The Unintended Consequences of Abundant Bank Liquidity

Erwan Quintin*
University of Wisconsin – Madison

July 15, 2026

Abstract

Most large central banks now saturate bank liquidity markets in order to exercise direct control on interbank lending rates. This has several operational advantages and mitigates spurious volatility in liquidity prices. But it also distorts the fundamental information and early warnings interbank liquidity markets convey to financial intermediaries and investors. In a rational expectation framework, I show that by supressing fundamental information abundant liquidity policies can in fact result in lower average output, deeper recessions, and lessen agents' ability to share idiosyncratic risk. The model also enables me to formally discuss the distinction between ample and abundant liquidity levels.

Keywords: Ample reserves framework; Interbank lending; Monetary policy
JEL codes: E52; G01; G21

*Erwan Quintin, 5151A Grainger Hall, 975 University Avenue, Madison, WI 53706. I thank Sangeeta Pratap, Enrique Martinez-Garcia, Cyril Monnet, Gadi Barlevy, Wenli Li, and seminar participants at the 2026 SAET meetings for helpful comments. All remaining errors are mine. All code is available for download at http://erwan.marginalq.com/index_files/wp.htm.

"Our ample reserves framework works well."

Jerome Powell, Chair of the Federal Reserve of the United States, 2025 [27]

"The costs of abundant reserve systems take considerable time to appear, are likely to grow and tend to be less visible . . . The central bank has taken over much of the intermediation in the overnight interbank market . . . The abundant reserve system kills the overnight interbank market."

Claudio Borio, Head of the Bank of International Settlements Monetary and Economic Department, 2023 [3]

1 Introduction

The Federal Reserve Bank, among other large central banks, has switched from a scarce reserve system to an ample reserve system. In an environment with scarce reserves, the framework that prevailed before 2008 in most countries, central banks typically control the interest rate on reserves via open-market operations. In an ample-reserve system, the supply of reserves saturates banks' precautionary liquidity needs by design, and central banks can in principle control the same interest rate simply by paying interest on reserves at the target level. In an ample reserve system, as Irhig et al., [19] explain "*banks do not find much benefit from holding additional reserves other than earning the interest the Fed pays on these balances.*"

One view is that the ample-reserve system was born out of necessity, as central banks found it ever more difficult to influence rates in an environment with unprecedented levels of excess reserves (see Chien and Stewart [5] for a discussion) in the wake of the Great Financial Crisis. However, central bankers have emphasized the new paradigm's operational advantages (see e.g. Powell [27] and Clouse et al. [6]) and made clear that the ample reserve system is the new normal of monetary policy implementation. In fact, some both in academia and policy circles had touted those operational advantages well before the crisis struck, as Ennis and Keister [9] discuss. Besides limiting if not suppressing unwanted volatility in policy rates, saturating reserve markets also means for instance that reserve requirements can be, as Ennis and Keister [9] put it, "*simplified or eliminated*," as they officially were in the United States in 2020, removing a source of "*administrative burdens on both banks and the central bank.*"

Central banks (see Afonso et al. [1] for a discussion) have also acknowledged the potential drawbacks of this new normal. The large balance sheet necessary to support ample reserves carries political and inflation risks and could crowd out lending (see e.g. Diamond et al. [8]). Furthermore, as Borio [3] discusses at length, by saturating bank liquidity markets, ample reserve frameworks distort the price of bank liquidity. As Afonso et al. [1] put it, "*abundant reserves discourage unsecured interbank trading and active reserve redistribution, hindering market discipline and price discovery.*"

My main goal in this paper is to formalize the consequences of liquidity market distortions in a rational expectation model of bank liquidity management that features aggregate credit quality shocks. In my model, normally functioning liquidity markets endogenously provide banks with advanced warnings of credit shocks which enable them to mitigate the consequences of those shocks. When the price signal is suppressed in liquidity markets as a consequence of monetary policy implementation loan formation, investment, and output fall, and the impact of negative credit shocks becomes more severe. To my knowledge, this constitutes the first formalization of the real consequences of the informational distortions caused by abundant liquidity in a fully articulated rational expectation framework.

When bank liquidity needs are saturated by central banks, an aggregate spike in liquidity demand no longer causes large spikes in the cost of liquidity because banks hold sufficient excess reserves to absorb these shocks. One salient manifestation of this is the collapse of unsecured interbank lending since 2008. Figure 1 shows that the volume of interbank transactions, measured as a ratio to GDP, has fallen markedly since 2010. This is also true of other markets banks used to tap routinely to meet their liquidity needs. Figure 1 shows for instance the fast shrinking of the financial commercial paper market since the crisis. It should not be surprising that banks have exploited the ease with which central banks are pushing liquidity on them which has resulted in less interbank intermediation. But interbank markets do not simply enable banks to manage liquidity shocks, they also convey information about

Figure 1: Traditional sources of liquidity for banks

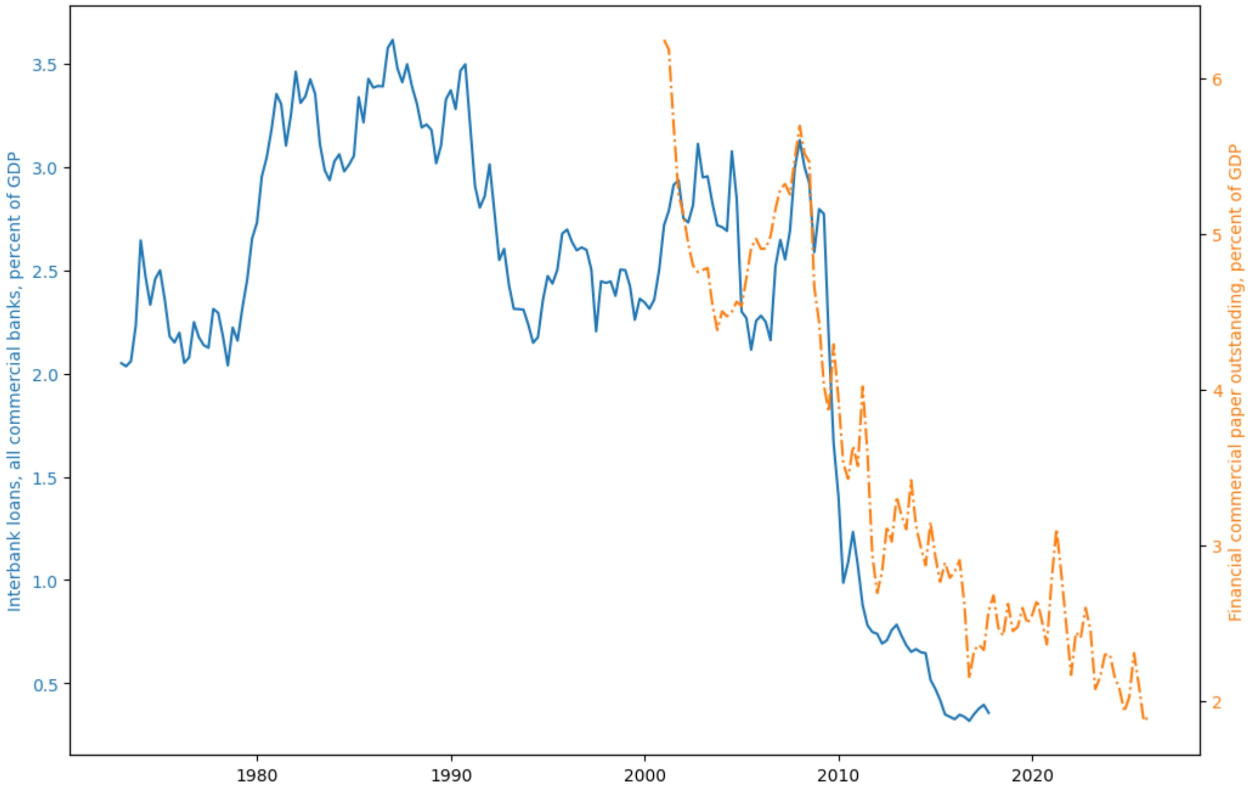

Sources and notes: Board of Governors, downloaded via FRED. Interbank transactions (blue) and financial CP volumes (orange) are series IBLACBW027NBOG and FINCP, respectively. The maximum length of data is shown for each series, resulting in different time scales.

the aggregate state of the banking industry. Sudden spikes in the demand for liquidity can be transient and innocuous, but they can also be canaries in the landmine in an industry whose core business involves exposure to credit risks and liquidity mismatches.

Historically, interbank market disruptions have often preceded lending contractions and recessions as was the case, notoriously, before the 2008 crisis (see, e.g., Taylor [32].) Flare-ups in interbank lending spreads tend to be followed by a tightening of lending standards and attempts by banks to bolster their balance sheet in preparation for possible disruptions. Those evasive measures can amplify economic crises by causing a destabilizing wave of panic sales of long-term assets. But they also enable banks to mitigate upcoming losses before they get worse.

Figure 2 shows that the volatility of liquidity spreads both in the interbank and the financial commercial paper markets has been significantly reduced. The bottom panel of the figures shows the compression of rolling one-year inter-quantile ranges. Standard statistical tests confirm that volatility in the post-2010 sample – measured either by standard deviation or inter-quantile range – is significantly lower than it was prior to 2008.[1] One way to look at this phenomenon is that the intended goals of the ample liquidity doctrine have been met. My simple point in this paper is that liquidity markets no longer provide as much fundamental information to intermediaries as they used to, and this is unlikely to be without consequences.[2]

The model I build to formalize the information content of bank liquidity prices borrows features from the canonical frameworks of Poole [26] and Diamond and Dybvig [7]. My main innovation in this respect is to add aggregate credit quality shocks to those standard models. When borrowers discover that the quality of their investment is bad they become more likely to request early liquidity payouts which introduces a correlation between liquidity calls and credit conditions. In the resulting model, interbank markets serve the standard purpose of enabling banks to share idiosyncratic liquidity and economic risk. But because those risks are aggregate in nature, interbank markets have the additional advantage of helping banks distinguish between pure liquidity and fundamental shocks. When the price of liquidity

[1] The post-2010 standard deviation of the TED spread is half of what it was pre-2008 in the time period plotted of the figure, with an F-stat for the difference in the two samples of over 650. The post-2010 standard deviation is 36% lower for the CP spread compated to pre-2008 data, with an F-stat of 203.

[2] One source of liquidity to banks that remains subject to well documented flare-ups is the secured overnight financing (a.k.a. repo) market. See Logan and Schulhofer-Wohl [25] for a discussion of those phenomena. Although the spread between the secured overnight financing rate and the fed funds rate has been confined to a 20bps range since 2022, noticeable volatility spikes have occurred as recently as in September of 2025. As Logan and Schulhofer-Wohl [25] discuss, those episodes appear caused by technical factors such as a large settlement of new securities and/or tax payment dates. Those technical events are unlikely to convey much information to banks about credit quality. The interbank and financial commercial paper markets are different in that they are primarily driven by bank fundamentals. Their progressive displacement by Federal Reserve backed liquidity causes potentially important information to be suppressed. Section 9 expands on the connection between the ample reserve framework and the information content of the cost of liquidity to banks.

Figure 2: Liquidity spreads

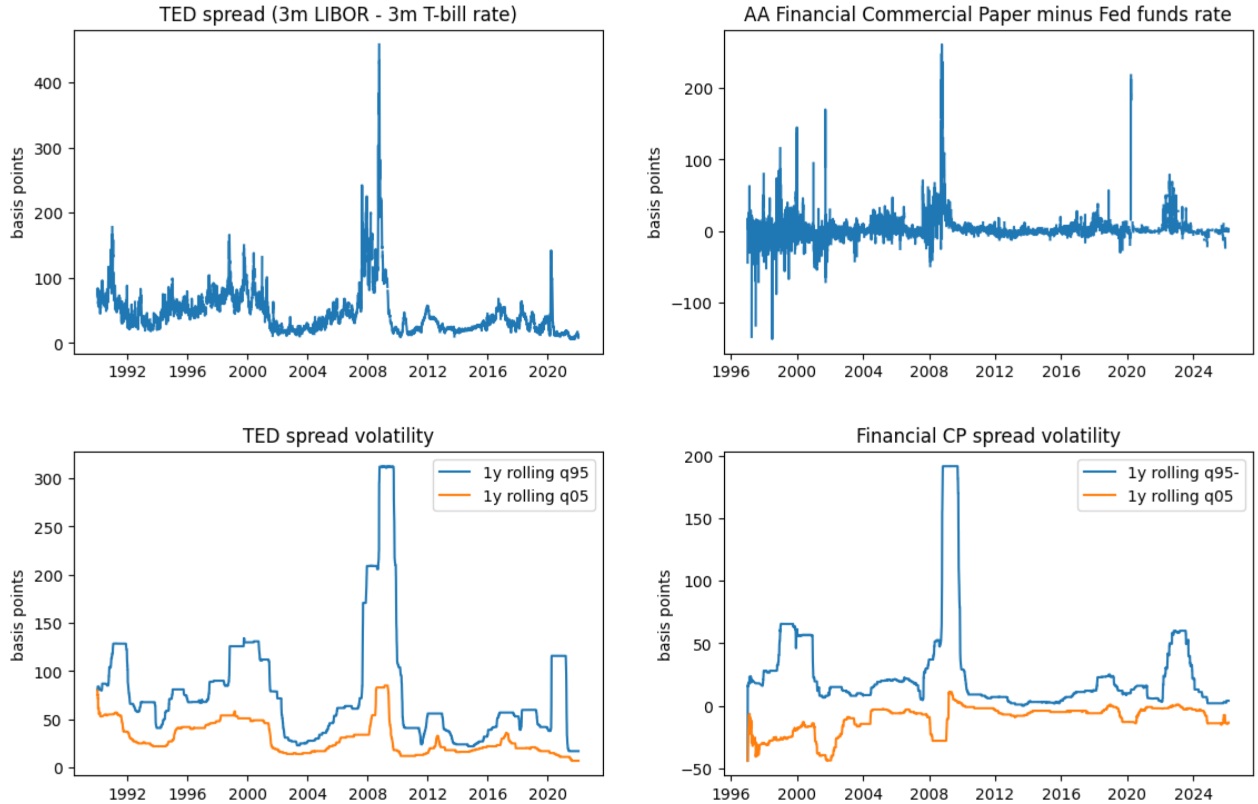

Sources and notes: Board of Governors, downloaded via FRED. The TED spread, AA Financial Commercial paper rate, and fed funds rate data are from series TEDRATE, DCPF1M, and DFF, respectively. The maximum length of data is shown for each series, resulting in different time scales.

suggests that average loan quality has become low, banks are able to take defensive measures that make the eventual consequences of those credit shocks smaller. When, instead, a central bank saturates liquidity markets so it can control the price of liquidity via administered rates, fundamental information is lost, and the impact of aggregate credit shocks becomes more severe. Another consequence of suppressing the information content of interbank markets is that it prevents banks from fully sharing idiosyncratic risk fully. Aggregate consumption, as a result, can become subject to both aggregate and idiosyncratic risk, whereas normally functioning interbank markets would largely eliminate the idiosyncratic component.

