Payment Fraud and Welfare*

Qi Li†

March 24, 2026

Click here for the latest version of the paper.

Abstract

Payment fraud is a central challenge for modern payment systems. This paper studies the welfare consequences of imperfect payment recognizability in a monetary search model in which sellers observe noisy signals about payment authenticity and decide whether to accept or reject transactions at posted prices. Imperfect fraud detection creates a trade-off: rejecting suspicious payments reduces fraud losses but may also lead to the loss of legitimate sales. I calibrate the model to match e-commerce payment fraud statistics, the empirical distribution of seller profit margins, and the money demand curve. The calibrated economy features a high but imperfect level of payment recognizability consistent with modern fraud detection technologies. Quantitatively, perfect recognizability raises online trading volume by about 10%, increases total profits of online sellers by roughly 27%, and generates a welfare gain equivalent to approximately 0.15% of GDP. Because the benefits of improved fraud detection accrue not only to sellers but also to buyers and issuers of the payment instrument, a private payment service provider that can charge only merchants may underprovide payment recognizability. The model also implies that payment recognizability is an important determinant of payment choice: introducing a competing digital payment instrument with a 10% improvement in recognizability leads to substantial adoption, capturing 5–16% of online payments even when its yield is 25–100 basis points below that of bank deposits.

Keywords: Money, Payment Fraud, Recognizability, Counterfeiting, Competitive Search, E-Commerce

JEL codes: E40, E41, D82, D83, D47

---

*The author did not receive any external funding for this research. I thank Guillaume Rocheteau for extensive feedback and guidance at various stages of this project. I also thank participants at the West Coast Search and Matching Workshop, the Renmin University of China Mini Workshop, the Fudan University seminar, and the Shanghai University of Finance and Economics seminar for helpful comments and discussions. All remaining errors are my own.

†Affiliation: Pennsylvania State University. Email: qili@psu.edu.

1 Introduction

Payment fraud is a central challenge for modern payment systems. Online merchants in particular face the risk that payments may be fraudulent, exposing them to losses through chargebacks. To mitigate this risk, merchants and payment service providers evaluate transactions in real time and reject those deemed high risk. Despite these efforts, fraud detection remains imperfect: fraudulent payments are sometimes accepted, while legitimate payments are occasionally declined. According to the *2025 Global eCommerce Payments & Fraud Report* published by the Merchant Risk Council, Visa Acceptance Solutions, and Verifi, approximately 3.4% of accepted e-commerce payments in North America in 2025 were fraudulent.[1] Moreover, false rejections of legitimate payments generate substantial costs for sellers through lost sales(Wedge, Kanter, Veeramachaneni, Rubio, and Perez, 2018; Wang, Chen, and Chen, 2019). Despite the payments industry's focus on fraud prevention, research on the economics of payment fraud remains limited. I develop a framework to study how imperfect payment recognizability affects equilibrium fraud, decentralized trade, and welfare.

The payments industry deploys sophisticated machine-learning systems that generate fraud risk signals for payments.[2] Merchants use these signals to determine whether to accept a payment. Because fraud detection is imperfect, acceptance decisions involve a trade-off between rejecting fraudulent payments and incorrectly declining legitimate ones. Importantly, these signals rarely affect transaction prices. In retail trade, prices are typically posted in advance, and merchants use fraud signals only to decide whether to accept payments at those prices. Industry guidance further suggests that acceptance decisions depend on sellers' profit margins: high-margin businesses can afford to approve more transactions because each legitimate sale is especially valuable.[3] These institutional features suggest that fraud signals affect trade primarily through accept–reject decisions rather than price adjustments.

Motivated by these institutional features, I develop a model of monetary exchange to study the economic consequences of imperfect payment recognizability in which acceptable means of payment are essential for trade in decentralized markets (Lagos and Wright, 2005; Rocheteau and Wright, 2005). Sellers post prices in advance and observe a noisy signal about payment authenticity when they encounter a buyer. Rather than adjusting prices based on the signal, sellers decide whether to accept the payment at the posted price. To capture the performance

[1] See https://info.merchantriskcouncil.org/hubfs/Documents/Reports/Fraud%20Reports/2025_Global_Fraud_and_Payments_Report.pdf.

[2] See https://www.mastercard.com/gateway/payment-solutions/secure-payments/fraud-protection.html and https://stripe.com/en-si/guides/primer-on-machine-learning-for-fraud-protection.

[3] See https://docs.stripe.com/radar/risk-settings?locale=en-GB and https://stripe.com/guides/state-of-online-fraud.

of modern fraud detection technologies, the model adopts a parsimonious information structure motivated by the Receiver Operating Characteristic (ROC) curve, a standard tool used to evaluate classification accuracy.

Imperfect payment recognizability affects decentralized trade through two channels. First, sellers may reject payments due to fraud concerns, preventing transactions with positive gains from trade. Second, imperfect detection allows some fraudulent payments to be accepted, exposing sellers to fraud losses. Although these losses themselves redistribute surplus rather than generating a direct welfare cost, they reduce seller profitability and discourage participation in the market. Both channels reduce trade in decentralized markets and therefore lower welfare.

To quantify the welfare consequences of imperfect payment recognizability, I calibrate the model to match e-commerce fraud statistics, the empirical distribution of seller profit margins, and the money demand curve. The calibrated economy features a high but imperfect level of recognizability, consistent with performance metrics reported in the machine-learning literature on fraud detection. I then compare the baseline economy with a counterfactual in which genuine and fraudulent payments can be perfectly distinguished. Perfect recognizability raises online trade by about 10% and increases total profits of online sellers by roughly 27%. Because decentralized trade accounts for only a small share of aggregate output, the impact on GDP is modest. Nevertheless, gains from trade in the decentralized market increase by about 2%, corresponding to a welfare gain of roughly 0.15% of GDP.

These results highlight the economic importance of payment recognizability in modern payment systems. Imperfect fraud detection distorts decentralized trade by reducing participation among online sellers and lowering the probability that trades are accepted. Improvements in fraud detection, therefore, generate sizable gains for sellers, particularly those with high profit margins and thus more exposed to fraud. At the same time, the benefits of improved payment recognizability extend beyond sellers. Buyers benefit from greater seller competition, while issuers of the payment instrument benefit from stronger demand for legitimate payment balances. Because private payment service providers typically charge only merchants for fraud detection services, they cannot fully internalize these broader benefits. As a result, private incentives to improve fraud detection accuracy may fall short of the socially optimal level, leading to underprovision of payment recognizability.

To assess the implications of payment innovations with superior fraud detection, I conduct a counterfactual exercise in which a competing payment asset (such as a stablecoin or CBDC) offers a 10% improvement in recognizability relative to bank deposits. Online sellers choose which instrument to request based on expected trading profits, generating competition between deposits and the new asset. Focusing on partial equilibrium, I show that improved fraud detection alone can induce meaningful substitution away from deposits: even when the

competing asset offers a lower yield, it captures a positive share of online transactions. At a yield disadvantage of 100 basis points, the new instrument attains about a 5% share of online payment volume; when the spread narrows to 25 basis points, its share rises to roughly 16%. These results highlight that payment recognizability, not just yield, is an important determinant of deposit demand for online transactions.

Related Work This paper contributes to the search-theoretic theory of money and credit, surveyed in Gu, Mattesini, and Wright (2016). The existing literature typically assumes that payments are either perfectly recognizable (Shi, 1995; Trejos and Wright, 1995; Lagos and Wright, 2005; Ferraris and Watanabe, 2008; Martin, Orlando, and Skeie, 2008; Telyukova and Wright, 2008; Bethune, Rocheteau, Wong, and Zhang, 2022; Chiu, Davoodalhosseini, Jiang, and Zhu, 2023) or, when focusing on counterfeit money or fraudulent assets, completely unrecognizable (Cavalcanti and Nosal, 2011; Li and Rocheteau, 2011; Li, Rocheteau, and Weill, 2012; Lester, Postlewaite, and Wright, 2012; Hu, 2013; Kang, 2017). A smaller set of papers, including Velde, Weber, and Wright (1999); Nosal and Wallace (2007); Quercioli and Smith (2015), allows for noisy signals of payment authenticity. However, their information structures generate piecewise-linear ROC curves, whereas modern machine-learning-based fraud detection systems exhibit smooth trade-offs between fraud detection and false rejections. This paper bridges this gap by embedding an empirically grounded information structure, consistent with ROC-based evaluation of modern fraud detection technologies, into a competitive search model of monetary exchange. This framework allows me to study how payment recognizability affects equilibrium fraud, decentralized trade, and welfare. The competitive search structure also generates predictions linking sellers' acceptance behavior to profit margins, consistent with industry practice, and enables the key recognizability parameter to be disciplined by observed fraud rates and seller margins.

The paper is also related to the broader literature on the economics of payments, surveyed in Nosal and Rocheteau (2006); Kahn and Roberds (2009); Koulayev, Rysman, Schuh, and Stavins (2016). Similar to the problem of counterfeiting or fraudulent assets in store-of-value payment systems, account-based payment systems are subject to identity theft and rely on service providers to verify identities when processing payments (Baxter, 1983). These providers exhibit significant economies of scale and network externalities, which raise concerns about market power and high fees in payment networks (Rochet and Tirole, 2002, 2003; Shy and Wang, 2011). The continued use of cash may partly reflect a desire to avoid such fees as well as concerns about identity theft and fraud in electronic payments (Alvarez and Argente, 2022; Shy, 2023). This paper is closely related to work on fraud and identity theft in account-based payment systems, such as Kahn and Roberds (2008). In that framework, genuine buyers are

always accepted, so the false positive rate is zero, and higher monitoring intensity reduces the false negative rate, the probability of incorrectly accepting fraudulent buyers, at the cost of higher per-transaction monitoring expenses. In contrast, this paper focuses on improvements in payment recognizability arising from technological advances, such as better data collection and more advanced machine learning algorithms, that do not increase per-transaction monitoring costs. Moreover, by incorporating an information structure that allows for both false positives and false negatives in payment fraud detection, and modeling merchants' trade-off between fraud losses and lost legitimate sales, the paper provides a tractable framework for analyzing payment fraud and links the payments literature to machine learning approaches to fraud detection in account-based systems.

Outline The paper proceeds as follows. Section 2 introduces the basic environment and defines equilibrium. Section 3 characterizes the partial equilibrium of the model, taking buyers' market utility as given. Section 4 determines buyers' market utility as the solution to a fixed-point problem. Section 5 extends the baseline model to allow for multiple means of payment and heterogeneous sellers. Section 6 calibrates the model to the data and conducts two counterfactual exercises to evaluate the welfare gain of perfect payment recognizability and the potential effect of introducing a competing payment instrument with superior recognizability. Section 7 concludes.

2 Model

2.1 Environment

Agents, preferences, and the payment asset Time is discrete and indexed by $t \in \mathbb{N}_0$. Each period consists of two sequential stages: a decentralized market (DM), followed by a centralized market (CM). In each stage, a consumption good is produced using labor as the sole input.

The economy is populated by a unit measure of sellers and a measure N of buyers. The labels "buyer" and "seller" refer to agents' roles in the DM. Buyers freely choose whether to participate in the DM and, conditional on participation, whether to behave as genuine or fraudulent buyers. This choice gives rise to three possible buyer types: genuine, fraudulent, and inactive, whose interpretations are clarified later.

Agents' utility functions are additively separable across stages within a period and across periods. There is no discounting between stages within the same period. The discount factor between periods is $\beta \in (0, 1)$. In the DM, buyers have unit demand for the DM good, which yields utility $u > 0$. Sellers produce the good at disutility cost c. In the baseline model, this

production cost is common across sellers. In Section 5 and Section 6, I relax this assumption by allowing sellers to be heterogeneous in their production costs.

In the CM of period t, the utility of each agent takes the form

$$U(x_t) - h_t,$$

where x_t is consumption and h_t is hours of work.[4] Production is linear: one unit of the CM good requires one unit of labor. In what follows, we use the CM good (or, equivalently, labor) as the numeraire. The utility function $U(x)$ is strictly increasing and strictly concave. Therefore, there exists a unique $x^* > 0$ satisfying

$$U'(x^*) = 1,$$

which characterizes the efficient level of CM consumption.

In the model, credit is infeasible in the DM because sellers cannot enforce repayment, and there is no technology to record past behavior.[5] However, an asset exists that can serve as a means of payment. Payments may be final, or they may occur under the mutual understanding that buyers will repurchase the asset in the CM. In the latter case, the asset effectively serves as collateral to guarantee repayment.

I assume that the payment asset is supplied elastically by an issuer to maintain a constant gross real return R, with $R \leq 1/\beta$. The issuer shares the same discount factor as private agents and values profits earned in the CM, including profits from offering payment assets and the associated seigniorage revenue. When the payment asset is money, including paper currency and zero-yield transactional deposits as in Section 6, it is issued by a consolidated banking sector comprising both the central bank and commercial banks. In this case, the real return equals the reciprocal of the gross inflation rate, so maintaining a constant R corresponds to a constant inflation rate.

I define the liquidity spread (or convenience yield) of the payment asset as

$$s \equiv \frac{1}{\beta R} - 1.$$

The spread $s \geq 0$ captures the proportional holding cost of liquidity. Relative to an illiquid bond that exactly compensates buyers for their time preference, acquiring one unit of purchasing power in the decentralized market (DM) using the payment asset entails an additional cost of

[4] Although one could restrict h_t to lie in $[0, \bar{h}]$, I assume this constraint is slack in equilibrium.

[5] It is possible to relax this assumption by introducing banks that can enforce repayment and by allowing some sellers to accept bank credit.

s.[6] This holding cost borne by buyers corresponds to seigniorage revenue for the issuer. In the case of money, this includes profits accruing to both the central bank and commercial banks.

