée des ménages belges sur les nouveaux bons d'état provoque une baisse record des dépôts bancaires.* NBB Blog. Retrieved from https://www.nbb.be/fr/blog/la-ruee-des-menages-belges-sur-les-nouveaux-bons-detat-provoque-une-baisse-record-des-depots (Accessed July 2026)

National Bank of Belgium. (2023b). *Rapport annuel 2023 — évolutions économiques et financières* (Rapport annuel). Bruxelles: National Bank of Belgium. Retrieved from https://www.nbb.be/doc/ts/publications/nbbreport/2023/fr/t1/rapport2023_tii_h7.pdf (Chapitre 7, Évolutions financières)

Neumark, D., & Sharpe, S. A. (1992). Market structure and the nature of price rigidity: Evidence from the market for consumer deposits. *Quarterly Journal of Economics*, *107*(2), 657–680.

Paraschiv, F., & Schürle, M. (2010). *Modeling client rate and volumes of non-maturing accounts* (Working paper). University of St. Gallen, Institute for Operations Research and Computational Finance.

Pesaran, M. H., Shin, Y., & Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics*, *16*(3), 289–326.

Ravn, M. O., & Uhlig, H. (2002). On adjusting the Hodrick–Prescott filter for the frequency of observations. *The Review of Economics and Statistics*, *84*(2), 371–376. doi: 10.1162/003465302317411604

von Feilitzen, H. (2011). *Modeling non-maturing deposits* (Unpublished master's thesis). KTH Royal Institute of Technology.

You, J., & Noh, Y. (2025). Deposit pricing under credit stress: Business-model heterogeneity and threshold effects. *Finance Research Letters*, *87*, 107650. doi: 10.1016/j.frl.2025.107650

Zanders. (2025). *Savings modelling series: Non-maturing deposits model concepts.* Zanders Insights. Retrieved from https://zandersgroup.com/en/insights/blog/savings-modelling-series-non-maturing-deposits-model-concepts (Accessed July 2026)

Appendix A

Data Coverage and Construction

This appendix documents the coverage of every series used in the thesis, its start, its length, and its source, together with the construction of a series built specifically for the analysis: the reallocation series used to validate the cyclical component (Section 5.1.4).

A.1 Coverage of the Series

The dataset is monthly and spans January 2003 to March 2026. The core rate and volume series (L23) are complete over the whole window; the swap curve begins in September 2004, which sets the effective start of any pillar-based specification and, with the 60-month moving-average window, means the pillar becomes usable from around 2009. Table A.1 reports each series.

Table A.1: Coverage of the series used in the thesis

Series	**Start**	**End**	***n***	**Source**
Core series				
Client rate (L23)	2003M1	2026M3	279	ECB (MIR)
Deposit volumes (L23)	2003M1	2026M3	279	ECB (BSI)
Market rates and drivers				
€STR / EONIA (spliced)	2003M1	2026M3	279	ECB
Euribor 3M	2003M1	2026M3	279	ECB
Euribor 6M	2003M1	2026M3	279	ECB
Euribor 1Y	2003M1	2026M3	279	ECB
Swap €STR 1Y	2004M9	2026M3	259	ECB
Swap €STR 3Y	2004M9	2026M3	259	ECB
Swap €STR 5Y	2004M9	2026M3	259	ECB
Swap €STR 7Y	2004M9	2026M3	259	ECB
Swap €STR 10Y	2004M9	2026M3	259	ECB
Swap €STR 15Y	2004M9	2026M3	259	ECB
OLO 1Y	2003M1	2026M3	279	NBB
OLO 2Y	2003M1	2026M3	279	NBB
Bon de caisse 3M	2003M1	2026M3	279	NBB
Bon de caisse 12M	2003M1	2026M3	279	NBB
Substitute instruments (reallocation, §5.1.4)				
Term deposits (L22)	2003M1	2026M3	279	ECB (BSI)

(continued)

