**Knowledge, Technology Adoption and
Financial Innovation**

Ana Fernandes

05-13

June 2005

DISCUSSION PAPERS

Knowledge, Technology Adoption and Financial Innovation

Ana Fernandes*

University of Bern and Stern School of Business

June 2005

Abstract

Why are new financial instruments created? This paper proposes the view that financial development arises as a response to the contractual needs of emerging technologies. Exogenous technological progress generates a demand for new financial instruments in order to share risk or overcome private information, for example. A model of the dynamics of technology adoption and the evolution of financial instruments that support such adoption is presented. Early adoption may be required for financial markets to learn the technology; once learned, financial innovation boosts adoption further. Financial learning emerges as a source of technological diffusion. The analysis identifies a causality link from technology to growth which is nonetheless consistent with empirical findings of a positive effect of current financial development on future growth.

JEL Codes: G20, N20, O30.

Keywords: Technology adoption, financial innovation, learning.

*Address: <http://anafernandes.name>. This paper previously circulated under the title “What Does the Walrasian Auctioneer Know? Technology Adoption and Financial Innovation.” I received valuable comments from Thorsten Koepl, Ramon Marimon, seminar participants at the Chicago, Minneapolis, New York and Richmond Federal Reserve Banks, the First Banca D’Italia/CEPR conference on Money, Banking and Finance, CREI Innovation and Macroeconomic Workshops, the 2004 Winter Meetings of the Econometric Society, the University of Illinois at Urbana-Champaign and Stern Macroeconomics seminar. Financial support from *Fundação Para a Ciência e a Tecnologia* is very gratefully acknowledged.

1 Introduction

In most of the work addressing the relationship between economic progress and financial development, there is reference to the divide amongst well-known economists concerning the nature and importance of the relationship between those phenomena.¹ For example, in his *Theory of Economic History*, John Hicks [8] argues that the development of financial markets in England was a pivotal condition for the industrialization process started in 18th century England. Other classical references on the topic of growth and financial development include Joseph Schumpeter [19] and Joan Robinson [17]. While the views of the former are qualitatively similar to those of Hicks (finance spurs growth), Robinson argues that economic entrepreneurship leads to financial innovation.

This paper proposes the view that financial development and economic growth are linked through the characteristics of technology, as follows. Perhaps the most obvious connection between technology and financial innovation emerges through risk-sharing. Technology is modeled as a distribution function over output values. Technological progress occurs when Nature makes new distribution functions available to economic agents. In choosing which technology to operate, agents simultaneously select the risk profile of their income source. While progress allows higher output values to be attained, it also changes the risk profile faced by economic agents. The financial sector provides risk intermediation among agents who face distinct risk profiles.

How does financial intermediation affect technology adoption and, as a consequence, growth? The arrival of a new technology creates demand for a new risk-sharing contract. Such a contract is “new” in the sense that it requires the pricing of a different risk, previously inexistent in the economy. If financial institutions (called banks) are less knowledgeable about the new technology relative to entrepreneurs, they may fail to offer insurance when the technology is first available, or may only be able to offer some form of second-best insurance contract. The market of risk-sharing arrangements becomes, therefore, at least temporarily incomplete. In turn, an entrepreneur faced with the option of continuing to operate an old, less productive but fully insured technology, and a more productive but more risky one may well choose not to switch.

While banks start out less knowledgeable relative to entrepreneurs, they may *learn* about the new technology in several alternative ways. The simplest way to learn is to observe the outcome of adopters, the quantities they produce. By doing so, banks will gradually be able to offer more sophisticated contracts and eventually converge to the

¹See Levine [13] for a survey.

first-best. (In the simple setup of our model, banks only need one period to learn the necessary features of the technology and offer first-best contracts.) Another learning alternative comes from the implementation of second-best incentive-compatible contracts. By designing contracts whose transfers depend on a report from the entrepreneur, banks “learn” from the entrepreneur’s report itself, which the contract design ensures is truthful.

Given that technologies are modeled simply as two-point support stochastic processes, first-best contracts resemble a conventional fire insurance contract: the entrepreneur makes a payment to the bank every period and, if output is low, the bank pays the entrepreneur back. These contracts depend only on the observation of output. Second-best contracts additionally rely on the entrepreneur’s report concerning the realization of output, whether it was high or low. More generally, insurance constrained contracts suggest an added dimensionality and complexity relative to the first-best case. It is conceivable that the economy may be well prepared to enforce the first-best type of contract but that its institutions lack the ability to enforce the transitional – and more complex – type. If banks are to offer second-best insurance to entrepreneurs as the technology becomes available, there must be institutions in the economy with the capacity to enforce second-best arrangements. Once learning is complete, banks will resort to first-best contracts.

Two scenarios are considered. In the first one, the economy lacks the ability to enforce second-best contracts. As such, entrepreneurs who adopt the new technology (if any) will have to do so without insurance. For early adoption to occur, there must be sufficiently high skilled entrepreneurs in the operation of the new technology. Here, skill multiplies output; therefore, under logarithmic utility, a more skilled person becomes less risk averse in absolute terms. If the economy’s most able individuals are skilled enough so as to prefer the uninsured technology to the old one, adoption will take place. The observation of the output of adopters will, in turn, allow banks to learn and offer first-best insurance subsequently. Better contracts boost adoption further, and all entrepreneurs switch to the more productive technology. There is a feedback from adoption to finance, through learning, and then back to adoption. If no one is sufficiently skilled to adopt without insurance, no adoption will take place.

When the economy has the ability to enforce second-best contracts, whether or not they will be used depends on the stringency of the information constraints. One form of incentive-compatible contract is offering no transfers to entrepreneurs across all states; this, of course, is equivalent to providing no insurance. If the stringency of the infor-

mation constraints is such that no better contract can be offered, adoption dynamics resemble the earlier scenario. If, however, second-best contracts improve upon the no insurance case, they will meet with positive demand and the mass of early adopters will exceed the case when no second-best insurance is provided. Second-best contracts make a crucial difference when no one is sufficiently skilled to adopt without some form of insurance. By providing at least partial insurance, these contracts may make adoption possible when it otherwise would not be.

The analysis provides a structural interpretation for the relationship between growth (technology adoption) and finance. Here, ongoing technological progress renders financial markets incomplete. Financial innovation is understood in this context as the progressive completion of financial markets. Technology adoption allows banks to learn about the technology and improve on the financial instruments offered. In turn this reinforces subsequent adoption. The model identifies one source of causality from technology to finance: new contracts are required only when technology evolves.² Financial innovation is necessary for technology adoption – at least for adoption to take full scale – but it is not sufficient: if markets are complete, no new combination of existing contracts will affect adoption decisions. The response of finance to the needs of technology seems to be perfectly summarized in Joan Robinson’s [17] claim that “where enterprise leads finance follows.” In the present context, enterprise is understood both as a new technology in a strict sense (that of scientific discovery), or as a new form of organizing one’s business or conducting trade. While contracts do evolve and enhance adoption, better finance is the results of previous technological innovation, in this model.

The model is also consistent with the positive correlation of measures of financial development and growth, as documented in King and Levine [10] and Levine and Zervos [12], for example. In fact, the intensity of usage of financial contracts is the flip side of the coin of technology adoption. Further, to the extent that empirical measures of financial development are correlated with the learning ability of the financial sector, the model also predicts that today’s financial development is positively correlated with the growth rate of output (or other indicators of economic development) in the future (see [10]). Arguably, measures of financial development could be capturing to some extent the financial sector’s ability to learn about new technologies and to offer second-best contracts, facilitating the adoption of new technologies. As such, countries with greater

²Since the rate of arrival of new technologies is exogenous, in this model financial markets do not influence the intensity of research and development activities; as such, there is no causality in the reverse direction.

financial development would experience faster technology adoption and higher growth rates, present and future. Although this paper is not in any way dismissive of the likely possibility that a well-functioning financial market may enhance the endogenous intensity of research and development, providing yet another explanation for the aforementioned empirical findings, it shows that such findings are also consistent with finance being a consequence of technological progress.

This paper also identifies a new channel for technological diffusion (adoption with a lag). The fact that banks are less informed than entrepreneurs and the inherent necessity for them to learn the technology is responsible for a lag between the period the technology is made available and first-best insurance contracts are offered (provided adoption does occur). This period of “financial learning” is responsible for the gradual adoption of new, more productive technologies.

The model analyzes the relationship between financial innovation and technology adoption from the point of view of risk-sharing arrangements. There are other dimensions of technology that link financial arrangements to technology adoption. One example is asymmetric information. To the extent that technology forces shareholders to delegate on a manager the ability to run their firm given his greater expertise, contracts must be designed to convey adequate incentives to the manager. The literature on corporate finance addresses precisely the properties of such contracts. Yet another example is the presence of indivisibilities (investments that require the commitment of capital for long periods of time), as they require the matching of the different liquidity requirements of investors and savers over time. We next discuss two historical episodes where we believe the relationship between financial innovation and growth suggested by our model is transparent. The Industrial Revolution is an example of the indivisibility/liquidity case, whereas the Chicago Board of Trade example is an instance of price risk.