Trading mechanism The trading mechanism differs across the two stages. In the centralized market (CM), sellers post contracts that open submarkets in the subsequent decentralized market (DM). Agents then trade the CM good and the asset in a Walrasian market.[7]

In the DM, buyers choose which submarket to enter. Each submarket is defined by a seller-posted contract ω, consisting of a price p, a recommended buyer–seller ratio n, and a recommended share of fraudulent buyers γ. I assume that the price is posted before the fraud signal is realized and therefore cannot depend on it. This captures a common feature of retail transactions: prices are fixed in advance, while fraud detection systems generate risk signals that guide accept-or-reject decisions after a payment attempt. The recommended market tightness and buyer composition are interpreted as seller recommendations and must be incentive-compatible for buyers to follow.

The measure of matches in a submarket with a buyer-seller ratio n is determined by a matching technology $\alpha(n)$. I assume a Leontief matching technology,

$$\alpha(n) = \sigma \min\{1, n\},$$

where $\sigma \in (0, 1]$ governs the intensity of meetings.

I distinguish between genuine (g), fraudulent (f), and inactive (0) buyers. Genuine buyers pay for DM transactions using legitimate payment assets, whereas fraudulent buyers commit fraud, either by using worthless fraudulent assets or by submitting payment instructions that are later reversed through chargebacks after receiving the DM good.[8] Inactive buyers do not participate in the DM. In each period, a type-$\varkappa$ buyer incurs a utility cost $k^{\varkappa}$. I assume $k^0 = k^g = 0$, i.e., both staying out of the DM and participating as a genuine buyer are costless. I assume $k^f \geq 0$, capturing a fixed cost of committing payment fraud (e.g., producing counterfeit currency, or stealing account information).

Once matched in the DM, a seller observes a noisy signal $z \in [0, 1]$ about the buyer's type, genuine (g) or fraudulent (f), and decides whether to accept the transaction. Sellers cannot commit to arbitrary acceptance rules in advance. Instead, any acceptance rule must be ex-post

[6] To obtain one unit of real purchasing power in the DM of period t, a buyer must acquire the payment asset in the centralized market (CM) of period $t-1$ by supplying $1/R$ units of labor, which has a present value of $1/(\beta R)$. By contrast, holding an illiquid bond yields one unit of purchasing power in the CM of period t, implying a cost of one. Hence, the additional intertemporal cost of using the payment asset is $1/(\beta R) - 1 = s$.

[7] Buyers observe posted contracts before choosing their asset holdings.

[8] If fraudulent assets are produced, they are detected and confiscated in the subsequent CM.

optimal given the seller's posterior belief about the buyer's type and the posted price. As in Lagos and Wright (2005), the DM captures trading opportunities where acceptable liquidity is essential. If a payment is rejected due to fraud concerns, the match dissolves, and the trading opportunity is lost.

Information structure In practice, modern fraud detection systems generate noisy signals about the authenticity of a payment. The performance of such classification systems is commonly evaluated using the Receiver Operating Characteristic (ROC) curve, which describes the trade-off between the true positive rate (the probability that a fraudulent payment is rejected) and the false positive rate (the probability that a genuine payment is rejected) as the decision threshold varies. A widely used summary measure of classification performance is the area under the ROC curve (ROC-AUC), which equals the probability that the signal ranks a randomly drawn fraudulent payment above a randomly drawn genuine payment and ranges from 0.5 (random guessing) to 1 (perfect classification). State-of-the-art machine learning algorithms achieve ROC-AUC values between 0.95 and 0.99 (Van Vlasselaer, Bravo, Caelen, Eliassi-Rad, Akoglu, Snoeck, and Baesens, 2015; Hashemi, Mirtaheri, and Greco, 2022; Nguyen, Duong, Chau, Nguyen, Trinh, Tran, and Ho, 2022). Although payment service providers' ROC-AUC statistics are typically proprietary, industry reports suggest that firms often target levels above 0.90–0.95.[9]

Motivated by the ROC-based evaluation of fraud detection systems, I adopt a parsimonious information structure that generates a realistic ROC curve. For fraudulent buyers, the signal is uniformly distributed:

$$F^f(z) = z.$$

For genuine buyers, the signal distribution is

$$F^g(z) = z^{1-e},$$

where $e \in (0, 1)$ measures payment recognizability. A higher value of e shifts the distribution of signals for genuine payments toward lower values of z, making them easier to distinguish from fraudulent payments.

The limiting cases are instructive. When $e = 1$, genuine payments are perfectly recognizable: the distribution of signals for genuine buyers collapses to a mass point at $z = 0$. A seller who accepts only when $z = 0$ admits all genuine buyers while rejecting fraudulent buyers with probability one. At the other extreme, when $e = 0$, we have $F^g = F^f$, so genuine and fraudulent

[9]See https://moldstud.com/articles/p-advanced-stripe-integration-with-machine-learning-for-enhanced-fraud-protection and https://galileo.ai/blog/auc-score.

payments are statistically indistinguishable.

Since the signal satisfies the monotone likelihood ratio property, the seller's optimal strategy is a threshold rule: a transaction is rejected whenever the signal z is above a cutoff $\bar{z}$. Varying this cutoff traces out an ROC curve:

$$\text{True Positive Rate} = 1 - (1 - \text{False Positive Rate})^{\frac{1}{1-e}}.$$

Here the true positive rate corresponds to the probability that a fraudulent payment is rejected, while the false positive rate corresponds to the probability that a genuine payment is rejected.[10]

As shown in Figure 1, the implied ROC curve closely resembles those observed in empirical studies of payment fraud detection algorithms.[11] Under this information structure, the parameter e maps directly into the ROC-AUC via the closed-form relation

$$\text{ROC-AUC} = \frac{1}{2-e}.$$

The parameter e therefore summarizes the accuracy of the fraud detection technology.

2.2 Equilibrium Definition

I assume that only genuine buyers choose to carry the asset across periods in the amount required for payment. If $R < 1/\beta$ (equivalently, $s > 0$), they are the only agents willing to incur the associated holding cost to use the asset in DM transactions. If $R = 1/\beta$ (equivalently, $s = 0$), agents are indifferent between carrying the asset and not carrying it, so this assumption is without loss of generality.[12]

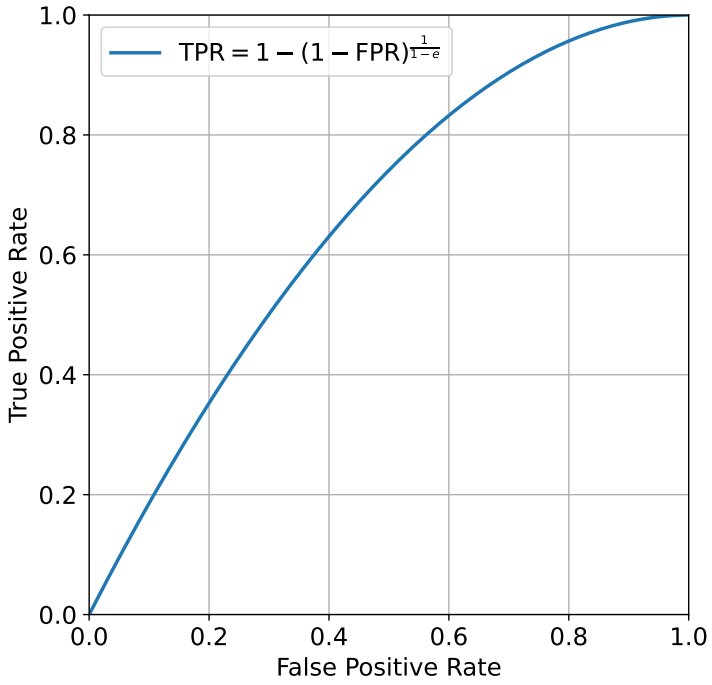
CM value functions Consider a seller entering the CM with a units of asset expressed in the CM good. As noted above, sellers do not carry assets across periods. The problem of a seller consists of posting a contract, $\omega = (p, n, \gamma)$, for the following DM composed of a price, $p \in \mathbb{R}^+$, a buyer-seller ratio $n \in \mathbb{R}^+$, and the share of fraudsters among buyers, $\gamma \in [0, 1]$. Let Ω denote the set of contracts that satisfy buyers' participation conditions. The value of a type-c

[10] Fraudulent payments are defined as the positive class to align with industry practice. Reversing the positive and negative classes reflects the ROC curve across the anti-diagonal ($x + y = 1$) but leaves the area under the curve unchanged.

[11] See Van Vlasselaer, Bravo, Caelen, Eliassi-Rad, Akoglu, Snoeck, and Baesens (2015); Hashemi, Mirtaheri, and Greco (2022); Nguyen, Duong, Chau, Nguyen, Trinh, Tran, and Ho (2022).

[12] See Rocheteau and Wright (2005) for a formal argument.

Figure 1: ROC Curve



Notes: The solid curve depicts the Receiver Operating Characteristic curve implied by the model information structure.

seller in the CM is given by

$$W(a) = \max_{x,h,\omega} \{U(x) - h + \beta V(\omega)\} \quad \text{s.t.} \quad x = a + h \text{ and } \omega \in \Omega.$$

The seller chooses his consumption, x, hours of work, h, and DM contract, ω, to maximize his period utility plus DM continuation value, $V(\omega)$, with discounting. Consumption is financed with the initial wealth and labor income. I substitute $h = x - a$ in the objective to obtain:

$$W(a) = a + U(x^*) - x^* + \beta \max_{\omega \in \Omega} V(\omega).$$

As is standard in New Monetarist models, this CM value function is linear in wealth.

I now turn to buyers. They choose their consumption, x, hours of work, h, a submarket ω among all open submarkets, O, and a participation decision, $\varkappa \in \{f, g, 0\}$. If $\varkappa = f$, the buyer participates in the DM as a fraudster; if $\varkappa = g$, she participates as a genuine buyer; and if $\varkappa = 0$, she remains inactive. As noted above, only genuine buyers carry assets into the DM, in the amount required for payment. Let $p(\omega)$ denote the price in submarket ω. The required real asset holdings brought into the DM are therefore

$$\mathbb{I}_{\{\varkappa = g\}} \frac{p(\omega)}{R},$$

where $\mathbb{I}_{\{\varkappa = g\}}$ is an indicator equal to 1 if the buyer is genuine and 0 otherwise.

The value of a buyer in the CM is given by:

$$W^b(a) = \max_{x,h,\varkappa,\omega} \{U(x) - h + \beta [V^{\varkappa}(\omega) - k^{\varkappa}]\} \qquad \text{s.t.} \quad x + \frac{\mathbb{I}_{\{\varkappa = g\}} p(\omega)}{R} = a + h.$$

The buyer chooses consumption x, labor h, a submarket ω, and a participation type $\varkappa$ to maximize current utility plus the discounted DM continuation value The continuation value is $V^{\varkappa}(\omega)$, the value of participating as type $\varkappa$ in submarket ω, net of the participation cost $k^{\varkappa}$. As in the case of sellers, the value function can be rewritten as

$$W^b(a) = a + U(x^*) - x^* + \max_{\varkappa,\omega} \left\{ -\frac{\mathbb{I}_{\{\varkappa = g\}} p(\omega)}{R} + \beta [V^{\varkappa}(\omega) - k^{\varkappa}] \right\}. \tag{1}$$

The CM value function of the buyer is also linear in wealth.

DM value functions Consider the value function of a seller in the DM who offers a contract $\omega = (p, n, \gamma)$. It solves

$$V(p,n,\gamma) \quad = \quad [1-\alpha(n)]W(0) + \alpha(n)\int_0^1 \max\left\{-c + [1-\hat{\gamma}(z)]\,W(p) + \hat{\gamma}(z)W(0), W(0)\right\} dF^{\gamma}(z),$$

where $F^{\gamma}(z) = \gamma F^{f}(z) + (1-\gamma)F^{g}(z)$ is the cumulative distribution of signals in that submarket, given the share of fraudulent buyers, and $\hat{\gamma}(z)$ is the probability of facing a fraudster, conditional of receiving the signal z. By Bayes's rule,

$$\frac{\hat{\gamma}(z)}{1-\hat{\gamma}(z)} = \frac{f^{f}(z)}{f^{g}(z)}\frac{\gamma}{1-\gamma} = \frac{1}{(1-a)z^{-e}}\frac{\gamma}{1-\gamma}. \tag{2}$$

With probability $1-\alpha(n)$, the seller does not meet a buyer and enters the next CM with zero asset. With probability $\alpha(n)$, he meets a buyer. Conditional on meeting, he observes a signal z drawn from $G(z)$ and decides whether to accept the trade. If he accepts, he incurs the disutility cost c and receives payment p only if the buyer is genuine, which occurs with posterior probability $1-\hat{\gamma}(z)$. If she rejects, she proceeds to the next CM with zero assets. Using the linearity of W, the Bellman equation can be rewritten as:

$$V(p,n,\gamma) = W(0) + \alpha(n)\int_0^1 \max\left\{-c + [1-\hat{\gamma}(z)]\,p, 0\right\} dF^{\gamma}(z).$$

Optimal acceptance A seller accepts the trade when

$$-c + \left[1-\hat{\gamma}(z)\right]p \geq 0. \tag{3}$$

When the condition holds with equality, the seller exactly breaks even. At this point, the seller's posterior belief that the buyer is fraudulent equals the seller's profit margin:

$$\hat{\gamma}(z) = 1 - \frac{c}{p}.$$

In the payment-fraud detection literature, the posterior probability at which the seller breaks even is known as the *break-even precision*.

