American Economic Review 2017, 107(4): 1005–1029
https://doi.org/10.1257/aer.20140782

Banks as Secret Keepers†

By Tri Vi Dang, Gary Gorton, Bengt Holmström, and Guillermo Ordoñez*

Banks produce short-term debt for transactions and storing value. The value of this debt must not vary over time so agents can easily trade it at par like money. To produce money-like safe liquidity, banks keep detailed information about their loans secret, reducing liquidity if needed to prevent agents from producing costly private information about the banks' loans. Capital markets involve information revelation, so they produce risky liquidity. The trade-off between less safe liquidity and more risky liquidity determines which firms choose to fund projects through banks and which ones through capital markets. (*JEL* D92, E51, G21, G31, G32)

The output of banks is short-term debt used for transactions and storing value (Gorton and Pennacchi 1990). A defining characteristic of privately produced, money-like securities is that agents accept them at par when transacting, because the agents expect to be able to redeem them at par. The value of the money is not in doubt when transacting and the value does not vary over time. In other words, bank money is not sensitive to information, either public or privately produced. But, how can banks produce such money when it must be backed by risky assets that require evaluation?

Our answer to this conundrum is that banks produce private money because they can keep the information that they produce about backing assets secret. By being

* Dang: Department of Economics, Columbia University, International Affairs Building, 420 W 118 Street, New York, NY 10027 (e-mail: td2332@columbia.edu); Gorton: School of Management, Yale University, 135 Prospect Street, Box 208200, New Haven, CT 06520, and NBER (e-mail: gary.gorton@yale.edu); Holmström: Department of Economics, Massachusetts Institute of Technology, 50 Memorial Drive, Building E52, Cambridge, MA 02142, and NBER (e-mail: bengt@mit.edu); Ordoñez: Department of Economics, University of Pennsylvania, 428 McNeil Building, 3718 Locust Walk, Philadelphia, PA, 19104, and NBER (e-mail: ordonez@econ.upenn.edu). Special thanks to Douglas Diamond who discussed the paper multiple times at various conferences. We also thank Tobias Adrian, Fernando Alvarez, Bruno Biais, Paul Beaudry, Ricardo Caballero, John Geanakoplos, Mike Golosov, Jonathan Heathcote, Christian Hellwig, Alessandro Pavan, Enrico Perotti, Adriano Rampini, Andrei Shleifer, Alp Simsek, Jean Tirole, Juuso Toikka, Vish Viswanathan, Michael Woodford, and seminar participants at Banque de France, Boston College, Chicago, Chinese University of Hong Kong, Columbia, Commodity Futures Trading Commission, EIEF, Federal Reserve Board, Harvard/MIT Applied Theory Workshop, Hong Kong Monetary Authority, IMF, LBS, MIT, Northwestern, NY Fed, NYU, Rutgers, Toulouse, UBC, Universite de Montreal, Wharton, Yale Cowles Foundation, the AEA Meetings in Boston, the SED Meetings in Cyprus, the LUISS Workshop on "Financial Frictions in Macro-Models," the Barcelona GSE Summer Forum on "Financial Intermediation, Risk and Liquidity Management," the ESSET Meetings at Gerzensee, the Sciences Po Paris Workshop on "Macro and International Finance," the 2013 Summer Workshop at the Chicago Fed, and the Northwestern Nemmers Conference in Honor of Jean Tirole. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

† Go to https://doi.org/10.1257/aer.20140782 to visit the article page for additional materials and author disclosure statement(s).

opaque, banks can produce bank money more efficiently. Opacity makes it costly for an expert investor to find out information about the details of the bank's balance sheet, eliminating the expert's informational advantage. Opacity also mutes the effects that public information may have on the value of the bank's assets. By keeping information symmetric among traders, opacity makes trading in bank money liquid.

We study a simple three-period model. A firm has an investment project at date 0, but no funds. The project pays off at date 2. There is a consumer (the *early consumer*) with an endowment at date 0, who could fund the project. The problem is that the early consumer has a liquidity need at date 1 and while he can store his endowment at no cost, it is insufficient to cover both the firm's investment need and the consumer's date-1 liquidity need. The early consumer could get help from a *late consumer*, who enters at date 1 with an endowment that is large enough to cover the early consumer's liquidity shortfall. If the two consumers were able to get together and write an optimal contract at date 0, the contract would provide both efficient investment and liquidity provision. This is impossible, however, because the late consumer enters only at date 1 (or equivalently, cannot commit to pay out of future endowments.)

We consider two institutions that can help resolve the dilemma and intermediate trade between the early and late consumers. The first institution is a capital market that raises funds at date 0 from the early consumer in exchange for tradable shares of the firm. At date 1 the early consumer can sell the shares to the late consumer in the market. The late consumer has the expertise to process information about the firm's project; for simplicity we assume that this allows him to determine the project's final payoff already at date 1. Competition between late consumers will then bring the date 1 market price equal to the final project payoff before it is realized. In the bad state, when the project's payoff is low, the early consumer will not be able to cover all his liquidity needs at date 1. Anticipating this possibility he will demand a premium for the risk of illiquidity. This is the cost of capital markets. There is liquidity, but it is *risky liquidity*.

The alternative is to fund the project through a bank. The bank can hide key details about the project so that the late consumer is unable to costlessly use her expertise to evaluate the project at date 1. We assume that the detailed information that the bank can hide is information that the bank does not know how to interpret. This assumption is one way to assure that the bank does not want to reveal any information to the late consumer at date 1. Opacity will facilitate risk sharing between the early and the late consumers as follows. Both consumers will deposit their endowments in the bank, the early consumer at date 0 and the late consumer at date 1. The bank uses the early consumer's deposit to fund the firm's investment. The bank promises to cover some or all of the early consumer's liquidity needs at date 1. Both consumers will get a share in the risky payoff of the project at date 2. Unlike capital markets, the payment to the early consumer at date 1 is noncontingent as in a deposit account. There is *safe liquidity*. In the case where the late consumer cannot acquire private information, banking can implement a first-best outcome.

Our main interest, however, is in the cost of producing bank money. It is determined by the cost of preventing expert investors from acquiring private information about the assets that are used to back bank money. Investors will not acquire information about assets that are too costly to evaluate (because they are complex or distant) or which offer low returns to information production (because they are unlikely to fail). Such information insensitive assets can be used to create bank money at no cost. However, when the bank relies on information sensitive assets to produce bank money there is a cost, because the bank has to scale down the amount of bank money that is being produced or its volume of loans in order to maintain opacity. The cost of producing money from more information sensitive assets implies more expensive terms for firms that seek bank financing. As we will see, firms with assets that are too risky or too transparent will find it cheaper to seek funding in capital/stock markets.

The key distinction between banks and capital markets is the way the two institutions process information. Capital markets aggregate private and public information into prices that fully reflect the information about projects. Efficient price discovery, due to competition among expert traders, results in symmetric information among all investors and liquid contingent claims. The downside of price discovery is that capital markets cannot deliver securities with a stable value; they cannot create money-like securities. This is the purview of opaque banks. On the other hand, banks cannot produce state-contingent claims that mimic capital markets, because they only deal with one investor at a time; bilateral transactions lack informational efficiency. Both banks and capital markets create liquid claims, but of different kinds. Capital markets create state-contingent liquidity through information-revealing price discovery, while banks create money-like liquidity through information concealing opacity. In the capital market the investment is risky, but realizable at all times. In the bank only part of the investment can be realized early, but that part is safe (like a deposit). This is the key trade-off between banking and capital markets.

Note that the allocation of projects between banks and financial markets does not rely on any comparative advantage that banks have in evaluating and overseeing its assets. Banks will invest in projects that are less risky and more opaque because there is an information complementarity between the production of private money and the assets that lower the cost of producing private money. Banks cannot compete with capital markets in funding risky, transparent projects, so they are more constrained in their investments than regular firms, which have the opportunity to choose between banks and capital markets in seeking funds. The production of private money makes banks special in our model.

Related Literature.—The idea that it may be optimal to keep information secret is not new. It was perhaps first articulated by Hirshleifer (1971), who showed that early release of information can destroy future insurance opportunities. This general idea also underlies Kaplan's (2006) study of a Diamond and Dybvig (1983) type model in which the bank acquires information before depositors do. Kaplan studies when the optimal deposit contract will be noncontingent. Breton (2011) views banks as a solution to information appropriability problems. Our focus is on the costs of preventing information acquisition by outsiders (the late consumer). The bank may have to reduce liquidity provision or ration lending to prevent information acquisition, leading to an endogenous sorting of firms between banks and capital markets.