The model also enables me to formally discuss the distinction between ample reserves and abundant reserves the Federal Reserve has insisted on making in monetary policy communications over the past few years. As Haubrich [17], Gagnon and Sack [11], or Afonso et al. [1] (among many others) discuss, defining that distinction precisely, not to mention estimating the level at which reserves are ample but not abundant, has been a challenge. In my model, there are two natural and related interpretations of the notion of ample reserves. The first is the lowest level of reserve supply such that the price of liquidity is at the administered interest rate on reserves for all possible aggregate demand shocks. At that threshold level of reserves, a small unexpected change in the supply of reserves – which could be caused in practice by real-time estimation error or a shock to the size of the central bank's balance sheet – causes a small response in liquidity rates. By operating at that threshold, the central bank minimizes the expected cost of remunerating reserves but has to accept potential deviations from the target rate if unexpected supply shocks occur.

The second interpretation involves setting a level of reserves which keeps the cost of liquidity at the target rate except if outlying aggregate demand shocks materialize. At that level of reserves, the cost of liquidity now reveals outlying credit shocks which means, in turn, that an ample reserve framework can help banks mitigate the real impact of those shocks. This promotes loan formation relative to an abundant reserve framework that suppresses the price signal altogether. This provides a rationale for trying to fine-tune the implementation of monetary policy by making sure that the supply of reserves is not higher than it needs to be to keep rate volatility low. In addition to mitigating the risks and costs associated with a large central bank balance sheet, allowing small deviation from target rates conveys early information about large credit shocks to banks.

On the technical front, this paper provides an existence argument for Radner equilibria in an environment with a continuum of agents and a continuum of signals. In this respect, I follow Sun et al. [31] in assuming that private signals are independent conditional on macroeconomic conditions. This circumvents the implementation and existence issues discussed by Kreps [23] and Glycopantis et al. [14] because agents cannot attempt to influence equilibrium prices in their favor. I provide a general existence argument in the context of a canonical model of bank liquidity which enables me to formalize the consequences of government interference with the price mechanism in those markets. In so doing I provide, to my knowledge, the first theoretically consistent formalization of the unintended informational consequences of ample bank liquidity policies.

2 Physical environment

As a useful benchmark, I will first lay out and characterize equilibria in a bank liquidity management model without informational frictions or aggregate credit shocks. The model merges the main features of the canonical models of Poole [26] and Diamond and Dybvig [7]. It provides a simple rationale for interbank liquidity markets when banks face idiosyncratic liquidity shocks. Once I add informational frictions and aggregate credit shocks, in sections 5 and 6, respectively, it will enable me to fully characterize the informational content of liquidity markets.

Consider an environment with three dates ($t \in \{0, 1, 2\}$) and a unit continuum of locations. Each location contains one bank and a unit continuum of borrowers. Borrowers at a given location are each endowed with a project in which they can invest any non-negative quantity k of capital at $t = 0$. The project generates payoff Ak^α at date 2 with probability $p > 0$ and nothing otherwise, with $\alpha \in (0, 1)$.

The capital committed to risky projects can be redirected at date 1 to a safe linear technology with a unit transformation rate of $0 \leq \delta \leq 1$. This option will play a key role in the upcoming analysis. It means that the bank is able to make fundamental adjustments to its portfolio when it learns new information at interim date 1.

Instead of operating their project, each borrower can devote their time to an outside activity that yields $w \geq 0$ for sure at date 1. With probability $\phi > 0$ borrowers value consumption only at date 1. With probability $1 - \phi$ they are indifferent between consuming at date 1 or at date 2. Abusing language slightly I will refer to the first type of borrower as early consumers while I call borrowers who are indifferent as to the timing of their consumption late consumers.

Consumption profiles for borrowers are

$$(c^E, c_1^L, \bar{c}_2^L, \underline{c}_2^L) \in I\!R_+^4$$

where c^E denotes their consumption at date 1 when they are early consumers, c_1^L is their consumption at date 1 if they are late consumers while $(\bar{c}_2^L, \underline{c}_2^L)$ is the consumption of late consumers at date 2 if the project succeeds or fails, respectively. Their expected utility as of date 0 given such a consumption plan is

$$\phi \ln c^E + (1-\phi) \ln (c_1^L + p\bar{c}_2^L + (1-p)\underline{c}_2^L).$$

I use this log-log specification of preferences for concreteness only, the key feature of these preferences for my purposes is that they satisfy Inada conditions for early consumers so that $c^E > 0$ in any of the equilibria I will study in this paper.

At each location, both liquidity and project success shocks satisfy a law of large number[3] so that the fraction of borrowers who get the early consumption draw is ϕ while exactly fraction p of projects succeed.

The second type of agent at each location, the bank, is risk-neutral, is endowed with an amount l of liquidity at date 0, and a storage technology that can carry liquidity risk-free across periods with a net return of 0. At $t = 0$, the bank decides what fraction of its liquidity to lend to borrowers.

[3] See Feldman and Gilles [10], Judd [20], or Sun [30] for a discussion of the technical difficulties I side-step with this assumption.

3 Loan contracts

At each location, the bank offers each borrower a contract $(k, s, c^E, c_1^L, \bar{c}_2^L, \underline{c}_2^L) \in [0, l]^2 \times I\!R_+^4$ that stipulates:

1. an investment $k \in [0, l]$ in the project at date 0;

2. a quantity $0 \leq s \leq k$ of capital redirected to the safe technology at date 1;

3. a consumption plan $(c^E, c_1^L, \bar{c}_2^L, \underline{c}_2^L) \in I\!R_+^4$.

The assumption that the bank offers the same, type-symmetric contract to each borrower is innocuous for my purposes since all borrowers are identical at $t = 0$.

To simplify the contracting problem I will assume that the date-two consumption the bank can offer each borrower is bounded above by the project's output, which implies $\underline{c}_2^L = 0$. This amounts to a standard limited commitment assumption. My key results only require that $c^E > \underline{c}_2^L$ which can be guaranteed for instance by introducing a private effort friction, but the above assumption reduces both the length of arguments and notation. For one thing, I can now write c_2^L for the consumption of late consumers at date 2, where it is now understood that it is contingent on the project succeeding.

A type-symmetric contract, then, is feasible if it satisfies the following resource constraints:

$$\begin{aligned} s &\leq k, \\ \phi(c^E + c_1^L) &\leq l - k, \\ c_2^L &\leq A\,(k - s)^\alpha\,, \end{aligned}$$

if it meets each borrower's participation constraint:

$$\phi \ln c^E + (1 - \phi) \ln (c_1^L + p c_2^L) \geq \ln (w),$$

and if late customers do not have an incentive to misrepresent their type (recall that late customers value early and late consumption equally):

$$pc_2^L \geq c^E.$$

Throughout this paper, I will find that the last requirement does not restrict the space of feasible contracts, but this will require verification in all versions of the economy I consider.

Two additional observations are useful towards understanding these conditions. First, the feasibility condition for c_2^L does not feature the probability of success because payment is contingent on project success and, therefore, any promise up to the amount generated by successful projects is feasible. Second, if divesting is interior, it is optimal to equalize the marginal effect of divestment across projects which implies that s is the same for all borrowers.

The bank's payoff associated with any feasible contract is

$$Ap(k-s)^\alpha + (l-k) + s\delta - \phi c^E - (1-\phi)(c_1^L + pc_2^L).$$

The bank can set $c_1^L = 0$ without any loss. Doing so weakens feasibility constraints and can be offset for the borrower by raising pc_2^L and or c^E if needed. This can be accomplished until all surplus is exhausted and so a positive $c_1^L = 0$ does not augment the set of feasible payoffs. In addition, until I introduce additional frictions below there is clearly no need to overinvest in capital only to divest it later at a rate of $\delta \leq 1$, and so for now set $s = 0$ without any loss.

It is then natural for the bank to consider the expected output maximizing level of capital

$$k^* = \arg\max_k Apk^\alpha + l - k = (p\alpha A)^{\frac{1}{1-\alpha}}.$$

If

$$\phi w \leq l - k^*$$

then the bank can promise $c_E = pc_2^L = w$ i.e. provide the borrower with full intertemporal insurance while delivering the value of their outside option. And it is optimal to do so since the bank cannot promise less than w in expected payoff to the borrower and would have to offer strictly more expected payoff if the borrower's consumption profile featured any risk.

(Note also for future reference that, regardless of how high l is, the bank can profitably offer a contract to the borrower only if

$$Ap(k^*)^\alpha \geq w.$$

Indeed, if that condition didn't hold, the bank would be better off storing their entire capital since they must promise an expected payoff of at least w to borrowers given that they are risk-averse.)

If, on the other hand, $\phi w > l - k^*$ then, at the profit maximizing solution, either $k < k^*$ or the borrower is not receiving full insurance. Standard variational arguments show that, in fact, both levers are used by the bank. When $k = k^*$ the marginal consequence of a slight decrease in capital is zero while, since in that case insurance is imperfect, there is a first order gain in insurance for the borrower which the bank can use to decrease the expected payoff to them. And the same argument holds in reverse when insurance is full.

4 Bank liquidity markets

Assume now that liquidity shocks are stochastic and heterogenous across locations in the sense that the fraction of borrowers that experience a liquidity shocks is drawn at date 1 at each location from a distribution Φ with continuous density $f > 0$ on $[0, 1]$ and mean $\bar{\phi} = \int \phi f(\phi) d\phi$. Once again, a law of large number holds so that liquidity risk is purely idiosyncratic and it is known ex ante that the average liquidity shock across locations will be $\bar{\phi}$.

Assume further that banks can now borrow or lend liquidity at date 1 at a net rate r that will be determined in equilibrium. The economy also contains agents willing to supply liquidity $S(r)$ when the rate is $r \geq 0$. In addition, S is increasing in r and $S(r) = 0$ when $r \leq 0$, which amounts to assuming that participants in liquidity markets have access to the same zero-return storage technology banks do. Then any equilibrium in which $r > 0$ is such that banks are net demanders of wholesale funding. Think of this new source of funds as representing non-bank liquidity lenders which, in practice, include corporations, government-sponsored enterprises, money-market funds, pension funds, and other institutional lenders.

Let $b(\phi)$ be the amount of liquidity the bank borrows or (if negative) lends given its realized liquidity shock. The optimal contracting problem becomes finding:

1. $k \geq 0$
2. $c^E : [0,1] \mapsto I\!R_+$
3. $c^L : [0,1] \mapsto I\!R_+$
4. $s : [0,1] \mapsto I\!R_+$
5. $b : [0,1] \mapsto I\!R$

to maximize:

$$\int_0^1 \left[\left(Ap\,(k - s(\phi))^\alpha + (l-k) + s(\phi)\delta - \phi c^E(\phi) - b(\phi)r - (1-\phi)pc^L(\phi)\right)\right] f(\phi)d\phi$$

subject to, for all $\phi \in [0,1]$:

$$k \leq l, \tag{4.1}$$

$$\int_0^1 \left[\phi \ln c^E(\phi) + (1-\phi)\ln\left(pc^L(\phi)\right)\right] f(\phi)d\phi \geq \ln w, \tag{4.2}$$

$$\phi c^E(\phi) \leq l - k + b(\phi), \tag{4.3}$$

$$c^L(\phi) \leq A\,(k - s(\phi))^\alpha\,, \tag{4.4}$$

$$s(\phi) \leq k, \tag{4.5}$$

$$c^E(\phi) \leq pc^L(\phi). \tag{4.6}$$

Notice that I have once again set $c_1^L = 0$ without any loss of generality.

This is a Mayer problem with both standard and isoperimetric (integral) constraints and Filipppov's existence theorem implies that a solution exists (see theorem 4.3 in Cesari [4].) Furthermore, the solution is generally unique since the choice set is convex while the objective is jointly strictly concave and in k and s.

The problem above implies, inter alia, liquidity policies $b(\phi, r)$ as a function of the realized rate ϕ of early consumer arrival and the price r of wholesale funds. Equilibrium in the market for liquidity means

$$\int b(\phi, r) f(\phi) d\phi \leq S(r),$$

with equality if $r > 0$. In and only in the case in which $r = 0$ banks who are net suppliers of liquidity are indifferent between lending to other banks and storing their funds and so equilibrium only requires an excess supply of funds.

In any equilibrium with $r = 0$, banks will first attempt to sustain $k = k^*$ while providing full insurance to borrowers, i.e. $c^E(\phi) = pc^L(\phi) = w$ and $s(\phi) = 0$ for all ϕ, which maximizes both total surplus and expected bank profits. To implement this, set

$$b(\phi) = \phi w - (l - k^*).$$

Doing so meets date 1 feasibility by construction as long as $k^* \leq l$. As for market clearing:

$$\int b(\phi) d\Phi = \int [\phi w - (l - k^*)] d\Phi = \bar{\phi} w - (l - k^*).$$

Therefore, at $r = 0$ we have an excess supply of liquidity according to the above proposal if and only if

$$\bar{\phi} w - (l - k^*) \leq 0 = S(0).$$

In summary:

Remark 1. *An interbank market equilibrium exists with $r = 0$ if and only if*

$$\bar{\phi}w \leq l - k^*.$$

When, on the other hand, $\bar{\phi}w > l - k^*$ then only equilibria with $r > 0$ can exist. To characterize equilibria in general, note first that constraint (4.5) cannot bind since the marginal product of capital becomes unboundedly large near zero. On the other hand, participation constraint (4.2) must bind at any optimum and I will write $\lambda > 0$ for the shadow price of that constraint. The following results provides a complete characterization of optimal bank policies. Since the proof is lengthy, I relegate it to the appendix.