Conditions (2) and (3) imply an acceptance threshold $\hat{z} \in [0,1]$ such that the seller accepts the trade if and only if $z \leq \hat{z}$. The threshold is given by

$$\hat{z}(p,\gamma) = \min\left\{\left[\frac{(1-\gamma)(1-e)}{\gamma}\left(\frac{p-c}{c}\right)\right]^{1/e}, \; 1\right\}.$$

The value of a seller can then be rewritten as

$$V(p,n,\gamma) = W(0) + \alpha(n)\left[(1-\gamma)F^g(\hat{z}(p,\gamma))(p-c) - \gamma F^f(\hat{z}(p,\gamma))c\right] \tag{4}$$

I now turn to buyers in the DM. The value of a buyer who does not participate is simply $V^0(\omega) = W^b(0)$. The value of a genuine buyer who participates in submarket $\omega = (p,n,\gamma)$ is

$$V^g(p,n,\gamma) = W^b(p) + \frac{\alpha(n)}{n} F^g(\hat{z}(p,\gamma))\left[u + W^b(0) - W^b(p)\right].$$

With probability $\alpha(n)/n$, the buyer meets a seller, and with probability $F^g(\hat{z})$, the seller observes a signal below $\hat{z}$ and accepts the trade. In that event, the buyer obtains utility u, pays p, and enters the next CM with zero asset rather than with assets worth p. Using the linearity of W^b I can rewrite the Bellman equation as:

$$V^g(p,n,\gamma) = W^b(p) + \frac{\alpha(n)}{n} F^g(\hat{z}(p,\gamma))\,(u-p)\,. \tag{5}$$

By a similar reasoning, the value of a fraudster in the DM is

$$V^f(p,n,\gamma) = W^b(0) + \frac{\alpha(n)}{n} F^f(\hat{z}(p,\gamma))u. \tag{6}$$

The key difference is that a fraudster does not make any genuine payment.

Buyer participation and eligible contracts Consider the decision to enter submarket $\omega = (p,n,\gamma)$ as a genuine buyer. Optimal participation requires

$$-\frac{p}{R} + \beta\left[V^g(\omega) - k^g\right] \;\leq\; \max_{\varkappa,\omega'}\left\{-\frac{\mathbb{I}_{\{\varkappa=g\}}\,p(\omega')}{R} + \beta\left[V^{\varkappa}(\omega') - k^{\varkappa}\right]\right\}, \text{ with equality if } n(1-\gamma) > 0.$$

Substituting for $V^g(\omega)$ from (5), this condition can be written as

$$-sp + \frac{\alpha(n)}{n}\,F^g\big(\hat{z}(p,\gamma)\big)\,(u-p) - k^g \;\leq\; \Delta, \quad \text{with equality if } n(1-\gamma) > 0, \tag{7}$$

where

$$\Delta \equiv \max_{\varkappa,\omega}\left\{-\frac{\mathbb{I}_{\{\varkappa=g\}}\,p(\omega)}{\beta R} + \left[V^{\varkappa}(\omega) - k^{\varkappa}\right]\right\} - W^b(0)$$

denotes the buyers' market utility from DM participation. We have $\Delta \geq 0$ because the maximization problem always includes the outside option of non-participation, $\varkappa = 0$, which yields exactly $W^b(0)$.

Similarly, the participation condition for a fraudulent buyer in submarket $\omega = (p, n, \gamma)$ is

$$\frac{\alpha(n)}{n} F^f\big(\hat{z}(p,\gamma)\big)\, u \; - k^f \leq \; \Delta, \quad \text{with equality if } n\gamma > 0. \tag{8}$$

I refer to contracts that satisfy the participation constraints in (7) and (8) as *eligible contracts*.

Optimal contract posting Among the set of eligible contracts, a seller chooses a contract $\omega = (p, n, \gamma)$ to maximize its value $V(p, n, \gamma)$. Using (4), the seller's problem can be written compactly as

$$O \subseteq \arg\max_{p,n,\gamma} \; v(p,n,\gamma) \quad \text{s.t.} \quad (7) \text{ and } (8), \tag{9}$$

where

$$v(p,n,\gamma) \equiv V(p,n,\gamma) - W(0) = \alpha(n)\Big\{(1-\gamma)\, F^g\big(\hat{z}(p,\gamma)\big)(p-c) \; - \; \gamma\, F^f\big(\hat{z}(p,\gamma)\big)c\Big\}$$

is the seller's expected trading profit in the submarket.[13]

Adding-up condition The equilibrium satisfies an adding-up condition: the total number of active buyers across all open submarkets does not exceed the fixed supply of buyers. This constraint binds whenever buyers earn a strictly positive market utility. Accordingly,

$$\int_O n(\omega)\, d\mu(\omega) \; \leq \; N, \qquad \text{with equality if } \Delta > 0. \tag{10}$$

Definition 1. An *equilibrium* consists of:

- a set of open submarkets O that solves the optimal contract problem (9);
- a probability measure μ on O describing the distribution of sellers across optimal contracts;
- a buyer market utility $\Delta \geq 0$;

such that the collection (μ, Δ) satisfies the adding-up requirement (10).

[13] The value of (9) is nonnegative. A seller can always choose a null contract (any contract with $n = 0$) and earn zero expected profit. Therefore, any solution to the problem is weakly optimal for the seller.

3 Partial Equilibrium

I analyze the equilibrium of the model in two steps. In this section, I take buyers' market utility in the decentralized market, Δ, as given and characterize the optimal contracts for sellers. Depending on the parameter configuration, the optimal contract can be separating, pooling, or null. In particular, pooling contracts arise when fraud costs are low and sellers' production costs are intermediate. In the limiting case in which committing fraud is costless, no separating contract can sustain a strictly positive payment, so trade requires pooling. In this sense, partial recognizability combined with low fraud costs transforms fraud from a mere off-path threat into a robust equilibrium outcome.

For notational simplicity, I define the effective participation costs as $\hat{k}^g \equiv \Delta$ and $\hat{k}^f \equiv k^f + \Delta$. It is useful to introduce

$$\psi \equiv \frac{F^g(\hat{z})}{F^f(\hat{z})} = \hat{z}^{-e} \geq 1, \tag{11}$$

which increases strictly in the threshold $\hat{z}$. When $e \in (0,1)$, as $\hat{z}$ rises from 0 to 1, I have ψ increasing smoothly from 1 to $+\infty$.

The variable ψ captures the ratio of the *true positive rate* to the *false positive rate*. A higher ψ implies a more *selective* acceptance rule: acceptance probabilities fall for both types, but decline more sharply for fraudulent buyers. As a result, the share of genuine buyers among accepted matches increases with ψ. Thus, sellers face a fundamental trade-off between *quantity* and *quality* when choosing their acceptance rule.

As in Li and Shimer (2025), the cases $\psi = 1$ and $\psi > 1$ correspond to two distinct regimes of optimal contracts. Proposition 1 shows that every optimal contract is either separating with $\psi = 1$ or pooling with $\psi > 1$ (unless the seller's expected trading profit in the DM, $\bar{v}$, is zero). A pooling contract with $\psi = 1$ cannot be optimal because it unravels (Guerrieri, Shimer, and Wright, 2010). In particular, if there exists an active pooling contract under which sellers always accept, then sellers strictly prefer an otherwise identical separating contract. In such a deviation, sellers will still always accept, leaving genuine buyers indifferent while strictly increasing the sellers' expected profit by eliminating fraud. In what follows, I use the term *pooling contracts* to refer specifically to contracts with a non-degenerate buyer pool and $\psi > 1$.

Proposition 1. *Let $\bar{v}$ denote the seller's expected trading profit in the DM, i.e., the maximized value of (9). When $\bar{v} > 0$, an optimal contract takes one of two forms:*

(i) Separating contract: *(p, n, γ), where $\gamma = 0$ and (p, n) solves*

$$\begin{aligned} \bar{v}^S = \max_{p,n} \quad & \alpha(n)(p-c) \\ s.t. \quad & -sp + \frac{\alpha(n)}{n}(u-p) = \hat{k}^g, \\ & \frac{\alpha(n)}{n}u \leq \hat{k}^f; \end{aligned} \tag{12}$$

(ii) Pooling contract: *(p, n, γ), where $\gamma = \gamma(\psi, p)$, $p > c$, $\psi > 1$, and (p, n, ψ) solves*

$$\begin{aligned} \bar{v}^P = \max_{p,n,\psi} \quad & \alpha(n)\Big\{(1-\gamma(\psi,p))\psi^{-\frac{1-e}{e}}(p-c) \; - \; \gamma(\psi,p)\psi^{-\frac{1}{e}}c\Big\} \\ s.t. \quad & -sp + \frac{\alpha(n)}{n}\psi^{-\frac{1-e}{e}}(u-p) = \hat{k}^g, \\ & \frac{\alpha(n)}{n}\psi^{-\frac{1}{e}}u = \hat{k}^f. \end{aligned} \tag{13}$$

Finally, the seller's expected profit in the DM is $\bar{v} = \max\{\bar{v}^S, \bar{v}^P, 0\}$.[14]

The partial equilibrium analysis proceeds by solving for the optimal contract in each regime and then comparing seller profit under an *optimal separating contract* (maximizing profit among separating contracts) and an optimal pooling contract (maximizing profit among pooling contracts with $\psi > 1$).

3.1 Optimal Separating Contract

In this subsection, I focus on the optimal separating problem (12). I begin by considering an environment in which buyers cannot enter as fraudsters and must pay with genuine payment assets. In this case, sellers solve (12) without the inequality constraint. Lemma 1 characterizes the first-best contract in this environment.

Lemma 1. *The first-best contract is given by*

$$p^{FB} = \frac{\sigma u - \hat{k}^g}{\sigma + s}, \qquad n^{FB} = 1,$$

provided that $c < p^{FB}$.

[14] If no triple (p, n, ψ) satisfies the constraints for the pooling problem, I define $\bar{v}^P = -\infty$.

In this case, the seller's expected trading profit equals

$$\bar{v}^{FB} = \sigma\big(p^{FB} - c\big).$$

If $c \geq p^{FB}$, the seller prefers not to trade (i.e., posts a null contract).

Proposition 2 characterizes the optimal separating contract. When the fraud cost is high, the participation constraint of fraudulent buyers does not bind. In this case, the optimal separating contract coincides with the first-best contract. When the fraud cost is low, the participation constraint of fraudulent buyers binds. The optimal contract then distorts the buyer-seller ratio upward, thereby reducing buyers' meeting probabilities so that fraud does not yield a strict gain. Relative to the first-best allocation, the price is lower to compensate genuine buyers for the lower trading probability. In both cases, sellers earn strictly positive expected profit if and only if $c < p^S$; otherwise, they prefer the null contract to $(p^S, n^S, 0)$.

Proposition 2. *Assume $\bar{v} > 0$. The separating problem (12) admits a unique optimal separating contract $(p^S, n^S, 0)$.*

(i) High fraud cost. *If $\hat{k}^f \geq \sigma u$, the participation constraint of fraudulent buyers does not bind, and the optimal separating contract coincides with the first-best contract:*

$$p^S = p^{FB}, \qquad n^S = n^{FB}.$$

(ii) Low fraud cost. *If $\Delta \leq \hat{k}^f < \sigma u$, the optimal separating contract sets*

$$p^S = \frac{\hat{k}^f - \hat{k}^g}{\hat{k}^f/u + s}, \qquad n^S = \frac{\sigma u}{\hat{k}^f}.$$

In either case, sellers' expected trading profit is

$$\bar{v}^S = \sigma\,(p^S - c),$$

which is strictly positive if $c < p^S$ and non-positive otherwise (in which case the seller prefers a null contract).

Special case: costless fraud In many environments, the cost of committing payment fraud can be negligible.[15] I therefore consider the case $k^f = 0$, in which committing payment fraud

[15]For example, with the shift from in-person to online payments, fraudsters no longer need to produce physical counterfeit bills, checks, or cards. Instead, they can submit stolen credentials electronically, substantially lowering the cost of fraud.

is costless.

Corollary 1. *If $k^f = 0$ and $\bar{v} > 0$, then $p^S \leq 0$ and hence $\bar{v}^S \leq 0$. Therefore, the equilibrium cannot be separating and must be pooling.*

The intuition for Corollary 1 is straightforward. When fraud is costless, no separating contract can sustain a strictly positive price: buyers strictly prefer to commit fraud in order to avoid both the holding cost, sp, and the expected payment in genuine assets, $\frac{\alpha(n)}{n}p$. Without a strictly positive price, sellers cannot earn strictly positive trading profits. Absent the possibility of pooling contracts, as in Lester, Postlewaite, and Wright (2012), the asset would then fail to serve as a means of payment and the equilibrium would collapse to autarky. In the present model, however, only separating contracts are eliminated; a pooling equilibrium with trade remains feasible.

3.2 Optimal Pooling Contract

In this subsection, I focus on the optimal pooling problem (13). Lemma 2 highlights that since the participation constraints tie price p and buyer-seller ratio n to the selectiveness level ψ, any eligible pooling contract can be fully indexed by ψ.

Lemma 2. *Each eligible pooling contract can be indexed by the selectiveness level ψ via*

$$p(\psi) = \frac{\hat{k}^f \psi - \hat{k}^g}{\hat{k}^f \psi / u + s}, \tag{14}$$

$$n(\psi) = \frac{\sigma u}{\hat{k}^f} \psi^{-\frac{1}{e}}, \tag{15}$$

$$\gamma(\psi) = \frac{(1-e)\psi\left[1 - c/p(\psi)\right]}{c/p(\psi) + (1-e)\psi\left[1 - c/p(\psi)\right]}.$$

Given Lemma 2, the pooling problem (13) reduces to a single-variable maximization over ψ, as shown in Lemma 3.