Dang, Gorton, and Holmström (2013) also study a setting in which the provision of money-like securities relies on information insensitiveness, which shields uninformed traders from losing money when they trade with agents who can become privately informed. They use a three-period model where the only means of saving funds from the first date to meet liquidity requirements arising at the interim date is to buy a claim issued on a given project that pays off on the last date. That claim is then used as collateral to raise funds for liquidity demands on the intermediate date. The main result is that the optimal claim to issue at the first date is debt, and the optimal claim to issue using that debt as collateral at the intermediate date is also debt. So that paper is entirely about the optimal structure of contracts, in a world with an exogenous project and no other savings instruments, to fight private information acquisition. There are no funding decisions (the project is already up and running) so information has no social value. Most importantly there is no comparison of the trade-off between bank funding and capital market funding. In our paper, a critical trade off is producing information for investment efficiency while simultaneously attempting to produce liabilities that are efficient. Our paper is about the choice of assets to fight private information acquisition.

Our paper offers a new explanation for the existence of financial intermediaries that relies on complementarities between the two sides of a bank's balance sheet. Most explanations look at just one side of the balance sheet. One line focuses on the role of banks in making loans. Banks are viewed as producing information about potential borrowers and/or monitoring borrowers after the loan has been made: see, e.g., Boyd and Prescott (1986). Another line looks at the liability side of banks. In Diamond and Dybvig (1983) the bank issues demand deposits to insure consumers against liquidity shocks, but the asset side is deterministic. In their paper the possibility of trading among consumers destroys insurance opportunities, because the insured consumers can cash out after they learn their liquidity shock (privately): also see, e.g., Jacklin 1987 and Haubrich and King 1990. In our case liquidity needs are deterministic while the asset side is stochastic. Bank secrecy is exactly what prevents direct trading between agents allowing for insurance opportunities.

There are a few papers that, like ours, show a complementarity between the two sides of banks' balance sheets. In Diamond (1984) the bank invests in a large number of independent projects that allows the bank to issue riskless debt (deposits) to investors. Because debt is riskless depositors do not need to monitor the bank, while the bank as the residual claimant, will monitor borrowers efficiently. Because no information is being produced publicly about the bank, the bank is opaque as a by-product of being fully diversified. In Diamond and Rajan (2001) banks monitor borrowers and can do so better than others because the design of their liability side gives them credibility in enforcing repayments more effectively. Demand deposits, which can be withdrawn at any moment, create the right incentives for investing in and collecting from borrowers. In Kashyap, Rajan, and Stein (2002) banks structure their balance sheets to take advantage of the imperfect correlation between deposit withdrawals and loan commitment draw-downs, tying the two banking activities together. In Breton (2007) investors invest in long-term projects which they monitor and therefore have private information about. The private information makes the projects illiquid; if they were sold on the market, they would be subject to adverse selection. Kept on the bank's balance sheet, the information about the projects will not leak out so depositors, now without private information, can be issued claims that are liquid in the market.

There is, finally, a large accounting literature on the potential costs of disclosure, to which our paper contributes. A large part of it focuses on firms' disclosure in stock markets, such as Diamond and Verrecchia (1991) who show that more information revelation reduces the firm's cost of capital, but it can have the opposite effect by reducing liquidity in the stock market. In a different context, Andolfatto (2010) shows that transparency can be socially costly in a monetary economy with search frictions.[1]

Banks have always been opaque, even with deposit insurance. Badertscher, Burks, and Easton (2015) study bank stock price reactions to the quarterly release of the *Call Reports*, which contain information the banks have submitted to bank regulators. They find significant (and economically meaningful) stock price and volume reactions upon release of the information, even when the *Call Reports* release follows a bank earnings call. Also, since the advent of deposit insurance banks are examined by government regulators. Examination results are kept secret, but are still informative. DeYoung et al. (2001) find that government examinations did produce new, value-relevant, information which is eventually revealed in bank subordinated debt prices. Berger and Davies (1998) find that information from unfavorable examinations is eventually revealed in banks stock prices. These results are consistent with banks being opaque and examiners uncovering secrets.[2]

In the next section we introduce the model, calculate the first best allocation, and then show that the first best can be implemented by banks, but not by capital markets. Section II is the heart of the paper. We study the case where the late consumer is tempted to produce information about the state of the firm's project, the bank's secret. If she learns that the project will not turn out well, she will not deposit in the bank (or will discount the early consumer's check if they trade directly). Anticipating this outcome, the bank may supply less bank money or ration credit ex ante. This leads to a trade-off between bank funding and market funding, which we explore in Section III. Section IV concludes.

I. Model

In this section we present the model. Then, we derive the first best allocations and the allocations achievable with capital markets and with a banking technology that allows secret keeping.

A. *Setting*

Preferences and Technologies.—Consider an economy with a single good (the numeraire), three dates, $t \in \{0, 1, 2\}$, and three sets of agents: a firm (F), a single

[1] See also the more recent contribution of Monnet and Quintin (2013).

[2] Other evidence includes Bessler and Nohel (1996), who study dividend cuts and find significantly stronger negative reactions for banks than for nonbanks. Hirtle (2006) examines the abnormal stock returns to 44 bank holding companies in response to the SEC mandate that CEOs certify the accuracy of their financial statements. This mandate resulted in no abnormal response in the case of nonfinancial firms, but bank holding companies did experience positive and significant abnormal returns. Hirtle also finds that the abnormal returns are related to measures of opacity. Also see Haggard and Haggard and Howe (2007); Morgan (2002); Iannota (2006); Jones, Lee, and Yeager (2012); Flannery, Kwan, and Nimalendran (2004, 2013); and Flannery (1998) for evidence of bank opacity.

early consumer (him, E) and $N > 1$ identical late consumers (her, L). Preferences and endowments are as follows:

$$U_F = \sum_{t=0}^{2} C_{Ft}, \qquad \omega_F = (0,0,0);$$

$$U_E = \sum_{t=0}^{2} C_{Et} + \alpha \min\{C_{E1}, k\}, \qquad \omega_E = (e,0,0);$$

$$U_L^l = \sum_{t=1}^{2} C_{Lt}^l + \alpha \min\{C_{L2}^l, k\}, \qquad \omega_E^l = (0,e,0) \quad \forall l \in \{1,\ldots,N\},$$

where C_{ht} denotes the consumption of agent $h \in \{F,E,L\}$ at date $t \in \{0,1,2\}$ and α and k are positive constants. Consumption is constrained to be nonnegative for all parties (limited liability). The early consumer is born in period $t = 0$ and each late consumer in $t = 1$. Both types of consumer have e units of the good as endowment when they are born and nothing at other dates, and they both have some urgency $(\alpha > 1)$ to consume up to k the period after they are born—in period $t = 1$ for the early consumer and in period $t = 2$ for each late consumer. We assume that all consumers can store the good for consumption at a later date at no cost, so no institutions or contracts are needed for savings purposes.

The urgency to consume the period after birth can be thought of as a demand for liquidity, for instance because there is a productive investment opportunity the period after birth that costs k and produces $k(1+\alpha)$ in the subsequent period. The need for liquidity leads to a kinked utility function, featuring risk-aversion globally, but risk-neutrality locally. This is convenient for several reasons. First, it highlights the effects of liquidity needs on generating risk aversion even in the presence of fundamental risk neutrality. Second, it simplifies the exposition without any loss in conclusions. In the online Appendix we show that qualitatively similar results are obtained with arbitrary risk-averse preferences.

The firm has two projects that need w at $t = 0$ to operate. One, which we call a *lemon*, never generates any output. The other, which we call *worthy*, generates $x > w$ at $t = 2$ (state g) with probability λ, and zero otherwise (state b), all measured in terms of the single good. The projects are linearly divisible, i.e., if the firm operates a fraction η of the worthy project, it costs ηw and generates ηx in case of success. We assume these projects' characteristics are common knowledge. We make the following further assumptions about the worthy project and the endowments.

ASSUMPTION 1 (Worthy Project and Endowments):

(*i*) *The worthy project is ex ante efficient*: $\lambda x > w$.

(*ii*) *The early consumer can fully cover either the investment in the worthy project or his liquidity need, but not both*: $e > k$ *and* $e > w$, *but* $e < k + w$.

(*iii*) *Combining the endowments of the early consumer and of a single late consumer is enough to cover both their liquidity needs and the investment in the worthy project*: $2e \geq 2k + w$.