Series	**Start**	**End**	***n***	**Source**
Retail government bonds	2003M1	2026M3	279	Debt Agency
Bank credit risk (App. C.6)				
CDS 5Y senior (3 banks)	2006M1	2026M3	243	Bloomberg
iTraxx Senior Fin. 5Y	2006M1	2026M3	243	Bloomberg
Macro-financial controls (App. C)				
GDP	2003M1	2026M1	277	NBB
Household confidence	2003M1	2026M3	279	NBB
HICP, inflation	2003M1	2026M3	279	Eurostat
Unemployment	2003M1	2026M3	279	Eurostat
EPU (Belgium)	2003M1	2026M3	279	policyuncertainty.com
Real-estate transactions	2003M1	2026M3	172	Statbel

Note. The spliced overnight rate combines EONIA until its discontinuation and the €STR thereafter, adjusted by the official 8.5-basis-point spread. GDP is quarterly, interpolated to monthly. Real-estate transactions contain an internal gap over the most recent months and are used only in the auxiliary test of Appendix C. CDS series begin in 2006, a shorter window than the core sample, and include Dexia SA to carry the 2008–2011 Belgian banking stress.

A.2 Construction of the Reallocation Series

The reallocation series of Section 5.1.4 measures, independently of the L23 balances, the money that left the regulated account for rate reasons. It combines two substitute instruments, each rendered as a cyclical quantity, and is constructed without any reference to the L23 series, its decomposition, or its smoothing parameter.

Household term deposits (L22). We take the ECB BSI outstanding-balance series for household deposits with agreed maturity and detrend it with the same Hodrick–Prescott filter and smoothing parameter used for the L23 balances (λ = 100,000), so that both series are cyclical deviations extracted on a common convention. The cyclical component of the L22 is the deviation of its outstanding balance from this trend.

Retail government bonds (*bons d'État*). Retail holdings are approximated from the issue amounts documented by the National Bank: a background stock of roughly €1 billion held before 2023, the €21.9 billion one-year issue of September 2023 (held over September 2023–August 2024), and the small €0.4 billion re-issue of September 2024. The series is therefore flat at its background level until the 2023 issue, steps up over the holding window of each *bon d'État*, and steps back down as it matures. This is a deliberately coarse, level-based proxy; what it needs to capture, for the validation of Section 5.1.4, is the timing and order of magnitude of the retail bond stock, not its exact path.

Aggregation. The two cyclical quantities are summed and sign-inverted: a movement of funds *into* the substitutes is an outflow *from* the regulated account, so the inverted sum is on the same sign convention as the L23 cyclical component against which it is compared (Section 5.1.4). No smoothing, interpolation or reference to the L23 enters at any step.

Appendix B

Client-Rate Model: Specification Search

This appendix documents the specification search behind the client-rate model of Chapter 4, in the order it was carried out. A preliminary exploration of the candidate market rates (Section B.1) motivates the pillar construction; the three candidate families are then compared on out-of-sample criteria, across 54 specifications, and the best of each reported (Section B.2); the retained family is finally examined more closely (Section B.3), including the exhaustive grid that confirms its optimum and the curve-slope factor, tested and discarded, that proves the most instructive of the refinements explored.

B.1 Preliminary Exploration

Before specifying the three families, we examined the candidate market rates individually. Univariate regressions of the client rate on each maturity (Euribor 3M–1Y, OLO 1Y–2Y, swaps 1Y–15Y, the policy rate and the *bon de caisse*), run over the full sample and within each of the three regimes, established which tenors track the client rate and how that tracking shifts across regimes. This exploration, together with rolling correlations and regime-by-regime betas, motivated the pillar construction that follows: the client rate co-moves most strongly with medium-maturity rates smoothed over several years, not with any single spot rate, which is the empirical content of a replicating-ladder pillar.

B.2 The Three Candidate Families

Section 4.3 compares three nested families of client-rate specification. We write them explicitly here. Throughout, m_{t-L} denotes a money-market rate (Euribor) at transmission lag L, and a *pillar* $p_t^{(M,W)}$ the W-month moving average of the maturity-M swap rate, the proxy for the yield of a replicating ladder of that maturity.