Proposition 2. *Optimal bank policies satisfy*

$$k(r) = \min\left\{\left(\frac{Ap\alpha}{1+r}\right)^{\frac{1}{1-\alpha}}, l\right\}$$

and, for almost all $\phi \in [0, 1]$,

1. $c^E(\phi, r) = \frac{\lambda}{(1+r)}$,
2. $c^L(\phi, r) = \min\left\{\frac{\lambda}{p}, Ak(r)^\alpha\right\}$,
3. $s(\phi, r) = 0$,
4. $b(\phi, r) = \frac{\phi\lambda}{1+r} + k(r) - l$.

Furthermore, if condition

$$w(1+r)^{\bar{\phi}+\frac{1}{1-\alpha}} \leq \min\left(A^{\frac{1}{1-\alpha}}(p\alpha)^{\frac{\alpha}{1-\alpha}}, Al^\alpha\right) \tag{4.7}$$

holds then $\lambda = w(1+r)^{\bar{\phi}}$ and, for all ϕ,

1. $c^E(\phi, r) = \frac{w}{(1+r)^{1-\bar{\phi}}}$,
2. $c^L(\phi, r) = \frac{w(1+r)^{\bar{\phi}}}{p}$

3. $b(\phi, r) = \frac{\phi w}{(1+r)^{1-\phi}} + k(r) - l.$

where $\bar{\phi} = \int \phi f(\phi) d\phi.$

Proof. See appendix. □

Condition (4.7) holds for instance when A is high enough given other parameters, or if w or r is low enough given other parameters. In the sequel, I will often find it convenient to focus on this more explicit case and I will do so for concreteness by assuming that A is as large as it needs to be to make those policies optimal.

Once again note that it is always the case at the optimum contract that $c^E(\phi, r) \leq pc^L$, with a strict inequality whenever $r > 0$, so that late consumers do not want to misrepresent their type. I will dispense with this observation henceforth because it holds in all those upcoming versions exactly as it does here.

The optimal level $k(r)$ of investment in risky projects is unique as long as $\delta < 1$. When $\delta = 1$, the policies listed in the proposition remain optimal but, given that divesting is costless it also becomes optimal when $r = 0$ to over-invest in projects early by setting, say, $k' > k(r)$ but then divest $k' - k(r)$ once liquidity shocks are realized. Consumption, output, and interbank lending policies, hence welfare, are all unaffected by moving to the limiting case, they all remain uniquely defined and as stated in the proposition.

Bank liquidity demand is strictly monotone in r both because higher rates tilt the consumption profile towards late consumption and because capital formation falls when r rises. As a result, a unique equilibrium exists and the cost of liquidity is positive if and only if the average liquidity shock is high enough.

5 Private information

This paper is about the informational content of bank liquidity markets. The mechanism I want to formalize is that bank liquidity demand may be correlated with credit quality shocks.

To generate this, assume that a fraction $\eta > 0$ of borrowers whose project is bad receives a perfectly informative signal of project quality immediately after capital is installed.

Given that consumption at date 2 is zero if the project fails, all borrowers who discover that their project is bad ask for early consumption and it becomes impossible for banks to distinguish between liquidity and fundamental shocks.[4] Resource feasibility constraints become for all ϕ:

$$\begin{aligned}(\phi + (1-\phi)(1-p)\eta)c^E(\phi) &\leq l - k + b(\phi), \\ c^L(\phi) &\leq A\left(k - s(\phi)\right)^{\alpha}.\end{aligned}$$

Indeed, borrowers who get the early consumption shock always ask for early consumption. Among the borrowers who do not experience that shock, mass $\eta(1-p)$ learn that their project is bad and also ask for early consumption even though they are indifferent between consuming early or late.

In principle, the bank could attempt to induce truth-telling by promising the same early consumption regardless of the signal value. The assumption I am making here is that the bank does not have the ability to commit to paying the borrower once it learns that the project is worthless.

Among the borrowers who opt for late consumption, that is among the borrowers who do not receive a signal, the average success rate is:

$$\hat{p} = \frac{p}{p + (1-p)(1-\eta)} > p.$$

From the point of view of the bank, however, since no projects are shut down early, the average success rate among surviving projects remains p.

[4] This is the specific juncture of the construction where I make use of the fact that preferences satisfy Inada conditions. This guarantees that promised consumption for people who report that they are early consumers is strictly positive. This means that agents who know their project is bad cannot be induced to report their signal.

For its part, the arrival rate of liquidity shocks is now

$$\hat{\phi} \equiv \phi + (1-\phi)\eta(1-p) \in [\eta(1-p), 1],$$

with the corresponding change of measure

$$\hat{f}(\hat{\phi}|p) \equiv f\left(\frac{\hat{\phi} - \eta(1-p)}{1-\eta(1-p)}\right),$$

for all $\hat{\phi} \in [\eta(1-p), 1]$. For example, if f is uniform on $[0,1]$ then $\hat{f}$ is uniform on $[\eta(1-p), 1]$. In the general case, as long as $\eta > 0$, $\hat{f}$ dominates f in a first-order sense.

The borrower's participation constraint becomes:

$$\int_0^1 \left[\hat{\phi} \ln c^E(\hat{\phi}) + (1-\hat{\phi}) \ln\left(\hat{p} c^L(\hat{\phi})\right)\right] \hat{f}(\hat{\phi}|p) d\hat{\phi} \quad \geq \quad \ln w.$$

In other words, the nature of the participation constraint is unchanged. The economy features uniformly higher liquidity shocks once private information is introduced. It also features a higher payoff rate for consumers who claim to be late consumers since some bad project owners get the signal to quit early.

The bank's objective function becomes:

$$\int_0^1 \left(\hat{A}\hat{p}\left(k - s(\hat{\phi})\right)^\alpha + (l-k) + s(\hat{\phi})\delta - b(\hat{\phi})r - \hat{\phi}c^E(\hat{\phi}) - (1-\hat{\phi})\hat{p}c^L(\hat{\phi})\right) \hat{f}(\hat{\phi}|p) d\hat{\phi},$$

where $\hat{A} = \frac{Ap}{\hat{p}}$.

If p is known or idiosyncratic therefore, the problem with private signals is a mere relabeling of the problem I studied in the previous section. It follows that policy functions and equilibria all have the same form as they do when all information is public. Private information does result in more early consumption and potentially more reliance on wholesale funding but that is entirely captured in the level of interest rates. Higher rates, in turn and per proposition 2, imply lower investment and output.

6 Aggregate shocks and Radner equilibria

The foregoing analysis assumes that the quality p of projects is known. When that quality is subject to aggregate shocks, an entirely new set of considerations emerges. Most obviously, both the liquidity shocks banks experience and the cost r of funds convey information about the state of the economy. This creates an opportunity for capacity adjustments at date 1.

To study this, consider an extension of the economy with private information in which p is drawn from a known distribution ν with full support on $[p_L, p_H]$ where $0 < p_L < p_H \leq 1$ at date 0. Write $\bar{p} = \int p d\nu$ for the expected value of the aggregate shock. The shock is aggregate in the sense that, unlike liquidity shocks, it is common to all locations.

No bank observes p directly but they each receive two pieces of information about aggregate conditions. First, as we discussed in the previous section, the arrival rate at a particular location is correlated with aggregate conditions. This signal, alone, would cause a particular bank to update their views of aggregate conditions to a distribution g of aggregate shocks that satisfies

$$g(p|\hat{\phi}) \propto \nu(p)\hat{f}\left(\hat{\phi}|p\right). \tag{6.1}$$

In this application of Bayes' rule I assume that all agents know ν and use it as their prior beliefs, and that all agents understand the nature of $\hat{f}$. Given the strictly decreasing relationship between aggregate shocks and the arrival rate of early consumers, a low realization of the arrival rate causes a posterior distribution of aggregate shocks that dominates in a first-order sense the posterior distribution that results from high arrival rates. A high rate of arrival causes banks to form pessimistic beliefs about aggregate shocks. In the next section, this updating step plays a key role and I will discuss it in detail.

Banks also receive a second, common signal of aggregate conditions, namely the price r of liquidity at the interim date. This opens the possibility that a Full Communication

Equilibrium (FCE) in the sense of Radner [28] may exist. Denote by $\mathcal{F}$ the set of distributions consistent with update condition (6.1). A given p induces a unique element of $\mathcal{F}$ and the map between p and $\hat{F} \in \mathcal{F}$ is a strictly monotonic function in the sense of first-order dominance. In keeping with Radner, I will call a FCE fully revealing if the mapping from p to the equilibrium price of liquidity is one-to-one. In such an equilibrium, observing market prices is equivalent to and sufficient for knowing $\hat{F}$ hence p itself. A FCE equilibrium, then, is a mapping $\hat{r}(p)$ from each possible aggregate shock $p \in [p_L, p_H]$ to liquidity prices which is one-to-one and such that markets clear for all p when banks understand this mapping. Since $\hat{r}$ is one-to-one in any FCE I will also write $p(r)$ for the aggregate shock that corresponds to each possible $r \in \hat{r}\left([p_L, p_H]\right)$.[5]

As usual, introducing aggregate shocks raises a number of technical difficulties in an optimal contracting environment. To simplify the analysis of the resulting economy, I will focus on equilibria in which wholesale rates may depend on the aggregate state but do not depend on extrinsic sources of uncertainty such as sunspots.

Denote by h the density of the distribution of interest rates induced by the forecast $\hat{r}$ convoluted with the distribution ν of aggregate shocks. Formally, h is the derivative of

$$\mathcal{H}(r) \equiv \int_{\{p:\hat{r}(p)\leq r\}} d\nu(p).$$

Assuming that $\hat{r}(p)$ is monotonic and continuous, two features of the rate schedule I will establish later as I build a FCE, h is well defined and continuous. Economically, each p induces a unique r in FCE and changes in p impact interest rate conditions continuously.

In keeping with the previous section, write

$$\hat{\phi} \equiv \phi + (1-\phi)\eta(1-p) \in [0,1]$$

[5] Unlike in Radner's seminal environment the distribution of private signals in my model is continuous. This raises additional technical difficulties which Sun et al. [31] deal with in details. For instance, in establishing injectivity I need to consider two distributions equivalent if they are almost surely the same.

for arrival rates given the distribution of the signal, with $\hat{f}(\bullet|p)$ the corresponding distribution of liquidity shocks. Then, in a FCE, optimal lending contracts are lists

$$\left\{k, s(\hat{\phi}, r), c^E(\hat{\phi}, r), c^L(\hat{\phi}, r), b(\hat{\phi}, r) : \hat{\phi} \in [0,1], r \in \hat{r}([p_L, p_H])\right\}.$$

Feasibility imposes, for all $(\hat{\phi}, r)$,

$$\begin{aligned} \hat{\phi} c^E(\hat{\phi}, r) &\leq l - k + b(\hat{\phi}, r), \\ \text{and } c^L(\hat{\phi}, r) &\leq A(k - s(\hat{\phi}, r))^\alpha, \end{aligned}$$

while borrower participation requires

$$\int_{r_L}^{r_H} \int_0^1 \left(\hat{\phi} \ln c^E(\hat{\phi}, r) + (1 - \hat{\phi}) \ln[\hat{p}(r)\, c^L(\hat{\phi}, r)]\right) \hat{f}(\hat{\phi}|p(r))h(r) d\hat{\phi} dr \geq \ln w,$$

where

$$\hat{p}(r) = \frac{p(r)}{p(r) + (1 - p(r))(1 - \eta)}.$$

The plan also needs to be such that late borrowers who do not receive a signal have no incentives to claim to be early consumers which means, for all $(\hat{\phi}, r)$:

$$c^E(\hat{\phi}, r) \leq \hat{p}(r) c^L(\hat{\phi}, r).$$

The bank's objective function now reads as follows:

$$\int_{r_L}^{r_H} \int_0^1 \left[Ap(r)(k - s(\hat{\phi}, r))^\alpha + \delta s(\hat{\phi}, r) + l - k - rb(\hat{\phi}, r) - \hat{\phi} c^E(\hat{\phi}, r) - (1 - \hat{\phi})\hat{p}(r) c^L(\hat{\phi}, r)\right] \hat{f}(\hat{\phi}|p(r))h(r) d\hat{\phi} dr.$$

In any FCE, banks discover the exact nature of aggregate conditions $p \in [p_L, p_H]$ and this gives them an opportunity to divest capital to the safe technology if p proves low. For all

$(k,p) \in [0,l] \times [p_L, p_H]$, define

$$H(k,p) = \max_{s\in[0,k]} Ap(k-s)^{\alpha} + s\delta.$$

Note that H is strictly increasing in both k and p and differentiable with respect to both arguments. It is also strictly concave in k up to $k^*(p,\delta) = \left(\frac{Ap\alpha}{\delta}\right)^{\frac{1}{1-\alpha}}$ and linear with slope δ past that threshold. Most importantly – this is a version of Hotelling's lemma since in this context p is isomorphic to an output price – H is convex in p and, in fact, strictly so for sufficiently large spreads in p. This is because while p enters linearly in the profit function, changes in p can also lead to updates in s so that the impact of changes in p on output becomes super-linear. (Figure 5 in the next section shows the shape of H in the parametric example I will use to illustrate the impact of ample liquidity.)

