

Platform Finance

Yingjie Huang^{*}

First Draft: November 1, 2025

This Draft: October 15, 2026

Abstract

Platforms increasingly partner with banks to finance users, but partnership structures differ across platforms and over time. We model a capital-constrained, informed platform that partners with a capital-rich, uninformed bank, choosing between intermediation (on-lending bank funds) and information provision (sharing borrower data). The tension is that the platform internalises network externalities from on-platform trade but the bank does not, and the degree of this conflict is endogenous to the partnership choice. Our main result is a regime switch in borrower pool size: information provision, if sufficiently credible, is optimal for small pools and intermediation for large pools, and the switching threshold increases with this credibility. Buy Now, Pay Later (BNPL) buyer financing makes information provision in merchant financing more attractive when the BNPL provider places greater value on enabled trades.

^{*}University of Amsterdam and Tinbergen Institute. Email: y.huang4@uva.nl. I gratefully acknowledge helpful comments and suggestions from Matthieu Bouvard, Catherine Casamatta, Alexandre de Cornière, Alexander Guembel, Natalya Martynova, Isabelle Michalski-Roland, Alessio Ozanne, Enrico Perotti, Sébastien Pouget, Magdalena Rola-Janicka, Spyros Terovitis, Vladimir Vladimirov, Yenan Wang, and all participants and discussants at the EEA-ESEM (Dublin), Econometric Society North American Summer Meeting (Atlanta), Bank of Italy-Bocconi-EIEF-CEPR Conference, IBEFA Summer Meeting (Denver), Asian Finance Association Annual Meeting (Seoul), Edinburgh FinTech Conference, Benelux Corporate Finance Days (Ghent), Bundesbank-Frankfurt Summer School, and University of Amsterdam.

1 Introduction

Platforms have emerged as major credit providers to their users.¹ Credit from BigTech platforms, such as Amazon and Apple, surpassed all other forms of FinTech credit in 2018 and reached \$700 bn in 2020, with recent growth rates three times those of bank lending in developing markets ([Barakova et al., 2024](#); [BIS, 2019](#)). Non-BigTech platforms show similar growth. Shopify originated \$4.2 bn in financing for its sellers in 2025, a 40% year-on-year increase.² This rapid credit expansion addresses a significant market need,³ so understanding how platforms organise lending is economically relevant.

Platforms typically lend in partnership with banks, but the structure of these partnerships is notably heterogeneous across platforms and changes over time even for a given platform. For example, Amazon initially lent to its merchants from its own balance sheet (2011–2016), then gradually moved to intermediating bank funding (2016–2023), and more recently stopped underwriting loans itself and began sharing merchant data with partner banks that make the lending decisions (2024–).⁴ Across platforms, Block lends to its merchants using capital raised from external lenders, whereas DoorDash refers its merchants to a partner lender that decides whether to finance their cash needs. While it is well understood that partnerships arise because platforms possess superior information and banks have cheaper access to capital ([Doerr et al., 2025](#)), we lack a theory to explain why these partnerships take different forms and how the choice of arrangement varies with platform characteristics and economic environments.

Platform financing differs from standard uninformed lending problems with a separation of information and capital. This is because platforms internalise network externalities from on-platform trade, whereas their bank partners do not. Indeed, financing enables a merchant to produce and, if production is successful, to trade with buyers on the platform. The service fees that buyers and the merchant pay on these trades accrue to the platform rather than the bank. The platform would therefore like to finance more merchants than the bank is willing to fund. This conflict of interest is the central friction that any partnership must resolve. More importantly, we show that the degree of this conflict is itself endogenous to the chosen partnership form, motivating a joint analysis of lending decisions and partnership arrangements.

We develop a theory of how a platform optimally chooses between two partnership

¹By platforms, we mean firms whose core business is to facilitate transactions between two sides of a market, such as online marketplaces, rather than peer-to-peer lending platforms dedicated to facilitating lending.

²See “[Shopify Inc. Form 10-K, 2025 \(SEC\)](#)”.

³In 2024, 46% of U.S. SMEs obtained less than half of the credit they sought from traditional lenders. This continues a multi-year trend of shifting applications from banks (–11%) to online creditors (+7%). See “[2025 Report on Employer Firms](#)” by Federal Reserve Banks.

⁴See [Fortune \(7 March 2024\)](#) on the start of Amazon’s lending in 2011 and the end of its own lending in March 2024, and “[Amazon has partnered with Bank of America for its lending program](#)” on the revolving credit facility that Amazon obtained in October 2016.

arrangements with a bank: (i) financial intermediation, where the platform on-lends bank funds to borrowers; and (ii) information provision, where the platform shares its data on borrowers with the bank, which lends to them directly.⁵ Our main result is that the choice depends on the size of the platform’s borrower pool and on the credibility of its data sharing technology. With a large pool of potential borrowers, the platform prefers intermediation; with fewer borrowers, it prefers information provision if sufficiently credible; and the pool size at which it switches increases with this credibility.

In our model, buyers have private valuations for a good, and merchants have privately known productivity, or type, for the good and require external capital for an upfront investment. The platform charges transaction fees on both sides for matchmaking, is capital constrained, and learns merchant types at an interim stage, whereas the bank has abundant capital but no information on merchants. We model intermediation as a credit line from the bank to the platform, which then lends to merchants. We model information provision as a pre-designed reporting rule that specifies the report the platform sends to the bank for each level of merchant productivity. The platform is bound by this rule only with some probability and can otherwise override the rule without the bank observing the deviation. We call this probability the credibility of data sharing.⁶ The sequence of events is as follows. The platform first sets the fees, then agrees on a partnership arrangement with the bank and, under information provision, announces its reporting rule. The platform then learns merchant types, the partnership arrangement is implemented, financing decisions are made, and production and trade occur.

We establish two benchmarks to guide the analysis. First, a monopoly benchmark assumes that the platform has sufficient capital and requires no partnership. The platform screens merchants and finances all those whose productivity exceeds a financing cutoff, including some of intermediate productivity. Loans to these intermediate merchants make a financial loss but are justified by the network benefits they generate. Second, a bank-lending benchmark illustrates the market failure of uninformed lending without a partnership. Unable to screen, the bank finances either all merchants or none, depending on whether average productivity exceeds its break-even cutoff, which is higher than the platform’s because the bank does not internalise network benefits. By lowering its merchant fee, the platform can pass these benefits to the bank through higher loan repayments, so the bank lends whenever average productivity exceeds the platform’s cutoff. We focus on the case in which average productivity is below this cutoff, so the bank does not lend, although some merchants are worth financing for the platform. These benchmarks

⁵These are the two major forms of platform-bank partnership (Wang & Charoenwong, 2025). A third form, co-lending, in which the platform and the bank jointly finance borrowers with a syndicated loan, is mainly observed in emerging markets such as China and India (Liu et al., 2022).

⁶Banks fund the lending of platforms through credit or liquidity facilities (Barakova et al., 2024); see footnote 4 for Amazon’s credit line. When platforms share merchant data with lenders, they do so through pre-designed interfaces such as APIs (see “Amazon steps up financing to fuel seller growth”), but can still alter the data sent on individual merchants. More discussions are provided in Appendix B.

provide upper and lower bounds on platform profit, and we examine whether intermediation and information provision improve on uninformed bank lending and how closely they approach the monopoly benchmark.

The intuition for our main result is as follows. Under intermediation, the platform and the bank agree on which merchants to finance first, as both prefer more productive merchants, but not on how many, since only the platform values the trades that financed merchants generate. Thus, while there is no lending without a partnership, the bank is willing to provide the platform with an endogenously capped credit line, knowing that the funds will go to the top-ranked merchants, whose expected productivity is sufficiently high. In a larger pool, a merchant's rank reveals more about its productivity, so selection becomes increasingly effective and profit per merchant rises towards the monopoly level. However, the credit line is fixed before the platform learns merchant types. The number of merchants worth financing may then be larger or smaller than the credit line allows. In the first case, some merchants worth financing remain unfunded; in the second, some merchants not worth financing are funded. Since the platform finances the best merchants first, these mistakes involve only merchants close to the financing cutoff. In a larger pool, there are more such mistakes, but each is less costly because the affected merchants lie closer to the cutoff. The two effects offset each other, so the total loss relative to the monopoly benchmark remains bounded and converges to a positive constant.

Under information provision, by contrast, the bank decides whether to finance each merchant based on the platform's report. When setting its fees, the platform bears the cost of lending. Since the bank must break even, recommending weaker merchants requires a lower merchant fee. The reporting rule therefore recommends exactly the merchants that the platform would finance as a monopolist. Once fees are set, however, the platform earns fees whenever a merchant is financed and bears no credit risk, so it would like every merchant to be financed. The conflict with the bank is therefore stronger than under intermediation, as it concerns not only how many merchants to finance but also which merchants. This conflict is disciplined only by the platform's commitment to its reporting rule. Indeed, whenever the platform can override the rule, it recommends the merchant regardless of productivity. Provided that the rule is sufficiently credible, recommended merchants are still productive enough on average for the bank to lend, which improves on the market shutdown under the bank-lending benchmark. Misallocation now takes the form of overfinancing. Since the reporting rule does not depend on pool size, profit per merchant is constant.

Together, these results mean that, in a small pool, intermediation gains little from screening, and with a single merchant it cannot improve on bank lending at all. Information provision, if sufficiently credible, can instead restart lending even then. In a large pool, by contrast, the total misallocation loss under intermediation remains bounded, while the overfinancing loss under information provision grows in proportion to pool size,

giving rise to the regime switch. The switching threshold increases with the credibility of data sharing, because greater credibility reduces overfinancing under information provision. With full credibility, information provision attains the monopoly outcome and dominates intermediation for every pool size. When data sharing is sufficiently credible, the threshold is also higher if trade is more valuable or projects are cheaper, since financing weaker merchants then costs less.

The two arrangements also differ in loan quality and welfare. In large pools, loans originated under intermediation default less often, because information provision overfinances weak merchants regardless of pool size. The platform’s choice, however, need not maximise welfare. Since the platform does not fully capture the surplus that buyers obtain from trade, it ignores part of the value created by a financed merchant. Even the monopoly benchmark therefore finances too few merchants, so overfinancing under information provision can raise welfare in large pools, even when the platform itself prefers intermediation.

We extend the model in two directions. First, we introduce buyer financing. Embedded finance products targeted at buyers have grown rapidly, with Buy Now, Pay Later (BNPL) as a leading example.⁷ Since BNPL stimulates trade from the demand side, we examine how it affects the platform’s choice on the supply side. We find that partnering with a BNPL provider makes information provision more attractive when the provider places greater value on the trades it enables, for example through advertising revenue, provided that data sharing is sufficiently credible. The provider pays the buyer fee on buyers’ behalf in exchange for a commission from the platform, and the more it gains from each trade, the lower the commission it is willing to accept. Each trade then becomes more valuable to the platform, so financing weaker merchants costs less and the loss under information provision falls. Second, we allow merchants to sell to buyers through their own sales channels, which limits what the platform can collect from them. A platform collects its loans by deducting repayments from merchants’ platform payouts, which direct sales bypass. Cheaper direct sales therefore favour information provision if banks can collect repayment from off-platform sales, but reduce credit under both arrangements if they cannot.

Our theory delivers several testable implications. First, in markets with weaker formal institutions, banks are less able to verify information provided by platforms, lowering the credibility of data sharing. Platforms should therefore switch to intermediation at a smaller pool size. Second, the platform loan securitisation market provides a direct test of our prediction on loan quality. Asset-backed securities (ABS) backed by merchant loans originated under intermediation by platforms with large pools should carry lower

⁷Embedded finance refers to the integration of financial services into non-financial platforms. Its market size is anticipated to reach \$7.2trn by 2030, equivalent to 5% of world GDP. See [“Why Banks Are Embracing Embedded Finance To Stay Competitive”](#) and IMF World Economic Outlook Database.

credit risk, reflected in tighter yield spreads and lower default rates, than those backed by loans originated under information provision. Third, our theory predicts a life cycle in platform credit strategy. Platforms with sufficiently credible data sharing initially rely on information provision, switch to intermediation as their merchant pool grows beyond a threshold, and may return to information provision as the credibility of data sharing improves. Amazon’s recent move from underwriting seller loans funded by a bank credit facility to sharing seller data with partner lenders, which then lend directly, is consistent with this last step. Finally, among platforms with sufficiently credible data sharing and merchant pools of the same size, those that offer BNPL through providers that value the trades they enable should rely more on information provision.

Our theory also has regulatory implications. First, rules that make platform data harder to manipulate, such as auditable data-access interfaces and merchants’ right to share their own platform data with lenders, increase the credibility of data sharing. These rules reduce overfinancing and shift partnerships towards information provision. Second, rules that impose the same arrangement on all platforms have different effects depending on pool size. Restricting data sharing with lenders forces platforms with small pools into intermediation, which can shut down lending. Preventing platforms from lending instead forces platforms with large pools into information provision, which produces lower-quality loans but can raise welfare. Third, rules that allow merchants to steer platform customers to their own sales channels can reduce merchant credit unless lenders can collect repayment across sales channels.

Related literature Closest to our work is theory on platform lending with some level of bank cooperation. [Wang & Charoenwong \(2025\)](#) build a dynamic model in which an informed BigTech originates loans to low-credit merchants to build up their credit histories, then sells the data so a capital-rich bank finances repeat borrowing. [Boualam & Yoo \(2022\)](#) study when banks and FinTechs partner, compete, or coexist, and show that declining borrower captivity shifts banks from competition towards coexistence and potentially full partnership, with this shift strengthened by banks’ funding-cost advantage, FinTechs’ enforcement advantage, and low partnership costs.

Relative to both papers, we model network externalities as a core friction for any cooperation to resolve and a defining platform feature ([Rochet & Tirole, 2006](#)), whereas their mechanisms hinge mostly on differences in the cost of capital and abstract from the platform’s core business. Compared to [Wang & Charoenwong \(2025\)](#), our framework pins partnership form to platform primitives rather than to building credit histories, and we compare the platform’s choice of partnership form with the one that maximises welfare. Relative to [Boualam & Yoo \(2022\)](#), we do not impose borrower captivity and we consider richer partnership forms.

Two other theoretical papers study platform lending solely as competition with banks.

[Bouvard et al. \(2022\)](#) analyse a bank-platform duopoly where platforms and banks have different strengths in monitoring merchants. They show that the platform can profit by adjusting transaction fees upwards and lending to riskier merchants that banks reject, thereby achieving implicit price discrimination against safe merchants with bank financing. [Li & Pegoraro \(2022\)](#) show that, by controlling access to its marketplace, a platform can enforce partial repayment and lend to risky merchants, but good borrowers then turn to banks to avoid this enforcement, inducing welfare losses. Our extension with direct sales channels shows that this enforcement advantage is limited when merchants can move sales off the platform, a form of marketplace leakage ([Hagi & Wright, 2024](#)). In contrast to these papers, we study the cooperative arrangements widely observed in reality and endogenise whether and how banks and platforms partner.

More broadly, our work relates to research on alternative lenders, including shadow banks and FinTechs, that extend credit using proprietary data. [Parlour et al. \(2022\)](#) study bank-FinTech competition in payments and show that FinTech entry weakens banks' use of payment data for credit scoring, reshaping loan pricing and offers; the net effect broadens inclusion but can disadvantage bank-loyal consumers. A strand of contemporary literature similarly analyses interactions between alternative lenders and banks, again largely through a competitive lens ([Agur et al., 2025](#); [Ahnert et al., 2025](#); [Chu & Wei, 2024](#); [He et al., 2023](#); [Vives & Ye, 2025](#)). Compared to this line of research, we introduce network externalities exclusively benefiting platforms as creditors and, most importantly, endogenise the platform's decision between different cooperative partnership models. Our platform also resembles a supplier that extends trade credit to its customers. A supplier may know its customers better than banks do ([Biais & Gollier, 1997](#)) or be less exposed to the diversion of funds ([Burkart & Ellingsen, 2004](#)), and it profits from the sales it finances. Unlike a supplier, however, our platform has no capital to lend, so the question is how it combines its information with the capital of a bank.

On the technical side, our paper relates to the delegation-versus-communication literature in organisational economics (see [Mookherjee \(2006\)](#) for a review). [Dessein \(2002\)](#) studies an uninformed principal facing an informed agent and shows that delegation is optimal when the agent's bias is small relative to environmental uncertainty: strategic communication (cheap talk) distorts information, whereas delegation preserves it at the cost of control. In our setting, intermediation corresponds to delegation and information provision to communication. We depart by introducing an endogenous conflict between the bank (principal) and the platform (agent) driven by network externalities, which is stronger under communication than under delegation as explained above. We also go beyond cheap talk and model communication as persuasion with limited commitment ([Lipnowski et al., 2022](#); [Min, 2021](#)), which nests cheap talk and Bayesian persuasion ([Kamenica & Gentzkow, 2011](#)). We further examine how the production environment and platform primitives shape the optimal choice between delegation and communication.

Finally, our focus on and modelling of platform financing is supported by a growing body of empirical literature that documents its rapid expansion as a supplement to bank credit (Cornelli et al., 2023; Frost et al., 2020; Liu et al., 2022); platforms’ restricted access to cheap capital along with an advantage over banks in possessing and using proprietary borrower data in lending decisions (Buchak et al., 2018; Gambacorta et al., 2022; Gertler et al., 2021; Hau et al., 2019); and their real economy effects, such as fostering entrepreneurial growth and enhancing firm stability (Chen et al., 2022; De Fiore et al., 2024; Hau et al., 2024).

