

The Insurance is the Lemon: Failing to Index Contracts*

Barney Hartman-Glaser

†

UCLA

Benjamin Hébert

‡

Stanford University

December 29, 2017

Abstract

We model the widespread failure of contracts to share risk using available indices. A borrower and lender can share risk by conditioning repayments on an index. The lender has private information about the ability of this index to measure the true state the borrower would like to hedge. The lender is risk-averse, and thus requires a premium to insure the borrower. The borrower, however, might be paying something for nothing, if the index is a poor measure of the true state. We provide sufficient conditions for this effect to cause the borrower to choose a non-indexed contract instead.

*The authors would like to thank Vladimir Asriyan, Philip Bond, Francesca Carapella, Eduardo Davila, Peter DeMarzo, Emmanuel Farhi, Valentin Haddad, David Hirschleifer, Christopher Hrdlicka, Roger Myerson, Batchimeg Sambalaibat, Jesse Shapiro, Alp Simsek, Amir Sufi, David Sraer, Sebastian Di Tella, Victoria Vanasco, Jeff Zwiebel, and seminar and conference participants at Stanford University, UC Berkeley Haas, University of Washington Foster, the Federal Reserve Board, and the Finance Theory Group Meetings at University of Minnesota. All remaining errors are our own.

†Hartman-Glaser: UCLA. Email: barney.hartman-glaser@anderson.ucla.edu.

‡Hébert: Stanford University. Email: bhebert@stanford.edu.

Barney Hartman-Glaser has nothing to disclose.

Benjamin Hébert has nothing to disclose.

1 Introduction

A central implication of the literature on financial contracting is that agents should structure contracts to share risk as efficiently as possible. In many financial markets, standard contracts are simple and do not include risk-sharing arrangements that condition payments on publicly available indices. A leading example of this phenomenon is the mortgage market. In this market, homeowners are exposed to the risk that their home will decline in value. Lenders are arguably better equipped to bear this risk and could insulate homeowners against declines in house prices by making mortgage repayment terms contingent on a house-price index. These types of mortgage contracts have been widely proposed as a solution to problems facing the mortgage market, such as the subprime default crises of 2007,¹ but have failed to supplant the standard mortgage. Two common explanations for this type of market failure are that either the space of feasible contracts is incomplete (Hart and Moore [1988]) or that implementing risk-sharing contracts entails high transaction costs. Neither explanation applies when there are indices available that would allow agents to share risk efficiently and appear costless to contract upon.

In this paper, we develop a model in which the failure to condition on indices and thus efficiently share risk is an equilibrium outcome resulting from asymmetric information. In our model, an agent, whom we call the borrower, seeks financing from a set of lenders. This financial contract must be written in view of potential conflicts of interest between the lender and the borrower related to an “internal”, or idiosyncratic, state. For example, this internal state could represent the hidden ability of a mortgage borrower to make payments to the lender. At the same time, there are potential risk sharing benefits between lenders and the borrower over some imperfectly measured

¹See, for example the “Shared Responsibility Mortgage” proposed in Mian and Sufi [2015] in which interest and principal payments are contingent upon local house price indices or the “Shared-Equity Mortgage” proposed by Caplin et al. [2007] in which a borrower receives a second mortgage where payment is only due upon sale of the house and is contingent on house value.

state (e.g. local area house prices). We call this state “external” to indicate that it is unaffected by the actions of the lenders and the borrower. The external state is not directly unobservable, and to realize any risk sharing benefits, contracts must condition on some potentially imperfect measurement of the state which we call an index (e.g. a house price index). Lenders know the true joint distribution of the index and the external state, i.e. the quality of the index, while the borrower does not. In effect, the borrower faces an adverse selection problem over basis risk when lenders offer an indexed contract.

At least two equilibria can arise in the model. In the first type of equilibrium, which we call the full-information optimal contracts equilibrium, all lenders offer a contract featuring the optimal amount of insurance conditional on the true quality of the index. The full-information optimal contracts equilibrium always exists when there is competition between lenders, and features no loss in efficiency due to asymmetric information about the index. In the second type of equilibrium, which we call the non-contingent-contracts equilibrium, all lenders offer a contract that does not condition on the index. To see why such an equilibrium can obtain, consider the borrower’s response when a single lender deviates and offers a contingent contract. To at least break even on such a contract, the lender must charge the borrower an insurance premium. At the same time, the borrower will be concerned that the index is in fact uncorrelated with the risk she is aiming to insure, i.e., the basis risk for the contract is too high to justify the premium. Lenders that know the basis risk is high are happy to offer insurance and charge a high premium since the insurance is cheap for them to provide (precisely because the basis risk is high). As a result, the borrower will reject the indexed contract in favor of a standard non-contingent contract. We note that the non-contingent-contracts equilibrium exists even though the contracting space allows the use of an index, there are no transactions costs, and lenders make competing offers.

To illustrate the intuition behind these two equilibria, suppose there are just two equally likely

external states, “good” and “bad.” Now suppose a borrower receives offers of 1 dollar of financing from several competing lenders. The borrower is risk-averse with respect to the external state, meaning that her expected marginal value of a dollar is $1/2$ in the good state and $3/2$ in the bad state. The lenders are also risk averse, but less so than the borrower. Their expected marginal value of a dollar is $3/4$ in the good state and $5/4$ in the bad state. Lenders can offer contracts that are contingent upon some index but not upon the true external state directly. The index can be “high quality,” in which case it is perfectly correlated with the true underlying state, or “low quality,” in which case it is entirely independent of the true state and hence unrelated to the either the borrower’s or lender’s preferences. The borrower believes these two cases are equally likely, but the lenders observe the quality of the index before making their offers. Finally, lenders cannot offer contracts that specify positive transfers from the lender to the borrower.

Suppose that lenders make the following offers, depending on the quality of the index. If the index is high quality, they offer a contract that calls for the borrower to repay $8/3$ dollars if the realization of the index indicates the good state and nothing otherwise. If the index is low quality, they offer a contract that calls for the borrower to repay 1 dollar regardless of the realization of the index. These offers constitute what we call a full-information optimal contracts equilibrium. To see why they can obtain in equilibrium, note that all lenders are earning weakly positive profits and could not possibly earn more by making different offers. Moreover, given that all lenders have common information, the borrower can perfectly infer the quality of the index by observing the contract that the lenders offer. In other words, it is not possible for a single lender to convince the borrower the index is high quality if all the other lenders offer a non-contingent contract. This same intuition carries over to the second type of equilibrium we describe above.

Now suppose that all lenders offer a contract that calls for the borrower to repay 1 dollar regardless of whether the index is high-quality or low-quality. These offers constitute what we

call the non-contingent-contracts equilibrium. Can a single lender gain by deviating and offering the best contingent-contract? Again the answer is no. If a single lender deviates by offering a contingent contract, then she will have to charge a premium for it to at least break even. In the case of the best contingent contract, that premium is $1/6$, or the difference between the expected utility from the payment ($4/3$) and the amount financed (1). If the index is low quality, this premium is pure profit because in that case, the realization of the index is unrelated to the lender's preferences and the lender is thus risk neutral with respect to the index. As a result, the lender would be at least as willing to make such an offer given a low-quality index as given a high-quality one, and as such, standard belief refinements imply that the borrower can believe that the index is low quality after observing this deviation. Given these beliefs, the borrower is strictly better off accepting one of the offers of a non-contingent contract. The failure of the agents to share risk in this case is closely related to the classic lemons market breakdown of Akerlof [1970].

Two elements are essential to the existence of the non-contingent contracts equilibrium in the simple example above. First, the lenders know the quality of the index, and second, the lenders are risk averse with respect to expected payoffs across external states. In particular, with regards to the first element, the borrower believes there is at least some chance that there are no gains from risk sharing using the index. The second element means that deviating from the non-contingent contracts equilibrium requires that a lender charge a premium for a contingent contract, which makes such a deviation more attractive when the index is low quality. Our main result rests on two conditions which generalize this simple example, and in particular these two elements.

Our general model also encompasses settings in which there is an additional security design problem concerning payoffs given idiosyncratic states. These security design problems are important for our results in that they determine the borrower's and lenders' indirect utility over securities and external states and thus the potential gains from indexation. For example, in our mortgage

example (section §6), the borrower needs incentives to repay the lender across idiosyncratic states. In that example, conditional on a particular external state, standard debt contracts are optimal. In principle, these debt contracts could allow for risk sharing over the external states by having a higher face value in a good external state than a bad one. However, the face value of a debt contract is not equivalent to its expected payoff; put differently, promises are not payoffs. A lender can prefer a higher debt level in a good external state simply because the debt is more likely to be repaid in good external states (due to correlation between the external and idiosyncratic states). At the same time, the lender has a lower marginal utility in the good external state. The key condition to generate a non-contingent contracts equilibrium becomes a tradeoff between the lender's decreasing marginal utility and the increasing value of promises as the external state improves. If the latter force dominates, then the lender will not need to charge a premium to insure the borrower against bad external states, and the non-contingent contracts equilibrium does not exist. A key condition for the existence of the non-contingent contracts equilibrium is that the lender be sufficiently "risk averse" over promises, a notion we formalize in our general model.

There is an important distinction between the type of adverse selection problem we consider and one in which lenders have information about the external state itself. In our model, lenders do not have better information about the distribution of the external state, only about the relationship between the index and the external state. In contrast, much of the literature on adverse selection (following Akerlof [1970]) assumes there is asymmetric information about something that is directly relevant to payoffs. For example, in the context of mortgages, lenders might know that local house prices are more likely to appreciate in the future than the borrower expects. In an extension of our model, we show that under our assumptions, the non-contingent contracts equilibrium does not exist if the index is known to be perfectly correlated with the external state, and the only adverse selection is about the distribution of the external state itself.

