BANK OF CANADA
BANQUE DU CANADA

Staff working paper / Document de travail du personnel—2026-41

Last updated: November 10, 2026

Instant Payments, Fraud, and Optimal Consumer Protection

Yuteng Cheng
Payments and Regulatory Oversight
Bank of Canada
ycheng@bank-banque-canada.ca

Charles Kahn
University of Illinois
cmkahn@illinois.edu

Francisco Rivadeneyra
Payments and Regulatory Oversight
Bank of Canada
frivadeneyra@bank-banque-canada.ca

Bank of Canada staff research is produced independently from the Bank's Governing Council and may support or challenge prevailing views. The views expressed in this paper are solely those of the authors and may differ from official Bank of Canada positions. No responsibility for them should be attributed to the Bank.

DOI: https://doi.org/10.34989/swp-2026-41 | ISSN 1701-9397

© 2026 Bank of Canada

Instant Payments, Fraud, and Optimal Consumer Protection

Yuteng Cheng
Bank of Canada

Charles Kahn
University of Illinois

Francisco Rivadeneyra
Bank of Canada

Abstract

Real-time payment systems increase fraud risks by compressing the detection window. We develop a framework to study two institutional levers: centralized versus decentralized fraud control, and the allocation of fraud losses through liability and reimbursement rules. Centralized governance does not by itself implement the planner's allocation, because user fees can inefficiently exclude small payments. In a decentralized regime, a simple liability split generally does not implement the first best, but the optimal split can be characterized by the relative adjusted curvature of sender- and receiver-side prevention technologies. Liability is generally shared, but becomes sender-heavy when receiver-side monitoring is sufficiently less flexible. With access and cap mandates, centralized governance dominates; without them, decentralization can yield higher welfare.

1 Introduction

While real-time retail payment systems have shifted the frontier of retail transfers, they have also amplified security risks. Systems such as FedNow and RTP in the United States, Faster Payments in the United Kingdom, Pix in Brazil, SEPA Instant in the euro area, and Canada's forthcoming Real-Time Rail (RTR) settle payments in seconds on a 24/7/365 basis. For example, in the EU, banks must process euro instant credit transfers within ten seconds and make funds available immediately.[1] Yet faster settlement also compresses the time window for intervention: once funds move, they can quickly leave the effective control of payment service providers (PSPs), making fraud monitoring and recovery more difficult. As a result, PSPs increasingly require upgraded tools and more complex, cooperative arrangements to protect customers and manage losses.

Authorized push payment (APP) scams—where a payer is manipulated into initiating a transfer to a fraudster—have become a leading manifestation of these risks on instant-payment rails. In the United Kingdom, UK Finance reports £459.7 million of APP fraud losses across 232,429 cases in 2023, roughly 0.017% of 2023 GDP, while the Payment Systems Regulator reports about £341 million of APP scam losses on Faster Payments alone in 2023, roughly 0.013% of GDP.[2] Similar concerns have emerged elsewhere. In Brazil, public reporting has pointed to large scam- and fraud-related losses on Pix, with estimates around R$ 4.9 billion in 2024, or about 0.04% of 2024 GDP.[3] These figures differ in scope and reporting conventions across jurisdictions, but they point to the same policy problem: once instant payments operate at scale, fraud losses can become economically significant even when they remain small relative to GDP and total transaction value.

Among current policy responses, two institutional choices stand out. Canada plans to launch one of the first instant-payment systems with mandatory centralized fraud services from day one.[4] The United Kingdom, meanwhile, has implemented a mandatory authorized push-payment reimbursement regime covering Faster Payments and CHAPS, under which

[1] European Central Bank, "What are instant payments?"; see also the EU Instant Payments Regulation.

[2] Reporting scopes and conventions differ across sources. UK Finance reports economy-wide APP aggregates based on its member reporting framework, while the PSR figures are compiled for Faster Payments and may differ in institutional coverage, instrument scope, and case-counting. The GDP ratios use 2023 U.K. GDP at current prices of approximately £2.687 trillion. We therefore treat these figures as indicative orders of magnitude rather than harmonized statistics.

[3] The Brazil ratio uses 2024 GDP of approximately R$ 11.7 trillion. As with the U.K. figures, the Pix fraud estimate and the GDP denominator are used to convey scale rather than to construct a harmonized cross-country fraud statistic.

[4] The RTR's centralized fraud framework comprises Central Fraud Analytics, a Central Risk List, Central Fraud Reporting, and Confirmation of Payee, and RTR participants will be required to use these services. Interac is an RTR delivery partner and is building the fraud service as well as the RTR exchange. See Payments Canada (2025a) and Payments Canada (2025b).

sending and receiving payment service providers split reimbursement costs 50/50, reimbursement is generally capped at £85,000 per claim, and most claims must be resolved within five business days.[5] Elsewhere, high nominal transaction limits are paired with granular bank-set caps and velocity controls, while name-check and confirmation-of-payee tools, as well as network-level analytics, are being introduced at different speeds.[6]

We develop a theoretical framework to analyze these institutional choices. The model focuses on two policy questions that map directly to current design debates. First, what are the welfare tradeoffs between centralized and decentralized fraud control on an instant rail? Second, in a decentralized regime with loss sharing, how should liability be allocated between sending and receiving banks once prevention incentives and selection on payment size are taken into account? The framework centers on instruments that mirror practice: payment-size caps, flat-fee access pricing, sender-side controls, receiver-side monitoring, and reimbursement or liability rules.

The model is deliberately parsimonious. An instant rail generates value from speed but also exposes payments to fraud: payers with heterogeneous transfer sizes choose between the rail and a slower outside option, and because the convenience of speed is concave in payment size while expected fraud losses scale linearly, efficient use of the rail follows a threshold rule. Fraud prevention is a team problem: sending banks screen payments at initiation, receiving banks monitor accounts and incoming funds, both controls are privately chosen, and a successful fraud must pass both. Policy operates under one of two governance structures: Under centralized governance, a scheme operator sets the access fee and eligibility cap and coordinates prevention on both sides; under decentralized governance, competitive sending banks post the fee and each bank privately chooses its own controls given the liability rule. Because the fee is flat, pricing determines not only how many payments use the rail but which ones, and this selection margin interacts with both governance structures.

The analysis yields three main results. First, centralization by itself does not deliver the planner's allocation. The difficulty is the pricing instrument: a flat fee charges every payment the average cost of the rail, so small payments—whose own fraud exposure is minimal but whose value of speed is high—can be priced off the rail, and the admitted region becomes a band rather than the planner's lower interval. Moreover, because different operator objectives select different fee–cap combinations, the composition of routed payments and hence fraud-control incentives vary with the operator's mandate. Both distortions disappear if regulation fixes the boundaries of access directly: with a zero minimum transaction size and an appropriate cap, a centralized operator implements the first best regardless of its

[5]See the Payment Systems Regulator's consolidated policy statement on the APP scam reimbursement requirement and its accompanying consumer guidance. (Payment Systems Regulator, 2025a,2)

[6]FedNow program and product updates (June–July 2025); EY (Jan. 2025), "EU Instant Payments Regulation: Five key hurdles for banks to clear."

objective.

Second, in a decentralized regime, a simple liability split between sending and receiving banks generally does not implement the planner's solution. Even if the admitted region is fixed efficiently, each side internalizes only its share of the marginal benefit from tighter controls, so prevention incentives remain split—an instance of the impossibility of efficient budget-balanced sharing under moral hazard in teams (Holmström, 1982). We note two implementation schemes that restore the planner's allocation, then study the optimal liability split as a second-best instrument. With the admitted region fixed, the optimal split depends on the relative adjusted curvature of the two prevention technologies: liability loads toward the side whose prevention margin is more flexible, and it becomes sender-heavy when receiver-side monitoring is sufficiently less flexible than sender-side screening. When access is priced, a second force appears. Sending banks price access while receiving banks do not, so liability assigned to the sender passes through the competitive fee and can exclude low-size payments, while liability assigned to the receiver does not. One implication is that symmetric prevention technologies do not imply a symmetric split: even when the two sides are identical, the optimal allocation tilts toward the receiving bank whenever sender liability raises the access price.

Third, centralized and decentralized regimes are not unambiguously ranked. With access and cap mandates, centralized governance dominates: it coordinates prevention on both sides while the mandates preserve the planner's admitted region, whereas a decentralized regime governed by a single loss share cannot align both prevention margins at once. Without those mandates, the same pass-through asymmetry that shapes the optimal liability split can reverse the ranking. A centralized operator must recover all system costs—including receiver-side monitoring costs—through the access fee, whereas the decentralized fee reflects only the costs internalized by sending banks. The centralized fee can therefore be much higher, especially when the sender's loss share is small. In our numerical illustration, this fee gap excludes enough low-value payments under centralization that decentralization yields higher welfare. Finally, we discuss two natural extensions, payer-side moral hazard and endogenous fraudster entry, and explain how they interact with the baseline access, eligibility, and prevention-incentive mechanisms.

These results speak directly to the two institutional choices that motivate the paper. For the United Kingdom's reimbursement regime, the framework identifies the conditions under which an equal split between sending and receiving banks performs well: the two sides' prevention technologies are similarly flexible at the margin, and sender-side liability does not materially pass through into the price of access. Where these conditions do not hold, the framework indicates how the split should adjust—toward the more flexible prevention margin, and toward the receiving bank when access pricing is sensitive to sender costs. For

jurisdictions adopting centralized fraud services, as Canada has mandated, the framework highlights a complementarity: centralized fraud control is most effective when paired with access and eligibility rules that keep low-value payments on the rail, since these are the payments for which immediacy is most valuable relative to fraud exposure. The general message is that liability rules, access pricing, and eligibility caps are complements in the design of instant rails, and that evaluating any one of them requires attention to the others.

Related Literature This paper bridges three strands of literature: the economics of information security, the design of payment systems, and recent policy assessments of fraud in instant payments.

First, we build on the economics of information security to model the misalignment of incentives. Anderson and Moore (2006) and Moore (2010) establish that security failures often arise not from technical inability, but because the party capable of prevention is not the one bearing the loss. We formalize this "security externality" in the context of a two-sided payment network. While standard models often assume a single defender, we show that decentralized payments involve split incentives between sending and receiving institutions, creating a specific type of under-investment that varies with the "stiffness" of the monitoring technology.

Second, the decentralized regime is an instance of moral hazard in teams. Holmström (1982) shows that when unobservable individual efforts produce a joint outcome, no budget-balanced sharing rule implements the efficient effort profile: each agent internalizes only their share of the marginal return, and efficiency would require every agent to face it in full. Our split-liability result is this impossibility specialized to the two transaction-side roles, with the liability share λ playing the role of the sharing rule over a joint loss. Our double-liability implementation (Scheme A) is the corresponding group-penalty solution: charging each bank the full expected loss breaks budget balance, and the scheme operator serves as the residual budget-breaker, absorbing the excess assessments through a compensation pool or rebates untied to a bank's own losses. Relative to that literature, our contribution is twofold: we characterize the optimal budget-balanced split as a second-best instrument, showing that it depends on the relative adjusted curvature of the two prevention technologies, and we embed the sharing rule in a market where the sending side also prices access, so that liability allocation affects the composition of routed payments and not only prevention incentives.

Third, our framework relates to the mechanism design literature on payment system tiering and settlement. Kahn and Roberds (2009) demonstrate that in tiered networks, restricting access can induce better monitoring by direct participants. We extend this logic to fraud control. While their focus is on credit and settlement risk, we apply similar incentive-compatibility constraints to "dissipation risk." In our model, the shape of admission—using

caps and fees to screen out inefficient payments—serves as a primary tool for managing this risk, parallel to how credit limits manage default risk in settlement systems.