The first part states that the worthy project has a positive NPV. The second part creates a trade-off between investment and saving for the early consumer. The early consumer faces the risk of not being able to consume k at $t = 1$ if he finances the full project at $t = 0$. The third part provides a motive for risk sharing between the early and at least one late consumer that could resolve the tradeoff if the late consumer could fund, directly or indirectly, a part of the project. Late consumers cannot contribute investment funds at $t = 0$ because they enter the economy at $t = 1$.[3]

Banks and Markets.—We consider two institutions that intermediate between the firm and the consumers: banks (B) and markets (M). The role of these institutions is to facilitate risk sharing between the early consumer and the late consumers so that the trade-off between the early consumer's need for liquidity and the firm's need for investment funds is alleviated. The firm can choose whether to raise funds from a bank or in the market, depending on which of the two offers better terms to the firm. In our model banks and markets are distinguished by their expertise in interpreting and disclosing information as we explain next.

Information Structure.—If the firm goes to a bank to borrow w in order to invest in a worthy project at $t = 0$ the bank receives a *file* of the projects that contains all financial statements that are needed to identify which project is worth investing in. The same file will also be presented to a market agent if the firm chooses to seek funding from the market. The file contains information that is relevant for evaluating the project, but also some information that reveals whether the project will fail, which requires *expertise* to understand. The bank has a low-tech information production technology, while late consumers (hedge funds for example) have a high-tech technology. We capture these differences in *interpreting* the file in the next assumption.

ASSUMPTION 2: *Based on the project's file*, *a bank and a market agent can determine which of the two firm's projects is worthy. Only late consumers in possession of the same file have the expertise to determine at* $t = 1$ *whether the project will be a success or a failure at* $t = 2$.

The early consumer knows that both the bank and the market agent can identify the worthy project at $t = 0$ and does not need to see the file to know that by lending w he is financing a worthy project. If the bank does not show the file of the project to any late consumer at $t = 1$, then the payoff of the project is not revealed until it is realized at $t = 2$. This is however a decision the bank can make. The next assumption shows the distinction between the bank and the market in terms of *disclosing* the file.

[3] Alternatively we could have assumed that late consumers enter at $t = 0$ but cannot contribute investment funds because they cannot pledge their future income due to information problems.

ASSUMPTION 3: *Banks can keep the files of projects secret if they choose to do so. Markets cannot keep the files secret.*

The event that causes variation in the project's valuation at $t = 1$, and that induces banks to hide information in order to avoid such variation, is the arrival of late consumers, who have the expertise to process the file's information better than banks, the early consumer or a market issuer.[4] While the file's information is beneficial to identify the worthy project at $t = 0$ (capturing this benefit of information is the single role of lemons in the model), the same information may cause valuation variation at $t = 1$ that hinders risk-sharing between the early and late consumers.

Here is an example to illustrate the information structure. A bank is making commercial real estate loans in order for the borrower to buy office buildings in large cities. Offices in each city will be leased out by a management company in that city. The bank is not an expert in commercial real estate management companies. It relies on appraisals and financial information about the borrower but does not know how to evaluate real estate management companies. From the bank's point of view the loans are all of AAA quality. However, a hedge fund that trades commercial real estate loans (there is a market for such loans), knows about management companies; some are good and some are bad. If the hedge fund knew which buildings were purchased by the borrower, and which management companies were managing each building (information available in the loan application) then the hedge fund could distinguish more finely the loans and determine whether the bank's portfolio will be successful at $t = 2$.

Finally we assume that markets operate in a centralized way with many late consumers participating simultaneously, while banks operate in a decentralized way one consumer at a time.

ASSUMPTION 4: *All late consumers interact in the market simultaneously. Only one late consumer* (*chosen at random*) *interacts with the bank.*

Our stylized informational assumptions are meant to capture the reality that banks cannot produce price information with the same integrity as markets, which are informationally efficient due to competition and information aggregation in a centralized environment. On the other hand, markets cannot keep information hidden the way banks can.

Before turning to a preliminary analysis of the equilibrium outcomes in markets and in banking, respectively, we describe two useful benchmarks: autarky and first best.

Autarky.—In this case early and late consumers just store their endowments and do not interact with the firm (no intermediation), so $E(U_F^A) = 0$ and $E(U_E^A) = E(U_L^{l,A}) = e + \alpha k$ for all late consumers l.

[4] An alternative assumption would have potentially relevant public information arrive at $t = 1$ which only the late consumer can interpret provided that he has access to the full project file. This more elaborate structure is more realistic, but is isomorphic to our model.

First Best.—If risk sharing is feasible the economy can do better than autarky. Consider the problem of an unconstrained social planner, who can make transfers across consumers; so the planner does not have to satisfy individual participation constraints. At $t = 0$ the planner transfers an amount w from the early consumer to the firm so that the worthy project is fully financed. The early consumer saves $z \equiv e - w < k$. Since it is optimal for the early consumer to consume k in period $t = 1$, the planner transfers a total of $k - z$ from the late consumers to the early consumer at $t = 1$, regardless of whether the project will be a success at $t = 2$. By Assumption 1 these allocations are feasible. Assuming that the social planner assigns the social surplus (relative to autarky) to the firm, the first best ex ante expected utilities are $E(U_F^{FB}) = \lambda x - w$ and $E(U_E^{FB}) = E(U_L^{l,FB}) = e + \alpha k$ for all late consumers l.

B. *Capital Markets*

Table 1 shows the sequence of events in capital markets. At $t = 0$ the firm approaches a *market agent*, who obtains the firm's file, identifies the worthy project and makes the file publicly available to investors. In this section we assume that the firm seeks funds to finance the project fully (in Section III we allow for fractional investment, $\eta^M < 1$). The firm offers a security that promises $s^M(y) \geq 0$, where the superscript M refers to the market. This security is contingent on the project outcome y at $t = 2$, where $y \in \{b, g\}$ can be *bad* or *good*. Formally, the firm's strategy is denoted by $f_0^F: \{s^M(b), s^M(g)\}$, subject to limited liability (i.e., $s^M(b) \leq 0$ and $s^M(g) \leq x$).

There is a *market maker* that intermediates between the firm and the consumers, who underwrites the security on behalf of the firm at $t = 0$, making a take-it-or-leave-it offer to the early consumer. The early consumer's strategy at $t = 0$ is whether or not to buy the security offered in the market. Formally, $f_0^E : \{s^M(b), s^M(g)\} \to \{$*buy*, *do not buy*$\}$.

At $t = 1$ the early consumer holding the security wants to consume k but has only saved $z < k$, so he wants to sell a fraction $\theta(y)$ of the security to late consumers, who have full information about the state $y \in \{b, g\}$ and play a Bertrand game that determines its unit price.

We take *subgame perfect equilibrium* as our equilibrium concept.[5] The strategies defined earlier maximize each agent's utility conditional on the information structure that characterizes capital markets.

PROPOSITION 1: *The equilibrium in capital markets displays fully revealing, state-contingent prices at* $t = 1$ *and, when the project is fully financed, it implements an allocation that generates a welfare loss relative to the first best of* $\min\{\alpha(1 - \lambda)(k - z), \lambda x - w\}$.

[5] The moves of the game are fully specified because a party always makes a take-it-or-leave-it offer.

Table 1—Timing: Model with Capital Markets

Date $t = 0$	The firm goes to the market agent, who identifies the worthy project from the firm's files. The firm raises w from the market agent by issuing a security that pays $s^M(b)$ in case of failure and $s^M(g)$ in case of success. The market agent obtains w by selling the security to E and making files public.
Date $t = 1$	All L observe the file and learn the state, b or g. E sells a fraction of his security in the market to L.
Date $t = 2$	Project payoffs are realized. The firm pays the security.

PROOF:

We proceed by backward induction. At $t = 1$, late consumers know the state and face Bertrand competition for the security, therefore prices are fully revealing. In the bad state, by limited liability, $s^M(b) = 0$. In the good state, the early consumer sells a fraction of the security to raise at least $\theta(g)s^M(g) = k - z > 0$, being indifferent between selling the rest or holding it until maturity at $t = 2$. Assumption 1 guarantees that a single late consumer does not face liquidity concerns at $t = 2$ and has enough funds to cover the extra liquidity needs of the early consumer. The fraction of security that each late consumer buys is indeterminate.