Our work is related to the literature on incomplete contracts, surveyed by Tirole [1999]. Papers focusing on incomplete contracts and asymmetric information include Spier [1992], Allen and Gale [1992], and Aghion and Hermalin [1990], among others.² Our model differs from most of this literature in several respects. First, our model emphasizes competitive markets, rather than bilateral negotiation. Second, our model is focused on asymmetric information about the quality of the index, rather than the “fundamentals.” This second difference is what allows us to generate non-contingent contracts in equilibrium without relying on transaction costs of using the index or arguing that the index is manipulable. Like some, but not all, of the incomplete contracts literature, we focus on equilibria with no indexation at all (as opposed to explaining why agents might use the index but not achieve perfect risk-sharing).

More significantly, our model differs from the incomplete contracts literature in its assumptions about what is contractible and what is observable. In the risk-sharing extension of Hart and Moore [1988], the agents can renegotiate after observing a non-verifiable state. A subsequent literature (Green and Laffont [1992], Dewatripont and Maskin [1995], Segal and Whinston [2002]) has found that, by altering the outside options or other aspects of the renegotiation process, the agents can share risks and perhaps even achieve first-best risk sharing despite their inability to contract on the state. In contrast, in our model, the index is both observable and verifiable, whereas the true external state is not observed by the agents until the end of the game, when renegotiation is no longer possible.³

Formally, the model is similar in some respects to Allen and Gale [1992], although the focus of that paper is the manipulability of the index. One can also view our model as related to models of

²Papers that endogenize contractual incompleteness, but do not emphasize asymmetric information, include Anderlini and Felli [1994], Battigalli and Maggi [2002], Bernheim and Whinston [1998], Dewatripont and Maskin [1995], Kvaløy and Olsen [2009], Tirole [2009].

³Relatedly, the Maskin and Tirole [1999] critique of the incomplete contracts literature applies when the agents are aware of the payoff-relevant states before actions are taken, and for this and other reasons is not directly applicable to our model.

insurance, in the vein of Rothschild and Stiglitz [1976]. The key difference between our model and those models is that our model places the information advantage and the competition on the same side of the market (with lenders), rather than on opposite sides of the market. Loosely speaking, the key intuition in our model is that the insurance itself might be a “lemon,” in the sense of Akerlof [1970].

A closely related paper to ours is Spier [1992]. She shows that asymmetric information can amplify the effect of transaction costs on the ability of agents to write contracts that condition on relevant information. In her model, an informed and risk-averse principal contracts with an uninformed and risk-neutral agent. If the principal offers a contract that insulates herself from risk, she must also signal her private information, which in turn reduces the benefits of risk sharing. This effect lowers the level of transaction costs needed to destroy risk sharing in equilibrium. However, in her model, if transaction costs are close enough to zero, asymmetric information alone does not eliminate risk sharing. In contrast, in our model, asymmetric information can lead to zero risk sharing without transaction costs. Another related paper is Asriyan [2015]. He shows that concern for future liquidity and private information can lead market participants to write very simple contracts. This intuition is that if the holder of a contract must liquidate at some future date, she will want hold a contract that is as informationally insensitive as possible. In contrast, we emphasize situations in which there are risk-sharing failures associated with simple contracts. In other words, the value of simple contracts is informationally sensitive in our model, and only by using the index could the agents minimize information sensitivity.

We also employ a general space of states and contracts. As a result, there is a great deal of scope for signaling, in contrast with the previous literature (in Spier [1992] and Aghion and Hermalin [1990], the contract space has one or two dimensions). As a consequence of this ability to signal, to generate our results, borrowers must be somewhat “suspicious,” in the sense that they place non-

zero probability on the index being irrelevant. Belief in this possibility, however unlikely, creates at least some chance that the index is not useful (and in this sense is reminiscent of the conditions of the Myerson and Satterthwaite [1983] theorem).

The failure of risk-sharing in our model can be thought of as a coordination failure, in the sense that there are multiple, Pareto-ranked equilibria. In the context of mortgages, we view this multiplicity as a feature. Mortgage contracts differ substantially across countries in ways that are difficult to explain with “fundamentals.” Relatedly, our model considers only a single index, but could naturally be extended to consider multiple indices (interest rates and house prices, for example). In this case, we expect that “partial indexing” equilibria (e.g. indexing to interest rates but not home prices, like an adjustable rate mortgage) exist. As a result of this multiplicity, there is the potential for policy to improve welfare in our model by ruling out undesirable equilibria. Our model does not feature any externalities as a result of this risk-sharing failure; the existence of such externalities (which are emphasized by Campbell et al. [2011], among others) would provide an additional motivation for policy interventions.

One motivating example of our model is the mortgage market, although the model is abstract and could easily apply to other settings. In the context of home ownership, as noted by Sinai and Souleles [2005], purchasing a house hedges a homeowner against changes in future rents. Nevertheless, homeowners are exposed to both price and rent risks, and these could be hedged through the mortgage contract. Of course, as discussed by Case et al. [1995] and Shiller [2008], homeowners could also hedge these risks through other financial markets, although this almost never occurs in practice. This failure to hedge might be explained by household’s limited access to such markets, or by the sophistication required to hedge in this manner. However, these arguments suggest that it would be profitable for a financial intermediary to provide hedging services, and mortgage lenders appear to be ideally situated to do this as part of mortgage contracts. Mortgages

that have a more equity-like claim on house value have been proposed (see, for example, Caplin et al. [2007]). Some of these early proposals made mortgage payments contingent on the sale price of the house, which clearly induces moral hazard for the borrower. More recent studies have pointed out that conditioning mortgage payments on an index of house prices avoids this problem. Piskorski and Tchistyi [2017] develop an equilibrium model of housing and mortgage markets and show that under many circumstances, the optimal mortgage design hedges the borrower against house price risk. Greenwald et al. [2017] provides a quantitative analysis of the general equilibrium effects of house priced indexed mortgage and show that using a local house price index improves financial stability. Proposals for mortgage reform after the recent financial crisis (e.g. Mian and Sufi [2015]) have advocated this approach. Although rare, shared appreciation mortgages are legal in the United States and used, for example, by Stanford University faculty who borrow from Stanford to purchase a house.⁴ We develop a stylized model of mortgage borrowing, building on Hart and Moore [1998], and show that the conditions of our general theorem apply in this model and thus can explain the lack of prevalence of shared appreciation mortgages by appealing to asymmetric information over the quality of house price indices.

We begin in section §2 by describing the indirect utility functions of the borrower and lenders in our model, with examples. We describe the market for loans, the asymmetric information problem, and the equilibrium concept in section §3. In section §4, we discuss the zero-profit condition that arises from competition in our model, and characterize the “best” equilibria, which features contingent contracts. In section §5, we discuss our most general results, which describe assumptions under which risk-sharing fails and non-contingent contracts arise in equilibrium. In section §6, we provide an example security design problem that satisfies the assumptions in the preceding setting. In section §7, we describe a number of variations and extensions to our basic framework. We

⁴Stanford mortgages are indexed to an appraisal, rather than a local house price index, and involve renegotiation when the homeowner makes major investments.

conclude in section §8.

2 The Indirect Utility Functions

In this section, we begin describing our general model. The “primitives” of our model are the indirect utility functions of the borrower and lender. In this section, we will describe enough of our model to define these functions, and then provide two examples. These examples are based on standard utility functions and the costly state verification (CSV) model of Townsend [1979]. An additional example, related to mortgages, can be found in section §6.

At date zero, a borrower wishes to raise $K > 0$ dollars to pursue a project (e.g. purchasing a home). After the borrower and lender agree to a contract and initiate the project, at date 1 an index $z \in Z$ is determined. This index is observable and verifiable, and related to the true external state $a \in A$. The true external state a is what enters the agents’ indirect utility functions (they have no particular concern for the value of the index), but the index is the only thing they observe and can contract on.

The index $z \in Z$ should be thought of an index based on the external state $a \in A$. For simplicity, we will assume that both A and Z are totally ordered sets. We will write $a \succ a'$ to denote the idea that the external state $a \in A$ is “better than” the external state $a' \in A$, and use the same notation for the index values. In the context of mortgages, the external state $a \in A$ might influence house prices, borrower income, and/or the cost of capital for lenders. The index $z \in Z$ is an index that, perhaps imperfectly, measures these things, such as a local area house price index, a wage index, or an interest rate. We assume that A and Z are finite sets, but nothing relies on this.

The external state $a \in A$ influences the distribution of the borrower’s idiosyncratic outcomes, $i \in I$. For a mortgage borrower, idiosyncratic outcomes could include the borrower’s particular

house price or income. The set I can be finite or infinite. A contract is a function $s : I \times Z \rightarrow \mathbb{R}^+$ that takes the idiosyncratic outcome i and index z and determines a payment from the borrower to the lender. We use the notation $s_z : I \rightarrow \mathbb{R}^+$ to refer to the “conditional contract,” which is the contract for a particular value of the index.

The idiosyncratic outcomes may or may not be observable or contractible, and might be influenced by the borrower’s behavior. Conditional on any particular index value $z \in Z$, the set of feasible contracts is S_I , which reflects the limits on the observability or verifiability of the idiosyncratic outcomes, and any additional restrictions on the contract space (e.g. limited liability, monotonicity in idiosyncratic outcomes). We will restrict our attention to conditional contracts that are “ex-post” efficient, appealing to notions of renegotiation-proofness after the index $z \in Z$ has been revealed. We define conditional ex-post efficiency below (condition 1). This condition requires that there exists a one-dimensional family of ex-post efficient conditional contracts, indexed by a parameter $d \in D$, which spans the space of ex-post efficient contracts. We define $d = 0$ as the contract that pays nothing, and assume it is in the feasible set. We use the notation $S_D \subseteq S_I$ to refer to conditional contracts in this parametric family. The feasible set of conditionally ex-post efficient contracts is S , which can be thought of as a map from values of the index to conditionally ex-post efficient contracts.