Finally, the most closely related paper to our work is the recent contribution by Balasubramaniam, Mosk, and Uettwiller (2025). They examine liability rules in the context of APP fraud through a dynamic lens, focusing on how fraudsters strategically adapt to detection technologies, which causes model performance to decay over time. They find that in competitive markets, placing full liability on payment service providers can induce inefficient overinvestment, as individual banks fail to internalize how their screening efforts accelerate market-wide fraudster adaptation. To resolve the free-rider problem between sending and receiving banks, they propose a first-best solution that assigns full reimbursement liability to both sides, complemented by a regulatory fee.

Our approach complements theirs by systematically comparing centralized and decentralized governance structures. Their main focus is a dynamic environment in which fraudsters adapt and detection models decay. We instead embed fraud prevention in a market where access is priced, so that the operator's pricing and eligibility decisions—the shape of the admitted set—become regulatory instruments that interact with liability rules. We also characterize the optimal liability split under decentralized governance.

Layout The rest of the paper is organized as follows. Section 2 introduces the baseline model of fraud on an instant payment rail, defining the technology of prevention and the instruments available to operators. Section 3 characterizes the social planner's problem and derives the first-best allocation. Section 4 analyzes the centralized regime, solving for the optimal contract under alternative operator objectives and identifying the conditions under which centralization implements the first best. Section 5 studies the decentralized regime, examining equilibrium under split liability and comparing it to the planner's benchmark. Section 6 discusses two natural extensions, payer-side moral hazard and endogenous fraudster entry, and explains how they interact with the baseline mechanisms. Section 7 concludes.

2 Baseline model

This section introduces a parsimonious model of retail instant payments in which speed generates value but also amplifies exposure to authorized push payment (APP) fraud. The model isolates a single organizing tension: the convenience gain from immediacy rises with payment size but exhibits diminishing marginal returns, whereas expected fraud harm scales linearly with size. This wedge implies that, absent sufficiently strong controls, it is inefficient to route arbitrarily large transfers through an instant rail. What makes this a model of instant payments is that once funds move in real time, there is much less opportunity to

stop, reverse, or recover a fraudulent transfer, so fraud control must rely more heavily on ex ante screening and monitoring.

The rail is characterized by two publicly observable policy instruments that capture common design choices in practice: a per-payment fee and a size-based eligibility cap. These instruments interact with two control margins: *front-end* controls at the sending side and *back-end* monitoring at the receiving side. The centralized regime interprets these choices as scheme-level policy (the operator selects pricing, eligibility, and effort targets), while the decentralized regime interprets them as the outcome of competitive pricing and privately chosen bank efforts under a given loss-sharing rule.

We keep the setup deliberately simple—a continuum of payers with heterogeneous transfer sizes and an outside option—so that later sections can cleanly compare the planner's benchmark with market outcomes under alternative reimbursement and liability allocations. The rest of this section states the primitives, timing, and regularity conditions used throughout.

Players. There is a unit continuum of payers $j \in [0,1]$; sending banks that originate instant payments; receiving banks that hold recipient accounts; and a scheme operator (regulator/utility). Fraudsters generate mule accounts and attempt to dissipate funds; recipients are passive account holders. Institutional governance varies by regime: in the *centralized* case the operator sets eligibility and pricing rules and chooses the control targets (δ, θ) to be implemented by participating banks; in the *decentralized* case competitive sending banks post prices and privately choose sender-side controls, while receiving banks privately choose monitoring given the loss-sharing rule.[7]

Payment sizes and convenience of immediacy. Each payer draws a required transfer size $m_j \sim \mathrm{U}[0, M]$, independently across j, and enjoys outside-option utility $\bar{u}$.[8] Routing a

[7]The distinction between sending and receiving banks is role-based rather than institutional: within a given transaction, sender-side screening and receiver-side monitoring are different control margins even if the same legal entity performs both functions in aggregate. The model is therefore best interpreted as capturing the incentives attached to the two transaction-side roles; allowing a subset of payments to be "on-us" would partially internalize the externality but would not eliminate the basic sending-bank–receiving-bank incentive problem on interbank transfers.

[8]The lower support point is chosen for notational convenience and is without loss of generality: all results hold verbatim with $m \sim U[m_{\min}, M]$ for any $m_{\min} \in [0, M)$, with $m_{\min}$ replacing 0 throughout. What matters economically is not the support but the value of its smallest payments: the convenience of immediacy has a fixed component ($\phi(0) > 0$), reflecting benefits such as meeting a deadline or obtaining immediate confirmation, so low-value payments are not low-value users of speed. Section 3 states the corresponding condition formally.

transfer of size m through instant payments yields a size-dependent convenience gain

$$\Delta(m) \equiv \phi(m), \qquad \phi : [0, M] \to \mathbb{R}_+,$$

where ϕ is continuous and increasing. We interpret $\phi(m)$ as the payer's reduced-form value of immediacy relative to a slower outside option, capturing benefits such as avoiding late-payment penalties or overdraft fees, obtaining immediate confirmation, and improving cash-flow timing; it is conceptually distinct from fraud risk and does not vary with whether the payment is ultimately fraudulent or successful, which enters separately through the expected-loss term $\delta\theta m$.[9] We further assume that ϕ is strictly concave: once the most urgent timing needs are met, the incremental value of faster settlement grows more slowly.

Fraud technology and control margins. If a payment is routed through instant payments, an APP-fraud event succeeds only if it passes sending-bank controls and the funds are then successfully withdrawn or dissipated on the receiving side. We summarize sending-bank controls by an *initiation probability* $\delta \in (0, 1]$: conditional on a fraud attempt, δ is the probability that the payment is initiated and passes front-end screening. Lower δ means tighter sending-bank controls. We summarize receiving-bank monitoring by a *withdrawal probability conditional on initiation* $\theta \in (0, 1]$: conditional on fraudulent initiation, θ is the probability that funds are successfully withdrawn or dissipated. Lower θ means tighter receiving-bank monitoring.

Per routed payment, sending and receiving banks incur resource costs $e(\delta)$ and $\psi(\theta)$, respectively. We assume that $e, \psi : (0, 1] \to \mathbb{R}_+$ are twice continuously differentiable and satisfy

$$e(1) \geq 0, \qquad e'(\delta) < 0, \qquad e''(\delta) > 0, \qquad \lim_{\delta \downarrow 0} e(\delta) = +\infty,$$

and

$$\psi(1) \geq 0, \qquad \psi'(\theta) < 0, \qquad \psi''(\theta) > 0, \qquad \lim_{\theta \downarrow 0} \psi(\theta) = +\infty.$$

Thus, tighter controls are increasingly costly, and perfect prevention requires unbounded resources. The corners $\delta = 1$ and $\theta = 1$ correspond to no sender-side screening and no receiver-side monitoring beyond the baseline. Together with the boundedness of ϕ, these conditions confine relevant optimal controls to a compact subset of $(0, 1]^2$. Conditional on routing a payment of size m, an APP-fraud event succeeds with probability $\delta\theta$, generating

[9] A payee may also value immediacy, for example because faster receipt improves liquidity or reduces delivery risk. One can therefore interpret $\phi(m)$ as a reduced-form net benefit to the payer after any bilateral adjustment with the payee. Modeling the joint choice of rail by payer and payee would require a bargaining or contracting game, potentially with private information and strategic interaction with fraudsters. We abstract from those margins here because the paper focuses on the policy design problem that arises once a payment is routed onto the instant rail: pricing, eligibility, fraud controls, and liability allocation.

expected loss $\delta\theta m$. Under centralization, the operator chooses (δ, θ); under decentralization, sending and receiving banks choose δ and θ, respectively, given the prevailing loss-sharing rule.

Pricing and eligibility instruments. The rail is characterized by a per-payment flat fee $f \geq 0$ and a size eligibility cap $m_p \in [0, M]$; only payments with $m \leq m_p$ are admitted to instant payments. In the centralized regime, (f, m_p) are set by the operator; in the decentralized regime, f is posted by competitive sending banks while m_p is a scheme rule. Receiving banks do not charge fees on incoming instant payments; we assume they fund monitoring costs and any liability share out of exogenous account-relationship revenue, so their participation is not at issue and their problem reduces to cost minimization.[10]

Loss sharing and reimbursement (decentralized). In the *decentralized* regime, upon a successful APP event on a routed payment of size m, expected loss $\delta\,\theta\,m$ is allocated according to

$$(\lambda_s, \lambda_r, \lambda_p) \in [0,1]^3, \qquad \lambda_s + \lambda_r + \lambda_p = 1,$$

borne by the sending bank, receiving bank, and payer, respectively. The case $\lambda_p = 0$ corresponds to full payer reimbursement. These shares shape private incentives only in the decentralized analysis; in the planner's problem they are pure transfers. When we study the centralized operator, reimbursement can be imposed at the scheme level and the precise split is immaterial for welfare except insofar as it affects budget balance.

Timing and information (by regime). (i) *Centralized:* the scheme operator posts (f, m_p) and chooses control targets (δ, θ) to be implemented by participating banks. (ii) *Decentralized:* the scheme posts m_p as a rule, competitive sending banks post f, and sending banks and receiving banks privately and simultaneously choose δ and θ; these choices are noncontractable. (iii) Each payer observes m_j and decides whether to use instant payments (subject to $m_j \leq m_p$) or the outside option. (iv) If routed to instant payments, the transaction executes immediately; with probability $\delta\theta$ a fraud event succeeds and losses are allocated according to the prevailing regime (decentralized $(\lambda_s, \lambda_r, \lambda_p)$ or scheme reimbursement). Resource costs $e(\delta)$ and $\psi(\theta)$ are incurred per processed instant payment.

[10] This reflects current practice on retail instant rails, where recipients are typically not charged per incoming payment. The assumption implies an asymmetry that matters later: sender-side costs pass through the competitive fee into the composition of routed payments, whereas receiver-side costs do not. Allowing receiving banks to price incoming payments, or introducing an interbank transfer akin to an interchange fee, would ~~re~~introduce pass-through on the receiving side; we abstract from this margin.

3 Planner

We begin with a planner's benchmark that abstracts from institutional details of pricing, bargaining, and reimbursement splits. Throughout this section, reimbursement and liability allocations are irrelevant for welfare except through feasibility or budget balance; they are pure transfers from the planner's perspective. This benchmark therefore isolates the efficient cap and effort margins that later sections seek to decentralize or implement via observable instruments such as fees and eligibility rules.

The planner chooses (i) which payment sizes are routed through instant payments and (ii) the intensity of fraud controls on the sending and receiving sides. Because the convenience gain from immediacy is concave while expected fraud harm is proportional to size, the efficient routing set takes a threshold form: instant payments is used for low- and mid-sized payments, while sufficiently large payments are diverted to the outside option.

An allocation is a triple (A, δ, θ), where $A \subseteq [0, M]$ is a measurable admission set of payment sizes routed through instant payments and (δ, θ) are control intensities.

For fixed (δ, θ), define

$$C = e(\delta) + \psi(\theta), \qquad R = \delta\theta,$$

and let the net contribution from admitting a payment of size m be

$$G(m; \delta, \theta) \equiv \phi(m) - C - Rm.$$

Welfare from admitting a measurable set A is

$$W(A, \delta, \theta) = \bar{u} + \frac{1}{M} \int_A G(m; \delta, \theta)\, dm.$$

Thus, for fixed efforts, the planner admits precisely those payment sizes for which net surplus is nonnegative.

Assumption 1 (Value of small payments)**.** *Let*

$$K := \inf_{(\delta,\theta)\in(0,1]^2} \Big\{ e(\delta) + \psi(\theta) + \delta\theta M \Big\}.$$

Then $\phi(0) > K$.

The condition evaluates prevention costs at the largest possible payment size M, providing a primitive bound that ensures positive net value at the lower boundary under any optimal control profile.

