Bocconi

Private Credit, SRTs and the Banking System: Mind the Protection Gap

Working Paper N. 274
May 2026

Riccardo Russo

BAFFI
CENTRE ON ECONOMICS, FINANCE AND REGULATION

Working Paper Series

Private Credit, SRTs and the Banking System: Mind the Protection Gap

Riccardo Russo*

This version: May, 2026

Abstract

Private credit funds have emerged as major players in corporate lending and protection sales to banks within the SRT market. This paper models a contagion loop between the banking system and private credit funds working through SRTs. The loop is activated by two elements whose interaction is scarcely formalised in the literature: the circular leverage embedded in SRT financing and the liquidity mismatch introduced by semi-liquid private credit funds. Together, they can result in shrinking protection for banks, i.e. a protection gap. Implemented through a simulation, the model provides evidence on the impact of this mechanism on banks' capital relief and lending capacity, and represents a useful risk-assessment tool to both private and public agents.

JEL codes: C15, G01, G21, G23, G28
Keywords: private credit, SRT, banking system, doom loop, systemic risk

* Baffi CAREFIN Research Centre, Bocconi University, r.russo@bocconialumni.it

1. Introduction

Private credit usually takes the form of floating-rate, covenant-intensive loans of medium-term maturity, extended by different forms of private credit funds (PCFs) to credit-constrained mid-sized enterprises, or to larger companies favouring the flexibility of private contracts (Chernenko et al., 2022; IMF, 2024; Cai and Haque, 2024; BdF, 2025; Ivashina, 2025). Together with other non-bank financial institutions (NBFIs), the sector has experienced an exceptional growth since the Global Financial Crisis (GFC), and even more so in recent years (Acharya et al., 2024; Fillat et al., 2025). At the global scale, PCFs' assets under management have increased from less than \$500 billion in 2010 to around \$2.5 trillion in 2025, and are expected to reach \$4.5 trillion by 2030 (Preqin data). This expansion has been fuelled by the post-2008 banking reforms and the search for yield in the prolonged low-rate environment after the crisis (Acharya et al., 2024; Block et al., 2024). Growth has been heavily concentrated in the US, followed by Europe and, to a lesser extent, Asia.

PCFs have historically been closed-end, with long-term institutional investors (limited partners) as capital providers (Chernenko et al., 2025). Until recently, this structure has made them less subject to the liquidity mismatches characterising other kinds of NBFIs (Matvos et al., 2026). Yet, the growth of semi-liquid PCFs, like non-traded business development companies (BDCs) in the US, some European long-term investment funds (ELTIFs) in the European Union, or long-term asset funds (LTAFs) in the UK, open to retail investors and allowing redemptions, has changed this picture.

Recent stress events have led several of them (from Morgan Stanley, Ares, Apollo, Blue Owl, Blackrock, Blackstone) to activate gates to surging redemption requests in the wake of generalised fears about the soundness of the sector (FT, 2026a,b,c,d). This has shown how retail investors can be subject to panics, leading to sudden redemption waves on the fund because of its semi-liquid structure. These can be problematic given the illiquid nature of PCFs' assets, made of their loan portfolio and uninvested capital (dry powder). "Semi-liquid" structures have thus reintroduced concerns on liquidity mismatches.

Besides direct loan origination, PCFs also heavily engage in synthetic risk transfer (SRT) as protection sellers to banks (Cai and Haque, 2024; EBA, 2025; Cortes et al., 2025). As protection buyers, banks benefit from SRTs through regulatory capital relief (BCBS, 2019; González and Triandafil, 2023; Chernenko et al., 2025). In unfunded SRTs, the risk weight applicable to the protected tranche becomes the one of the protection seller, which should be lower than the original. In funded SRTs, which are the most common, the protection seller posts collateral upfront, making the resulting risk weight on the protected tranche essentially null. At the same time, the bank benefits from a reduced risk weight on the retained tranches. For protection sellers, participating in SRT transactions gives access to bank-originated loan portfolios, underwritten to bank credit standards, without the burden of origination, servicing and direct borrower relationships (IACPM, 2025; BCBS, 2026).

This link between PCFs and banks is often coupled with other bilateral exposures, where PCFs post their SRT notes as collateral to obtain debt financing from the same banking system (EBA, 2024; Cai and Haque, 2024; Acharya et al., 2025; Cortes et al., 2025; Berg and Hoon Lee, 2026; BCBS, 2026; FSB, 2026). Overall, the PCF is then tied to the banking sector via the SRT note on the asset side, and via debt financing (which can take the form of credit lines, term loans or repo facilities) on the liability side. This implies that a fraction of the apparent risk transfer of the SRT is back-funded by the protection buyer, and that the net reduction in the banking sector's exposure to the underlying portfolio is smaller than the gross capital relief obtained in the first place (FSB, 2026).

Despite this, it is still rational for banks to lend to PCFs for several reasons (Acharya et al., 2025; Fillat et al., 2025; FSB, 2026). First, credit lines can generate more revenues than those stemming from the lost portion of freed-up capital. Second, the amount of regulatory capital needed for the credit lines can be reasonably lower than the one needed for the underlying SRT assets. Third, each bank is incentivised to offer such credit lines to evict their competitors. Finally, providing the credit line can strengthen the business relationship between protection seller and protection buyer on the parallel SRT transaction, paving the way for banks to the sell other "bundled" products in the future (Acharya et al., 2024). Throughout the paper, I adopt this last view, allowing for simplifying assumptions.

Together with the concerns introduced by semi-liquid PCFs, this interconnection with the banking system can represent a risk-amplification channel. If a PCF's retail investors observe a sequence of defaults on the fund's exposures, like the loans underlying an SRT, or are pushed to doubt about its creditworthiness, they may run. With illiquid loans and only limited dry powder on the asset side, the PCF's liquidity may be eroded severely. As a consequence, and since SRTs are usually rolled over at maturity, the PCF could not be able resell the same amount of protection to the bank,[1] resulting in a capital relief loss besides the one given by the debt financing to the fund. For capital-constrained banks, this would translate into a lending contraction which, through the standard credit channel (Bernanke and Gertler, 1995; Kashyap and Stein, 2000, Peek and Rosengren, 2000; Jiménez et al., 2017; Couaillier et al., 2025) would amplify default rates and restart the sequence in a perverse loop.