At $t = 0$ the early consumer chooses to buy the security (finance the project) or not. If the early consumer does not buy the security, he stores his endowment and obtains utility of $U_{E|Store} = e + \alpha k$. If the early consumer buys the security, then he faces a lottery as the project represents a risk, which will be reflected in the price at $t = 1$. For any $\theta(g)$, the early consumer's expected utility if buying the security is

$$U_{E|Finance} = (1+\alpha)z + \lambda[s^M(g) + \alpha(k-z)].$$

The early consumer buys the security if and only if $U_{E|Finance} \geq U_{E|Store}$, which together with limited liability $s^M(g) \leq x$ implies

$$s^M(g) = \min\left\{\frac{w}{\lambda} + \frac{\alpha(1-\lambda)}{\lambda}(k-z), x\right\}. \tag{1}$$

Now that we have characterized the equilibrium, we can evaluate welfare. As is clear from equation (1), $\lambda s^M(g) > w$ because the firm has to compensate the early consumer for taking the risk of not consuming as much as desired at $t = 1$. The welfare loss relative to the first best outcome is then

$$E(U_F^{FB}) - E(U_F^M) = \lambda s^M(g) - w = \min\{\alpha(1-\lambda)(k-z), \lambda x - w\}. \blacksquare$$

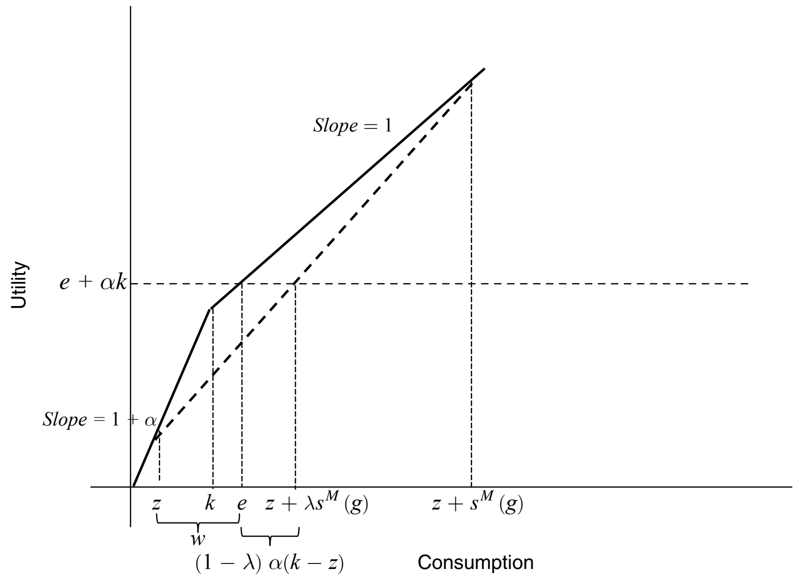

FIGURE 1. COSTS OF CAPITAL MARKETS

Figure 1 illustrates the equilibrium in Proposition 1. The figure depicts the early consumer's contract and consumption at date 1. In the bad state $s^M(b) = 0$, since the market price in this fully-revealing equilibrium cannot be positive in the bad state. In this state the early consumer only consumes $z < k$, which is what he saved after having funded the project. This pins down the left end of the dotted line. The right end of dotted line is given by the early consumer's total consumption in the good state $z + s^M(g)$ and is pinned down by its payment in the good state, $s^M(g)$. This payment must be set so that the weighted average between the end points of the dotted line fall on the early consumer's reservation utility line, the horizontal line $e + \alpha k$. For now, we assume the payment in the good state is feasible.

The figure shows that the early consumer is risk averse with respect to the payment $s^M(g)$, because $z < k$. For this reason, the early consumer has to be paid in expectation more than the loan size, w. The risk premium is the segment $(1 - \lambda)\alpha(k - z) > 0$. Note that the risk premium is smaller the higher is z. If z were equal to or greater than k, the early consumer would not be concerned about the risk. The welfare loss stems from the need to pay the early consumer a risk premium. If the risk premium of fully funding the project requires that $s^M(g) > x$, then full scale investment is infeasible. Once we allow for fractional investments in Section III, the solution is for the early consumer to save k and invest only $e - k$ at date 0 (this is, just to finance a fraction $\eta^M = \frac{e-k}{w} < 1$ of the project). Because the early consumer behaves risk-neutrally above k and the project is constant-returns to scale, a scaled-down investment is always preferred to not financing the project.

C. *Banks*

The problem with capital markets is that they reveal too much information too early. In the fully-revealing equilibrium markets provide *risky liquidity*. We now show that banks can achieve the first-best outcome, providing the early consumer with *safe liquidity*, always guaranteeing consumption k and full financing. So banks will dominate markets. The key assumption is that a bank can hide its information from the late consumer. In the next section we turn to the main case of interest, one where the late consumer can acquire private information (at a cost) and get access to the firm's file also in the banking regime. In this case, markets will sometimes dominate banks.

Table 2 shows the sequence of events with bank financing. At $t = 0$ the firm approaches a bank which obtains the firm's file, identifies the worthy project and can *hide* the file. The firm seeks to finance the worthy project fully by issuing a contingent security that pays $s^B(b)$ in case of failure and $s^B(g)$ in case of success at $t = 2$.

To obtain funds for investing in the firm, the bank turns to the early consumer, making him a take-it-or-leave-it offer. The offer asks E to deposit his endowment e with the bank in exchange for receiving a non-contingent payment at date 1 and a state-contingent payment at date 2. In addition, the contract specifies what the bank commits to offer to the (single) late consumer (still unknown) who will arrive at the bank at date 1.[6] We denote the promises by r_1^E at $t = 1$ and contingent claims $r_2^i(y)$ at $t = 2$, for $i \in \{E, L\}$ in good and bad states $(y \in \{b, g\})$.[7] Formally, $f_0^B : \{s^B(b), s^B(g)\} \rightarrow \{r_1^E, r_2^i(y)\}$ for $i \in \{E, L\}$ and $y \in \{b, g\}$ at $t = 0$.

At $t = 1$ the bank decides whether to reveal the file (rev) to the late consumer or keep it secret (sec). Conditional on the banking contract to both consumers and the bank's revelation decision, the late consumer decides whether to deposit e in the bank, that is $f^L : \{r_1^E, r_2^i(y)\} \times \{sec, rev\} \rightarrow \{deposit,\ do\ not\ deposit\}$. At $t = 1$, the early consumer withdraws r_1^E from the bank provided the late consumer deposits (if not, she has rejected the bank's offer so he gets paid nothing). At $t = 2$, the outcome of the project is revealed to all parties and everyone is paid according to contract.

Note that the early consumer trades indirectly with the late consumer by withdrawing r_1^E from the bank. Alternatively, and equivalently, the early consumer could trade directly with the late consumer by writing her a check or using a bank note issued by the bank. The key is that neither consumer observes the information that the bank obtained at $t = 0$ unless the bank reveals the file. The bank, by hiding this information permits efficient risk-sharing between the consumers, covering the

[6] The first-best contract is non-contingent at date 1 and we will show that this contract can be implemented by a bank.

[7] Even though the late consumer is not present at $t = 0$, one can interpret our assumption that the bank can commit to an offer to the late consumer as part of the contract with the early consumer, as the bank's deposit policy toward depositors entering at date 1. That the contract commits to payments to the late consumer at date 2 discourages banks to show her the file at $t = 2$ and it is motivated by the presumption that it is easier for a consumer to detect changes in contract provisions than to detect communication of information.

TABLE 2—TIMING: MODEL WITH FINANCIAL INTERMEDIARIES

Date $t = 0$	E deposits e with the bank. The bank promises a non-contingent payment r_1^E at $t = 1$ and a conditional payment $r_2^E(g)$ if the worthy project succeeds and $r_2^E(b)$ if it fails at $t = 2$. The firm goes to the bank, who identifies the worthy project from the firm's files. The bank lends w to the firm with a loan contract, where the firm pays $s^B(b)$ if the project fails and $s^B(g)$ if is succeeds. The bank chooses whether to keep the firm's files secret.
Date $t = 1$	L deposits e and the bank promises a contingent payment $r_2^L(g)$ if the project succeeds and $r_2^L(b)$ if the project fails at $t = 2$. E withdraws r_1^E and consumes.
Date $t = 2$	Project payoffs are realized. The firm pays the security and the bank pays to the consumers.

early consumer's liquidity needs at date 1 regardless of the state. This is a variation of Hirshleifer (1971), applied to a banking setting.

The next proposition shows that banks can implement the first-best allocation.

PROPOSITION 2: *There is a subgame perfect equilibrium in which the bank, by keeping the firm's file secret, permits first-best implementation: the firm is fully funded and the early consumer's liquidity needs are fully covered.*

PROOF:

We proceed to construct a subgame perfect equilibrium which implements the first-best allocation and gives all the surplus net of autarky to the firm. Note that the bank, by accepting the early consumer's deposits at $t = 0$ commits itself to a contract with the late consumer at $t = 1$. With this later offer fixed, the bank has nothing to gain by showing her the firm's file. So, it keeps the firm's file secret and there is no new information revealed at $t = 1$.