Lemma 1 (Lower-interval admission and reduction)**.** *Suppose that* ϕ *is continuous, increasing, and strictly concave on* $[0, M]$*, Assumption 1 holds, and the prevention-cost conditions in the model setup are satisfied. Then*

$$A^\star = [0, m_p] \qquad \text{for some } m_p \in (0, M].$$

If $m_p < M$, then m_p is the unique solution to

$$\phi(m_p) = e(\delta^\star) + \psi(\theta^\star) + \delta^\star \theta^\star m_p.$$

Let $\Phi(x) \equiv \int_0^x \phi(t)\,dt$. With a uniform distribution of payment sizes, welfare on the lower interval $[0, m_p]$ is

$$W(m_p, \delta, \theta) = \bar{u} + \frac{1}{M}\Big\{\Phi(m_p) - \big(e(\delta) + \psi(\theta)\big)m_p - \tfrac{1}{2}\,\delta\,\theta\,m_p^2\Big\}. \tag{1}$$

Then the lemma reduces the planner's problem to the choice of (m_p, δ, θ) to maximize (1).

The optimal policy balances three margins: (i) how far up the size distribution instant payments should be used (the admission cap m_p); (ii) how stringent the sender's front-end controls should be; and (iii) how stringent the receiver's monitoring should be. At an interior solution, these trade-offs are summarized by

$$\phi(m_p) - \big(e(\delta) + \psi(\theta)\big) - \delta\,\theta\,m_p = 0, \tag{2}$$

$$e'(\delta) + \tfrac{1}{2}\,\theta\,m_p = 0, \tag{3}$$

$$\psi'(\theta) + \tfrac{1}{2}\,\delta\,m_p = 0. \tag{4}$$

Condition (2) pins down the admission boundary: at the top of the routed range, the marginal convenience gain from speed equals the marginal social cost, which consists of resource costs $e(\delta) + \psi(\theta)$ and expected fraud harm $\delta\theta m_p$. Conditions (3)–(4) determine how the planner allocates preventive intensity across the two sides of the transaction: marginal tightening costs are equated to marginal expected fraud loss saved. The factor $\frac{1}{2}$ reflects that, with $m \sim U[0, m_p]$, a marginal change in either control affects expected routed losses evaluated at the average admitted size, $\mathbb{E}[m \mid m \le m_p] = m_p/2$.

4 Centralized regime

Having characterized the efficient benchmark, we now analyze a *centralized governance* regime in which a scheme operator (regulator/utility) sets the rail's publicly observable access rules and coordinates fraud-prevention standards on both sides of the transaction. We treat this regime as a *benchmark* rather than a literal description of any single real-world rail: many systems delegate important decisions to participating banks, but scheme-level rules and common standards remain central design levers. The conceptual appeal of coordination is that by choosing (δ, θ) jointly, the operator internalizes cross-bank spillovers in fraud prevention that are at the core of the decentralized problem.

Precisely because effort is coordinated in this regime, the main economic friction is *implementation*. Unlike the planner, who can directly select an arbitrary measurable admission

set $A \subseteq [0, M]$, the operator must induce usage through a restricted and publicly posted menu: a size cap m_p and a flat per-payment access price F. Payers then self-select based on their realized transfer size m. While the efficient benchmark typically features a lower interval $[0, m_p]$, a flat fee can induce a band $[m_0, m_p]$ by screening out low-value transactions.

We proceed in two steps. First, we map a posted policy (F, m_p, δ, θ) into the induced instant payments usage region and expected surplus. Second, we solve for the operator's optimal policy under alternative objectives and compare it to the planner's cutoff rule, isolating which wedges arise from the restricted pricing and access instruments rather than from decentralization.

In the centralized regime, a single system operator commits to a contract $\langle F, \delta, \theta, m_p \rangle$ before transactions occur. Payments up to a size cap $m_p \in [0, M]$ are eligible for the rail and face a flat per-payment fee $F \geq 0$. [11] The operator also determines effort levels (δ, θ) and bears the resource costs $e(\delta)$ and $\psi(\theta)$. Throughout this section we assume the operator imposes *full reimbursement to the payer* in case of a successful APP event; the resource loss $\delta\,\theta\,m$ is still incurred in production.

Given F and m_p, a payer with transfer size m adopts the rail iff $F \leq \phi(m)$ and $m \leq m_p$. With monotone ϕ, usage is an interval $[m_0, m_p]$ with $m_0 \equiv g(F)$, where g is the (left–continuous) generalized inverse of ϕ on $[0, M]$, with the conventions $g(F) = 0$ if $F \leq \phi(0)$ and $g(F) = M$ if $F > \phi(M)$.

The expected unit resource cost per processed instant payment is

$$c(\delta, \theta; F, m_p) \;=\; e(\delta) + \psi(\theta) + \delta\,\theta\,\frac{m_p + g(F)}{2}, \tag{5}$$

and budget balance (when imposed) requires

$$F \;=\; c(\delta, \theta; F, m_p). \tag{6}$$

4.1 Operator objectives

The planner benchmark provides a welfare yardstick, but an actual scheme operator may face mandates or governance constraints that induce objectives other than welfare maximization. Therefore we consider outcomes when the the operator works with alternative objective functions that are standard in policy and industry discussions. We focus on three benchmarks: (i) maximizing adoption subject to break-even, (ii) maximizing payer surplus subject to break-even, and (iii) profit maximization without a break-even constraint. In each case, the policy (F, m_p) determines the induced usage region through $m_0 = g(F)$, and when break-even is imposed the fee must satisfy the fixed point in (6). We use these mappings to

[11] We interpret F as the effective price of access: it need not be a literal surcharge paid by the consumer, but can represent interbank or scheme-level fees or other shadow costs that are passed through by banks.

characterize the operator's optimal choices and to compare the resulting allocation to the planner's.

4.1.1 Adoption maximization (break-even)

This objective captures a scheme whose primary mandate is diffusion of the fast-payment technology, but that must still recover operating and expected fraud costs. Because the operator cannot price by size, expanding adoption can be pursued either by lowering the access price F or by relaxing fraud controls (δ, θ) to reduce unit cost. The break-even constraint therefore creates a direct trade-off between volume and security.

Formally, the operator maximizes the fraction of payers who use instant payments subject to break-even:

$$\text{(A)} \quad \max_{F,\delta,\theta,m_p} \; s(F,\delta,\theta,m_p) = \frac{m_p - g(F)}{M} \quad \text{s.t.} \quad F = c(\delta,\theta;F,m_p).$$

4.1.2 Payer-surplus maximization (break-even)

This objective corresponds to a public-utility operator that prices at cost and chooses policy to maximize payers' expected surplus from using the instant rail:

$$\text{(PS)} \quad \max_{F,\delta,\theta,m_p} \; S(F,\delta,\theta,m_p) \quad \text{s.t.} \quad F = c(\delta,\theta;F,m_p), \quad m_p \in [0,M],$$

where

$$S(F,\delta,\theta,m_p) = \bar{u} + \frac{1}{M}\int_{g(F)}^{m_p} \Big[\phi(m) - F\Big]\, dm.$$

Using $\Phi(x) = \int_0^x \phi(t)\, dt$ and imposing break-even yields

$$S(F,\delta,\theta,m_p) = \bar{u} + \frac{1}{M}\Big[\Phi(m_p) - \Phi(g(F)) - c(\delta,\theta;F,m_p)\,(m_p - g(F))\Big].$$

4.1.3 Profit maximization

This objective is a reduced-form benchmark for a platform-like operator or a coalition of banks that can set a markup over cost. It highlights how market power interacts with the same instrument restriction: because demand is screened by a flat fee, a markup raises F, shifts the lower cutoff upward, and can exclude socially valuable low-size transactions.

Formally,

$$\text{(P)} \quad \max_{F,\delta,\theta,m_p} \; D(F;m_p)\,\big[F - c(\delta,\theta;F,m_p)\big], \qquad m_p \in [0,M],$$

with demand $D(F;m_p) = \dfrac{(m_p - g(F))_+}{M}$.

4.2 Solution

Two wedges recur across centralized objectives. First, a *flat-fee entry wedge*: the fee generally induces the admitted set $[m_0, m_p]$ rather than the planner's lower interval $[0, m_p]$. Second, an *objective-induced scale wedge*: different operator objectives choose different fee–cap pairs (F, m_p), hence different usage bands $[m_0, m_p]$ and different average routed sizes $\bar{m} = (m_0 + m_p)/2$. Because break-even and effort conditions depend on $\bar{m}$, fraud controls vary with the operator's objective. Thus, even under centralized coordination, alternative operator objectives generally do not implement the planner's solution.

Under break-even, $F = c(\delta, \theta; F, m_p)$ and adoption is $s = (m_p - g(F))/M$. The next result shows that the operator chooses efforts as if the relevant scale is the average admitted payment size.

Proposition 1 (Adoption maximization)**.** *Let ϕ be strictly increasing with $g = \phi^{-1}$. At any solution to (A) with positive adoption, the admitted set is $[m_0, m_p]$ with $m_0 = g(F)$ and*

$$F = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,\big(m_p + m_0\big),$$

and the first-order conditions are

$$e'(\delta) + \tfrac{1}{2}\,\theta\big(m_p + m_0\big) = 0, \qquad \psi'(\theta) + \tfrac{1}{2}\,\delta\big(m_p + m_0\big) = 0, \tag{7}$$

$$\phi'(m_0) = \delta\,\theta, \qquad m_0 = g\Big(e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\big(m_p + m_0\big)\Big).$$

A notable implication is that maximizing adoption need not preserve access for the smallest payments. Under flat break-even pricing, extending eligibility to larger payments raises average fraud losses and therefore the common fee, so the operator may sacrifice low-value payments at the bottom to expand usage at the top.

The payer-surplus objective internalizes the payer-side gains from speed on the admitted region, but it remains subject to the same pricing technology. The next proposition isolates a clean case in which this objective collapses to the planner's problem: when the operator can guarantee zero minimum transaction size.

Proposition 2 (When payer-surplus implements the planner)**.** *If $F \le \phi(0)$ so that $m_0 = g(F) = 0$, then payer-surplus maximization under break-even is equivalent to the planner's problem:*

$$\max_{\delta,\theta,m_p}\ \bar{u} + \frac{1}{M}\Big[\Phi(m_p) - \big(e(\delta) + \psi(\theta)\big)m_p - \tfrac{1}{2}\,\delta\,\theta\,m_p^2\Big],$$

with interior FOCs

$$\phi(m_p) - \big(e(\delta) + \psi(\theta)\big) - \delta\,\theta\,m_p = 0, \quad e'(\delta) + \tfrac{1}{2}\,\theta\,m_p = 0, \quad \psi'(\theta) + \tfrac{1}{2}\,\delta\,m_p = 0.$$

When $m_0 = 0$, the flat fee no longer distorts the shape of usage, and payer surplus coincides with welfare over the planner-relevant interval $[0, m_p]$. In that case, the centralized operator selects the same cap and effort levels as the planner.

Under profit maximization, the operator trades off margin and volume. The same band logic applies, but now the operator is free to choose a markup and therefore typically sets a higher fee and a higher lower cutoff.

Proposition 3 (Profit maximization)**.** *At any interior optimum with $D > 0$, the admitted set is an upper band $[m_0, m_p]$ with $m_0 = g(F)$, and the effort conditions are*

$$e'(\delta) + \tfrac{1}{2}\,\theta\,\big(m_p + m_0\big) = 0, \qquad \psi'(\theta) + \tfrac{1}{2}\,\delta\,\big(m_p + m_0\big) = 0.$$

The cap satisfies

$$F - c(\delta, \theta; F, m_p) = \tfrac{1}{2}\,\delta\,\theta\big(m_p - m_0\big),$$

and, unless the fee is constrained to satisfy $F \le \phi(0)$, the operator generally chooses $m_0 > 0$.

Relative to the break-even cases, profit maximization increases the flat-fee entry wedge. The difference across objectives lies in how the fee–cap pair is determined: under adoption maximization the fee is pinned down by break-even, while under profit maximization it is chosen to trade off margin against the volume of admitted transactions. But in both cases the induced usage band $[m_0, m_p]$ feeds back into incentives through the average routed size.