According to Hicks [8], the core feature of modern industry, born in England’s Industrial Revolution, is the fact that fixed capital takes center stage and replaces circulating capital in the production process. In turn, financing fixed capital required the commitment of sizeable investments for long periods of time. Partly as a consequence of the Revolution of 1688, “that established Parliament as the key agency in managing national fiscal affairs,”³ as well as the need to finance the British warfare, financial markets in England experienced significant developments. The financial revolution was centered around the crucial reforms regulating the placement of public debt, taking place between the first placement of government-backed annuities, in 1693-94, and the establishment

³Baskin and Miranti [3].

of the ‘Three Per Cent Consol,’ of 1751. From as early as 1698, “[...] prices of all the major government securities as traded on the London stock exchange were printed twice weekly and circulated widely through London.”⁴

The crucial role played by a liquid and transparent market, whose prices were widely circulated, is also a subtle one. The entrepreneurs’ needs for external financing, arising from the need to build up their fixed capital stocks, were not – at this time – directly met by issuing equity on the firm at the London stock market. Entrepreneurs relied first on the net supply of their circulating capital, then on credit from friends or institutions to meet their demand for fixed capital. However, the wide availability and reliability of public debt made it the reference asset for all other financial operations, namely those associated with credit. Intense exchange in the public debt market also set the stage for savers to seek other – more profitable – alternatives.

Credit to firms was granted using as support a number of other instruments also introduced in England at the time of the Glorious Revolution. The foreign bill of exchange, for example, became the major source of credit for British merchants engaged in the growing trade with North Atlantic colonies. The inland bill of exchange, which, just like the foreign, could be used both as a means of payment as well as a credit instrument, became the dominant form for the medium of exchange in Britain. Neal [15] provides a vivid account of the web of credit that linked institutions, entrepreneurs and savers, as well as the financial instruments used as support for the credit operations. These credit instruments were the financial contracts which the technological processes underlying the Industrial Revolution needed for adoption to take place. The financial revolution, driven primarily by the necessity to regulate the financing of British warfare, provided financial institutions with the knowledge to operate contracts required for the operation of manufactures. In the case of Britain, the set of contracts financial institutions were able to operate was exogenously enlarged by these technology-unrelated reforms; financial learning did not stem from early adoption as adoption was not taking place for lack of the appropriate financial instruments. Had there been no financial revolution, the Industrial Revolution might not have taken place. Referring to the extensive “web of credit” that formed at the time, Neal [15] argues:

‘[...] this web of credit was anchored securely in the City of London. Without this anchor, it is very doubtful whether the British economy could have made the structural changes in techniques, products and markets that characterized its transformation from

⁴Neal [15].

1760 to 1850. The financial revolution was necessary, even if not sufficient, for the industrial revolution.”

Another example of the relationship between the emergence of new financial arrangements and economic development was the establishment of trade in forward contracts, at the inception of the Chicago Board of Trade.⁵ Immediately after native Indians were forced to sell their ancestral lands, in 1833, immigrants from the East moved to the region. Many of the newcomers were New England and New York farmers who settled on the fertile soils of northern Illinois and southern Wisconsin. By 1860, “the Old Northwest was the nation’s granary made so by a mighty immigration from Europe and the eastern United States in the preceding decade.”⁶ Uncertainty about the final price of grain was an important factor in this trade. Early in the settling and trading process, as described in historical accounts of the Chicago Board of Trade, [4], the transportation of grain to the East was a process which depended on weather conditions. River merchants bought crop proceeds from farmers in the early Fall and needed to ship it to processors. Shipping required a not too cold Winter, so that the river would not freeze and shipments could sail away, and low humidity, in order not to damage the cereal. Often, as these two conditions failed them, river merchants ended up storing the grain all Winter. Early transactions in forward contracts were sought by these intermediaries. They provided insurance against the price variability between the time they purchased the grain, in the Fall, until the time of the final sale, by June of the following year. Here, the new “technology” is the long-distance grain trade. Early immigrants correspond to the early adopters, in the model, who started revealing to financial intermediaries the potential of the technology. Forward contracts were the financial instrument required for the grain trade to continue, instead of an easily forecastable collapse into local production that would have ensued had those contracts not been used. Notice that forward contracts were not formally new, there are many records of their previous usage in history. The novelty revealed by early adoption was the necessity of relying on forward contracts for the technology (the grain trade) to succeed.

2 Related Literature

This paper proposes the first model where technology affects the way in which finance evolves; it is also the first work where financial contracts and economic activity (adop-

⁵See Ferris [6].

⁶Ferris [6].

tion) evolve simultaneously and interact. The multiple aspects of the analysis make it have points of contact with a number of different literatures: growth and finance, I.O. (diffusion), common knowledge and asymmetric information between market participants. On the theory side of the relationship between growth and financial development, the closest links are Acemoglu and Zilibotti [1], Greenwood and Jovanovic [7], Bencivenga and Smith [2] and Saint-Paul [18]. Acemoglu and Zilibotti focus on conditions under which a more productive but riskier technology can be adopted and its implications for the volatility of output throughout the process of development. At early stages of development, the minimum size requirements of the risky project prevent its generalized adoption; as these projects bear idiosyncratic risk, the more projects adopted, the lower the aggregate risk for the economy. In their work, however, the set of technologies is fixed, and, conditional on the adoption decision, financial markets are complete. Similarly, in Greenwood and Jovanovic, the set of technologies is given and adoption depends on whether or not individual investors have become sufficiently wealthy to bear the fixed cost that financial intermediation entails. The current paper focuses on the rather different question of the implications of market incompleteness for technology adoption and on the way in which adoption influences contracts.

Bencivenga and Smith's work is centered around the comparative advantage of the financial system as a provider of liquidity to economic agents. In an economy with financial intermediation, individual needs to hold on to liquid but unproductive assets are reduced and the economy will grow faster as more funds are devoted to a more productive technology. Although the provision of liquidity to economic agents is one important dimension in which technology and financial innovation are related, as argued above, in this paper only the risk-sharing dimension of financial innovation is explored.

Saint-Paul has a model where productivity growth occurs through the specialization of labor. Firms determine their degree of specialization by selecting a particular technology from a given set. Higher specialization exposes the firm to greater (demand) uncertainty. Financial markets allow firms to insure against uncertainty, leading to a higher degree of specialization and, consequently, to greater productivity gains. This paper takes the opposite direction of Saint-Paul's approach. It asks the question of how the performance of financial markets will affect the adoption of exogenous technological progress.

A survey on the theoretical and empirical developments on the topic of growth and finance can be found in Levine [13]. As argued above, this paper has implications for the interpretation of the positive correlation found empirically between indicators of

financial development and growth. King and Levine [10], Levine and Zervos [12], Rajan and Zingales [16], are some examples of the prolific body of literature on the topic.

Concerning the lags in adoption of new technologies, this paper proposes “financial learning” as a new channel for diffusion. Stoneman [20] provides an overview of the diffusion literature. While several dimensions of learning have been emphasized concerning diffusion (Jovanovic and Rousseau [9] is one such instance), the novelty here is financial learning, that is, the need for financial institutions to adjust to technological change.

The paper also relates to issues concerning equilibria in markets where traders and some trade centralizing institution (a market maker) are asymmetrically informed about the value of objects being traded. Examples of this literature are Kyle [11], and Milgrom and Stokey [14]. This paper relates to those in that the fundamental friction in the economy is the information asymmetry between market participants (entrepreneurs) and financial institutions. The focus on technology adoption leads us to pursue a simpler informational setup.

3 The Model

There are three types of agents in the economy: savers, entrepreneurs and financial intermediaries called banks. Savers own a perfectly safe technology which grants them a constant income y , $y > 0$; they are risk-neutral and maximize expected discounted utility. Entrepreneurs run a risky project and face endowment uncertainty; they are strictly risk-averse expected utility maximizers, with Bernoulli utility function $u(\cdot) = \log(\cdot)$.⁷ Entrepreneurs and savers are infinitely lived and discount the future with discount factor $\beta \in (0, 1)$. There is an identical mass of savers, banks and entrepreneurs, n , which we normalize to unity. Banks perform financial intermediation between entrepreneurs and savers. Their technology will be described below.

Entrepreneurs operate a risky technology. Technologies are characterized by a probability density function (pdf) over output values (the positive real numbers). For simplicity, and without loss of generality, they will assume a very simple structure: in the model, pdfs have a two-point support in $\mathbb{R}_+$, and a common probability $q \in (0, 1)$ for the low output realization. For example, the pair $(10, 15)$ describes one technology where the probability of 10 occurring is q , and the probability of 15 is $1 - q$. The set $\mathcal{O}$ contains

⁷The assumption of risk-neutrality for savers is made for simplicity; the results would be qualitatively similar if one considered risk-averse savers, instead, but risk-aversion would come at a substantial cost in terms of the tractability of the model.

all pairs of ordered positive real numbers; the support set of a technology is an element of $\mathcal{O}$. $\mathcal{F}$ denotes the set of all technologies that could possibly be operated (currently known and to be discovered). At a given moment in time, only a strict subset of $\mathcal{F}$ is known. $\mathcal{F}_t$ denotes the technologies that entrepreneurs know how to operate in period t . As time passes, Nature reveals new technologies to economic agents.⁸ It follows that $\mathcal{F}_{t+1} \supseteq \mathcal{F}_t$.