In the equilibrium both consumers are asked to deposit their endowments e – the early consumer at $t = 0$ and the late consumer at $t = 1$. This permits the firm to invest fully in the project. Because consumers are risk neutral once their interim period consumption k has been guaranteed, there are many alternative contracts that make consumers indifferent. We use this flexibility to fix $r_1^E = k$ and $r_2^E(b) = 0$ for reasons that will become clear in the next section. The firm's limited liability constraint implies that $s^B(b) = 0$.

We proceed to determine by backward induction the rest of the payments using aggregate resource constraints, the bank's and the consumers' reservation utilities and the maximization of the firm's expected utility.

If the project fails, the bank's assets at $t = 2$ are

$$A_b \equiv e + z - k \qquad \text{where} \quad z = e - w.$$

By Assumption 1, A_b is greater than k and smaller than e. If the project succeeds the bank's assets at $t = 2$ are $A_b + s^B(g)$. Because we chose $r_1^E = k$ and $r_2^E(b) = 0$, the early consumer is willing to deposit e if and only if $(1+\alpha)k + \lambda r_2^E(g) \geq e + \alpha k$. The early consumer's break even constraint then implies

$$r_2^E(g) = \frac{e-k}{\lambda}. \tag{2}$$

Since $r_2^E(b) = 0$ and the bank breaks even, feasibility in the bad state together with Assumption 1 (part iii) implies

$$r_2^L(b) = A_b > k. \tag{3}$$

The late consumer deposits if and only if $(1+\alpha)k + (1-\lambda)(A_b - k) + \lambda(r_2^L(g) - k) \geq e + \alpha k$. Therefore her break-even constraint implies

$$r_2^L(g) = e + \frac{(1-\lambda)}{\lambda}[e - A_b] > e > k. \tag{4}$$

This completes the characterization of the banking contract that induces consumers to deposit. The contract is feasible and, as equations (3) and (4) show, the late consumer also gets his liquidity needs covered in both states as required of a first-best contract.

Finally, we can determine the firm's strategy. Because the firm makes a take-it-or-leave-it offer to the bank it can extract all the surplus from the bank (and, via the bank, both consumers' surplus as well). We have established that the banking contract is feasible in the bad state (derivation of equation (3)). It will also be feasible in the good state whenever $r_2^E(g) + r_2^L(g) \leq A_b + s^B(g)$, or

$$s^B(g) = \frac{w}{\lambda}.$$

Notice that the firm's limited liability constraint is not binding as the project has a positive NPV implying that the project can be implemented at full scale.

In this equilibrium the firms' surplus is $E(U_F^B) = \lambda x - \lambda s^B(g) = \lambda x - w$, hence the bank contract with secret keeping implements the first best. ∎

When the bank keeps the firm's information secret, it delays information revelation about the project's returns, which assures that the early consumer receives a non-contingent payment at date 1 according to the promise the bank made at date 0. Banks provide *safe liquidity*, that is, *money*-like securities in contrast to markets, which provide *risky liquidity*.[8]

[8] In our three period model, late consumers are effectively bank equity holders, while early consumers are part depositors and part equity holders. In the online Appendix we extend the model to one with an overlapping generations (still living for three periods) structure, with a (stochastic) delay between the time a project is financed and the time it pays off. In this extension all but the last two generations before projects mature (which is unknown) will receive payments that do not depend on the projects' outcome. This corresponds more closely to standard demand deposits whose payments are independent of the performance of the bank portfolio.

II. Private Information Acquisition

We are now in a position to study the main case of interest: the late consumer can privately learn the bank's information about the firm (have access to its files) by exerting costly effort γ in terms of consumption. What are the benefits to the late consumer of finding out the bank's information? If the late consumer does not acquire private information and just deposits in the bank according to the banking contract of Proposition 2, her expected utility is

$$(1+\alpha)k + \lambda(r_2^L(g) - k) + (1-\lambda)(r_2^L(b) - k).$$

If the late consumer deviates and acquires information at a cost γ she will find out with certainty whether the project is successful. She prefers to deposit in the bank if and only if the state is good, obtaining $r_2^L(g) > e$ at $t = 2$ (from equation (4)). If she finds out that the state is bad, she prefers to store her endowment e rather than deposit in the bank, and to consume $e > r_2^L(b)$ at $t = 2$ (from equation (3)). The expected payoff if the late consumer acquires information is therefore

$$(1+\alpha)k + \lambda(r_2^L(g) - k) + (1-\lambda)(e-k) - \gamma.$$

Comparing the two payoffs, the late consumer deposits her endowment without acquiring information if and only if

$$r_2^L(b) \geq e - \frac{\gamma}{1-\lambda}. \tag{5}$$

Equation (5) shows that conditional on guaranteeing the early consumer liquidity k at $t = 1$ and the minimum reservation utility, the bank wants to distribute his payments so the late consumer's consumption in the bad state is maximal, because this will minimize her incentives to acquire information. This is also why we chose $r_2^E(b) = 0$ in the first-best contract to be implemented in Proposition 2.

Substituting $r_2^L(b) = A_b = e + z - k$ (equation (3) from the banking contract that implements the first best) into equation (5), we can write the incentive constraint as

$$(1-\lambda)(k-z) \leq \gamma. \tag{6}$$

The left-hand side is the expected value of acquiring information: if the late consumer learns that the state is bad she will avoid the loss from depositing, which happens with probability $1-\lambda$. The late consumer will not acquire information if this value is less than the information cost. If we define the incentive for acquiring information as

$$\Psi \equiv k - z - \frac{\gamma}{1-\lambda}, \tag{7}$$

the banking contract that implements the first-best allocation in Proposition 2 is feasible if and only if $\Psi \leq 0$. If this condition is not fulfilled, banks cannot credibly promise to pay k to the early consumer at $t = 1$ as the late consumer would have an incentive to learn about the state, and not deposit if the state is bad. In that case the bank could not guarantee liquidity k for the early consumer at date 1 and the first-best allocation would not be implemented.

From equation (6) we see that banks are more likely to implement the first-best allocation when: (i) projects have a low probability of default (high λ), (ii) they are difficult to monitor (high γ), (iii) they are relatively small (low w), (iv) the liquidity needs are relatively small (low k) or (v) the early consumer is relatively rich (high e). That is, relatively safe, small and complex projects are more likely to be observed in the portfolios of banks.

When a late consumer can privately acquire information, the banking contract needs to satisfy the additional restriction (6) in order to implement the first-best contract of Proposition 2. If (6) is violated by the first-best allocation, the bank has various options to adjust its contract so that the late consumer no longer has the incentive to acquire information. We study two such possibilities in the next subsections. In the first, banks can distort investment by scaling down its level (firms are not fully financed). In the second, banks distort money provision (by promising a lower level of safe liquidity at $t = 1$). Having analyzed the best options for the bank to fend off private information acquisition, we go on to study (in Section III) how the costs of such measures affect the firm's choice of whether to finance the investment through banks or capital markets.

A. *Banks Distorting Investment*

The bank can relax the incentives for the late consumer to acquire information by investing in just a fraction η^B of the project and storing the rest (for example in Treasury bonds or other safe assets). The reason is that that the bank can promise a higher payment $r_2^L(b)$ in the bad state, while maintaining the promise to pay $r_1^E = k$ at date 1, as we will show below.

LEMMA 1: *Suppose the late consumer's incentive to acquire information is strictly positive* ($\Psi > 0$). *Banks can prevent information acquisition by scaling down its investment in the project to a fraction* $\eta = 1 - \frac{\Psi}{w}$. *The resulting allocation generates a welfare loss relative to first best of* $\left(\frac{\lambda x}{w} - 1\right)\Psi$.

PROOF:

From equation (5), the bank can discourage information acquisition by promising the late consumer (no less than)

$$r_2^L(b) = e - \frac{\gamma}{1-\lambda}. \tag{8}$$

The bank stores (publicly) a fraction $(1 - \eta^B)$ of the full cost of the project, making this amount available for disbursement at $t = 2$. If the bank continues to supply full insurance ($r_1^E = k$), the late consumer can be paid in the bad state at most

$$r_2^L(b) = 2e - \eta^B w - k.$$

Substituting this expression into equation (8) gives us the maximal investment scale that still allows the early consumer to consume k in each state at date 1:

$$\eta^B = \frac{e-k}{w} + \frac{\gamma}{w(1-\lambda)} = 1 - \frac{\Psi}{w} < 1. \tag{9}$$

Since the rest of the original first-best contract and the utilities of the two consumers remain unchanged (that is, equal to their utility in autarky), the firm's loss relative to the first best is

$$E(U_F^{FB}) - E(U_F^I) = (1-\eta^B)(\lambda x - w) = \frac{\Psi}{w}(\lambda x - w).$$

This expression shows that the opportunity cost of the funds Ψ that have to be guaranteed in the bad state to discourage the late consumer from acquiring information is the expected rate of return of the project, $\frac{\lambda x}{w} - 1$. ∎

B. *Banks Distorting Money Provision*

An alternative way to discourage the late consumer from acquiring information about the firm is to distort money provision, which manifests itself as less safe liquidity (a reduction in the amount available for withdrawal by the early consumer at $t = 1$). Because the early consumer will not be able to cover all his liquidity needs at date 1, the firm has to compensate him for the loss by paying him more at date 2, when he values consumption less. This will lead to higher funding costs.