Given a particular state $a \in A$ and conditional contract $s_z \in S_I$, the borrower’s indirect utility function is $\phi_B(s_z, a)$. We refer to this as an indirect utility function because it summarizes the borrower’s payoff, given some underlying relationship between the external state a , conditional contract s_z , and the distribution of idiosyncratic outcomes. Similarly, we denote the lender’s payoff as $\phi_L(s_z, a)$. In both cases, these functions should be understood as expected utilities conditional on a , and do not imply that the borrower or lender knows a .

We treat the indirect utility functions as primitives that satisfy several properties. First, we

assume that both of these functions are continuous in d . Formally, letting $s_d \in S_D$ denote the ex-post efficient conditional contract associated with the parameter d , $\phi_B(s_d, a)$ and $\phi_L(s_d, a)$ are continuous for $d \in D$. Second, we assume that $\phi_L(s_d, a)$ is zero for the contract that pays nothing. Third, the borrower's utility function satisfies a monotonicity property: if $d' > d$ for some $d, d' \in D$, then $\phi_B(s_{d'}, a) \leq \phi_B(s_d, a)$ for all $a \in A$.

We will refer to the parameter d as a “promise,” and in our leading examples it will correspond to the face value of a debt claim. Intuitively, if the borrower makes a larger promise to the lender, she is worse off. For the lender, this property does not necessarily hold; promises will not necessarily be paid, and demanding excessive repayment can result in lower expected utility for the lender. We use debt as our leading example, but these conditions could be describing other families of securities as well. Examples include the set of fixed payments of varying size, the set of 100% equity claims less a fixed payment of varying size, and the set of equity shares of varying shares. The first two of these examples could be motivated by risk-sharing type problems, and the third by security design problems resulting in equity as the optimal security design.

We have assumed that the family of contracts indexed by the parameter $d \in D$ span the space of ex-post efficient contracts. We will now define this formally. Consider the “social welfare function,” conditional on a particular external state $a \in A$, security $s \in S_I$, and Pareto weight $\lambda \geq 0$,

$$U(s, a; \lambda) = \phi_B(s, a) + \lambda \phi_L(s, a). \quad (1)$$

We impose the condition below on the contracts that maximize the expected value of these social welfare functions. This condition is motivated by security design problems that have been studied in the literature, and in particular the examples we will describe below. These security design problems share a common feature, which is that, holding the external state $a \in A$ fixed, the set of

Pareto-optimal securities is or includes a family of security designs indexed by a single parameter. For example, in many models of security design (e.g. Hart and Moore [1998], Innes [1990], Townsend [1979]), regardless of the parameters, the set of Pareto-optimal securities is or includes the set of debt contracts.

Condition 1. There exists a family of contracts $S_D \subseteq S_I$, indexed by the parameter $d \in D$, such that:

1. For all probability distributions $\pi \in \Delta(A)$ and Pareto weights $\lambda \in [0, \infty)$, there exists a $d \in D$ such that $s_d \in \arg \max_{s \in S_I} \sum_{a \in A} \pi(a)U(s, a; \lambda)$, and
2. For any $s' \in \arg \max_{s \in S_I} \sum_{a \in A} \pi(a)U(s, a; \lambda)$ with $s' \notin S_D$, there exists an $s_d \in S_D$ such that, for all $a \in A$, $\phi_B(s_d, a) = \phi_B(s', a)$ and $\phi_L(s_d, a) = \phi_L(s', a)$, and
3. D is a convex subset of the real line that includes $d = 0$.

The first part of this condition says that there is always a contract in S_D that is ex-post efficient. The second part refers to the “spanning” the space of ex-post efficient contracts with the elements of S_D . In costly state verification models (Townsend [1979], Gale and Hellwig [1985]), there are many contracts that result in the same payoffs (most involving false reports of the state). Debt contracts are optimal, but not uniquely so; however, debt contracts “span” the space of efficient payoffs in the sense required by this condition.

We next turn to examples of these indirect utility functions. Our first example is the case of utility functions.

Example 1. In this example, the set I is a singleton, and s_z is the payment made by the borrower to the lender if the index takes on the value $z \in Z$. There is no distinction, in this context, between the set of all conditional contracts (S_I) and the set of ex-post efficient contracts (S_D), and the parameter

d is simply equal to the payment, $s_d = d$. To simplify further, let us suppose that both agents are risk-neutral. Then

$$\phi_B(s, a) = y - \beta_B(a)s,$$

$$\phi_L(s, a) = \beta_L(a)s,$$

where $\beta_B(a) > 0$ and $\beta_L(a) > 0$ are the marginal utilities of the borrower and lender, and y is a strictly positive endowment. We will assume limited liability for the borrower, which requires that $s \in [0, y]$. This example satisfies the assumptions above: these functions are continuous in d , ϕ_L is weakly positive, and ϕ_B is monotone decreasing in security payments. We will refer to this as example as the “utility function example.”

The second example is based on the costly state verification models of Townsend [1979] and Gale and Hellwig [1985].

Example 2. In this example, each idiosyncratic state $i \in I$ is a triple (x, y, y') , where y represents the borrower’s verifiable income, x represents the borrower’s non-verifiable income⁵, and y' represents the borrower’s report of her verifiable income, all of which are weakly positive reals. Let $c(y'; s) = \bar{c} > 0$ if there exists a y_1, y_2 such that the conditional contract s offers different payments for (y_1, y') and (y_2, y') , and zero otherwise. Borrowers are risk-averse, with utility function $u(\cdot)$, and lenders are risk-neutral over idiosyncratic states, but have marginal utility $\beta_L(a)$ that depends on the external state. The indirect utility functions are

$$\phi_B(s, a) = \max_{\omega(y'|x, y) \in \Omega(s)} E \left[\int_0^\infty u(x + y - s(y, y')) \omega(y'|y, x) dy' | a \right],$$

$$\phi_L(s, a) = \beta_L(a) E \left[\int_0^\infty (s(y, y') - c(y'; s)) \omega^*(y'|y) dy' | a \right],$$

⁵We discuss our reasons for including non-verifiable income in Appendix section §A.

where $\omega(y'|y,x)$ is a (possibly mixed) reporting strategy for the borrower, $\Omega(s)$ is the set of feasible reporting strategies, and ω^* is the borrower's optimal reporting strategy. Limited liability requires that, for all reports y' , either $s(y_1,y') = s(y_2,y')$ for all y_1,y_2 and $0 \leq s(y',y') \leq y'$ (the non-verification case), or $0 \leq s(y,y') \leq y$ for all y (the verification case). We restrict the reporting strategies $\omega(y'|x,y)$ to place support only on y' for which the reports are feasible, meaning that if $s(y_1,y') = s(y_2,y')$ for all y_1,y_2 , then $\omega(y'|x,y) = 0$ if $y < s(y',y')$. In other words, for reports that do not trigger verification, the borrower must have the funds to repay the loan. Townsend [1979] and Gale and Hellwig [1985] demonstrate in an almost identical setting that debt contracts are optimal (but not uniquely optimal) in this model. The continuity in debt, weak positivity, and monotonicity in debt assumptions follow immediately. We return to a discussion of the costly state verification model in Appendix A.

We provide an additional example, related to mortgages, in section §6.

3 The Model

The dependence of the borrower and lender's marginal utilities (with respect to contract payoffs) on the external state is the force that causes them to want to condition their contract on the external state. By assumption, they cannot directly condition their contract on the external state, only on the index. In this section, we will first describe the relationship between the index and the external state. We will then describe the market for loans at date zero, and finally discuss the definition of equilibrium in the model.

3.1 Types

We define $\theta(a, z)$ as the joint distribution of the external state and the index. This joint distribution is common knowledge amongst the lenders, but is not known to the borrower; it is the “type” in our adverse selection problem. The types θ are drawn from a convex set Θ , which we define as the set of all joint distributions that have the same marginal distributions for $a \in A$ and $z \in Z$, which we denote $p(a)$ and $q(z)$, respectively. Without loss of generality, we assume these marginal distributions have full support over A and Z , respectively. Let $\theta_0(a, z) = p(a)q(z)$ denote an “uninformative type” (the type with an index that is independent of the external state).

The borrower’s prior belief over these types is μ_0 . In effect, the borrower is uncertain about the relationship between the index and the external state. A homeowner, for example, might not be certain how the S&P Case-Shiller index for his metro area is related to the price of his particular house. We assume that the borrower is aware of the marginal distributions, to abstract from the problems generated by that type of asymmetric information and focus on the borrower’s doubt about the relevance of the index (we revisit this in our extensions, section §7). We do not require that the beliefs μ_0 have full support on Θ , but will impose assumptions on the support, which we describe below.

Having defined the type space, we next describe the market for loans.

3.2 The Market for Loans

Let L denote the set of lenders, with $|L| \geq 3$, each of whom can post a contract. After these lenders post contracts, the borrower can pick whichever one she prefers, or choose to forgo the investment opportunity. The outside options for both borrower and lender are zero. Note that, from the borrower’s perspective, lenders are perfect substitutes.

Let S^L be the multi-set containing the contracts offered by each lender at date zero. From lender l 's perspective, the payoff of offering a contract $s^l \in S$, when the other lenders offer contracts S^{-l} (and hence the menu is $S^L = S^{-l} \uplus \{s^l\}$) and the common type is θ , is

$$\sigma(s^l, S^L) \left\{ -K + \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_z^l, a) \right\}, \quad (2)$$

where $\sigma(s^l, S^L)$ is the probability that the buyer accepts the contract s^l , given the contracts posted. This notation implicitly assumes that the buyer's decision does not depend on the identity of the lender, only on the contract that the lender offers. We will assume this in the equilibria we study, and note that it is consistent with the assumption that the borrower's utility does not depend on the lender she chooses, only on the design of the contract.