4.3 First-best implementability

Centralized coordination alone is not sufficient for first-best implementation: with a flat fee, usage is generally a band $[m_0, m_p]$, and alternative operator objectives choose different bands and hence different incentives. The next result shows that these distortions disappear if regulation directly fixes the two margins that the flat-fee regime otherwise distorts: the lower edge of usage and the upper cap.

Suppose the regulator enforces zero minimum transaction size $m_0 = 0$ and mandates the planner's cap $m_p = m_p^{\mathrm{FB}}$. Then the admitted set is mechanically fixed at $[0, m_p^{\mathrm{FB}}]$, so the average routed size is $\bar{m} = m_p^{\mathrm{FB}}/2$. The fee F can still redistribute surplus or satisfy budget balance, but it no longer affects who uses instant payments or the scale on which fraud controls operate.

Proposition 4 (Centralized implementation under fixed access and cap)**.** *Suppose the regulator enforces zero minimum transaction size $m_0 = 0$ and mandates $m_p = m_p^{\mathrm{FB}}$. Under centralized governance, the planner's effort profile $(\delta_{\mathrm{FB}}, \theta_{\mathrm{FB}})$ is implemented under payer-surplus maximization subject to break-even and under profit maximization.*

Once zero minimum size and the planner's cap are imposed, the usage set is fixed at the planner-relevant interval $[0, m_p^{\mathrm{FB}}]$ and the access price F affects only transfers and, where relevant, budget balance, not efficiency. The flat-fee entry wedge disappears, and operator objectives no longer generate different usage bands or different average routed sizes. The remaining efficiency margin is the choice of fraud-prevention efforts. Under payer-surplus maximization and profit maximization, this margin reduces to minimizing the cost of serving the admitted set. Under adoption maximization, usage is already pinned down by the mandates, so the operator is indifferent across effort profiles unless a tie-breaking criterion is imposed; with least-cost tie-breaking, it selects the same cost-minimizing effort profile. Thus, in all three cases, the effort problem collapses to the planner's effort subproblem, and pricing affects only transfers and financing rather than the implemented allocation.

The centralized implementation result should be read as an allocation result, not as a claim that the access fee can always finance the rail. Enforcing $m_0 = 0$ requires the effective usage price to satisfy $F \leq \phi(0)$. If the break-even cost of operating the rail and providing fraud controls exceeds this bound, the missing revenue must be raised through an instrument that does not distort the marginal routing decision, for example a fixed participation assessment on banks, membership fees, or external funding. Conversely, if cost recovery must be imposed payment-by-payment through a flat usage fee, first-best implementation may fail.

4.4 Size-contingent fees

The preceding analysis shows that a flat usage fee can distort centralized implementation by inducing a band of users rather than the planner's lower interval. This raises a natural question: is the centralized wedge merely an artifact of flat pricing? This subsection shows that the answer is partly yes and partly no. Size-contingent fees remove the flat-fee entry wedge, and they allow first-best implementation under objectives that internalize total surplus. However, they do not make adoption maximization equivalent to welfare maximization when the operator is subject only to aggregate break-even. The reason is that aggregate break-even allows cross-subsidization across payment sizes.

For fixed (δ, θ), define

$$G(m; \delta, \theta) = \phi(m) - e(\delta) - \psi(\theta) - \delta\theta m.$$

This is the social contribution from admitting a payment of size m. The planner admits payment size m if and only if

$$G(m; \delta, \theta) \geq 0.$$

The key distinction is therefore between pointwise surplus and aggregate feasibility. With size-contingent fees, the operator can make participation depend on the payment size, so

the flat-fee selection wedge is removed. But if break-even is imposed only in aggregate, the operator can also transfer surplus across payment sizes. The following proposition shows how this distinction affects different centralized objectives.

Proposition 5 (Size-contingent fees under aggregate break-even)**.** *Suppose the centralized operator can set a size-contingent fee $F(m)$, subject to*

$$0 \leq F(m) \leq \phi(m)$$

for any admitted payment size m, and suppose cost recovery is imposed only in aggregate.

(i) Under payer-surplus maximization subject to aggregate break-even, the operator implements the planner's allocation.

(ii) A profit-maximizing operator implements the planner's allocation if it can perfectly price discriminate.

(iii) An adoption-maximizing operator subject to aggregate break-even generally does not implement the planner's allocation. Starting from the planner's allocation, it can increase adoption by using inframarginal surplus to cross-subsidize payments with negative marginal social contribution.

Size-contingent cost-based pricing removes the flat-fee selection wedge. Payer-surplus maximization and perfect price discrimination then align the operator's objective with total surplus. Adoption maximization is different: its objective is volume rather than surplus. The planner stops where marginal surplus is zero, while the adoption-maximizing operator expands beyond that point. Therefore, size-contingent fees eliminate the flat-fee wedge, but they do not eliminate the objective wedge.

5 Decentralized fraud control and liability

We now turn to a decentralized governance regime in which fraud prevention and pricing are chosen at the bank level rather than coordinated by a scheme operator. The key distinction from Section 4 is incentive alignment: because each bank internalizes only its own share of expected fraud losses, liability allocation becomes a policy lever that shapes equilibrium effort on both sides of the transaction.

The timing is as follows. The scheme fixes an eligibility cap $m_p \in (0, M]$ as an exogenous rule. Taking m_p as given, the sending bank posts a flat per-payment fee f and chooses sender-side controls $\delta \in (0, 1]$, while the receiving bank chooses receiver-side monitoring $\theta \in (0, 1]$. As in the centralized section, f should be interpreted as the effective user-facing

cost of instant payments access: even when retail apps appear "free," usage typically entails implicit charges (e.g., account pricing, cross-subsidies, data monetization, or user frictions) that limit the processing of very small or net-negative transactions.

Given (f, m_p), a payer with transfer size m routes through instant payments if and only if $m \leq m_p$ and $f \leq \phi(m)$. With ϕ increasing, the induced usage region again takes an interval form: the admitted set is the band $[m_0, m_p]$, where $m_0 \equiv g(f)$ is the (generalized) inverse cutoff of ϕ. The adoption share is therefore $s = (m_p - m_0)_+/M$. Conditional on adoption, the average routed payment size is $\bar{m} = (m_p + m_0)/2$. Positive adoption requires the upper-edge participation condition $f \leq \phi(m_p)$; because m_p is imposed by the scheme rather than chosen by the banks, this condition need not bind in equilibrium.

If a routed payment is successfully misdirected, expected loss equals $\delta\,\theta\,m$. In the decentralized regime we assume the payer is reimbursed, and a constant liability parameter $\lambda \in [0,1]$ splits the loss across banks: the sender bears $\lambda\,\delta\,\theta\,m$ and the receiver bears $(1-\lambda)\,\delta\,\theta\,m$.[12]

5.1 Equilibrium

This subsection characterizes the decentralized outcome when the scheme fixes the eligibility cap m_p and competitive sending banks set a flat user fee while privately choosing sender-side controls. Two forces jointly determine equilibrium usage. First, Bertrand competition drives the fee to expected private cost. Second, because the fee is flat, payers self-select by transfer size, which endogenously determines the lower cutoff m_0 and therefore the average size of routed payments. The resulting fixed point is central: the fee affects the composition of users, and that composition feeds back into average exposure and pricing. Throughout the decentralized analysis, we restrict attention to parameter configurations for which the equilibrium is well defined and regular.[13]

[12] Allowing the payer to bear a positive share only redistributes private payoffs and does not affect the sending-bank–receiving-bank incentive structure that drives the results below.

[13] For each liability share $\lambda \in [0,1]$ and average routed payment size $\bar{m} \in (0, M]$, we assume that the sender–receiver effort subgame admits a unique equilibrium $\big(\delta(\lambda,\bar{m}), \theta(\lambda,\bar{m})\big)$, that the equilibrium varies continuously with $(\lambda,\bar{m})$, and that it is continuously differentiable wherever both effort choices are interior. Near the liability endpoints, the corresponding effort choice lies at the upper boundary: $\delta = 1$ for sufficiently small λ and $\theta = 1$ for sufficiently large λ, with the interior first-order conditions replaced by the corresponding Kuhn–Tucker conditions.

When the lower cutoff is endogenous, we additionally assume that the fee–composition map

$$T_{\lambda,m_p}(m_0) = \min\left\{ g\left(f\left(\lambda, \frac{m_p + m_0}{2} \right) \right), m_p \right\}$$

has a unique fixed point on $[0, m_p]$. These conditions make equilibrium efforts, pricing, composition, and welfare well defined and continuous in λ, ensuring the existence of an optimal liability share.

Under Bertrand competition over the flat fee, the sender prices at average private cost,

$$f \;=\; c_S(\delta, \theta; \lambda, m_p, m_0), \qquad c_S(\delta, \theta; \lambda, m_p, m_0) = e(\delta) + \lambda\, \theta\, \delta\, \bar{m},$$

with

$$\bar{m} = \mathbb{E}[m \mid m \in [m_0, m_p]] = \frac{m_p + m_0}{2}, \qquad m_0 = g(f) = g\big(c_S(\delta, \theta; \lambda, m_p, m_0)\big).$$

The last equality emphasizes the equilibrium feedback: a higher fee raises the lower cutoff $m_0 = g(f)$, which increases the average routed size $\bar{m}$ and, through expected fraud exposure, feeds back into the sender's cost and hence into the fee.

Taking the scheme-set cap m_p as given, an interior decentralized outcome $(\delta^*, \theta^*, m_0^*, f^*)$ satisfies

$$e'(\delta^*) = -\,\lambda\, \theta^*\, \frac{m_p + m_0^*}{2}, \tag{S}$$

$$\psi'(\theta^*) = -(1-\lambda)\, \delta^*\, \frac{m_p + m_0^*}{2}, \tag{R}$$

$$f^* = e(\delta^*) + \lambda\, \theta^*\, \delta^*\, \frac{m_p + m_0^*}{2}, \qquad m_0^* = g(f^*). \tag{P}$$

Conditions (S)–(R) equate marginal tightening costs to marginal privately internalized fraud exposure, evaluated at the average routed payment size $\bar{m} = (m_p + m_0^*)/2$ and scaled by the liability shares λ and $1-\lambda$. The pricing condition (P) then pins down f^* given these private choices.

We assume the scheme chooses m_p so a positive-volume outcome is feasible. Given $m_0^* = g(f^*)$, the admitted set is $[m_0^*, m_p]$; adoption is positive iff $m_0^* < m_p$.

The decentralized equilibrium departs from the planner for two distinct reasons. First, flat-fee pricing creates an entry wedge: unless $f^* \le \phi(0)$, the lower edge $m_0^* = g(f^*)$ is strictly positive. Second, effort incentives are shaped by the liability split. The best-response conditions (S)–(R) allocate marginal fraud harm according to $(\lambda, 1-\lambda)$, so each bank internalizes only its assigned share. As a result, even if zero minimum transaction size were imposed ($m_0^* = 0$), the decentralized effort conditions generally do not coincide with the planner's.

A primitive sufficient condition for the determinant condition used in the interior comparative statics is

$$\left(\inf_{\delta \in (0,1]} e''(\delta)\right)\left(\inf_{\theta \in (0,1]} \psi''(\theta)\right) > \frac{M^2}{4},$$

because it implies

$$e''(\delta)\psi''(\theta) - \lambda(1-\lambda)\bar{m}^2 > 0$$

throughout the domain. A contraction condition on T_{λ, m_p} is sufficient for uniqueness of the fee–composition fixed point. We maintain these regularity conditions to focus on the allocation of prevention incentives rather than on equilibrium multiplicity.

5.2 Two implementations under decentralization

We now ask whether observable policy instruments can implement the planner's allocation when banks' security efforts are non-contractible. Relative to the first best, decentralization introduces two wedges. First, flat-fee pricing generates a distorted admitted region $[m_0, m_p]$, rather than the planner's $[0, m_p]$. Second, the liability rule divides fraud losses between the two banks, so each bank internalizes only part of the benefit from its security effort.