In this paper, we will consider the process of technology adoption as a new technology becomes available, and how such a process is affected by the nature and depth of financial intermediation. At the beginning of time, only f_1 is known: $\mathcal{F}_0 = \{f_1\}$. Technology f_1 has support over $\Theta_1 = \{\theta_1, \theta_3\}$. Later, in period $t > 0$, as a result of technological progress, f_2 becomes available: $\mathcal{F}_t = \{f_1, f_2\}$. Technology f_2 has support over $\Theta_2 = \{\theta_2, \theta_4\}$. Technologies f_1 and f_2 are independent.

The output draws $\theta \in \Theta_j$, $j = 1, 2$, represent sector-wide shocks. That is, all entrepreneurs operating a given technology f_l whose support is (θ_i, θ_j) will face a common output draw (all receive θ_i , or else all experience θ_j); there is no idiosyncratic risk in this economy.

Entrepreneurs are identical concerning their ability to run the old technology f_1 . If they choose to operate f_1 , their income will be θ_1 , with probability q , and θ_3 , with probability $1 - q$. Concerning f_2 , entrepreneurs are characterized by a skill level $s_i \in S$, $S = [1, \bar{s}]$. The number s_i is the marginal product of individual i in the risky sector. Comparing individuals i and j for whom $s_i > s_j$, if both adopt f_2 , output for individual i will be $s_i\theta_l$, for $\theta_l \in \Theta_2$, and only $s_j\theta_l$ for individual j . The distribution of skill over S is given by the pdf $g(\cdot)$. We interpret $g(s)$ as the number of entrepreneurs whose skill is s . Let $\mathcal{G}$ be the set of all pdfs with support in S . The density $g(\cdot)$ and the support set Θ_2 are jointly drawn by Nature in period t , when f_2 is made available, according to pdf $h(\cdot) : \mathcal{G} \times \mathcal{O} \rightarrow [0, 1]$. Further, $g(s) > 0$ for all $s \leq \bar{s}$.

We assume that new technologies dominate old ones in the first-order stochastic sense; this is true even for the least able entrepreneur operating f_2 . One implication of this assumption is that the expected value of output attained under f_2 exceeds that under f_1 for all entrepreneurs, irrespective of how skilled they turn out to be in the operation of f_2 .

The idea of technological progress suggests that newer technologies should allow

⁸Agents do not affect the rate of technological progress. That is, the stochastic process by which Nature reveals new technologies in period t relative to those known in period $t - 1$ is known by the agents but unaffected by their choices.

higher output levels to be attained. This will be the case provided $\theta_4 > \theta_3$. Technological progress need not be associated with first-order stochastic dominance, however, but as this assumption makes it more likely that entrepreneurs prefer f_2 relative to f_1 , it makes the results sharper.

Since savers and entrepreneurs experience different risk profiles, they could engage in mutually beneficial insurance arrangements. Insurance is provided by banks through the posting of contracts. Consider the case when only f_1 is known. Then, financial intermediaries complete the insurance markets by posting prices over the contingencies in Θ_1 and buying or selling claims over these contingencies. That is, insurance markets are complete provided Arrow-Debreu claims are traded on the output contingencies made relevant by the particular technology (or set of technologies) being operated.⁹ The objective of banks is to maximize expected profits from intermediation. We assume intermediation costs are zero and that banks operate under perfect competition: individual banks are atomistic with respect to the market for financial intermediation. There are no entry or exit costs in the market of financial intermediaries. Therefore, banks will operate if the expected profits from doing so are (weakly) positive and offer no insurance otherwise.

There is no intertemporal transfer of resources (borrowing or saving). The trade of contingent claims is therefore restricted to occur on a period-by-period basis. We will discuss later the implications of this assumption.

4 Equilibrium with One Technology

We start at time 0, when only f_1 is known. We assume all agents (entrepreneurs, savers and banks) to be fully and symmetrically informed about the features of this technology (that is, everybody knows the support of f_1 and how productive entrepreneurs are in its operation).

As mentioned above, banks provide intermediation by posting contracts. Specifically, banks buy or sell contingent claims over states of the world, which they do after posting a price vector on such contingencies. For example, at time 0, since only f_1 is known and operated, banks trade contingent claims over Θ_1 after posting the price vector p :

$$p : \Theta_1 \rightarrow \mathbb{R}_+.$$

⁹Later, we will discuss the reason for the representation of insurance contracts in a state-space format and corresponding trade in Arrow-Debreu contingencies.

The interpretation of $p(\theta)$ is the price at which banks promise to trade contingent claims on the state of the world θ . At price $p(\theta)$, they will buy or sell any amount of contingent claims that entrepreneurs or savers demand of them. After uncertainty is resolved, if θ materializes, banks will give one unit of the consumption good to an agent who bought one contingent claim on the state of the world θ , and will collect an identical amount from agents who hold short positions on the same contingency. The price vector $p(\cdot)$ is chosen to maximize the expected profits from the bank's activity. It is assumed that one-period contracts are perfectly enforceable in this economy; that is, banks can always collect the transfers corresponding to short-positions in a given state of the world and, likewise, entrepreneurs can always collect transfers promised by banks in states over which they hold long positions.

Let $c(\theta) = \{c_e(\theta), c_s(\theta)\}$ denote the consumption of entrepreneurs and savers when the state of the world is θ . Likewise, $a(\theta) \equiv \{a_e(\theta), a_s(\theta)\}$ represents the quantity of contingent claims bought by each type of agent as a function of θ . A consumption allocation $c(\theta)$ is feasible if $c_j(\theta) \geq 0$, for $j = e, s$.

Entrepreneurs and savers maximize utility taking the price vector $p(\theta)$ as given. When technology f_1 is used, individual entrepreneurs solve:

$$\max_{c_e(\theta)} \{q \log(c_e(\theta_1)) + (1 - q) \log(c_e(\theta_3))\}$$

subject to:

$$\sum_{\theta \in \Theta_1} p(\theta) c_e(\theta) \leq \sum_{\theta \in \Theta_1} p(\theta) y_e(\theta),$$

where $y_e(\theta) = \theta$. Given $\theta \in \Theta_1$, the optimal amount of contingent claims $a_e(\theta)$ is given by the difference $c_e(\theta) - y_e(\theta)$. Savers solve an identical problem but with $\log(\cdot)$ replaced by linear utility and where $y_s(\theta) = y$.

Definition 1 *An equilibrium in the economy where only f_1 is known is a feasible allocation $\{c_e(\theta), c_s(\theta)\}$ and a price vector $p(\theta)$, $\theta \in \Theta_1$, such that: given the price vector, the consumption allocation maximizes the utility of entrepreneurs and savers, and banks maximize profits; the securities' market clears.*

The linearity in the utility of savers forces bankers to set the relative price $p(\theta_3)/p(\theta_1)$ equal to the ratio of probabilities across the corresponding states: $(1 - q)/q$. Otherwise, savers would buy or sell short infinite amounts of contingent claims on one of the states $\theta \in \Theta_1$, and lead banks to very large negative expected profits.

We use the normalization $p(\theta_1) = 1$ to define:

$$p \equiv \frac{p(\theta_3)}{p(\theta_1)} = \frac{1-q}{q}.$$

Solving the problem of entrepreneurs, we find:

$$c_e(\theta_1) = q(\theta_1 + \theta_3 p), \quad c_e(\theta_3) = (1-q) \left(\frac{\theta_1}{p} + \theta_3 \right),$$

whereas the corresponding demand functions for Arrow securities are:

$$\begin{aligned} a_e(\theta_1) &= c_e(\theta_1) - \theta_1 = q\theta_3 p - (1-q)\theta_1 \\ a_e(\theta_3) &= c_e(\theta_3) - \theta_3 = (1-q) \frac{\theta_1}{p} - q\theta_3. \end{aligned}$$

Substituting in the price ratio for p , we get:

$$c_e(\theta_1) = \bar{\theta}_1 = c_e(\theta_3),$$

where $\bar{\theta}_1$ equals the expected income of entrepreneurs:

$$\bar{\theta}_1 = q\theta_1 + (1-q)\theta_3.$$

Under the equilibrium relative price p , risk-averse entrepreneurs face an actuarially fair insurance contract; consequently, they choose to fully insure against the volatility of f_1 . Entrepreneurs have constant consumption in every period, equal to the expected value of their income process. There is a unit mass of buyers of claims contingent on θ_1 and a unit mass of short-sellers of claims contingent on θ_3 . Since savers are risk-neutral, under the equilibrium price p they are indifferent between participating or staying out of the market for financial claims. Without loss of generality, we assume banks are able to find a saver for each entrepreneur they sell contracts to, ensuring that the bank's expected profits are zero.

Finally, we check the feasibility requirement by verifying that the income y of savers is enough to meet the entrepreneurs' insurance demand in any state of the world. Feasibility will only be of concern in the bad state of the world for f_1 entrepreneurs (when θ_1 occurs). We need:

$$y \geq a_e(\theta_1) \Leftrightarrow y \geq (\theta_3 - \theta_1)(1-q). \quad (1)$$

We assume equation (1) is satisfied.