LEMMA 2: *Suppose the late consumer's incentive to acquire information is strictly positive* ($\Psi > 0$). *Banks can prevent information acquisition by reducing the early consumer's date* 1 *safe liquidity to* $r_1^E = k - \Psi$. *The resulting allocation generates a welfare loss relative to first best of* $\alpha\Psi$.

PROOF:

To discourage information acquisition, the bank has to promise the late consumer at least $r_2^L(b)$ in equation (8). At full scale ($\eta^B = 1$), equation (3) implies that $r_2^L(b) = e + z - r_1^E = e - \frac{\gamma}{1-\lambda}$, or

$$r_1^E = \frac{\gamma}{1-\lambda} + z = k - \Psi. \tag{10}$$

To make the early consumer whole, the bank has to offer him a larger payment at $t = 2$ in the case the project succeeds. In particular the early consumer will deposit in the bank if and only if $(1+\alpha) r_1^E + \lambda r_2^E(g) \geq e + \alpha k$. Replacing r_1^E (from equation (10)) above, implies

$$r_2^E(g) = \frac{e-k}{\lambda} + \frac{(1+\alpha)}{\lambda}\Psi. \tag{11}$$

Similarly, the late consumer will deposit her endowment in the bank if and only if $(1+\alpha)k + \lambda(r_2^L(g) - k) + (1-\lambda)(r_2^L(b) - k) = e + \alpha k$. Since $r_2^L(b) = e - \frac{\gamma}{1-\lambda}$, we have

$$r_2^L(g) = e + \frac{\gamma}{\lambda}. \tag{12}$$

Next we check that these payments are feasible in the good state, which requires

$$r_2^E(g) + r_2^L(g) \leq e + z - r_1^E + s^B(g).$$

The firm will extract the maximum surplus by selling the bank a security that promises

$$s^B(g) = \frac{w}{\lambda} + \frac{\alpha}{\lambda}\Psi. \tag{13}$$

Since the investment is at full scale and the utilities of the two consumers are unchanged by construction, the firm's loss relative to the first best is

$$E(U_F^{FB}) - E(U_F^{LP}) = \lambda s^B(g) - w = \alpha\Psi. \blacksquare$$

Finally, we specify conditions under which banks distort money provision rather than investment. Combining Lemmata 1 and 2, we get Corollary 1.

COROLLARY 1: *Banks distort money provision instead of investment if and only if*

$$\frac{\lambda x}{w} - 1 \geq \alpha.$$

Intuitively, banks distort money provision rather than investment when the welfare costs in terms of liquidity needs (captured by α) are lower than the welfare costs in terms of investment (captured by the NPV of the project $\frac{\lambda x}{w} - 1$). The bank is more likely to distort money provision when liquidity needs are small (low α), when the unit cost of the project is small (low w) and when a project is more likely to succeed or generate a high payoff in case of success (high λ or high x).

Could the bank do better by simultaneously distorting investment and money provision? The answer is no. Suppose that the bank has chosen to reduce money provision to prevent information acquisition ($r_1^E < k$). If the bank now reduces the investment scale by one dollar, it can use the dollar to reduce money distortion by one dollar (i.e., raising r_1^E one dollar), without affecting the late consumer's incentive to acquire information, which only depends on $r_2^L(b)$ (equation (5)). So the question is only whether investment distortion is desirable regardless of money provision. If it is desirable, we know that it is optimal to scale down the investment to the point where the early consumer's liquidity needs are fully covered. Consequently, there is no need to reduce money provision when investment distortion is the more efficient option.[9]

III. Banks versus Markets

A. *When Do Markets Provide Cheaper Funding than Banks?*

In this section we discuss conditions under which firms prefer to raise funds from capital markets rather than from banks. Recall that banks dominate markets whenever they do not have to distort contracts to prevent information acquisition, that is, whenever $k - z < \frac{\gamma}{1-\lambda}$. In order to have a fair comparison between banks and markets when banks distort investment, we need to give markets the option to fund projects at a reduced scale. The next lemma allows for a partial financing and generalizes Proposition 1, showing the conditions under which capital markets prefer to finance only a fraction of the project.

LEMMA 3: *Consider the possibility of financing a fraction of the project in capital markets. If* $\frac{\lambda x}{w} - 1 \geq \alpha(1-\lambda)$ *the project is financed at full scale. Otherwise only the fraction* $\frac{e-k}{w}$ *of the project is financed. The equilibrium in the capital market generates a welfare loss relative to the first best of* $\Omega \equiv \min\left\{\alpha(1-\lambda), \frac{\lambda x}{w} - 1\right\} \times (k-z)$.

PROOF:

Suppose the early consumer saves k and only invests the balance $e - k < w$ in the project. The scale of the project is then $\eta^M = \frac{e-k}{w}$. Since the early consumer is not facing any liquidity risk at date 1, he is willing to make this investment if he gets in return its expected value $\eta^M \lambda s^M(g)$, where $s^M(g) = \frac{w}{\lambda}$ is the equilibrium market price in the good state.[10]

The firm's expected profit from financing the fraction η^M of the project in the capital market is $\eta^M \lambda x - \lambda s^M(g|\eta^M)$ where, from Proposition 1,

$$\lambda s^M(g|\eta^M) = \eta^M w + \alpha(1-\lambda)\min\{k - e + \eta^M w, 0\}.$$

Maximizing the firm's profit with respect to η^M, we see that because of linearity the solution is bang-bang. The firm fully finances the project if and only if $\frac{\lambda x - w}{w} \geq \alpha(1-\lambda)$. Otherwise the firm invests the fraction $\frac{e-k}{w}$ defined above. In the first case, the loss from funding in the capital market is $\alpha(1-\lambda)(k-z)$ due to riskiness of the investment. In the second case, with fractional investment, there is no loss due to risk but instead the loss comes from underinvestment and is $(1-\eta^M)(\lambda x - w) = \frac{\lambda x - w}{w}(k-z)$. The loss in capital markets is then the lower of these two expressions. ∎

[9] This argument relies on the (piecewise) linearity of the utility functions and the linear divisibility of the project. It would not go through with general concave utility and production functions.

[10] Note that the early consumer does not need to sell his share of the project at date 1 as he saves enough to cover his liquidity needs.

Combining Lemmata 1, 2, and 3, the next Proposition states the conditions under which a firm would prefer to fund projects in capital markets. A necessary condition is that $k - z > \frac{\gamma}{1-\lambda}$, that is, banks have to distort to prevent information acquisition.

PROPOSITION 3: *Suppose* $k - z > \frac{\gamma}{1-\lambda}$. *The firm prefers to fund the project in the capital market if and only if one of the following two conditions hold:* (*i*) $\frac{\lambda x}{w} - 1 < \alpha$ *and* $\Omega < \left(\frac{\lambda x}{w} - 1\right)\Psi$ *or* (*ii*) $\frac{\lambda x}{w} - 1 \geq \alpha$ *and* $\Omega < \alpha\Psi$. *If additionally* $\frac{\lambda x}{w} - 1 < \alpha(1-\lambda)$, *the capital market does not invest fully. If neither condition holds, the firm funds the investment through a bank, which scales down the investment in case* $\frac{\lambda x}{w} - 1 < \alpha$ *and reduces money provision if* $\frac{\lambda x}{w} - 1 \geq \alpha$.

Figure 2 illustrates why market funding can dominate bank funding for the set of parameters under which banks and capital markets would rather not reduce investment, that is $\frac{\lambda x}{w} - 1 > \alpha$. The figure is similar to Figure 1, displaying the early consumer's utility from reduced money provision (the 45 degree line in red). The two utility curves differ with respect to the location of the kink and the slope of the expected utility above the kink.