Family 1 — single pillar plus transmission. One funding anchor summarises the stable base; the lagged short rate carries the policy stance:

$$r_t = \alpha + \beta_p\, p_t^{(M,W)} + \beta_s\, m_{t-L} + \varepsilon_t. \tag{B.1}$$

Family 2 — two pillars. A real replicating portfolio mixes short and long tranches; this family weights a short and a long anchor separately:

$$r_t = \alpha + \beta_1\, p_t^{(M_1,W_1)} + \beta_2\, p_t^{(M_2,W_2)} + \beta_s\, m_{t-L} + \varepsilon_t, \qquad M_1 < M_2. \tag{B.2}$$

Family 3 — multi-rate combination. Following von Feilitzen (2011), a direct weighted combination of staggered market rates, without pillars:

$$r_t = \alpha + \sum_k \beta_k \, m_t^{(k)} + \varepsilon_t, \tag{B.3}$$

the $m_t^{(k)}$ spanning the curve (Euribor 3M, swap 1Y, swap 5Y).

Table B.1 reports the best specification of each family on the two out-of-sample criteria. The interpretation of this comparison, why the two-pillar family fails on identification, why the multi-rate family is dominated, and why the parsimonious single-pillar structure is retained, is given in Section 4.3; here we record only the figures behind it.

Table B.1: The three candidate families: best specification of each

	Best specification	**RMSE**	**Expanding**	$t_{\min}$	**Verdict**
F1	5Y pillar (60M) + Euribor 3M (L=4)	0.063	0.073	9.8	**retained**
F2	3Y/36M + 7Y/84M pillars + Euribor 6M	0.102	0.052	1.4	collinear
F3	Euribor 3M + swap 1Y + swap 5Y	0.176	0.120	3.3	dominated

Note. RMSE: out-of-sample error of the recursively simulated rate beyond 2019. Expanding: median RMSE across eleven expanding-window origins. $t_{\min}$: smallest HAC t-statistic among the market regressors. Source: author's computations.

B.3 Finer Analysis of the Retained Family

Family 1 having emerged as the best of the three, we subjected it to two further checks. First, an exhaustive grid over the single-pillar structure, three short rates, four transmission lags, four swap maturities and four moving-average windows, 192 combinations in all, confirmed that the retained specification sits on a plateau rather than at an isolated peak: neighbouring windows and lags yield near-identical out-of-sample errors, so the choice is robust to small perturbations. Second, we tested whether the shape of the yield curve carries information beyond the pillar and the short rate, by extracting the principal components of the swap curve on the in-sample window and adding the second component (PC2, the curve slope) to the specification Litterman and Scheinkman (1991). The curve-slope factor is significant at a three-month lag but loses its significance once the transmission lag is set to its retained value of four months (Section 4.3.1), which identifies it as a proxy for a misspecified lag rather than an independent driver.

Appendix C

Volume Model: Decomposition, Specification Search and Robustness

This appendix gathers the supporting material of Chapter 5. It documents why the trend–cycle decomposition is statistical rather than economic (Section C.1) and why the parametric trends standard in the deposit-modelling literature are likewise rejected on these data (Section C.3); the calibration of the smoothing parameter and the robustness of the decomposition, including the audit of its end-of-sample bias (Sections C.2 and C.4); the functional-form search behind the retained outflow model (Section C.5); and the two auxiliary channels tested and not retained, bank credit risk and policy uncertainty (Section C.6).

C.1 A Fundamentals-Based Trend: The Failed Economic Decomposition

The most natural candidate for the structural trend of Equation (5.1) is macroeconomic: if the stable core of deposits grows with the economy, the trend should be recoverable from a cointegrating relationship between deposit volumes and GDP, and the cyclical component would be the residual of that relationship, a decomposition with direct economic content and no filtering choice to defend. We tested this route in both real and nominal terms.

It fails on formal grounds. With real GDP, the Engle–Granger test does not reject the null of no cointegration ($p = 0.19$); the residual of the candidate long-run relation, $\ln \mathrm{Vol}_t = -23.7 + 2.50 \ln \mathrm{GDP}_t$ ($R^2 = 0.76$), is itself non-stationary (ADF $p = 0.06$) and carries a significant residual trend ($t = +2.2$). With nominal GDP the failure is starker (Engle–Granger $p = 0.34$; elasticity 1.12; residual ADF $p = 0.15$). The elasticities above unity identify the substantive problem: over 2003–2026, Belgian regulated deposits grew persistently faster than the economy, a financialisation of household savings that a GDP anchor cannot absorb, so that the residual of such a relation mixes a slow trend divergence with the cycle it is meant to isolate.