5 Technological Progress

We now go to period t , the period when Nature makes f_2 known to entrepreneurs: $\mathcal{F}_t = \{f_1, f_2\}$. The fundamental friction in this economy is the asymmetry of information between entrepreneurs and the remaining agents of the economy concerning new technologies. Before proceeding, we start with a simple numerical example to illustrate how the adoption of f_2 , a technology more productive than f_1 , may require the provision of insurance for at least a subset of entrepreneurs.

Example. Let $\{\theta_1, \theta_3\} = \{1, 4\}$, $\{\theta_2, \theta_4\} = \{1.3, 4.15\}$ and $q = 1/3$. Average output under f_1 , $\bar{\theta}_1$, is $1/3 + (2/3)4 = 3$. Average output under f_2 , $\bar{\theta}_2$, is $1.3/3 + 8.3/3 = 3.2$. Given that q is the common probability of low output under both technologies, the assumed change in support clearly indicates that f_2 dominates f_1 in the first-order stochastic sense. Suppose now that entrepreneurs had to choose between adopting f_2 without insurance but could still operate f_1 under full insurance. Would they choose to switch? We consider the case of the least skilled entrepreneur, whose skill is unity. For this entrepreneur, the expected utility under f_2 and no insurance is

$$\frac{\ln(1.3)}{3} + 2\frac{\ln(4.15)}{3} = 1.0362.$$

The expected utility under f_1 and insurance, invariant to skill, is

$$\ln(3) = 1.0986 > 1.0362.$$

Consequently, despite a more productive technology being available, f_2 will not be adopted by the least skilled entrepreneurs unless some form of insurance is provided.

Through the rest of the paper, we will assume that f_2 and f_1 are such that the result from the previous example remains true. Specifically, that there is a subset of entrepreneurs with positive mass for which f_1 with full insurance is strictly preferred to f_2 without it. Given that $g(s) > 0$ for all $s \in [1, \bar{s}]$, this assumption is formalized by imposing the following condition:

Condition 2 The values of θ in Θ_1 and Θ_2 and the probability $q \in (0, 1)$ are such that:

$$\ln(\bar{\theta}_1) > q \ln(\theta_2) + (1 - q) \ln(\theta_4).$$

When f_2 is drawn by Nature, entrepreneurs learn its support, Θ_2 , the distribution of skill, $g(\cdot)$, as well as their own skill type, $s \in S$. Banks and savers are less knowledgeable

about f_2 relative to entrepreneurs: although they know the prior $h(\cdot)$ from which Θ_2 and $g(\cdot)$ are drawn, as well as the ratio $r = \theta_4/\theta_2$, they do not know the particular realizations of the skill distribution and the technology support set. Further, banks and savers cannot tell apart the skill level of different entrepreneurs. These dimensions in which entrepreneurs are better informed relative to other economic agents reflect two realistic features of the interaction between the financial system and entrepreneurs: as a new technology becomes available, the latter typically know more about its profitability. Further, information asymmetries and adverse selection concerning the talent and ability of individual entrepreneurs are well-known to affect the functioning of credit markets.

Once f_2 becomes available, the set of relevant contingencies (states of the world) becomes

$$\Theta \equiv \{(\theta_1, \theta_2), (\theta_1, \theta_4), (\theta_3, \theta_2), (\theta_3, \theta_4)\}.$$

Although entrepreneurs know Θ , at t , banks and savers only know Θ_t :

$$\Theta_t = \{\{(\theta_1, \theta_i), (\theta_1, \theta_j)\}, \{(\theta_3, \theta_i), (\theta_3, \theta_j)\}\},$$

for $(\theta_i, \theta_j) \in \mathbb{R}_+$, with $\theta_j/\theta_i = r$, and $(\theta_i, \theta_j) \neq (\theta_1, \theta_3)$. That is, banks are still able to distinguish high from low output under technology f_1 but they are unable to do so under technology f_2 .

Here, it is the appearance of a new technology that causes the market for risk-sharing arrangements to become (at least temporarily) incomplete. Even if banks remain providers of insurance for f_1 , the new technology creates a market for a new insurance contract to emerge. Insurance for f_2 is “new” in the sense that it consists of the pricing of a new risk: banks will have to post prices and trade contingencies over θ_2 and θ_4 in order to provide such insurance; however, the form of the insurance contract itself (trade of contingent claims over output contingencies) is identical to the insurance contract under f_1 given the identical nature of the two technologies. (In the real world, new technologies are likely to require also an element of contract innovation – a contractual form that differs from previously used financial arrangements.) In transition, nonetheless, contracts differ formally from the kind that is operated when the technology is known by banks.

In this paper, *financial innovation* is understood as the progressive completion of insurance markets, rendered incomplete by the emergence of new technologies. How contracts and adoption evolve over time, and how they reinforce each other, is the object of the analysis below.

Given the asymmetries of information between entrepreneurs and banks, two related questions emerge. First, what types of contracts could banks offer that were of interest

to entrepreneurs operating f_2 while simultaneously satisfying the banks' objective of attaining zero expected profits? The immediate difficulty in providing f_2 insurance is the fact that, unlike under f_1 operation, banks cannot distinguish high from low output just by observing the results of f_2 adopters. As such, they don't know when to make transfers to entrepreneurs or when to collect goods from them. The second and related question is: As time goes by, how do banks learn about the new technology?

Regarding the first question, a possible taxonomy of contracts is as follows. Contracts are simply sets of transfers between the bank and entrepreneurs to which both parties agree. Banks could either i) offer uncontingent contracts, ii) offer contracts contingent on output, or iii) offer contracts contingent on output and on a report from the agent concerning the state of nature. In case i), contracts consist simply of a transfer to be made from one party to the other, regardless of the realization of output. Clearly, such contracts will never be implemented since entrepreneurs would not agree to a contract that would specify a negative transfer and banks would not agree to one that specified a positive one. The only feasible unconditional transfer is zero. Regarding the second option, we could have banks specify a threshold such that, if output goes above that value, they collect a given amount from entrepreneurs; otherwise, they make a transfer to entrepreneurs. These transfers and thresholds would be such that, given the banks priors on $g(\cdot)$ and Θ_2 , and their impact on the skill levels for whom this contract would be accepted, banks still made zero profits in expectation. The contract would also have to deliver high enough expected utility to entrepreneurs so as to confirm their decision to adopt f_2 . In what follows, we will effectively assume that the space of suitable transfers to satisfy both these requirements is empty. Given the enormity of the requirements put on acceptable transfers, this is a reasonable assumption.

The last option, contracts contingent on the observation of output and on a report from the agent, represents the standard approach from contract theory. Given that information asymmetries develop previous to the contractual relationship, screening is the most natural environment to consider. Under screening, banks offer a menu of contracts to entrepreneurs specifying what transfers will be made for high and low output realizations. If the entrepreneur accepts to enter this relationship, he must report what the low output realization will be – which is equivalent to the selection of an insurance contract.¹⁰ Naturally, transfers will be designed in an incentive compatible way.

Contract forms under items i) and ii) above have important features in common.

¹⁰Since the report is made previous to output being observed, it is immaterial whether the report concerns the high or low output realization.

Specifically, transfers are functions of the bank's knowledge Θ_t , only, and the verification of whether or not the contract was fulfilled depends exclusively on the observation of output.¹¹ These features are also shared by first-best insurance contracts. However, screening contracts differ formally from these in that they involve one additional dimension: the agent's report over the realization of output.¹² Information constrained contracts are therefore more complex than first-best contracts. It is conceivable that the economy may lack the ability to enforce screening contracts. (Say, for example, that the agent's report – given to the bank – is not verifiable by outside parties.) In what follows, we consider two alternative scenarios, depending on whether or not, in the economy we study, such enforcement ability exists. In both cases, in order to understand how contracts evolve over time and how adoption decisions change with them, we first need to address the question posed earlier regarding the banks' learning process.

Since insurance contracts are mappings from what banks know (and possibly a report) into the real numbers, we need to specify how banks learn over time. One form of learning is through the third type of contract, the screening kind. In fact, by requesting information from the entrepreneur concerning the actual amount of the low output realization, information which incentive compatible transfers ensure is truthful, banks learn to identify the alternative output states in Θ_2 . Posing the question more sharply, we need to know how banks learn from the observation of output alone. This will allow us to understand how banks who may not be providing any form of insurance in the first period of adoption learn from the observation of f_2 adopters.

If banks observe the output generated by f_2 entrepreneurs, they learn in a Bayesian way from it. We assume that banks get to observe the entire distribution of output generated by early f_2 adopters. Once banks are able to identify output as being high or low, they will be able to do so in the future irrespective of adoption decisions. In other words, when banks first identify an entrepreneur's output as high or low, they learn to identify the *state-of-nature* that is associated with high or low; they do not depend on the future adoption decision of this particular entrepreneur in order to retain their output identification ability. So, should there be independent adoption of f_2 , or should there be adoption through any form of contract offered to entrepreneurs, banks observe the resulting distribution of output and update their priors in a Bayesian way.