2 Model

We consider an economy with four types of agents—buyers, merchants, a platform, and a bank—interacting over three dates $t \in \{0, 1, 2\}$. All agents are risk-neutral and do not discount payoffs over time.

Buyers There is a unit mass of buyers indexed by $b \in (0, 1)$. Each buyer b privately observes a valuation $v_b \sim \text{i.i.d. } F(v)$ for a tradable good. The cdf F has support $[0, \bar{v}]$, where we allow $\bar{v} \rightarrow \infty$, and density f .

Merchants There are M merchants, where $M \in \mathbb{N}_{>0}$, indexed by $m = 1, 2, \dots, M$.⁸ Each merchant requires a fixed investment I to produce the tradable good. Production succeeds with probability θ_m , which represents the merchant’s productivity or credit type. Successful production generates per-transaction net surplus $\pi > 0$ for the merchant, while unsuccessful production generates zero.⁹ Types θ_m are independently drawn from $G(\theta)$ with log-concave density $g(\theta)$ on $[0, 1]$ and are privately observed by merchants. All merchants are cash-constrained and cannot self-finance I .

We impose the following standard regularity condition to ensure that relevant optimisation problems are well behaved.

Assumption 1. $F(v)$ and $G(\theta)$ are twice continuously differentiable on their supports. The hazard rate $\frac{f(v)}{1-F(v)}$ is non-decreasing.

Platform A monopolistic platform intermediates trade between buyers and merchants. It charges per-transaction fees P^M and P^B on the merchant and buyer sides respectively.¹⁰

⁸We model buyers as a unit mass for streamlined analysis, since that side is not the focus. By contrast, we keep the number of merchants countable at M to study how the platform’s behaviour varies with this number. The continuum case arises as the limit $M \rightarrow \infty$.

⁹We abstract from the price of the good, so that π is the merchant’s margin per transaction and buyer b ’s valuation v_b is net of this price.

¹⁰Restricting attention to per-transaction fees is without loss of generality in the current setting (Rochet & Tirole, 2006). A membership fee is equivalent to a transaction fee times the equilibrium

The platform does not possess internal capital for financing merchants. At an interim stage, however, it receives a perfectly informative signal of all merchant types.¹¹

Bank A single bank provides external financing for merchants with total investable capital MI . The bank has no information about merchant types beyond what may be revealed by the platform.

Merchant financing Merchant financing takes the standard debt form $\{I, D\}$, where I is the loan and D the repayment conditional on success. With binary production outcomes, restricting attention to debt contracts is without loss of generality.

Following standard accounting practice, we assume the per-transaction fee P^M , as an operating expense, settles prior to any debt service. Hence, limited liability requires $D \leq N^B(\pi - P^M)$, where N^B denotes the number of buyers on the platform and thus determines the number of transactions conducted by a financed merchant.¹² We also impose $D \geq 0$ to prevent inefficient destruction of output (Innes, 1990). We abstract from merchant moral hazard.

Platform-bank arrangements We model two possible arrangements that combine the platform’s information and the bank’s capital.

Intermediation arrangement. The platform borrows from the bank and on-lends to merchants without transmitting type information. Since we abstract from other uses of funds by the platform, the platform-bank contract is a credit line nI , a discrete multiple of the per-merchant investment I , where $n \in \{0, 1, \dots, M\}$ is the maximum number of merchants that can be financed, and $n = 0$ means no credit line. The bank funds the loans made under this credit line and has a claim only on their repayments made by the financed merchants.¹³ Without loss of generality, we assume that the platform passes these repayments to the bank in full. Indeed, a merchant pays both its loan repayment and the merchant-side fee P^M only when it succeeds. Any repayment retained by the platform could therefore instead be collected through a higher P^M , without changing any party’s payoff.

number of trades. In Appendix D, we consider an extension in which merchants and buyers can trade through a direct channel that bypasses the platform.

¹¹All results remain qualitatively robust under imperfect platform information, provided the distribution of success probabilities conditional on its information satisfies our assumptions. In that case, θ_m is interpreted as merchant m ’s conditional success probability and G as its distribution. Corollary 4 then studies the effects of a more informative platform signal, represented by a mean-preserving spread in G .

¹²A merchant that produces successfully trades with every buyer on the platform since $v_b \geq 0$ and $\pi > 0$. We treat N^B as contractible. In practice, platforms maintain time-stamped user and transaction logs for billing and compliance. A court or arbitrator can verify N^B ex post from audited platform records or payment-processor statements.

¹³In practice, the bank’s loan to the platform is earmarked for merchant financing, not general operations, and is secured by the merchant loans it funds. For example, [Amazon’s revolving credit facility](#) was “secured by certain seller receivables”.

Information provision arrangement. The platform shares merchant-type information with the bank, which lends directly to merchants. Let S denote the report space. Since merchant types lie in $[0, 1]$, we set $S = \mathcal{B}([0, 1])$, the Borel σ -algebra, without loss of generality. Before observing merchant types, the platform announces a reporting rule, a measurable mapping $\sigma : [0, 1] \rightarrow S$, where $\sigma(\theta)$ is the report for a type- θ merchant. Each report is therefore a subset of $[0, 1]$.

We introduce a commitment parameter $\alpha \in [0, 1]$, which measures the credibility of data sharing. For each merchant independently, the bank receives the report specified by the reporting rule with probability α ; with probability $1 - \alpha$, the platform can override the rule and send any report of its choice. The bank observes the report but not whether the rule was overridden. Thus, $\alpha = 0$ corresponds to cheap talk (Crawford & Sobel, 1982), $\alpha = 1$ to Bayesian persuasion (Kamenica & Gentzkow, 2011), and $0 < \alpha < 1$ to persuasion with limited commitment (Lipnowski et al., 2022; Min, 2021). The Bayesian-rational bank decides whether to offer each merchant a financing contract on the basis of the (possibly null) report on that merchant. We discuss this technology further in Appendix B.

Timing At $t = 0$, buyer and merchant types are drawn independently and privately observed; the platform sets per-transaction fees $\{P^M, P^B\}$.¹⁴ At $t = 1a$, the platform and the bank decide whether to sign a contract for a credit line nI . If no contract is signed, the platform announces a reporting rule σ . At $t = 1b$, the platform observes a perfectly informative signal of all merchant types. If a platform-bank contract is in place, the platform offers financing contracts to merchants. Otherwise, the bank receives a report on each merchant, as prescribed by the reporting rule with probability α and chosen by the platform with probability $1 - \alpha$. At $t = 1c$, in the absence of a platform-bank contract, the bank offers financing contracts to merchants. At $t = 2$, buyers decide whether to join the platform. Financed merchants produce, trades occur, and outcomes and payoffs are realised. Table 1 summarises the sequence.

2.1 Monopoly benchmark

We first study a benchmark in which the platform is assumed to hold sufficient internal capital MI to finance merchants itself, thereby monopolising both matchmaking and financial services, and the bank is redundant.¹⁵ The platform then finances merchants

¹⁴We assume the platform commits to the fee schedule $\{P^M, P^B\}$ once posted. This is a standard assumption in static models of multi-sided markets where platforms post prices before users participate (Armstrong, 2006; Hagiu & Wright, 2015; Rochet & Tirole, 2006). It is further grounded in legal and regulatory constraints, reputational stakes with both creditors and users, and non-trivial adjustment costs for platforms. We discuss this assumption in Appendix E.

¹⁵The bank cannot profitably compete in this benchmark. Indeed, the platform can set $P^M = \pi$, which leaves no repayment for a bank loan ($D \leq N^B(\pi - P^M) = 0$) and, as shown below, does not affect

Date $t = 0$
• Buyer and merchant types are independently drawn and privately observed. • The platform announces per-transaction fees $\{P^M, P^B\}$.
Date $t = 1a$
• The platform and the bank decide whether to enter into a credit line contract nI. • If no contract is signed, the platform announces a reporting rule σ.
Date $t = 1b$
• The platform observes a perfectly informative signal of all merchant types. – <i>Intermediation</i>. If a platform-bank contract was signed, the platform offers financing contracts to merchants, and $t = 1$ ends. – <i>Information provision</i>. If a platform-bank contract was not signed, the bank receives a report on each merchant, specified by σ with probability α and chosen by the platform with probability $1 - \alpha$.
Date $t = 1c$
• <i>Information provision</i>. The bank offers financing contracts to merchants.
Date $t = 2$
• Buyers decide whether to join the platform. Production occurs, outcomes realise, trade is executed, and payoffs are settled.

Table 1: Timeline of events

above a threshold determined by their total contribution to platform revenue, including transaction fees generated by network externalities. It therefore finances high-quality borrowers, but may also finance some lower-quality borrowers whose network externalities offset expected loan losses.

$t = 2$: Buyers joining decision Buyer b joins the platform if and only if $v_b \geq P^B$, so the number of buyers on the platform is $N^B = 1 - F(P^B)$. Indeed, buyers evaluate each transaction independently and face no fixed cost of joining, so there are no cross-side network externalities from merchants to buyers. We adopt this formulation for tractability and relax it in Appendix C. There are, however, cross-side network externalities from buyers to merchants, as shown below.

$t = 1b$: Optimal merchant financing If financed under the debt contract $\{I, D\}$, merchant m expects a payoff of $\theta_m [1 - F(P^B)](\pi - P^M) - D$, which takes into account the buyer participation level. As the sole financier, the platform sets $D = [1 - F(P^B)](\pi - P^M)$, leaving the merchant with zero expected surplus.

Financing a type- θ merchant then yields the expected loan profit $\theta[1 - F(P^B)](\pi - P^M) - I$ and the expected fee income $\theta[1 - F(P^B)](P^M + P^B)$. The total payoff from financing the merchant is thus non-negative if and only if the merchant's type exceeds a

the platform's payoff.

financing cutoff $\underline{\theta}^{\text{PF}}(P^B)$:

$$\theta \geq \frac{I}{[1 - F(P^B)](\pi + P^B)} \equiv \underline{\theta}^{\text{PF}}(P^B). \quad (1)$$

The intuition for the threshold is twofold. First, the merchant-side fee P^M is inessential, since the monopolist captures the full surplus π and is indifferent to whether it collects it as a fee or as a loan repayment. Second, the buyer-side fee P^B is essential, as it represents the internalised cross-side network externalities: whenever a financed merchant produces successfully, each resulting trade yields P^B to the platform, relaxing the financing constraint.

$t = 0$: Optimal platform pricing Since P^M is irrelevant, the platform chooses the buyer-side fee P^B to maximise its total profit from financed merchants:

$$\max_{P^B} M \int_{\underline{\theta}^{\text{PF}}(P^B)}^1 \left[\theta [1 - F(P^B)](\pi + P^B) - I \right] g(\theta) d\theta, \quad (2)$$

which yields $P^{B*} = \frac{1 - F(P^{B*})}{f(P^{B*})} - \pi$. That is, the buyer-side fee follows the Ramsey inverse-elasticity rule, adjusted downwards by the surplus π that the marginal buyer generates on the merchant side. This fee maximises $[1 - F(P^B)](\pi + P^B)$, the revenue generated by a successful merchant, which is the same for all merchant types. Since a merchant is worth financing if and only if $\theta [1 - F(P^B)](\pi + P^B) \geq I$, the same fee also minimises the financing cutoff $\underline{\theta}^{\text{PF}}(P^B)$. The platform therefore faces no trade-off between pricing and financing, since the fee that maximises revenue from each successful merchant also makes the largest set of merchants worth financing.

2.2 Bank-lending benchmark

We now consider a second benchmark, in which the uninformed bank finances merchants and the platform only matches buyers and merchants, with no partnership arrangement. Credit is then misallocated for two reasons. First, the bank does not receive the platform's fees and therefore applies a stricter lending cutoff than the platform would prefer. The platform can partly correct this conflict by lowering its merchant fee, which increases the amount merchants can repay to the bank. Second, the bank cannot screen merchants and therefore lends either to all of them or to none, meaning that the above correction works only when average merchant quality is sufficiently high and still cannot be applied selectively across merchants, so platform profit remains below the monopoly benchmark.

$t = 1c$: Optimal merchant financing The bank sets $D = (1 - F(P^B))(\pi - P^M)$ as the only financier. Without information about merchant types, its expected payoff from

lending to a merchant is $\mathbb{E}[\theta]D - I$, which is non-negative if and only if

$$\mathbb{E}[\theta] \geq \frac{I}{[1 - F(P^B)](\pi - P^M)} \equiv \underline{\theta}^{\text{Bank}}(P^M, P^B). \quad (3)$$

The condition does not depend on the merchant's type, so the bank finances either all M merchants or none.¹⁶ Moreover, as long as the platform earns positive fees ($P^M + P^B > 0$), the bank's threshold exceeds the platform's, $\underline{\theta}^{\text{PF}}(P^B) < \underline{\theta}^{\text{Bank}}(P^M, P^B)$, since the platform internalises matchmaking revenues that are unavailable to the bank. This creates a conflict of interest between the platform and the bank over whether merchants should be financed.

$t = 0$: Optimal platform pricing When no merchant is financed, all agents earn zero payoffs. If the platform implements full financing, it solves

$$\max_{P^M, P^B} [1 - F(P^B)]M\mathbb{E}[\theta](P^M + P^B) \quad \text{s.t.} \quad \mathbb{E}[\theta] \geq \frac{I}{[1 - F(P^B)](\pi - P^M)}. \quad (\text{PC})$$

The platform can resolve the conflict of interest through the merchant-side fee. Lowering P^M by one unit reduces the platform's fee income from a successful merchant by $1 - F(P^B)$, the number of buyers that trade with the merchant, and increases the repayment the merchant can make to the bank by the same amount. Thus, P^M moves surplus between the platform and the bank one for one. The platform binds (PC), and its problem becomes

$$\max_{P^B} M [[1 - F(P^B)]\mathbb{E}[\theta](\pi + P^B) - I], \quad (4)$$

with solution

$$P^{B*} = \frac{1 - F(P^{B*})}{f(P^{B*})} - \pi, \quad P^{M*} = \pi - \frac{I}{[1 - F(P^{B*})]\mathbb{E}[\theta]}.$$

The platform thus holds the bank to zero profit and extracts the full surplus from merchant production, as if it lent to merchants itself. It implements full financing if and only if $\mathbb{E}[\theta] \geq \underline{\theta}^{\text{PF}}(P^{B*})$, so its own financing cutoff, rather than the bank's threshold, becomes decisive for lending. However, strategic pricing corrects the conflict of interest only for the pool as a whole, since the platform's information remains unused.

¹⁶The bank cannot screen merchants using a menu of contracts. A type- θ merchant's expected payoff from a contract $\{I, D\}$ offered with probability ρ is $\rho\theta[(1 - F(P^B))(\pi - P^M) - D]$, so all merchant types rank contracts in the same way.

2.3 Platform profit across benchmarks

Platform profit in the two benchmarks is

$$\begin{aligned}\Pi^{\text{Mono}}(M) &= M \int_{\underline{\theta}^{\text{PF}}(P^{B*})}^1 \left[\theta [1 - F(P^{B*})] (\pi + P^{B*}) - I \right] g(\theta) d\theta, \\ \Pi^{\text{Bank-lend}}(M) &= \begin{cases} M \left([1 - F(P^{B*})] \mathbb{E}[\theta] (\pi + P^{B*}) - I \right) & \text{if } \mathbb{E}[\theta] \geq \underline{\theta}^{\text{PF}}(P^{B*}), \\ 0 & \text{otherwise.} \end{cases}\end{aligned}$$

The gap between the two reflects the loss from separating capital and information: the monopolist finances merchants selectively, targeting high-type merchants and some lower-type merchants whose externalities justify loan losses, whereas the uninformed bank finances all merchants or none. This loss is largest when low average merchant quality shuts down lending altogether. In the next two sections, we examine whether intermediation and information provision improve on the bank-lending benchmark, and how closely each approaches the monopoly benchmark. To make this comparison stark, we assume that average merchant quality is too low for any lending in the bank-lending benchmark, while some merchants are still worth financing for the platform.

Assumption 2. $\mathbb{E}[\theta] < \underline{\theta}^{\text{PF}}(P^{B*}) < 1$.

From now on, we write $\underline{\theta}^{\text{PF}}$ for $\underline{\theta}^{\text{PF}}(P^{B*})$ when no confusion arises. By (1), $[1 - F(P^{B*})](\pi + P^{B*}) = I/\underline{\theta}^{\text{PF}}$, so financing a type- θ merchant yields $\frac{I}{\underline{\theta}^{\text{PF}}}(\theta - \underline{\theta}^{\text{PF}})$ to the platform at the optimal buyer-side fee. Hence $\Pi^{\text{Mono}}(M) = M \Pi^{\text{Mono}}(1)$, where

$$\Pi^{\text{Mono}}(1) = \frac{I}{\underline{\theta}^{\text{PF}}} \int_{\underline{\theta}^{\text{PF}}}^1 (\theta - \underline{\theta}^{\text{PF}}) g(\theta) d\theta \quad (5)$$

is the monopoly profit per merchant.

3 Intermediation: Platform as delegated financier

In this section, we study intermediation, where the platform borrows through a credit line from the bank and on-lends the funds to merchants. The key mechanism is that the platform and the bank agree on which merchants to finance first, since both prefer more productive merchants, but disagree on how many to finance. The bank is therefore willing to provide an endogenously capped credit line and leave merchant selection to the platform, which finances merchants from the highest type down. This can restart lending even when the bank would not lend to any merchant as in the bank-lending benchmark.