The consequences for the implied cyclical component are disqualifying. Its 2023 trough comes out at €−83.6 billion (real GDP) or €−68.4 billion (nominal), three times the documented field outflow of roughly €22 billion; its correlation with the filter-based cycle is modest (0.41–0.46); and the threshold model, re-estimated on it, degenerates: the decoupling coefficient loses significance or takes the wrong sign (t between −1.3 and +1.5 depending on the variant). A trend that is economically appealing but statistically absent produces a cyclical component that corresponds to no measured event; this negative result motivates the statistical decomposition of Section 5.1.2 and the two safeguards imposed on it (external calibration of λ; economic validation of the extracted cycle).

C.2 Robustness of the Statistical Decomposition

Section 5.1.2 sets the smoothing parameter at λ = 100,000 and shows that the frequency rule of Ravn and Uhlig (2002) independently selects the same region. This appendix reports the underlying grid, the alternative filters and, the point that matters most, what the choice of λ does and does not determine. Figure C.1 first shows the decomposition this discussion refers to: the observed L23 balances split into a smooth structural trend and a cyclical component whose two arbitrage episodes, the 2008 crisis and the 2023 *bon d'État*, the threshold model of Section 5.2 explains.

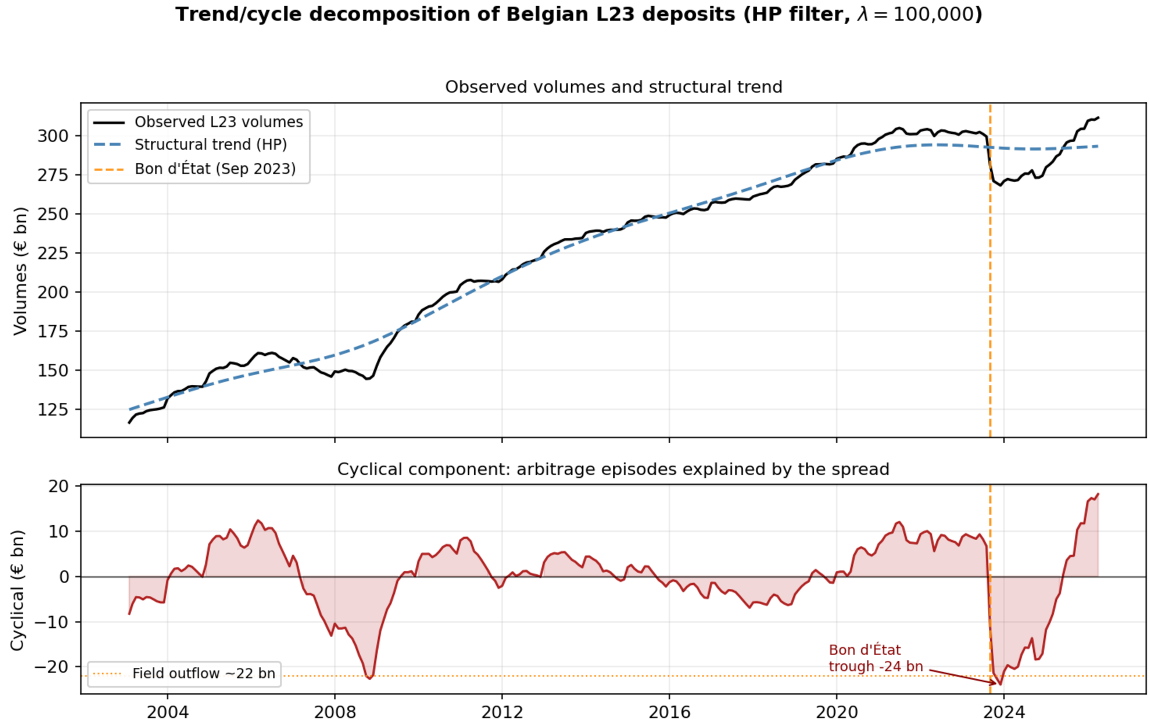

Figure C.1: Trend/cycle decomposition of Belgian L23 deposits (HP filter, λ = 100,000). **Top:** observed volumes and structural trend. **Bottom:** cyclical component; the 2023 trough reaches −23.9 billion against a field outflow of some 22 billion. Source: author's computations.