Policy work has focused on borrowers' fragility, opaque valuations, hidden leverage and unclear interconnections among market participants (IMF, 2024; EBA, 2025; ESRB, 2025; BdF, 2025; Fed, 2025). The press has echoed these concerns in light of the sector's stress episodes, including the bankruptcies of FirstBrands and Tricolors (FT, 2026a,b,c,d; Bloomberg, 2026a,b). Data on private credit is indeed scarce, and the definition of private credit is not even fully agreed upon among different regulators and supervisors. This has limited the formal modelling of private credit and its interconnections with the traditional banking system, both on the policy and academic sides. Yet, as shown by the recent events, such lack of formalisation may create blind spots where risk would concentrate (FSB, 2026).

With this paper, I propose a model of contagion between PCFs and the banking system, working in a loop through SRTs, and potentially affecting the real economy through the credit channel. By capturing the joint effects of back-funded leverage and the liquidity mismatch introduced by semi-liquid PCFs, the model allows (i) to point to the critical areas where policy intervention should focus and (ii), when coupled with banks' or PCFs' proprietary data, to quantify their potential losses. As such, it provides these agents, as well as regulators, with a useful risk-assessment tool which is currently lacking.

[1] In a funded SRT, this stems from the inability of the PCF to post the fresh collateral needed to roll over the SRT. In an unfunded SRT, this can stem from downgrade of the PCF creditworthiness (FT, 2026b).

To be more specific, starting from a given situation linking a representative bank and PCF, the model delivers estimates of the "protection gap" that can emerge if SRTs are not fully rolled over, of the capital relief loss banks would endure as a consequence, and of the lending contraction that may trigger higher default rates. Overall, this structure resembles a top-down stress test. As in this type of exercises, it is assumed that no corrective action is taken, or can be taken, while events happen, providing a clear picture of what backlashes may result from the initial shock. Therefore, in the context of this paper, the loop cannot be stopped for good, only for full rollover failure.

The rest of this paper is structured as follows: section 2 reviews the related literature from which this work draws, and highlighting its contributions to it. Section 3 develops the model linking PCFs and banking system via SRTs. Given the lack of public data, section 4 proposes an implementation of the model based on a simulation of a two-round loop. A discussion of results is presented in section 5. Section 6 concludes.

2. Related Literature

This paper draws from three strands of the literature: (i) the one analysing private credit as an emerging segment of corporate financing, (ii) the one focusing on SRTs as an instrument of bank capital relief and (iii) the broader strand on the bank-NBFI nexus. As regards the first strand, Chernenko et al. (2025) explicitly link the post-2009 growth of private credit to bank capital regulations, while Fillat et al. (2025) and Acharya et al. (2025) show how bank credit facilities have supported its expansion. Cai and Haque (2024) characterise the typical borrower, loan structure and risk profile of private credit transactions. Block et al. (2024) provide stylised facts on the size, geographical distribution and pricing of the asset class.

On the risk side, the literature is populated by institutional contributions (IMF, 2024; EBA, 2025; ESRB, 2025; BdF, 2025; Fed, 2025), focusing on the sector's macro risks and its opaqueness, but not providing a quantifiable formalisation that can be replicated or adapted to specific entities. Among academic contributions, a recent study by Matvos et al. (2026) highlights features in the individual balance-sheet structure of PCFs that would shield them from external shocks. Yet, this micro-oriented approach can be problematic in spotting risks at the systemic level, especially when generated by interconnections rather than individual fragilities (Allen and Gale, 2000; Acemoglu et al., 2015; Cetorelli et al., 2023).

From the second strand on SRTs, Acharya et al. (2013) studied the use of securitisation as a vehicle for regulatory arbitrage in the period preceding the GFC, which supports the view that interconnections via these instruments can contribute to the build-up of financial instability. More recent work has focused on post-Basel III SRTs (BCBS, 2019, 2026; EBA, 2024, 2025). González and Triandafil (2023) provide an assessment of the European market based on supervisory data, documenting typical structures, reference pools and capital-relief outcomes. Cortes et al. (2025) develop an analytical framework for the evaluation of synthetic securitisations and their capital-relief implications, while Osberghaus and Schepens (2026) use European supervisory data to quantify the extent to which credit risk is genuinely transferred outside the banking sector through these instruments, finding that a non-trivial share of risk circles back via bank-funding of protection sellers.

The third and final strand relates to the systemic implications between banks and NBFIs in general. Acharya et al. (2024) provide a comprehensive picture of bank-NBFI exposures in the US, showing again that growth of NBFIs is supported by bank financing. Cetorelli et al. (2023) develop a fire-sale framework that propagates portfolio shocks across NBFI segments and back to banks, generalising the asset-manager fire sale literature (Cetorelli et al., 2016). On the open-end fund segment, Goldstein et al. (2017) and Falato et al. (2021) document fragility in corporate-bond funds and identify a first-mover advantage in redemptions that can rationalise run-like behaviour. Bernstein et al. (2018) provide evidence on fire-sale discounts in distressed assets disposals. The credit channel mechanism through which bank capital constraints translate into reductions in lending, and ultimately in real activity, is a long-established result in the literature (Bernanke and Gertler, 1995; Kashyap and Stein, 2000, Peek and Rosengren, 2000), with more recent confirmations from European credit registry data (Jiménez et al., 2017; Couaillier et al., 2025).