Lemma 3. *Let*

$$\psi_{\min} \equiv \frac{\hat{k}^g + cs}{\hat{k}^f\left(1 - \frac{c}{u}\right)}, \qquad \psi_{\max} \equiv \left(\frac{\sigma u}{\hat{k}^f}\right)^e.$$

Then the pooling problem (13) reduces to the one–dimensional program

$$\bar{v}^P = \max_{\psi} \ \tilde{v}(\psi) \quad s.t. \quad \max\{1, \psi_{\min}\} \leq \psi \leq \psi_{\max}, \tag{16}$$

where

$$\tilde{v}(\psi) \equiv v(p(\psi), n(\psi), \gamma(\psi)) = \frac{e}{1-e}\, \sigma\, \gamma(\psi)\ \psi^{-\frac{1}{e}}\, c. \tag{17}$$

Under the information structure, the recognizability parameter e governs the degree of *cross-subsidization* in pooling contracts. The expression for $\tilde{v}(\psi)$, before multiplying by $e/(1-e)$, captures the seller's expected loss from trading with fraudulent buyers. The seller's overall profit is then scaled by the factor $e/(1-e)$ because the expected gain from genuine buyers is exactly $1/(1-e)$ times the expected loss from fraudulent buyers. Therefore, the net profit is

$$\left(\frac{1}{1-e} - 1\right) = \frac{e}{1-e}$$

times the expected loss.

Proposition 3 establishes the conditions under which an optimal pooling contract exists and characterizes its structure. Feasibility requires that the lower bound on ψ, which ensures the production cost is weakly below the price, does not exceed the upper bound on ψ that is consistent with delivering the required market utility to buyers. When this condition is satisfied, the feasible set is non-degenerate, and the seller's profit under the optimal pooling contract is strictly positive.

Proposition 3. *Assume $\bar{v} > 0$. The pooling problem (16) admits at most one optimal pooling contract (p^P, n^P, γ^P).*

(i) High fraud cost. *If $\hat{k}^f \geq \sigma u$, no triple (p, n, ψ) satisfies the constraints of the pooling problem. Hence, an optimal pooling contract does not exist.*

(ii) Low fraud cost. *If $\Delta \leq \hat{k}^f < \sigma u$, the simplified pooling problem (16) admits a unique solution*

$$\psi^* = \min\{\psi^0,\ \psi_{\max}\},$$

where ψ^0 denotes the interior solution to the first-order condition. The associated optimal pooling contract is

$$(p^P, n^P, \gamma^P) = \big(p(\psi^*),\ n(\psi^*),\ \gamma(\psi^*, p(\psi^*))\big),$$

and the seller's expected trading profit is

$$\bar{v}^P = \tilde{v}(\psi^*),$$

which is strictly positive if and only if $\psi_{\min} < \psi_{\max}$, equivalently, if $c < p(\psi_{\max})$.

Comparative static with respect to e In Corollary 2, I study exogenous changes in the asset's recognizability parameter. Such changes may reflect a technological arms race between fraud techniques and the fraud-detection systems designed to counter them.

Corollary 2. *Holding buyers' market utility Δ fixed, increasing e raises the equilibrium selectiveness ψ^* and the price $p(\psi^*)$.*

Higher payment recognizability increases seller selectiveness for two reasons. First, more accurate fraud detection lowers the marginal cost of selectiveness: increasing selectiveness reduces payment acceptance probabilities by less. Second, higher recognizability allows sellers to be more selective while still delivering the required market utility to buyers, thereby relaxing the upper bound on feasible selectiveness. As a result, sellers optimally choose greater selectiveness as recognizability improves. Greater selectiveness, in turn, allows sellers to charge a higher price without providing excessive utility to fraudulent buyers, thereby increasing sellers' profit margins and raising the break-even precision.

In equilibrium, higher recognizability increases the acceptance probability for a given level of selectiveness. However, sellers optimally respond by raising selectiveness, so the net effect on acceptance probability is ambiguous. As a result, the effect of higher recognizability on equilibrium buyer-seller ratio $n(\psi^*)$is also ambiguous.

The effect on the prior share of fraudulent buyers is likewise ambiguous. Holding selectiveness and break-even precision fixed, higher recognizability requires the prior odds of fraudulent buyers to decline. At the same time, the induced increases in selectiveness and break-even precision tend to raise the prior odds of fraudulent buyers. Because these forces operate in opposite directions, the overall effects are indeterminate.

Comparative statics with respect to k^f In Corollary 3, I study exogenous changes in the cost of committing payment fraud. Such changes may reflect, for example, shifts in the payment environment from in-person to online transactions, which can reduce the cost of fraud by enabling identity theft and unauthorized payments without the need to produce physical counterfeits (such as fake checks or cards).

Corollary 3. *Holding buyers' market utility Δ fixed, an increase in k^f lowers the equilibrium selectiveness ψ^* and raises the equilibrium price $p(\psi^*)$. It also increases the share of fraudsters among accepted matches,*

$$\bar{\gamma} \;=\; \frac{\gamma(\psi^*)}{(1-\gamma(\psi^*))\psi^* + \gamma(\psi^*)}.$$

Higher payment recognizability and a higher fraud cost both weaken sellers' incentives to engage in ex-ante screening through lower prices. However, the two shocks have opposite implications for ex-post screening through selective acceptance. Higher recognizability reduces the cost of selectiveness and therefore increases equilibrium selectiveness In contrast, a higher fraud cost makes deterring fraud easier, weakening the incentive to screen through selective acceptance and thereby reducing equilibrium selectiveness.

Under the information structure, the posterior odds of fraud at the acceptance threshold and the odds of fraud among accepted matches are linked to the prior odds of fraud as

$$\frac{\hat{\gamma}}{1-\hat{\gamma}} = \frac{1}{(1-e)\psi}\frac{\gamma}{1-\gamma}, \qquad \frac{\bar{\gamma}}{1-\bar{\gamma}} = \frac{1}{\psi}\frac{\gamma}{1-\gamma}. \tag{18}$$

Because a higher fraud cost raises the equilibrium price, it increases sellers' profit margins and therefore the break-even precision $\hat{\gamma}$. As a result, the fraud rate among accepted matches also increases.

Comparative static with respect to s In Corollary 4, I consider exogenous changes in the liquidity spread s. When the asset is fiat money, such changes can be interpreted as changes in the inflation rate.

Corollary 4. *Holding the buyers' equilibrium market utility Δ fixed, an increase in s raises equilibrium selectiveness ψ^*, while reducing the equilibrium price $p(\psi^*)$ and buyer-seller ratio $n(\psi^*)$. It also decreases the break-even precision $\hat{\gamma}$ and lowers the share of fraudsters among accepted matches $\bar{\gamma}$.*

Holding selectivity fixed, higher liquidity spreads harm genuine buyers and force sellers to lower payments to maintain participation, thereby reducing seller profits. Increasing selectivity attenuates the impact of these payment reductions on profits. As a result, greater selectivity mitigates the adverse effect of liquidity spreads, leading sellers to optimally choose higher selectivity as liquidity spreads rise. As the acceptance probability falls, keeping fraudulent buyers participating requires sellers to reduce the buyer-seller ratio.

Because a higher liquidity spread lowers payments, it compresses sellers' profit margins and therefore reduces the break-even precision $\hat{\gamma}$. Because a higher liquidity spread lowers payments, it compresses sellers' profit margins and therefore reduces the break-even precision and the fraud rate among accepted matches.

By contrast, the effect of liquidity spreads on the prior share of fraudulent buyers is generally ambiguous. On the one hand, lower odds of fraud among accepted matches require lower prior odds of fraud. On the other hand, greater selectiveness requires a higher prior odds of fraud. Because these forces operate through distinct equilibrium margins, the net effect on the prior share of fraudulent buyers is indeterminate.

3.3 Equilibrium Regimes

Proposition 4. *Fix buyers' market utility Δ with $0 \leq \Delta \leq \sigma u$. A partial equilibrium exists and is payoff-unique.*

(i) High fraud cost. *If $\hat{k}^f \geq \sigma u$, then:*

- *If $0 < c < p^{FB}$, sellers post the optimal separating (first-best) contract.*
- *If $c > p^{FB}$, the equilibrium exhibits no trade.*

(ii) Low fraud cost. *If $\Delta \leq \hat{k}^f < \sigma u$, then there exist two thresholds*

$$0 \leq c_1(\hat{k}^f) \leq c_2(\hat{k}^f) < p^{FB},$$

with both inequalities strict when $\hat{k}^f > \Delta$, such that:

- *If $0 < c < c_1(\hat{k}^f)$, sellers post the optimal separating contract.*
- *If $c_1(\hat{k}^f) < c < c_2(\hat{k}^f)$, sellers post the optimal pooling contract.*
- *If $c > c_2(\hat{k}^f)$, the equilibrium exhibits no trade.*

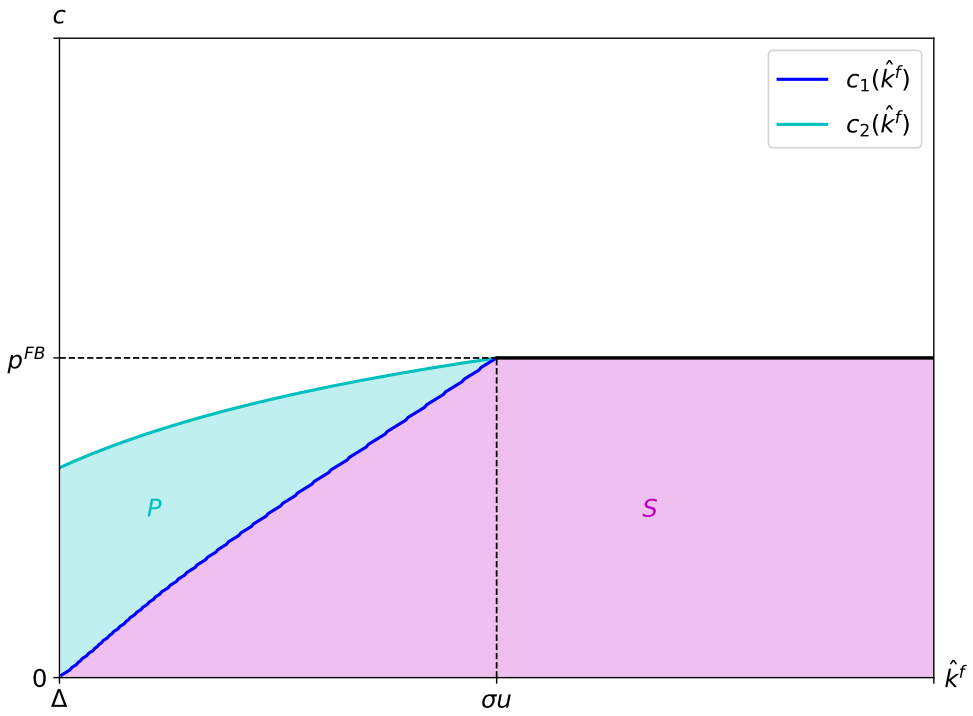
In the special case $\hat{k}^f = \Delta$, we have $c_1(\hat{k}^f) = 0$, so the equilibrium is either pooling or autarkic. If, in addition, $\Delta = 0$, then $c_2(\hat{k}^f) = 0$, and the equilibrium is autarkic for any $c > 0$.

Proposition 4 establishes the existence and payoff-uniqueness of the partial equilibrium. When the fraud cost is high, the equilibrium is efficient: sellers post the first-best contract provided their production cost is below the first-best price. If production cost exceeds the first-best price, the equilibrium exhibits no trade. When the fraud cost is low, the equilibrium depends on sellers' production cost. If production cost is sufficiently low, the equilibrium is separating. For intermediate production costs, the equilibrium is pooling. If production cost is too high, the equilibrium is autarkic. The optimal contract is generically unique, except when parameters lie on the boundaries between equilibrium regimes. In the special case of costless fraud, separating equilibria cannot be sustained for any positive production cost. The equilibrium is then either pooling or autarkic. Figure 2 illustrates the separating, pooling, and autarkic regimes in the $(\hat{k}^f, c)$ plane.

3.4 Limit Cases: Perfectly Recognizable Asset ($e = 1$) or Completely Unrecognizable Asset ($e = 0$)

When $e = 1$, the asset is perfectly recognizable, and the model goes back to the setting in Lagos and Wright (2005); Rocheteau and Wright (2005). Perfect recognizability eliminates the informational friction underlying payment fraud: the seller can attain the first-best expected trading profit with a contract isomorphic to the first-best contract, except that payments are accepted if and only if $z = 0$.

Figure 2: Equilibrium Regimes



Notes: The figure illustrates the equilibrium regimes in the $(\hat{k}^f, c)$ plane. The magenta shaded region, defined either by $\hat{k}^f \geq \sigma u$ and $0 < c < p^{FB}$, or by $\Delta \leq \hat{k}^f < \sigma u$ and $0 < c < c_1(\hat{k}^f)$, corresponds to optimal separating contracts. The cyan shaded region, given by $c_1(\hat{k}^f) < c < c_2(\hat{k}^f)$, corresponds to the optimal pooling contract with selective acceptance. The unshaded region represents no trade. All parameters other than $\hat{k}^f$ and c are fixed at $\sigma = 0.25$, $u = 1.0$, $s = 0.05$, $\Delta = 0.10$ and $e = 0.6$.