Sensitivity to the smoothing parameter. Varying λ from 14,400 to 1,600,000 moves the estimated 2023 trough within a bounded range: €−19.0 billion at λ = 14,400, €−23.9 billion at the retained 100,000, and €−24.3 billion at the rule-based 129,600. Two orders of magnitude of smoothing shift the trough by some five billion euros, and the qualitative result, a large, well-dated outflow consistent with the documented episode, is intact throughout.

The Hamilton filter. Hamilton (2018) argues that the HP filter can induce spurious dynamics and proposes a regression-based alternative. Applied to our series, it estimates a 2023 trough of €−39.6 billion in February 2024: two-thirds larger than the HP estimate and three months late. The documented episode adjudicates decisively against it, on both magnitude and timing.

The overstatement and the delay are precisely the defects the recent literature attributes to the Hamilton filter on series with sharp movements: its trend, lagged by construction, amplifies the estimated cycle (Moura, 2024).

The band-pass filter. The filter of Christiano and Fitzgerald (2003) isolates a frequency band rather than penalising curvature, and its asymmetric variant is designed to remain reliable at the end of the sample. The band must be specified rather than estimated: following the business-cycle convention, we retain fluctuations with a period between 18 and 96 months.[1] The two methods yield highly concordant components (correlation 0.94 over the full sample, 0.95 post-2023) and agree on the trough magnitude (€−23.9 against €−24.0 billion); the band-pass dates it five months later (April 2024 against November 2023), the known cost of its end-of-sample robustness and the reason the HP filter, which dates the documented episode correctly, is retained as the primary decomposition.

What λ does and does not determine. A distinction worth making explicit, because it bounds the reach of the whole discussion. The smoothing parameter governs the *amplitude* of the extracted cycle; it does not govern the *mechanism* estimated on it. The attractiveness threshold is a property of where the volume–spread relationship breaks, not of the scale of the cyclical series, and it is accordingly stable across the entire exercise: $\tau = 1.15$–1.35% across values of λ, across decomposition methods (HP, band-pass) and across end-of-sample truncations (samples cut at 2022, 2023 and 2024), with δ_{excess} between -14.1 and -14.7. What the external anchoring secures is the correspondence between a cyclical amplitude and a measured euro figure; what the data deliver independently of it are the threshold and the decoupling.

C.3 The Parametric Trend of the NMD Literature: A Rejected Standard

Section C.1 rejects an economic (GDP-anchored) trend. A second, distinct candidate deserves the same scrutiny, because it is the standard of the non-maturity-deposit literature itself: a *parametric* trend in time, estimated *jointly* with the rate response in a single equation. Kalkbrener and Willing (2004) model deposit volumes as an Ornstein–Uhlenbeck process with a linear trend, correlated with market and deposit rates; Bardenhewer (2007) splits the volume into a trend component, linear, quadratic or exponential, and an unexpected component, alongside moving averages of market rates. This approach has two attractions over a filter: it requires no smoothing parameter, and it estimates the trend and the arbitrage response simultaneously, so that no pre-filtered series is carried into a second-stage regression. We

[1] The cut-off periods are a modelling choice, not an estimated quantity: they define what is meant by "cyclical". We verified that the extracted cycle and the resulting threshold are insensitive to moderate variations of these bounds.

therefore tested it directly, in the form

$$\mathrm{Vol}_t = a + \sum_{k=1}^{K} b_k\, t^k + \delta_{\text{excess}} \max(0,\ \text{spread}_{t-5} - \tau) + u_t, \tag{C.1}$$

for $K = 1$ (linear), $K = 2$ (quadratic) and $K = 3$ (cubic), with τ estimated by the same grid search as in the retained model.