Platform profit therefore exceeds the bank-lending benchmark, but is still lower than the monopoly benchmark because the credit line fixes the number of financed merchants ex ante, while the number worth financing ex post is random. As the merchant pool

grows, a merchant's rank reveals more about its type, making selection more effective. Profit per merchant therefore converges to the monopoly level.

The total gap to the monopoly benchmark, however, converges to a positive constant reflecting the loss from fixing the credit line ex ante. This loss remains bounded in larger merchant pools because two effects offset each other: more merchants are misallocated when the realised number worth financing differs from the credit line, but each misallocated merchant lies closer to the financing cutoff and therefore causes a smaller loss.

3.1 $t = 1b$: Optimal merchant financing

Given a credit line from the bank, the platform finances merchants in descending order of quality θ . This is because the repayment from a merchant financed under the credit line goes to the bank, as introduced in Section 2, while the platform earns transaction fees whenever the merchant succeeds. The platform's payoff from financing a merchant is therefore strictly increasing in its quality. It starts with the best merchant and, as long as it earns positive fees on each trade, continues down the ranked list until the credit line is exhausted.

Lemma 1. *Fix the prices (P^B, P^M) and the credit line nI , and suppose that the platform earns positive fees from a successful merchant ($P^M + P^B > 0$). Then the platform finances the n merchants with the highest types.*

3.2 $t = 1a$: Optimal credit line features endogenous cap on n

Anticipating the platform's monotone financing strategy, the bank agrees only to a credit line that caps the number of financed merchants at $n^*(M)$. This cap arises because the bank must break even and can be repaid only from merchants' realised output, and it does not benefit from the network externalities generated on the platform. As the platform expands lending to progressively lower-quality merchants, the average quality of the financed portfolio of merchants falls. The cap $n^*(M)$ is therefore the largest portfolio size for which average merchant quality remains above the bank's lending threshold $\underline{\theta}^{\text{Bank}}(P^M, P^B)$.

Formally, under a credit line nI , the bank receives the repayment of each financed merchant that succeeds. Since the platform keeps none of these repayments, the credit line can specify the largest repayment allowed by limited liability, $D = [1 - F(P^B)](\pi - P^M)$. The bank's participation constraint therefore becomes

$$\sum_{a=0}^n p_a(n) a [1 - F(P^B)] (\pi - P^M) \geq nI$$

$$\Longleftrightarrow \underbrace{\frac{1}{n} \sum_{a=0}^n p_a(n) a}_{\text{average quality of the top } n \text{ merchants}} \geq \underbrace{\frac{I}{[1 - F(P^B)](\pi - P^M)}}_{\text{bank threshold } \underline{\theta}^{\text{Bank}}},$$

where $p_a(n)$ is the probability that exactly a of the top- n (out of M) merchants succeed. Since the average quality of the top n merchants is strictly decreasing in n , there is a maximum portfolio size $n^*(M)$ that the bank is willing to finance.¹⁷

Lemma 2. *There exists a largest integer $n^*(M) \in \{0, 1, \dots, M\}$ such that the bank's participation constraint holds.*

3.3 $t = 0$: Optimal platform pricing

As in the bank-lending benchmark, the platform uses the merchant-side fee to resolve the conflict between $\underline{\theta}^{\text{Bank}}(P^M, P^B)$ and $\underline{\theta}^{\text{PF}}(P^B)$, but now for a selected portfolio of merchants rather than the entire pool. Formally, the platform solves

$$\begin{aligned} \max_{P^M, P^B, n \in \{0, \dots, M\}} \quad & \sum_{a=0}^n p_a(n) a [1 - F(P^B)] (P^M + P^B) & (6) \\ \text{s.t.} \quad & \sum_{a=0}^n p_a(n) a [1 - F(P^B)] (\pi - P^M) \geq nI, & (\text{PC}) \\ & (P^M + P^B) \sum_{j=M-n+1}^M \theta_{j:M} \geq (P^M + P^B) \sum_{m \in T} \theta_m \quad \forall T : |T| \leq n. & (\text{IC}) \end{aligned}$$

where $\theta_{j:M}$ denotes the j -th lowest of the M merchant types. The objective is the platform's expected fee income from financed merchants, each of which repays $D = [1 - F(P^B)](\pi - P^M)$ to the bank when successful, as in the bank's participation constraint. The incentive constraint reflects that the bank delegates merchant selection to the platform: for every realisation of merchant types, the platform must prefer financing the n merchants with the highest types to financing any other set $T \subseteq \{1, \dots, M\}$ of at most n merchants. By Lemma 1, (IC) holds whenever $P^M + P^B > 0$, so the platform and the bank agree on which merchants to finance. The platform would, however, like to finance as many merchants as possible, since its payoff $\theta[1 - F(P^B)](P^M + P^B)$ from each financed merchant is positive at the equilibrium fees. The bank therefore controls the number of financed merchants through the size n of the credit line.

Since P^M moves surplus between the platform and the bank one for one, the platform binds the bank's participation constraint, and its problem reduces to choosing the buyer-

¹⁷The size of the credit line does not depend on the timing of contracting. If the credit line were agreed after the platform observes merchant types, the platform would request the largest line that the bank is willing to provide regardless of the realised types, so the request would reveal no information.

side fee and the size of the credit line:

$$\max_{P^B, n \in \{0, \dots, M\}} \sum_{a=0}^n p_a(n) [1 - F(P^B)] (\pi + P^B) a - nI. \quad (7)$$

The optimal buyer-side fee remains the adjusted Ramsey fee P^{B*} . The optimal credit-line size is

$$n^*(M) = \sum_{j=1}^M \mathbf{1}\{\mathbb{E}[\theta_{j:M}] \geq \underline{\theta}^{\text{PF}}\}.$$

That is, the platform finances a merchant if and only if the expected type at its rank is at least $\underline{\theta}^{\text{PF}}$. The optimal merchant-side fee

$$P^{M*} = \pi - \frac{I}{[1 - F(P^{B*})] \cdot \frac{1}{n^*(M)} \sum_{a=0}^{n^*(M)} p_a(n^*(M)) a}$$

leaves the bank just enough repayment to break even, and at this fee $n^*(M)$ is the largest credit line that the bank accepts. Finally, since $\sum_{a=0}^n p_a(n) a = \sum_{j=M-n+1}^M \mathbb{E}[\theta_{j:M}]$ and $[1 - F(P^{B*})](\pi + P^{B*}) = I/\underline{\theta}^{\text{PF}}$ as in (1), platform profit under intermediation is

$$\Pi^{\text{Interm}}(M) = \frac{I}{\underline{\theta}^{\text{PF}}} \sum_{j=1}^M \max\{\mathbb{E}[\theta_{j:M}] - \underline{\theta}^{\text{PF}}, 0\}. \quad (8)$$

3.4 Platform profit behaviour

Since the platform can always choose not to contract with the bank, $\Pi^{\text{Interm}}(M) \geq \Pi^{\text{Bank-lend}}(M) = 0$ under Assumption 2. Platform profit, however, remains strictly below the monopoly benchmark. The reason is that the credit line fixes the number of financed merchants at $n^*(M)$ ex ante, while the number of merchants with $\theta \geq \underline{\theta}^{\text{PF}}$ that a monopolist would finance depends on the realised merchant types. Whenever these numbers differ, intermediation either leaves a merchant with $\theta > \underline{\theta}^{\text{PF}}$ unfunded or finances a merchant with $\theta < \underline{\theta}^{\text{PF}}$.

Lemma 3. $\Pi^{\text{Interm}}(M) < \Pi^{\text{Mono}}(M), \forall M$.

The following result describes how platform profit changes with the merchant pool size M .

Proposition 1 (Platform profit as intermediary). *Platform profit under intermediation has the following properties.*

- (i) $\Pi^{\text{Interm}}(1) = \Pi^{\text{Bank-lend}}(1) = 0$.
- (ii) Profit per merchant $\Pi^{\text{Interm}}(M)/M$ is non-decreasing in M and converges to the monopoly level, $\lim_{M \rightarrow \infty} \Pi^{\text{Interm}}(M)/M = \Pi^{\text{Mono}}(1)$.

(iii) $\Pi^{\text{Interm}}(M)$ is non-decreasing in M . The total profit gap to the monopoly benchmark converges to a strictly positive constant C that depends on I , $\underline{\theta}^{\text{PF}}$ and G :

$$\lim_{M \rightarrow \infty} [\Pi^{\text{Mono}}(M) - \Pi^{\text{Interm}}(M)] = C > 0.$$

The intuition is as follows. Under intermediation, the platform finances the top $n^*(M)$ merchants in the pool, so financing depends on a merchant's rank rather than directly on its type. A larger pool makes ranks more informative about types. With only two merchants, for example, ranking first reveals only that a merchant is better than the other. With many merchants, a merchant's rank reveals its percentile in the type distribution, and hence its type, almost exactly. Selection therefore becomes more effective as the pool grows, so profit per merchant is non-decreasing in M and converges to the monopoly level, at which the platform finances exactly those merchants with $\theta \geq \underline{\theta}^{\text{PF}}$.

It follows immediately that total platform profit under intermediation is also non-decreasing in M . The total gap to the monopoly benchmark, however, converges to a positive constant C , capturing the loss from fixing the credit line ex ante. This loss nevertheless remains bounded in larger merchant pools. Since the platform finances the best merchants first, any mismatch between the credit line and the realised number of merchants worth financing affects only merchants close to the financing cutoff $\underline{\theta}^{\text{PF}}$. As M grows, this mismatch involves more merchants, but each lies closer to the cutoff and therefore causes a smaller loss. These two effects offset each other in the limit. Appendix A provides a closed-form representation of C and discusses its further properties.

4 Communication: Platform as information provider

In this section, we study information provision, where the platform does not lend but shares information with the bank through a reporting rule. The key mechanism is that the platform would like the bank to lend to every merchant, since it earns fees on each financed merchant and bears no credit risk, but recommending every merchant would leave the bank's belief unchanged. The rule therefore recommends only merchants above a cutoff, raising the bank's posterior belief about recommended merchants and restarting lending, provided that the rule is sufficiently credible. As in the bank-lending benchmark, the platform uses the merchant fee to resolve its conflict with the bank. It therefore chooses a rule that recommends exactly the merchants it would finance as a monopolist.

Platform profit therefore exceeds the bank-lending benchmark but remains below the monopoly benchmark. Whenever the platform can override the rule, it recommends every merchant, so some merchants below its financing cutoff are also financed. Unlike under intermediation, a larger merchant pool brings no selection gains. The platform applies the same reporting rule independently to each merchant, so profit is linear in pool size and

the profit gap per merchant remains constant. The total gap to the monopoly benchmark therefore grows in proportion to the pool.

4.1 $t = 1c$: Optimal merchant financing

After observing the (possibly null) report s_m on merchant m , the bank forms a posterior belief $\mathbb{E}[\theta_m \mid s_m]$ about the merchant's quality and lends if and only if

$$\mathbb{E}[\theta_m \mid s_m] \geq \underline{\theta}^{\text{Bank}}(P^M, P^B).$$

Relative to the bank-lending benchmark (3), the platform's report replaces the prior $\mathbb{E}[\theta]$ with a posterior as the basis for the bank's decision.

4.2 $t = 1a$: Optimal reporting rule

The platform chooses its reporting rule to maximise fee income, which increases in the number of financed merchants, since the platform receives only transaction fees and no loan repayments. For the same reason, whenever it can override the rule at $t = 1b$, it sends the report that induces the bank to lend, regardless of the merchant's type.

The optimal reporting rule takes a cutoff form: every merchant with type above θ^* receives the report " $\theta \geq \theta^*$ ", upon which the bank lends, while every other merchant receives a null report. A single cutoff applied to all merchants suffices for three reasons. First, merchant types are i.i.d. and lending decisions are merchant-specific, so the same rule applies to every merchant. Second, the bank has only two actions (lend/not), so a favourable report and its absence are sufficient to generate the two corresponding posteriors. Third, for a given lending threshold of the bank, recommending an upper interval of types $[\theta^*, 1]$ maximises both the bank's posterior and the platform's fee income. Appendix A provides the formal argument.

Given the reporting rule, a favourable report is either specified for a merchant with $\theta \geq \theta^*$ or results from an override, which occurs with probability $1 - \alpha$ for every type. The bank's posterior on a recommended merchant is therefore

$$\mathbb{E}[\theta \mid s'(\theta^*)] = \frac{\int_{\theta^*}^1 \theta g(\theta) d\theta + (1 - \alpha) \int_0^{\theta^*} \theta g(\theta) d\theta}{1 - \alpha G(\theta^*)}, \quad \text{where } s'(\theta^*) = [\theta^*, 1]. \quad (9)$$

The choice of θ^* trades off the reach and persuasiveness of the report. Raising θ^* initially filters out low types and raises the posterior. If θ^* is too high, however, favourable reports specified by the rule become rare and increasingly come from overrides on low types, lowering the posterior.

Lemma 4 (Optimal reporting rule). *The optimal reporting rule specifies the report $s'(\theta^*)$ for merchants with $\theta \geq \theta^*$ and a null report otherwise, where θ^* is the smallest $\theta' \in [0, 1]$*

satisfying

$$\mathbb{E}[\theta \mid s'(\theta')] \geq \underline{\theta}^{\text{Bank}}(P^M, P^B),$$

for any fixed prices P^M and P^B with $P^M + P^B > 0$. Whenever it can override the rule, the platform reports $s'(\theta^*)$. If no such θ' exists, no reporting rule induces lending.

4.3 $t = 0$: Optimal platform pricing

The platform sets P^M so that the bank, given its posterior, just breaks even on a recommended merchant, which again resolves the conflict between $\underline{\theta}^{\text{Bank}}(P^M, P^B)$ and $\underline{\theta}^{\text{PF}}(P^B)$. Formally, the platform solves

$$\max_{P^M, P^B, \theta'} M [1 - F(P^B)] (P^M + P^B) \left[\int_{\theta'}^1 \theta g(\theta) d\theta + (1 - \alpha) \int_0^{\theta'} \theta g(\theta) d\theta \right] \quad (10)$$

$$\text{s.t. } \mathbb{E}[\theta \mid s'(\theta')] \geq \frac{I}{[1 - F(P^B)](\pi - P^M)}, \quad (\text{PC})$$

$$\theta' = \min \left\{ \tilde{\theta} \in [0, 1] : \mathbb{E}[\theta \mid s'(\tilde{\theta})] \geq \frac{I}{[1 - F(P^B)](\pi - P^M)} \right\}. \quad (\text{IC})$$

The objective is the platform's expected fee income from financed merchants. A merchant with $\theta \geq \theta'$ is recommended under the reporting rule, while a merchant with $\theta < \theta'$ is recommended with probability $1 - \alpha$, when the platform overrides the rule. Relative to intermediation, both constraints take a different form. The bank's participation constraint now depends on the posterior quality of a recommended merchant rather than on the average quality of a portfolio. The incentive constraint now concerns the reporting rule rather than merchant selection: since the platform bears no credit risk and would like every merchant to be financed, it chooses the lowest cutoff at which the bank is still willing to lend as in Lemma 4.

Substituting the binding (PC) into the objective reduces the platform's problem to choosing the buyer-side fee P^B and the reporting cutoff θ' :

$$\max_{P^B, \theta'} M \left[[1 - F(P^B)] (\pi + P^B) \left[\int_{\theta'}^1 \theta g(\theta) d\theta + (1 - \alpha) \int_0^{\theta'} \theta g(\theta) d\theta \right] - I[1 - \alpha G(\theta')] \right]. \quad (11)$$

The solution is $P^B = P^{B*}$ and $\theta^* = \underline{\theta}^{\text{PF}}$, with

$$P^{M*} = \pi - \frac{I}{[1 - F(P^{B*})]\mathbb{E}[\theta \mid s'(\theta^*)]}.$$

That is, the reporting rule recommends exactly the merchants that the platform would finance as a monopolist, while merchants below $\underline{\theta}^{\text{PF}}$ are overfinanced with probability $1 - \alpha$. A higher cutoff would leave merchants just above $\underline{\theta}^{\text{PF}}$ unfunded with probability α , while a lower cutoff would finance merchants just below $\underline{\theta}^{\text{PF}}$ with certainty, even though

they are not worth financing.

4.4 Platform profit behaviour

Since merchants below $\underline{\theta}^{\text{PF}}$ are financed with probability $1 - \alpha$, platform profit under information provision remains below the monopoly benchmark.

Proposition 2 (Platform profit as information provider). *Platform profit under information provision is linear in M , $\Pi^{\text{Info}}(M) = M \Pi^{\text{Info}}(1)$, where*

$$\Pi^{\text{Info}}(1) = \max \left\{ \Pi^{\text{Mono}}(1) - (1 - \alpha) \frac{I}{\underline{\theta}^{\text{PF}}} \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta, 0 \right\}. \quad (12)$$

Moreover, $\Pi^{\text{Info}}(1) > 0$ if and only if $\alpha > \underline{\alpha} \equiv (\underline{\theta}^{\text{PF}} - \mathbb{E}[\theta]) / \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta$, where $\underline{\alpha} \in (0, 1)$.

The intuition for linearity is straightforward. The platform applies the same reporting rule independently to each merchant, so total profit equals profit per merchant multiplied by the number of merchants. Unlike under intermediation, a larger merchant pool brings no selection gains. Profit per merchant is positive if and only if the reporting rule is sufficiently credible, $\alpha > \underline{\alpha}$, so that the average type of a recommended merchant exceeds the platform's financing cutoff.