Assuming the borrower chooses to borrow, his expected payoff for contract s is

$$\sum_{\theta' \in \Theta, a \in A, z \in Z} \mu(\theta'; S^L) \theta'(a, z) \phi_B(s_z, a), \quad (3)$$

where $\mu(\theta'; S^L)$ denotes the borrower's beliefs about the distribution of the lender's common type θ' after observing the menu S^L . The beliefs $\mu(\theta'; S^L)$ are central to our theory. The borrower does not observe the lender's common type θ ; initially, she has prior μ_0 over the set of types Θ , but might refine these beliefs based on the menu of securities offered. It is important to note that, because the type θ is common across lenders, an optimal mechanism could allow the borrower to solicit this information and then negotiate a contract (Cremer and McLean [1988]). The market structure we impose, which we believe is realistic in many contexts, prevents the buyer from conducting this sort of auction.⁶

⁶The mechanism of Cremer and McLean [1988] also requires commitment, and hence is inconsistent with our ex-post efficiency assumption.

Having discussed the basic structure of the model, we next describe the equilibrium concept and the refinements for off-equilibrium beliefs that we employ.

3.3 Equilibrium Definition

The basic equilibrium concept we use is perfect Bayesian. Given the strategies of the other lenders (S^{-l*}) and the buyer (σ^*), and the common type θ , we require that lender l posts

$$s^l \in \arg \max_{s \in \mathcal{S}} \sigma^*(s, S^{-l*} \cup \{s\}) \{-K + \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_z^l, a)\}, \quad (4)$$

if that strategy yields weakly positive profits, and otherwise does not participate. That is, each lender's choice of contract maximizes her utility, given the strategies of the other lenders and borrower.

If the borrower is offered any contracts, she must choose a strategy $\sigma(s^l, S^L)$ such that, given posterior beliefs $\mu(\cdot; S^L)$, if $\sigma(s^l, S^L) > 0$, then

$$s^l \in \arg \max_{s \in S^L} \sum_{\theta' \in \Theta, a \in A, z \in Z} \mu(\theta'; S^L) \theta'(a, z) \phi_B(s_z^l, a), \quad (5)$$

and

$$\sum_{\theta' \in \Theta, a \in A, z \in Z} \mu(\theta'; S^L) \theta'(a, z) \phi_B(s_z^l, a) \geq 0. \quad (6)$$

In words, the borrower must maximize his utility given the menu of contracts being offered.

The equilibrium strategies of the lenders create a multi-set valued function $S^*(\theta)$ that describes the menu of securities that might be offered, given the common type. If the buyer observes a menu S^L for which there exists a type θ' such that $S^L = S^*(\theta')$, then she must update her beliefs according

to Bayes' rule:

$$\mu(\theta; S^L) = \frac{\mu_0(\theta)1(S^L = S^*(\theta))}{\sum_{\theta' \in \Theta} \mu_0(\theta')1(S^L = S^*(\theta'))}. \quad (7)$$

This does not, of course, pin down what the buyer believes when he observes some menu S^L that could not have been generated from the equilibrium strategies $S^*(\theta)$, for any $\theta \in \Theta$ with $\mu_0(\theta) > 0$. For the purpose of determining if a conjectured set of strategies is an equilibrium, we only need to consider menus S^L that differ from a menu $S^*(\theta')$ for a single lender.

The result we are building towards is that there are many equilibria. This would be expected in the absence of refinements for off-equilibrium beliefs. Without refinements, the borrower can in effect dictate the contract by forming pessimistic beliefs when offered any other contract, justifying rejection. For this reason, we employ two refinements. The first refinement requires that the borrower believe the minimal number of lenders have deviated from equilibrium play. For concreteness, suppose the true common type is θ , and that all but one of the lenders offer an equilibrium contract for that type. The other lender deviates by offering another security that is not offered by type θ in equilibrium. Moreover, suppose the resulting menu could not have arisen from the equilibrium strategies of any type. Absent this refinement, the borrower could believe that multiple lenders have deviated. Imposing our refinement, and using the fact that there are at least three lenders, the borrower must instead correctly identify the deviating lender.

The second refinement we employ is the *D1* equilibrium refinement (Banks and Sobel [1987]). This refinement captures the intuition that, if confronted with a “deviating” contract, the borrower should believe the lender is of a type that would benefit from this deviation. Under our first refinement, the borrower is able to identify the deviating lender (when there is only a single deviating lender), and it is to the security offered by this lender that we apply the *D1* refinement. We believe our results are robust to using other refinements (aside from *D1*) that provide a similar intuition.

We use the standard definition of *D1*, and think of the borrower's “strategy” as consisting of

an acceptance probability ρ . A lender of type θ offering contract s' , instead of the equilibrium contract s , would benefit, given that the buyer accepts the deviating contract with probability ρ , if

$$\rho \left\{ -K + \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s'_z, a) \right\} \geq \sigma^*(s, S^*(\theta)) \left\{ -K + \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_z, a) \right\}. \quad (8)$$

The types for whom the set of $\rho \in [0, 1]$ satisfying this condition is maximal are the types with positive support in the buyer's beliefs following this deviation, $\mu(\cdot; S^{-l*}(\theta) \uplus \{s'\})$.

Looking ahead, we will show that in equilibrium, lender profits are zero, due to the effects of competition. As a result, the *D1* refinement will simply state that the buyer must place the support of her beliefs on types that would weakly profit from offering the deviating contract, if that contract were accepted and such a type exists. The buyer cannot believe the deviating lender is of a type such that the lender would lose money if the buyer accepted the deviating contract, unless every lender type would lose money if the contract were accepted (and in this case, the deviating contract would never be offered).

Our analysis will focus on a particular set of equilibria, symmetric pure-strategy equilibria. These equilibria are pure strategy equilibria and symmetric in the sense, for all types $\theta \in \Theta$, either all of the lenders offer the same security with certainty, $s(\theta)$, or none of the lenders offer a security. They are also symmetric in the sense that the borrower, faced with a menu of identical securities, chooses each lender with probability $|L|^{-1}$.

4 Preliminary Analysis

We begin our analysis by focusing on the effects of competition. Consider a symmetric pure-strategy equilibrium, and imagine that the lenders' profits from offering the contract $s(\theta)$ are strictly positive. Intuitively, this could not be an equilibrium. Suppose a lender offered a devi-

ating contract $s' \in S$ such that, for each index value $z \in Z$, the associated promise d'_z was less than the promise associated with the original contract, d_z . The buyer would be better off regardless of her beliefs, and therefore accept the contract with probability one. The lender, by sacrificing some profit, would capture the entire market, and be better off. Because of the monotonicity property of the buyer's indirect utility function and the continuity property of the lender's indirect utility function, standard Bertrand competition effects apply, and profits must be zero in equilibrium.

Lemma 1. *In any symmetric pure-strategy equilibrium, lender profits must be zero.*

Proof. See the appendix, section B.1. □

We next introduce an assumption to ensure that there are contracts which can satisfy both the lender and borrower's participation constraints.

Assumption 1. *There exists a contract $s \in S_I$ that offers weakly positive utility to the borrower, while satisfying the lender's participation constraint. That is, the problem*

$$\max_{s \in S_I} \sum_{a \in A, z \in Z} \theta_0(a, z) \phi_B(s, a)$$

subject to the constraint $\sum_{a \in A, z \in Z} \theta_0(a, z) \phi_L(s, a) \geq K$ is feasible and has a weakly positive solution.

Because we have assumed that the marginal distributions are the same for all types $\theta \in \Theta$, this assumption is sufficient to ensure that for any type, there is a contract that both the borrower and lender would be willing to accept under full information.

Next, we discuss the existence of a “best” equilibrium. Consider a symmetric, pure-strategy equilibrium, described by an offer of the contract $s(\theta)$. Suppose that the mapping between types

θ and securities $s(\theta)$ is one-to-one. In this case, in equilibrium, the borrower knows the lenders' common type. Define a full-information optimal contract as

$$\bar{s}(\theta) \in \arg \max_{s \in \mathcal{S}} \sum_{a \in A, z \in Z} \theta(a, z) \phi_B(s_z, a), \quad (9)$$

subject to the constraint that $\sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_z, a) = K$. By assumption 1, the solution to the above maximization can offer the buyer weakly positive utility for all types $\theta \in \Theta$.

A set of full-information optimal contracts is on the Pareto frontier for all θ , and offers the lender zero profit. As a result, for any deviating contract a lender might be willing to offer, if the borrower correctly inferred the lenders' true type, the borrower would weakly prefer the full-information optimal contract being offered. The D1 refinement in our model allows the borrower to make this inference, and the presence of a competing lender allows the borrower to choose the equilibrium full-information optimal contract instead of the deviating contract. The following proposition summarizes this logic:

Proposition 1. *The pure-strategy symmetric equilibrium $s(\theta) = \bar{s}(\theta)$ exists.*

Proof. See the appendix, section B.2. □

The above proposition describes a “best” pure-strategy symmetric equilibrium, in which a full-information optimal contract is offered. Our main results describe the conditions under which another type of pure-strategy symmetric equilibrium exists. This alternative equilibrium is notable because it uses a non-contingent contract, is a pooling equilibrium, and is Pareto-inferior to the “best” equilibrium, from an ex-ante perspective.

We say a contract is “non-contingent” if $s_z = s_{z'}$ for all $z, z' \in Z$; that is, the contract does not make use of the index. We will consider the existence of a non-contingent contract pooling

equilibrium, in which, for all $\theta \in \Theta$ with $\mu_0(\theta) > 0$,

$$s_z(\theta) = s^* \in \arg \max_{s \in S_D} \sum_{a \in A} p(a) \phi_B(s, a) \quad (10)$$

subject to the constraint that $\sum_{a \in A, z \in Z} p(a) \phi_L(s, a) = K$. By assumption 1 and condition 1, this contract can offer the buyer weakly positive utility for all types $\theta \in \Theta$.