The next result shows that banks able to observe the distribution of output will offer first-best insurance no later than the period after f_2 is first adopted.

¹¹Both these contracts can be represented by functions mapping Θ_t into $\mathbb{R}$.

¹²Screening contracts are mappings of the form $\Theta_t \times H \rightarrow \mathbb{R}$, where H denotes the set of reports.

Lemma 3 *If the new technology f_2 is operated for one period, banks that are able to observe the full distribution of output values generated by f_2 adopters will provide f_2 first-best insurance contracts starting in the following period.*

Proof. *If f_2 is adopted and banks observe the full distribution of output values, they will be able to observe the output generated by the least able entrepreneur. For this entrepreneur, it must be the case that the expected utility from adopting exactly equals the alternative from fully insured f_1 ; the expected utility from f_2 adoption will depend on several things, namely on the skill level of the entrepreneur and on any insurance arrangements provided by banks in period one. In all cases, however, the adoption rule for the least able entrepreneur will take the general form:*

$$d(s\theta_2, r, q, \beta) = \ln(\bar{\theta}_1). \quad (2)$$

In words, the expected utility from adoption will be a function $d(\cdot)$ entailing a computation of expected utility; since banks know the utility function of entrepreneurs, they also know the functional form of the function $d(\cdot)$ as well as the parameters it depends on (the probability q , the discount factor β , the ratio r between high and low output). Further, being a computation of expected utility, $d(\cdot)$ is continuous and strictly monotonic in its first argument. Since f_1 is known, banks can compute the right-hand side of (2) independently. From the observation of output, they see either $s\theta_2$ or $s\theta_4$. From the strict monotonicity of $d(\cdot)$, given the banks' knowledge of r , q and β , equation (2) will only be satisfied at equality for one of the two possible values $s\theta_2$ or $s\theta_4$. Banks are therefore able to identify whether high or low output took place and this is all they need to know in order to offer first-best insurance contracts. ■

5.1 Adoption Without Screening Contract Enforcement

We now assume that the economy does not have the ability to enforce screening contracts – where transfers would be contingent on both observed output and the entrepreneur's report.

Before defining and characterizing the adoption equilibria for the economy, we first define the following objects. Bank intermediation in period $t + j$ is summarized by the price vector $p_{t+j} : \Theta_{t+j} \rightarrow \mathbb{R}_+$. This notation emphasizes contracts as reflecting the banks knowledge, given by Θ_{t+j} .

Let

$$c_{e,s,t+j}^l : \Theta_t \times \Theta \rightarrow \mathbb{R}_+$$

denote the consumption in period $t + j$ of an entrepreneur with skill s , who chooses to operate technology l in that period. Consumption depends on the state of the world perceived by banks as well as on the true output contingencies. Likewise, we define

$$c_{s,t+j} : \Theta_t \rightarrow \mathbb{R}_+$$

as the consumption in period $t + j$ of a saver. Savers' consumption depends on the knowledge of banks, only. Feasibility in consumption requires that consumption of savers and entrepreneurs be positive in all time periods and states of the world. Further, we denote by

$$\nu_{s,t+j} : \Theta_{t+j} \rightarrow \{1, 2\}$$

the technology choice of an entrepreneur whose skill is s . If $\nu_{s,t+j} = 1$, this entrepreneur chose to operate technology f_1 in period $t + j$.

We define an equilibrium in this economy as follows.

Definition 4 *An equilibrium in the economy without screening contract enforcement is a feasible consumption allocation $\{c_{e,s,t+j}^1, c_{e,s,t+j}^2, c_{s,t+j}\}_{j=0}^\infty$, a sequence of adoption decisions $\{\nu_{s,t+j}\}_{j=0}^\infty$, and a sequence of prices $\{p_{t+j}\}_{j=0}^\infty$ such that: given the prices, the consumption allocation and adoption decisions maximize the utility of entrepreneurs; given the prices, the consumption allocation maximizes the utility of savers and banks maximize profits; the securities' market clears.*

Definition 5 *A steady-state in the economy without screening contract enforcement is an equilibrium where $c_{e,s,t+j}^1 = c_{e,s}^1$, $c_{e,s,t+j}^2 = c_{e,s}^2$, $c_{s,t+j} = c_s$, $p_{t+j} = p$, $\forall j \geq 0$, and where the adoption decisions of individual entrepreneurs remain constant over time.*

We define the following thresholds. Let s_i denote the skill level of an entrepreneur who is indifferent between the per-period expected payoff from f_2 uninsured and the per-period payoff under f_1 with full insurance:

$$q \ln(s_i \theta_2) + (1 - q) \ln(s_i \theta_4) = \ln(\bar{\theta}_1).$$

Entrepreneurs whose skill is at least s_i are willing to adopt f_2 regardless of the provision of insurance.

We also define $\tilde{s}_i$, with $\tilde{s}_i < s_i$, as the skill level of an entrepreneur for whom the expected utility from adoption under the knowledge that first-best insurance will be provided the period after (and onwards) exactly matches the expected utility from operating f_1 for ever:

$$q \ln(\tilde{s}_i \theta_2) + (1 - q) \ln(\tilde{s}_i \theta_4) + \frac{\beta}{1 - \beta} \ln(\bar{\theta}_2) = \frac{\ln(\bar{\theta}_1)}{1 - \beta}.$$

This would be the adoption condition if there were to be a single adopter. For skill type $\tilde{s}_i$, it is indifferent to adopt f_2 in period 1 (and forever) or sticking to f_1 , under the knowledge that today's adoption will trigger the provision of first-best insurance from tomorrow onwards.

There may be three different equilibria in this economy, depending on how $\tilde{s}_i$ and s_i compare with $\bar{s}$. We outline each of the possible cases below.

Case 1: $s_i \in (1, \bar{s}]$. The dynamics are as follows. Entrepreneurs whose skill is at least s_i adopt f_2 at t . No other entrepreneur wishes to adopt as entrepreneurs whose skill is below s_i experience a per-period utility cost associated with the lower expected utility they get under f_2 uninsured relative to f_1 with full insurance.

The mass of early adopters, n_0 , is given by:

$$n_0 = \int_{s_i}^{\bar{s}} g(s) ds.$$

Banks observe the output produced by these entrepreneurs. Given Lemma 3, they offer first-best insurance contracts at the outset of period $t + 1$, leading all entrepreneurs to switch to f_2 from that period onwards. The steady-state mass of adopters, n_t , is identical to the mass of entrepreneurs (unity).

Consumption of f_1 adopters during the learning period is always identical to $\bar{\theta}_1$. Those who adopt f_2 before intermediation is provided simply consume the output they generate. The steady-state is reached in period $t + 1$, where all entrepreneurs adopt f_2 starting then and forever, and consume $\bar{\theta}_2$ in every period. Feasibility is assumed throughout.

Case 2: $s_i > \bar{s} \geq \tilde{s}_i$. Early adoption will be carried out by one entrepreneur only, whose skill is at least $\tilde{s}_i$. Given that one entrepreneur is adopting the new technology, and since early adoption entails a per-period utility cost previous to intermediation being offered, all other entrepreneurs prefer to stick to f_1 until f_2 is learned, as the adoption of one entrepreneur is all that is required for banks to learn f_2 . In period $t + 1$, all entrepreneurs adopt f_2 . The adoption dynamics are qualitatively similar to the previous case, although now early adoption is restricted to the bare minimum. There are multiple equilibria as far as the identity of the early adopter is concerned. However, all equilibria have the same outcome concerning adoption decisions (a single entrepreneur adopts f_2 in period t).

Case 3: $\tilde{s}_i > \bar{s}$. In this case, there will be no adoption: the skill level that would make an entrepreneur indifferent between adopting f_2 or sticking to f_1 , $\tilde{s}_i$, is too steep relative to the economy's endowment, $\bar{s}$.

Discussion. Screening contracts are more complex than the first-best type in that they require the ability to verify the agent's report, in addition to output. This added complexity provides a justification for why the economy may lack the ability to enforce such arrangements: institutions such as courts may be well prepared for the enforcement of steady-state first-best insurance contracts, where no report is involved, but may lack the knowledge or even the legal authority to enforce the second-best kind. We interpret the inability to enforce screening contracts as a measure of the quality (or lack thereof) of a country's legal institutions.

The results in this section illustrate the complementarity between technology adoption and financial innovation, the latter understood as the progressive completion of insurance markets. Despite the advantages of the new technology (recall that f_2 dominates f_1 in the first-order stochastic sense), entrepreneurs whose skill is too low will not adopt f_2 prior to financial intermediation being offered. Only once learning occurs and insurance is offered will the bulk of adoption take place. In case 1, there exist sufficiently high skilled entrepreneurs such that, even prior to being offered insurance, their skill raises the value of output by a sufficient amount so as to compensate for the increased volatility of their consumption profile. This group of very able individuals forms the mass of early f_2 adopters. In case 2, however, no such individuals are around. In this case, early adoption is undertaken by an entrepreneur who experiences an expected utility loss in the first period of f_2 operation, since he has not yet been granted insurance; however, the skill level of an early adopter is high enough for the probability of future insurance to prompt adoption. In this scenario, it is the prospect of future insurance that justifies early adoption. Finally, in case 3, despite a better technology being available, no entrepreneur is sufficiently skilled to willingly bear the cost of early adoption. As a consequence, no learning takes place and insurance will never be offered. In turn, this causes adoption not to take place at all.