Relative to the first strand, the contribution of this paper is given by the formalisation of a major channel through which risks can move from single institutions to a systemic level, that can be adopted and replicated by individual institutions, either public or private, in their risk management and assessment processes. Relative to the second strand, this paper identifies SRTs as a specific channel through which PCFs and the banking system can be linked, and models how, under certain circumstances, they may fuel shock propagation and financial instability. This same contribution applies to the third and last strand, which until now has not focused on formal modelling of interconnections between banking and private credit sectors, nor it has modelled the repercussions in terms of financing of the economy.

3. Model

The model consists of a three-stage loop that repeats itself in rounds. Rounds are indexed with the time subscript t, where $t = 0$ is an ex-ante moment defined as "inception". Inception is not part of the loop and only takes place once. Its aim is to present the pre-existing links between a representative bank and a semi-liquid PCF: an SRT and a credit line. The loop starts at $t = 1$. Following a series of defaults on the SRT's protected tranche and the expiration of the credit line, the first stage delivers the amount of liquid assets left available to the PCF and the behaviour of its retail investors. The second stage shows how these dynamics translate into a protection gap for the bank, which, in turn, can result into a contraction in its lending capacity in stage 3. A new default rate, fuelled by lending contraction, restarts the loop.

3.1. Inception

The economy's financial sector is made of a representative bank and a semi-liquid PCF, with retail investors as part of its capital providers. At inception, bank and PCF are signed in two contracts: first, a funded SRT, where the bank buys protection from the PCF; second, a credit line, through which the bank provides financing to the PCF accepting the SRT note as collateral. The bank holds a corporate loan portfolio of value A_0, with an associated risk weight ω_0. Via the SRT, the bank buys protection on its portfolio's junior tranche, of thickness $0 < j_0 < 1$, while retaining the senior tranche.[2] The value of the SRT note, equal to the junior tranche, is therefore $P_0 = j_0 A_0$.

The *gross* capital relief for the bank stems from two channels: first, transferring the credit risk of the junior tranche frees the regulatory capital to be posted otherwise. This amounts to $\rho_0 = k\omega_0 P_0$, with $0 < k < 1$ the fixed bank's minimum capital ratio. Second, the bank benefits from an improved risk weight $\widehat{\omega}_0 < \omega_0$ on the remaining senior tranche (BCBS, 2019; 2026), obtaining a further relief $r_0 = k(\omega_0 - \widehat{\omega}_0)(A_0 - P_0)$. The overall gross relief is then:

$$G_0 = \rho_0 + r_0. \tag{1}$$

The SRT being funded, the PCF must post collateral $C_0 = P_0$ upfront. To classify as collateral, assets must be liquid, so that, if defaults materialise on the protected tranche, the bank draws from them. It is assumed that all the liquid assets of the PCF are posted as collateral.

The credit line is in turn collateralised by the SRT. Therefore, its value only equals a fraction $0 < \lambda_0 < 1$ of the note, i.e. $L_0 = \lambda_0 P_0$. This setting implies that a fraction λ_0 of ρ_0 is back-funded by the bank, and that a component $\lambda_0 \rho_0$ of G_0 has not genuinely left the banking sector. The *net* capital relief of the bank is thus equal to:

$$N_0 = G_0 - \lambda_0 \rho_0. \tag{2}$$

The rationale for the credit line to be extended despite the reduction in net capital relief is given by the "business relationship" view (Acharya et al., 2024), where the bank trades a part of current protection for future business opportunities.

3.2. The Loop

3.2.1. Stage 1: Distress

As first stage of round t begins, two events happen in sequence: first, a default rate $0 \leq e_t \leq 1$ materialises on the junior tranche of the SRT. Its value thus decreases by $e_t P_{t-1}$, and the bank reclaims the same amount from the collateral posted in the previous round.[3] Second, the SRT matures and the credit line expires. At this point, the PCF has to reimburse $L_{t-1} = \lambda_{t-1} P_{t-1}$ – the value of the credit line from the previous round – to the bank, after receiving from it the

[2] For simplicity, I consider a two-tranche (junior and senior) SRT – most common in the US – rather than the three-tranche (junior, mezzanine and senior) alternative – more common in Europe. This choice does not affect the meaningfulness of results. The SRT is assumed to be funded, this structure being predominant both in Europe and the US (BCBS, 2026).

[3] P_{t-1} is the SRT value at $t-1$. The collateral draw by the bank is orderly and pre-contracted and, since collateral was posted upfront, does not entail any change in the PCF's net asset value (NAV).

value of the SRT net of defaults, $J_t = (1 - e_t)P_{t-1}$.[4] Since all the PCF's liquid assets had been posted as collateral for the credit line in the previous round, the net value of these two transactions will give the PCF's liquid assets in t as:[5]

$$Z_t = \max[0, J_t - L_{t-1}]. \tag{3}$$

PCF's retail investors observe the default rate, or are anyway able to infer it. For instance, through the press (FT, 2026a,b,c,d; Bloomberg, 2026a,b). If e_t is below a certain stress threshold $\bar{e}$ – assumed constant across rounds –, they will hold their shares and no redemption request will follow.[6] If the rate exceeds $\bar{e}$, these investors will run, exploiting the semi-liquid structure of the PCF. To honour redemption requests in the first place, the PCF can only sell liquid assets. The share of Z_t that will be used to cover redemptions is $0 \leq \mu_t \leq 1$. Of course, $\mu_t = 0$ for $e_t < \bar{e}$ since there are no redemptions to service. Instead, as long as $e_t > \bar{e}$, $\mu_t > 0$ since at least part of liquid assets will have to be used to service investors. Furthermore, it can be reasonable to assume that μ_t is not exogenous, but is rather an increasing function of e_t (Goldstein and Pauzner, 2005; Brunnermeier and Pedersen, 2009; Bernstein et al., 2018; Falato et al., 2021). Thus, I define it as $\mu_t = \max\{0, \min[\alpha(e_t - \bar{e}), 1]\}$, with $\alpha > 1$ and constant across rounds.