To expand on this key point, let

$$s(r) = \arg\max_{s\in[0,k]} Ap(r)(k-s)^{\alpha} + s\delta.$$

Unconstrained, the bank would always select divestment policy $s(r)$ upon discovering r hence, in an FCE, p. The one limit the bank potentially faces in implementing this policy is the constraint that $c^L(\hat{\phi}, r) \leq A(k - s(\hat{\phi}, r))^{\alpha}$. Provided A is high enough (in a sense my next lemma will make precise) given other parameters, that constraint cannot bind hence the bank is able to attain H in all possible states at the interim date. To simplify the description of FCEs I will focus on that case.

One immediate simplification that follows is that the bank's objective function becomes:

$$\int_{r_L}^{r_H}\int_0^1 \left[H(k,p(r)) + (K-k) - \hat{\phi}c^E(\hat{\phi},r) - r\,b(\hat{\phi},r) - (1-\hat{\phi})[\hat{p}(r)\,c^L(\hat{\phi},r)]\right]\hat{f}(\phi|p(r))h(r)d\phi dr.$$

With this notation at hand, I can now fully characterize optimal policies in any FCE.

Lemma 3. *In any FCE, provided A is high enough given other parameters, optimal lending* *contracts satisfy:*

1. $\int_{r_L}^{r_H} H_1(k, p(r))h(r)dr \geq \int_{r_L}^{r_H}(1+r)h(r)dr$, *with equality if* $k < l$;
2. $s(\hat{\phi}, r) = \arg\max_{s\in[0,k]} Ap(r)(k-s)^{\alpha} + s\delta$;
3. $c^E(\hat{\phi}, r) = \frac{\lambda}{(1+r)}$;
4. $c^L(\hat{\phi}, r) = \frac{\lambda}{\hat{p}(r)}$;
5. $b(\hat{\phi}, r) = \frac{\hat{\phi}\lambda}{1+r} + k - l$;

where $\lambda > 0$ *is the shadow price of the borrower participation constraint.*

Proof. See appendix. □

A key part of this result is that in a FCE capital formation does not depend on the arrival rate $\hat{\phi}$ because the opportunity cost of investment is pinned down by r and because once r is known so is the aggregate shock. Nothing further needs to be inferred from the arrival rate for the purpose of optimal investment. Another important consequence of this characterization is that b is strictly monotonic in the arrival rate, which in turn is strictly decreasing in p since more borrowers realize that they are sitting on a bad investment. This means that an FCE exists in this environment under mild conditions, as I will now establish.

To start the search for an FCE, start with a guess $\bar{r}$ for the average rate of return $\int_{r_L}^{r_H}(1+r)h(r)dr$ in equilibrium. Because in FCE the mapping from r to p is one-to-one, we must have:

$$\int_{r_L}^{r_H} H_1(k, p(r))h(r)dr = \int_{p_L}^{p_H} H_1(k, p)d\nu(p)$$

and so a guess for $\bar{r}$ implies a unique k in FCE by the first item of lemma 3. Furthermore, the mapping from the guess $\bar{r}$ to k is continuous and decreasing.

Now, at the interim stage, having pinned down k, given any realization of r, recall that a liquidity market equilibrium given p requires

$$\int b(\hat{\phi}, \hat{r}(p))d\hat{f}(\hat{\phi}|p(r)) \leq S(\hat{r}(p)), \tag{6.2}$$

with equality whenever $\hat{r}(p) > 0$. In any FCE, the bank's forecasting function p must be correct so consider any pair (p_1, p_2) of aggregate shocks $p_L < p_1 < p_2 < p_H$, which banks are able to perfectly deduct from liquidity market conditions. Let $\hat{r}(p_1)$ and $\hat{r}(p_2)$ be the only non-negative values of the price of liquidity that meet condition (6.2) given the unique k implied by $\bar{r}$, after fixing $p(r)$ at p_1 and p_2, respectively. There is only one such value per aggregate shock because, after fixing p and given the k implied by my guess for $\bar{r}$, supply is strictly increasing in r while demand is strictly decreasing due to the strictly monotonic relationship between aggregate shocks and early arrival rate. In other words, the demand schedule for liquidity is uniformly higher at p_1 than at p_2 because $\hat{f}(\bullet|p_1)$ first-order stochastically dominates by $\hat{f}(\bullet|p_2)$. Therefore, as long as $\hat{r}(p_1) > 0$, we must have $\hat{r}(p_1) > \hat{r}(p_2)$.

A first condition for our initial guess for $\bar{r}$ to be correct and lead to a FCE is that

$$\int_{p_L}^{p_H} r(p)d\nu(p) = \bar{r}.$$

If

$$\int_{p_L}^{p_H} r(p)d\nu(p) \geq \bar{r}$$

I can keep raising my guess for $\bar{r}$ which lowers k, eventually strictly. This can only decrease $r(p)$ for all p and we can do this continuously until the desired fixed point is found. Now note that if we start with a guess of $\bar{r} = 0$ we must in fact have $\int_{p_L}^{p_H} r(p)d\nu(p) \geq 0$ so that there has to be such a fixed point.

A second requirement for this process to produce a FCE is that the resulting interest rate schedule be strictly monotonic. As discussed above, this must be the case if $\hat{r}(p) > 0$ for all p. Making the productivity A of projects high enough given other parameters is one parametric assumption that guarantees that $\hat{r}(p) > 0$. Indeed, with A high enough k must approach and in fact eventually exceed l which by proposition 3 means $b(\hat{\phi}, r) > 0$ for all $\hat{\phi}$ as long as $w > 0$ hence $\lambda > 0$. The following result, which I prove in the appendix, summarizes this discussion.

Proposition 4. *An FCE exists provided A is high enough given other parameters.*

Proof. See appendix. □

7 The real consequences of ample liquidity

The FCE with aggregate shocks is characterized by an interval $[\hat{r}(p_H), \hat{r}(p_L)]$ of interest rates. Banks infer aggregate conditions from liquidity market conditions and, as a result, are able to adjust the quantity of loans they make and mitigate losses when conditions warrant adjustments.

Now introduce a central bank who, at $t = 0$, can add to bank liquidity endowments as much as they wish. Denote by l^* the liquidity level they set.[6] At the same time, the central bank offers an administered rate r^* on reserve deposits (in other words, on reserves held at the central bank) at $t = 1$.

Assume that banks forecast that the price of liquidity at $t = 1$ will be r^* with probability one. This is the equilibrium that the central bank seeks to support under an ample liquidity framework. Indeed, the central bank's goal under this framework is to control the market price of bank liquidity simply by changing the rate it offers to pay on reserves.[7] The goal of this section is to show that the central bank can always implement this outcome and characterize the resulting equilibrium.

When banks forecast that the interest rate will be constant at r^*, and when that forecast materializes, they no longer get any information from the realization of the cost of liquidity. On the other hand, they still get information from observing the arrival rate $\hat{\phi}$ of early consumers. Given this arrival rate at a particular location, they update their beliefs to a distribution $\hat{\nu}$ with support in $[p_L, p_H]$. As I discussed in the previous section, letting g be

[6] An equivalent interpretation is that the central bank can transfer τ^* in reserves to banks at $t = 1$ where $\tau^* = l^* - l$.

[7] An other way for the central bank to implement that equilibrium is to borrow and lend as much as needed via open-market operations or liquidity facilities at $t = 1$ so that markets clear at r^*. The point of the ample reserve approach, however, is to accomplish the same market outcome by saturating banks with liquidity.

the density function of this distribution and given arrival rate $\hat{\phi}$, Bayes' rule implies that

$$g(p|\hat{\phi}) \propto \nu(p)\hat{f}\left(\hat{\phi}|p\right),$$

and I will write

$$\bar{p}(\hat{\phi}) = \int p dg(p|\hat{\phi})$$

for the expected value of the aggregate shock following the arrival rate. Given the strictly decreasing relationship between aggregate shocks and the arrival rate of early consumers, a low realization of the arrival rate causes a posterior distribution of aggregate shocks that dominates in a first-order sense the posterior distribution that results from high arrival rates. As a consequence, $\bar{p}(\hat{\phi})$ is strictly decreasing in the arrival rate. I will also write

$$\bar{p} = \int \bar{p}(\hat{\phi})\hat{f}(\hat{\phi}|p)d\nu(p)$$

for the unconditional expected quality of projects before banks discover the arrival rate.

To illustrate these ideas, assume for example that both ν is uniform on $[p_L, p_H]$ while f are uniform on $[0, 1]$. In that case Bayes' rule boils down to

$$g(p|\hat{\phi}) \propto \frac{1}{1-(1-\eta)p}$$

with $p > 1 - \frac{\hat{\phi}}{\eta}$. It follows that

$$g(p|\hat{\phi}) = \begin{cases} \dfrac{\eta}{\ln(1-\eta(1-p_H)) - \ln(1-\eta(1-p_L))} \cdot \dfrac{1}{1-\eta(1-p)}, & p \in [p_L, p_H] \quad \text{if } \hat{\phi} \geq \eta(1-p_L) \\ \dfrac{\eta}{\ln(1-\eta(1-p_H)) - \ln(1-\hat{\phi})} \cdot \dfrac{1}{1-\eta(1-p)}, & p \in \left[1 - \dfrac{\hat{\phi}}{\eta},\ p_H\right] \quad \text{if } \hat{\phi} < \eta(1-p_L) \end{cases}.$$

The top panel of figure 3 illustrates those calculations in the case in which $\eta = 10\%$ and for various possible levels of arrival rates. The bottom panel shows the corresponding $\bar{p}(\hat{\phi})$. In

Figure 3: Bank beliefs upon observing arrival rates

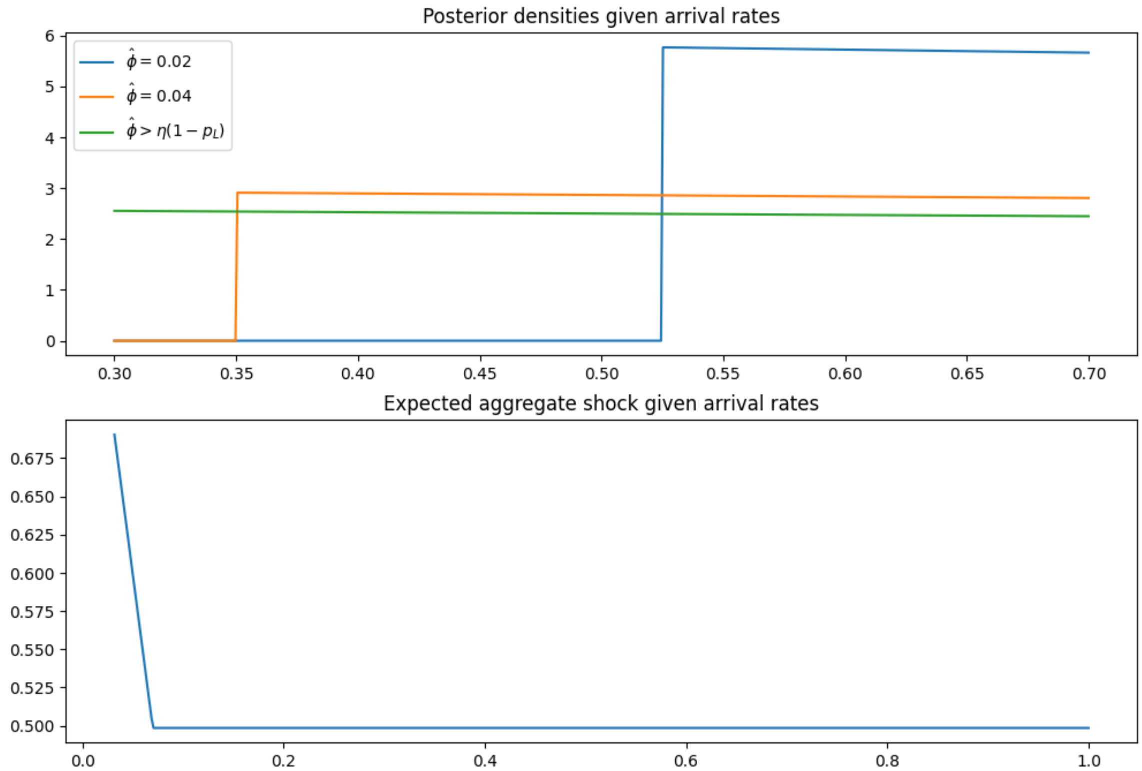

Sources and notes: This assumes that the distribution ν of aggregate shocks is uniform on $[p_L, p_H] = [0.3, 0.7]$ and that the distribution f of arrival rates is uniform on $[0, 1]$. Fraction $\eta = 10\%$ of borrowers with bad projects receive a signal at $t = 0$.

this simple case, when $\hat{\phi} > \eta(1 - p_L)$, additional increases in arrival rate do not reveal more information. Unusually low arrival rates, on the other hand, contain a lot of information about the aggregate state of the economy and cause the bank to form optimistic beliefs. All told, local arrival rates provide at best coarse information about aggregate credit conditions.