5 Regime switch, comparative statics, welfare

In this section, we compare the two arrangements and show how the platform's choice depends on the size of its merchant pool. We then examine how the switching threshold depends on the credibility of data sharing, the production environment via the platform's financing cutoff, and the dispersion of merchant types. Finally, we compare credit quality and welfare under the two arrangements.

5.1 Regime switch

There is a unique threshold M^* that determines the platform's optimal choice between the two arrangements. Profit per merchant is non-decreasing in M under intermediation (Proposition 1) and constant under information provision (Proposition 2). Since intermediation profit per merchant converges to the monopoly level, which strictly exceeds profit per merchant under information provision when $\alpha < 1$, it eventually overtakes information provision and does so only once.

If information provision is sufficiently credible, $\alpha > \underline{\alpha}$, it dominates in small pools because it can restart lending regardless of pool size. Intermediation, by contrast, gains little from selection in a small pool and cannot improve on bank lending at all with a

single merchant. As M grows, however, selection under intermediation becomes more effective, raising profit per merchant until it exceeds the constant profit per merchant under information provision.

Proposition 3 (Regime switch). *Fix $\alpha < 1$. There exists a finite M^* , such that $\Pi^{\text{Interm}}(M) \geq \Pi^{\text{Info}}(M)$ for all $M \geq M^*$, and $\Pi^{\text{Interm}}(M) < \Pi^{\text{Info}}(M)$ for all $M < M^*$. Furthermore, if $\alpha > \underline{\alpha}$, then $M^* \geq 2$.*

Remark 1 (Commitment and the regime switch). With full commitment ($\alpha = 1$), corresponding to Bayesian persuasion, information provision attains the monopoly benchmark, $\Pi^{\text{Info}}(M) = \Pi^{\text{Mono}}(M)$, and therefore strictly dominates intermediation for every M by Lemma 3. The regime switch thus arises from the platform’s limited commitment to its reporting rule. This is related to institutional requirements on platform-bank partnerships that require banks to verify information provided by platforms, and we show that such requirements help discipline platforms precisely when commitment is imperfect.¹⁸

5.2 Comparative statics

Credibility α A higher α raises platform profit under information provision and leaves it unchanged under intermediation, so the switching threshold M^* increases.

Corollary 1 (Effect of credibility α). *The threshold M^* is non-decreasing in α , and $M^* \rightarrow \infty$ as $\alpha \rightarrow 1$.*

Intuitively, a higher α means that the platform overrides its reporting rule less often, reducing overfinancing and making a favourable report more informative to the bank. As α approaches one, information provision approaches the monopoly benchmark, which intermediation never attains by Lemma 3. The switching threshold therefore grows without bound.

We can also fix the merchant pool size and ask how credible data sharing must be for the platform to choose information provision. For a given M , profit under intermediation does not depend on α , whereas profit under information provision increases with α . There is therefore a credibility level at which the platform is indifferent between the two arrangements.

Corollary 2. *Fix M . There exists a boundary $\alpha^*(M)$ such that information provision yields strictly higher profit than intermediation if and only if $\alpha > \alpha^*(M)$. The boundary $\alpha^*(M)$ is non-decreasing in M , equals $\underline{\alpha}$ for $M = 1$, is below one, and converges to one as $M \rightarrow \infty$.*

¹⁸See, for example, the [Interagency Guidance on Third-Party Relationships: Risk Management](#) (Federal Reserve, FDIC, and OCC, 2023).




Figure 1: Regime switch and credibility

Notes: Merchant types are uniform on $[0, 1]$. Panel (a) shows profit per merchant, in units of $I/\underline{\theta}^{\text{PF}}$, under intermediation, $\Pi^{\text{Interm}}(M)/M$, and under information provision, $\Pi^{\text{Info}}(1)$, for $\underline{\theta}^{\text{PF}} = 0.6$. Panel (b) shows the boundary $\alpha^*(M)$ from Corollary 2 for $\underline{\theta}^{\text{PF}} = 0.6$ (solid) and $\underline{\theta}^{\text{PF}} = 0.55$ (dashed). Information provision is optimal above the boundary and intermediation below it. The arrows illustrate a platform with $\alpha = 0.95$ that switches to intermediation as its pool grows from 3 to 30 merchants and switches back to information provision when credibility rises to 0.99.

Intuitively, as the merchant pool grows, intermediation becomes more profitable, so information provision must be more credible to remain preferred. With a single merchant, information provision is preferred whenever it is profitable, while in large pools credibility must approach one. Corollaries 1 and 2 therefore describe the same comparison from two sides: the platform chooses information provision if and only if $\alpha > \alpha^*(M)$, or equivalently $M < M^*$.

Figure 1 illustrates these results for uniform merchant types and $\underline{\theta}^{\text{PF}} = 0.6$. In panel (a), profit per merchant under intermediation increases with pool size towards the monopoly level, whereas under information provision it is constant. With merchant pool size growing from $M = 1$, the platform switches to intermediation once the former reaches the latter, at $M^* = 7$ when $\alpha = 0.9$ and at $M^* = 13$ when $\alpha = 0.95$. Panel (b) plots the boundary $\alpha^*(M)$, that is, for each M , the credibility at which information provision yields the same profit as intermediation in panel (a). Information provision is optimal above the boundary and intermediation below it. The arrows illustrate a platform with $\alpha = 0.95$ that starts with information provision, switches to intermediation as its pool grows from 3 to 30 merchants, crossing the boundary at $M^* = 13$, and switches back when its credibility rises to 0.99.

Remark 2. Suppose that, before choosing the arrangement, the platform can pay to increase the credibility of its data sharing, for example by building an interface that the bank can audit. A platform with a large merchant pool has little incentive to do so, for two reasons. First, the gain is bounded. Since the platform can always intermediate and full credibility yields at most the monopoly profit, the gain cannot exceed $\Pi^{\text{Mono}}(M) -$

$\Pi^{\text{Interm}}(M)$, which converges to the constant C by Proposition 1. Second, the gain is all or nothing. An intermediating platform benefits only if its credibility rises above $\alpha^*(M)$, which approaches one as the merchant pool grows by Corollary 2. A return to information provision, as in panel (b) of Figure 1, is therefore more likely to follow improvements in credibility that the platform does not pay for, such as rules governing data access.

Financing cutoff $\underline{\theta}^{\text{PF}}$ The investment cost I , the trade surplus π , and the distribution F of buyer valuations affect the platform's choice only through the financing cutoff $\underline{\theta}^{\text{PF}} = I / \max_{P^B} [1 - F(P^B)](\pi + P^B)$. The cutoff increases with I , decreases with π , and weakly decreases with buyer demand, for example following an upward shift in buyer valuations in the sense of first-order stochastic dominance. When information provision is credible, a lower cutoff favours information provision.

Corollary 3 (Effect of the financing cutoff). *For any two cutoffs $\underline{\theta}^{\text{PF}} < \hat{\underline{\theta}}^{\text{PF}}$ in $(\mathbb{E}[\theta], 1)$, M^* is larger under $\underline{\theta}^{\text{PF}}$ for all α sufficiently close to one.*

The intuition is as follows. When information provision is credible, the threshold M^* is large, so the comparison depends on how the two arrangements perform in large pools. Under information provision, the overfinancing loss per merchant is proportional to $\int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta = \mathbb{E}[\max\{\underline{\theta}^{\text{PF}} - \theta, 0\}]$, the expected amount by which a merchant's type falls below the cutoff. As the cutoff falls, fewer merchants lie below it and those that do are closer to it, so this loss falls quickly. Under intermediation, the loss in large pools comes only from merchants near the cutoff and therefore depends on the density of types around it. Log-concavity of g ensures that the overfinancing loss falls proportionally faster than the intermediation loss as the cutoff decreases. Hence, when information provision is credible, a lower cutoff, due to cheaper projects, more valuable trade, or stronger buyer demand, favours information provision, while a higher cutoff favours intermediation. In panel (b) of Figure 1, the boundary for the lower cutoff $\underline{\theta}^{\text{PF}} = 0.55$ (dashed) therefore lies below that for $\underline{\theta}^{\text{PF}} = 0.6$ (solid).

Merchant type distribution $G(\cdot)$ An increase in merchant quality dispersion, represented by a mean-preserving spread of $G(\cdot)$, weakly raises the platform's profits under both arrangements. The resulting net effect on the threshold M^* is ambiguous.

Corollary 4 (Effect of merchant type dispersion). *When the merchant type distribution $G(\cdot)$ is replaced by a mean-preserving spread, both $\Pi^{\text{Interm}}(M)$ and $\Pi^{\text{Info}}(M)$ weakly increase for every M . The change in the threshold M^* is ambiguous.*

Intuitively, a mean-preserving spread shifts probability mass towards both tails of the type distribution. Under intermediation, the spread raises the expected types of the best merchants, allowing the platform to form a better portfolio for any given credit

line. Under information provision, since the mean type is unchanged, the spread raises the expected gain from merchants above the cutoff and the expected loss from those below it by the same amount. Because merchants below the cutoff are financed only with probability $1 - \alpha$, the gain dominates and profit rises. Both arrangements therefore benefit from the spread, but neither benefits more in general, so the effect on M^* is ambiguous.

5.3 Credit quality

The two arrangements also differ in the quality of the loans they generate. Under information provision, the average success probability of financed merchants is the bank's posterior $\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})]$, which is independent of M . Under intermediation, it is the average expected type of the $n^*(M)$ best merchants, $\frac{1}{n^*(M)} \sum_{j=M-n^*(M)+1}^M \mathbb{E}[\theta_{j:M}]$.

Corollary 5 (Credit quality). *As $M \rightarrow \infty$, the average success probability of merchants financed under intermediation converges to $\mathbb{E}[\theta \mid \theta \geq \underline{\theta}^{\text{PF}}]$, which strictly exceeds $\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})]$ when $\alpha < 1$. Hence, for sufficiently large M , loans originated under intermediation default less often than those originated under information provision.*

The reason is that, under intermediation, misallocation occurs only around the cutoff and vanishes per merchant as the pool grows. Under information provision, by contrast, merchants far below the cutoff are financed with probability $1 - \alpha$ regardless of pool size.

5.4 Welfare

Welfare in our model is the expected total surplus of buyers, merchants, the platform, and the bank. In both arrangements, the bank breaks even and merchants obtain no surplus, so welfare equals platform profit plus buyer surplus. We compare this sum under the two arrangements, each as analysed in the previous sections. A buyer who joins the platform obtains $v - P^{B*}$ from each trade. A financed merchant that succeeds therefore generates total surplus $[1 - F(P^{B*})](\pi + \mathbb{E}[v \mid v \geq P^{B*}])$, of which the platform captures only $[1 - F(P^{B*})](\pi + P^{B*})$. Financing a type- θ merchant therefore raises welfare if and only if

$$\theta \geq \frac{I}{[1 - F(P^{B*})](\pi + \mathbb{E}[v \mid v \geq P^{B*}])} \equiv \underline{\theta}^{\text{W}}, \quad (13)$$

where we call $\underline{\theta}^{\text{W}}$ the welfare cutoff. Since the platform does not internalise the surplus of buyers who trade with a financed merchant, $\underline{\theta}^{\text{W}} < \underline{\theta}^{\text{PF}}$.

Hence, from a welfare perspective, the monopoly benchmark finances too few merchants, which means that overfinancing under information provision is not necessarily inefficient. It finances merchants with $\theta \in [\underline{\theta}^{\text{W}}, \underline{\theta}^{\text{PF}})$, which increase welfare, as well as merchants with $\theta < \underline{\theta}^{\text{W}}$, which reduce it.

Under information provision with $\alpha > \underline{\alpha}$, a merchant with $\theta \geq \underline{\theta}^{\text{PF}}$ is financed, while a merchant with $\theta < \underline{\theta}^{\text{PF}}$ is financed with probability $1 - \alpha$. Since financing a type- θ merchant adds $\frac{I}{\underline{\theta}^{\text{W}}}(\theta - \underline{\theta}^{\text{W}})$ to welfare by (13), welfare per merchant is

$$\frac{I}{\underline{\theta}^{\text{W}}} \left[\int_{\underline{\theta}^{\text{PF}}}^1 (\theta - \underline{\theta}^{\text{W}}) g(\theta) d\theta + (1 - \alpha) G(\underline{\theta}^{\text{PF}}) (\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}] - \underline{\theta}^{\text{W}}) \right]. \quad (14)$$

Under intermediation, total welfare is $\frac{I}{\underline{\theta}^{\text{W}}} \sum_{j=M-n^*(M)+1}^M (\mathbb{E}[\theta_{j:M}] - \underline{\theta}^{\text{W}})$.

Proposition 4 (Welfare). *Let $\underline{\alpha} < \alpha < 1$.*

- (i) *For $M = 1$, information provision yields strictly higher welfare than intermediation.*
- (ii) *If $\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}] > \underline{\theta}^{\text{W}}$, information provision yields strictly higher welfare than intermediation for all sufficiently large M . If $\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}] < \underline{\theta}^{\text{W}}$, intermediation yields strictly higher welfare for all sufficiently large M .*

Part (i) holds because, with a single merchant, the bank provides no credit line under intermediation by Proposition 1, while information provision raises welfare whenever it is profitable for the platform. The intuition for part (ii) is as follows. In large pools, intermediation finances almost exactly the merchants above the platform's cutoff, as the monopolist does. Information provision finances these merchants as well and, with probability $1 - \alpha$, also finances merchants below the cutoff, whose average type is $\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}]$. This overfinancing raises welfare if the average type of these merchants is above $\underline{\theta}^{\text{W}}$ and lowers it if the average type is below $\underline{\theta}^{\text{W}}$.

6 Buyer liquidity constraints & BNPL

We now extend the baseline model by assuming that buyers cannot pay the platform fee P^B at the time of trade, so the platform partners with a Buy Now, Pay Later (BNPL) provider that pays the fee on their behalf. We show that the regime switch continues to hold and that, when information provision is credible, the threshold M_{BNPL}^* increases with the value that the BNPL provider derives from the trades it enables.

6.1 Extended model setting

Buyers All buyers are cash constrained at the start of $t = 2$. Each buyer receives a sufficiently large endowment at the end of $t = 2$ with probability $q > 0$, and no endowment with probability $1 - q$. Endowment draws are independent across buyers.

BNPL The platform partners with a representative competitive BNPL provider that has sufficient capital to pay the platform's fee on behalf of buyers at the time of trade. Any buyer on the platform can use BNPL. The provider charges the platform a commission ϕ

per transaction, so the platform receives $P^B - \phi$ from the buyer side of each transaction. A negative commission means that the provider pays the platform for access to its buyers. For simplicity, we assume that the platform has full bargaining power in setting ϕ .

The BNPL provider, in turn, derives value from enabling trades. This value can be pecuniary (for example, advertising revenue, cross-selling) and non-pecuniary (for example, data to improve its algorithms) (Bergemann et al., 2022; Bian et al., 2024; Hagiu & Wright, 2023).¹⁹ We denote by $w \geq 0$ the value that the provider derives from each trade it enables, which captures its ability to monetise its lending services.²⁰

Timing The sequence of events at $t = 2$ is updated as follows. Financed merchants produce. Buyers who wish to trade request BNPL, the provider pays the platform $P^B - \phi$ for each trade with a successful merchant, and trades occur. Buyers' stochastic endowments are then realised, and those who receive an endowment repay the BNPL provider (we assume no cash-diversion technology).

6.2 Regime switch

The BNPL provider breaks even on a transaction if

$$\phi + w \geq (1 - q)P^B, \quad (\text{PC}_{\text{BNPL}})$$

since it bears expected default losses $(1 - q)P^B$, receives a commission ϕ , and derives value w from each transaction it enables. With full bargaining power, the platform sets the commission so that $(\text{PC}_{\text{BNPL}})$ binds and therefore receives $P^B - \phi = qP^B + w$ per transaction from the buyer side. The commission is negative when $w > (1 - q)P^B$, that is, when the provider's value from a trade exceeds its expected default loss. Buyer participation is $1 - F(qP^B)$, since a buyer effectively pays the per-transaction fee P^B only with probability q , rather than with certainty as in the baseline model.

Intermediation & information provision BNPL affects the two arrangements only by changing the revenue generated by a successful merchant. Let $x \equiv qP^B$ denote a buyer's expected payment per transaction. The platform then receives $x + w$ from the buyer side of each transaction, while the merchant side generates π as before. A successful merchant therefore generates revenue $[1 - F(x)](\pi + w + x)$.

The platform's problems are given by (6) and (10), including constraints (PC) and (IC), with P^B replaced by x , π by $\pi + w$, and P^M by $P^M + w$. The last replacement adds

¹⁹We refer to Berg et al. (2025) for more institutional details of BNPL.

²⁰For tractability, we assume that this value is proportional to the number of trades. If data externalities instead made it convex in trade volume, the BNPL provider would value an additional trade more on platforms with more trades, introducing a scale effect that we abstract from.

w to the platform's revenue while leaving the merchant's pledgeable income unchanged. The financing cutoff therefore becomes

$$\underline{\theta}_{\text{BNPL}}^{\text{PF}} \equiv \frac{I}{\max_x [1 - F(x)] (\pi + w + x)}. \quad (15)$$

It is independent of q , since buyers and the platform care only about the expected payment x , while the commission absorbs the provider's default losses. When $w = 0$, the cutoff is the same as in the baseline model. A positive w lowers the cutoff because the provider passes the value it derives from each trade to the platform through a lower commission. BNPL therefore acts like a higher surplus from trade.