The kink in either utility function is determined by the level of consumption that the early consumer is guaranteed at $t = 1$, which is z in capital markets and r_1^E in banks. The slope of the expected utility above the kink depends on the marginal utility from the lottery that the early consumer faces. With a bank, the random payments are realized at $t = 2$ and therefore have a marginal utility of 1. In capital markets, the random payments are realized at $t = 1$, so the extra consumption in case of a good state $(k - z)$ provides marginal utility $1 + \alpha$. This implies that the slope of the utility payoff above the kink is in expectation larger for capital markets than for banks.

As in Figure 1, the welfare cost of financing the project is given by the gap between k and the consumption level for an expected utility that matches autarky. In Figure 2 this is the point at which the expected utility function cuts the horizontal line $e + \alpha k$. We see that the bank dominates the capital market whenever the solid red line crosses $e + \alpha k$ at lower consumption levels than the dashed black line. Naturally, when there are no distortions and $r_1^E = k$, banks dominate capital markets (the benchmark in Section I). At the other extreme, if $r_1^E = z$ (a large distortion), the solid line always crosses $e + \alpha k$ at larger consumption levels than capital markets (recall the expected utility has a smaller slope), and capital markets dominate banks. More generally, the less date-1 consumption (money-like assets) the bank is able to provide the more costly bank funding becomes. Proposition 3 provides the exact cut-off point where capital market funding becomes less costly than bank funding.

B. *A Graphical Representation of Proposition 3*

In this section we illustrate the set of parameters under which projects are financed by capital markets, by banks without distortions, and banks that distort either money provision or investment.

Figure 3 displays these regions as a function of the project characteristics λ_i and γ_i. Projects with negative NPV are not financed. This is the region in which $\lambda_i \in \left[0, \frac{w}{x}\right)$ at the left of the figure. Banks do not need to distort and can implement

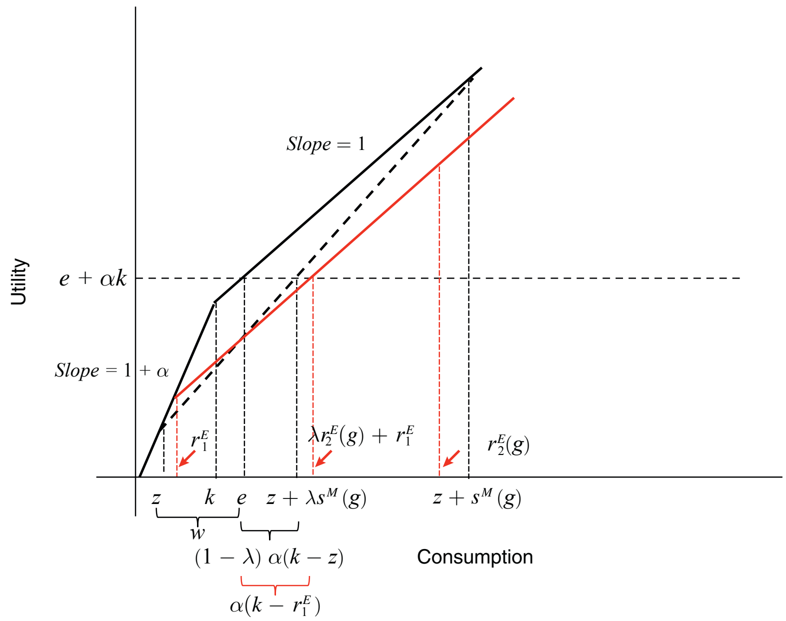

Figure 2. Case in which Capital Markets Dominate Banks

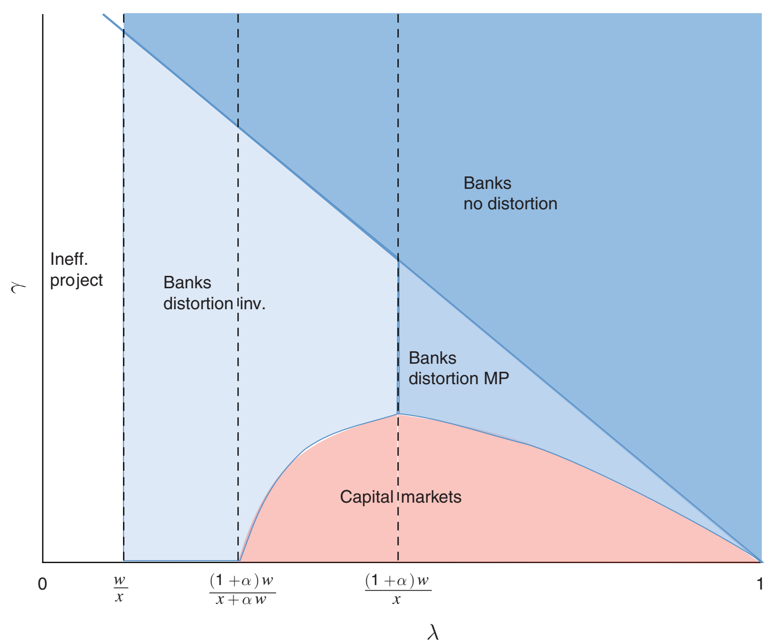

Figure 3. Regions of Financing

the first best allocation whenever the cost of information acquisition is large relative to the probability that the project defaults. This is the region in which $\Psi < 0$ (or $\gamma > (1-\lambda)(k-z)$) that corresponds to the upper right corner of the figure, with high γ and high λ. For the rest of the figure, banks have to introduce distortions, either to money provision or to investment. From Proposition 1 the banking distortion only depends on λ_i. Banks distort investment when investments have a relatively low NPV, when $\lambda_i \in \left[\frac{w}{x}, \frac{(1+\alpha)w}{x}\right)$, and distort money provision otherwise, when $\lambda_i \in \left[\frac{(1+\alpha)w}{x}, 1\right)$.[11]

When banks distort investment we can identify two subregions. In the first subregion $\left(\text{when } \lambda_i \in \left[\frac{w}{x}, \frac{(1+\alpha)w}{x+\alpha w}\right)\right)$ projects are not very likely to succeed and there is investment rationing both by capital markets and by banks. Investments, however, are rationed more in capital markets than in banks. This is because the fraction of the project that is not financed in capital markets is proportional to $(k-z)$, while the fraction of the project that is not financed by the bank is proportional to $\Psi = \left(k - z - \frac{\gamma}{1-\lambda}\right)$. Therefore, in this subregion of investment distortion banks strictly dominate capital markets whenever $\gamma > 0$.

In the second subregion $\left(\text{when } \lambda_i \in \left[\frac{(1+\alpha)w}{x+\alpha w}, \frac{(1+\alpha)w}{x}\right)\right)$ there is investment rationing by banks but not by capital markets. In capital markets welfare losses come from a lack of liquidity while in banks welfare losses come from distorted investments. In capital markets, as λ increases, the welfare loss from the lack of liquidity always declines, as the probability that the early consumer does not cover his liquidity needs is $(1-\lambda)$. In banks, as λ increases, the welfare loss from distorting investment has two opposing effects. On the one hand, there are more projects that are financed which reduces the losses. On the other hand, the projects that are not financed have a higher NPV, which increases the losses. This is why, for a fixed γ, as λ increases, capital markets dominate banks for a larger set of projects.

In the region of money-provision distortion, as λ increases, banks distort in proportion $\Psi = \left(k - z - \frac{\gamma}{1-\lambda}\right)$ while capital markets induce risks to the early consumer for $(1-\lambda)(k-z)$. As welfare losses decline in both cases but faster for banks capital markets dominate banks that distort money provision for a smaller set of projects.

One interpretation of γ is that it measures the opacity of the asset (e.g., the asset is hard to evaluate) and then is related to size and/or age, with larger and/or older firms having lower γ, which seems consistent with the life cycle of firms. According to the figure, this interpretation implies that when firms are young they are usually financed through banks. At this stage they are small, so it is relatively costly to produce information about them. As a firm grows there is more publicly available information (relations with suppliers, advertisements, etc.) that reduces γ to a point at which information is so widely available that the firm is better off going public.

[11] In the online Appendix we introduce the possibility of many aggregate states and show how their distribution changes these regions.

C. *Replication Possibilities*

Figure 3 highlights two regions. In one, banks that keep secrets dominate capital markets in terms of welfare, even when distorting. In the other, capital markets dominate banks that keep secrets. Since firms get the social surplus, they obtain funds at a lower rate from banks in the first region and from markets in the second region.

Can markets replicate banks in the region where banks dominate? The answer is no. Because detailed information is available publicly at $t = 1$, late consumers can always interpret the information and compete for the claims in the project, the market equilibrium at $t = 1$ will necessarily feature state-contingent prices.