Results. The specification fails on the criteria that matter. Table C.1 reports the estimates. The linear and quadratic variants, the forms proposed by Kalkbrener and Willing (2004) and Bardenhewer (2007), leave *non-stationary residuals* (ADF $p = 0.084$ and 0.076), the diagnostic that the trend is mis-specified. The polynomial does not track the 2003–2026 trajectory of Belgian regulated deposits, whose growth combines a secular rise with the 2008 crisis and the pandemic savings surge. Only the cubic variant produces stationary residuals ($p = 0.004$).

Table C.1: One-step joint estimation: parametric trend and threshold

Trend	τ	δ_{excess}	Residual ADF	Implied 2023 trough	Verdict
Linear (Kalkbrener & Willing, 2004)	0.60	−20.3	$p = 0.084$	−51.3	mis-specified
Quadratic (Bardenhewer, 2007)	0.70	−18.8	$p = 0.076$	−45.6	mis-specified
Cubic	1.00	−19.4	$p = 0.004$	−41.3	implausible trough
Retained (HP, two-step)	1.35	−14.7	—	−23.4	**retained**

Note. Troughs in € billion, against a documented field outflow of roughly €22 billion. $\tau = 0.60$ is a corner solution at the lower bound of the estimation grid. The R^2 of the one-step variants (0.986–0.992) is not comparable with that of the retained model: it is computed on the volume level, where a time polynomial mechanically explains almost all of the variance of a series rising from €125 to €290 billion. Source: author's computations.

Why it fails, and what the failure shows. The implied outflows settle the matter: €−41 to €−51 billion for the 2023 episode, roughly twice the documented figure. The pathology is the same as that of the GDP-anchored trend of Section C.1 (which implied €−68 to €−84 billion), and it has the same cause: when the trend is too rigid to follow the series, the threshold term absorbs the trend mis-specification in addition to the arbitrage it is meant to measure. The symptoms are diagnostic: the threshold is pulled down (to a corner solution at 0.60 for the linear variant) so that the term activates early enough to soak up the residual trend, and the slope is inflated accordingly (−19 to −20 against −14.7).

The result is not an argument against the literature but a statement about this product. The HP filter is a penalised smoothing spline whose trend tends to a straight line as $\lambda \to \infty$; the linear and polynomial trends of Kalkbrener and Willing (2004) and Bardenhewer (2007) are therefore the rigid limiting cases of the decomposition retained in Chapter 5. What these tests establish is that those limiting cases are rejected on Belgian regulated savings, formally, by the stationarity of their residuals, and substantively, by the implausibility of the outflows they imply. A finite smoothing parameter is not a convenience but a requirement of the data; and since Section C.2 shows that its value is pinned down by two independent criteria, the flexibility it buys is not paid for in arbitrariness.

Two clarifications, in fairness to the alternative. The one-step approach remains preferable *in principle*: it avoids carrying a pre-filtered series into a second stage, and its failure here is one of trend specification, not of the joint-estimation logic; a one-step model with a sufficiently flexible trend (a spline, say, which is what the HP filter effectively supplies) would be the natural synthesis, and is left as a methodological perspective. And the cubic variant, the only well-specified one, places the threshold at $\tau = 1.00$, below our 1.35, but in the direction the contamination predicts, since absorbing residual trend pulls the threshold down. Read that way, it corroborates rather than contradicts the retained estimate.

C.4 The End-of-Sample Bias of the HP Cycle: An Audit

The HP filter is least reliable at the end of the sample, where no future observations discipline the trend. We quantify this weakness with two complementary measurements, whose conclusions are used in Sections 5.1.4 and 5.2.3.

Against the edge-robust control. Over the core of the sample (before 2024), the HP and band-pass cyclical components are essentially interchangeable: mean gap €−0.1 billion. Over the final stretch (≥2024), the HP component sits systematically *above* the edge-robust one, by €+3.2 billion on average in 2024, +4.8 in 2025 and +8.0 over the first months of 2026 (+4.5 on average over the whole boundary window). The gap widens as the boundary approaches, the signature of an edge artefact inflating an otherwise real reflux rather than of an economic phenomenon, which the band-pass filter would also record.