We next impose two assumptions that setup the ‘‘puzzle’’ this paper addresses in the context of the model. The first of these assumptions ensures that, for all θ in the support of μ_0 , except the uninformative type θ_0 , the index $z \in Z$ is related to the external state $a \in A$. The second assumption describes the apparent failure of the agents to share risks under the non-contingent contract.

The first assumption imposes structure on the support of μ_0 . There are a variety of ways of defining ‘‘more interrelated’’ in the context of joint probability distributions with identical marginal distributions. For two variables (i.e. $a \in A$ and $z \in Z$), many of these orders are equivalent (Meyer and Strulovici [2012]). We assume that every element in the support of μ_0 has more interrelatedness between $a \in A$ and $z \in Z$ than the uninformative type θ_0 . We can express this in terms of ‘‘lower orthant dominance,’’ meaning that for all $\bar{a} \in A$ and $\bar{z} \in Z$, and all $\theta \in \Theta$ with $\mu_0(\theta) > 0$,

$$\sum_{a \in A: a \preceq \bar{a}} \sum_{z \in Z: z \preceq \bar{z}} (\theta(a, z) - \theta_0(a, z)) \geq 0. \quad (11)$$

Intuitively, lower orthant dominance captures the notion that (under θ) lower values of the external state tend to coincide with lower values of the index. Because the marginal distributions are identical, lower orthant dominance implies upper orthant dominance, meaning that high values of a tend to coincide with high values of z . In other words, the index is relevant for all $\theta \neq \theta_0$.

Assumption 2. *The support of the prior μ_0 includes only types θ that lower-orthant dominate θ_0 . This assumption implies (by Atkinson and Bourguignon [1982] and others) that $\sum_{a \in A, z \in Z} (\theta(a, z) -$*

$\theta_0(a, z))f(a, z) \geq 0$ for all super modular functions f and all θ in the support of μ_0 .

We will impose additional conditions on the type space as part of our “resolution” of the puzzle, which we discuss in the next section.

Our second assumption is that the non-contingent contract s^* is not the full-information optimal contract $\bar{s}(\theta)$ for at least some types. That is, there are risk-sharing opportunities (if $\theta \neq \theta_0$) that are not being exploited under the non-contingent contract. This assumption is, in essence, a statement that the puzzle motivating this paper exists in the model.

Assumption 3. *There exists a $\theta \in \Theta$ for which $\mu_0(\theta) > 0$ such that $\bar{s}(\theta)$ is not equal to any non-contingent optimal contract s^* .*

This assumption ensures that our results are not trivial. The “non-contingent” equilibrium (if it exists) is ex-ante Pareto-inferior to the “best” equilibrium described previously, in the sense that the expected payoff to the borrower under the prior μ_0 is strictly lower.⁷

These assumptions are largely unrelated to the indirect utility functions (and hence to our examples). Assumption 1 is a standard feasibility assumption, of the sort that appears in Innes [1990] and many other contracting papers. Assumption 2 is an assumption about the external state and index, and hence not directly connected to the indirect utility functions. Assumption 3 implies that the probability distribution of the external state $a \in A$ affects, in some way, the optimal security design, and that there is some type θ for which the index z is usefully interrelated to the external state.

⁷In both equilibria, the lenders earn zero profits, regardless of whether they are selected by the borrower or not.

5 Risk-Sharing Failure in Equilibrium

In this section, we provide sufficient conditions for the existence of a “non-contingent” equilibrium. This equilibrium will exist despite its ex-ante Pareto-inferiority to the “best” equilibrium discussed above.

Our second condition is defined using the variable λ^* , which is the Pareto-weight associated with the non-contingent contract s^* :

$$s^* \in \arg \max_{s \in S_D} \sum_{a \in A} p(a) U(s, a; \lambda^*). \quad (12)$$

This Pareto-weight is the particular Pareto-weight that delivers an expected payoff of K to the lender. Our condition requires that the marginal value of a promise to the lender is higher in bad external states than in good external states, but the marginal social value of a promise to lender is lower in bad external states than in good external states. This condition captures the notion that the borrower is more risk-averse than the lender (and is consistent with assumption 3). It also captures the notion that the lender is risk-averse with respect to promises.

There are two competing forces that will determine whether the lender is risk-averse with respect to promises. The lender is risk-averse about external states, and hence has higher marginal utility in worse external states. However, promises are more likely to be fulfilled in better external states. In the context of debt contracts, repayment of the full face value is more likely in good external states. If the first of these forces weakly dominates the second, the lender have a marginal value of promises that is higher in bad states. In this case, we will say that the lender is risk-averse over promises with respect to the external state.

Condition 2. For all $d', d \in D$ with $d' > d$, $\phi_L(s_{d'}, a) - \phi_L(s_d, a)$ is weakly decreasing in a , and $U(s_{d'}, a; \lambda^*) - U(s_d, a; \lambda^*)$ is weakly increasing in a . Equivalently, $\phi_L(s_d, a)$ is sub modular in

(d, a) , and $U(s_d, a; \lambda^*)$ is super modular in (d, a) .

Note that this condition applies only to security designs in the Pareto-optimal family (for example, debt securities); it does not impose restrictions on sub-optimal securities.

This condition can be understood as consisting of several claims. The first claim is that the “marginal benefit of debt” to the lender, $\phi_L(s_{d'}, a) - \phi_L(s_d, a)$, is monotone in the aggregate state, regardless of the levels of debt involved. The first part of this claim can be thought of as defining the order on the aggregate states— up to this point, nothing has depended on that order. The second part (“regardless of the level of debt”) is the key point. The second claim is that the “marginal cost of debt” to the borrower, $\phi_B(s_d, a) - \phi_B(s_{d'}, a)$, is monotone and increases in the same direction as the marginal benefit of debt to the lender.⁸ In other words, states in which the lender would really like larger promises are also states in which the borrower would really prefer not to make larger promises. The third claim is that the borrower is “more risk averse” than the lender in this sense. That is, in states in which the lender would really like a large promise, and the borrower would really prefer a small promise, the latter effect dominates, and under the Pareto weight λ^* , it is more efficient to have smaller promises when both “marginal cost” and “marginal benefit” are high. In other words, the optimal contract would involve the lender insuring the borrower, and because preferences are aligned, this is costly for the lender.

As suggested by this description, our results do not really depend on the ordering over the external states. That is, the proof of proposition 2 below would go through almost unchanged if we imposed, instead of condition 2, that $\phi_L(s_d, a)$ was super modular and that $U(s_d, a; \lambda^*)$ was sub modular.

Finally, our third condition requires that the type space be sufficiently rich and include the uninformative type. Recall, by assumption 2, that every type θ in the support of μ_0 weakly lower-

⁸Condition 2 implies that ϕ_B is the difference of a super modular and sub modular function, and hence super modular.

orthant dominates the type θ_0 .

Condition 3. The support of μ_0 includes in the uninformative type θ_0 . Moreover, for all θ in the support of μ_0 and all $\theta' \in \Theta$ that lower-orthant dominate θ_0 , if θ lower-orthant dominates θ' , then θ' is in the support of μ_0 .

This condition ensures that for any type, there is a rich set of “less informative” types, which limits the lender’s simultaneous ability to signal her type while providing risk-sharing benefits.

Under these conditions, we prove a general result.

Proposition 2. *Under assumptions 1, 2, and 3, and conditions 1, 2, and 3, there exists a symmetric pure-strategy equilibrium in which $s(\theta) = s^*$.*

Proof. See the appendix, section B.3. The proof relies on results from Meyer and Strulovici [2015]. □

This proposition establishes that the conditions given above are sufficient for the existence of a non-contingent equilibrium. Intuitively, if it is not efficient for the borrower to hedge the lender, the deviations necessary to separate from the uninformative type are never welfare-improving. Our conditions are designed to ensure that this is the case. Given a particular specification for the indirect utility functions, our conditions can be checked, and used to demonstrate the existence of a risk-sharing failure. We next turn to our examples, and describe when our conditions will be satisfied.

The key condition in our result is condition 2, lender risk aversion over promises with respect to the external state. In example 1 (utility functions), promises are payments, and hence the assumption requires that $\beta_L(a)$ is weakly decreasing in a . In other words, the lender is weakly risk-averse with respect to the external state. To generate super-modularity in the social welfare function, we must have $\lambda^* \beta_L(a) - \beta_B(a)$ weakly increasing in a .

In example 2 (the CSV example), we generate lender risk-aversion through a strictly decreasing $\beta_L(a)$. Lenders must be strictly risk-averse in this example to offset the effect of promises being more likely to be repaid in better external states. We also endogenously generate borrower risk-aversion over external states, using a utility function and non-verifiable income for the borrower. For details, see Appendix section §A.

In the next section of the paper, we discuss a standard security design problem that satisfies the assumptions of our general model.

6 Failing to Index Mortgages

In this section, we discuss a simple example of mortgage lending that illustrates the intuition behind the results of the previous section. Specifically, we consider a setting in which a mortgage borrower has hidden information about her endowment. As a result, the borrower can only raise financing through debt-like contracts, in which the house serves as collateral, along the lines of Hart and Moore [1998]. These contracts sometimes cause inefficient liquidation in equilibrium. Moreover, the degree of inefficiency of liquidation depends on the external state. As a result, there are benefits to writing contracts in which the face value of the debt depends on the external state. We will show that under mild assumptions, this model of mortgage lending satisfies assumption 1 and conditions 1 and 2.⁹ Therefore, there exists an equilibrium in which the face value of the debt does not depend on the external state, despite the benefits of such contracts.