Liquid assets can be exhausted if, even after imposing gates, Z_t does not suffice to cover for all redemption requests. In this case, the PCF will have to sell illiquid assets at a fire sale discount, resulting in NAV markdowns. Sometimes with a lag, these are eventually observable by investors, fuelling generalised mistrust about the creditworthiness of the PCF (Brunnermeier and Pedersen, 2009; Bernstein et al., 2018; FT, 2026a,b,c,d; Bloomberg, 2026a,b; JPM, 2026). In turn, this can trigger further redemption requests (Goldstein et al., 2017; Falato et al., 2021), worsening the run. Overall, the first stage can thus deliver three different outcomes: no run with no liquid assets depleted to service redemption (O1.1: $e_t < \bar{e}$, $\mu_t = 0$), a run with sufficient liquid assets to meet investors' demand (O1.2: $e_t > \bar{e}$, $\mu_t < 1$), or a run with insufficient liquid assets to pay all redemptions at once (O1.3: $e_t > \bar{e}$, $\mu_t = 1$).

3.2.2. Stage 2: Contagion

At the beginning of the second stage, bank and PCF aim to renew the two contracts, SRT and credit line, following the same logic as at inception. I assume no replenishment of the bank's loan portfolio, and that the bank will seek protection on a junior tranche equal to the one from

[4] Coherently with the business relationship approach, I adopt the simplifying assumption that the bank was not required to pay premia to the PCF throughout the lifetime of the SRT, and that the PCF was not required to pay interest on the credit line. Equivalently, it can be assumed that the value of the former equals the value of the latter, netting one another out.

[5] Of course, liquid assets must be non-negative. Yet, $J_t - L_t = (1 - e_t - \lambda_{t-1})P_{t-1}$ becomes negative for $e_t > 1 - \lambda_{t-1}$. In this case, the PCF's new liquid assets will be null – as returned by equation 3 –, but the fund will also have to sell part of its illiquid portfolio if it aims to repay the bank. Even though stemming from a different channel, this is equivalent to the outcome O1.3 presented later on in this section. Therefore, assuming that the PCF's only liquid assts stem from the two transactions, but at the same time imposing the lower bound of 0 to Z_t, does not entail any loss of generality. Another observation is that $Z_t \geq 0$ does not imply that the bank is a net "loser" from the two transactions. The capital relief obtained through the SRT would have generated profits elsewhere, and the business relationship established with the PCF could open further sale opportunities in the future.

[6] Recent stress events suggest a threshold between 6% and 10% (FT, 2026a,b,c,d).

the previous round, net of defaults which have taken place in the meantime. This gives the bank's demand for protection as J_t. On the other hand, I also assume no replenishment of the PCF's liquid assets, giving its maximum protection supply as:

$$P_t = (1 - \mu_t)Z_t. \tag{4}$$

It is easy to see that under O1.1 (4) attains its peak, $P_t = Z_t$. Under O1.2, P_t will be lower than for O1.1, but still larger than for O1.3, which is the most severe outcome with $P_t = 0$.

The PCF commits again all of its available liquidity to collateralise the SRT. Still, under no outcome this would suffice to fully cover the bank's current demand. The size of protection that can be sold on the junior tranche is capped by the PCF's protection capacity, $P_t < J_t$.[7] As a result, a positive wedge emerges between demand and supply of protection, equal to $W_t = J_t - P_t$. Despite this, it will be rational for the bank to proceed with a partial rollover of the SRT equal to the PCF's capacity, benefitting from partial rather than no protection on the tranche. P_t thus becomes the new SRT value in round t.

Together with the devaluation of the junior tranche due to defaults, $e_t P_{t-1}$, W_t contributes to the wider "protection gap" between the amount of protection sold in $t-1$ and t. The protection gap is defined as $\Gamma_t = e_t P_{t-1} + W_t = P_{t-1} - P_t$, and corresponds to the amount by which the SRT note shrinks round after round.[8] It stems from three effects which gradually constrain PCF's liquidity: first, the realisation of defaults, which reduces the reimbursement of collateral by the bank; second, the credit line, needing to be repaid; third, the further drain potentially implied by redemption requests. The protection gap is the largest under O1.3, where the absence of available liquid assets prevents any rollover of the SRT, and the smallest under O1.1 where, together with O1.2, at least a partial rollover is possible.

It follows that the maximum value of the credit line, collateralised by the new SRT note, will be:

$$L_t = \lambda_t P_t, \tag{5}$$

where, for simplicity, I assume $\lambda_t = \lambda_{t-1}$.[9] Together with P_t, L_t will also decrease as rounds proceed, pro-cyclically reducing the PCF's liquidity right as it would need it the most. Therefore, it is assumed that the full credit capacity will be used, and that the new value of the line corresponds to (5).

Under any of the outcomes, the protection gap translates into a loss of net capital relief for the bank. This stems from the same channels that provided it in the first place: on the one hand, the relief on the junior tranche, net of back-funding, is reduced proportionally to the smaller tranche size; on the other hand, the pre- and post-SRT risk-weight differential on the senior tranche is reduced, as the attachment point of this tranche decreases. For simplicity, I assume this second effect is also proportional to the shrinking of the junior tranche. So, the overall loss

[7] P_t is at most equal to Z_t, and Z_t is smaller than J_t. So, P_t is necessarily smaller than J_t.
[8] By definition, $J_t \leq P_{t-1}$. Therefore, P_t is also smaller than P_{t-1}.
[9] Alternatively, it may be assumed $\lambda_t < \lambda_{t-1}$ because of the devaluation of the SRT serving as collateral. For the same reason, $\lambda_t > \lambda_{t+1}$ would be hard to justify. Imposing $\lambda_t = \lambda_{t-1}$ is then an accommodating assumption.

occurring in round t can be expressed as the part of the junior tranche that cannot be rolled over, times the original capital relief, i.e.:

$$F_t = N_{t-1}\frac{\Gamma_t}{P_{t-1}}. \tag{6}$$

Following the same dynamics as PCF's capacity and protection gap, F_t reaches its maximum for $P_t = 0$ (i.e. under O1.3), implying $F_t = N_{t-1}$. Instead, it is at its minimum for $P_t = Z_t$ (i.e. under O1.1). Of course, $N_t = N_{t-1} - F_t$.