Real-time revision. A second measurement dispenses with any alternative filter. For each month of the boundary window, we re-estimate the HP cycle using only the data available at that date and compare the resulting end-point value with the full-sample estimate. The end-of-sample values are revised *downward* by €6.3 billion on average once later data arrive (extremes: −11.7 to +3.9). The final months of the extracted cycle are, in short, systematically inflated in real time, in a direction and magnitude consistent with the gap to the band-pass control.

Both measurements imply that the 2024–2026 observations of the HP cycle overstate the true cyclical level by several billion euros. This is the basis for the qualitative reading of the recent rebound in Section 5.1.4, for the exclusion argument of Section 5.2.3, and for the conservativeness of the estimated decoupling slope noted in Section 5.1.3.

C.5 The Functional Form: The Below-Threshold Slope Examined

The retained model of Equation (5.2) is flat below the threshold. The natural enrichment adds a below-threshold slope:

$$\text{Cycle}_t = \alpha + \beta_{\text{base}}\,\text{spread}_{t-5} + \delta_{\text{excess}}\,\max(0,\ \text{spread}_{t-5} - \tau) + u_t. \tag{C.2}$$

Estimated on the full sample, the variant yields $\tau = 1.20\%$, $\beta_{\text{base}} = +2.74$ ($t = 1.8$, significant at the 10% but not the 5% level) and $\delta_{\text{excess}} = -18.95$ ($t = -5.8$), with $R^2 = 0.60$ against 0.56 for the flat form. Three examinations led to its rejection.

Leverage. The significance of β_{base} rests on the end-of-sample observations whose upward bias Section C.4 quantifies. Excluding the boundary window (≥2024), the coefficient falls to $t = 1.1$ and the information-criterion advantage of the richer form collapses (from ΔAIC = 21 to 5); excluding, in addition, the 2022 transition months leaves the same picture ($t = 1.5$, ΔAIC = 10). Within the model's declared perimeter (excluding 2008–2009) but including the boundary, the coefficient is significant ($t = 2.2$), significance that the leverage analysis attributes to the boundary points themselves.

Sign. Under the substitution reading that motivates the model, a below-threshold slope should be *negative*: as the outside option gains ground, the account should retain less. The estimated slope is positive, the cyclical component is highest precisely when the alternative is most attractive short of the threshold, which no substitution mechanism can produce.

Interpretation. Three candidate readings of the positive slope were tested against the spread on the below-threshold subsample (HAC standard errors throughout). A *remuneration-level* effect (the level of the client rate driving accumulation) explains almost nothing alone ($R^2 = 0.03$, $t = 0.7$) and dies entirely in a horse-race with the spread ($t = -0.1$); a structural variant placing the rate level below the threshold and the spread beyond it is rejected on information criteria (ΔAIC = +29) with an insignificant level coefficient ($t = 0.5$). A *rate-momentum* effect (the three-month change in the client rate) explains nothing ($R^2 = 0.01$, $t = -0.6$) and is annihilated by the spread ($t = 0.0$). A *market-level* effect (the level of the competing rate) is observationally inseparable from the spread below the threshold, the client rate being quasi-flat there, the spread is the market level up to a constant, and neither survives joint inclusion ($t = 0.8$ and 0.1). None of the three readings, in short, gives the slope an economic content distinct from the boundary artefact quantified in Section C.4. The flat form is therefore retained; it is also the form that nests into the multi-strike extension of Chapter 6.

C.6 Two Auxiliary Channels: Bank Credit Risk and Policy Uncertainty

The threshold model attributes cyclical movements to the market-to-deposit spread alone. Two other channels have a theoretical claim on the same variation, and both were tested rather than merely acknowledged. Neither is retained, for different reasons, but the contrast is instructive.

Bank credit risk. The 2008–2009 discrepancies documented in Section 5.3.4 suggest a credit-risk channel: with the deposit guarantee capped at €100,000, a rise in the perceived risk of bank failure gives holders of large balances an incentive to move funds independently of the rate on offer. We tested it with five-year senior CDS spreads of the major Belgian deposit banks and the iTraxx Senior Financials index (correlation 0.88 between the two measures), on the CDS-matched subsample (2006–2026), in both directions in which it could enter, and throughout on the retained flat specification.