With r^* set and publicly known, banks select $\left\{k, s(\hat{\phi}), c^E(\hat{\phi}), c^L(\hat{\phi}), b(\hat{\phi}) : \hat{\phi} \in [0, 1]\right\}$ to maximize

$$\int_0^1 \left\{ \int_0^1 \left[Ap(k - s(\hat{\phi}))^\alpha + \delta s(\hat{\phi}) + l^* - k - \hat{\phi} c^E(\hat{\phi}) - r^* b(\hat{\phi}) - (1 - \hat{\phi})\hat{p}(p) c^L(\hat{\phi}) \right] g(p|\hat{\phi}) dp \right\} \hat{f}(\hat{\phi}) d\hat{\phi}$$

subject to:

$$\int_0^1 \left\{ \int_0^1 \left[\hat{\phi} \ln c^E(\hat{\phi}) + (1-\hat{\phi}) \ln \left[\hat{p}(p) c^L(\hat{\phi}) \right] \right] g(p|\hat{\phi}) dp \right\} \hat{f}(\hat{\phi}) d\hat{\phi} \geq \ln w$$

and, for all $\hat{\phi}$,

$$\begin{aligned} k &\leq l^*, \\ \hat{\phi} c^E(\hat{\phi}) &\leq b(\hat{\phi}) + K - l^*, \\ c^L(\hat{\phi}) &\leq Ak \left(\hat{\phi} - s(\hat{\phi}) \right)^\alpha, \\ c^E(\hat{\phi}) &\leq \hat{p}(\hat{\phi}) c^L(\hat{\phi}). \end{aligned}$$

As before, I use hat notation

$$\hat{p}(p) \equiv \frac{p}{p + (1-p)(1-\eta)}$$

for the expected success rate of late consumers who do not receive the signal. Abusing notation somewhat, I also write

$$\hat{p}(\hat{\phi}) \equiv \int \hat{p}(p) dg(p|\hat{\phi}),$$

for the expected success rate of projects for late consumers when the arrival rate is $\hat{\phi}$. With all this notation at hand, similar algebra as in previous sections gives:

Proposition 5. *Under an ample liquidity framework with policy rate $r^* \geq 0$ and as long as A is high enough, bank policies satisfy for almost all $\hat{\phi} \in [0,1]$,*

1. $\int_0^1 H_1(k, \bar{p}(\hat{\phi})) d\hat{\phi} \geq 1 + r^*$, *with equality if $k < l^*$;*
2. $s(\hat{\phi}) = \arg\max_{s \in [0,k]} A\bar{p}(\hat{\phi})(k-s)^\alpha + s\delta$;
3. $c^E(\hat{\phi}) = \frac{\lambda}{(1+r^*)}$,
4. $c^L(\hat{\phi}) = \frac{\lambda}{\bar{p}(\hat{\phi})}$,
5. $b(\hat{\phi}) = \frac{\hat{\phi}\lambda}{1+r^*} + k - l^*$,

where $\lambda > 0$ is the shadow price of the borrower participation constraint.

Proof. See appendix. □

Equilibrium in the interim market for liquidity requires

$$\int b(\hat{\phi}, r)\hat{f}(\hat{\phi})d\hat{\phi} \leq S(r),$$

with equality if $r > r^*$. The key distinction with the previous version of this condition is that now the supply of liquidity by banks vanishes at any $r < r^*$ since they know that they can park liquidity at the central bank at that rate. It follows that $r = r^*$ is an equilibrium given a realized value p of the aggregate shock if and only if

$$\int b(\hat{\phi}, r^*)\hat{f}(\hat{\phi}|p)d\hat{\phi} \leq S(r^*).$$

Since $b(\hat{\phi}) = \frac{\hat{\phi}\lambda}{1+r^*} + k - l^*$, by making l^* high enough, the inequality above hold for all p. In fact, overkill though this would be, the central bank could drive $\int b(\hat{\phi}, r^*)\hat{f}(\hat{\phi}|p)d\hat{\phi}$ to zero or negative for all p by making l^* large enough.

To economize as best it can on the cost of paying interest on reserves, the central bank can pick the lowest value of l^* such that the inequality holds even when $p = p_L$. Given that all policy functions are continuous, l^* must be such that

$$\int b(\hat{\phi}, r^*)\hat{f}(\hat{\phi}|p_L)d\hat{\phi} = S(r^*).$$

This is an ample reserve equilibrium, in the Federal Reserve Bank's language, in the sense that l^* is the smallest level of reserves that guarantees that the market price of liquidity is r^* regardless of demand conditions. In summary,

Proposition 6. *There exists a unique level $l^* > 0$ of bank liquidity such that:*

1. *The market for liquidity clears at r^* when bank liquidity supply is set to l^*;*

2. *At any liquidity level lower that l^*, no equilibrium exists in which the liquidity market clears at r^* for all possible values of the aggregate shocks.*

One practical consequence for a central bank of operating at exactly l^* is that any error in implementing or estimating the threshold must result in deviations from r^*. As long as the error is small, so is the deviation. The central bank can avoid this risk entirely by overshooting l^* by a wide margin – by making reserves abundant, that is – at the cost of a larger balance sheet and a higher interest bill. As I will discuss in the next section, this aspect of my model corresponds to the the way the Federal Reserve describes the distinction between ample and abundant reserves.

The key distinction between an ample reserve equilibrium and a FCE is that banks are now constrained to make divestment decisions based on the expected quality of projects $\bar{p}(\hat{\phi})$ rather than based on a precise knowledge of aggregate conditions. As a result, they operate on an inefficient scale in many cases. Ex-ante, this lowers the expected marginal product of capital hence results in lower investment. The following result studies a version of the economy in which this impact of suppressing information on output can be fully described.

Proposition 7. *Provided η is small enough, capital formation and average output fall when an ample reserve framework is implemented at a rate r^* as long as*

$$r^* \geq \int_{r_L}^{r_H} (1+r)h(r)dr.$$

Proof. Because I have assumed that the density f of preference shocks is continuous,[8] as $\eta \mapsto 0$, $\bar{p}(\hat{\phi}) \mapsto \bar{p}$ uniformly. In other words, the discovery of arrival rates leads to negligible updates in bank views about aggregate conditions. So, for any given k expected output (and average output across locations) converges almost surely to $H(k, \bar{p})$ when η becomes

[8]If the distribution of preference shocks is binary, for instance, even minute variations in observed arrival rates can lead to perfect updates.

negligible. To see why, this follows because banks choose s given k to maximize

$$\int Ap(k-s)^{\alpha}dp + s\delta = A\bar{p}(k-s)^{\alpha}dp + s\delta.$$

So all banks, in particular, choose the same divestment policy, namely the policy $s(\bar{p})$ that would be optimal if the aggregate shock was known to be $\bar{p}$. Then the law of large numbers implies that average output across locations is in fact $H(k,\bar{p})$.

On the other hand,

$$\begin{aligned} H(k,\bar{p}) &= \int_{p_L}^{p^H} H(k,\bar{p})dp \\ &\leq \int_{p_L}^{p^H} H(k,p)dp \end{aligned}$$

because H is convex in p. In fact, if $s(k,p) > 0$ for a positive measure of p, then the inequality is strict. (Economically, the simple intuition here is that the bank finds it profitable to adjust capacity for some realizations of the aggregate shock if it were able to discover it, as it does in a FCE.) So it follows that for any k the bank selects, expected output is always higher in an FCE than in an ample liquidity framework.

To argue that capital formation can only fall when an ample liquidity framework is implemented at $r^* \geq \int_{r_L}^{r_H}(1+r)h(r)dr$, note that for all $k \geq 0$:

$$\int_{p_L}^{p^H} H_1(k,p)dp \geq \int_{p_L}^{p^H} Ap\alpha k^{\alpha-1}dp = A\bar{p}\alpha k^{\alpha-1}.$$

So the marginal product of capital is uniformly higher at all k in a FCE than in an ample liquidity equilibrium. □

The intuition for this result is simple: because banks learn little about aggregate conditions they are not able to make adjustments that reflect aggregate conditions. As a a result the marginal product of capital from the point of view of $t = 0$ is uniformly lower in an ample liquidity equilibrium than in a FCE which depresses capital formation. In turn, output falls both because of this capital formation effect and because of the lack of adjustment at $t = 1$.

The condition that $r^* \geq \int_{r_L}^{r_H}(1+r)h(r)dr$ is necessary because when policy makers lower the rate further than what they are on average in an FCE they can lower the opportunity cost of capital at will. In order to measure the full benefit of those easy money policies, one would need to measure the consequences of all the distortions associated with creating the necessary liquidity.

While the premise that $\eta \mapsto 0$ leads to transparent mathematics, the same intuition is operative in cases when $\eta > 0$. To see this and by way of numerical illustration of my main message, take the parametrization of shocks I used in producing figure 3 (namely, ν is uniform on $[0.3, 0.7]$, arrival rates are uniform on $[0, 1]$, and $\eta = 10\%$) and further assume that $\alpha = 0.5$, $A = 10$, $l^* = 10$, $w = 4$, and $\delta = 1$. Finally, let S be linear with slope 2. None of the patterns I am about to display depend critically on these parameters but this constellation of parameters does meet with room to spare all the assumptions I have made in the previous two sections.

Given this parametrization, figure 4 shows the shape of output and adjustment decisions given an aggregate shock level in the right-hand panel and given a capital level in the right-hand panel. It confirms the concavity of H in capital, its convexity with respect to project quality, with convexity strict where $s > 0$. It also confirms that whenever capital is too high given the aggregate shock, banks redirect capital to storage.

To compute FCE equilibria, I follow the algorithm that guides the proof of proposition 4. I start with guesses for the shadow price λ and for average interest rates, calculate all policy functions and the rate schedule $\hat{r}$ that clears markets given those guesses, and then update those guesses until the borrower participation constraint is met exactly and average rates are correct in equilibrium. Finally, I verify that given other parameters the upper bound on c^L never binds. The code that implements this algorithm and produces all the figures I show in this paper is available for download at http://erwan.marginalq.com/index_files/wp.htm.

The main equilibrium object in any FCE is the interest rate schedule which must be strictly monotonic. Figure 6 shows that this is the case in this parameterization. In the

Figure 4: Output given capital and aggregate shocks

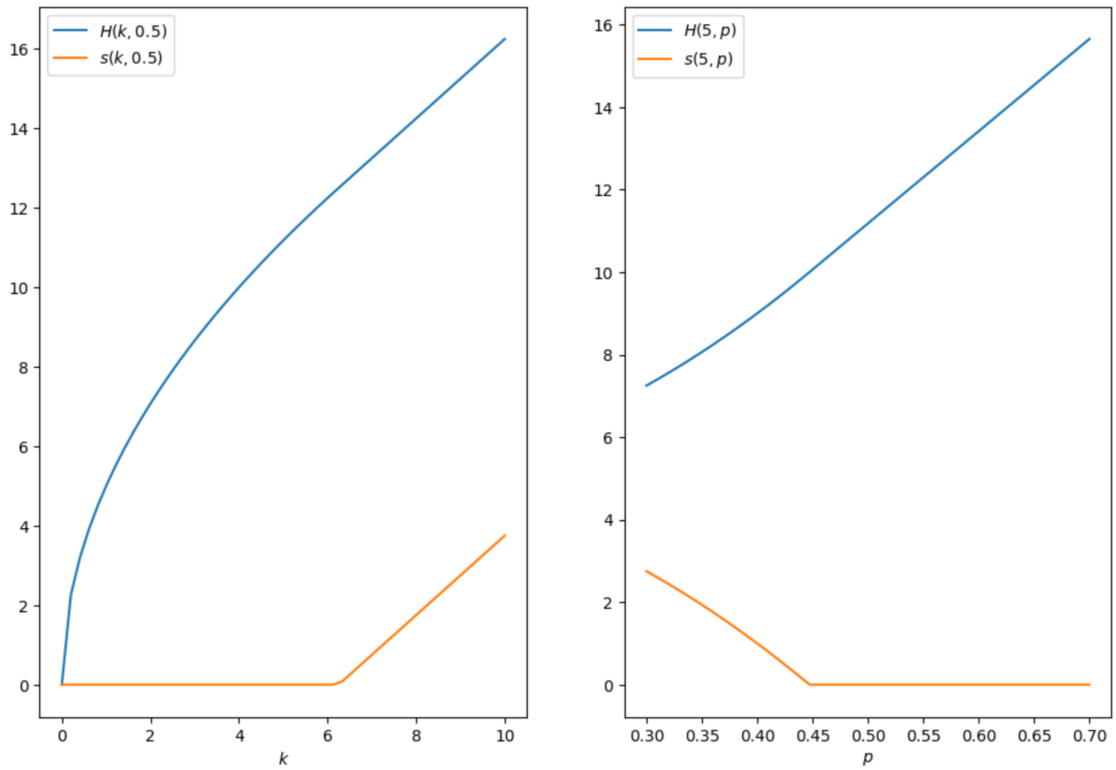

Sources and notes: This assumes $A = 10$, $\alpha = 0.5$, and $\delta = 1$. The left-hand panel holds p to 0.5 and varies the capital stock, while the right-hand panel sets $k = 5$ and varies the aggregate shock.

ample liquidity equilibrium, I assume that the central bank implements a constant rate equal to the unconditional average cost of liquidity in the FCE. Consistently with the intuition behind proposition 7, this implies that capital formation is significantly lower. Specifically, investment k falls by almost 40% in equilibrium from 8.01 with ample reserves to 5.84 in an FCE. This, alone, causes average output to be significantly higher in FCE than under an ample liquidity framework. This effect is compounded by the fact that in a FCE banks are able to fully infer the aggregate shock from liquidity markets at $t = 1$ and adjust scale accordingly.

To illustrate those effects, figure 6 tracks interest rates, output, and divestment decisions for a time-series of one hundred independently drawn aggregate shocks. Naturally, output is

Figure 5: Liquidity cost vs aggregate shocks

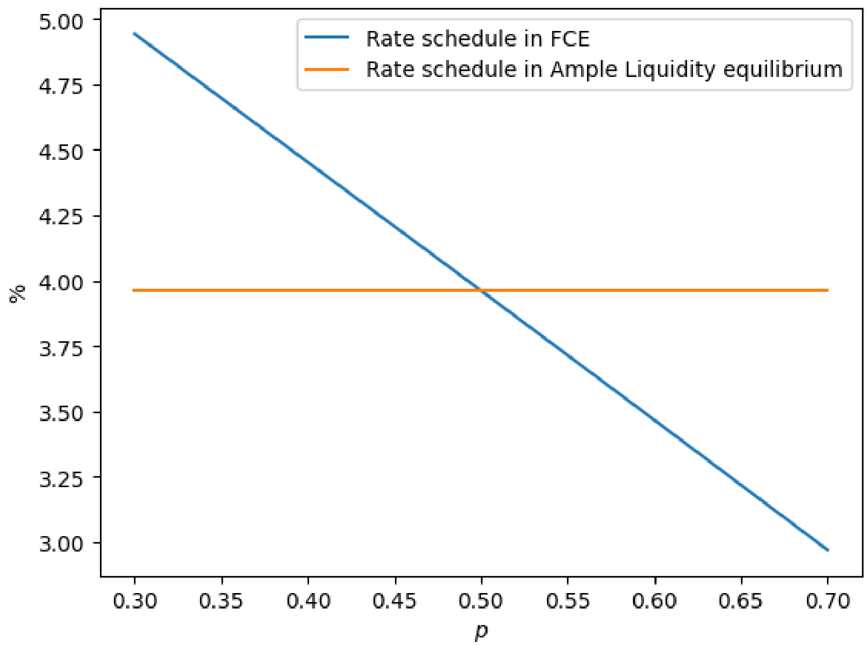

pro-cyclical while divestment decisions are counter-cyclical in FCE. Rates for their part are strictly counter-cyclical in FCE because as aggregate conditions worsens, more borrowers ask for their money early. This strict counter-cyclicality is precisely what enables banks to make optimal divestment decisions.[9]

In an ample liquidity equilibrium, rates are fully acyclical and banks only have arrival rates to inform their beliefs about aggregate conditions. As the bottom panel of figure 3 however, because the fraction of borrowers who receive a signal is low in this parameterization, banks do not receive sufficiently precise information to ever be confident that aggregate conditions are sufficiently bad to warrant divesting capital at $t = 1$. So, in equilibrium in this parameterization, they simply ride things out every period at the capital level they selected at $t = 0$. Both because capital formation is lower in the ample liquidity equilibrium than in

[9]Output inherits the lack of persistence of the driving aggregate shocks in this experiment. This, of course, is a feature of every single real business cycle model. Generating recessions and expansions of a given persistence requires aggregate shocks of basically the same persistence.