3.2.3. Stage 3: Transmission

The last stage consists in the transmission of the bank's relief loss to its lending capacity. If the bank is capital constrained, the loss can force it to deleverage. In the first place, the contraction will concern riskier borrowers (Chodorow-Reich, 2014; Iyer et al., 2014; Cingano et al., 2016; Jiménez et al., 2017; Greenstone et al., 2020; De Jonghe et al., 2020), i.e. those making up the junior tranche of the SRT. I define this contraction as:

$$\Lambda_t = \frac{\max[0, F_t - B_t]}{k\omega_t}, \tag{7}$$

where B_t is the bank's excess capital buffer. Λ_t does not include the credit line to the PCF, whose dynamics are determined in stage 2. Importantly, because of the degradation of the loan portfolio due to defaults, $\omega_t > \omega_{t-1}$.[10]

If $F_t < B_t$, the bank's excess capital is sufficient to compensate the lost relief. There is no lending contraction besides for the PCF. For $F_t > B_t$, instead, the reduction in lending to firms in the junior tranche is positive, making default rates increase through the bank lending channel (Bernanke and Gertler, 1995; Kashyap and Stein, 2000; Peek and Rosengren, 2000). The share of lending contraction over the shrinking junior tranche determines the rate of defaults at the beginning of the subsequent round:

$$e_{t+1} = v_{t+1} + d\frac{\Lambda_t}{J_t}, \tag{8}$$

where v_{t+1} is a stochastic component with given mean and variance, and $0 < d < 1$ is the pass-through rate from lending contraction to additional defaults.[11]

Each subsequent round starts with the realisation of the default rate determined at the end of the previous round. The remaining sequence of events follows accordingly, as described above. Of course, this applies as long as the first stage of a round delivers O1.1 or O1.2. If it

---

[10] To model the evolution of the risk weight as the quality of the loan portfolio deteriorates, I assume it increases proportionally to the junior tranche reduction, i.e. $\omega_t = \omega_{t-1}\left(1 + \frac{\Gamma_t}{P_{t-1}}\right)$.

[11] This implies that the realised rate of defaults will be fully exogenous only in round 1. For $t > 1$, it will be composed of an exogenous component, v_t, and an endogenous one, the lending-contraction effect.

delivers O1.3, the sequence of events stops at the end of that round, as a full rollover failure of the SRT has taken place.

4. Simulation

To show how the model could be implemented, a simulation can be of use. I run a two-round loop, for a loan portfolio of initial value 5.000 Mn$ and junior tranche thickness 10% (thus, of value 500 Mn$). These and the remaining values of the parameters are presented in table A.1 in the appendix. As outputs of the simulation, I consider the cumulative (two-stage) (i) protection capacity of the PCF ($P_1 + P_2$), (ii) protection gap ($\Gamma_1 + \Gamma_2$), (iii) net capital relief loss ($F_1 + F_2$) and (iv) lending contraction ($\Lambda_1 + \Lambda_2$). Figure 1 shows the simulated values of (i) and (ii) for different initial default rates (e_1). Figure 2 does the same for (iii) and (iv). Each figure can be divided into three regions: a first region where e_1 is below $\bar{e}$ (8%), a second region between $\bar{e}$ and the default rate such that O1.3 results from the first round (58%, i.e. where μ_1 hits 1), and a third region where e_1 is above this second limit.

Figure 1. Protection capacity and protection gap

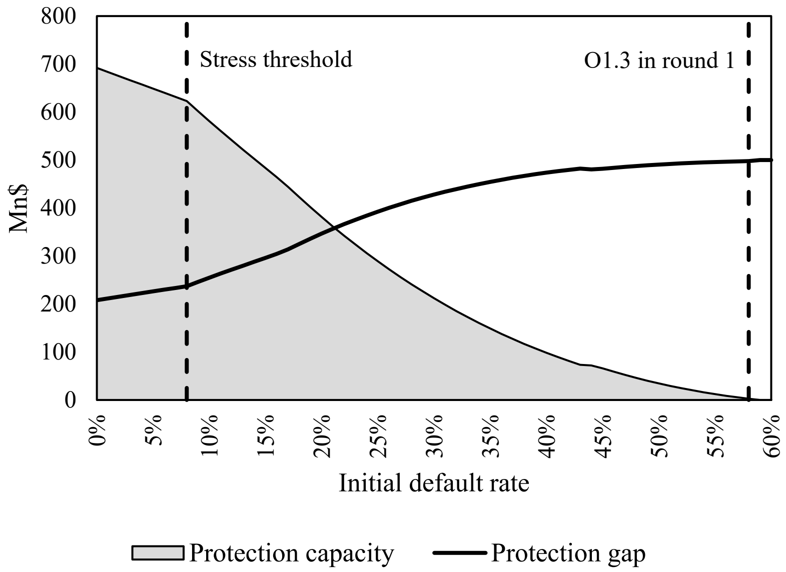

Note. The values shown in the figure are cumulative of the two rounds. The boundaries of the three regions are as given in the text. To ease the presentation, values of e_1 above 60% are not shown.

In the first region of figure 1, the PCF's protection capacity decreases linearly. In either round, this is due – together with the reimbursement of L_{t-1} – to the impact of defaults on J_t, through which e_t enters Z_t and then P_t. It does so linearly and with a negative contribution. In the second region, protection capacity becomes convex, showing the amplifying effect of the PCF's semi-liquid structure. Beyond the stress threshold, redemption requests enter into the picture. The effect of $1 - \mu_t$ on P_t is also linear and decreasing in e_t, but adds to the initial impact. Then, the overall effect on protection capacity becomes quadratic.

For this same mechanism, the protection gap is linear below the stress threshold, and becomes concave in the second region. Concavity is reinforced by the contribution of the second round, which decreases as e_t increases and the size of the junior tranche shrinks. In other words, for high e_1, in the second round there is little to lose protection on. In the third region, the values of either variable stay fixed at their value for e_1 = 58%, 0 and 500Mn$ (the entire junior tranche) respectively, as a full rollover failure has taken place.