When $e = 0$, the asset is completely unrecognizable and the model becomes analogous to Li and Rocheteau (2011). Proposition 5 establishes the existence and payoff-uniqueness of the partial equilibrium in this case. Pooling contracts cannot be sustained, so the equilibrium is either separating or autarkic. In particular, if fraud is costless, the asset ceases to function as a means of payment, and the equilibrium exhibits no trade for any positive production cost. Figure 3 illustrates the equilibrium regimes in the $(\hat{k}^f, c)$ plane when $e = 0$.

Proposition 5. *Assume the asset is completely unrecognizable ($e = 0$). Fix buyers' market utility $0 \leq \Delta \leq \sigma u$. A partial equilibrium exists and is payoff-unique.*

(i) High fraud cost. *If $\hat{k}^f \geq \sigma u$, then:*

- *If $0 < c < p^{FB}$, sellers post the optimal separating (first-best) contract.*
- *If $c > p^{FB}$, the equilibrium exhibits no trade.*

(ii) Low fraud cost. *If $\Delta \leq \hat{k}^f < \sigma u$, then:*

$$0 \leq \bar{c}(\hat{k}^f) < p^{FB},$$

with the inequality strict when $\hat{k}^f > \Delta$, such that: such that:

- *If $0 < c < \bar{c}(\hat{k}^f)$, sellers post the optimal separating contract.*
- *If $c > \bar{c}(\hat{k}^f)$, the equilibrium exhibits no trade.*

In the special case $\hat{k}^f = \Delta$, we have $\bar{c}(\hat{k}^f) = 0$, so the equilibrium s autarkic for any $c > 0$.

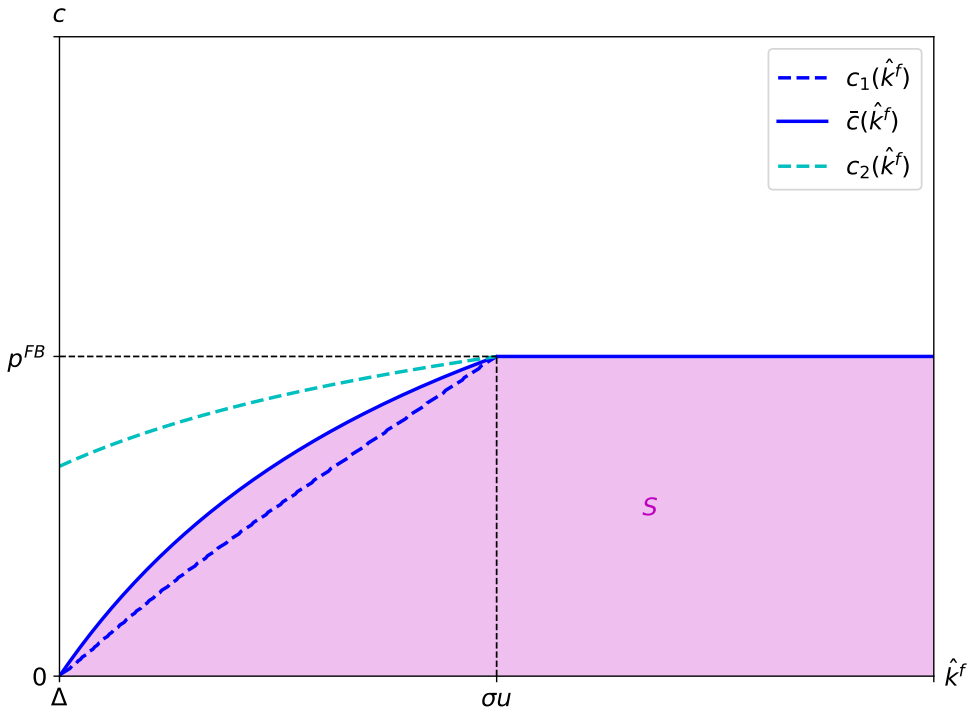
In both limit cases, fraudulent buyers do not participate in equilibrium, and there is therefore no payment fraud on the equilibrium path. Partial recognizability, i.e., $0 < e < 1$, is therefore essential for generating equilibrium payment fraud in the environment.

4 General Equilibrium

In the previous section, I characterized the set of optimal contracts for sellers, denoted by $O(\Delta)$, taking buyers' market utility Δ as given. For a given Δ, the optimal contract is unique, except when parameters lie exactly on the boundary between equilibrium regions.

In this section, I then analyze the general equilibrium. General equilibrium requires adding up: the total measure of active buyers across all open submarkets cannot exceed the fixed supply of buyers, with equality whenever $\Delta > 0$. This condition, together with the partial

Figure 3: Equilibrium Regimes for Completely Unrecognizable Assets



Notes: The figure illustrates the equilibrium regimes in the $(\hat{k}^f, c)$ plane. The magenta shaded region, defined either by $\hat{k}^f \geq \sigma u$ and $0 < c < p^{FB}$, or by $\Delta \leq \hat{k}^f < \sigma u$ and $0 < c < \bar{c}(\hat{k}^f)$, corresponds to optimal separating contracts. All parameters other than $\hat{k}^f$ and c are fixed at $\sigma = 0.25$, $u = 1.0$, $s = 0.05$, and $\Delta = 0.10$.

equilibrium characterization, determines Δ as a fixed point. I find that a general equilibrium exists, but it need not be unique.

A general equilibrium is a fixed point in Δ satisfying the adding-up condition (10) for some probability measure μ supported on $O(\Delta)$. When $O(\Delta)$ is a singleton, μ is degenerate at that contract; when $O(\Delta)$ contains multiple optimal contracts, μ may be any convex combination of them.

Proposition 6 establishes the existence of the general equilibrium. A general equilibrium need not be unique. In particular, distinct values of Δ may induce optimal pooling contracts with an identical buyer-seller ratio, leading to multiple fixed points.

Proposition 6. *A general equilibrium exists; that is, there exists at least one value of* $\Delta \in [0, \sigma u]$ *satisfying the adding-up condition.*

5 Extensions

5.1 Endogeneous Payment Recognizability

In the baseline model, payment recognizability is exogenous. As an extension, suppose an economy can raise fraud detection accuracy for all sellers from a baseline level e_0 starting in period t by paying a fixed upfront cost $C(e)$ in the CM of period $t-1$, with $C(e_0) = 0$ and $C'(e) > 0$. This setup captures upfront investments in fraud detection technology that improve payment recognizability, including data infrastructure, large-scale data collection across merchants, and the development and deployment of improved machine-learning models. Unlike Kahn and Roberds (2008), which models costly monitoring intensity at the transaction level, this paper focuses on the adoption of improved payment technologies and abstracts from per-transaction monitoring costs and their dependence on payment recognizability.

The socially efficient level of payment recognizability maximizes social welfare $\mathcal{W}(e)$ net of the fixed cost $C(e)$. In particular, welfare evaluated at the beginning of CM in any period, is given by

$$\mathcal{W}(e) = A(e) + \frac{1}{1-\beta}\big(U(x^*) - x^*\big)(1+N) + \frac{\beta}{1-\beta}\Big[(u-c)\int_{O_e} \alpha\big(n(\omega)\big)\, F_e(\omega)\, d\mu_e(\omega) - k^f \int_{O_e} n_e(\omega)\, \gamma(\omega)\, d\mu_e(\omega)\Big],$$

where O_e and μ_e denote the general equilibrium objects under payment recognizability a, A(e)=

$$\int_{O_e} n(\omega)\,\big[1-\gamma(\omega)\big]\, p(\omega)\, d\mu_e(\omega) \tag{19}$$

is the total holding of the payment asset and

$$F_e(\omega) \;\equiv\; (1-\gamma(\omega))\, F_e^g\big(\hat{z}(p(\omega),\gamma(\omega))\big) \;+\; \gamma(\omega)\, F^f\big(\hat{z}(p(\omega),\gamma(\omega))\big)$$

is the average acceptance probability in submarket ω. Because the economic environment is stationary, aggregate welfare is discount sum of contemporaneous CM utilities, which are independent of e, and discount sum of DM gains from trade across all agents, which depend on e. Notably, the holding cost of the payment asset is not a net welfare loss, as it reflects a transfer from buyers to the banking sector rather than a real resource cost.

The socially efficient level of payment recognizability maximizes social welfare $\mathcal{W}(e)$ net of the fixed cost $C(e)$. Welfare, evaluated at the beginning of the CM in any period, is given by

$$\mathcal{W}(e) = A(e) + \frac{1}{1-\beta}\underbrace{\big(U(x^*)-x^*\big)(1+N)}_{\text{net CM utilities}} + \frac{\beta}{1-\beta}\underbrace{\left[(u-c)\int_{O_e}\alpha\big(n(\omega)\big)\, F_e(\omega)\, d\mu_e(\omega) - k^f\int_{O_e} n(\omega)\,\gamma(\omega)\, d\mu_e(\omega)\right]}_{\text{DM gains from trade}},$$

where O_e and μ_e denote the equilibrium set of submarkets and the associated distribution under payment recognizability e. The term

$$A(e) = \int_{O_e} n(\omega)\,\big[1-\gamma(\omega)\big]\, p(\omega)\, d\mu_e(\omega) \tag{20}$$

denotes total holdings of the payment asset, which determines the issuer's profits from supplying the payment asset.

$$F_e(\omega) \equiv (1-\gamma(\omega))\, F_e^g\big(\hat{z}(p(\omega),\gamma(\omega))\big) + \gamma(\omega)\, F^f\big(\hat{z}(p(\omega),\gamma(\omega))\big)$$

is the average acceptance probability in submarket ω. Because the environment is stationary, welfare equals the discounted sum of CM utilities, which are independent of e, and DM gains from trade, which depend on e. Importantly, the holding cost of the payment asset is not a net welfare loss, as it represents a transfer from buyers to the issuer rather than a real resource cost.

Would the socially efficient level of payment recognizability be chosen by a profit-maximizing private payment service provider? Improvements in fraud detection generate benefits not only for sellers but also for buyers and the issuer. However, these groups typically cannot be excluded from enjoying these benefits, so the provider cannot charge them and they effectively free ride.

Even if the provider could fully extract the increase in sellers' profits, this would not capture the increase in gains from trade in the DM across all agents. As a result, the level of payment recognizability chosen by a private provider generally differs from the social optimum.

5.2 Multiple Means of Payment and Heterogeneous Sellers

In the baseline model, I assume a single means of payment and homogeneous (representative) sellers. Before turning to the application, it is useful to extend the framework to allow for multiple payment assets and seller heterogeneity.

Each means of payment is indexed by its liquidity spread s. I impose a technological constraint under which each seller can accept only one type of payment assets, s. Sellers may also differ in their ability to recognize fraudulent payments, captured by e, and in their production cost c. A seller is therefore characterized by type $t = (s, e, c)$.

A partial equilibrium requires that, for each seller type $t \in T$, both the optimal acceptance condition and the optimal contract-posting condition are satisfied, defined analogously to the baseline model. Consequently, each seller type has its own optimal acceptance rule, its own set of feasible contracts, and its own optimal contract choice. Because each seller type independently designs its optimal contract subject to type-specific eligibility constraints that depend only on buyers' market utilities, the partial equilibrium of the extended model is obtained by stacking the baseline partial equilibria across seller types. Existence and payoff uniqueness therefore follow directly from the baseline analysis.

Let G denote the distribution of sellers over types $t \in T$, and let μ_t, supported on O_t, denote the distribution of contracts posted by type-t sellers. Adding-up requires that aggregate buyer demand across all submarkets opened by sellers of different payment types and production costs does not exceed the total mass of buyers N:

$$\int_{t \in T} \int_{O_t} n(\omega)\, d\mu_t(\omega)\, dG(t) \leq N, \qquad \text{with equality if } \Delta > 0.$$

As in the baseline model, a general equilibrium exists but need not be unique.

6 Quantitative Analysis

6.1 Calibration

For the quantitative analysis, I assume there is a single payment asset with zero nominal yield, capturing both paper currency and transactional deposits issued by banks (bank IOUs). To make a payment using deposits, a genuine buyer acquires bank deposits in the previous CM

and, in the DM, instructs her bank to electronically transfer funds from her account to the seller's account. Fraudulent buyers, by contrast, submit payment instructions that are later reversed through chargebacks after receiving the DM good; in such cases, banks do not honor the IOUs, and sellers bear the resulting losses.

I assume that a fraction α of sellers operate online and face imperfect payment recognizability ($e < 1$), while the remaining sellers operate in person and can recognize payments perfectly ($\bar{e} = 1$). For online transactions, in the absence of Strong Customer Authentication mechanisms such as 3-D Secure, which has an adoption rate of roughly 3% in the United States, fraud liability typically rests with the seller.[16] Assuming imperfect payment recognizability in the online environment enables the model to capture the effect of fraud risk on payment acceptance decisions and trading volume. For in-person transactions, I assume perfect payment recognizability because counterfeit paper currency is now rare, and card-present transactions in the United States are dominated by EMV chip-enabled payments (including tap-to-pay), in which issuers bear fraud liabilities. As a result, fraud risk does not materially affect sellers' acceptance decisions.[17]

I calibrate the model at an annual frequency. I normalize $u = 1$ and assume $U(x) = B \log x$. The density of production costs, g, prior to normalization, is specified as a quadratic function based on Legendre polynomials and truncated below at zero. Production costs are assumed to be independent of whether a seller operates online or in person.