Figure 6: Aggregate shock simulation

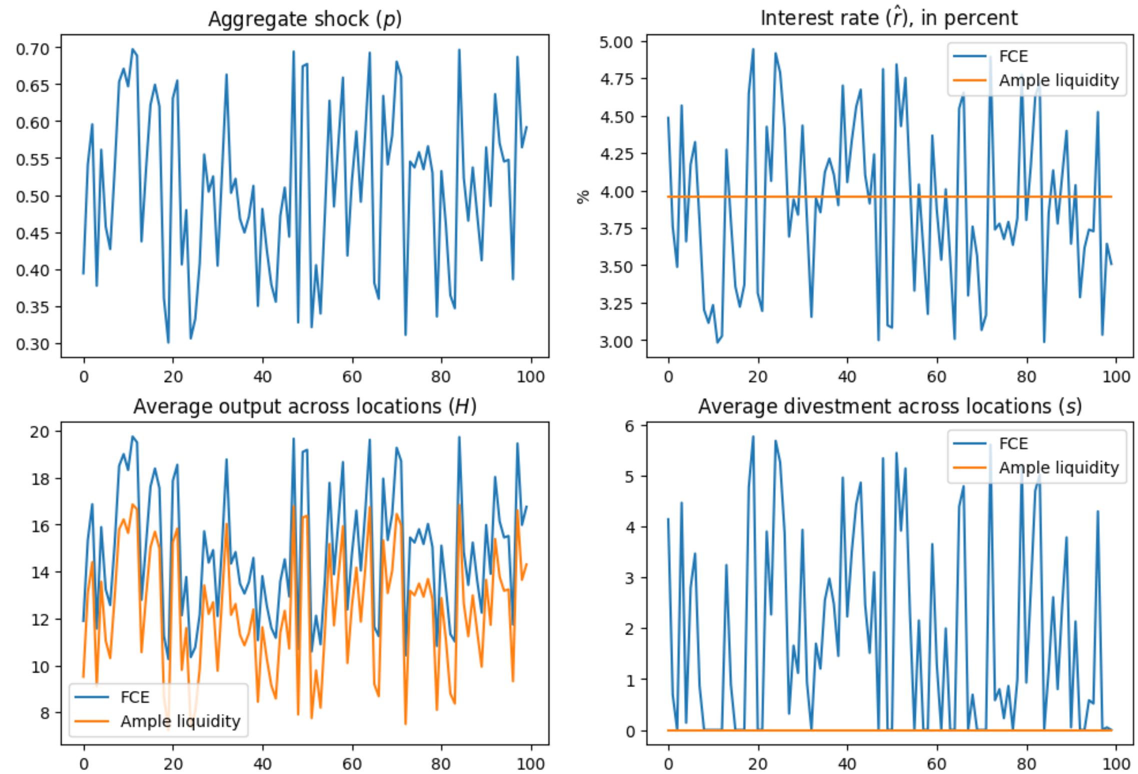

Sources and notes: The top left panel shows 100 uniformly and independently drawn aggregate shocks from $[0.3, 0.7]$. *The other three panels show the corresponding values of equilibrium variables in FCE and in ample liquidity equilibrium.*

FCE and because of the lack of fundamental adjustments, output is uniformly lower in the ample liquidity equilibrium.

Quantitatively, in a neoclassical accounting sense, the capital formation effect accounts for almost all (over 90% of) the gap between the two output series. But this does not mean that the adjustment margin is not a key driver of the gap. Per the logic behind proposition , it is because banks cannot adjust optimally to aggregate conditions that the expected marginal product of capital is low in the ample liquidity framework and that, in turn, capital formation collapses by almost 40%.

One potential gain from the point of view of a policy maker could be that dampening liquidity cost volatility – which is the expected opportunity cost of investment – could reduce output volatility. In addition, knowing that they will have a fully operating margin enable banks to take on more exposure to the risky technology instead of storing capital safely. However, these effects are offset by the fundamental adjustments banks are able to make in FCE. In an ample liquidity framework banks know what liquidity costs they will face but they almost surely operate at an inefficient scale which magnifies the effect of shocks. In the simulations displayed in figure 6, these effects largely cancel each other, with the standard deviation of output roughly 2% smaller in the ample liquidity scenario.

One final difference between FCE and ample liquidity equilibria is that whereas all location risk is eliminated in FCE, different locations may make different adjustments at date $t = 1$ hence operate at a different scale. This does not happen in the simulations above because with η small the information contained in arrival rates is limited. Only pushing η to high level, which amounts to eliminating project risk for borrowers, would cause incomplete insurance in some cases in equilibrium.

8 Ample vs abundant reserves

In official communications about its new operational framework, the Federal Reserve (the Fed, henceforth) has been careful to point out that it seeks to implement monetary policy via ample rather than abundant levels of reserve. As Afonso et al. [1] explain, in principle, the abundant reserves level begins where "*banks demand for reserves is satiated*" and "*the demand curve is flat*." Ample reserves refers to a much murkier level of reserves where the demand for reserve is, as Afonso et al. [1] put it "*gently sloped*." In practice, in their effort to estimate the threshold between ample and abundant, Afonso et al. [1] look for the reserve level at which the slope of the reserve demand curve becomes significantly different from zero. As the authors discuss and acknowledge, this sort of estimation is plagued with deep identification and endogeneity issues. For one challenge, the demand curve is obviously not stable over time – my model shows this clearly: aggregate shocks cause structural shifts in b

– which means that using time series evidence to pin down the elasticity of a demand curve at a particular point in time is heroic, at best.

In the model I have laid out so far, one interpretation of the official Fed language is that the threshold between ample and abundant reserves corresponds to the smallest level l^* past which the price of liquidity at $t = 1$ is r^* regardless of aggregate liquidity demand conditions. Proposition 6 makes this idea precise. Operating at that precise threshold, however, exposes the central bank to potential liquidity market disruptions if it is unable to implement or estimate l^* precisely. This could happen for instance if unforeseen or rare macroeconomic shocks require quantitative tightening. Small supply shocks, i.e small deviations from the threshold level l^* of reserves, or unexpected shocks in non-central bank liquidity supply, would cause small deviations from the target rate. (Subsection 9.2 discusses this effect in more detail.) It follows that attempting to operate precisely at l^* means accepting at least small variation in the market price of liquidity.

On the other hand, keeping reserves abundant would build in a buffer in l^* that would mean that small supply shocks do not alter the liquidity market equilibrium. Gagnon and Sack [11] explicitly advocate for such a buffer: "*[The] minimum level of reserves is conceptually murky, impossible to estimate, and likely to vary over time. The best approach is to steer well clear of it, especially since maintaining a higher level of reserves as a buffer has no meaningful cost.*"

A simpler interpretation of the distinction between ample and abundant reserves is that under an ample reserve framework the Fed is willing to tolerate small deviations from its target r^* when liquidity market conditions change. In the context of my model, can the Federal Reserve improve upon fully satiating reserve demand by tolerating small deviations from r^*? The answer is a clear yes: deviations can be used to convey fundamental information, as I have repeatedly emphasized in this paper. I have established that allowing full and free variation in liquidity costs promotes loan formation and output. Partial or small deviations, by the same logic, can also be used to promote economic activity. To make this point as simply as possible, I will consider the case where ν has finite support.

Proposition 8. *Assume ν has finite support $\{p_L, p_2, \dots p_H\}$. Then the central bank can increase loan formation and output by allowing the price of liquidity to deviate from r^* at $t = 1$ in some aggregate states.*

Proof. See appendix $\square$

Allowing for small deviations from the target rate can enables the policy maker to reveal outlying credit conditions thereby increasing the expected marginal product of capital and output. The fine-tuning of reserves in this fashion requires a precise understanding of the nature fundamental shocks. Therefore, it exposes the central bankers to mistakes hence potentially unwanted volatility liquidity costs.

9 Discussion of main simplifications

In this section, I elaborate on the mapping between my model and the way monetary policy is implemented in practice in the United States. And I discuss two straightforward extensions of the model. The first introduces technical, exogenous supply shocks. Then I will briefly discuss the potential consequences of persistent aggregate shocks.

9.1 Bank liquidity vs bank reserves

The Federal Open Market Committee (FOMC) has made it clear that its intent is not to exert control on just the federal funds rate (FFR) via administered rates, but also on other money market rates. For instance, in the Statement Regarding Monetary Policy Implementation which has accompanied every recent FFR decision, the FOMC states (emphasis added): "*The Committee intends to continue to implement monetary policy in a regime in which an ample supply of reserves ensures that control over the level of the federal funds rate* ***and other short-term interest rates*** *is exercised primarily through the setting of the Federal Reserve's administered rates, and in which active management of the supply of reserves is not required.*" As I discussed in the introduction, in fact, spreads between other money market rates and the

FFR have shown significantly less volatility since ample excess reserves have been the norm in the United States.

The practical implementation of ample reserves in the United States consists of several elements. See Irhrig et al. [19] for a detailed discussion. First, since 2008, the Fed has paid interest on the liquidity banks hold in their Central Bank account. The rate banks earn at the Fed constitutes in principle a lower bound on what banks must earn in the federal funds market, a market in which they exchange reserves with one another. When reserves are abundant, there is little precautionary motive to hold additional reserves, hence the demand for reserves by banks becomes flat at or near the rate the Fed pays on bank balances. In other words, the precautionary motive for holding liquidity at the Fed has been suppressed – banks hold so many excess reserves that they no longer need to take any action to meet liquidity shocks – and so what they can earn on liquidity balances is pinned down by the rate the Fed offers. By setting the supply of reserves high enough, the Fed can all but insure that the FFR is near the desired target even if variation around that high level occurs say because the Fed finds it necessary to add or substract from its quantitative easing asset purchases. This is exactly what proposition 6 formalizes.

Because non-banks also participate in the market for reserves (Government-Sponsored Enterprises and Money Market Funds, e.g.) and are not eligible to receive interest on reserves, it is possible in practice for the FFR to fall below the administered rate on reserve. The combat that issue the Fed has introduced an Overnight Reverse Repurchase Agreement (ON RRP) where non-bank suppliers of short-term liquidity can earn a rate set by the Fed. This has resulted in a FFR mostly trapped between the interest rate on reserves and the ON RRP rate, which the Fed usually sets a mere 10 basis points apart from one another. Conversely, the Federal reserve operates a repo facility (the Standard Repo Rate facility) where depository institutions and primary dealers can borrow at a rate typically 25 basis points above the ON RRP rate using highly rated securities as collateral.

In my model, to keep things tractable, I have modeled one combined market for liquidity at $t = 1$ where banks interact with other providers and demanders of liquidity. I have also modeled liquidity contracts as risk-free so one should think of the non-bank agents in the liquidity market as institutional-quality borrowers and lenders, institutional grade corporations, and broker-dealers who can provide treasury-quality collateral.

In practice, banks rely on at least four different options to manage their liquidity needs: their balances at the Federal Reserve, the Federal Funds Market, the Financial Commercial Paper market, and the overnight repo market. My model assumes, in effect, that variation in rates in those various options and markets no longer provide as much information about credit quality once an ample reserve framework is in place. The rate on Fed account balances is directly administered and does not respond to bank activity. As I discussed above, the principal goal of ample reserves is to keep the Fed Fund rates near those administered rates as well. In the Financial Commercial Paper market, banks borrow primarily from institutions that have access to Fed facilities including the ON RRP facility which once again mostly keeps transactions within the range imposed by administered rates. Furthermore, as I already discussed, with banks having full access to liquidity sources at the Federal Reserve, the Financial Commercial Paper market has shrunk to a fraction of its former size.

As for the traditional repo market, the same logic as above should apply. Since the Fed is willing to borrow and lend against high quality collateral, the repo rate should remain close to administered rates. As Logan and Schulhofer-Wohl [25], among others, discuss, the repo markets has nonetheless experienced flare ups and spikes beyond the bounds of administered rates as recently as September 2025. The key question from the point of view of my model is do those spikes contain information about loan quality? Logan and Schulhofer-Wohl [25] attributes those spikes mostly to technical factors such as tax and treasury settlement events, events that presumably contain little fundamental information about credit conditions.

In summary, ample liquidity has meant that traditional sources of liquidity for banks tend to trade at stable and small spreads to administered rates. When noticeable spikes in spreads do occur as they have recently, they appear to be caused by events that are largely unrelated to liquidity demand conditions. My next subsection discuss the role of exogenous shocks more precisely.