Consider a household seeking to borrow K dollars for the purchase of a home from one of set of competitive lenders. The borrower has a uniformly distributed¹⁰ random endowment $e \sim F(e)$

⁹The other assumptions/conditions for our main result apply to the type space, and are assumed for the purposes of this example.

¹⁰We assume a uniform distribution to make the underlying conditions as transparent as possible.

with full support on $[0, \bar{e}]$. The realization of e is hidden to lenders and non-verifiable; the borrower can make a report \tilde{e} . In the notation of our general model, the idiosyncratic states I are pairs (e, \tilde{e}) . In addition to uncertainty in the borrower's endowment, there is also uncertainty in some external state $a \in A$ that will affect the value the borrower's cost of defaulting in a manner we describe below. As in our general model, the external state is unobservable and can only be contracted upon via an index $z \in Z$. This index could represent a local house price index that summarizes the external conditions of the housing market affecting the value of the home. A security is a map from these idiosyncratic states and the index to payments, $Z \times I \rightarrow \mathbb{R}^+$, as in the general model. The set of admissible contracts conditional on the index value, S_I , is restricted to depend only on the report, \tilde{e} , and not on the actual endowment e . For expositional convenience, we also restrict the set S_I to be weakly increasing in the report, \tilde{e} , although nothing depends on this assumption.

In this example, the contract payment is not necessarily the payment the lender will receive. The borrower has an inalienable option to sell the house, which results in proceeds $L \in [0, K]$.¹¹ The lender has priority to the proceeds from the sale of the house, meaning that L first goes towards paying off the contract, and anything remaining goes to the borrower. In other words, the house serves as collateral for the contract between the lender and the borrower. However, if L is insufficient to cover the payment demanded by the contract, the lender has no additional recourse.

If the borrower chooses to sell the house, she bears a private, monetary cost $C(a) \geq 0$. This cost is the way in which the external state enters the problem. We assume that the private cost of liquidation is large relative to the endowment, $C(a) \geq \bar{e}$ for all $a \in A$, and that liquidation is inefficient, $C(a) \geq L$ for all $a \in A$. If the borrower is unable to pay the required repayment out of her endowment, she must liquidate. The borrower is risk-neutral over the value of her endowment net of the contract repayments and/or liquidation costs. Given an endowment e , a contract s , a

¹¹We assume that the liquidation value is constant across states of the world to simplify our exposition.

realization of the index z , and an external state a , the borrower will choose an optimal report $\tilde{e}^*(e; s_z, a)$ and liquidation probability $\rho^*(e; s_z, a)$ to solve

$$(\tilde{e}^*(e; s_z, a), \rho^*(e; s_z, a)) = \arg \min_{\tilde{e}, \rho} \{(1 - \rho)s_z(\tilde{e}) + \rho(\min\{s_z(\tilde{e}), L\} - L + C(a))\} \quad (13)$$

subject to the constraint that, if the borrower does not liquidate, repayment must be feasible:

$$e - (1 - \rho)s_z(\tilde{e}) - \rho(\min\{s_z(\tilde{e}), L\} - L) \geq 0. \quad (14)$$

In other words a feasible strategy is one for which the borrower has enough cash, either via her endowment or from the proceeds from liquidation, to cover the required repayment.

Because $s(\tilde{e})$ is weakly increasing, it is without loss of generality to assume that $\tilde{e}^*(e; s_z, a) = 0$ in all states. Moreover, the borrower will default only when she is unable to make the required repayment:

$$\rho^*(e; s_z, a) = \mathbf{1}(e \geq s_z(0)). \quad (15)$$

In words, the borrower's payment is always weakly increasing in her report, regardless of her liquidation decision, and as such she will always report the minimum possible endowment. Moreover, since liquidation is inefficient, the borrower will only default when her endowment is insufficient to pay the lowest possible amount she could pay given s_z . Thus, any security in this model implies an allocation that is equivalent to a security that does not depend on the report $s_z(\tilde{e}) = d$ for all \tilde{e} . Such securities are essentially defaultable mortgages as they require the payment of some face value d . If repayment is not made, liquidation (i.e., foreclosure) occurs and the lender receives the minimum of the liquidation value and the face value. In this way, this simple model gives rise to the common structure of mortgage lending that we observe in practice.

We denote the lender's marginal utility in external state a as $\beta_L(a) > 0$. For the borrower, the

marginal utility is endogenous to the problem. The lender, who is presumed to be connected to the broader financial markets, is risk-neutral with respect to the borrower's endowment but not with respect to external risk. We normalize the marginal utility so that $\sum_{a \in A} p(a) \beta_L(a) = 1$. The indirect utility function of the lender is given by

$$\phi_L(s_d, a) = \beta_L(a)(1 - F(d))d + \beta_L(a)F(d) \min\{d, L\}. \quad (16)$$

while the borrower's indirect utility function is given by

$$\begin{aligned} \phi_B(s_d, a) &= E[e] - (1 - F(d))d - F(d)(\min\{d, L\} + C(a) - L) \\ &= E[e] - \beta_L(a)^{-1} \phi_L(s_d, a) - F(d)(C(a) - L). \end{aligned} \quad (17)$$

We next discuss our assumptions on the magnitudes and orderings of the variables that allow us to apply our main results from the general model above.

Assumption 4. *The lender's marginal utility, $\beta_L(a)$, is weakly decreasing in a , the required investment is sufficiently small,*

$$K \leq \frac{(\bar{e} + L)^2}{4\bar{e}},$$

and the quantity

$$C(a) - \psi \beta_L(a),$$

is weakly decreasing in a , where

$$\psi = \bar{e} \frac{\bar{e} + \sum_{a' \in A} p(a') C(a')}{\bar{e} + L}.$$

The first part of this assumption requires that the lender be risk averse over aggregate states,

in the sense that the lender's marginal utility is higher in worse aggregate states. The second part ensures that trade is feasible. The final part of the assumption effectively ensures that the borrower is "more risk averse" than the lender in the appropriate sense. The borrower is endogenously risk averse due to the inefficiencies associated with liquidation, and these inefficiencies are larger in worse states of the world ($C(a)$ is weakly decreasing in a). The final part of the assumption states that not only is the borrower endogenously risk averse, in this sense, but that the borrower is sufficiently "more risk averse" than the lender.

We begin by noting that flat contracts are always Pareto-optimal, and that trade is feasible.

Lemma 2. *Under assumption 4, the model described in this section satisfies condition 1 (debt is always optimal) and assumption 1 (feasibility). The set D is the interval $[0, \frac{\bar{e}+L}{2}]$.*

Proof. See the appendix, section B.4. □

This lemma establishes the feasibility of trade (the lender can earn at least K) and the ex-post efficiency of debt. The result only depends on the second part of assumption 4, that the required investment is sufficiently small. The next lemma establishes the super/sub-modularity condition (condition 2) using this result and the other parts of assumption 4.

Lemma 3. *Assumption 4 implies that condition 2 holds.*

Proof. See the appendix, section B.5. □

Because the indirect utility functions in this example satisfy the conditions of our general theorem, we have the following corollary:

Corollary 1. *In the mortgage model described in this section, there exists an equilibrium characterized by non-contingent debt contracts.*

The result in corollary 1 summarizes our answer to the question, “why aren’t mortgages indexed to house prices?” It illustrates a concrete example in which debt contracts tied to an index would be welfare-improving compared to debt contracts that are not indexed, but the latter arise in equilibrium due to adverse selection about the quality of index.

7 Variations and Extensions

In this section, we discuss modifications and extensions to the model. We begin by discussing a model with positive profits for lenders, that nevertheless retains the competition between lenders. In this case, our results go through essentially unchanged. We then discuss what would happen with a single, monopoly lender. We will see that there is no “full-information optimal contracts” equilibrium with a monopoly lender, but there is still a non-contingent equilibrium. Finally, we will discuss how to extend our results to settings in which there is adverse selection about the marginal distribution of the index (q), or only about the distribution of the external states and not about the index quality.

7.1 Profitable Lending

In this extension, we describe a model in which lenders make positive profits in equilibrium, but nevertheless face competition. We introduce profits into the economy by assuming that each lender faces a convex cost in the number of loans she makes, and that there is a unit mass of borrowers.¹² Let Q_l be the number of loans made by lender l . A lender of type θ who makes Q loans using contract s earns

$$\Pi(s, Q, \theta) = Q \left\{ \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_z, a) \right\} - C(Q), \quad (18)$$

¹²Introducing profits in this way is an old idea, described in the textbook of Tirole [1988].

where $C(Q)$ is a convex, twice differentiable function with $C(0) = 0$ and $C'(|L|^{-1}) = K$.

In this analysis, the $D1$ refinement can be more complicated than in our baseline model. If one considers a deviation in which the lender offers a single, marginal borrower a different contract, then the criteria the same as in our main analysis, because (in equilibrium) the marginal profit of each lender is zero. If however, the lender contemplates a deviation in which he offers a deviating contract to all borrowers, then substantial profits could be at stake, because the average profits of lenders are positive.

In this case, the $D1$ refinement requires that the borrower place her beliefs on the lender type who would break-even under the smallest amount of the demand for the deviating contract. This is equivalent to saying that the borrower must believe the lender is of a type for whom the difference between the marginal profit of the deviating contract and the marginal profit of the equilibrium contract is maximal.

Surprisingly, perhaps, our non-contingent equilibrium exists under the same conditions in this model. The intuition comes from the proof description in section §5. When a lender with a “good” index offers a contract that insures the borrower, the lender requires a higher expected value of repayments to be indifferent between the deviating contract and the non-contingent contract. However, a lender with an irrelevant index could offer the same deviating contract at a profit, and therefore (in the case of profitable lending) the borrower must believe that the lender is of this type, or of a type that is even worse from the perspective of the borrower.