Concerning cases 1 and 2, it is worthwhile to notice the feedback from adoption to finance and then back to adoption. Early adoption enables learning, making it possible for financial institutions to offer insurance; in turn, insurance boosts adoption further as entrepreneurs can take full advantage of the new technology.

We next consider the case when legal institutions are able to enforce screening con-

tracts.

5.2 Adoption with Screening Contracts

We solve for a menu of transfers to be given out when $s\theta_2$ and $s\theta_2r$ are observed. Before operating the technology, the entrepreneur reports the value of $s\theta_2$. As will be confirmed below, upon observation of $s\theta_2$, the bank transfers resources to the agent. When high output is observed instead, the entrepreneur makes a net transfer to the bank.

Since net transfers are received when low output is realized, there is a potential incentive for the entrepreneur to lie and state that, in fact, $s\theta_4$ is the *low* output value. If he does so, with probability $1 - q$ observed output will be $s\theta_4$ and he will collect a transfer from the bank without having to make any payment. Since output is observed by the bank, the contract must take into account the possibility that a lie is detected. We assume that, if a lie is detected (which would happen with probability q if $s\theta_4$ is reported as the low output), the bank has the ability to collect a fraction τ of the entrepreneur's resources. Making τ sufficiently large eliminates the agency problem ($\tau = 1$ is the obvious example). Here, we assume $\tau = 0$: if a lie is detected, the bank does not make nor receive transfers from the agent. The "penalty" suffered by an agent who lies and is spotted will come not in the form of a fine on his second period income but from the absence of an otherwise positive transfer he would have received had a lie not been detected.

The contract must also specify what transfers will be given to the entrepreneur should he lie and not be detected. The possibility of a lie going undetected occurs when the entrepreneur reports $s\theta_4$ as the low income and, in fact, $s\theta_4$ is the income that turns out to be observed once the technology is operated. In this situation, the bank will assume it is facing an entrepreneur whose output possibilities are $s\theta_4$ and $sr\theta_4$ (instead of $s\theta_2$ and $sr\theta_2$). As such, upon seeing the realization $s\theta_4$, the bank will transfer to the entrepreneur the same amount that would be received by an entrepreneur whose low income were $s\theta_4$. In order to handle this "nesting" of contracts (transfers to a type appear in the incentive compatibility constraint of a lower type), we guess and verify that the optimal transfers are linear functions of output. For example, for an entrepreneur whose low income is $s\theta_2$, the transfer received should the low income be observed will be $ls\theta_2$, with $l > 0$. If high income is observed, he will receive $hsr\theta_2$, with $h < 0$. The slopes l and h are constant across types. We next formulate the optimal contract problem.

Perfect competition across banks forces them to maximize the agent's expected utility

from the contract, subject to the transfers yielding the bank zero profits in expectation. The optimal slopes l and h are the solution to:

$$\max_{l,h} \{q \ln(s\theta_2(1+l)) + (1-q) \ln(s\theta_4(1+h))\}$$

subject to the incentive compatibility constraint,

$$q \ln(s\theta_2(1+l)) + (1-q) \ln(s\theta_4(1+h)) \geq q \ln(s\theta_2) + (1-q) \ln(s\theta_4(1+l)),$$

and the break-even condition for the bank:

$$s(ql\theta_2 + (1-q)h\theta_4) \leq 0.$$

It is immediate to notice that s drops out of the objective function as the term $\ln(s)$ can be isolated from the rest. Further, it also cancels out on the two constraints. The problem simplifies to:

$$\max_{l,h} \{q \ln(1+l) + (1-q) \ln(1+h)\} \quad (3)$$

s.to:

$$-(1-2q) \ln(1+l) + (1-q) \ln(1+h) \geq 0 \quad (4)$$

and

$$ql\theta_2 + (1-q)h\theta_4 = 0. \quad (5)$$

This confirms the guess that the slopes l and h will not depend on the value of output, $s\theta_2$ or $s\theta_4$, and are therefore identical for all entrepreneurs for the same state of the world. The fact that the term in $\ln(s)$ drops out of the optimization problem also makes apparent that consumption is proportional to skill: an entrepreneur whose skill is s will enjoy higher consumption than another with skill s' , $s' < s$, by a factor s/s' . As a consequence, total utility for an entrepreneur whose skill is s equals $\ln(s) + u^*$, where u^* is the utility attained by an entrepreneur whose skill is unity.

Equations (3) through (5) describe a concave programming problem when $q \geq 1/2$ (concave objective function and convex constraints). Therefore, first-order conditions are necessary and sufficient for a global maximum for this range of values of q .

Let γ and μ be the multipliers of the incentive-compatibility and break-even con-

straints, respectively. First-order conditions are:

$$\frac{1}{\theta_2(1+l)} - \frac{\gamma}{\theta_2} \left(\frac{1-2q}{q} \frac{1}{(1+l)} \right) = \mu \quad (6)$$

$$\frac{1+\gamma}{\theta_4(1+h)} = \mu \quad (7)$$

$$l = -h \frac{\theta_4(1-q)}{\theta_2 q} \quad (8)$$

$$-(1-2q) \ln(1+l) + (1-q) \ln(1+h) \geq 0. \quad (9)$$

The properties of the optimal screening contract are as follows:

Theorem 6 *Under the optimal screening contract, $l \geq 0$ and $h \leq 0$. There are parameter configurations for which the incentive compatibility constraint (IC) binds. When the IC constraint binds, h takes values in $[-1, 0]$. If this constraint binds for $q \leq 1/2$, the optimal contract is given by $l = h = 0$.*

Theorem 6, proved in the Appendix, confirms the earlier guess that entrepreneurs receive transfers when the low output state occurs and that banks collect transfers otherwise.

Adoption with Second-Best Insurance. Equilibria in this economy will depend crucially on the magnitude of u^* . If the incentive compatibility constraint did not bind, first-best insurance would be provided and adoption would be immediate and full: all entrepreneurs would switch to f_2 the moment it became available. The more stringent the incentive compatibility constraint, the lower the utility entrepreneurs will get from the second-best contract.

Definition 7 *An equilibrium in the economy with multi-period contract enforcement is a feasible consumption allocation $\{c_{e,s,t+j}^1, c_{e,s,t+j}^2, c_{s,t+j}\}_{j=0}^\infty$, a sequence of adoption decisions $\{\nu_{s,t+j}\}_{j=0}^\infty$, a second-best insurance contract (l, h) , and a sequence of prices $\{p_{t+j}\}_{j=1}^\infty$ such that: given the second-best contract and the prices, the consumption allocation and adoption decisions maximize the utility of entrepreneurs; given the second-best contract and the prices, the consumption allocation maximizes the utility of savers and banks maximize profits; the securities' market clears.*

Definition 8 *A steady-state in the economy with multi-period contract enforcement is an equilibrium where: $c_{e,s,t+j}^1 = c_{e,s}^1$, $c_{e,s,t+j}^2 = c_{e,s}^2$, $c_{s,t+j} = c_s$, $p_{t+j} = p$, $\forall j \geq 0$, and where the adoption decisions of individual entrepreneurs remain constant over time.*

When parameter values are such that the optimal contract is $h = l = 0$, the equivalent of no insurance, adoption dynamics are the same as those of section 5.1. We next focus on the cases when nontrivial second-best insurance is offered, so that both h and l are strictly different from zero.

For an entrepreneur whose skill is $s > s_i$, the decision to accept second-best insurance depends on whether the expected utility attained under the contract exceeds that from independent adoption:

$$\begin{aligned} \ln(s) + u^* &\geq q \ln(s\theta_2) + (1 - q) \ln(s\theta_4) \\ \iff u^* &\geq q \ln(\theta_2) + (1 - q) \ln(\theta_4). \end{aligned} \tag{10}$$

This inequality does not depend on s , indicating that, if it holds for s_i , it will also hold for all skill levels $s \in [s_i, \bar{s}]$. So, provided (10) holds, entrepreneurs who would adopt even if there was no insurance will now adopt and attain higher expected utility. This, however, is the group for which second-best insurance is less relevant as they would adopt regardless of its existence. We turn to lower skill levels, for whom insurance is necessary for adoption to occur.

The adoption condition for those whose skill is below s_i is¹³:

$$\ln(s) + u^* \geq \ln(\bar{\theta}_1).$$

The left-hand side is increasing in skill, whereas the right-hand side is not. Consequently, there will be a cutoff point $\underline{s}_i$ such that, for skill levels below $\underline{s}_i$, entrepreneurs stick to f_1 , while those above it adopt f_2 and resort to second-best insurance. The interesting case (as compared to section 5.1), is when $\underline{s}_i < \bar{s} < \tilde{s}_i$. In words, there would be no adoption whatsoever without second-best insurance, since $\tilde{s}_i$ exceeds the highest skill level in the economy. But since the expected utility under second-best insurance is high enough relative to the old alternative of f_1 , there will be adoption of f_2 . The ability to enforce screening contracts makes a crucial difference in this case: adoption is only possible because of the early provision of insurance.