Can banks replicate markets in the region where markets dominate? This would be possible only if the bank offers to repay at $t = 2$ whatever the random late consumer deposits at $t = 1$ and if the bank could reveal the detailed information to the late consumer at no cost. There would be replication in this case as the late consumer would only deposit funds in the bank if the state is good, in which case the bank could compensate the early consumer at $t = 1$ and repay the late consumer with the proceedings of the firm's claims at $t = 2$, exactly as in capital markets.

This is however a knife-edge situation. As long as there is at least one "naive" late consumer who is not able to interpret the file, she would never deposit in the bank, as the bank would use those funds to pay the early consumer at $t = 1$ and then not have enough resources to repay her in a bad state at $t = 2$. This slight departure from our assumptions would not change the allocations of capital markets (as the "expert" late consumers would compete for the project's claims and generate fully-revealing state contingent prices at $t = 1$) or the allocations of secret keeping banking (as it does not matter whether the random late consumer is naïve when the bank keeps the file in secret), but would prevent banks from replicating capital markets. In other words, as banks interact with a single random late consumer it is less effective in generating information with the integrity of markets, in which many late consumers aggregate information as they compete based on such information.

But even assuming that banks can reveal information at no cost and no late consumers are naive, bankers still cannot replicate markets if there is an opportunity cost of setting up a bank. In the online Appendix we study this possibility, in which an agent B can choose wether to set up a bank or endeavor in an alternative activity that pays ϕ. As the banking contract has to compensate agent B for becoming a bank, the welfare gains from keeping secrets have to compensate that cost. When there are no gains from secrets, there is no reason for setting up a bank that replicates what markets can do at no (opportunity) cost.

IV. Conclusions

Banks and capital markets are fundamentally different. There is price discovery in capital markets as competition among expert traders aggregates information efficiently. But price discovery is not conducive to producing securities that have a stable value and can be used for transactions and storing value. Capital markets produce *risky liquidity*. Banks cannot replicate the price discovery in capital markets, because there is no centralized trading which would aggregate information. So banks produce securities with a stable value, that is, *safe liquidity* which is useful for transactions and storing value. Note that in our model there is nothing unique about the banks' activities on the asset side (though banks may be screening and monitoring). But there is an important complementarity between bank assets and bank liabilities. In order to produce money, banks select assets to minimize information leakage and sensitivity to public and private information.

We argue that banks exist to produce money and this dictates the nature of bank assets, not the other way around as in Diamond (1984), for instance. The costs (of information production) and the benefits (with respect to the riskiness of the loan) determine the sorting of borrowers between capital markets and banks. The costs of private information production may be such that banks cannot produce the efficient amount of safe liquidity. Then banks reduce the amount of liquidity (or loans) that they produce. This may lead to banks altering the nature of their loans. For example, in US history banks have lengthened the maturity of their loans (Summers 1975).

That banks, by their very nature as money producers, are opaque is consistent with what we observe about banking throughout history. Entry into banking has always been restricted and banks have always been overseen by governments. While we do not model bank runs, banker moral hazard or other agency problems, our model provides a new rationale for opaque banking examinations, reserve requirements, capital requirements and deposit insurance, in terms of maintaining opacity. For example, reserve requirements create a fund that can be used in the bad state and relax the incentives to acquire information. One implication of this is that the bank can then make loans to borrowers with lower costs of information production. These are interesting and important issues, but beyond the scope of the paper.

The recent financial crisis further illustrates the opacity of shadow banking. The money produced by shadow banks, i.e., sale and repurchase agreements (repo) and asset-backed commercial paper (ABCP), are short-term stores of value. This money, when not backed by US Treasuries, is backed by asset-backed securities and mortgage-backed securities (ABS/MBS). ABS/MBS are complicated and opaque; and importantly have no traded equity which would reveal information. So, ABS/MBS are useful for backing repo and ABCP. And, the responses to the crisis maintained bank anonymity; emergency lending facilities were carefully designed to keep bank borrowers from having their identities revealed. Further, the SEC imposed short sale constraints to prevent information from being revealed.

REFERENCES

Andolfatto, David. 2010. "On the Social Cost of Transparency in Monetary Economies." Federal Reserve Bank of St. Louis Working Paper 2010–001.

Badertscher, Brad, Jeffrey Burks, and Peter Easton. 2015. "Day 30: The Tacit Quarterly Information Event in the Banking Industry." Unpublished.

Berger, Allen N., and Sally M. Davies. 1998. "The Information Content of Bank Examinations." *Journal of Financial Services Research* 14 (2): 117–44.

Bessler, Wolfgang, and Tom Nohel. 1996. "The Stock-Market Reaction to Dividend Cuts and Omissions by Commercial Banks." *Journal of Banking and Finance* 20 (9): 1485–508.

Boyd, John H., and Edward C. Prescott. 1986. "Financial Intermediary-Coalitions." *Journal of Economic Theory* 38 (2): 211–32.

Breton, Regis. 2007. "Monitoring and the Acceptability of Bank Money." Unpublished.

Breton, Regis. 2011. "A Smoke Screen Theory of Financial Intermediation." Banque de France Working Paper 356.

Dang, Tri Vi, Gary Gorton, and Bengt Holmström. 2013. "Ignorance, Debt and Financial Crises." Unpublished.

DeYoung, Robert, Mark J. Flannery, William W. Lang, and Sorin M. Sorescu. 2001. "The Information Content of Bank Exam Ratings and Subordinated Debt Prices." *Journal of Money, Credit, and Banking* 33 (4): 900–25.
Diamond, Douglas W. 1984. "Financial Intermediation and Delegated Monitoring." *Review of Economic Studies* 51 (3): 393–414.
Diamond, Douglas W., and Philip H. Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91 (3): 401–19.
Diamond, Douglas W., and Raghuram G. Rajan. 2001. "Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking." *Journal of Political Economy* 109 (2): 287–327.
Diamond, Douglas W., and Robert E. Verrecchia. 1991. "Disclosure, Liquidity, and the Cost of Capital." *Journal of Finance* 46 (4): 1325–59.
Flannery, Mark. 1998. "Using Market Information in Prudential Bank Supervision: A Review of the U.S. Empirical Evidence." *Journal of Money, Credit, and Banking* 30 (3): 273–305.
Flannery, Mark J., Simon H. Kwan, and M. Nimalendran. 2004. "Market Evidence on the Opaqueness of Banking Firms' Assets." *Journal of Financial Economics* 71 (3): 419–60.
Flannery, Mark J., Simon H. Kwan, and Mahendrarajah Nimalendran. 2013. "The 2007–2009 Financial Crisis and Bank Opaqueness." *Journal of Financial Intermediation* 22 (1): 55–84.
Gorton, Gary, and George Pennacchi. 1990. "Financial Intermediaries and Liquidity Creation." *Journal of Finance* 45 (1): 49–71.
Haggard, Stephen, and John Howe. 2007. "Are Banks Opaque?" Unpublished.
Haubrich, Joseph G., and Robert G. King. 1990. "Banking and Insurance." *Journal of Monetary Economics* 26 (3): 361–86.
Hirshleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61 (4): 561–74.
Hirtle, Beverly. 2006. "Stock Market Reaction to Financial Statement Certification by Bank Holding Company CEOs." *Journal of Money, Credit, and Banking* 38 (5): 1263–91.
Iannota, Giulano. 2006. "Testing for Opaqueness in the European Banking Industry: Evidence from Bond Credit Ratings." *Journal of Financial Services Research* 30 (3): 287–309.
Jacklin, Charles. 1987. "Demand Deposits, Trading Restrictions, and Risk Sharing." In *Contractual Arrangements for Intertemporal Trade*, edited by Edward Prescott and Neil Wallace, 26–47. Minneapolis: University of Minnesota Press.
Jones, Jeffrey S., Wayne Y. Lee, and Timothy J. Yeager. 2012. "Opaque Banks, Price Discovery, and Financial Instability." *Journal of Financial Intermediation* 21 (3): 383–408.
Kaplan, Todd R. 2006. "Why Banks Should Keep Secrets." *Economic Theory* 27 (2): 341–57.
Kashyap, Anil K., Raghuram Rajan, and Jeremy C. Stein. 2002. "Banks as Liquidity Providers: An Explanation for the Coexistence of Lending and Deposit-Taking." *Journal of Finance* 57 (1): 33–73.
Monnet, Cyril, and Erwan Quintin. 2013. "Rational Opacity." Unpublished.
Morgan, Donald P. 2002. "Rating Banks: Risk and Uncertainty in an Opaque Industry." *American Economic Review* 92 (4): 874–88.
Summers, Bruce. 1975. "Loan Commitments to Business in U.S. History." *Economic Review* 61 (September/October): 15–23.