The parameters governing the distribution of production costs, together with the remaining partial-equilibrium structural parameters $(\alpha, e, k^f, \sigma, B)$ and the buyers' market utility gain Δ, are jointly calibrated to match two sets of moments. At a fixed liquidity spread of $s = 4.5\%$, set equal to the average 3-month Treasury bill yield over October 2024–January 2025, we target: (i) a payment acceptance rate of 94.8% and a fraud rate of 3.4% among accepted payments, as reported in the *2025 Global eCommerce Payments & Fraud Report*; (ii) a share of online payments of 23%, as reported in the *2025 Findings from the Diary of Consumer Payment Choice*; and (iii) the empirical distribution of seller profit margins from the *2025 State of the Amazon Seller* (Jungle Scout). Separately, we compute the model-implied demand for the zero-yield payment asset relative to total output value at different levels of the liquidity spread, holding the real interest rate fixed at 2%, and match these predictions to the binned means of the money demand relationship between M1/GDP and the 3-month Treasury bill yield documented by Lucas Jr and Nicolini (2015). The buyer mass N is then determined from the adding-up condition in (10). The resulting calibrated parameter values are reported in Table 1.

The calibrated model closely matches the targeted fraud statistics, including the fraud

[16] See https://justt.ai/blog/3ds-2-0-what-is-it/.

[17] See https://stripe.com/resources/more/what-are-emv-chip-cards.

Table 1: Calibrated Parameter Values

Parameter	Description	Value
α	Online seller share	0.2452
e	Recognizability	0.9849
k^f	Fraud cost	1.2×10^{-8}
σ	Matching efficiency	0.2834
B	CM utility scale	0.2585
b_1	Linear term in $g(c)$	1.5428
b_2	Quadratic term in $g(c)$	0.7940
Δ	Seller market utility	0.1044
N	Buyer mass	0.1441

rate and the payment acceptance probability, with relative errors of only 0.002% and 0.054%, respectively. Figure 4 compares the model-implied distribution of profit margins to the data, and Figure 5 compares the model-implied money demand curve to its empirical counterpart.

The model's ability to accommodate heterogeneous sellers and match the empirical distribution of seller profit margins is crucial because this distribution identifies the recognizability parameter e. In a representative seller model, e can be recovered in closed form from

$$\frac{\bar{\gamma}}{1-\bar{\gamma}} = (1-e)\frac{\hat{\gamma}}{1-\hat{\gamma}}, \tag{21}$$

where $\hat{\gamma}$, the break-even precision, equals the seller's profit margin under optimal acceptance. Hence, the observed margin, together with the fraud rate, directly pins down e. This relationship carries over to the case in which profit margins differ across sellers but only the aggregate fraud rate is observed. In that case, identifying e requires taking expectations on both sides of the equation above. As a result, recovering e requires matching the cross-sectional distribution of online sellers' profit margins, or the relevant sufficient statistics of that distribution.[18]

The calibrated recognizability parameter $e = 0.985$ implies an ROC-AUC of 98.7%. This estimate falls within the range of 0.95 to 0.99 reported in academic studies using state-of-the-art machine-learning algorithms for payment fraud detection and is broadly consistent with industry anecdotes. The close alignment between the calibrated value and external evidence lends credibility to the model and supports its use for counterfactual analysis.

[18]In partial equilibrium, a heterogeneous-seller economy is observationally equivalent in the aggregate to a representative-seller economy whose profit margin matches the average of the transformed term $\frac{\hat{\gamma}}{1-\hat{\gamma}}$. Hence, a sufficient statistic for identification is the mean of this nonlinear transformation, rather than the mean profit margin itself. Simply matching the average margin $\mathbb{E}[\hat{\gamma}]$ would be incorrect: by Jensen's inequality, the nonlinear transformation implies that using the mean margin understates the model-implied aggregate fraud rate and therefore leads to an underestimation of the recognizability parameter e.

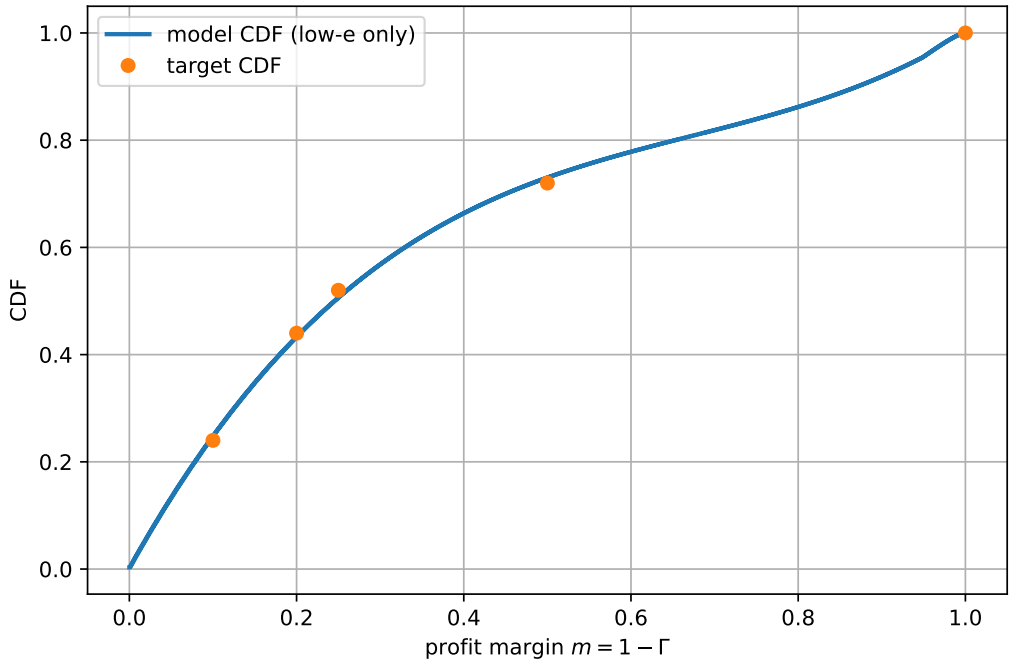


Figure 4: Model Fit: Seller Profit Margin Distribution

Since the calibrated value of k^f is minuscule, the calibrated economy is close to the special case of costless fraud. As a result, in equilibrium, almost all active sellers choose pooling contracts. Figure 6 plots, for each active seller type, the equilibrium level of selectiveness, the posted price, the buyer-seller ratio, and the prior share of fraudulent buyers, as well as the implied fraud rate among accepted payments and the payment acceptance probability.

6.2 Impact of Perfect Payment Recognizability

In this subsection, I quantify the welfare consequences of imperfect payment recognizability by comparing the calibrated economy to a counterfactual with perfect payment recognizability. Because the model does not explicitly capture the trade-off between identity verification intensity and fraud prevention accuracy, as in Kahn and Roberds (2008), the improvement in payment recognizability is interpreted as a technological advance that holds the verification burden fixed.

I proceed in two steps. First, I solve the counterfactual in partial equilibrium (PE) by holding the buyers' market utility, Δ, fixed at its calibrated level. This exercise isolates how perfect recognizability affects the set of active sellers and the optimal contract chosen by each seller type. Second, I allow Δ to adjust endogenously and solve for the counterfactual general equilibrium (GE) by recomputing the fixed point in Δ. I then compare the output, money

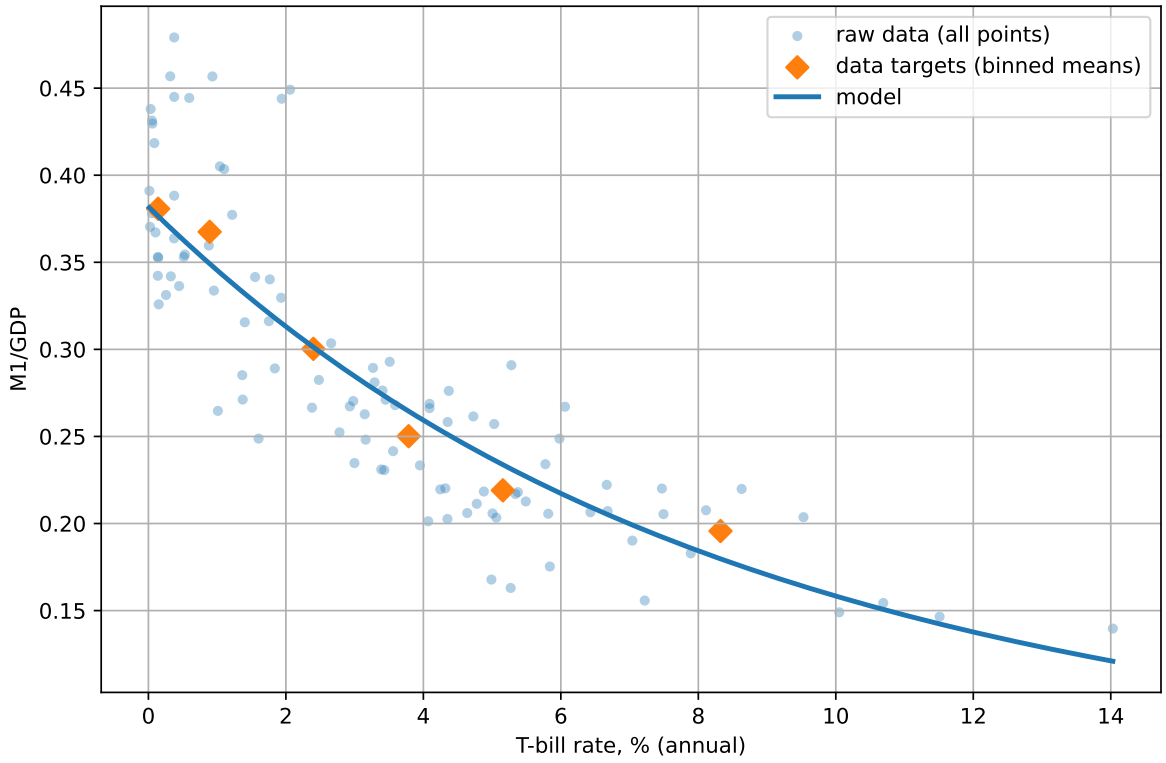


Figure 5: Model Fit: Money Demand Curve

demand, the equilibrium value of Δ, and the distribution of sellers' expected profits across the calibrated economy and the counterfactual.

Counterfactual PE Following the analysis in Subsection 3.4, I solve for the PE under perfect payment recognizability, holding fixed (e, k^f, σ, B) as well as the buyers' market utility Δ at their calibrated values. Figure 7 compares the set of active online sellers in the calibrated economy and in the counterfactual PE, as well as the posted price chosen by each active seller type.

In the calibration, the equilibrium price schedule under imperfect recognizability ($e < 1$) is increasing in sellers' production costs and is capped at a level that is very close to the perfect-recognizability counterfactual ($e = 1$). The reason is that, for sufficiently high-cost types, the optimal pooling contract is associated with the upper bound on selectiveness, $\psi_{\max} = (\sigma u/\hat{k}_f)^e$. Evaluated at this cap, the pooling price is

$$p(\psi_{\max}) = \frac{\hat{k}_f \psi_{\max} - \hat{k}_g}{\hat{k}_f \psi_{\max}/u + s} = \frac{\sigma u (\hat{k}_f/\sigma u)^{1-e} - \hat{k}_g}{\sigma (\hat{k}_f/\sigma u)^{1-e} + s}.$$