¹³This condition assumes there will be more than one adopter. Should there be a single adopter, the condition would be:

$$\ln(s) + u^* + \beta \frac{\ln(s\bar{\theta}_2)}{1 - \beta} \geq \frac{\ln(\bar{\theta}_1)}{1 - \beta}.$$

If there is a single adopter, this entrepreneur is the one revealing information to the financial sector; he cannot rely on other entrepreneurs conveying that information. As such, he weighs the benefit from adopting f_2 and getting full insurance from the second period onwards against getting the utility from f_1 forever.

Whenever $\underline{s}_i < \bar{s}$, so that second-best insurance finds positive demand by entrepreneurs, the economy goes through a sequence of contracting stages. In the first period, entrepreneurs resort to second-best screening contracts. In the second period, as banks learn to distinguish high from low output values, insurance markets are completed and first-best insurance is offered; screening contracts no longer find positive demand.

Discussion. The relevance of second-best insurance is crucial for adoption to take place when $\tilde{s}_i > \bar{s}$. In this case, as seen in the previous section, without financial intermediation no entrepreneur is sufficiently skilled to adopt the new technology on his own, not even in the knowledge that he will be offered insurance once learning is complete. As a consequence, no adoption takes place whatsoever. When there is the ability to enforce multi-period contracts, on the other hand, second-best insurance will induce early adoption and, as a consequence, learning, which will further reinforce adoption.

The starting point of the model is the emergence of new technologies and, as an immediate consequence, the fact that financial markets become temporarily incomplete. From the analysis, we saw that financial innovation (the progressive completion of insurance markets) is necessary for adoption, at least for it to take full scale; it is never sufficient. Stated differently, if insurance markets were complete at a moment in time and banks offered greater variety of financial contracts, they would not modify adoption decisions. This paper therefore identifies a causality channel between technology and finance, where causality goes from the former to the latter. It is the emergence of new technologies that creates the demand for new contracts here.

The model predicts also that, provided institutions are capable of enforcing more complex contracts than first-best contracts, exclusively used in the long-run, when the former meet with positive demand, the economy operates different generations of contracts. Initially, only second-best insurance is offered, later first-best insurance takes over the market for risk-sharing arrangements. In reality, different contracts also seem to be provided by different institutions. For example, at the outset of the internet bubble of the 1990's, venture capitalists and other financial intermediaries provided financing to startup companies in the Bay area. As time went by, some of these companies resorted to more conventional forms of financing (such as going public). One could interpret venture capitalists as institutions that have a comparative advantage in learning about new technologies. Going back to the model, we could think of ranking different financial institutions in terms of their knowledge of the new technology and on how this knowledge

affects the stringency of the incentive compatibility constraint. The more knowledgeable an institution, the closer the contract will be to the first-best and the greater the mass of early adopters.

This paper also has implications for cross-country differences in technology adoption, closely tied to the learning merits of a country's institutions. The ability to enforce more complex contracts may be decisive for technology adoption; in the very least, it enhances early adoption. Further, if a spectrum of different institutions is required for the learning phase (venture capitalists, etc.), countries with institutions that are better at learning will adopt a broader set of new technologies than others where such institutions do not operate. Although the relevance of institutions for growth and development is not a new topic, the learning channel for financial contracts had not been explored to date.

The model shows that the intensity of adoption and the use of financial contracts evolve together, predicting that empirical measures of financial development and growth should display positive correlation. In addition, to the extent that such measures of financial development are correlated with the learning ability of the financial sector, the results here are consistent with the finding that financial development today has a positive impact on the growth rate of output (or other indicators of economic development) in the future. The results in [10] and [12], among many other examples in the Macro empirical literature, document both predictions. Although the direction of causality in this paper goes from technology to finance, the joint dynamics of technology adoption and financial development are consistent with results usually interpreted as evidence of finance to growth causality.

Generalizations. The model assumed technologies have a two-point support. Would things change if the support had more mass points? Suppose there were three possible values of θ in Θ_2 . The main difference would be the fact that the learning time would be stochastic: banks would need to observe two different output realizations before they could learn to distinguish high from medium or low output. The learning time would be of at least two periods, and the exact moment of learning would be a random variable. It would correspond to the first period when two different output values were first observed. Of course banks would be updating their priors concerning the likelihood that observed output values corresponded to each of the three possibilities. As a consequence of learning, it could be the case that some contracts could be offered before learning were fully complete, of the uncontingent kind. In turn, this would reinforce earlier adoption. Qualitatively, the results are similar. If the new technology had a continuous support,

then banks would never learn the full support but, as time and learning went on, some partial insurance contracts would start being offered.

We ruled out borrowing and saving from the start. Entrepreneurs know the likelihood that f_2 will become available at some future random date. Do they wish to borrow or save in view of future adoption? Consider the case when no second-best insurance is available. Allowing for saving would leave our results virtually unchanged. Suppose an entrepreneur knew the skill he will have at the time f_2 becomes available. If this skill level exceeds s_i , there would be no incentive to save, whatsoever. Entrepreneurs with skill higher than s_i will adopt regardless of insurance and, as such, will experience higher expected utility from f_2 than they do under f_1 . If they could, they would like to borrow against future income. Except for the case when there will be a single adopter (section 5.1, case 3), all entrepreneurs whose skill leads them to adopt once f_2 becomes available (either because their skill exceeds s_i , and they adopt without insurance, or because it exceeds $\underline{s}_i$ and they adopt under second-best insurance), adoption will never entail a utility cost. From the moment they switch to f_2 , they will experience higher utility than under f_1 and, therefore, for these entrepreneurs optimal saving would be negative. For the case when there is a single adopter, it is conceivable that this entrepreneur would want to save (assuming no borrowing for the moment) to smooth consumption at the time of transition. The incentive to save would be lowest the further ahead in time the new technology were expected to appear. If, more generally, people did not know for sure their future skill level, entrepreneurs may wish to save a little for the possibility they might be the single adopter; but the magnitudes of saving would likely be trivial. With or without saving, the qualitative dynamics would be identical to the analysis above. (Saving would also change the previously defined thresholds s_i and $\tilde{s}_i$ marginally, as higher income from borrowing would make entrepreneurs less risk averse in absolute terms.)

What if entrepreneurs could borrow? Since the new technology provides higher income to all entrepreneurs, all would like to borrow against their future income. Consider period t when f_2 is revealed. Suppose entrepreneurs could borrow any amount b and repay $b(1+r)$, in the following period, with the interest rate equal to the rate of time preference $(1/\beta - 1)$. All would want to do so to take advantage of the higher future income. Adoption thresholds s_i and $\tilde{s}_i$ would change marginally again, but the qualitative dynamics would not be altered: those with high skills would adopt early on, if any; the others would wait another period and adopt then. In addition, if the arrival rate of new technologies is very high, any relevance of borrowing or saving decisions would be

further reduced.¹⁴

In the model, we confined contracts to agreements between a bank and a single entrepreneur; more importantly, the agreed upon transfers were independent of the performance or reports of other entrepreneurs. Since output shocks are common to all entrepreneurs, there is the potential for significant improvement in risk-sharing by allowing transfers to depend on the statements of multiple agents. Consideration of contracts that depend on a single individual's report was preferred for several reasons. First, transfers dependent on several reports entail a dimension of contract complexity that appears extremely unrealistic at the early stages of technological discovery. Second, while the assumption of sectorial shocks and no idiosyncratic risk served the purpose of simplicity, it is not a fully realistic assumption either. More realistically, the output generated by an entrepreneur would likely be affected by an idiosyncratic component. This being the case, while there may remain risk-sharing improvements from allowing contracts to depend on the reports and/or output of different entrepreneurs, they are not as obvious as in the current setting.

Some of the results relied on very strong information and rationality assumptions. For example, banks are assumed able to observe the entire distribution of output produced by entrepreneurs. Since banks are atomistic, this is perhaps too strong an assumption. We could instead consider the case when banks would only observe a fraction of the output produced by adopters, a fraction increasing in the mass of adopters. Banks would use the observed adoption outcomes to update their priors on the technology and offer progressively better contracts. Here, the *intensity* of early adoption would also contribute to learning, whereas in the model, *time* is the only learning factor (provided there is early adoption, one period is all that is required for banks to learn, irrespective of the intensity of early adoption). The completion of insurance markets would likely be smoother: progressively better contracts would be offered as a function of past adoption intensity. Further, another more realistic implication would emerge: different generations of contracts (first- and second-best) would likely overlap, as banks would have asymmetric knowledge of the technology as a function of the specific subset of adopters they got to observe.

Going back to an interpretation matter stressed in the introduction, although this

¹⁴If second-best contracts are offered simultaneously with saving, provided saving is observable by banks, our results remain qualitatively unchanged. If saving is not observable by banks, a number of nontrivial issues arise. We conjecture that, because of the possibility that lying may be detected with positive probability, a modified incentive compatibility constraint will hold and incentives for thurthtelling may still be implemented.

paper has dealt with the insurance aspects of technological progress, the relationship between financial innovation and technology adoption must be understood as broadly applying to all the features defining the implementation of new technologies. The message of this paper may be summarized by saying that “technology equals contracts.” The paper focussed on technologies that involved risk, only, but the message should not be limited by this. The job of financial institutions is that of finding contracts that accommodate all aspects of a new technology, in order to allow for its adoption. Examples of other dimensions that contracts should handle are private information (managers and shareholders have potentially conflicting objectives) and indivisibilities (investors and entrepreneurs have different liquidity needs over time that need to be matched).