Figure 2. Net relief loss and lending contraction

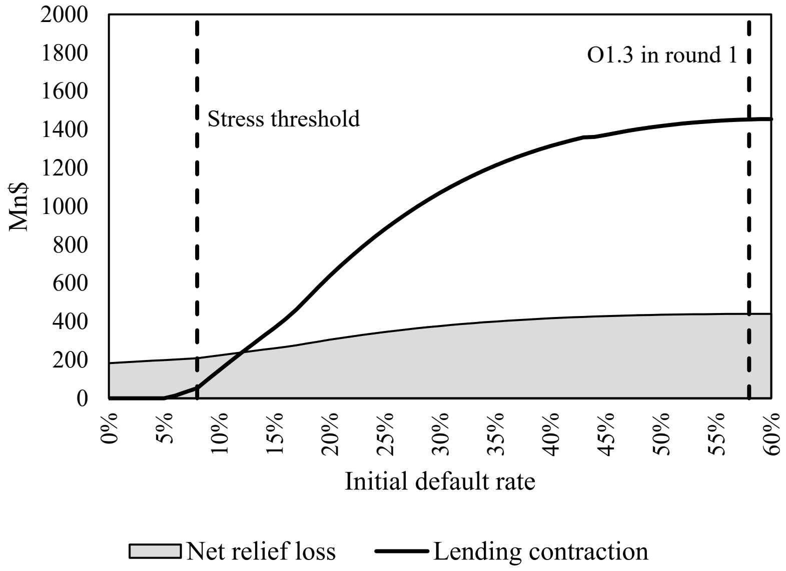

Note. The values shown in the figure are cumulative of the two rounds. The boundaries of the three regions are as given in the text. To ease the presentation, values of e_1 above 60% are not shown.

As regards figure 2, the same mechanism as for the protection gap applies to the net capital relief loss. The dynamics for lending contraction are slightly different, as they also depend on the excess capital buffer: as long as the loss does not exceed the buffer (200 Mn$ in the simulation), there is no lending contraction and the curve stays flat at 0. Once the buffer is exhausted (here, at around 6% of the initial default rate) lending contraction starts increasing. It can do so linearly – as long as the increase takes place in the first region – or in a concave fashion – if the increase happens in the second region.

Importantly, lending contraction can happen before the stress threshold, as redemption requests are amplifiers of a mechanism (devaluations following defaults and reimbursements of credit lines) anyway in place. This simulation, as shown in figure 2, captures both. Lending contraction starts to increase linearly between 6% and 8%, and then moves to a concave shape between 8% and 58%. Still as before, in the third region the two variables stay fixed at their respective values for e_1 = 58%. The net capital relief loss will equal the initial net capital relief. In the simulation, 440 Mn$. Cumulative lending contraction will be 1454 Mn$.

5. Discussion

Back-funding of SRTs and semi-liquid structures of PCFs emerge as the key elements letting defaults on a protected portfolio turn into a perverse loop, ultimately leading to a potential lending contraction to most needing firms. Importantly, the initial default rate is *not* the trigger of the loop. It is the trigger of potential runs. The loop is triggered by the insufficient amount of liquid assets resulting from each round's stage 1, which translates in the protection gap in stage 2. This is caused by three elements: the SRT's back-funding, the part of liquid assets possibly used to meet redemption requests, *and* the initial default rate.

The partial back-funding of SRTs has three effects on the bank's capital relief: first, it reduces the net relief at inception; second, it creates a protection gap in the second stage, which leads a further reduction in both gross and net reliefs if the SRT is partially rolled over. Finally, the rollover of the SRT comes with a new line of financing that, even if smaller, has the same effects. Therefore, as the loop repeats itself in rounds, the financing provided by the bank carries over and over its negative impact until the SRT cannot be rolled over at all. This is still coherent with the business relationship view: since the bank trades current protection for future business opportunities, in any given round of the loop it will be rational for it to extend the line, despite its adverse consequences in the short term.

The semi-liquid structure of the PCF acts as an amplifier of this mechanism. If the PCF was closed-end, the first stage of each round could only deliver O1.1, the most favourable outcome in terms of PCF's liquidity which minimises the protection gap. There would be no retail investors, or they would not be able to request redemptions anyways. These would not add to the depletion of liquid assets necessary to repay the bank's credit line, leaving it available to reduce the protection loss stemming from the latter. Because of defaults and back-funding, the gap would still be present and carried over round after round, but it would grow at a lower rate. Ultimately, this would translate into a slower contraction of bank lending, pushing forward the moment where a given round delivers O1.3 and there is a complete rollover failure.

A similar effect would stem from assuming that the PCF had further liquid assets available, such as dry powder or a positive inflow of bank premia, besides the net amount between SRT repayment and reimbursement of the credit line. Unlike the imposition of a closed-end structure, this would not ensure that O1.1 is realised. The first stage could still deliver the three outcomes depending on the realised default rate, but relatively more liquidity would be available. Yet, as long as no replenishment of such assets is envisioned, this would only slow down the process. Given the highly-illiquid nature of PCFs (FSB, 2026), considering a limited availability of liquid assets is not a strong assumption. Furthermore, the model assumes that neither bank nor PCF can fail. In terms of the outcomes considered in the paper, the failure of the PCF would not be different from O1.3. At the same time, despite their potentially-disruptive effects, assuming a bank's failure because of SRTs would not be realistic. This amount of information can be informative for regulators and supervisors.