9.2 Supply shocks

While the times-series volatility in liquidity rates and spreads has fallen, it has not been completely eliminated as figure 2 illustrates. Some of that volatility comes from the fact that administered rates do change in practice, albeit at a low frequency, as central banks respond to changes in aggregate conditions. In my model, it would suffice to specify a loss function for the central bank that depends on the economy's fundamentals (such as the level A of aggregate productivity) and allow those fundamentals to vary in time to generate movements in r^*. As I discussed in section 8 another source of rate volatility in the United States is that the Fed is trying to operate in a region where bank demand is "gently slopped" which means it sets a level of reserves that allows for some limited movement in the Fed Funds rate.

But figure 2 shows that money market spreads to administered rates also display some volatility. Significant spikes in spreads have been well documented in the repo market as well. As I discussed in the previous subsection, Logan and Schulhofer-Wohl [25] attribute them mostly to technical factors such as tax and treasury settlement events, events that presumably contain little fundamental information about credit conditions.

To make room for exogenous supply shocks in my model, assume that the supply schedule by non-banks in liquidity markets is given by $S(r) + \epsilon$ when the cost of liquidity is r where ϵ is stochastic, uncorrelated with r, and has zero mean. The central bank is able to observe the demand schedule by banks, or, isomorphically, is able to fully ascertain what aggregate conditions are at $t = 1$. But it must set the level of reserves before knowing the realization of ϵ.

In general, the way the central bank would deal with those spurious factors depends on its specific loss function. Here, for simplicity, I will assume that the central bank sets the supply or liquidity so that the rate is r^* if $\epsilon = 0$.

With those assumptions, it remains the case that the realization of interest rates contains no information about aggregate condition since its entire volatility is driven by ϵ. It is then immediate given my earlier arguments that, provided A is high enough given other parameters, optimal loan formation by banks now obeys:

$$\int H_1(k, \bar{p}(\hat{\phi}))d\hat{\phi} \geq \int (1+r)h^{\epsilon}(r)dr,$$

with equality if $k < l$, where h^{ϵ} is the density of interest rates implied by the combination of the central bank targeting r^* and the presence of spurious shocks. The qualitative consequences of ample liquidity remain exactly as I discussed them in the previous section, but there is now residual volatility in interest rates.

9.3 Persistent aggregate shocks

The simulations I performed in the previous section assume that aggregate shocks are drawn independently across periods and that the endowment of banks is constant at l or l^* across time. Both assumptions are easy to relax. Starting with the endowment l of capital, it is enough to introduce dynastic banks that can survive multiple periods. In such a model, the consequences of ample liquidity are no longer just static but become persistent as banks accumulate more funds under an FCE than when information is suppressed. This will also exacerbate the consequences of imperfect risk-sharing as banks who operate further from an optimal scale due to an outlying realization of arrival rates are unable to accumulate as much capital as other banks.

As for persistent aggregate shock, if one assumes that p follows an autoregressive process, different histories of arrival rates will mean that banks enter a given period with different prior views about aggregate shocks. This will once again impede risk-sharing as banks now opt for different loan volumes even if their starting capital is the same.

All told, persistent aggregate shocks are likely to magnify the consequences of ample liquidity I have described in this paper.

10 Conclusion

The ample reserve framework adopted by many central banks in the wake of the great financial crisis has a number of operational advantages. For instance, it gives central banks the ability to engage in quantitative easing or tightening measures without compromising their ability to control the level of money market rates. But, as central banks recognize, ample reserve frameworks have clear drawbacks, including the potential consequences of interfering with the price signal liquidity markets provide. This paper provides, to my knowledge, the first formalization of these consequences in a fully articulated rational expectations model. Borio [3] makes the case that central banks underestimate the importance of those consequences. My hope is that this paper is a useful step towards measuring that importance precisely.

References

[1] Afonso, G., Giannone, D., La Spada, J., and J. Williams (2025), "Scarce, Abundant, or Ample? A Time-Varying Model of the Reserve Demand Curve," Federal Reserve Bank of New-York Staff Report 1019.

[2] Afonso G., Kim, K., Martin, A., Nosal. E., Potter, S., and S. Schulhofer-Wohl (2020), "Monetary Policy Implementation with an Ample Supply of Reserves," Federal Reserve Bank of Chicago Working Paper 2020-02.

[3] Borio, C. (2023), "Getting Up From the Floor," BIS Working Papers No 1100.

[4] Cesari, L. (1983), Optimization: Theory and Applications, Springer-Verlag, New York.

[5] Chien, Y. and A. Stewart (2025), "Monetary Policy Implementation with Ample Reserves," *Federal Reserve Bank of St. Louis Review*, Second Quarter 2025, Vol. 107, No. 8, pp. 1-11.

[6] Clouse, J., Infante, S., and Z. Senyuz (2025), "Market-Based Indicators on the Road to Ample Reserves," FEDS Notes, Board of Governors of the Federal Reserve, 2025.

[7] Diamond D. and P. Dybvig (1983), "Bank runs, deposit insurance, and liquidity," *Journal of Political Economy* 91, pp. 401–419.

[8] Diamond, W., Jiang, Z., and Ma, Y. (2024), "The Reserve Supply Channel of Unconventional Monetary Policy," *Journal of Financial Economics* 159, pp. 1-32.

[9] Ennis, U. and T. Kiester (2008), “Understanding Monetary Policy Implementation,” *Federal Reserve Bank of Richmond Economic Quarterly* 94, pp. 235–263.

[10] Feldman, M. and C. Gilles (1985), “An Expository Note on Individual Risk without Aggregate Uncertainty,” *Journal of Economic Theory*, pp. 26-32.

[11] Gagnon, J. and B. Sack (2014), “Monetary Policy with Abundant Liquidity: A New Operating Framework for the Federal Reserve,” Peterson Institute for Internatioal Economics, Policy Brief 14-4.

[12] Goodfriend, M. (2002), “Interest on Reserves and Monetary Policy,” *Federal Reserve Bank of New York Economic Policy Review* 8, pp. 77–84.

[13] Gorton, G., and G. Ordoñez (2014), “Collateral Crises,” *American Economic Review*, 104, pp. 343–78.

[14] Glycopantis, D., Muir, D. and Yannelis, N. (2005), “Non-implementation of Rational Expectations as a Perfect Bayesian Equilibrium,” *Economic Theory* 26, pp. 765-791.

[15] Haubrich, J. (2025), “QT, Ample Reserves, and the Changing Fed Balance Sheet,” Federal Reserve Bank of Cleveland Economic Commentary.

[16] Hildenbrand, W. and Kirman, A. P. (1988), “Equilibrium Analysis,” North Holland, Amsterdam, The Netherlands.

[17] Hirshleifer, J. (1971), “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review*, 61, pp. 561-574.

[18] Hryckiewicz, A. and L. Kozlowski (2018), “The consequences of Liquidity Imbalance: When Net Lenders Leave Interbank Markets,” *Journal of Financial Stability*, 36, pp. 82-97.

[19] Ihrig, J, Senyuz, Z., and G. C. Weinbach (2020), “The Fed’s Ample-Reserves Approach to Implementing Monetary Policy,” Finance and Economics Discussion Series, Board of Governors of the Federal Reserve.

[20] Judd. K. (1985), “The Law of Large Numbers with a Continuum of IID Random Variables,” *Journal of Economic Theory* 35, pp. 19-25.

[21] Keister, T., Martin, A. and J. McAndrews (2008), “Divorcing Money from Monetary Policy,” *Federal Reserve Bank of New York Economic Policy Review* 14, pp. 41–56.

[22] Keister, T. and J. McAndrews (2009), ”Why Are Banks Holding So Many Excess Reserves?,” Federal Reserve Bank of New-York Staff Report no. 380.

[23] Kreps, M. (1977), "A Note on Fulfilled Expectations Equilibria," *Journal of Economic Theory* 14, pp. 32-43.

[24] Lacker, J. (2006), "Central Bank Credit in the Theory of Money and Payments," Speech, Federal Reserve Bank of Richmond.

[25] Logan, L. and S. Schulhofer-Wohl (2025), "Options for Modernizing the FOMC's Operating Target Interest Rate," Federal Reserve Bank of Dallas September 2025 In Depth.

[26] Poole, W. (1968), "Commercial Bank Reserve Management in a Stochastic Model: Implications for Monetary Policy," *Journal of Finance* 23, pp. 769-791.

[27] Powell, J. (2025), "Understanding the Fed's Balance Sheet," Presentation to the 67th Annual Meeting of the National Association for Business Economics, Philadelphia, Pennsylvania.

[28] Radner, R. (1979), "Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices," *Econometrica* 47, pp. 655-678.

[29] Remache, J. (2026) "Monetary Policy Implementation in an Ample Reserves Regime," Federal Reserve Bank of New-York, February 2026 Speech.

[30] Sun, Y. (2006), "The Exact Law of Large numbers via Fubini Extension and Characterization of Insurable Risks," *Journal of Economic Theory* 126, pp. 31-69.

[31] Sun, Y., Wu, L., and N. Yannelis (2012), "Existence, Incentive Compatibility and Efficiency of the Rational Expectations Equilibrium," *Games and Economic Behavior,* 76, pp. 329-339.

[32] Taylor, J. (2013), "A Review of Recent Monetary Policy," Testimony Before The Subcommittee on Monetary Policy and Trade Committee on Financial Services, U.S. House of Representatives.

[33] Woodford, M. (2000), "Monetary Policy in a World Without Money," *International Finance* 3, pp. 229–60.

A Proofs

A.1 Proof of proposition 2

I will drop the requirement that $c^E(\phi) \leq pc^L(\phi)$ and arrive at policies that satisfy it so that ignoring it is without loss of generality. The resulting Lagrangian for the bank's problem is:

$$\mathcal{L} = \int_0^1 \left[Ap(k - s(\phi))^\alpha + (l - k) + \delta s(\phi) - \phi c^E(\phi) - r\, b(\phi) - (1-\phi)p\, c^L(\phi)\right] f(\phi)\, d\phi$$
$$+ \lambda \left[\int_0^1 \left(\phi \ln c^E(\phi) + (1-\phi) \ln p\, c^L(\phi)\right) f(\phi)\, d\phi - \ln w\right]$$
$$+ \int_0^1 \gamma(\phi) \left[l - k + b(\phi) - \phi\, c^E(\phi)\right] d\phi + \int_0^1 \nu(\phi) \left[A(k - s(\phi))^\alpha - c^L(\phi)\right] d\phi$$
$$+ \int_0^1 \eta(\phi)\, [k - s(\phi)]\, d\phi - \mu(k - l),$$

with λ, γ, ν, η and ν the shadow price of each constraint. I will also attach shadow price $\rho(\phi)$ to the constraint that $s(\phi) \geq 0$ for all $\phi \in [0, 1]$.

The first order condition for b is

$$r f(\phi) + \gamma(\phi) = 0$$

which implies

$$\gamma(\phi) = r f(\phi)$$

for all ϕ. In turn, the first order condition for c^E gives

$$-\phi f(\phi) + \frac{\lambda\, \phi\, f(\phi)}{c^E(\phi)} - \gamma(\phi)\, \phi = 0.$$

Substituting $\gamma(\phi) = r f(\phi)$ and dividing by $\phi f(\phi)$ gives

$$c^E(\phi) = \frac{\lambda}{1 + r}.$$

The first-order condition for c^L is

$$-(1-\phi)p\, f(\phi) + \frac{\lambda(1-\phi) f(\phi)}{c^L(\phi)} - \nu(\phi) = 0,$$

which gives

$$c^L(\phi) = \min\left\{\frac{\lambda}{p},\ A(k - s(\phi))^\alpha\right\}.$$

Indeed, if the upper bound $A(k - s(\phi))^\alpha$ is not binding for some ϕ then $\nu(\phi) = 0$ by complementary slackness and it is then immediate that $c^L(\phi) = \frac{\lambda}{p}$.

For scrapping policy s, the necessary first-order condition is:

$$\left[\delta - Ap\alpha(k - s(\phi))^{\alpha-1}\right] f(\phi) = \nu(\phi)\, A\alpha(k - s(\phi))^{\alpha-1} + \eta(\phi) - \rho(\phi)$$

for almost all ϕ where, recall, $\rho(\phi)$ is the shadow price of the non-negativity constraint on s. Note first that with $\alpha < 1$ the constraint $s(\phi) \leq k$ cannot bind as long as $Ap > 0$, and so $\eta(\phi) = 0$ for all ϕ. Rearranging the necessary condition for $s(\phi)$ then gives:

$$Ap\alpha(k - s(\phi))^{\alpha-1} f(\phi) = \delta f(\phi) - \nu(\phi) A\alpha(k - s(\phi))^{\alpha-1} + \rho(\phi).$$

It follows that if $s(\phi) > 0$ then $Ap\alpha\,(k - s(\phi))^{\alpha-1} < \delta$. Assume by way of contradiction that $s(\phi^*) > 0$ for some positive measure of ϕ^*. Then in any state in which $s(\phi) = 0$ we have

$$Ap\alpha\,(k - s(\phi))^{\alpha-1} = Ap\alpha k^{\alpha-1} \leq Ap\alpha\,(k - s(\phi^*))^{\alpha-1} < \delta.$$

It follows that if $s(\phi) > 0$ for any ϕ then $Ap\alpha(k - s(\phi))^{\alpha-1} < (1 + r)$ for all ϕ.

Now consider the necessary condition for optimal capital k:

$$\int_0^1 Ap\alpha\,(k - s(\phi))^{\alpha-1} f(\phi)\, d\phi = (1 + r) + \mu.$$

With the assumption that

$$k^* = (p\alpha A)^{\frac{1}{1-\alpha}} < l,$$

the upper bound on capital does not bind and $\mu = 0$. This rules out the case that $Ap\alpha(k - s(\phi))^{\alpha-1} < (1 + r)$ for all ϕ hence, per the logic of the previous paragraph, rules out the case $s(\phi) > 0$ for some ϕ. So we now know that as long as $\delta < 1$ $s(\phi) = 0$ for all ϕ. (When $\delta = 1$, an indeterminacy arises because the bank can start with excess capital in place only to divest the excess at no net cost at $t = 1$, but setting $s(\phi) = 0$ remains weakly optimal in that case too.)