6 Conclusion

This paper proposes a novel link between financial innovation and growth through technology adoption. The properties of technology (risk-profile, indivisibilities, private information) determine an optimal set of contracts that allow economic agents to share the surplus associated with technology adoption. The friction in this environment is the asymmetry of information regarding new technologies between entrepreneurs and financial intermediaries. The financial sector has an important role in making the adoption of new technologies possible by progressively completing financial markets, rendered incomplete by the emergence of new technologies. In order to do so, however, financial institutions need to learn about the new technologies. Only the link between risk-sharing and technology adoption has been explored here, although the implications of the analysis generalize in a straightforward way to other dimensions.

The current work is to be interpreted as a first step in what seems to be an area of research with very broad implications. Perhaps more important than the growth to finance causality link that the model supports is the notion that changes in financial contracts and markets are driven by technological advancements, and that the intensity of technology adoption and the types of financial contracts that support such adoption evolve jointly and interact with each other. These issues are the object of ongoing research.

References

- [1] Acemoglu, Daron; and Fabrizio Zilibotti “Was Prometheus Unbound By Chance? Risk, Diversification and Growth.” *Journal of Political Economy* 105(4), 1997, pp.709-751.
- [2] Bencivenga, Valerie R.; and Bruce D. Smith “Financial Intermediation and Endogenous Growth.” *Review of Economic Studies* 58, 1991, pp. 195-209.
- [3] Baskin, Jonathan B.; and Paul J. Miranti, Jr. *A History of Corporate Finance*. Cambridge University Press, 1999.
- [4] *Action in the Marketplace*, Board of Trade of the City of Chicago, Chicago, IL, 2000.
- [5] Cronon, William *Nature’s Metropolis: Chicago and the Great West*, (New York: W. W. Norton, 1991).
- [6] Ferris, William G. *The Grain Traders: The Story of the Chicago Board of Trade*. East Lansing: Michigan State University Press.
- [7] Greenwood, Jeremy; and Boyan Jovanovic “Financial Development, Growth and the Distribution of Income.” *Journal of Political Economy* 98(5), 1990, pp.1076-1107.
- [8] Hicks, John *A Theory of Economic History*. Oxford: Clarendon Press, 1969.
- [9] Jovanovic, Boyan, and Peter L. Rousseau “Why Wait? A Century of Life Before IPO.” *AEA Papers & Proc.*, May 2001.
- [10] King, Robert G. and Ross Levine “Finance and Growth: Schumpeter Might be Right.” *Quarterly Journal of Economics* 108(3), August 1993, pp. 717-37.
- [11] Kyle, Albert S. “Continuous Auctions and Insider Trading.” *Econometrica* 56(3), 1985, pp. 1315-36.
- [12] Levine, Ross and Sara Zervos “Stock Markets, Banks, and Economic Growth.” *American Economic Review* 88(3), June 1998, pp. 537-558.
- [13] Levine, Ross “Financial Development and Economic Growth: Views and Agenda.” *Journal of Economic Literature* 35, June 1997, pp. 688-726.

- [14] Milgrom, Paul, and Nancy L. Stokey “Information, Trade and Common Knowledge.” *Journal of Economic Theory*, 1982.
- [15] Neal, Larry “The Finance of Business During the Industrial Revolution,” in *The Economic History of Great Britain: 1700-1860*, Vol. 1, 2nd Ed., R. Floud and D.N. McCloskey, eds. (Cambridge: Cambridge University Press, 1993).
- [16] Rajan, Raghuram G., and Luigi Zingales “Financial Dependence and Growth.” *American Economic Review* 88, 1998, pp. 559-586.
- [17] Robinson, Joan “The Generalization of the General Theory,” in *The Rate of Interest, and Other Essays*. London: Macmillan, 1952, pp. 67-142.
- [18] Saint-Paul, Giles “Technological Choice, Financial Markets and Economic Development.” *European Economic Review* 36, 1992, pp. 763-781.
- [19] Schumpeter, Joseph A. *Theorie der Wirtschaftlichen Entwicklung* [The Theory of Economic Development.] Leipzig: Dunker & Humblot, 1912; translated by Redvers Opie. Cambridge, MA: Harvard U. Press, 1934.
- [20] Stoneman, Paul “Technological Diffusion.” in *Handbook of the Economics of Innovation and Technological Change*, Paul Stoneman ed. (Cambridge: Blackwell Publishers, 1995).

A Proof of Theorem 6

We prove the statements in the order they are listed in the theorem.

Suppose that the optimal contract is given by (l_0, h_0) , with $l < 0$ and $h > 0$. We proceed to show that $(l, h) = (0, 0)$ outperforms (l_0, h_0) and is incentive-compatible, contradicting the optimality of (l_0, h_0) . We resort to a graphical proof. Plot the log function (or any concave function) against output. Identify the absciss points θ_2 and θ_4 , as well as the expected value of output under $f_2, \bar{\theta}_2$. Draw the line that connects $\log(\theta_2)$ to $\log(\theta_4)$; call this segment AB . The expected utility attained under $(l, h) = (0, 0)$ is identified by the value on the y -axis of the point corresponding to the intersection of AB and a vertical line $x = \bar{\theta}_2$. To identify the expected utility attained under (l_0, h_0) , proceed as follows. Identify the values $\theta_2(1+l)$ and $\theta_4(1+h)$ on the x -axis. Since, by assumption, $l < 0$, $\theta_2(1+l)$ is smaller than θ_2 ; conversely, since $h > 0$, $\theta_4(1+h)$ exceeds θ_4 . Draw the line that connects $\log(\theta_2(1+l))$ to $\log(\theta_4(1+h))$, call it CD . Given that $l < 0$, $h > 0$, CD is everywhere below AB . Since (l_0, h_0) satisfies the break-even condition, the expected value of output under this insurance contract coincides with the expected value of output under $f_2, \bar{\theta}_2$. Therefore, the expected utility attained under (l_0, h_0) is given by the value on the y -axis of the point corresponding to the intersection of CD and the vertical line $x = \bar{\theta}_2$. Since CD is everywhere below AB , it is also the case that the point at which AB and $x = \bar{\theta}_2$ intersect is strictly above the point at which CD and $x = \bar{\theta}_2$ intersect. Further, notice that $(0, 0)$ satisfies both the incentive compatibility condition, (9), and the break-even condition for banks, (8). This shows that $(l, h) = (0, 0)$ is feasible and gives strictly higher expected utility to entrepreneurs, contradicting the optimality of (l_0, h_0) .

Suppose the incentive compatibility constraint does not bind. Then, consumption across both states would be identical and equal to $\bar{\theta}_2$. When θ_2 is observed, the transfer made to the entrepreneur would be such that:

$$\begin{aligned} q\theta_2 + (1-q)\theta_4 &= \theta_2(1+l) \iff \\ (1-q)(r-1) &= l. \end{aligned} \tag{11}$$

If the incentive compatibility binds, when given full insurance, the entrepreneur will prefer to lie, violating the constraint. For this constraint to bind, it must be the case that:

$$q \ln(\theta_2) + (1-q) \ln(r\theta_2(1+l)) > \ln(q\theta_2 + (1-q)\theta_4).$$

We next verify that, with $\theta_2 = 1$, $r = 4$ and $q = 0.6$, the previous inequality is satisfied, indicating a violation of the incentive compability constraint.

Using the expression for l from (11), the left-hand side totals

$$\begin{aligned} & q \ln(\theta_2) + (1 - q) \ln(r\theta_2(1 + (1 - q)(r - 1))) \\ = & q \ln(\theta_2) + (1 - q) \ln(\theta_4) + (1 - q) \ln(1 + (1 - q)(r - 1)) \\ = & 0.4 \ln(4) + 0.4 \ln(1 + 1.2) = 0.8699, \end{aligned}$$

whereas the right-hand side equals:

$$\ln(0.6 + 0.4 * 4) = 0.7885.$$

Therefore, given these parameter values, the incentive-compatibility constraint binds. (For these parameter values, the optimal contract is $(l, h) = (0.4533, -0.17)$).

When the incentive-compatibility constraint binds, combining (8) and (9) at equality, we get an equation in h , only:

$$\left(1 - hr \frac{(1 - q)}{q}\right)^{2q-1} (1 + h)^{1-q} = 1. \quad (12)$$

For the second term in brackets to take real values, it must be the case that $h + 1 \geq 0$, or $h \geq -1$. Since, under an optimal contract, $h \leq 0$, this implies that the optimal values of h must lie in the interval $[0, 1]$. When $q \leq 1/2$, the left-hand side of (12) becomes a strictly increasing function of h . The only value of h over the interval $[-1, 0]$ that satisfies (12) is $h = 0$ which, together with (8), shows that l must also equal 0. ■