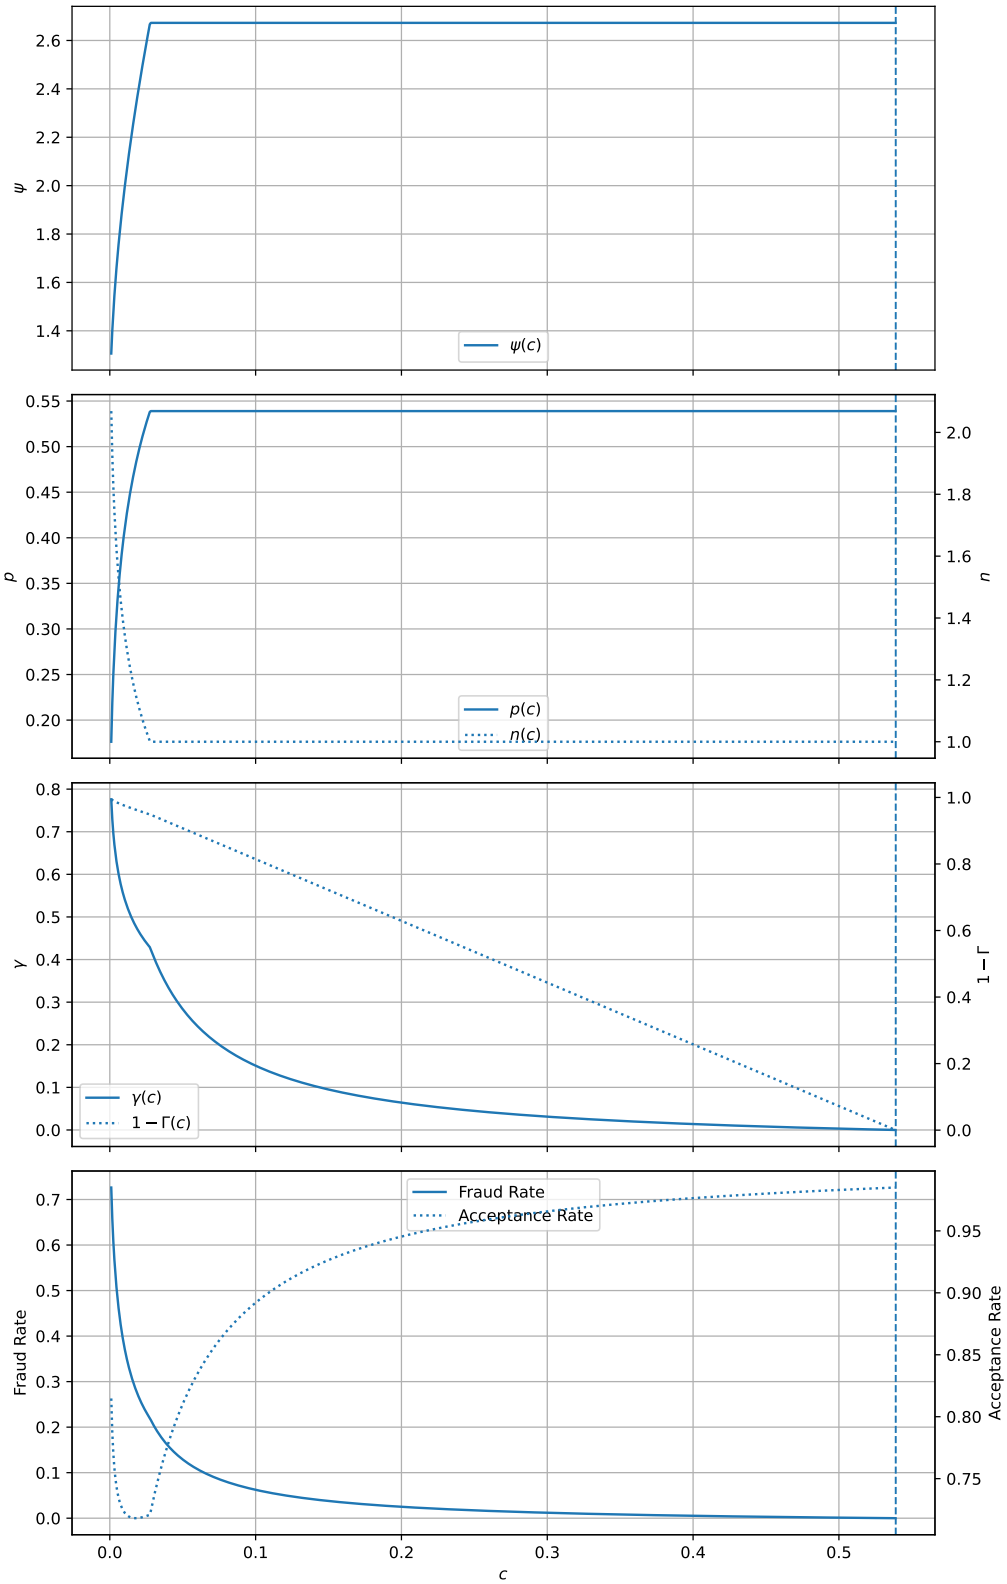


Figure 6: Optimal Contracts by Sellers

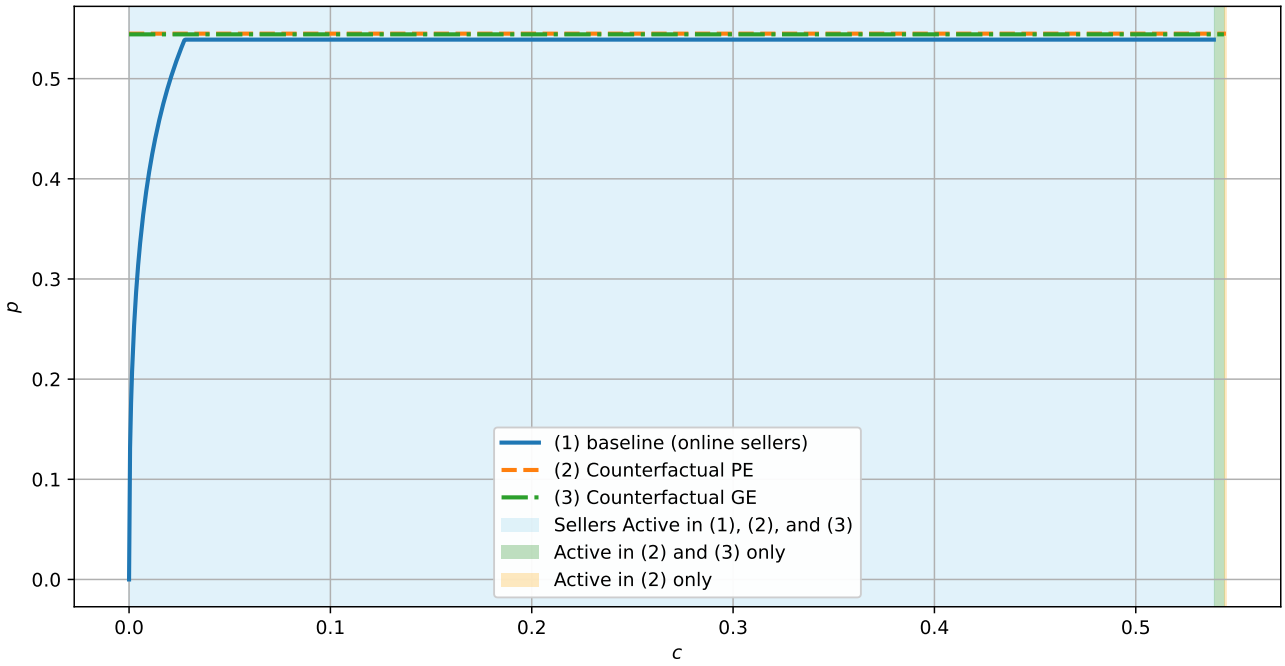


Figure 7: Active Online Sellers and Optimal Prices in the Counterfactual PE and GE

When e is close to one, the term $(\hat{k}_f/\sigma u)^{1-e}$ is only slightly below unity. As a result, the price ceiling in the calibrated economy lies just below $(\sigma u - \hat{k}_g)/(\sigma + s)$, which is the constant posted price under perfect recognizability. Consequently, the two price schedules differ appreciably only for low-cost sellers for whom the optimal selectiveness is interior. For higher-cost sellers, equilibrium prices in the calibrated economy are nearly identical to those in the counterfactual PE. Since posted prices for higher-cost sellers are so similar across the two scenarios, the mass of active sellers in the counterfactual PE is only 2.5% larger than in the calibrated economy.

Counterfactual GE I now turn to the counterfactual general equilibrium, in which the buyers' market utility Δ is determined endogenously. Figure 8 plots the right-hand side of (10), which gives the required buyer mass as a function of Δ in the counterfactual economy. The intersection of this curve with N pins down the equilibrium value of Δ in the counterfactual general equilibrium, which is approximately 0.22% higher than in the calibrated economy. The increase in Δ reflects the fact that perfect recognizability expands the set of active sellers, raising the required buyer mass in the DM. To clear the market given the fixed buyer mass N, buyers' market utility must adjust upward.

Given the resulting value of Δ, I solve for the set of active sellers and their posted prices in the counterfactual GE. These outcomes are also reported in Figure 7, alongside their counterparts in the calibrated baseline economy and in the counterfactual partial equilibrium. Relative to

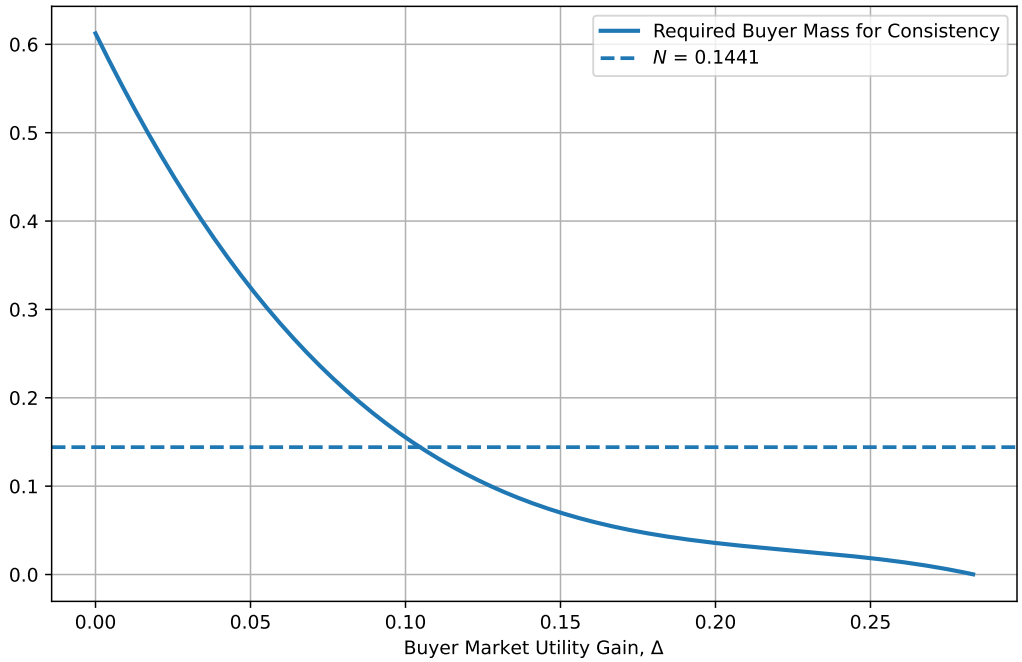


Figure 8: Determination of Buyer Market Utility in the Counterfactual GE

the counterfactual PE, posted prices are lower and the set of active sellers is smaller in the counterfactual GE because buyer market utility is higher due to endogenous adjustment.

Table 2 compares the counterfactual partial and general equilibria to the calibrated baseline. Perfect payment recognizability expands participation on the extensive margin among online sellers: the mass of active online sellers rises by about 3.1% in partial equilibrium and 2.8% in general equilibrium, while the mass of active in-person sellers slightly shrinks (-0.4%) in general equilibrium. The online trading volume increases by about 10.8% in partial equilibrium and 10.3% in general equilibrium. Because DM activity remains a small share of aggregate production, the implied effects on aggregate GDP are modest, with GDP increasing by 0.2% in partial equilibrium and 0.1% in general equilibrium.

The reduction in payment frictions provides substantial benefits to sellers. Total profits of online sellers increase by approximately 27.7% in partial equilibrium and 27.2% in general equilibrium. Buyer market utility rises modestly, by about 0.2% in general equilibrium. Overall gains from trade in the decentralized market, including seigniorage revenue earned by the Federal Reserve and commercial banks and defined as the liquidity spread multiplied by total money demand, increase by about 1.9%. This corresponds to a welfare gain equivalent to roughly 0.15% of GDP. For comparison, the U.S. payments industry accounted for about 1.2% of GDP in 2024, implying that the welfare gain from perfect payment recognizability is economically large, roughly one-eighth the size of the payments industry.

Perfect payment recognizability increases real money balances (M1) by 2.1%. This reflects both higher decentralized-market trade and, more importantly, stronger demand for legitimate payment balances once fraud is eliminated. Put differently, money velocity is higher under imperfect payment recognizability because some transactions are mediated by fraudulent payments rather than by genuine money holdings.

Table 2: Effects of Perfect Payment Recognizability

	Baseline	Counterfactual PE	Counterfactual GE
Active sellers mass (online)	0.0344	103.14	102.76
Active sellers mass (in person)	0.109	100.00	99.62
Total seller profit (online)	0.00147	127.68	127.19
Total seller profit (in person)	0.00578	100.00	99.61
Buyer market utility	0.104	100.00	100.23
DM gains from trade	0.0258	102.15	101.88
DM trading volume (online)	0.00494	110.89	110.32
DM trading volume (in person)	0.0169	100.00	99.49
CM trading volume	0.296	100.05	100.00
GDP	0.318	100.21	100.13
M1	0.0787	102.61	102.09

Notes: The baseline column reports levels. Counterfactual outcomes are expressed relative to baseline (baseline = 100).

Figure 9 illustrates how imperfect payment recognizability affects sellers heterogeneously across production costs. The profit gains from perfect recognizability are larger for low-cost sellers than for high-cost sellers. Intuitively, low-cost sellers have higher profit margins, implying a higher break-even precision and greater exposure to payment fraud.

Efficiency of Private Provision of Payment Recognizability Suppose a monopolistic payment service provider can pay a fixed cost to improve recognizability from its baseline level of 0.985 to 1. Even if the provider can engage in first-degree price discrimination and extract the entire increase in profits to online sellers, the maximum revenue it can obtain from this improvement is only

$$0.00147 \times 27.2\% = 0.00039984.$$

By contrast, the associated social welfare gain is

$$0.0258 \times 1.9\% = 0.0004902.$$

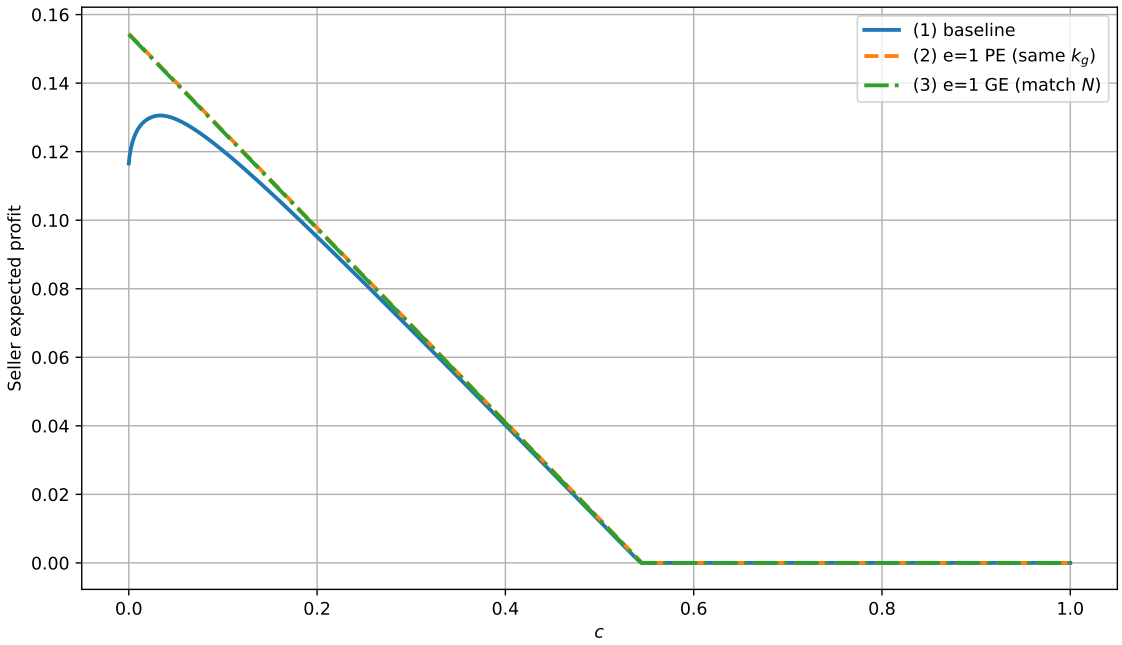


Figure 9: Effects of Perfect Payment Recognizability: Seller Profit by Type

Therefore, if the per-period equivalent of the upfront cost lies between these two values, a private provider will not undertake the investment, even though it is socially efficient. In this sense, payment recognizability is underprovided in equilibrium.