Regime switch & comparative statics Our results on platform profit and the regime switch therefore carry over with the cutoff in (15).

Proposition 5. *Suppose $\mathbb{E}[\theta] < \underline{\theta}_{\text{BNPL}}^{\text{PF}} < 1$ and $\alpha < 1$. Propositions 1–3 continue to hold with $\underline{\theta}^{\text{PF}}$ replaced by $\underline{\theta}_{\text{BNPL}}^{\text{PF}}$. In particular, there exists a finite M_{BNPL}^* such that $\Pi^{\text{Interm}}(M) \geq \Pi^{\text{Info}}(M)$ for all $M \geq M_{\text{BNPL}}^*$ and $\Pi^{\text{Interm}}(M) < \Pi^{\text{Info}}(M)$ for all $M < M_{\text{BNPL}}^*$.*

Corollary 6. *The cutoff $\underline{\theta}_{\text{BNPL}}^{\text{PF}}$ strictly decreases with w . Hence, for any $w_1 < w_2$ for which the conditions of Proposition 5 hold, M_{BNPL}^* is larger under w_2 than under w_1 for all α sufficiently close to one.*

Intuitively, a BNPL provider that values the trades it enables makes each successful merchant more valuable to the platform, which is therefore willing to finance lower-type merchants. This lowers the cost of overfinancing under information provision, since fewer merchants lie below the financing cutoff and those that do are closer to it. By Corollary 3, this favours information provision when data sharing is sufficiently credible.

7 Empirical and regulatory implications

Our model yields several empirically testable and regulatory implications. In this section, we first develop four empirical predictions, outlining the relevant institutional setting and testable hypotheses for each. We then discuss the regulatory implications of our results.

7.1 Empirical implications

First, our model predicts that the optimal partnership form, captured by the market-size threshold at which a regime switch occurs, is related to local conditions. This is well-suited for cross-country testing, as countries differ substantially in institutional quality

and digital infrastructure. Many emerging markets combine weak formal credit institutions, such as limited enforcement and credit bureaus ([Demirguc-Kunt et al., 2018](#); [Djankov et al., 2003](#)), with rapid digital finance adoption that generates rich transaction data outside the banking system ([Frost et al., 2020](#)). In such settings, information provision is less profitable because it relies on local banks operating under weaker institutions. Banks are less able to verify the information platforms provide, while legal sanctions for misreporting are weaker, lowering the credibility α of information provision. Intermediation, by contrast, does not rely on bank verification because the platform screens and lends itself. By Corollary 1, these forces lower the switching threshold, making intermediation optimal at smaller market sizes.

Implication 1. *In emerging markets with weaker formal institutions, platforms are predicted to switch to intermediation at a smaller merchant pool size than in developed economies.*

This prediction can be tested across countries by relating each platform’s lending model, identified from company reports, to institutional quality, its interaction with digital infrastructure, and the size of the platform’s merchant base.²¹ Such a test would reveal how platform-credit organisation responds to local conditions and clarify the interaction between platform technology and traditional institutions.

Second, the growing securitisation of platform loans provides a more direct test. Asset-backed securities (ABS) backed by platform loans can be classified by how the underlying loans were originated. Under intermediation, the platform selects borrowers and originates the loans using bank funding. Under information provision, the bank selects borrowers based on the platform’s information and originates the loans. Merchant loans sold directly to investors outside securitisations allow the same comparison. By Corollary 5, loans originated under intermediation by platforms with large merchant pools default less often than those originated under information provision.

Implication 2. *ABS created under intermediation by platforms with large merchant pools have lower credit risk, as reflected in tighter issuance yield spreads and lower ex-post default rates, than those created under information provision.*

This prediction can be tested by regressing issue-level yield spreads and default rates on the loan origination structure, hand-collected from platform reports, interacted with the size of the originating platform’s merchant base.

Third, our theory predicts a life cycle for a platform’s credit strategy. When its merchant pool is small and data sharing is sufficiently credible, information provision is optimal, allowing the platform to support trade with bank-originated credit without using

²¹These datasets include, for example, the [Worldwide Governance Indicators](#) and the [International Telecommunication Union \(ITU\) Data and statistics on digital transformation](#).

its own balance sheet. As the pool grows beyond the threshold M^* , selection becomes more effective and intermediation becomes more profitable. A later return to information provision can occur if the credibility of data sharing rises above the boundary $\alpha^*(M)$ (Corollary 2). Amazon’s move from seller loans funded partly through a bank credit facility to partner lenders with which sellers share their Amazon selling data is consistent with the last step of this prediction.

Implication 3. *The platform’s credit strategy is dynamic and non-monotonic: it begins with information provision if its data sharing is sufficiently credible, switches to intermediation as its merchant pool grows, and may return to information provision as the credibility of its data sharing improves.*

Finally, by Corollary 6, information provision is more attractive for platforms that offer BNPL financing to buyers, provided that the BNPL provider values the trades it enables ($w > 0$) and data sharing is credible. This prediction can be tested by relating each platform’s lending model to the presence of a BNPL partner at checkout, controlling for the number of active merchants.

Implication 4. *Among platforms with sufficiently credible data sharing and merchant pools of the same size, information provision is more prevalent among those that offer BNPL through providers that value the trades they enable.*

7.2 Regulatory implications

Our analysis also has implications for the regulation of platform-bank partnerships, which differs substantially across jurisdictions.

Credible data sharing Regulators can increase the credibility α of data sharing by making the data shared by platforms harder to manipulate. Examples include standardised and auditable data-access interfaces; rules allowing merchants to share their own data with lenders, as open banking does for bank account data (He et al., 2023) and Article 6(10) of the Digital Markets Act does for data generated on gatekeeper platforms;²² and the requirement in China that platforms provide personal credit information to financial institutions only through licensed credit reporting agencies.²³ These rules do not replace the platform’s assessment of merchants, but they provide lenders with additional data against which to assess the information received from the platform, reducing the scope for undetected overrides.

²²See “Regulation (EU) 2022/1925”.

²³See Yi Gang, “Speech at the Hong Kong FinTech Week 2021”: “Fintech companies are also required to carve out their personal credit information units” and “provide such services to financial institutions through licensed entities.”

In our model, a higher α reduces overfinancing under information provision and raises the switching threshold, as shown in Corollary 1. Such measures therefore improve the quality of loans underwritten using platform information and shift partnerships towards information provision. They can also substitute for investments in credibility, which platforms with large merchant pools have little incentive to make themselves, as discussed in Remark 2. The welfare effect is less clear. As shown in Section 5.4, for a platform that remains an information provider, a higher α raises welfare if and only if merchants below the platform’s financing cutoff are, on average, not worth financing from a welfare perspective.

Heterogeneous regulatory treatments Since the optimal arrangement depends on merchant pool size, rules that impose the same arrangement on all platforms have different effects across platforms. Restrictions on sharing merchant data with lenders effectively force platforms into intermediation, which is costly for platforms with small merchant pools. As shown in Proposition 1, with a single merchant, intermediation shuts down lending entirely, which also reduces welfare by Proposition 4. Conversely, restricting platforms to information provision, for example by separating commerce from lending, forces platforms with large merchant pools into an arrangement that overfinances low-quality merchants and produces lower-quality loans, as shown in Corollary 5. Such a restriction can nevertheless raise welfare in large pools if information provision is profitable and merchants below the platform’s financing cutoff are, on average, worth financing from a welfare perspective, as shown in Proposition 4. Regulation of platform lending should therefore take merchant pool size into account.

Direct sales channels Rules that allow merchants to steer platform customers to their own sales channels, such as Article 5(4) of the Digital Markets Act, reduce the cost of selling to these customers off the platform. Our extension in Appendix D shows that their effect on merchant credit depends on how repayment is collected. If banks can collect repayment from merchants’ off-platform sales, these rules shift partnerships towards information provision, under which banks lend directly to merchants, without reducing the credit available under this arrangement. If repayment can be collected only from merchants’ on-platform sales, they raise the financing cutoff and reduce credit under both arrangements. Allowing lenders to collect repayment across sales channels can therefore prevent this contraction.

8 Conclusion

We develop a theory of platform-bank partnerships in which the platform has information about merchants and the bank provides capital. The platform also benefits from the

trades generated by financed merchants, while the bank does not. Our main result is that the platform's choice between intermediating bank funding and providing information to a lending bank depends on the size of its borrower pool and the credibility of its data sharing.

Under intermediation, the platform selects the best merchants from the pool. As the pool grows, selection becomes more effective, so profit per merchant approaches the monopoly level, while the total loss remains bounded. Under information provision, the platform recommends merchants through a reporting rule. Limited commitment leads to overfinancing whenever the platform can override the rule. Profit per merchant is therefore constant, so the total loss grows with the pool. As a result, when data sharing is sufficiently credible, information provision is optimal for small borrower pools and intermediation for large pools. The switching threshold increases with credibility and grows without bound under full commitment. The switching threshold also depends on the platform's financing cutoff, which in turn depends on investment costs, trade surplus, and buyer demand, as well as on the distribution of merchant types.

The platform's preferred arrangement need not maximise welfare because it does not internalise buyer surplus. In large pools, information provision can therefore yield higher welfare even when the platform prefers intermediation. We also show that BNPL financing on the buyer side can make information provision more attractive for merchant lending, and that direct sales channels can shift lending from platforms to banks when banks can collect repayment outside the platform.

Our results suggest that the form of a platform-bank partnership can change as the platform grows. They also imply that regulation should take merchant pool size and the credibility of data sharing into account. Future work could allow banks and platforms to hold different information, introduce credit risk retained by the platform, or study how platform capital, data, and partnership choice evolve over time.

Appendix A Proofs

By (1), financing a type- θ merchant at the fee P^{B*} yields $\frac{I}{\underline{\theta}^{\text{PF}}}(\theta - \underline{\theta}^{\text{PF}})$ to the platform when the bank breaks even.

Proof for Lemma 1. Under the credit line, the repayment of each financed merchant goes to the bank, while the platform earns the fees $[1 - F(P^B)](P^M + P^B)$ when the merchant succeeds. Financing a type- θ_m merchant therefore yields the platform $\theta_m[1 - F(P^B)](P^M + P^B)$, which is positive, strictly increasing in θ_m , and independent of which other merchants are financed. Hence the platform uses the full credit line and finances merchants in descending order of type. $\square$

Proof for Lemma 2. The expected number of successes among the n best merchants is the sum of their expected types,

$$\sum_{a=0}^n p_a(n) a = \sum_{j=M-n+1}^M \mathbb{E}[\theta_{j:M}],$$

so the left-hand side of the bank's participation constraint is the average expected type of the n best merchants. Since $\mathbb{E}[\theta_{j:M}]$ strictly increases in j , adding the $(n+1)$ -th best merchant lowers this average. The left-hand side therefore strictly decreases in n , whereas the right-hand side, $\underline{\theta}^{\text{Bank}}(P^M, P^B)$, does not depend on n . Hence the constraint holds if and only if $n \leq n^*(M)$, where $n^*(M) = 0$ if $\mathbb{E}[\theta_{M:M}] < \underline{\theta}^{\text{Bank}}(P^M, P^B)$. $\square$

Proof for Lemma 3. For any realisation of types, the monopolist earns $\frac{I}{\underline{\theta}^{\text{PF}}} \sum_m \max\{\theta_m - \underline{\theta}^{\text{PF}}, 0\}$. Under intermediation, the platform finances the set T of the $n^*(M)$ merchants with the highest types, and by (7) it earns the expectation of $\frac{I}{\underline{\theta}^{\text{PF}}} \sum_{m \in T} (\theta_m - \underline{\theta}^{\text{PF}})$. For every realisation, $\sum_{m \in T} (\theta_m - \underline{\theta}^{\text{PF}}) \leq \sum_m \max\{\theta_m - \underline{\theta}^{\text{PF}}, 0\}$, with strict inequality whenever the number of merchants with $\theta_m > \underline{\theta}^{\text{PF}}$ differs from $n^*(M)$: the platform then either leaves a merchant with $\theta_m > \underline{\theta}^{\text{PF}}$ unfunded or finances one with $\theta_m < \underline{\theta}^{\text{PF}}$. This number is binomial with parameters M and $1 - G(\underline{\theta}^{\text{PF}}) \in (0, 1)$, so it differs from $n^*(M)$ with positive probability. Taking expectations gives $\Pi^{\text{Interm}}(M) < \Pi^{\text{Mono}}(M)$. $\square$

Proof for Proposition 1. Part (i). For $M = 1$, $\mathbb{E}[\theta_{1:1}] = \mathbb{E}[\theta] < \underline{\theta}^{\text{PF}}$ by Assumption 2, so $n^*(1) = 0$ and $\Pi^{\text{Interm}}(1) = 0 = \Pi^{\text{Bank-lend}}(1)$.

Part (ii). Expected order statistics satisfy the recurrence

$$\mathbb{E}[\theta_{j:M-1}] = \frac{j}{M} \mathbb{E}[\theta_{j+1:M}] + \frac{M-j}{M} \mathbb{E}[\theta_{j:M}], \quad j = 1, \dots, M-1,$$

which follows from deleting one of M draws at random (David & Nagaraja, 2003). Since

the function $x \mapsto \max\{x - \underline{\theta}^{\text{PF}}, 0\}$ is convex,

$$\max\{\mathbb{E}[\theta_{j:M-1}] - \underline{\theta}^{\text{PF}}, 0\} \leq \frac{j}{M} \max\{\mathbb{E}[\theta_{j+1:M}] - \underline{\theta}^{\text{PF}}, 0\} + \frac{M-j}{M} \max\{\mathbb{E}[\theta_{j:M}] - \underline{\theta}^{\text{PF}}, 0\}.$$

Summing over j , each term $\max\{\mathbb{E}[\theta_{k:M}] - \underline{\theta}^{\text{PF}}, 0\}$ on the right-hand side receives the total weight $\frac{k-1}{M} + \frac{M-k}{M} = \frac{M-1}{M}$. By (8), $\Pi^{\text{Interm}}(M-1) \leq \frac{M-1}{M} \Pi^{\text{Interm}}(M)$, so profit per merchant is non-decreasing in M . Its limit follows from part (iii), since $\Pi^{\text{Mono}}(M)/M = \Pi^{\text{Mono}}(1)$.

Part (iii). Since profit per merchant is non-decreasing and non-negative, total profit is non-decreasing. For the limit, note that $\sum_j \mathbb{E}[\max\{\theta_{j:M} - \underline{\theta}^{\text{PF}}, 0\}] = M \mathbb{E}[\max\{\theta - \underline{\theta}^{\text{PF}}, 0\}]$, since the order statistics are a rearrangement of the M types. By (5) and (8),

$$\Pi^{\text{Mono}}(M) - \Pi^{\text{Interm}}(M) = \frac{I}{\underline{\theta}^{\text{PF}}} \sum_{j=1}^M \left(\mathbb{E}[\max\{\theta_{j:M} - \underline{\theta}^{\text{PF}}, 0\}] - \max\{\mathbb{E}[\theta_{j:M}] - \underline{\theta}^{\text{PF}}, 0\} \right). \quad (16)$$

Each term is the gap in Jensen's inequality for the convex function $\max\{x - \underline{\theta}^{\text{PF}}, 0\}$, which is positive only if $\theta_{j:M}$ can fall on either side of $\underline{\theta}^{\text{PF}}$. For large M , this is the case only for ranks j with j/M close to $G(\underline{\theta}^{\text{PF}})$. For these ranks, $\theta_{j:M}$ is approximately normal with standard deviation $s = \sqrt{G(\underline{\theta}^{\text{PF}})[1 - G(\underline{\theta}^{\text{PF}})]/M / g(\underline{\theta}^{\text{PF}})}$, and the means of consecutive order statistics are $1/[Mg(\underline{\theta}^{\text{PF}})]$ apart (David & Nagaraja, 2003). For a normal random variable X with mean μ and standard deviation s ,

$$\mathbb{E}[\max\{X - \underline{\theta}^{\text{PF}}, 0\}] - \max\{\mu - \underline{\theta}^{\text{PF}}, 0\} = s \psi(|\mu - \underline{\theta}^{\text{PF}}|/s),$$

where $\psi(z) \equiv \mathbb{E}[\max\{Z - z, 0\}]$ for a standard normal Z . Summing over ranks amounts to integrating over μ and dividing by the spacing $1/[Mg(\underline{\theta}^{\text{PF}})]$. Since $\int_0^\infty \psi(z) dz = 1/4$, the sum in (16) converges to

$$Mg(\underline{\theta}^{\text{PF}}) \int_{-\infty}^\infty s \psi(|\mu - \underline{\theta}^{\text{PF}}|/s) d\mu = 2Mg(\underline{\theta}^{\text{PF}}) s^2 \int_0^\infty \psi(z) dz = \frac{G(\underline{\theta}^{\text{PF}})[1 - G(\underline{\theta}^{\text{PF}})]}{2g(\underline{\theta}^{\text{PF}})}.$$

Hence

$$C = \frac{I}{\underline{\theta}^{\text{PF}}} \cdot \frac{G(\underline{\theta}^{\text{PF}})[1 - G(\underline{\theta}^{\text{PF}})]}{2g(\underline{\theta}^{\text{PF}})} > 0. \quad (17)$$

The term $G(\underline{\theta}^{\text{PF}})[1 - G(\underline{\theta}^{\text{PF}})]$ is the variance of the indicator that a merchant is worth financing, which determines how much the number of such merchants fluctuates around the credit line. The term $1/g(\underline{\theta}^{\text{PF}})$ determines how far misallocated merchants lie from the cutoff. Hence C is larger when the cutoff is closer to the median type and when few merchants have types close to the cutoff. $\square$

Proof for Lemma 4. Write $\underline{\theta}^{\text{Bank}}$ for $\underline{\theta}^{\text{Bank}}(P^M, P^B)$. Since types are independent and the bank decides on each merchant on the basis of the report on that merchant, it suffices to

consider a single merchant.