We proceed in two steps. First, we eliminate the usage distortion by mandating the planner's admitted region. Holding that region fixed, we then address the partial internalization of security benefits and show how the planner's effort choices can be implemented.

Step 1: Mandate the admitted region. The regulator fixes the lower edge at zero minimum transaction size and the upper edge at the planner's boundary by imposing

$$m_0 = 0 \quad \text{and} \quad m_p = m_p^{\mathrm{FB}}.$$

Operationally, $m_0 = 0$ requires a cap on the usage fee, $f \leq \phi(0)$, while m_p is a scheme rule.

With $m_0 = 0$ and $m_p = m_p^{\mathrm{FB}}$, the routed average size is $\bar{m} = m_p^{\mathrm{FB}}/2$. The sending bank sets

$$f^* = \min\Big\{ c_S(\delta, \theta; \lambda),\ \phi(0) \Big\}, \qquad c_S(\delta, \theta; \lambda) = e(\delta) + \lambda\, \delta\, \theta\, \frac{m_p^{\mathrm{FB}}}{2}.$$

Because m_0 and m_p are fixed, $\bar{m}$ is constant and the effort first-order conditions do not involve f:

$$\text{(Sending bank) } e'(\delta) = -\lambda\, \theta\, \frac{m_p^{\mathrm{FB}}}{2}, \qquad \text{(Receiving bank) } \psi'(\theta) = -(1-\lambda)\, \delta\, \frac{m_p^{\mathrm{FB}}}{2}.$$

Hence the fee cap guarantees $m_0 = 0$ without distorting effort choices.

Step 2: Restore effort incentives. We now present two implementation schemes that collapse the split-incentive problem and deliver the planner's allocation on the mandated usage region.

Proposition 6 (FB implementation under decentralization: two schemes)**.** *Suppose the regulator enforces zero minimum transaction size $m_0 = 0$ and mandates $m_p = m_p^{\mathrm{FB}}$. Then the planner's allocation on $[0, m_p^{\mathrm{FB}}]$ is implemented by either:*

- *Scheme A (double liability). Charge* each *bank the full expected loss $\delta\, \theta\, m$ on every fraudulent instant payment.*
- *Scheme B (single-sided liability with one mandated effort).*

  (B1) Set $\lambda = 1$ and mandate $\theta = \theta_{\mathrm{FB}}$.

(B2) Set $\lambda = 0$ and mandate $\delta = \delta_{\mathrm{FB}}$.

Either scheme removes the split-incentive problem so that private effort choices satisfy the planner's first-order conditions. The flat fee affects only surplus redistribution and is inessential for welfare; the scheme operator acts as the budget-breaker, absorbing the excess assessments.

As in the centralized case, this result implements the first-best allocation but does not explain how its costs are financed. Admitting even the smallest payments requires $f \leq \phi(0)$. If this fee does not cover costs, the shortfall must be financed separately in a way that preserves banks' incentives to prevent fraud.

5.3 Optimal liability sharing

The previous subsection showed that the first best can be implemented if the regulator can combine several tools. We now study the natural second-best instrument: a simple liability split that reallocates fraud losses between the sending bank and the receiving bank. To isolate the incentive role of liability, we take the admitted region as fixed. Specifically, suppose the regulator mandates zero minimum transaction size, $m_0 = 0$, and fixes the eligibility cap m_p. The cap need not equal the planner's first-best cap. Then access and scale are fixed:

$$\bar{m} = \frac{m_p}{2} \equiv \mu, \qquad s = \frac{m_p}{M},$$

so liability affects welfare only through decentralized prevention incentives.

For a sender liability share $\lambda \in [0, 1]$, the decentralized equilibrium (δ, θ, f) solves

$$\text{(Sending bank)} \quad e'(\delta) + \lambda\,\theta\,\mu = 0, \tag{S_0}$$

$$\text{(Receiving bank)} \quad \psi'(\theta) + (1-\lambda)\,\delta\,\mu = 0, \tag{R_0}$$

$$\text{(Pricing)} \quad f = \min\{e(\delta) + \lambda\,\delta\,\theta\,\mu,\ \phi(0)\}\,. \tag{P_0}$$

Because m_0 and m_p are fixed, pricing decouples from the effort margins.

The regulator chooses λ to maximize welfare. With the admitted region fixed, average convenience and adoption are constant. Hence the regulator's problem is equivalent to minimizing the per-routed-payment social cost

$$C(\lambda) = e(\delta(\lambda)) + \psi(\theta(\lambda)) + \delta(\lambda)\theta(\lambda)\mu.$$

Equivalently,

$$\lambda^\star \in \arg\min_{\lambda \in [0,1]} C(\lambda).$$

We first characterize the optimal split and then discuss when the solution is sender-heavy or close to a corner.

Proposition 7 (Interior characterization of optimal liability)**.** *Suppose* $\lambda^\star \in (0,1)$ *and the associated effort choices are interior. Then the optimal liability split satisfies*

$$\lambda^\star e''(\delta^\star)(\delta^\star)^2 = (1-\lambda^\star)\psi''(\theta^\star)(\theta^\star)^2,$$

Thus,

$$\lambda^\star = \frac{\psi''(\theta^\star)(\theta^\star)^2}{\psi''(\theta^\star)(\theta^\star)^2 + e''(\delta^\star)(\delta^\star)^2}.$$

The optimal liability split depends on the relative adjusted curvature of the two prevention technologies. Because fraud risk depends on the product $\delta\theta$, the relevant objects are not raw curvatures alone, but

$$e''(\delta^\star)(\delta^\star)^2 \qquad \text{and} \qquad \psi''(\theta^\star)(\theta^\star)^2.$$

If receiver-side monitoring is locally more convex in this adjusted sense, the optimal sender liability share is higher. If sender-side screening is locally more convex, the optimal sender liability share is lower. Thus liability is loaded toward the side whose prevention margin is more flexible. A 50/50 split is optimal only when the two adjusted curvatures are equal. More generally, the optimal split is interior whenever neither prevention technology is sufficiently more rigid than the other.

Corollary 1 (Sender-heavy liability)**.** *If receiver-side monitoring is sufficiently more convex than sender-side screening at the optimum, in the sense that*

$$\psi''(\theta^\star)(\theta^\star)^2 \gg e''(\delta^\star)(\delta^\star)^2,$$

then the optimal sender liability share is close to one:

$$\lambda^\star \approx 1.$$

The sender-heavy case corresponds to environments in which receiver-side monitoring is the stiffer margin. This interpretation is natural in instant payments. Sending-bank controls are often transaction-level measures applied at payment initiation, such as authentication, warning screens, payment holds, and transaction-specific risk scoring. Receiving-bank controls are more naturally interpreted as account- and network-level monitoring, including account vetting, mule-account detection, and monitoring of incoming funds. These activities may be harder to adjust at the margin. In such an environment, shifting liability toward the sending bank reallocates prevention incentives toward the more flexible prevention margin. The opposite limiting case is symmetric: when sender-side screening is much stiffer than receiving-bank monitoring, the optimal sender liability share approaches zero, $\lambda^\star \approx 0$.

Remark 1 (Fraud losses and welfare)**.** *The liability split that minimizes fraud losses need not coincide with the liability split that maximizes welfare. Per-routed-payment fraud dollars are*

$$L(\lambda) = \delta(\lambda)\theta(\lambda)\mu.$$

A higher sender liability share may reduce fraud losses by inducing tighter sender-side controls, but it may also raise resource costs through excessive prevention effort. Thus fraud losses are an important outcome variable, but they are not sufficient for choosing the liability rule. The optimal split should be characterized from welfare, which trades off fraud reduction against the resource costs of prevention.

Remark 2 (Endogenous admitted set and pass-through asymmetry)**.** *When the admitted set is endogenous, liability affects not only prevention incentives but also payment selection through the user-facing fee. Because only the sending bank's private costs enter the competitive fee directly, changes in sender liability are passed through to users differently from changes in receiver liability. Consequently, even with symmetric prevention technologies, $e \equiv \psi$, the optimal liability split need not be* 50/50*, because it also depends on how fee pass-through changes access to the payment system.*

5.4 Centralized vs. decentralized: a discussion

Finally, we compare centralized and decentralized governance in terms of when one regime delivers higher welfare than the other. The earlier sections isolate distinct wedges: under centralization, flat-fee pricing can distort the admitted region; under decentralization, liability sharing distorts effort incentives. Here we combine those forces in two benchmarks. First, we evaluate a best-case comparison in which the scheme mandates zero minimum transaction size and fixes the cap at the planner's boundary, so the two regimes differ only in how they implement effort. Second, we drop these mandates and use a numerical illustration to show that the welfare ranking can reverse: the selection distortion induced by flat-fee centralization can outweigh the incentive gains from coordinated effort, yielding parameter regions where decentralization performs better.

5.4.1 Benchmark 1: mandated access and a planner cap

Suppose the scheme mandates zero minimum transaction size ($m_0 = 0$) and fixes the eligibility cap at the planner's boundary ($m_p = m_p^{\mathrm{FB}}$). Under these mandates, the admitted region is identical across regimes, and the remaining question is whether governance can align effort incentives.

On the centralized side, Proposition 4 shows that the operator can implement the first best: once the scheme fixes access and the cap, the operator's coordinated choice of (δ, θ) coincides with the planner's, and the fee only redistributes surplus subject to budget balance.

By contrast, under decentralization with a single liability share, each bank internalizes only its assigned fraction of marginal fraud harm when choosing controls. Even though the admitted region matches the planner by construction, the equilibrium effort conditions remain split, so the allocation generically differs from the planner's. Hence, under the mandated benchmark, centralized governance weakly dominates and is typically strictly welfare superior.

5.4.2 Benchmark 2: a numerical illustration without access/cap mandates

To illustrate that the welfare ranking can reverse when zero-minimum-access and planner-cap mandates are unavailable, we consider a simple parametric environment and compute welfare under each governance regime on a grid of parameters. Throughout the exercise, we fix a common exogenous eligibility cap $m_p = \bar{m}_p$ across regimes. The goal is a possibility result: even though a centralized operator coordinates (δ, θ), its flat break-even fee can induce a selection distortion that is large enough for decentralization to dominate.