On the *rate* side, the hypothesis is that a stressed bank over-remunerates to retain deposits, which would explain the positive residual of the client-rate model in 2008. Augmenting the long-run rate equation with each CDS measure, the coefficient enters with the wrong (negative) sign, is never significant (t between -0.6 and -1.6), adds nothing to fit ($\Delta R^2 \leq 0.001$) and leaves the 2008 residual intact. The hypothesis is rejected.

On the *volume* side, a CDS term carries the expected flight-to-safety sign on all three measures, roughly €1.5–2.2 billion of additional cyclical outflow per +100 basis points of CDS, and reduces the 2008–2009 mean absolute error (from €6.8 billion to €5.2–6.3 billion). Its identification, however, is not robust: the coefficient is significant on the iTraxx index ($t = -2.1$) and on the Dexia–ING variant ($t = -2.3$) but not on the Belgian bank CDS average ($t = -1.8$), so whether the channel clears the 5% threshold depends on which risk measure is used. A stress-regime variant, activating the term only above a CDS level κ, fares no better: below $\kappa = 40$ bps the coefficient stays insignificant ($t \approx -1.8$) while classifying more than half the sample as "stressed", depriving the regime of any stress content, and significance re-emerges only for $\kappa \geq 170$ bps, where fewer than twenty observations identify it. In every variant the arbitrage mechanism is untouched, δ_{excess} moving by less than one unit (from -14.9 to between -15.3 and -15.8). The channel carries the right sign but cannot be robustly identified on aggregate data.

Policy uncertainty. The second channel is *precautionary saving*: when economic and policy uncertainty rises, households are expected to save more and to favour liquid, safe assets, of which the regulated account is the archetype. This predicts a *positive* coefficient on an uncertainty measure, the opposite sign to the credit-risk channel. Augmenting the retained volume equation with the Belgian Economic Policy Uncertainty index in level, over the model's declared perimeter (excluding the 2008–2009 confidence crisis), the index enters with the expected positive sign but falls short of significance ($\hat{\theta} = +0.021$, $t = 1.51$, HAC standard errors, $n = 242$). It is essentially orthogonal to the spread (correlation -0.08), so its modest explanatory contribution is at least not an artefact of collinearity with the arbitrage signal; and it operates, if at all, through the *level* of uncertainty rather than its transitory movements: a shock measure, the deviation from a twelve-month moving average, is insignificant and of the wrong sign ($t = -0.96$), consistent with precautionary saving being a slow adjustment to a persistent environment rather than a reflex to a surprise.

Two grounds settle its exclusion. First, the level effect is not robustly identified: it does not clear the 5% threshold on the retained specification, and its inclusion leaves the decoupling coefficient essentially unchanged (a shift of about a third of a unit), so it neither carries independent weight nor improves the central mechanism. Second, and decisively for the purpose of this thesis, the index is not *projectable*. The coupling chain maps an interest-rate scenario into a volume path; policy uncertainty is not a function of that scenario and would have to be held at an arbitrary constant in simulation, where it would contribute nothing. A variable that cannot be conditioned on the scenario has no place in a scenario-consistent model, whatever its in-sample merit.

Why the aggregate sees neither. The two channels fail for related reasons: both carry the expected sign, neither is robustly identified on the aggregate, and the second is, in addition, unusable in a scenario-consistent chain. They also share a revealing property, visible when the perimeter is widened. Over the *full* sample, including 2008–2009, the already-marginal uncertainty coefficient weakens further ($t = 0.91$, against 1.51 on the declared perimeter). The reason is economic rather than statistical: during the crisis, uncertainty and bank credit risk moved together but pushed deposits in *opposite* directions, uncertainty towards the safe, liquid account, and the perceived risk of bank failure away from it for balances above the guarantee ceiling. The two forces partly offset, and neither is separately identifiable once they are summed. It is the same lesson the thesis meets at the asymmetry of the client rate (Section 4.4.2) and at the choice of competing instrument (Section 5.2.2): the aggregate reveals that something happened, and dissolves the structure of what it was.

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