7.2 Monopoly Lending

In this extension, we consider what type of equilibrium can obtain when the lender has monopoly power. Specifically, we assume that a single lender can make a take it or leave it offer to the borrower and that if the borrower rejects this offer, she receives her outside option. Neither of the

full-information optimal contract or the non-contingent contract defined in section §4 are equilibria, because both offer positive surplus to the borrower and zero surplus to the lender.

To study the monopoly case, imagine parameterizing both the full-information optimal contract and the non-contingent contract by the required investment. Additionally, suppose that there exists a $\bar{K} > K$ such that the full-information optimal contract, $\bar{s}(\theta, \bar{K})$, results a payoff for the borrower equal to her outside option. Likewise, suppose that there exists a $K^* > K$ such that the non-contingent contract, $s^*(K^*)$, also results in a payoff for the borrower equal to her outside option. In this sub-section, we will ask whether there exist equilibria with the contracts $\bar{s}(\theta, \bar{K})$ and $s^*(K^*)$. We will continue to impose the *D1* refinement on off-equilibrium contract offers.

The answer is no for the full-information contract, and yes for the non-contingent contract. The existence of the non-contingent contracts equilibria follows from the proof of proposition 2—nothing in that proof depended on the specific value of K . The only effect of competition was to allow the borrower to choose a contract from another lender. Although the type θ is common to all lenders, because the non-contingent contract’s payoff for the borrower does not depend on θ , the borrower’s inference about θ does not change the appeal of the non-contingent contract. It is as-if the borrower had a fixed outside option instead, which is what is assumed in the monopoly case.

However, for the full-information contract, competition is essential. For the uninformative type (θ_0), the full-information contract is identical to a non-contingent contract. For some other type (by assumption 3), the full-information contract is contingent, and by condition 2, this contract offers a higher payoff to the uninformative type than the non-contingent contract. As a result, the uninformative type is tempted to deviate. When there are other lenders, the borrower can use their offers to determine the common type, and avoid being “tricked” by this deviation. With a monopoly lender, this is not possible, and as a result there is no full-information contract equilibrium. In summary, competition is necessary for the existence of the “best” equilibrium, but the

non-contingent equilibrium always exists.

7.3 Adverse Selection about Marginal Distributions

Throughout the paper, we have assumed that the set Θ contained only joint distributions of the external state and index with marginal distributions $p(a)$ and $q(z)$. Suppose we relax this, and require only that the marginal distribution over external states, $p(a)$, be the same for all types. Under this assumption, there is no adverse selection about the true external state, only about the index, as in the main part of the paper. Intuitively, adding additional dimensions of adverse selection cannot improve the situation, and should only reinforce the non-contingent contracts equilibrium.

Formally, let $q(z; \theta)$ denote the marginal distribution of the index associated with type θ , and let $\Theta(q)$ be the set of all joint distributions with marginals $q(z)$ and $p(a)$. Let \mathcal{Q} be the set of all marginal distributions for the index, and let Θ be the union of all $\Theta(q)$ for each $q \in \mathcal{Q}$. Modify our condition 3 (rich type space) and assumption 3 (a full-information optimal contract is contingent) so that they apply to each $\Theta(q)$ such that $\mu(\theta) > 0$ for some $\theta \in \Theta(q)$. In other words, there is a “rich” type space and an opportunity for risk-sharing for each possible marginal distribution of the index. Under this condition, the proof of proposition 2 is essentially unchanged, and the result holds.

7.4 Adverse Selection on External States

In this extension, we modify the model of the main text to consider the case in which the index is known to be perfect, but there is adverse selection about the marginal distribution of the aggregate state. This sort of adverse selection is closer to the problems studied in the literature (e.g. Spier [1992], Asriyan [2015]). We build on the notation used in the previous extension. We assume that the set Z is identical to the set A , and that each $\Theta(q)$ is a singleton, containing only the joint

distribution

$$\theta(a, z) = \delta(a, z)q(z), \quad (19)$$

where $\delta(a, z)$ is one if $a = z$ and zero otherwise. Adverse selection occurs because, in this context, there are types in Θ with different values of $q(z; \theta)$.

In this setting, the “rich type space” condition (condition 3) is irrelevant, as is the assumption about lower orthant dominance (assumption 2). We continue to impose our feasibility assumption (assumption 1) for each $\Theta(q)$, and the condition on the security space (condition 1). Our result in this section does not depend on a super modularity condition (like condition 2), and therefore we will not discuss how to adapt that condition to this setting. As in the previous extension, we assume that assumption 3 applies for each $\Theta(q)$. In this extension, this assumption implies that every full-information optimal contract is non-contingent.¹³

For technical reasons, we assume that the support of the prior beliefs $\mu_0(\cdot)$ is a closed set, which was not required in the main text. We also assume that all types $\theta \in \Theta$ for which $\mu_0(\theta) > 0$ are associated with marginal distributions that have full support. In other words, $q(z; \theta) > 0$ for all $z \in Z$ and $\theta \in \Theta$ such that $\mu_0(\theta) > 0$. This generalizes the full support assumption of the main text.

Define the mapping $\Theta^*(\theta)$ as a set-valued function

$$\Theta^*(\theta) = \{\theta' \in \arg \max_{\theta'' \in \Theta: \mu_0(\theta'') > 0} \sum_{a \in A, z \in Z} \theta''(a, z) \phi_L(\bar{s}_z(\theta), a)\}, \quad (20)$$

where $\bar{s}_z(\theta)$ is the full-information optimal contract associated with the type θ . The set $\Theta^*(\theta)$ is the set of types in the support of $\mu_0(\cdot)$ that would earn the highest payoff from offering the security $\bar{s}_z(\theta)$.

¹³We discuss how to weaken this assumption below.

The following lemma (which is based on standard fixed-point arguments) states that there is a fixed point to this mapping.

Lemma 4. *There exists a θ^* such that $\theta^* \in \Theta^*(\theta^*)$. For any such θ^* , for all $\theta' \in \Theta^*(\theta^*)$, $\bar{s}(\theta') = \bar{s}(\theta^*)$.*

Proof. See the appendix, section B.6. □

This type, θ^* , can essentially “prove itself” by offering its full-information optimal contract. When a lender of type θ^* offers the contract $\bar{s}(\theta^*)$, it receives a payoff of K . All other types either receive a payoff strictly less than K , or receive a payoff equal to K and have an identical full-information optimal contract. Hence, under the $D1$ refinement, the borrower must believe that she is being offered a full-information optimal contract.

By assumption, every full-information optimal contract is not equal to a non-contingent contract. By the Pareto-optimality of the full-information optimal contract, the borrower must be willing to accept this contract. Therefore, there cannot be an equilibrium in which a non-contingent contract is employed.¹⁴

What makes this setting different than the one studied in the main text? The key is the idea that there is no type for which the full-information optimal contract is equal to the non-contingent contract. When there is adverse selection about the distribution of external states, this makes sense; the only way a non-contingent contract could be optimal is if some type of lender knew with certainty what the ex-post “fair” (payoff equal to K) level of debt was. In contrast, in the case emphasized in the main text, a non-contingent contract can be optimal so long as it is possible that

¹⁴This argument does not really depend on every full-information optimal contract being contingent. It only requires that the type(s) with payoff equal to K from a non-contingent contract, when all other types have lower payoffs from that contract, have non-contingent full-information optimal contracts. Usually, this will mean that the type with the “best” $q(z)$ in a first-order stochastic dominance sense has a non-contingent contract.

the index is irrelevant; perfect foresight about the external state is not required for a non-contingent contracts equilibrium.

In some sense, this result can be viewed as pointing to the necessity of an assumption like condition 3 in the main text. If borrower knew the index was at least somewhat relevant, the type with the least relevant index could “prove herself” and eliminate the non-contingent equilibrium. In this case, an equilibrium with a minimal (and ex-ante sub-optimal) level of indexing would exist. Non-contingency could be restored in this case by, following Spier [1992], by introducing a fixed cost of using the index in addition to asymmetric information. In this case, a non-contingent equilibrium would exist so long as the “worst” type was sufficiently bad, relative to the fixed cost.

8 Conclusion

We have introduced a theory to explain the widespread lack of indexation observed in contracts. Intuitively, when a borrower is offered a contract that includes insurance, she is concerned that the insurance is not actually relevant for the risks she faces. Under the conditions described in our model, this effect is strong enough to allow the borrower to reject that offer, and choose instead a contract without insurance from a different lender. As a result, equilibria that feature little or no risk-sharing can arise, even though they are ex-ante Pareto-dominated by equilibria that feature full risk-sharing.

References

Philippe Aghion and Benjamin Hermalin. Legal restrictions on private contracts can enhance efficiency. *Journal of Law, Economics, & Organization*, 6(2):381–409, 1990.

- George A Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- Franklin Allen and Douglas Gale. Measurement distortion and missing contingencies in optimal contracts. *Economic theory*, 2(1):1–26, 1992.
- Luca Anderlini and Leonardo Felli. Incomplete written contracts: Undescribable states of nature. *The Quarterly Journal of Economics*, pages 1085–1124, 1994.
- Vladimir Asriyan. Balance sheet recessions with informational and trading frictions. *Available at SSRN 2565603*, 2015.
- Anthony B Atkinson and Francois Bourguignon. The comparison of multi-dimensional distributions of economic status. *The Review of Economic Studies*, pages 183–201, 1982.
- Jeffrey S Banks and Joel Sobel. Equilibrium selection in signaling games. *Econometrica: Journal of the Econometric Society*, pages 647–661, 1987.
- Pierpaolo Battigalli and Giovanni Maggi. Rigidity, discretion, and the costs of writing contracts. *The American Economic Review*, 92(4):798–817, 2002.
- B Douglas Bernheim and Michael D Whinston. Incomplete contracts and strategic ambiguity. *American Economic Review*, pages 902–932, 1998.
- John Y Campbell, Stefano Giglio, and Parag Pathak. Forced sales and house prices. *The American Economic Review*, 101(5):2108–2131, 2011.
- Andrew Caplin, James H Carr, Frederick Pollock, Zhong Yi Tong, Kheng Mei Tan, and Trivikraman Thampy. Shared-equity mortgages, housing affordability, and homeownership. *Housing Policy Debate*, 18(1):209–242, 2007.