Payment convenience is

$$\phi(m) = \Delta_0 + \beta\big(1 - e^{-m}\big),$$

sender and receiver effort costs are

$$e(\delta) = \bar{e} + \alpha(1-\delta)^2, \qquad \psi_\kappa(\theta) = \kappa\Big[\bar{\psi} + \tfrac{1}{2}\big(\ln(\theta_0/\theta)\big)^2\Big],$$

and payment sizes satisfy $m \sim U[0, M]$.[14] Unless otherwise stated we set

$$(\Delta_0, \beta, \bar{e}, \alpha, \bar{\psi}, \theta_0, M, \bar{m}_p) = (0.14,\ 0.44,\ 0.04,\ 2.5,\ 0.02,\ 0.4,\ 1,\ 0.8).$$

The parameter $\lambda \in [0,1]$ is the sender's loss share under decentralization. The parameter $\kappa > 0$ scales the receiver-side monitoring cost: a larger κ raises both the per-payment level of receiver-side cost and its curvature, so that stiffer monitoring is also costlier to operate. In this specification θ_0 is the dissipation probability under baseline monitoring, and the calibration sets α large so that sender-side screening is inflexible at the margin. This choice

[14] The parametric cost functions are chosen for tractability and depart from the maintained assumptions of Section 2 in two inessential ways. First, $e(\delta)$ is bounded as $\delta \downarrow 0$ and $\psi_\kappa(\theta)$ is minimized at $\theta_0 < 1$ rather than decreasing on all of $(0,1]$; the divergence and monotonicity conditions of Section 2 are used only for the existence result in Lemma 1, and all equilibria computed here are interior in δ and satisfy $\theta \le \theta_0$, where ψ_κ is decreasing. Second, θ_0 is interpreted as the dissipation probability under baseline monitoring, so the receiver's no-liability corner is $\theta = \theta_0$ rather than $\theta = 1$; this is the corner configuration described in footnote 13, relabeled.

deliberately mutes the incentive channel of liability—equilibrium efforts respond little to λ—in order to isolate the selection channel of Remark 2: liability then operates primarily through the composition of the user-facing fee.[15]

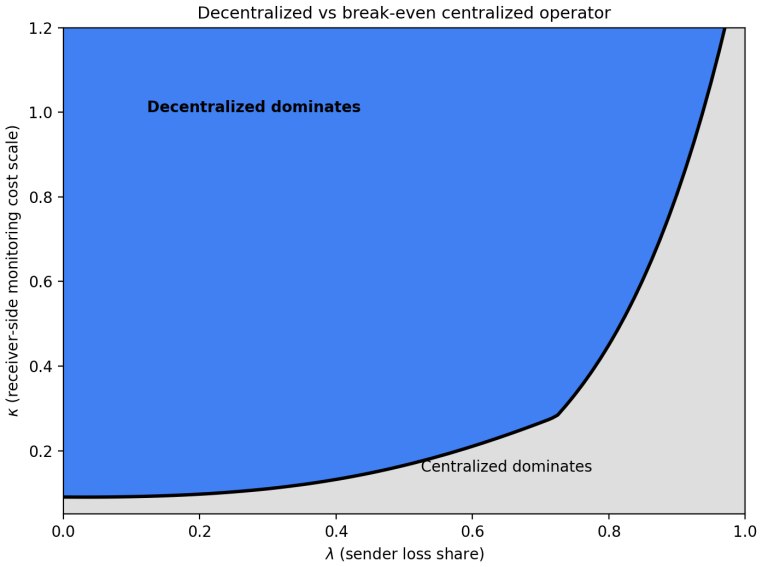


Figure 1: Dominance map in the (λ, κ) plane comparing a decentralized regime to a centralized break-even operator (equivalently, adoption- or payer-surplus-maximizing given the fixed cap). The black curve is the zero-gap locus; decentralization yields higher welfare above and to the left of the boundary. The figure is generated from the numerical solution described in the text.

Figure 1 plots the dominance regions in the (λ, κ) plane for the comparison between decentralization and a centralized break-even operator. With the cap fixed, adoption maximization and payer-surplus maximization coincide—both reduce to minimizing the break-even fee—so the comparison is against the best break-even flat-fee centralized regime; a profit-maximizing operator only enlarges the decentralized-dominance region.

The decentralized-dominance region is concentrated at high κ and low λ. The result reflects a pricing-selection effect. As κ rises, the receiver-side monitoring cost that a break-even centralized fee must recover grows, raising the lower cutoff m_0^C. Under decentralization,

[15]Fixing $m_p = \bar{m}_p$ isolates the governance-driven differences in effort incentives and fee-induced selection; allowing each regime to choose m_p endogenously yields the same qualitative message in our numerical experiments. The calibration restricts $\kappa \leq 1.2$: for sufficiently large κ the centralized break-even fee exceeds $\phi(\bar{m}_p)$ and the centralized rail attracts no users, so the comparison becomes degenerate. The fee–composition fixed point was verified to be unique at every grid point, consistent with the regularity conditions of footnote 13. Because the calibration makes both prevention margins inflexible, it illustrates a different parameter region from the sender-heavy case of Corollary 1, in which sender-side screening is the flexible margin.

receiver-side costs never enter the user-facing fee, and when λ is low the fee also excludes most of the expected fraud losses, so m_0^D remains at or near zero. Decentralization therefore preserves access for low-size payments that centralized pricing excludes, and in this region of the calibration the access gain dominates the loss from weaker prevention incentives.[16]

5.4.3 Policy takeaway

Two messages emerge. First, cap and access governance are first-order. If the scheme can mandate zero minimum transaction size and set the eligibility cap at the planner's boundary, then centralized governance implements the planner by construction: coordinated standards deliver the efficient effort profile, while the fee is relegated to cost recovery and surplus redistribution. Under the same mandates, a decentralized regime governed only by a single loss share generally does not implement the planner because effort incentives remain split across banks.

Second, when these mandates are not available and pricing must be implemented through simple instruments, the welfare ranking can reverse. The key reason is selection shape. A centralized break-even fee must recover all costs of the rail, including receiver-side monitoring costs, and can therefore exclude small payments whose value of speed is high. The decentralized access price, by contrast, contains only the costs the sending bank internalizes, so it can remain low enough to preserve access for these payments. When this access gain outweighs the benefits of coordinated effort, decentralization dominates.

6 Extensions and discussion

This section discusses two natural extensions of the model: payer-side moral hazard and endogenous fraudster entry. The goal is not to develop these extensions fully, but to clarify how they would interact with the baseline mechanisms. In both cases, the extensions add new incentive margins while preserving the paper's central message: access rules, eligibility caps, and prevention incentives must be designed jointly.

6.1 Payer care and reimbursement caps

The baseline model treats payers as passive. In that environment, reimbursement only reallocates ex post losses and does not affect the probability of successful fraud. In practice,

[16]The calibration is chosen to make this selection channel operative: sender-side screening is inflexible (α large), so lowering λ sacrifices little prevention, and κ raises the level of receiver-side cost, so stiffness translates into fee-relevant cost for the centralized operator. With a flexible sender margin, the dominance region instead shifts toward the second-best split of Section 5.3; the location of the region varies with the calibration, but its existence—the possibility result—does not depend on it.

however, reimbursement often depends on payer conduct: for example, the U.K. APP-fraud regime allows reimbursement to be denied in cases of gross negligence. Once payers can take costly care, the reimbursement cap becomes an incentive instrument. It affects both the payer's precaution and, in market implementation, the payer's willingness to use the instant rail.

Suppose that, conditional on using the instant rail, a payer chooses care $x \geq 0$. Care scales the residual hit probability through $q(x)$, where

$$q'(x) < 0, \qquad q''(x) > 0,$$

and entails cost $c(x)$, with

$$c'(x) > 0, \qquad c''(x) > 0.$$

If the reimbursement cap is L, then a payer facing a routed transfer of size m solves

$$x^*(m, L) \in \arg\min_{x \geq 0} \{c(x) + \delta\theta q(x)(m - L)_+\}.$$

At an interior solution,

$$c'\big(x^*(m, L)\big) = -\delta\theta q'\big(x^*(m, L)\big)(m - L)_+.$$

Thus, for payments above the cap, a higher reimbursement cap reduces the payer's residual exposure and weakens care.

Consider first a constrained planner that can choose the admitted set directly but cannot choose payer care. In this benchmark, reimbursement has no role in inducing participation because routing is chosen by the planner. If reimbursement is a pure transfer and payers are risk neutral, the constrained-efficient reimbursement cap is

$$L^{CFB} = 0.$$

The reason is that (L=0) makes the payer internalize the full monetary loss (m), so the payer's private care condition coincides with the social care condition, conditional on the planner's choices of (δ, θ) and the admitted set. This is not a policy recommendation, but a benchmark that isolates the payer-care incentive force.

This benchmark is generally not implementable in a non-planner regime. In centralized or decentralized implementation, payers choose whether to use the instant rail. Let

$$H(m, L; \delta, \theta) := \min_{x \geq 0} \{c(x) + \delta\theta q(x)(m - L)_+\}$$

denote the payer's minimized private exposure cost. Given a user-facing fee F, the payer participates if

$$\phi(m) - F - H(m, L; \delta, \theta) \geq 0.$$

A higher L lowers H and therefore relaxes participation. But the same increase in L weakens care and raises effective fraud risk. A zero cap preserves care incentives but may under-support adoption; a positive cap supports adoption but distorts care. Therefore, once payer-side moral hazard is present, even centralized coordination of bank-side controls need not implement the constrained-efficient allocation.

The conclusion changes if reimbursement has a direct consumer-protection value. For example, fraud may impose nonpecuniary harm on victims, such as stress, time costs, or loss of trust in the payment system. If this harm depends on the unreimbursed loss, write it as

$$D\big((m-L)_+\big), \qquad D' > 0.$$

Then reimbursement is no longer a pure transfer. Expected social harm from a successful fraud event becomes

$$m + D\big((m-L)_+\big).$$

A higher L directly reduces nonpecuniary harm, but it still weakens payer care. In that case, the optimal reimbursement cap can be positive, and possibly interior, because it balances consumer protection against payer moral hazard.

The broader lesson is that payer negligence does not make reimbursement a standalone solution. Instead, the reimbursement cap becomes an additional and imperfect policy tool. It can support participation and consumer protection, but it also weakens private care. Thus, when payer moral hazard matters, reimbursement rules must be interpreted as part of the broader implementation package, together with access, eligibility, and bank-side prevention incentives.

6.2 Endogenous fraudster entry

The baseline model treats the intensity of fraud attempts as exogenous. A natural extension is to allow fraudsters to choose whether to attack. This adds a deterrence margin: prevention affects not only the probability that an attempted fraud succeeds, but also the mass of fraudsters who find it profitable to enter. This channel reinforces the baseline emphasis on prevention incentives, because weak controls can increase both the success rate of fraud and the incidence of attempted fraud.

To see the additional margin, suppose fraudsters choose whether to attack a routed payment. A fraudster who attacks a payment of size m obtains expected gross payoff

$$\delta\theta m$$

and must pay an entry cost c, drawn from distribution Γ on $\mathbb{R}_+$. The fraudster attacks if

$$\delta\theta m \geq c.$$

Thus the mass of fraudsters who attack is

$$\Gamma(\delta\theta m),$$

and expected fraud loss becomes

$$\Gamma(\delta\theta m)\,\delta\theta m.$$

Relative to the baseline, prevention now has two effects. It lowers the probability that a given attack succeeds, and it also lowers the mass of fraudsters willing to attack. For example, differentiating expected fraud loss with respect to δ gives

$$\frac{\partial}{\partial\delta}\left\{\Gamma(\delta\theta m)\,\delta\theta m\right\} = \Gamma(\delta\theta m)\,\theta m + \gamma(\delta\theta m)(\theta m)\delta\theta m,$$

where γ is the density of Γ. The first term is the baseline conditional-success effect. The second term is the new entry-deterrence effect: tighter controls reduce the expected payoff from attacking and therefore reduce the mass of fraudsters who enter.

This additional deterrence channel strengthens the value of prevention. A decentralized bank that bears only a share of realized fraud losses internalizes only the corresponding share of both the conditional-success effect and the entry-deterrence effect. Thus endogenous fraudster behavior can make simple split-liability rules less effective as instruments for aligning private prevention incentives with social losses.

At the same time, the baseline access and eligibility wedges remain. A centralized operator may internalize the deterrence margin when choosing (δ, θ), but if usage is still induced through a flat fee, the admitted set can remain a band $[m_0, m_p]$ rather than the planner's lower interval $[0, m_p]$. Similarly, a decentralized regime still faces split prevention incentives unless liability or other instruments make each side internalize the relevant marginal harm.

Therefore, endogenous fraudster entry adds a new prevention margin, but the main implementation message is unchanged. Stronger controls deter attacks as well as reduce the success probability conditional on attack, making prevention more valuable. But first-best implementation still requires aligning both the admitted region and prevention incentives. Centralization becomes more valuable when deterrence is important, yet centralization alone remains insufficient if access and eligibility are distorted.

7 Conclusion

This paper studies two core institutional levers in instant-payment systems: the governance of fraud analytics, centralized versus decentralized, and the allocation of fraud losses through liability and reimbursement rules. Our analysis yields three broad conclusions. First, centralization by itself is not sufficient for efficiency. Even when a centralized operator coordinates

sender-side controls and receiver-side monitoring, flat-fee pricing can distort the admitted region away from the planner's lower interval, and different operator objectives can therefore induce different fraud-control choices. Second, decentralized liability is a powerful but imperfect instrument. A simple loss split between sending and receiving banks generally does not implement the first best, because each side internalizes only part of the benefit from tighter controls. The optimal liability split depends on the relative adjusted curvature of sender-side and receiver-side prevention technologies: liability is generally shared, but becomes sender-heavy when receiver-side monitoring is sufficiently less flexible than sender-side screening. Third, without mandates on access and eligibility, centralized and decentralized regimes need not be ranked uniformly: the gains from coordinated effort under centralization can be offset by distortions in the composition of admitted payments.