In terms of assumptions, the model relies upon a balanced set of conditions on the accommodative and conservative sides. For instance, as just said, the PCF is assumed not to replenish liquid assets on the one hand, or to have other sources of liquidity besides the net value of SRT repayment and reimbursement of the credit line. Yet, the bank is also assumed not to replenish its loan portfolio on the other hand. The former would reduce the protection gap in the second stage and across rounds, while latter would widen it. The bank is also assumed not to replenish its excess capital if part of it is used in the third stage, but, again, the net relief loss has also been partially reduced by the fact that no new loans have entered the bank's portfolio. The back-funding ratio is assumed constant, but the shrinking of the credit line stemming from the devaluations of the SRT is ensured by the size effect of the latter.

In terms of usefulness for private agents, simulating the model on its internal data can help a bank answer two questions: given a default scenario, what is the capital loss from these transactions? Or, running the exercise in reverse, what default rate exhausts a given buffer? This would allow the bank to pre-position the excess capital buffer held against SRT exposures to PCFs. As an example, for the parameters used here, the net capital relief loss begins to erode the bank's capital position once the realised default rate exceeds 6%, providing the bank with an early-warning level at which portfolio monitoring and credit-line covenants should be tightened or, alternatively, capital should start to be rebuilt.

Another option is forward-mapping: to avoid any lending contraction under, say, a 20% default scenario, the bank would need to hold at least 305 Mn$ of buffer against this single transaction, up to 440 Mn$ for defaults at the boundary of full rollover failure. Lending contraction could be used to estimate the business cost to absorb: 53 Mn$ (a bit more than 1% of the reference portfolio) at the stress threshold, 637 Mn$ (12.7%) at a 20% default rate, up to 1.455 Mn$ (29%) at full rollover failure. These figures can inform the bank whether the SRT is worth entering based on its estimation of realised defaults, how much to lend back to the PCF, and how to set internal limits on the concentration of protection bought from PCFs.

For the PCF, the same outputs can determine liquidity and structuring decisions. Estimates of the protection gap can help quantify the liquid reserves the PCF should hold to ensure rollover under stressed conditions. For instance, at a 15% default rate, the cumulative gap reaches 296 Mn$ on a 500 Mn$ junior tranche. An effective liquidity management should cover it, at least partially, by calibrating redemption policies, the amount of dry-powder and other liquidity to hold before entering the transaction, based on the PCF's estimation of possible default rates. For a PCF weighing the commercial appeal of semi-liquid structures against its risk, these estimates provide a quantifiable basis to best assess their strategy.

6. Conclusion

This paper proposes a tractable model of how distress between PCFs and the banking system can emerge and be fuelled through the SRT market. Its main contribution is formalisation. The literature on the bank-NBFI nexus has documented the sector's rapid growth and its reliance on bank financing, while the work on SRTs has shown that non-trivial share of transferred risk circles back to the originating banks. An explicit account of how these two observations interact has been missing. By bringing back-funded leverage and the liquidity concerns of semi-liquid PCFs into a single, replicable framework, the model turns an interconnection that has so far been treated qualitatively into a measurable channel of shock propagation.

The framework is deliberately parsimonious, and this is what can make it useful in practice. It requires little beyond the parameters of a bilateral relationship but, in spirit, behaves like a top-down stress test. Calibrated on proprietary data, it can support a bank quantify the capital it should set aside against the protection it buys from PCFs, identify the default rate at which a given buffer is consumed, and inform its limits on the concentration of such exposures. On the other hand, it can let a PCF translate the same dynamics into liquidity and structuring decisions such as how much dry powder to hold, or how to calibrate redemption policies before entering a transaction (or the business in general).

Finally, for regulators and supervisors, it offers a transparent way to investigate the impacts of interconnections while further ad-hoc modelling and more data become available. The current setting can already point to some corrective actions to be taken. The rise of semi-liquid PCFs has opened the door to investors that are excessively sensitive to market fluctuations relative to the long-term horizon that these investments should envision. Back-funding of leverage and the semi-liquid structure of PCFs amplify isolated shocks, accelerating their transformation into self-feeding contractions in credit.

As shocks per se cannot be avoided, acting on funds' design should be the way forward to mitigate their effects. Limiting back-funding of leverage could be an option, but an ambitious one given the liquidity needs of PCFs as well as the protection requests of banks. Acting on the open-end vs closed-end of PCFs is instead a viable. The AIFM2 Directive in the European Union is a good example in this sense, preventing loan-originating funds to be open-end unless they can prove that an adequate liquidity management system is in place. Efforts to replicate the spirit of this Directive should be encouraged, especially in those jurisdictions, like the US, where the growth of private credit has been more prominent than in Europe.

Appendix

Even though based as much as possible on background regulations and literature, these values should of course be considered as indicative, and cannot replace the level of information that a specific bank, PCF or regulator could input into the model. Yet, they serve the scope of showing how the model can be implemented and interpreted.

Table A.1. Simulation parameters

Parameter	Value	Unit	Background
A_0	5000	Mn$	IACPM (2025)
ω_0	100	%	BCBS (2019, 2026)
$\widehat{\omega}_0$	20	%	BCBS (2019, 2026)
j_0	10	%	IACPM (2025)
k	11	%	BCBS (2019), Fed (2023)
λ	20	%	Acharya et al. (2025); Fillat et al. (2025); Osberghaus and Schepens (2026)
α	200	%	Falato et al. (2021)
$\hat{e}$	8	%	FT (2026a,b,c,d)
B_1	200	Mn$	Fed (2024); EBA (2024,2025)
d	5	%	Bernanke and Gertler (1995), Kashyap and Stein (2000), Peek and Rosengren (2000)
e_1	varying	%	FT (2026a,b,c,d)

Notes. For the sake of the simulation, v_2 is assumed at 7%. The risk weight applicable to equation 7 is capped at 1.5. Background documents and literature do not necessarily provide an exact estimate of the associated parameter, especially for non-regulatory ones, but can be informative to select one in the absence of further data.

References

Acemoglu, D., Ozdaglar, A. and A. Tahbaz-Salehi (2015). "Systemic risk and stability in financial networks". *American Economic Review*, 105(2), 564-608.

Acharya, V. V., Gopal, M., Jager, M. and S. Steffen (2025). "Shadow always touches the feet: Implications of bank credit lines to non-bank financial intermediaries". *NBER Working Paper No. 33590*.