Two reports suffice. Take any reporting rule and the bank's best response to it. Merging all reports after which the bank lends into a single report s' , and all other reports into the null report, leaves the bank's decisions unchanged, since the posterior after s' is an average of posteriors at which the bank lends. When it can override the rule, the platform sends s' , since it earns expected fees $\theta[1 - F(P^B)](P^M + P^B) > 0$ from a financed merchant. A reporting rule can therefore be described by the probability $\sigma(\theta) \in [0, 1]$ with which it specifies s' for type θ , and type θ receives s' with probability $\alpha\sigma(\theta) + 1 - \alpha$.

The cutoff. The platform maximises its fee income, which is proportional to the expected type of a financed merchant times the probability of financing, subject to the bank lending after s' :

$$\max_{\sigma} \int_0^1 \theta [\alpha\sigma(\theta) + 1 - \alpha] dG(\theta) \quad \text{s.t.} \quad \int_0^1 (\theta - \underline{\theta}^{\text{Bank}}) [\alpha\sigma(\theta) + 1 - \alpha] dG(\theta) \geq 0.$$

For any σ , let θ' satisfy $1 - G(\theta') = \int_0^1 \sigma(\theta) dG(\theta)$. The cutoff rule $\mathbf{1}\{\theta \geq \theta'\}$ specifies s' with the same total probability but for higher types, since

$$\int_{\theta'}^1 \theta dG(\theta) - \int_0^1 \theta \sigma(\theta) dG(\theta) = \int_0^1 (\theta - \theta') [\mathbf{1}\{\theta \geq \theta'\} - \sigma(\theta)] dG(\theta) \geq 0.$$

It therefore raises both the objective and the left-hand side of the constraint. Among cutoff rules, the objective decreases in the cutoff, and by (9) the constraint is $\mathbb{E}[\theta \mid s'(\theta')] \geq \underline{\theta}^{\text{Bank}}$. Hence the optimal cutoff is the smallest θ' that satisfies this condition, and if none does, no reporting rule induces lending.

Shape of the posterior. Differentiating (9),

$$\frac{\partial \mathbb{E}[\theta \mid s'(\theta')]}{\partial \theta'} = \frac{\alpha g(\theta')}{1 - \alpha G(\theta')} \left(\mathbb{E}[\theta \mid s'(\theta')] - \theta' \right).$$

For $\alpha < 1$, the term in brackets equals $\mathbb{E}[\theta] > 0$ at $\theta' = 0$ and $\mathbb{E}[\theta] - 1 < 0$ at $\theta' = 1$, and its derivative equals -1 wherever it is zero, so it changes sign exactly once. The posterior therefore first increases and then decreases in θ' , and it is increasing at every θ' with $\mathbb{E}[\theta \mid s'(\theta')] > \theta'$. For $\alpha = 1$, the posterior is $\mathbb{E}[\theta \mid \theta \geq \theta']$, which increases in θ' . $\square$

Proof for Proposition 2. Binding (PC) gives $P^M = \pi - I / \{[1 - F(P^B)] \mathbb{E}[\theta \mid s'(\theta')]\}$, and substituting into (10) yields (11), which can be written as

$$M \int_0^1 \left(\theta [1 - F(P^B)](\pi + P^B) - I \right) [\alpha \mathbf{1}\{\theta \geq \theta'\} + 1 - \alpha] dG(\theta).$$

For any θ' , this is increasing in $[1 - F(P^B)](\pi + P^B)$, so $P^B = P^{B*}$, and the integrand

becomes $\frac{I}{\underline{\theta}^{\text{PF}}}(\theta - \underline{\theta}^{\text{PF}}) [\alpha \mathbf{1}\{\theta \geq \theta'\} + 1 - \alpha]$. The objective is therefore maximised by $\theta' = \underline{\theta}^{\text{PF}}$, which puts the additional weight α on exactly the types with $\theta \geq \underline{\theta}^{\text{PF}}$. Integrating by parts, the maximum per merchant is

$$\frac{I}{\underline{\theta}^{\text{PF}}} \int_{\underline{\theta}^{\text{PF}}}^1 (\theta - \underline{\theta}^{\text{PF}}) dG(\theta) - (1 - \alpha) \frac{I}{\underline{\theta}^{\text{PF}}} \int_0^{\underline{\theta}^{\text{PF}}} (\underline{\theta}^{\text{PF}} - \theta) dG(\theta) = \Pi^{\text{Mono}}(1) - (1 - \alpha) \frac{I}{\underline{\theta}^{\text{PF}}} \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta.$$

If this value is negative, the platform prefers no lending, which it obtains by setting P^M so high that the bank never lends. This gives (12) and linearity in M .

The value above equals $[1 - \alpha G(\underline{\theta}^{\text{PF}})] \frac{I}{\underline{\theta}^{\text{PF}}} (\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})] - \underline{\theta}^{\text{PF}})$. When it is positive, $\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})] > \underline{\theta}^{\text{PF}}$, so by the proof of Lemma 4 the posterior is increasing on $[0, \underline{\theta}^{\text{PF}}]$, and $\underline{\theta}^{\text{PF}}$ is indeed the smallest cutoff that satisfies (PC), as (IC) requires. Moreover, $P^{M*} + P^{B*} = \frac{I}{1 - F(P^{B*})} \left(\frac{1}{\underline{\theta}^{\text{PF}}} - \frac{1}{\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})]} \right) > 0$, as Lemma 4 assumes.

Finally, since $\int_0^1 G(\theta) d\theta = 1 - \mathbb{E}[\theta]$, we have $\int_{\underline{\theta}^{\text{PF}}}^1 (\theta - \underline{\theta}^{\text{PF}}) dG(\theta) = \int_{\underline{\theta}^{\text{PF}}}^1 [1 - G(\theta)] d\theta = \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta - (\underline{\theta}^{\text{PF}} - \mathbb{E}[\theta])$. The value per merchant therefore equals $\frac{I}{\underline{\theta}^{\text{PF}}} [\alpha \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta - (\underline{\theta}^{\text{PF}} - \mathbb{E}[\theta])]$, which is positive if and only if $\alpha > \underline{\alpha}$. By Assumption 2, $\underline{\theta}^{\text{PF}} > \mathbb{E}[\theta]$, so $\underline{\alpha} > 0$, and $\underline{\theta}^{\text{PF}} < 1$ gives $\int_{\underline{\theta}^{\text{PF}}}^1 [1 - G(\theta)] d\theta > 0$, so $\underline{\alpha} < 1$. $\square$

Proof for Proposition 3. By Proposition 2, intermediation yields weakly higher profit if and only if

$$\Pi^{\text{Interm}}(M)/M \geq \Pi^{\text{Info}}(1).$$

By Proposition 1(ii), the left-hand side is non-decreasing in M and converges to $\Pi^{\text{Mono}}(1)$, which exceeds $\Pi^{\text{Info}}(1)$ for $\alpha < 1$ by (12). Hence the inequality holds for all M at or above a finite threshold M^* and fails for all M below it. If $\alpha > \underline{\alpha}$, then $\Pi^{\text{Interm}}(1) = 0 < \Pi^{\text{Info}}(1)$, so $M^* \geq 2$. $\square$

Proof for Corollary 1. By (12), $\Pi^{\text{Info}}(1)$ is non-decreasing in α , whereas $\Pi^{\text{Interm}}(M)$ does not depend on α . Since M^* is the smallest M with $\Pi^{\text{Interm}}(M)/M \geq \Pi^{\text{Info}}(1)$, it is non-decreasing in α . As $\alpha \rightarrow 1$, $\Pi^{\text{Info}}(1) \rightarrow \Pi^{\text{Mono}}(1)$, which exceeds $\Pi^{\text{Interm}}(M)/M$ for every M by Lemma 3. Hence, for any $\bar{M}$, $M^* > \bar{M}$ for α close enough to one. $\square$

Proof for Corollary 2. Let $A(\underline{\theta}^{\text{PF}}) \equiv \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta$ and

$$\alpha^*(M) \equiv 1 - \frac{\Pi^{\text{Mono}}(1) - \Pi^{\text{Interm}}(M)/M}{\frac{I}{\underline{\theta}^{\text{PF}}} A(\underline{\theta}^{\text{PF}})}.$$

By (12), $\Pi^{\text{Info}}(1) = \max\{\Pi^{\text{Mono}}(1) - (1 - \alpha) \frac{I}{\underline{\theta}^{\text{PF}}} A(\underline{\theta}^{\text{PF}}), 0\}$. Since $\Pi^{\text{Interm}}(M) \geq 0$, $\Pi^{\text{Info}}(1) > \Pi^{\text{Interm}}(M)/M$ if and only if $\Pi^{\text{Mono}}(1) - (1 - \alpha) \frac{I}{\underline{\theta}^{\text{PF}}} A(\underline{\theta}^{\text{PF}}) > \Pi^{\text{Interm}}(M)/M$, that is, $\alpha > \alpha^*(M)$. By Proposition 1(ii), $\alpha^*(M)$ is non-decreasing in M and converges to one, and by Lemma 3, it is below one. For $M = 1$, $\Pi^{\text{Interm}}(1) = 0$ and, by the proof of Proposition 2, $\Pi^{\text{Mono}}(1) = \frac{I}{\underline{\theta}^{\text{PF}}} [A(\underline{\theta}^{\text{PF}}) - (\underline{\theta}^{\text{PF}} - \mathbb{E}[\theta])]$, so $\alpha^*(1) = \underline{\alpha}$. $\square$

Proof for Corollary 3. The fee P^{B*} maximises $[1 - F(P^B)](\pi + P^B)$, whose maximum equals $I/\underline{\theta}^{\text{PF}}$. By (8) and (12), both profits equal $I/\underline{\theta}^{\text{PF}}$ times functions of $\underline{\theta}^{\text{PF}}$, G and α only, so M^* depends on I , π and F only through $\underline{\theta}^{\text{PF}}$. The cutoff increases in I . It decreases in π , since, by the envelope theorem, the maximum increases in π at the rate $1 - F(P^{B*}) > 0$. A first-order stochastically dominant shift in F weakly raises the revenue $[1 - F(P^{B*})](\pi + P^{B*})$ at the original fee, since $\pi + P^{B*} > 0$, and hence the maximum, so it weakly lowers $\underline{\theta}^{\text{PF}}$.

With $A(\underline{\theta}^{\text{PF}})$ as in the proof of Corollary 2, let $\Delta(M) \equiv \Pi^{\text{Mono}}(M) - \Pi^{\text{Interm}}(M)$ and $m(\underline{\theta}^{\text{PF}}) \equiv G(\underline{\theta}^{\text{PF}})[1 - G(\underline{\theta}^{\text{PF}})]/[2g(\underline{\theta}^{\text{PF}})A(\underline{\theta}^{\text{PF}})]$. We first show that $(1 - \alpha)M^* \rightarrow m(\underline{\theta}^{\text{PF}})$ as $\alpha \rightarrow 1$. For $\alpha > \underline{\alpha}$, by (12), $\Pi^{\text{Interm}}(M) \geq \Pi^{\text{Info}}(M)$ if and only if $\Delta(M)/M \leq (1 - \alpha)\frac{I}{\underline{\theta}^{\text{PF}}}A(\underline{\theta}^{\text{PF}})$, where $\Delta(M)/M = \Pi^{\text{Mono}}(1) - \Pi^{\text{Interm}}(M)/M$ is non-increasing. Since $M^* \geq 2$ for α close to one, the definition of M^* gives

$$\frac{\Delta(M^*)}{M^*} \leq (1 - \alpha)\frac{I}{\underline{\theta}^{\text{PF}}}A(\underline{\theta}^{\text{PF}}) < \frac{\Delta(M^* - 1)}{M^* - 1}.$$

Multiplying by M^* , $(1 - \alpha)M^*\frac{I}{\underline{\theta}^{\text{PF}}}A(\underline{\theta}^{\text{PF}})$ lies between $\Delta(M^*)$ and $\frac{M^*}{M^* - 1}\Delta(M^* - 1)$. As $\alpha \rightarrow 1$, $M^* \rightarrow \infty$ by Corollary 1, so both bounds converge to C by Proposition 1(iii). Hence, by (17), $(1 - \alpha)M^* \rightarrow C/[\frac{I}{\underline{\theta}^{\text{PF}}}A(\underline{\theta}^{\text{PF}})] = m(\underline{\theta}^{\text{PF}})$.

Next, m is strictly decreasing. Indeed,

$$\frac{d \ln m(\underline{\theta}^{\text{PF}})}{d \underline{\theta}^{\text{PF}}} = \left[\frac{g(\underline{\theta}^{\text{PF}})}{G(\underline{\theta}^{\text{PF}})} - \frac{G(\underline{\theta}^{\text{PF}})}{A(\underline{\theta}^{\text{PF}})} \right] - \left[\frac{g(\underline{\theta}^{\text{PF}})}{1 - G(\underline{\theta}^{\text{PF}})} + \frac{g'(\underline{\theta}^{\text{PF}})}{g(\underline{\theta}^{\text{PF}})} \right].$$

Since g is log-concave, so is G (Bagnoli & Bergstrom, 2005), so $G(\theta) \leq G(\underline{\theta}^{\text{PF}}) \exp\{-(\underline{\theta}^{\text{PF}} - \theta)g(\underline{\theta}^{\text{PF}})/G(\underline{\theta}^{\text{PF}})\}$ for $\theta \leq \underline{\theta}^{\text{PF}}$. Integrating over $[0, \underline{\theta}^{\text{PF}}]$ gives $A(\underline{\theta}^{\text{PF}}) < G(\underline{\theta}^{\text{PF}})^2/g(\underline{\theta}^{\text{PF}})$, so the first bracket is negative. The second bracket is the derivative of the log of the hazard rate $g/(1 - G)$, which is non-decreasing since g is log-concave (Bagnoli & Bergstrom, 2005), so the second bracket is non-negative.

Finally, consider two cutoffs $\underline{\theta}^{\text{PF}} < \widehat{\underline{\theta}}^{\text{PF}}$ in $(\mathbb{E}[\theta], 1)$, with thresholds M^* and $\widehat{M}^*$. Since $(1 - \alpha)M^* \rightarrow m(\underline{\theta}^{\text{PF}})$, $(1 - \alpha)\widehat{M}^* \rightarrow m(\widehat{\underline{\theta}}^{\text{PF}})$ and $m(\underline{\theta}^{\text{PF}}) > m(\widehat{\underline{\theta}}^{\text{PF}})$, we have $M^* > \widehat{M}^*$ for all α sufficiently close to one. $\square$

Proof for Corollary 4. Let H be a mean-preserving spread of G , so that H has the same mean and $\int \varphi dH \geq \int \varphi dG$ for every convex function φ (Rothschild & Stiglitz, 1970). The cutoff $\underline{\theta}^{\text{PF}}$ does not depend on the type distribution.

Under intermediation, by (8), $\Pi^{\text{Interm}}(M) = \frac{I}{\underline{\theta}^{\text{PF}}} \max_n \{\mathbb{E}[Y_n] - n \underline{\theta}^{\text{PF}}\}$, where Y_n is the sum of the n highest of the M types. As the maximum of the sums of any n types, Y_n is convex in each type. Replacing the types drawn from G one at a time by independent draws from H therefore weakly raises $\mathbb{E}[Y_n]$ for every n , and hence $\Pi^{\text{Interm}}(M)$.

Under information provision, by the proof of Proposition 2,

$$\Pi^{\text{Info}}(1) = \max \left\{ \frac{I}{\underline{\theta}^{\text{PF}}} \left[\alpha \int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta - (\underline{\theta}^{\text{PF}} - \mathbb{E}[\theta]) \right], 0 \right\}.$$

Since $\int_0^{\underline{\theta}^{\text{PF}}} G(\theta) d\theta = \mathbb{E}[\max\{\underline{\theta}^{\text{PF}} - \theta, 0\}]$ is the expectation of a convex function, it weakly rises under H , while $\mathbb{E}[\theta]$ is unchanged. Hence $\Pi^{\text{Info}}(M)$ weakly rises.