The policy implication is that improving outcomes on instant rails requires treating access rules, eligibility caps, and liability allocation as a package rather than as independent design choices. A centralized fraud service may be highly valuable, but its welfare performance depends on how the rail admits transactions. Likewise, reimbursement and liability rules matter not only for ex post consumer protection, but also for ex ante prevention incentives.

The framework is deliberately parsimonious, and several extensions remain important. We treat payer-side moral hazard only briefly, abstract from richer bargaining between payer and payee over the choice of rail, and model fraud-prevention technologies in reduced form. Extending the analysis along these margins would help connect the theory more closely to institutional practice. But the central lesson is already clear: in fast-payment systems, the design of fraud control cannot be separated from the design of access.

References

Ross Anderson and Tyler Moore. The economics of information security. *Science*, 314(5799): 610–613, 2006.

Vimal Balasubramaniam, Thomas Mosk, and Antoine Uettwiller. Who pays for payment fraud? optimal detection and liability rules. Technical report, Working Paper, 2025. Available at SSRN.

Bengt Holmström. Moral hazard in teams. *Bell Journal of Economics*, 13(2):324–340, 1982.

Charles M. Kahn and William Roberds. Payments settlement: Tiering in private and public systems. *Journal of Money, Credit and Banking*, 41(5):855–884, 2009.

Tyler Moore. The economics of cybersecurity: Principles and policy options. *International Journal of Critical Infrastructure Protection*, 3(3–4):103–117, 2010.

Payment Systems Regulator. App fraud reimbursement protections. `https://www.psr.org.uk/information-for-consumers/app-fraud-reimbursement-protections/`, 2025a. Accessed July 15, 2026.

Payment Systems Regulator. App scams reimbursement requirement: Consolidated policy statement. Technical Report PS25/5, Payment Systems Regulator, May 2025b. URL `https://www.psr.org.uk/media/rhelv4op/ps25-5-app-scams-reimbursement-consolidated-policy-statement-may-2025.pdf`.

Payments Canada. Rtr fraud management — companion document (public consultation), 2025a. URL `https://cdn-static.bizzabo.com/bizzabo.file.upload/dnhfVVEbS1OKd0seB887_Webinar%204%20%20-%20Companion%20Document%20-%20RTR%20Fraud%20Management.pdf`.

Payments Canada. The real-time rail: An update on our approach to delivery and beyond. Payments Canada, Episode 36, June 2025b. Published June 25, 2025.

Appendix A Proofs

Proof of Lemma 1. For $(\delta, \theta) \in (0,1]^2$, define

$$G(m; \delta, \theta) = \phi(m) - e(\delta) - \psi(\theta) - \delta\theta m.$$

Conditional on (δ, θ), the planner admits every payment size for which $G(m; \delta, \theta) > 0$, excludes every size for which $G(m; \delta, \theta) < 0$, and is indifferent when $G(m; \delta, \theta) = 0$. Hence the optimized value conditional on efforts is

$$V(\delta, \theta) = \bar{u} + \frac{1}{M} \int_0^M \left[G(m; \delta, \theta)\right]_+ dm.$$

The function V is continuous on $(0,1]^2$: the efforts enter G additively and through the coefficient of m, so G is continuous in (δ, θ) uniformly in $m \in [0, M]$, and this continuity passes through $[\,\cdot\,]_+$ and the integral. Assumption 1 implies that there exists $(\hat{\delta}, \hat{\theta}) \in (0,1]^2$ such that

$$e(\hat{\delta}) + \psi(\hat{\theta}) + \hat{\delta}\hat{\theta}M < \phi(0).$$

Since ϕ is increasing, this implies

$$G(m; \hat{\delta}, \hat{\theta}) > 0 \qquad \text{for every } m \in [0, M],$$

and therefore the planner's value is strictly greater than $\bar{u}$. Moreover, if

$$e(\delta) + \psi(\theta) \geq \phi(M),$$

then $G(m; \delta, \theta) \leq 0$ for all m, so $V(\delta, \theta) = \bar{u}$. By the coercivity conditions in the model setup, effort profiles sufficiently close to either zero therefore cannot be optimal. The maximization of V can consequently be restricted to a compact subset of $(0,1]^2$, on which continuity implies existence of an optimum.

Consider any solution

$$(A^\star, \delta^\star, \theta^\star).$$

Because the planner's value exceeds $\bar{u}$, the admitted set has positive measure. Define

$$\bar{m} = \frac{1}{|A^\star|} \int_{A^\star} m \, dm > 0.$$

Holding $A^\star$ fixed, optimality requires

$$(\delta^\star, \theta^\star) \in \arg \min_{(\delta,\theta)\in(0,1]^2} \left\{e(\delta) + \psi(\theta) + \delta\theta\bar{m}\right\}.$$

Since the minimum is attained at $(\delta^\star, \theta^\star)$ and $\delta^\star\theta^\star\bar{m} > 0$,

$$\begin{aligned} e(\delta^\star) + \psi(\theta^\star) &= \min_{(\delta,\theta)\in(0,1]^2} \{e(\delta) + \psi(\theta) + \delta\theta\bar{m}\} - \delta^\star\theta^\star\bar{m} \\ &< \min_{(\delta,\theta)\in(0,1]^2} \{e(\delta) + \psi(\theta) + \delta\theta\bar{m}\}. \end{aligned}$$

Since $\bar{m} \leq M$ and $\delta\theta > 0$, the objective with $\bar{m}$ is pointwise no larger than the objective with M, so

$$\min_{(\delta,\theta)\in(0,1]^2} \{e(\delta) + \psi(\theta) + \delta\theta\bar{m}\} \leq \inf_{(\delta,\theta)\in(0,1]^2} \{e(\delta) + \psi(\theta) + \delta\theta M\} = K < \phi(0).$$

Thus

$$G(0; \delta^\star, \theta^\star) > 0.$$

For the optimal effort profile, G is continuous and strictly concave in m, and is positive at $m = 0$. Its nonnegative set is therefore a lower interval

$$[0, m_p] \qquad \text{for some } m_p \in (0, M].$$

If $m_p < M$, continuity gives

$$G(m_p; \delta^\star, \theta^\star) = 0,$$

or equivalently,

$$\phi(m_p) = e(\delta^\star) + \psi(\theta^\star) + \delta^\star\theta^\star m_p.$$

Strict concavity implies that this down-crossing is unique. Hence every optimal admitted set coincides with $[0, m_p]$ up to a null set.

Finally, substituting $A = [0, m_p]$ into welfare gives (1), so the value of the scalar problem over $(m_p, \delta, \theta) \in (0, M] \times (0, 1]^2$ is at most the planner's value. Conversely, every planner optimum coincides, up to a null set, with $([0, m_p], \delta^\star, \theta^\star)$ for some m_p, and null sets do not affect welfare, so the planner's value is attained by the scalar problem. The two problems therefore have the same value, and solutions correspond. □

Proof of Proposition 1. We maximize $m_p - m_0$ subject to $m_0 = g(F)$ and

$$F = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,(m_p + m_0).$$

Use the Lagrangian

$$\mathcal{L} = (m_p - m_0) + \nu\Big\{F - e(\delta) - \psi(\theta) - \tfrac{1}{2}\,\delta\,\theta\,(m_p + m_0)\Big\},$$

with $m_0 = g(F)$. Note that m_0 depends on F but not directly on (δ, θ, m_p).

Effort FOCs. Effort enters only through unit cost:

$$\partial_\delta \mathcal{L} = -\nu\Big[e'(\delta) + \tfrac{1}{2}\,\theta\,(m_p + m_0)\Big] = 0,$$

$$\partial_\theta \mathcal{L} = -\nu\Big[\psi'(\theta) + \tfrac{1}{2}\,\delta\,(m_p + m_0)\Big] = 0,$$

which yields

$$e'(\delta) + \tfrac{1}{2}\,\theta\,(m_p + m_0) = 0, \qquad \psi'(\theta) + \tfrac{1}{2}\,\delta\,(m_p + m_0) = 0.$$

Cap FOC.

$$\partial_{m_p} \mathcal{L} = 1 - \nu\,\tfrac{1}{2}\,\delta\,\theta = 0 \quad \Rightarrow \quad \nu = \frac{2}{\delta\,\theta}.$$

Fee / lower-edge slope. Using $m_0 = g(F)$ and $g'(F) = 1/\phi'(m_0)$, we obtain

$$\frac{\partial \mathcal{L}}{\partial F} = -\,g'(F) + \nu\Big[1 - \tfrac{1}{2}\,\delta\,\theta\,g'(F)\Big] = 0.$$

Substitute $\nu = 2/(\delta\theta)$ to get

$$-\,g'(F) + \frac{2}{\delta\theta}\Big[1 - \tfrac{1}{2}\,\delta\theta\,g'(F)\Big] = 0 \quad \Rightarrow \quad g'(F) = \frac{1}{\delta\theta}.$$

Since $g'(F) = 1/\phi'(m_0)$, this implies

$$\phi'(m_0) = \delta\,\theta.$$

Putting this back into the break-even condition gives

$$F = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,(m_p + m_0), \qquad m_0 = g(F) = g\Big(e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,(m_p + m_0)\Big),$$

which is the fixed point for m_0. □

Proof of Proposition 2. Impose zero minimum transaction size, so $m_0 = 0$ and the admitted region is $[0, m_p]$. Then

$$F = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,m_p.$$

Under break-even, payer surplus is

$$S(\delta, \theta, m_p) = \bar{u} + \frac{1}{M}\Big[\Phi(m_p) - F\,m_p\Big] = \bar{u} + \frac{1}{M}\Big[\Phi(m_p) - \big(e(\delta) + \psi(\theta)\big)m_p - \tfrac{1}{2}\,\delta\,\theta\,m_p^2\Big],$$

which coincides with the planner's objective on $[0, m_p]$ in (1).

Differentiating with respect to m_p, δ, θ yields the planner's FOCs:

$$\phi(m_p) - \big(e(\delta) + \psi(\theta)\big) - \delta\,\theta\,m_p = 0, \quad e'(\delta) + \tfrac{1}{2}\,\theta\,m_p = 0, \quad \psi'(\theta) + \tfrac{1}{2}\,\delta\,m_p = 0.$$

Hence payer-surplus maximization implements the planner under zero minimum transaction size. □

Proof of Proposition 3. Write $\pi = D\,[F - c]$ with

$$D(F; m_p) = \frac{(m_p - g(F))_+}{M}, \qquad c(\delta, \theta; F, m_p) = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\,\theta\,(m_p + m_0), \quad m_0 = g(F).$$

FOC in F.

$$\partial_F \pi = D\,[1 - \partial_F c] + [F - c]\,\partial_F D = 0.$$

Since $\partial_F c = \frac{1}{2}\,\delta\,\theta\,g'(F)$ and $\partial_F D = -(g'(F))/M$, we obtain, after rearranging,

$$\frac{F - c}{F} = \frac{1}{\varepsilon}, \qquad \varepsilon \equiv -\frac{F}{D}\frac{\partial D}{\partial F} = \frac{F\,g'(F)}{m_p - g(F)}.$$

Effort FOCs (holding F fixed). Only c depends on efforts:

$$\partial_\delta c = e'(\delta) + \tfrac{1}{2}\,\theta\,(m_p + m_0), \qquad \partial_\theta c = \psi'(\theta) + \tfrac{1}{2}\,\delta\,(m_p + m_0).$$

So at an interior optimum with $D > 0$,

$$e'(\delta) + \tfrac{1}{2}\,\theta\,(m_p + m_0) = 0, \qquad \psi'(\theta) + \tfrac{1}{2}\,\delta\,(m_p + m_0) = 0.$$

FOC in m_p.

$$\partial_{m_p}\pi = \underbrace{\partial_{m_p} D}_{=1/M}\,(F - c) + D\Big(- \partial_{m_p} c\Big) = 0.$$

Now $\partial_{m_p} c = \frac{1}{2}\,\delta\,\theta$, so with $D = (m_p - m_0)/M$,

$$\frac{1}{M}(F - c) - \frac{m_p - m_0}{M} \cdot \frac{1}{2}\,\delta\,\theta = 0 \quad \Rightarrow \quad F - c = \Big(\tfrac{1}{2}\,\delta\,\theta\Big)\,(m_p - m_0).$$

Thus all the claimed conditions follow. □

Proof of Proposition 4. Fix $(m_0, m_p) = (0, m_p^{\mathrm{FB}})$. Then the admitted set is fixed at $[0, m_p^{\mathrm{FB}}]$, and the mass of routed payments is fixed at

$$D = \frac{m_p^{\mathrm{FB}}}{M}.$$

For any effort profile (δ, θ), the per-payment social cost of serving this fixed set is

$$c(\delta, \theta; m_p^{\mathrm{FB}}) = e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\theta m_p^{\mathrm{FB}}.$$

Thus, once access and the cap are fixed, the only remaining efficiency-relevant choice is the choice of (δ, θ).