Now, provided c^L is interior for all ϕ, we can plug the resulting expression for c^L along with c^E into the borrower participation constraint to obtain:

$$\lambda = w(1 + r)^{\bar{\phi}}.$$

For c^L to be interior we need

$$w(1 + r)^{\bar{\phi}} < Ak(r)^{\alpha}.$$

Plugging in the expression for k we derived above, the sufficient condition stated in the proposition follows, and this completes the proof.

A.2 Proof of lemma 3

Note first that for A high enough there must a exist a profitable contract for the bank and so I will dispense with verifying that profits are non-negative.

Consider a version of the bank's problem amended to drop the upper bound restriction on $c^L(\hat{\phi}, r)$. (I will eventually verify that if A is high enough then in fact that constraint cannot bind at any solution.) As in the previous proof I will also ignore the requirement that $c^E(\hat{\phi}, r) \leq \hat{p}(r)c^L(\hat{\phi}, r)$ for all $(\hat{\phi}, r)$ and arrive at policies that satisfy that requirement anyway. The corresponding Lagrangian is:

$$\mathcal{L} = \int_{r_L}^{r_H}\int_0^1 \left[H(k, p(r)) + (l-k) - \hat{\phi}c^E(\hat{\phi}, r) - r\, b(\hat{\phi}, r) - (1-\hat{\phi})[\hat{p}(r)\, c^L(\hat{\phi})]\right] \hat{f}(\hat{\phi}|p(r))h(r)d\hat{\phi}dr$$
$$+ \lambda\left[\int_{r_L}^{r_H}\int_0^1 \left(\hat{\phi}\ln c^E(\hat{\phi}, r) + (1-\hat{\phi})\ln[\hat{p}(r)\, c^L(\hat{\phi}, r)]\right) \hat{f}(\hat{\phi})h(r)d\hat{\phi}dr - \ln w\right]$$
$$+ \int_{r_L}^{r_H}\int_0^1 \gamma(\hat{\phi}, r)\left[l - k + b(\hat{\phi}, r) - \hat{\phi}c^E(\hat{\phi}, r)\right] d\hat{\phi}dr + \mu(l-k)$$

with λ, γ, and μ are the shadow prices of the remaining constraints. I can now proceed exactly as in the case without aggregate shocks. The first order condition with respect to $b(\hat{\phi}, r)$ is

$$-r\hat{f}(\hat{\phi}|p(r))h(r) + \gamma(\hat{\phi}, r) = 0.$$

The first-order condition with respect to $c^E(\hat{\phi}, r)$, for its part, gives:

$$-\hat{\phi}\hat{f}(\hat{\phi}|p(r))h(r) - \hat{\phi}\gamma(\hat{\phi}, r) + \frac{\hat{\phi}\lambda\hat{f}(\hat{\phi}|p(r))h(r)}{c^E(\hat{\phi}, r)} = 0.$$

Combining those two first-order conditions gives

$$c^E(\hat{\phi}, r) = \frac{\lambda}{1+r},$$

exactly as in the case without aggregate shocks, and we likewise continue to obtain

$$c^L(\hat{\phi}, r) = \frac{\lambda}{\hat{p}(r)}$$

in this problem without upper bound restriction on c^L. In other words, the bank continues to offer full insurance against liquidity shock but does not provide insurance against the aggregate shock.

For k, a necessary condition is that, if $k < l$,

$$\begin{aligned}\int_{r_L}^{r_H}\int_0^1 [H_1(k, \hat{p}(r)) - 1]\, \hat{f}(\hat{\phi}|p(r))h(r)d\hat{\phi}dr &= \int_{r_L}^{r_H}\int_0^1 \gamma(\hat{\phi}, r)d\hat{\phi}dr \\ \Leftrightarrow \int_{r_L}^{r_H} [H_1(k, p(r)) - 1]\, h(r)dr &= \int_{r_L}^{r_H}\int_0^1 r\hat{f}(\hat{\phi}|p(r))h(r)d\hat{\phi}dr = \int_{r_L}^{r_H} rh(r)dr\end{aligned}$$

and the result for optimal k follows. The above uses the expression I derived above for γ. It also uses the fact that $[H_1(k,p(r))-1]$ is independent of $\hat{\phi}$ so that $\int_0^1 \hat{f}(\hat{\phi}|p(r))d\hat{\phi}=1$ for all r integrates out of the expression.

Now I only need to argue that if A is large enough then the upper bound for c^L cannot bind in any state. Plugging the expression for c^E and c^L we obtained above in the participation constraint, which must hold as an equality at any bank optimum, gives a value for λ that does not depend on A. So I can increase A without changing c^L at the proposed solution. On the other hand, k rises with A holding all other parameters the same, while s eventually decreases towards zero for all possible value of p. (Recall that $p_L>0$ so that $p(r)>p_L>0$ in any FCE.) And it now follows that the upper bound on c^L cannot bind in any state, which completes the proof.

A.3 Proof of proposition 4

In light of lemma 3, provided A is high enough, equilibrium allocations are fully described by a value $\lambda>0$ for the shadow price of the participation constraint and the average level $\bar{r}=\int_{r_L}^{r_H}(1+r)h(r)dr$ of interest rate. So start by fixing λ and $\bar{r}$. As I discussed before stating the proposition, associated with this joint guess is a capital level k and a unique schedule $\hat{r}(p,\bar{r})$ for all possible $p\in[p_L,p_H]$ that clears liquidity markets. In equilibrium, we need $\int \hat{r}(p)d\nu(p)$ to coincide with the starting guess for $\bar{r}$. Define, for all $n\in I\!N$ and $(\lambda,\bar{r})\in[0,\lambda_n]\times[0,r_n]$

$$E_{r,n}(\lambda,\bar{r})=\int \hat{r}(p,\bar{r})d\nu(p)-\bar{r}$$

to be the gap between the average of the rate schedule implied by $\bar{r}$ and $\bar{r}$, where λ_n and r_n are sequences of arbitrary upper bounds that grow to $+\infty$.

We also need the guess for λ to be correct. So define for all $n\in I\!N$

$$E_{\lambda,n}(\lambda,\bar{r})=\ln w-\int_{r_L}^{r_H}\int_0^1\left(\hat{\phi}\ln\left(\frac{\lambda}{1+\hat{r}(p,\bar{r})}\right)+(1-\hat{\phi})\ln\lambda\right)\hat{f}(\hat{\phi}|p(r))h(r)d\hat{\phi}dr$$

evaluated at the unique $\hat{r}$ (hence $p(r)$) schedule implied by guesses $(\lambda,\bar{r})\in[0,\lambda_n]\times[0,r_n]$.

I can now define correspondence $T_n:(E_{\lambda,n}\times E_{r,n})([\frac{1}{n},\lambda_n]\times[0,r_n])\mapsto[0,\lambda_n]\times[0,r_n]$ by:

$$T_n(E_\lambda,E_r)=\arg\max_{(\lambda,\bar{r})\in[\frac{1}{n},\lambda_n]\times[0,r_n]}\lambda E_\lambda+\bar{r}E_r.$$

Intuitively, this mapping increases λ when it is not high enough to deliver $\ln w$ to the borrower and raises $\bar{r}$ when it results in a schedule $\hat{r}$ whose average is too high. Classical arguments (see for instance Hildenbrand and Kirman [18]) imply that for all n T_n is non-empty, upper-hemicontinous, and convex-valued. It follows that $T_n\times(E_{r,n}\times E_{\lambda,n})$ has a fixed point for all n. If, for some n, at the fixed point, $(\lambda,\bar{r})$ is interior then both E_λ and E_r are zero, by construction of the T mapping. So I will rule out all possible corner outcomes for n large enough.

If λ is close enough to zero then $E_\lambda < 0$ and so the lower bound on λ cannot bind. Likewise if it gets large enough $E_\lambda > 0$ and the upper bound cannot bind either at a fixed point. So for n large enough the λ part of the fixed point is in fact interior. Turning now to $\bar{r}$, as it grows large, k converges to zero. At the same time, at any fixed point, this means that $\hat{r}(p)$ must also grow without bound for all of p. But this must eventually violate market clearing for all p. It follows that the upper bound on n does not bind.

This only leaves the possibility that $\bar{r} = 0$ eventually as n gets large. This is only compatible with $E_r \leq 0$ eventually, which in turn requires $\hat{r}(p) = 0$ for all p, which can be part of an equilibrium (with, possibly, excess supply of liquidity at $r = 0$.)

So now we have a candidate equilibrium consistent with optimality and market clearing. To constitute a FCE, however, in other words to guarantee that the $\hat{r}$ schedule is strictly monotonic, we also need $\hat{r}(p) > 0$ at all p. But as A grows large k must eventually reach l, and so we must in fact have $\hat{r}(p) > 0$ at all p. This completes the proof.

A.4 Proof of proposition 5

I will follow the same strategy as in the FCE case by ignoring the upper bound on c^L at first and then show that the constraint does not bind at the solution when A is high enough. The Lagrangian of interest, in this case, is:

$$\mathcal{L} = \int_0^1 \left\{ \int_0^1 \left[Ap(k - s(\hat{\phi}))^\alpha + \delta s(\hat{\phi}) + l^* - k - \hat{\phi} c^E(\hat{\phi}) - r^* b(\hat{\phi}) - (1-\hat{\phi})\hat{p}(p) c^L(\hat{\phi}) \right] g(p|\hat{\phi}) dp \right\} \hat{f}(\hat{\phi}) d\hat{\phi}$$
$$+ \lambda \left[\int_0^1 \left\{ \int_0^1 \left[\hat{\phi} \ln c^E(\hat{\phi}) + (1-\hat{\phi}) \ln [\hat{p}(p) c^L(\hat{\phi})] \right] g(p|\hat{\phi}) dp \right\} \hat{f}(\hat{\phi}) d\hat{\phi} - \ln w \right]$$
$$+ \int_0^1 \gamma(\hat{\phi}) \left[l^* - k + b(\hat{\phi}) - \hat{\phi} c^E(\hat{\phi}, r) \right] d\hat{\phi} + \mu(l^* - k)$$

with λ, γ, and μ are the shadow prices of the remaining constraints. As before, start by stating the first-order condition for $b(\hat{\phi})$:

$$r^* \hat{f}(\hat{\phi}) \int_0^1 g(p|\hat{\phi}) dp = r^* \hat{f}(\hat{\phi}) = \gamma(\hat{\phi}).$$

The FOC for c^E, in turn, is

$$\hat{\phi} \hat{f}(\hat{\phi}) \int_0^1 g(p|\hat{\phi}) dp = \hat{\phi} \hat{f}(\hat{\phi}) = \lambda \hat{f}(\hat{\phi}) \frac{\hat{\phi}}{c^E(\hat{\phi})} - \gamma(\hat{\phi})(\hat{\phi})$$

and the result for c^E follows given the expression for $\gamma(\hat{\phi})$ I obtained from the previous FOC. The algebra for c^L is likewise almost identical as in the FCE case.

The key part of the proposition involves optimal capital choice. To simplify this part of the problem, notice that the the bank's objective can be separated into a part that involves production decisions:

$$\int_0^1 \left\{ \int_0^1 \left[Ap(k - s(\hat{\phi}))^\alpha + \delta s(\hat{\phi}) + (l^* - k(\hat{\phi})) \right] g(p|\hat{\phi}) dp \right\} \hat{f}(\hat{\phi}) d\hat{\phi}$$

and a part that does not involve k or s in any way. But since production is linear in p we can rewrite this production block simply as:

$$\int_0^1 \left[A\bar{p}(\hat{\phi})(k - s(\hat{\phi}))^\alpha + \delta s(\hat{\phi}) + K - k \right] \hat{f}(\hat{\phi}) d\hat{\phi},$$

which, upon maximizing over s for each possible arrival rate, becomes:

$$\int_0^1 \left[H(k, \bar{p}(\hat{\phi})) + l^* - k \right] \hat{f}(\hat{\phi}) d\hat{\phi}.$$

The results for k and s then follow.

This only leaves proving that when A is high enough given other parameters the constraint on c^L cannot bind but that argument is exactly the same as in the FCE case. This completes the proof.

A.5 Proof of proposition 8

Consider first an altered version of my environment in which the bank observes p_L but cannot distinguish between the other possible values $\{p_2, \dots p_H\}$ of the aggregate shock. Assume the price of liquidity is r^* for all possible values of the aggregate shocks. Compute the bank policy functions k and s. These policies are exactly as in proposition 5 except that $\bar{p}(\hat{\phi}) = p_L$ for all $\hat{\phi}$ when $p = p_L$ while, as before,

$$g(p|\hat{\phi}) \propto \nu(p) \hat{f}\left(\hat{\phi}|p\right),$$

and

$$\bar{p}(\hat{\phi}) = \int p dg(p|\hat{\phi})$$

for $p > p_L$ where g now has support restricted to $p > p_L$. Now find the smallest l^* so that the price of liquidity is r^* for all aggregate shocks. Since arrival rates are highest on average when $p = p_L$, it must be that

$$\int b(\hat{\phi}, r^*) \hat{f}(\hat{\phi}|p_L) d\hat{\phi} = S(r^*)$$

at this value of l^*.

By setting p_L low enough I can ensure that $s(p_L) > 0$ at the resulting plan. This means that investment and output are higher when the bank learns p_L than when it does not, holding the price of liquidity at r^* and liquidity levels at l^*.

Consider then a small decrease from l^*. Liquidity markets still clear at existing policies (they had slack at l^* since $p_i > p_L$ for all $i \geq 2$.) But they no longer clear at $p = p_L$ since that constraint was tight. The intermediate value theorem implies that there is a value $r_L \neq r^*$ for the price of liquidity in a neighborhood of r^* such that markets clear when $p = p_L$ holding the rate at r^* for all other possible values of the aggregate shock. For a small enough change in l^*, liquidity markets continue to clear for all $p > p_L$. The resulting equilibrium is such that $r = r^*$ in all aggregate states, except for the worst. For small enough changes in l^* output and investment and output must be near their level in the partially revealing economy. But those are higher than what they are in an equilibrium in which the central bank impose $r = r^*$ in all states. This establishes that allowing for a small deviation from r^* in the worst possible aggregate state can raise investment and output.