The following numerical examples show that the effect on M^* is ambiguous. Symmetric Beta distributions $\text{Beta}(a, a)$ with $a \geq 1$ have log-concave densities and mean $1/2$, and a lower a gives a mean-preserving spread. With $\underline{\theta}^{\text{PF}} = 0.6$ and $\alpha = 0.95$, $M^* = 9$ under $\text{Beta}(3, 3)$ and $M^* = 11$ under $\text{Beta}(2, 2)$. With $\underline{\theta}^{\text{PF}} = 0.76$ and $\alpha = 0.98$, $M^* = 9$ under $\text{Beta}(3, 3)$ and $M^* = 8$ under $\text{Beta}(2.5, 2.5)$. $\square$

Proof for Corollary 5. For large M , $\mathbb{E}[\theta_{j:M}]$ is close to $G^{-1}(j/M)$, so the ranks financed under intermediation, those with $\mathbb{E}[\theta_{j:M}] \geq \underline{\theta}^{\text{PF}}$, are approximately those with $j/M \geq G(\underline{\theta}^{\text{PF}})$. Hence $n^*(M)/M \rightarrow 1 - G(\underline{\theta}^{\text{PF}})$, and the average expected type of the financed merchants converges to

$$\frac{1}{1 - G(\underline{\theta}^{\text{PF}})} \int_{G(\underline{\theta}^{\text{PF}})}^1 G^{-1}(p) dp = \mathbb{E}[\theta \mid \theta \geq \underline{\theta}^{\text{PF}}].$$

By (9), $\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})]$ is a weighted average of $\mathbb{E}[\theta \mid \theta \geq \underline{\theta}^{\text{PF}}]$ and $\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}]$, with weight $(1 - \alpha)G(\underline{\theta}^{\text{PF}})/[1 - \alpha G(\underline{\theta}^{\text{PF}})]$ on the latter, which is positive for $\alpha < 1$. Since $\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}] < \mathbb{E}[\theta \mid \theta \geq \underline{\theta}^{\text{PF}}]$, $\mathbb{E}[\theta \mid s'(\underline{\theta}^{\text{PF}})] < \mathbb{E}[\theta \mid \theta \geq \underline{\theta}^{\text{PF}}]$. $\square$

Proof for Proposition 4. By (13), financing a type- θ merchant adds $\frac{I}{\underline{\theta}^{\text{W}}}(\theta - \underline{\theta}^{\text{W}})$ to welfare. Under intermediation, welfare is $W^{\text{Interm}}(M) = \frac{I}{\underline{\theta}^{\text{W}}} \sum_{j=M-n^*(M)+1}^M (\mathbb{E}[\theta_{j:M}] - \underline{\theta}^{\text{W}})$, and under information provision, welfare $W^{\text{Info}}(M)$ is M times (14).

Part (i). Since $n^*(1) = 0$, $W^{\text{Interm}}(1) = 0$. For $\alpha > \underline{\alpha}$, the proof of Proposition 2 gives

$$\int_0^1 (\theta - \underline{\theta}^{\text{PF}}) [\alpha \mathbf{1}\{\theta \geq \underline{\theta}^{\text{PF}}\} + 1 - \alpha] dG(\theta) > 0.$$

Since $\underline{\theta}^{\text{W}} < \underline{\theta}^{\text{PF}}$, replacing $\theta - \underline{\theta}^{\text{PF}}$ by $\theta - \underline{\theta}^{\text{W}}$ in this integral makes it larger, and the result equals $\frac{\underline{\theta}^{\text{W}}}{I} W^{\text{Info}}(1)$. Hence $W^{\text{Info}}(1) > 0$.

Part (ii). Rearranging terms,

$$\begin{aligned} \frac{\underline{\theta}^{\text{W}}}{I} [W^{\text{Info}}(M) - W^{\text{Interm}}(M)] &= (1 - \alpha) M G(\underline{\theta}^{\text{PF}}) (\mathbb{E}[\theta \mid \theta < \underline{\theta}^{\text{PF}}] - \underline{\theta}^{\text{W}}) \\ &\quad + \frac{\underline{\theta}^{\text{PF}}}{I} [\Pi^{\text{Mono}}(M) - \Pi^{\text{Interm}}(M)] \\ &\quad + (\underline{\theta}^{\text{PF}} - \underline{\theta}^{\text{W}}) [M(1 - G(\underline{\theta}^{\text{PF}})) - n^*(M)]. \end{aligned}$$

Since the second term converges to $\frac{\theta^{\text{PF}}}{I}C$ by Proposition 1(iii) and $n^*(M)/M \rightarrow 1 - G(\theta^{\text{PF}})$ by the proof of Corollary 5, the second and third terms divided by M converge to zero. The difference in welfare per merchant therefore converges to $\frac{I}{\theta^{\text{W}}}(1 - \alpha)G(\theta^{\text{PF}})(\mathbb{E}[\theta \mid \theta < \theta^{\text{PF}}] - \theta^{\text{W}})$, which has the sign of $\mathbb{E}[\theta \mid \theta < \theta^{\text{PF}}] - \theta^{\text{W}}$. $\square$

Proof for Proposition 5 and Corollary 6. The platform's revenue decreases in ϕ , so it sets ϕ such that $(\text{PC}_{\text{BNPL}})$ binds, and $P^B - \phi = qP^B + w$. Let $x \equiv qP^B$. Since $q > 0$, choosing P^B is equivalent to choosing x . Buyer participation is $1 - F(x)$, the platform earns $[1 - F(x)](P^M + w + x)$ from the trades of a successful merchant, and the merchant's income pledgeable to a lender is $[1 - F(x)](\pi - P^M)$. The extended model therefore coincides with the baseline model once π , P^B and P^M are replaced by $\pi + w$, x and $P^M + w$, respectively. Since P^M is chosen freely, the proofs of Propositions 1–3 apply with θ^{PF} replaced by $\theta_{\text{BNPL}}^{\text{PF}}$.

For the corollary, let $w_1 < w_2$, and let x_1 maximise $[1 - F(x)](\pi + w_1 + x)$. Since $\theta_{\text{BNPL}}^{\text{PF}} < 1$ under w_1 , this maximum is positive, so $1 - F(x_1) > 0$. Hence

$$\max_x [1 - F(x)](\pi + w_2 + x) \geq [1 - F(x_1)](\pi + w_2 + x_1) > [1 - F(x_1)](\pi + w_1 + x_1),$$

and the cutoff is strictly lower under w_2 . Corollary 3 then implies that M_{BNPL}^* is larger under w_2 for all α sufficiently close to one. $\square$

Appendix B Information provision technology

Practical interpretation Our formulation reflects how platforms share merchant information with partner lenders in practice. Platforms often determine which merchants are invited to apply and what data lenders can access, while lenders make their own credit decisions.²⁴ These eligibility and data-sharing rules motivate the reporting rule in our model.

The platform cannot fully commit to this rule. It controls the interface, eligibility lists, and data used to generate recommendations, and can therefore change what is reported after observing merchant types. Hard transaction data could instead be transmitted through an interface that cannot be changed ex post, but the platform's advantage also comes from soft assessments of merchant prospects (Liberti & Petersen, 2019). Banks therefore verify the information they receive, but such verification need not detect every

²⁴Amazon stopped underwriting term loans for its U.S. and U.K. sellers in March 2024. Sellers now check Seller Central for financing invitations from “trusted third-party providers”, agree to share their Amazon selling data with the provider, and may be asked for additional information. See Fortune (7 March 2024) and “Amazon Lending”. Similarly, for DoorDash Capital, “Parafin determines eligibility, offer amounts, and other terms of each advance”, while “DoorDash does not provide or fund merchant cash advances or make approval decisions”; see “DoorDash Capital”.

deviation.²⁵ We therefore interpret α as the credibility of data sharing. Full commitment, $\alpha = 1$, is a special case and eliminates the trade-off studied in Section 5.1.

Partially verifiable communication For the equilibrium studied in Section 4, our reporting technology is outcome-equivalent to partially verifiable communication. Under this alternative formulation, the platform chooses reports after observing merchant types. A false report, meaning one that does not contain the merchant’s type, is detected with probability α and then not delivered to the bank. In either formulation, a merchant below the cutoff of the equilibrium report receives a favourable report with probability $1 - \alpha$. The bank therefore has the same posterior beliefs and makes the same lending decisions.

When $\alpha = 1$, this formulation becomes fully verifiable disclosure (Milgrom, 1981). It is related to the literature on partially verifiable communication starting with Green & Laffont (1986). Subsequent work studies particular limits on verification. Glazer & Rubinstein (2004), for example, consider a receiver who can inspect only part of the sender’s information, while Silva (2024) studies verification through comparison of the message with an external signal. We use the commitment formulation because it directly links the bank’s beliefs to the reporting rule chosen before merchant types are observed.

Merchant-level reporting We assume that the bank’s lending decision for each merchant depends only on that merchant’s report. Allowing the reporting rule to randomise, or to make a merchant’s report depend on other merchants’ types, does not change our results as long as lending remains merchant-specific. Such rules affect a merchant only through the probability of receiving a favourable report conditional on its type, which the proof of Lemma 4 already allows for.

This restriction matters if one lender observes all reports and can condition lending jointly across merchants. For example, a rule recommending the n best merchants could make an override detectable from the total number of recommendations, since an override only adds a recommendation. If the bank instead capped the number of recommended merchants it finances, the platform would recommend its best merchants up to the cap, since it earns fees only on successful merchants, and the arrangement would coincide with intermediation with a credit line for that number of merchants. This form of cross-merchant inference is less relevant when applications are underwritten separately or when

²⁵Bank verification of third-party borrower information is required by regulatory guidance. In the Amazon–Goldman Sachs partnership, for example, Goldman applied its own credit process using seller data provided by Amazon; see “Amazon, Goldman will offer small-business lines of credit”. U.S. banking agencies similarly state that “the use of third parties does not diminish or remove banking organizations’ responsibilities”; see “Interagency Guidance on Third-Party Relationships: Risk Management”. Verification can nevertheless be imperfect. According to the U.S. Department of Justice, LendingClub made misrepresentations to WebBank in 2009–2010 to increase loan volume, and WebBank originated more than 200 loans to borrowers who did not meet its credit requirements. LendingClub paid a \$2 million civil penalty in 2018 without admitting liability; see “Lending Club agrees to pay \$2 million penalty”.

a platform refers merchants to several lenders, each observing only its own applicants.²⁶

Appendix C Two-way cross-sided network externalities

In the baseline model, buyers join the platform regardless of how many merchants are financed (Section 2.1). We now introduce a fixed cost $\kappa > 0$ for buyers to join the platform.²⁷ Buyer b then joins if and only if $(v_b - P^B)S \geq \kappa$, where S is the expected number of successful merchants on the platform, so that

$$N^B = 1 - F\left(P^B + \frac{\kappa}{S}\right). \quad (18)$$

Financing more merchants, or better ones, now attracts more buyers, which raises the platform's revenue from every successful merchant. This introduces a second network externality, from merchants to buyers.

Let

$$R(S) \equiv \max_{P^B} \left[1 - F\left(P^B + \frac{\kappa}{S}\right)\right](\pi + P^B)$$

denote the platform's maximal revenue from a successful merchant. $R(S)$ increases in S , lies below its baseline counterpart I/θ^{PF} , and converges to it as $S \rightarrow \infty$. Under intermediation, a credit line nI gives $S_n = \sum_{j=M-n+1}^M \mathbb{E}[\theta_{j:M}]$, and the platform solves

$$\begin{aligned} \max_{P^M, P^B, n \in \{0, \dots, M\}} \quad & S_n \left[1 - F\left(P^B + \frac{\kappa}{S_n}\right)\right] (P^M + P^B) \\ \text{s.t.} \quad & S_n \left[1 - F\left(P^B + \frac{\kappa}{S_n}\right)\right] (\pi - P^M) \geq nI, \end{aligned} \quad (\text{PC})$$

$$(P^M + P^B) \sum_{j=M-n+1}^M \theta_{j:M} \geq (P^M + P^B) \sum_{m \in T} \theta_m \quad \forall T : |T| \leq n. \quad (\text{IC})$$

As in the baseline model, (IC) holds whenever $P^M + P^B > 0$, and the platform binds (PC) through P^M , so that

$$\Pi_2^{\text{Interm}}(M) = \max_n \{S_n R(S_n) - nI\}.$$

Under information provision, a reporting cutoff θ' gives $S = M Q(\theta')$, where $Q(\theta') \equiv$

²⁶Amazon, for example, offers sellers financing from several providers, including Lendistry, Parafin, QuickBooks Capital, and Uncapped; see “[Amazon steps up financing to fuel seller growth](#)”.

²⁷The fixed cost κ can be interpreted as monetary costs (e.g., internet access) or cognitive costs, such as the time and effort required to learn the platform, extract information, or switch from another provider (Bouvard et al., 2022).

$\int_{\theta'}^1 \theta dG(\theta) + (1 - \alpha) \int_0^{\theta'} \theta dG(\theta)$, and the platform solves

$$\begin{aligned} \max_{P^M, P^B, \theta'} \quad & M Q(\theta') \left[1 - F\left(P^B + \frac{\kappa}{MQ(\theta')}\right) \right] (P^M + P^B) \\ \text{s.t.} \quad & \mathbb{E}[\theta \mid s'(\theta')] \left[1 - F\left(P^B + \frac{\kappa}{MQ(\theta')}\right) \right] (\pi - P^M) \geq I, \\ & \theta' = \min \left\{ \tilde{\theta} \in [0, 1] : \mathbb{E}[\theta \mid s'(\tilde{\theta})] \left[1 - F\left(P^B + \frac{\kappa}{MQ(\tilde{\theta})}\right) \right] (\pi - P^M) \geq I \right\}. \end{aligned} \tag{PC}$$

$$\tag{IC}$$

Binding (PC) through P^M again gives

$$\Pi_2^{\text{Info}}(M) = \max \left\{ M \max_{\theta'} [Q(\theta') R(MQ(\theta')) - I [1 - \alpha G(\theta')]], 0 \right\},$$

and (IC) holds at the solution: since $Q(\theta')$ decreases in θ' , any lower cutoff at which the bank still lends would raise fee income above this maximum. In the bank-lending benchmark, the bank still lends to all merchants or none, and since $R(S) < I/\underline{\theta}^{\text{PF}}$, Assumption 2 still rules out lending.

Since a larger pool attracts more buyers, profit per merchant now increases with M under information provision as well, so the argument for a single threshold in Section 5.1 no longer applies. The comparison between small and large pools, however, is unchanged.

Proposition 6. *With two-way network externalities, let $\alpha < 1$.*

(i) $\Pi_2^{\text{Interm}}(1) = 0$.

(ii) *As $M \rightarrow \infty$, $\Pi_2^{\text{Interm}}(M)/M \rightarrow \Pi^{\text{Mono}}(1)$ and $\Pi_2^{\text{Info}}(M)/M \rightarrow \Pi^{\text{Info}}(1)$, the limits in the baseline model.*

Hence there exists a finite $\bar{M}$ such that intermediation yields strictly higher profit than information provision for all $M \geq \bar{M}$, whereas information provision yields strictly higher profit for $M = 1$ whenever $\Pi_2^{\text{Info}}(1) > 0$.

Proof for Proposition 6. For part (i), with a single merchant, $S_1 = \mathbb{E}[\theta]$, and $\mathbb{E}[\theta] R(\mathbb{E}[\theta]) - I < \mathbb{E}[\theta] \frac{I}{\underline{\theta}^{\text{PF}}} - I < 0$ by Assumption 2, so the platform does not finance the merchant.

For part (ii), since $R(S) \leq I/\underline{\theta}^{\text{PF}}$, $\Pi_2^{\text{Interm}}(M) \leq \Pi^{\text{Interm}}(M) \leq \Pi^{\text{Mono}}(M)$. Conversely, with the baseline credit line $n^*(M)$, $S_{n^*(M)}$ grows in proportion to M , so

$$\frac{\Pi_2^{\text{Interm}}(M)}{M} \geq \frac{\Pi^{\text{Interm}}(M)}{M} - \frac{S_{n^*(M)}}{M} \left[\frac{I}{\underline{\theta}^{\text{PF}}} - R(S_{n^*(M)}) \right] \rightarrow \Pi^{\text{Mono}}(1)$$

by Proposition 1(ii). Under information provision, $Q(\theta') \geq (1 - \alpha)\mathbb{E}[\theta] > 0$ for every θ' , so $R(MQ(\theta')) \rightarrow I/\underline{\theta}^{\text{PF}}$ uniformly in θ' , and profit per merchant converges to its baseline value $\Pi^{\text{Info}}(1)$.

Since $\Pi^{\text{Info}}(1) < \Pi^{\text{Mono}}(1)$ for $\alpha < 1$, part (ii) implies that $\Pi_2^{\text{Interm}}(M) > \Pi_2^{\text{Info}}(M)$ for all M large enough, and part (i) gives the comparison for $M = 1$. $\square$

Appendix D Direct sales channels

In the baseline model, merchants can only sell through the platform. We now let a financed merchant complete the trades it obtains on the platform through a direct channel, such as its own web shop, at a cost $\delta > 0$ per transaction. This cost captures the platform services the merchant forgoes, such as payment processing, fulfilment and dispute resolution, as well as any penalties the platform imposes on circumvention.²⁸ On the direct channel, the merchant pays no fee to the platform, and the platform cannot deduct any payment from the merchant's sales proceeds. We assume that buyers pay P^B for each merchant they are matched with on the platform, so that the direct channel affects only the merchant side. A merchant uses the direct channel only if it strictly gains from doing so.

Who finances merchants now depends on how repayment is collected. A platform collects the repayment of its loans by deducting it from merchants' payouts, which gives it an enforcement advantage over banks (Li & Pegoraro, 2022) but works only for sales made on the platform. For bank loans under information provision, we consider two cases. In case (a), the bank's claim is enforceable on all of the merchant's sales proceeds, on and off the platform, as with a standard loan. In case (b), the bank also collects repayment by deduction from platform payouts, as with advances repaid from a fixed share of platform sales.²⁹ In both cases, the bank can pay the platform a referral fee $\tau \geq 0$ for each loan it makes, as when lenders pay platforms for referrals.