Under payer-surplus maximization subject to break-even, fee revenue only covers costs, so maximizing payer surplus is equivalent to minimizing $c(\delta, \theta; m_p^{\mathrm{FB}})$. Under profit maximization, demand is fixed by the mandates, so maximizing profit is also equivalent to minimizing

$c(\delta, \theta; m_p^{\mathrm{FB}})$, taking the access mandate as given. Under adoption maximization, adoption itself is already fixed by the mandates; hence the operator is indifferent across effort profiles that deliver the mandated admitted set. If the operator breaks ties in favor of the least-cost implementation, it again chooses the minimizer of $c(\delta, \theta; m_p^{\mathrm{FB}})$.

The resulting problem,

$$\min_{\delta,\theta} \left\{ e(\delta) + \psi(\theta) + \tfrac{1}{2}\,\delta\theta m_p^{\mathrm{FB}} \right\},$$

is exactly the planner's effort subproblem conditional on the admitted set $[0, m_p^{\mathrm{FB}}]$. Its solution is therefore

$$(\delta, \theta) = (\delta_{\mathrm{FB}}, \theta_{\mathrm{FB}}).$$

Thus, with zero minimum access and the planner's cap imposed, pricing affects only transfers and financing, not the implemented allocation. □

Proof of Proposition 5. For fixed (δ, θ), write

$$G(m; \delta, \theta) = \phi(m) - e(\delta) - \psi(\theta) - \delta\theta m.$$

For an admitted set A, aggregate break-even requires

$$\int_A F(m)\, dm \geq \int_A \left[e(\delta) + \psi(\theta) + \delta\theta m \right] dm.$$

Because payers participate only if $F(m) \leq \phi(m)$, the maximum revenue that can be raised from admitted users is

$$\int_A \phi(m)\, dm.$$

Hence an admitted set A is feasible under aggregate break-even if and only if

$$\int_A G(m; \delta, \theta)\, dm \geq 0.$$

Part (i): payer-surplus maximization. For any fixed admitted set A and efforts (δ, θ), payer surplus is

$$\int_A \left[\phi(m) - F(m) \right] dm.$$

Since the operator maximizes payer surplus, it chooses fees that satisfy aggregate break-even with equality whenever the admitted set is feasible. Therefore,

$$\int_A F(m)\, dm = \int_A \left[e(\delta) + \psi(\theta) + \delta\theta m \right] dm.$$

Substituting this into payer surplus gives

$$\int_A [\phi(m) - e(\delta) - \psi(\theta) - \delta\theta m]\, dm = \int_A G(m; \delta, \theta)\, dm.$$

Thus the payer-surplus operator chooses the admitted set and prevention efforts to maximize the same objective as the planner. Hence it implements the planner's allocation.

Part (ii): profit maximization with perfect price discrimination. Suppose the operator can perfectly price discriminate. For any admitted payment size m, the profit-maximizing fee is

$$F(m) = \phi(m).$$

The profit contribution from admitting payment size m is therefore

$$\phi(m) - e(\delta) - \psi(\theta) - \delta\theta m = G(m; \delta, \theta).$$

Hence total profit from admitted set A is

$$\int_A G(m; \delta, \theta)\, dm.$$

The profit-maximizing operator therefore chooses the same admitted set and prevention efforts as the planner. Thus perfect price discrimination implements the planner's allocation.

Part (iii): adoption maximization under aggregate break-even. The adoption-maximizing operator does not maximize

$$\int_A G(m; \delta, \theta)\, dm.$$

Instead, it maximizes the measure of A subject to the feasibility constraint

$$\int_A G(m; \delta, \theta)\, dm \geq 0.$$

This generally differs from the planner's problem.

To see the distinction, evaluate the adoption-maximizing operator at the planner's allocation $(A^{FB}, \delta^{FB}, \theta^{FB})$, where

$$A^{FB} = [0, m_p^{FB}].$$

At the first-best cap, the marginal admitted payment has zero social contribution:

$$G(m_p^{FB}; \delta^{FB}, \theta^{FB}) = 0.$$

However, inframarginal payments generally generate positive surplus, so

$$\int_0^{m_p^{FB}} G(m; \delta^{FB}, \theta^{FB})\, dm > 0.$$

By continuity of G, there exists $\varepsilon > 0$ such that

$$\int_0^{m_p^{FB}+\varepsilon} G(m; \delta^{FB}, \theta^{FB})\, dm \geq 0.$$

Therefore, the enlarged admitted set

$$[0, m_p^{FB} + \varepsilon]$$

is feasible under aggregate break-even and has strictly higher adoption than the planner's admitted set.

The additional payments above m_p^{FB} have negative marginal social contribution, but they can be financed by charging some inframarginal users above their payment-specific cost while still below their willingness to pay. Thus an adoption-maximizing operator can use inframarginal surplus to cross-subsidize inefficient marginal payments. Hence size-contingent fees do not generally make adoption maximization equivalent to welfare maximization. □

Proof of Proposition 6. With $m_0 = 0$ and $m_p = m_p^{\mathrm{FB}}$ fixed, the routed average size is

$$\bar{m} = \frac{m_p^{\mathrm{FB}}}{2},$$

and does not depend on pricing.

Scheme A (double liability). Charging both banks the full loss $\delta\, \theta\, m$ on each fraudulent payment makes each side internalize the full marginal loss. For the sender,

$$e'(\delta^*) + \theta^*\, \bar{m} = 0,$$

and for the receiver,

$$\psi'(\theta^*) + \delta^*\, \bar{m} = 0.$$

These are exactly the planner's FOCs at m_p^{FB}, so $(\delta^*, \theta^*) = (\delta_{\mathrm{FB}}, \theta_{\mathrm{FB}})$.

Scheme B (single-sided liability with one mandated effort).

(B1) With $\lambda = 1$ and $\theta = \theta_{\mathrm{FB}}$ mandated, the sender's FOC is

$$e'(\delta) + \theta_{\mathrm{FB}}\, \bar{m} = 0,$$

which yields $\delta = \delta_{\mathrm{FB}}$.

(B2) With $\lambda = 0$ and $\delta = \delta_{\mathrm{FB}}$ mandated, the receiver's FOC is

$$\psi'(\theta) + \delta_{\mathrm{FB}}\, \bar{m} = 0,$$

so $\theta = \theta_{\mathrm{FB}}$.

In both schemes, $m_0 = 0$ and $m_p = m_p^{\mathrm{FB}}$ ensure the admitted set is $[0, m_p^{\mathrm{FB}}]$; pricing only redistributes surplus and does not affect efforts. □

Proof of Proposition 7. Let

$$\mu = \frac{m_p}{2}.$$

For a given liability share λ, the decentralized effort choices solve

$$e'(\delta) + \lambda\theta\mu = 0, \tag{S$_0$}$$

$$\psi'(\theta) + (1-\lambda)\delta\mu = 0. \tag{R$_0$}$$

With the admitted region fixed, the regulator's problem is equivalent to minimizing

$$C(\lambda) = e(\delta(\lambda)) + \psi(\theta(\lambda)) + \delta(\lambda)\theta(\lambda)\mu.$$

Differentiate (S$_0$)–(R$_0$) with respect to λ. This gives

$$e''(\delta)\delta_\lambda + \lambda\mu\theta_\lambda = -\mu\theta,$$

and

$$(1-\lambda)\mu\delta_\lambda + \psi''(\theta)\theta_\lambda = \mu\delta.$$

Let

$$D(\lambda) = e''(\delta)\psi''(\theta) - \lambda(1-\lambda)\mu^2.$$

By the maintained regularity conditions (footnote 13),

$$D(\lambda) = e''(\delta)\psi''(\theta) - \lambda(1-\lambda)\mu^2 > 0,$$

and Cramer's rule gives

$$\delta_\lambda = -\frac{\mu\left[\psi''(\theta)\theta + \lambda\mu\delta\right]}{D(\lambda)} < 0,$$

and

$$\theta_\lambda = \frac{\mu\left[e''(\delta)\delta + (1-\lambda)\mu\theta\right]}{D(\lambda)} > 0.$$

Now differentiate the social cost:

$$C_\lambda = e'(\delta)\delta_\lambda + \psi'(\theta)\theta_\lambda + \mu\left(\theta\delta_\lambda + \delta\theta_\lambda\right).$$

Using (S$_0$)–(R$_0$),

$$e'(\delta) = -\lambda\theta\mu, \qquad \psi'(\theta) = -(1-\lambda)\delta\mu.$$

Therefore,

$$\begin{aligned} C_\lambda &= -\lambda\theta\mu\delta_\lambda - (1-\lambda)\delta\mu\theta_\lambda + \mu\left(\theta\delta_\lambda + \delta\theta_\lambda\right) \\ &= \mu\left[(1-\lambda)\theta\delta_\lambda + \lambda\delta\theta_\lambda\right]. \end{aligned}$$

Substituting the expressions for δ_λ and θ_λ, we obtain

$$C_\lambda = \frac{\mu^2}{D(\lambda)} \left[\lambda e''(\delta)\delta^2 - (1-\lambda)\psi''(\theta)\theta^2 \right].$$

At an interior optimum $\lambda^\star \in (0,1)$, the first-order condition is

$$C_\lambda(\lambda^\star) = 0.$$

Since $\mu > 0$ and $D(\lambda^\star) > 0$, this is equivalent to

$$\lambda^\star e''(\delta^\star)(\delta^\star)^2 = (1-\lambda^\star)\psi''(\theta^\star)(\theta^\star)^2.$$

Rearranging gives

$$\frac{\lambda^\star}{1-\lambda^\star} = \frac{\psi''(\theta^\star)(\theta^\star)^2}{e''(\delta^\star)(\delta^\star)^2},$$

and hence

$$\lambda^\star = \frac{\psi''(\theta^\star)(\theta^\star)^2}{\psi''(\theta^\star)(\theta^\star)^2 + e''(\delta^\star)(\delta^\star)^2}.$$

This proves the proposition. □

Proof of Corollary 1. By Proposition 7,

$$1-\lambda^\star = \frac{e''(\delta^\star)(\delta^\star)^2}{\psi''(\theta^\star)(\theta^\star)^2 + e''(\delta^\star)(\delta^\star)^2}.$$

Thus, as

$$\frac{e''(\delta^\star)(\delta^\star)^2}{\psi''(\theta^\star)(\theta^\star)^2} \to 0,$$

we have

$$1-\lambda^\star \to 0, \qquad \text{so} \qquad \lambda^\star \to 1.$$

□