The source of the inefficiency is that the benefits of better fraud detection extend beyond online sellers. Buyers also benefit, as do the issuers of the payment instrument, including the Federal Reserve and commercial banks, through higher demand for legitimate payment balances. Yet a private payment service provider typically cannot charge these parties for the gains they receive. As a result, the provider captures only part of the social benefit created by better recognizability.

Although in-person sellers may be harmed by an improvement in online payment recognizability, this negative effect is quantitatively small and does not materially alter the main conclusion: the social benefit to improving payment recognizability exceeds the private benefit, implying underprovision by a for-profit payment service provider.

6.3 Impact of a Competing Payment Asset with Superior Recognizability on Deposit Demand

Payment innovations built on distributed ledger technology (DLT), such as payment stablecoins and central bank digital currencies, may offer stronger fraud detection relative to tradi-

tional bank payment rails. First, in token-based digital payment architectures, sellers need not store sensitive customer payment credentials, reducing vulnerabilities associated with database breaches. Second, these instruments can enable near-instantaneous settlement, shortening the window during which fraudulent transactions can be reversed or exploited compared with traditional systems that rely on deferred settlement or batch processing. Third, DLT-based systems operating on unified or coordinated ledgers can facilitate consolidated transaction monitoring, potentially enhancing fraud detection relative to traditional banking and card networks that rely on fragmented infrastructures with limited cross-institutional data sharing.

Despite these promising features, systematic evidence on payment fraud under these emerging digital payment assets remains limited. In this subsection, I conduct a counterfactual exercise in which a hypothetical competing payment asset is introduced with a 10% improvement in payment recognizability for online transactions, while maintaining the same fraud cost as the existing payment instrument. I then examine how its introduction affects the demand for bank deposits.

In this counterfactual analysis, online sellers choose whether to post contracts requesting payment in bank deposits or in the competing instrument, depending on which instrument yields higher expected trading profits. Consider a seller with production cost c. The seller accepts deposits if

$$\bar{v}(s, e, c) > \bar{v}(s_{new}, e_{new}, c),$$

where $\bar{v}(s, e, c)$ and $\bar{v}(s_{new}, e_{new}, c)$ denote the seller's expected DM trading profit when requesting payment in deposits or the competing instrument, respectively.

Because the previous subsection shows that general equilibrium effects are quantitatively small even under the extreme case of fully eliminating imperfect payment recognizability, I focus on the partial equilibrium in this exercise and abstract from general equilibrium feedback effects.

I fix the nominal yield on deposits at zero and set the T-bill yield to $i_{\text{bond}} = 4.5\%$. Let i_{new} denote the nominal yield of the new competing payment instrument. Its liquidity spread is

$$s_{\text{new}} = \frac{i_{\text{bond}} - i_{\text{new}}}{1 + i_{\text{new}}}.$$

Figure 10 plots the demand for deposits and the competing instrument in online transactions (relative to GDP) as a function of i_{new}. Despite offering a lower yield than deposits, the competing instrument captures positive market share due to superior fraud detection. For example, with a yield 100 basis points below deposits, its market share is about 5.5%; when the spread narrows to -25 basis points, the share rises to roughly 15.9%.

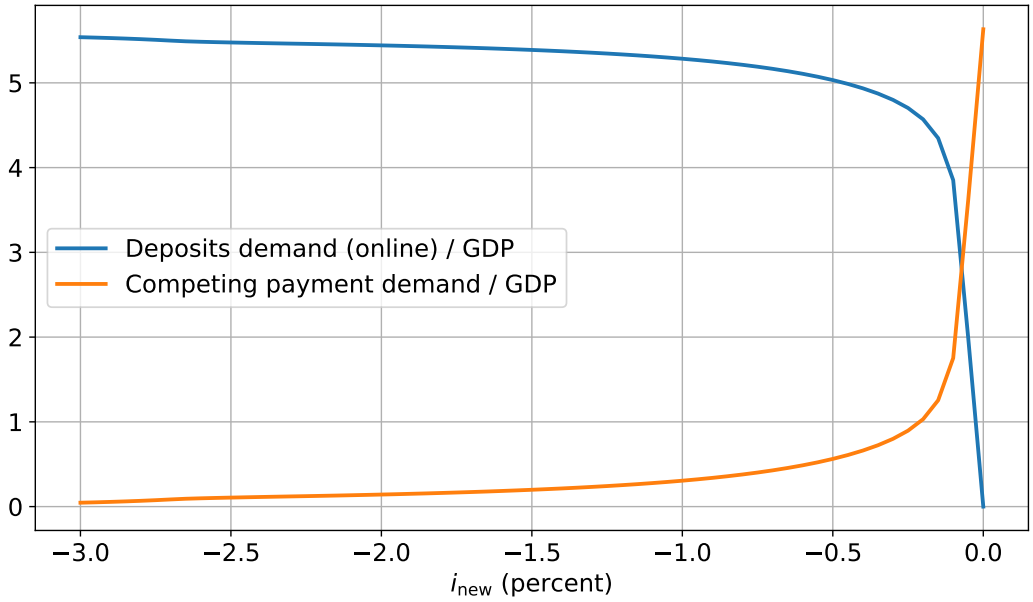


Figure 10: Deposit vs New Payment Demand as i_{new} Varies (PE)

7 Conclusion

This paper introduces imperfect payment recognizability into a monetary search model of decentralized exchange. Sellers observe a noisy signal about payment authenticity and decide whether to accept payments at posted prices, capturing the accept–reject choice that characterizes fraud detection in modern payment systems. When fraud is inexpensive to perpetrate and seller production cost is intermediate, equilibrium features pooling contracts with both mistaken acceptances of fraudulent transactions (false negatives) and mistaken rejections of legitimate payments (false positives). By embedding an information structure that maps directly into the ROC-AUC statistic used in practice, the model provides a tractable bridge between modern fraud detection technologies and monetary theory.

Calibrating the model to observed payment fraud rates, seller profit margins, and money demand implies a high but imperfect degree of recognizability and near costless fraud in online payments. Under these conditions, fraud becomes pervasive, with most sellers adopting pooling contracts and experiencing equilibrium fraud. Moving to perfect recognizability would substantially increase online trade and seller profits, generating meaningful welfare gains even though decentralized trade represents a small share of aggregate output.

The analysis also shows that payment recognizability is an economically important dimension of competition across payment instruments. A competing digital payment asset with superior fraud detection can attract market share even at a yield disadvantage, implying that

recognizability, in addition to yield, is an important determinant of payment demand in online transactions.

References

Alvarez, F., and D. Argente (2022): "On the effects of the availability of means of payments: The case of uber," *The Quarterly Journal of Economics*, 137(3), 1737–1789.

Baxter, W. F. (1983): "Bank interchange of transactional paper: Legal and economic perspectives," *The Journal of Law and economics*, 26(3), 541–588.

Bethune, Z., G. Rocheteau, T.-N. Wong, and C. Zhang (2022): "Lending relationships and optimal monetary policy," *The Review of Economic Studies*, 89(4), 1833–1872.

Cavalcanti, R., and E. Nosal (2011): "Counterfeiting as private money in mechanism design," *Journal of Money, Credit and Banking*, 43, 625–636.

Chiu, J., S. M. Davoodalhosseini, J. Jiang, and Y. Zhu (2023): "Bank market power and central bank digital currency: Theory and quantitative assessment," *Journal of Political Economy*, 131(5), 1213–1248.

Ferraris, L., and M. Watanabe (2008): "Collateral secured loans in a monetary economy," *Journal of Economic Theory*, 143(1), 405–424.

Gu, C., F. Mattesini, and R. Wright (2016): "Money and credit redux," *Econometrica*, 84(1), 1–32.

Guerrieri, V., R. Shimer, and R. Wright (2010): "ADVERSE SELECTION IN COMPETITIVE SEARCH EQUILIBRIUM," *Econometrica*, 78(6), 1823–1862.

Hashemi, S. K., S. L. Mirtaheri, and S. Greco (2022): "Fraud detection in banking data by machine learning techniques," *Ieee Access*, 11, 3034–3043.

Hu, T.-W. (2013): "Imperfect recognizability and coexistence of money and higher-return assets," *Economic Theory*, 53(1), 111–138.

Kahn, C. M., and W. Roberds (2008): "Credit and identity theft," *Journal of Monetary Economics*, 55(2), 251–264.

——— (2009): "Why pay? An introduction to payments economics," *Journal of Financial Intermediation*, 18(1), 1–23.

Kang, K.-Y. (2017): "Counterfeiting, screening and government policy," *Journal of Economic Theory*, 172, 26–54.

Koulayev, S., M. Rysman, S. Schuh, and J. Stavins (2016): "Explaining adoption and use of payment instruments by US consumers," *The RAND Journal of Economics*, 47(2), 293–325.

Lagos, R., and R. Wright (2005): "A unified framework for monetary theory and policy analysis," *Journal of political Economy*, 113(3), 463–484.

Lester, B., A. Postlewaite, and R. Wright (2012): "Information, liquidity, asset prices, and monetary policy," *The Review of Economic Studies*, 79(3), 1209–1238.

Li, Q., and R. Shimer (2025): "Screening with Price and Data: Adverse Selection and Information Heterogeneity in Competitive Search Markets," *Available at SSRN 5251260*.

Li, Y., and G. Rocheteau (2011): "On the threat of counterfeiting," *Macroeconomic Dynamics*, 15(S1), 10–41.

Li, Y., G. Rocheteau, and P.-O. Weill (2012): "Liquidity and the threat of fraudulent assets," *Journal of Political Economy*, 120(5), 815–846.

Lucas Jr, R. E., and J. P. Nicolini (2015): "On the stability of money demand," *Journal of Monetary Economics*, 73, 48–65.

Martin, A., M. J. Orlando, and D. Skeie (2008): "Payment networks in a search model of money," *Review of Economic Dynamics*, 11(1), 104–132.

Nguyen, N., T. Duong, T. Chau, V.-H. Nguyen, T. Trinh, D. Tran, and T. Ho (2022): "A proposed model for card fraud detection based on Catboost and deep neural network," *Ieee Access*, 10, 96852–96861.

Nosal, E., and G. Rocheteau (2006): "The economics of payments," *FRB of Cleveland Policy Discussion Paper*, (14).

Nosal, E., and N. Wallace (2007): "A Model of (the Threat of) Counterfeiting," *Journal of Monetary Economics*, 54(4), 994–1001.

Quercioli, E., and L. Smith (2015): "The economics of counterfeiting," *Econometrica*, 83(3), 1211–1236.

Rochet, J.-C., and J. Tirole (2002): "Cooperation among competitors: Some economics of payment card associations," *Rand Journal of economics*, pp. 549–570.

——— (2003): "Platform competition in two-sided markets," *Journal of the european economic association*, 1(4), 990–1029.

Rocheteau, G., and R. Wright (2005): "Money in search equilibrium, in competitive equilibrium, and in competitive search equilibrium," *Econometrica*, 73(1), 175–202.

Shi, S. (1995): "Money and prices: a model of search and bargaining," *Journal of Economic Theory*, 67(2), 467–496.

Shy, O. (2023): "Cash is alive: How economists explain holding and use of cash," *Journal of Economic Literature*, 61(4), 1465–1520.

Shy, O., and Z. Wang (2011): "Why do payment card networks charge proportional fees?," *American Economic Review*, 101(4), 1575–1590.

Telyukova, I. A., and R. Wright (2008): "A model of money and credit, with application to the credit card debt puzzle," *The Review of Economic Studies*, 75(2), 629–647.

Trejos, A., and R. Wright (1995): "Search, bargaining, money, and prices," *Journal of political Economy*, 103(1), 118–141.

Van Vlasselaer, V., C. Bravo, O. Caelen, T. Eliassi-Rad, L. Akoglu, M. Snoeck, and B. Baesens (2015): "APATE: A novel approach for automated credit card transaction fraud detection using network-based extensions," *Decision support systems*, 75, 38–48.

Velde, F. R., W. E. Weber, and R. Wright (1999): "A model of commodity money, with applications to Gresham's law and the debasement puzzle," *Review of Economic Dynamics*, 2(1), 291–323.

Wang, D., B. Chen, and J. Chen (2019): "Credit card fraud detection strategies with consumer incentives," *Omega*, 88, 179–195.

Wedge, R., J. M. Kanter, K. Veeramachaneni, S. M. Rubio, and S. I. Perez (2018): "Solving the false positives problem in fraud prediction using automated feature engineering," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 372–388. Springer.