A merchant bypasses the platform if the payments that the platform would collect from it on each transaction exceed δ . Under intermediation, these payments consist of the fee P^M and the repayment $D/[1 - F(P^B)]$ of the platform's loan per transaction, so no bypass requires

$$P^M + \frac{D}{1 - F(P^B)} \leq \delta. \quad (\text{NB})$$

Together with limited liability, (NB) limits what the platform collects from a merchant per transaction to $\min\{\pi, \delta\}$, in addition to the buyer-side fee. Intermediation is therefore equivalent to the baseline model with π replaced by $\min\{\pi, \delta\}$ in the platform's revenue, and its financing cutoff becomes

$$\underline{\theta}_\delta^{\text{PF}} \equiv \frac{I}{\max_{P^B} [1 - F(P^B)] (\min\{\pi, \delta\} + P^B)} \geq \underline{\theta}^{\text{PF}}.$$

Under information provision in case (a), only the fee passes through the platform, so no bypass requires $P^M \leq \delta$, and the referral fee transfers the bank's surplus to the platform.

²⁸The direct channel is a form of marketplace leakage, in which sellers move transactions off the platform to avoid its fees (Hagiu & Wright, 2024).

²⁹For DoorDash Capital, "Payments are based on a fixed percentage of your DoorDash gross sales"; see "DoorDash Capital".

Information provision is then unaffected by the direct channel. In case (b), the bank's repayment also passes through the platform, (NB) applies with D the bank's repayment, and information provision, too, is equivalent to the baseline model with π replaced by $\min\{\pi, \delta\}$.

Proposition 7 (Direct sales channel). *Let $\alpha < 1$, and let $M^* = \infty$ if intermediation is preferred for no M .*

- (i) *In case (a), $\Pi^{\text{Info}}(M)$ does not depend on δ , while $\Pi^{\text{Interm}}(M)$ is non-decreasing in δ for every M . Hence M^* is non-increasing in δ .*
- (ii) *In case (b), the model is the baseline model with π replaced by $\min\{\pi, \delta\}$. A lower δ weakly raises the cutoff $\underline{\theta}_\delta^{\text{PF}}$, and weakly lowers $\Pi^{\text{Interm}}(M)$, $\Pi^{\text{Info}}(M)$ and the expected number of financed merchants under both arrangements. For any $\delta_1 < \delta_2$ with $1 > \underline{\theta}_{\delta_1}^{\text{PF}} > \underline{\theta}_{\delta_2}^{\text{PF}}$, M^* is smaller under δ_1 for all α sufficiently close to one.*

Proposition 7 shows that a cheap direct channel undermines lending only to the extent that the lender relies on platform payouts to collect. In case (a), the bank's claim survives the direct channel, whereas the platform's claim does not. The direct channel therefore shifts the partnership towards information provision: the bank finances the merchants, and the platform earns referral fees from the bank in addition to the fees that merchants still pay on the platform. Intermediation may then never be chosen. In case (b), no lender can collect on off-platform sales. The direct channel then caps the surplus that can be pledged per transaction at δ , which raises the financing cutoff, reduces credit under both arrangements and, by Corollary 3, favours intermediation when information provision is credible.

Proof for Proposition 7. Under intermediation, a successful merchant pays the platform $[1 - F(P^B)]P^M + D$, subject to limited liability, $D \leq [1 - F(P^B)](\pi - P^M)$, and (NB). Hence the platform collects at most $[1 - F(P^B)] \min\{\pi, \delta\}$ from the merchant, and it can collect exactly this amount, with D pledged to the bank as in the baseline model. The problem is then (7) with π replaced by $\min\{\pi, \delta\}$. Its value,

$$\max_n \left\{ \max_{P^B} [1 - F(P^B)](\min\{\pi, \delta\} + P^B) \sum_{j=M-n+1}^M \mathbb{E}[\theta_{j:M}] - nI \right\},$$

is non-decreasing in δ .

Under information provision in case (a), let the bank collect $D = [1 - F(P^B)](\pi - P^M)$ from all proceeds and pay the referral fee $\tau = \mathbb{E}[\theta \mid s'(\theta')] D - I$ per loan, which leaves it with zero expected profit. The platform's expected payoff per merchant is then

$$[1 - F(P^B)](P^M + P^B) \int_0^1 \theta [\alpha \mathbf{1}\{\theta \geq \theta'\} + 1 - \alpha] dG(\theta) + \tau [1 - \alpha G(\theta')]$$

$$= \int_0^1 \left(\theta [1 - F(P^B)](\pi + P^B) - I \right) [\alpha \mathbf{1}\{\theta \geq \theta'\} + 1 - \alpha] dG(\theta),$$

where $\alpha \mathbf{1}\{\theta \geq \theta'\} + 1 - \alpha$ is the probability that type θ is recommended, as in the proof of Proposition 2. This does not depend on P^M , so the platform can set

$$P^M = \min \left\{ \delta, \pi - \frac{I}{[1 - F(P^B)] \mathbb{E}[\theta \mid s'(\theta')]} \right\},$$

which satisfies $P^M \leq \delta$ and gives $\tau \geq 0$. Hence $\Pi^{\text{Info}}(M)$ equals its baseline value and does not depend on δ . Since $\Pi^{\text{Interm}}(M)$ is non-decreasing in δ , the set of M with $\Pi^{\text{Interm}}(M) \geq \Pi^{\text{Info}}(M)$ weakly grows with δ , and its smallest element M^* weakly falls. This proves part (i).

In case (b), all payments by the merchant pass through the platform, so the argument for intermediation applies to information provision as well: the platform and the bank together collect at most $[1 - F(P^B)] \min\{\pi, \delta\}$ per successful merchant, and the referral fee only divides this amount between them. The model is therefore the baseline model with π replaced by $\min\{\pi, \delta\}$. A lower δ lowers $\max_{P^B} [1 - F(P^B)](\min\{\pi, \delta\} + P^B)$ weakly, which raises the cutoff weakly. For every reporting cutoff and credit line, profit is non-decreasing in this maximum, so both profits weakly fall. Under intermediation, the number of financed merchants $n^*(M)$ counts the ranks with $\mathbb{E}[\theta_{j:M}] \geq \underline{\theta}_\delta^{\text{PF}}$, which weakly falls when the cutoff rises. Under information provision, the expected number of financed merchants is $M[1 - \alpha G(\underline{\theta}_\delta^{\text{PF}})]$ if $\Pi^{\text{Info}}(M) > 0$ and zero otherwise, which also weakly falls, since $\Pi^{\text{Info}}(M)$ weakly falls. The last claim follows from Corollary 3, applied to the cutoffs $\underline{\theta}_{\delta_1}^{\text{PF}} > \underline{\theta}_{\delta_2}^{\text{PF}}$, which lie in $(\mathbb{E}[\theta], 1)$ since $\underline{\theta}_\delta^{\text{PF}} \geq \underline{\theta}^{\text{PF}} > \mathbb{E}[\theta]$. $\square$

Appendix E Commitment in pricing

Under both arrangements, the platform uses the merchant-side fee P^M to resolve the conflict of interest with the bank, which stems from network externalities that only the platform captures. This requires that the platform commits to its fees, as we assume in Section 2. In this section, we first show that both arrangements unravel without such commitment and then discuss the institutional practices that make it plausible.

Commitment to fee structure is crucial Without commitment to its fees, the platform can hold up the bank.

Under intermediation, if the platform could change P^M after the bank agrees to a credit line at $t = 1a$, it would raise its fee to $P^M = \pi$, leaving no funds to repay the credit line. Anticipating this, the bank would refuse to provide the credit line in the first place.

A similar logic applies to the information provision arrangement. The platform could wait until the bank has financed merchants based on its reports, at $t = 1c$, before raising P^M to π . This would again benefit the platform ex post but result in merchant defaults. Foreseeing this, the bank would expect any lending to be unprofitable, whatever the reports, and would decline to finance any merchant. In both arrangements, the platform can then no longer use its fees to resolve the conflict of interest with the bank, and lending breaks down.

Institutional practices The assumption that the platform commits to price setting is supported by applicable legal and regulatory constraints, the platforms’ reputational stakes when interacting with creditors and users, and non-trivial adjustment costs.

First, platforms cannot arbitrarily rewrite user fees mid-contract. European regulations, for instance, explicitly prohibit this by requiring platforms to provide significant advance notice, ranging from 15 days to two months, before any changes to their terms. These rules also grant business partners the right to terminate their contract in response.³⁰ Mandatory notice and the right to exit therefore fix the fee structure over a financing cycle and protect the cash flows that borrowers need for repayment. Competition laws further require platforms to keep their pricing consistent and transparent.³¹

Commitment can also be self-enforcing, sustained by the long-run value of the platform’s relationships. A platform that deviates on posted fees makes a one-off gain but jeopardises the present value of future profits from interactions with both creditors and users. Our assumption can therefore be microfounded by a dynamic model in which the platform keeps its posted fees whenever the discounted continuation value exceeds the short-run gain from deviating (Baker et al., 2002; Macchiavello, 2022; MacLeod, 2007). This continuation value is large for platforms whose business models rely on network effects and trust.

Finally, adjusting fees is costly. These costs extend beyond physical “menu costs” (Nakamura & Steinsson, 2008), and the largest are managerial and customer-facing (Zbaracki et al., 2004). This helps explain why platforms post fee schedules that they update infrequently, consistent with our assumption that fees are fixed over the financing horizon.

³⁰See “Regulation (EU) 2019/1150” and “Changes in conditions of the framework contract (Article 54, PSD2)”.

³¹See, for example, “Apple changes App Store rules in EU to comply with antitrust order” for a case on Apple found in breach of the EU Digital Markets Act.

References

- Agur, I., Ari, A., & Dell’Ariccia, G. (2025). Bank competition and household privacy in a digital payment monopoly. *Journal of Financial Economics*, 166.
- Ahnert, T., Hoffmann, P., & Monnet, C. (2025). Payments and privacy in the digital economy. *Journal of Financial Economics*, 169.
- Armstrong, M. (2006). Competition in two-sided markets. *The RAND Journal of Economics*, 37(3), 668–691.
- Bagnoli, M. & Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory*, 26(2), 445–469.
- Baker, G., Gibbons, R., & Murphy, K. J. (2002). Relational contracts and the theory of the firm. *The Quarterly Journal of Economics*, 117(1), 39–84.
- Barakova, I., Ehrentraud, J., & Leposke, L. (2024). A two-sided affair: banks and tech firms in banking. *Bank for International Settlements, Financial Stability Institute*.
- Berg, T., Burg, V., Keil, J., & Puri, M. (2025). The economics of “buy now, pay later”: A merchant’s perspective. *Journal of Financial Economics*, 171, 104093.
- Bergemann, D., Bonatti, A., & Gan, T. (2022). The economics of social data. *The RAND Journal of Economics*, 53(2), 263–296.
- Biais, B. & Gollier, C. (1997). Trade credit and credit rationing. *The Review of Financial Studies*, 10(4), 903–937.
- Bian, W., Cong, L. W., & Ji, Y. (2024). *The Rise of E-Wallet Super-Apps and Buy-Now-Pay-Later*. Working Paper 33178, National Bureau of Economic Research.
- BIS (2019). Big tech in finance: opportunities and risks. Chapter 3, Annual Economic Report, June.
- Boualam, Y. & Yoo, P. (2022). Fintech disruption, banks, and credit (dis-)intermediation: When do foes become friends? *SSRN Electronic Journal*.
- Bouvard, M., Casamatta, C., & Xiong, R. (2022). Lending and monitoring: Big tech vs banks. Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting 2023.
- Buchak, G., Matvos, G., Piskorski, T., & Seru, A. (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics*, 130(3), 453–483.
- Burkart, M. & Ellingsen, T. (2004). In-kind finance: A theory of trade credit. *American Economic Review*, 94(3), 569–590.
- Chen, T., Huang, Y., Lin, C., & Sheng, Z. (2022). Finance and firm volatility: Evidence from small business lending in China. *Management Science*, 68(3), 2226–2249.
- Chu, Y. & Wei, J. (2024). Fintech lending and credit market competition. *Journal of Financial & Quantitative Analysis*, 59(5), 2199–2225.
- Cornelli, G., Frost, J., Gambacorta, L., Rau, P. R., Wardrop, R., & Ziegler, T. (2023). Fintech and big tech credit: Drivers of the growth of digital lending. *Journal of*

- Banking & Finance*, 148, 106742.
- Crawford, V. P. & Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6), 1431–1451.
- David, H. A. & Nagaraja, H. N. (2003). *Order Statistics*. Hoboken, NJ: Wiley, 3rd edition.
- De Fiore, F., Cornelli, G., Gambacorta, L., & Manea, C. (2024). Big tech, financial intermediation and the macroeconomy. BIS Working Paper.
- Demirguc-Kunt, A., Klapper, L., Singer, D., Ansar, S., & Hess, J. (2018). *The Global Findex Database 2017: Measuring financial inclusion and the fintech revolution*. World Bank Publications.
- Dessein, W. (2002). Authority and communication in organizations. *The Review of Economic Studies*, 69(4), 811–838.
- Djankov, S., La Porta, R., Lopez-de Silanes, F., & Shleifer, A. (2003). Courts. *The Quarterly Journal of Economics*, 118(2), 453–517.
- Doerr, S., Frost, J., Gambacorta, L., & Shreeti, V. (2025). Big techs in finance. In *The Oxford Handbook of Banking: 4th Edition*. Oxford University Press.
- Frost, J., Gambacorta, L., Huang, Y., Shin, H. S., & Zbinden, P. (2020). BigTech and the changing structure of financial intermediation. *Economic Policy*, 34(100), 761–799.
- Gambacorta, L., Huang, Y., Li, Z., Qiu, H., & Chen, S. (2022). Data versus collateral. *Review of Finance*, 27(2), 369–398.
- Gertler, P. J., Green, B., & Wolfram, C. (2021). Digital collateral. *NBER Working Paper No. W28724*.
- Glazer, J. & Rubinstein, A. (2004). On optimal rules of persuasion. *Econometrica*, 72(6), 1715–1736.
- Green, J. R. & Laffont, J.-J. (1986). Partially verifiable information and mechanism design. *The Review of Economic Studies*, 53(3), 447–456.
- Hagiu, A. & Wright, J. (2015). Multi-sided platforms. *International Journal of Industrial Organization*, 43, 162–174.
- Hagiu, A. & Wright, J. (2023). Data-enabled learning, network effects, and competitive advantage. *The RAND Journal of Economics*, 54(4), 638–667.
- Hagiu, A. & Wright, J. (2024). Marketplace leakage. *Management Science*, 70(3), 1529–1553.
- Hau, H., Huang, Y., Lin, C., Shan, H., Sheng, Z., & Wei, L. (2024). Fintech credit and entrepreneurial growth. *The Journal of Finance*, 79(5), 3309–3359.
- Hau, H., Huang, Y., Shan, H., & Sheng, Z. (2019). How fintech enters China’s credit market. *AEA Papers and Proceedings*, 109, 60–64.
- He, Z., Huang, J., & Zhou, J. (2023). Open banking: Credit market competition when borrowers own the data. *Journal of Financial Economics*, 147(2), 449–474.
- Innes, R. D. (1990). Limited liability and incentive contracting with ex-ante action

- choices. *Journal of Economic Theory*, 52(1), 45–67.
- Kamenica, E. & Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6), 2590–2615.
- Li, J. & Pegoraro, S. (2022). Borrowing from a BigTech platform. *SSRN Electronic Journal*.
- Liberti, J. M. & Petersen, M. A. (2019). Information: Hard and soft. *The Review of Corporate Finance Studies*, 8(1), 1–41.
- Lipnowski, E., Ravid, D., & Shishkin, D. (2022). Persuasion via weak institutions. *Journal of Political Economy*, 130(10), 2705–2730.
- Liu, L., Lu, G., & Xiong, W. (2022). The big tech lending model. *NBER Working Paper Series 30160*.
- Macchiavello, R. (2022). Relational contracts and development. *Annual Review of Economics*, 14, 337–362.
- MacLeod, W. B. (2007). Reputations, relationships, and contract enforcement. *Journal of Economic Literature*, 45(3), 595–628.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, 12(2), 380–391.
- Min, D. (2021). Bayesian persuasion under partial commitment. *Economic Theory*, 72(3), 743–764.
- Mookherjee, D. (2006). Decentralization, hierarchies, and incentives: A mechanism design perspective. *Journal of Economic Literature*, 44(2), 367–390.
- Nakamura, E. & Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4), 1415–1464.
- Parlour, C. A., Rajan, U., & Zhu, H. (2022). When fintech competes for payment flows. *The Review of Financial Studies*, 35(11), 4985–5024.
- Rochet, J.-C. & Tirole, J. (2006). Two-sided markets: a progress report. *The RAND Journal of Economics*, 37(3), 645–667.
- Rothschild, M. & Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, 2(3), 225–243.
- Silva, F. (2024). Information transmission in persuasion models with imperfect verification. *Theoretical Economics*, 19(3), 1001–1026.
- Vives, X. & Ye, Z. (2025). Fintech entry, lending market competition, and welfare. *Journal of Financial Economics*, 168.
- Wang, Y. & Charoenwong, B. (2025). The BigTech-bank lending ecosystem. INSEAD Working Paper 2025/11/FIN.
- Zbaracki, M. J., Ritson, M., Levy, D., Dutta, S., & Bergen, M. (2004). Managerial and customer costs of price adjustment: Direct evidence from industrial markets. *The Review of Economics and Statistics*, 86(2), 514–533.