Acharya, V. V., Cetorelli, N. and B. Tuckman (2024). "Where do banks end and NBFIs begin?". *NBER Working Paper No. 32316*.

Acharya, V. V., Schnabl, P. and G. Suarez (2013). "Securitization without risk transfer". *Journal of Financial Economics*, 107(3), 515-536.

BCBS (2026). "Synthetic risk transfers". Basel Committee on Banking Supervision.

BCBS (2019). "The Basel framework". Basel Committee on Banking Supervision.

BdF (2025). "Financial stability report". Banque de France, December.

Berg, T. and J. Hoon Lee (2026). "Measuring counterparty exposures to private credit". *OFR Brief Series No. 26-02*.

Bernanke, B. and M. Gertler (1995). "Inside the black box: The credit channel of monetary policy transmission". *Journal of Economic Perspectives*, 9(4), 27-48.

Bernstein, S., Colonnelli, E. and B. Iverson (2018). "Asset allocation in bankruptcy". *Journal of Finance*, 74(1), 5-53.

Block, J. H., Soo Jang, Y., Kaplan, S. N. and A. Schulze (2024). "A survey of private debt funds". *Review of Corporate Finance Studies*, 13(2), 335-383.

Bloomberg (2026a). "Private credit's CLO machine ramps up in push to raise more cash". Bloomberg, 2nd April.

Bloomberg (2026b). "Why you should care about private credit". Bloomberg, 6th April.

Brunnermeier, M. K. and L. H. Pedersen (2009). "Market liquidity and funding liquidity". *Review of Financial Studies*, 22(6), 2201-2238.

Cai, F. and S. Haque (2024). "Private credit: Characteristics and risks". *FEDS Notes*.

Cetorelli, N., Landoni, M. and L. Lu (2023). "Non-bank financial institutions and bank's fire-sale vulnerabilities". *New York Fed Staff Report No. 1057*.

Cetorelli, N. Duarte, F. M. and T. M. Eisenbach (2016). "Are asset managers vulnerable to fire sales?". *Liberty Street Economics*.

Chernenko, S., Ialenti, R. and D. Scharfstein (2025). "Bank capital and the growth of private credit". Harvard Business School.

Chernenko, S., Erel, I. and R. Prilmeier (2022). "Why do firms borrow directly from nonbanks?". *Review of Financial Studies*, 35(11), 4902-4947.

Cortes, F., Fernandez Dionis, G., Li, Y., Ramirez, S. and X. Zhang (2025). "Recycling risk: Synthetic risk transfers". *IMF Working Paper No. WP/25/200*.

Couaillier, C., Reghezza, A., Rodriguez d'Acri, C. and A. Scopelliti (2025). "How to release capital requirements in an economic downturn? Evidence from euro area credit register". *Journal of Financial Intermediation*, 63.

EBA (2025). "Risk assessment report". European Banking Authority, June.

EBA (2024). "Risk assessment report". European Banking Authority, July.

ESRB (2025). "EU non-bank financial intermediation risk monitor". European Systemic Risk Board, September.

Falato, A., Goldstein, I. and A. Hortaçsu (2021). "Financial fragility in the Covid-19 crisis: The case of investment funds in corporate bond markets". *Journal of Monetary Economics*, 123, 35-52.

Fed (2024). "Dodd-Frank Act stress test". Board of Governors of the Federal Reserve System.

Fed (2023). "Frequently asked questions on Regulation Q". Board of Governors of the Federal Reserve System.

Fillat, J. L., Landoni, M., Levin, J. D. and J. C. Wang (2025). "Could the growth of private credit pose a risk to financial system stability?". *Boston Fed Current Policy Perspectives*, 25-8.

FSB (2026). "Report on vulnerabilities in private credit". Financial Stability Board, May.

FT (2026a). "Who cares if private credit goes kaput?" Financial Times, 18th March.

FT (2026b). "Morgan Stanley and Cliffwater limit private credit withdrawals". Financial Times, 11th March.

FT (2026c). "JP Morgan marking down loan portfolios of private credit groups". Financial Times, 11th March.

FT (2026d). "Private credit retail honeypot is getting less sweet". Financial Times, 5th March.

Goldstein, I., Jiang, H. and D. T. Ng (2017). "Investor flows and fragility in corporate bond funds". *Journal of Financial Economics*, 126(3), 592-613.

Goldstein, I. and A. Pauzner (2005). "Demand-deposit contracts and the probability of bank runs". *Journal of Finance*, 60(3), 1293-1327.

González, F. and C. M. Triandafil (2023). "The European significant risk transfer securitization market". *ESRB Occasional Paper No. 23.*

IACPM (2025). "Global SRT bank survey 2016-2024". International Association of Credit Portfolio Managers.

IMF (2024). "Global financial stability report". International Monetary Fund, October.

Ivashina,V. (2025). "The role of private debt in the financial ecosystem". *NBER Working Paper No. 34426*.

Jiménez, G., Ongena, S., Peydró J-L. and J. Saurina (2017). "Macroprudential policy, countercyclical bank capital buffers, and credit supply: Evidence from the Spanish dynamic provisioning experiments". *Journal of Political Economy*, 125(6).

JPM (2026). "Private credit under the microscope – Separating headlines from fundamentals". JP Morgan Private Bank, March.

Kashyap, A. and J. Stein (2000). "What do a million observations on banks say about the transmission of monetary policy?". *American Economic Review*, 90(3), 407-428.

Matvos, G., Piskorski, T. and A. Seru (2026). "Private credit, balance sheets and financial stability". *NBER Working Paper No. 34991*.

Osberghaus, A. and G. Schepens (2026). "Synthetic, but how much risk transfer?". *ECB Working Paper No. 3210*.

Peek, J. and E. Rosengren (2000). "Collateral damage: Effects of the Japanese bank crisis on the real activity in the United States". *American Economic Review*, 90(1), 30-45.