- Karl E Case, Robert J Shiller, and Allan N Weiss. Mortgage default risk and real estate prices: the use of index-based futures and options in real estate. Technical report, National Bureau of Economic Research, 1995.
- Jacques Cremer and Richard P McLean. Full extraction of the surplus in bayesian and dominant strategy auctions. *Econometrica: Journal of the Econometric Society*, pages 1247–1257, 1988.
- Mathias Dewatripont and Eric Maskin. Contractual contingencies and renegotiation. *The RAND Journal of Economics*, pages 704–719, 1995.
- Douglas Gale and Martin Hellwig. Incentive-compatible debt contracts: The one-period problem. *The Review of Economic Studies*, 52(4):647–663, 1985.
- Jerry R Green and Jean-Jacques Laffont. Renegotiation and the form of efficient contracts. *Annales d'Economie et de Statistique*, pages 123–150, 1992.
- Daniel Greenwald, Tim Landvoigt, and Stijn Van Nieuwerburgh. Financial fragility with sam? 2017.
- Oliver Hart and John Moore. Incomplete contracts and renegotiation. *Econometrica: Journal of the Econometric Society*, pages 755–785, 1988.
- Oliver Hart and John Moore. Default and renegotiation: A dynamic model of debt. *The Quarterly Journal of Economics*, 1998.
- Robert D Innes. Limited liability and incentive contracting with ex-ante action choices. *Journal of economic theory*, 52(1):45–67, 1990.
- Ola Kvaløy and Trond E Olsen. Endogenous verifiability and relational contracting. *The American Economic Review*, 99(5):2193–2208, 2009.

- Eric Maskin and Jean Tirole. Unforeseen contingencies and incomplete contracts. *The Review of Economic Studies*, 66(1):83–114, 1999.
- Margaret Meyer and Bruno Strulovici. Increasing interdependence of multivariate distributions. *Journal of Economic Theory*, 147(4):1460–1489, 2012.
- Margaret Meyer and Bruno Strulovici. Beyond correlation: Measuring interdependence through complementarities. 2015.
- Atif Mian and Amir Sufi. *House of debt: How they (and you) caused the Great Recession, and how we can prevent it from happening again*. University of Chicago Press, 2015.
- Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983.
- Tomasz Piskorski and Alexei Tchisty. An equilibrium model of housing and mortgage markets with state-contingent lending contracts. *Working Paper, National Bureau of Economic Research*, 2017.
- Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. In *Foundations of Insurance Economics*, pages 355–375. Springer, 1976.
- Ilya Segal and Michael D Whinston. The mirrlees approach to mechanism design with renegotiation (with applications to hold-up and risk sharing). *Econometrica*, 70(1):1–45, 2002.
- Robert J Shiller. Derivatives markets for home prices. Technical report, National Bureau of Economic Research, 2008.

Todd Sinai and Nicholas S Souleles. Owner-occupied housing as a hedge against rent risk. *Quarterly Journal of Economics*, 120(2), 2005.

Kathryn E Spier. Incomplete contracts and signalling. *The RAND Journal of Economics*, pages 432–443, 1992.

Jean Tirole. *The theory of industrial organization*. MIT press, 1988.

Jean Tirole. Incomplete contracts: Where do we stand? *Econometrica*, 67(4):741–781, 1999.

Jean Tirole. Cognition and incomplete contracts. *The American Economic Review*, 99(1):265–294, 2009.

Robert M Townsend. Optimal contracts and competitive markets with costly state verification. *Journal of Economic theory*, 21(2):265–293, 1979.

A Costly State Verification

In this appendix section, we discuss circumstances under which a version of the costly state verification model we discuss in Example 2 (Townsend [1979], Gale and Hellwig [1985], and others) satisfies the conditions of our main theorem, in particular condition 2. We introduce some functional forms in order to apply the results of our general model as simply as possible. The purpose of this exercise is to endogenously generate risk-aversion over external states for the borrower (we essentially assume it for the lender). The key modifications to the standard CSV model are the introduction of an external state, risk-aversion for the borrower, and the existence of non-verifiable income for borrower. Better external states induce a better (in a monotone-likelihood-ratio property sense) distribution of both verifiable and non-verifiable income. The presence of the non-verifiable

income, when combined with risk-aversion, means that the borrower has low expected marginal utility in good external states, even if she has promised a very high level of debt repayment. Consequently, the borrower is risk-averse over aggregate states, in the sense required by condition 2.

We now describe the specifics the modifications we make to Example 2. Recall that each idiosyncratic state $i \in I$ is a triple (x, y, y') , where x is the borrower's non-verifiable income, y is the borrower's verifiable income, and y' represents the borrower's report of her verifiable income, all of which are weakly positive reals. The conditional contract $s(y, y')$ can depend on both the true verifiable income and the report.¹⁵ The borrower's utility is

$$u(x + y - s(y, y')),$$

where $u(\cdot)$ is the borrower's strictly increasing, twice-differentiable, concave utility function.

The lender is risk-neutral with respect to the borrower's idiosyncratic state, but risk-averse with respect to the external state. Let $\beta_L(a)$ denote the lender's marginal utility given the external state. If the conditional contract differs depending on the true value, for a given value of the report, there is a verification cost paid by the lender. Let $c(y'; s) = \bar{c} > 0$ if there exists a y_1, y_2 such that the conditional contract s offers different payments for (y_1, y') and (y_2, y') , and zero otherwise.

Let $f(y|a)$ denote the distribution of y given a , and suppose it has the following functional form:

$$f(y|a) = q(y) \exp(a \ln(y) - \psi(a)).$$

In other words, the distributions $f(\cdot|a)$ are an exponential family whose sufficient statistic is the expected log verifiable income. The function $\psi(a)$ ensures that each $f(y|a)$ integrates to one. We

¹⁵It is without loss of generality to assume that it does not depend on a report of the borrower's non-verifiable income, since the borrower would also report the non-verifiable income level that minimized her repayment.

suppose (for tractability) that the non-verifiable income x is equal to

$$x = \mu(y, a) + \varepsilon,$$

where ε is a weakly positive random variable independent of a and y and $\mu(y, a)$ is a weakly positive function. The random variable ε plays essentially no role in our argument, but serves to emphasize that the borrower does not ever need to know the external state a . Let $g(\varepsilon)$ denote the distribution of ε .

Let $\omega(y'|y, x)$ denote a (possibly mixed) reporting strategy by the borrower, and let $\omega^*(y'|y, x)$ be the optimal reporting strategy. The indirect utility functions are

$$\phi_B(s, a) = \max_{\omega(y'|y, x) \in \Omega(s)} \int_0^\infty \int_0^\infty \int_0^\infty u(\mu(y, a) + \varepsilon + y - s(y, y')) \omega(y'|y, \mu(y, a) + \varepsilon) g(\varepsilon) f(y|a) dy' d\varepsilon dy,$$

$$\phi_L(s, a) = \beta_L(a) \int_0^\infty \int_0^\infty \int_0^\infty (s(y, y') - c(y'; s)) \omega^*(y'|y, (y, a) + \varepsilon) g(\varepsilon) f(y|a) dy' d\varepsilon dy,$$

where $\Omega(s)$ denotes the constraints in the reporting strategy, which we describe shortly. We impose limited liability, meaning that, for all reports y' , either $s(y_1, y') = s(y_2, y')$ for all y_1, y_2 and $0 \leq s(y', y') \leq y'$ (the non-verification case), or $0 \leq s(y, y') \leq y$ for all y (the verification case). We restrict the reporting strategies $\omega(y'|y, x)$ to place support only on y' for which the reports are feasible, meaning that if $s(y_1, y') = s(y_2, y')$ for all y_1, y_2 , then $\omega(y'|y, x) = 0$ if $y < s(y', y')$. In words, for reports that do not trigger verification, the borrower must have the funds to repay the loan.

Although this model is different from Townsend [1979] and Gale and Hellwig [1985], the arguments for the optimality of a debt contract are essentially unchanged. Fixing some distribution over external states $\pi(a)$, and integrating the indirect utility functions, it is immediately apparent

that the model is exactly that of Townsend [1979], except with random non-verifiable income. However, the argument for the optimality of debt depends only on non-satiation and that utility at zero verifiable income net debt repayments is not infinite. Therefore, assuming either that $u'(0) > -\infty$ or that $x > 0$ with probability one, debt will be optimal for all $\pi(a)$. It follows that our condition 1 is satisfied.

Specializing the indirect utility functions to a debt contract, which induces truthful reporting,

$$\begin{aligned}\phi_B(s_d, a) &= \int_d^\infty \int_0^\infty u(\mu(y, a) + \varepsilon + y - d)g(\varepsilon)f(y|a)d\varepsilon dy \\ &\quad + \int_0^d \int_0^\infty u(\mu(y, a) + \varepsilon)g(\varepsilon)f(y|a)d\varepsilon dy,\end{aligned}$$

and

$$\begin{aligned}\phi_L(s_d, a) &= \beta_L(a) \int_d^\infty df(y|a)dy \\ &\quad + \beta_L(a) \int_0^d (y - \bar{c})f(y|a)dy.\end{aligned}$$

We note that these indirect utility functions satisfy the assumption we have imposed. In particular, they are both differentiable (and hence continuous) in d and the lender's indirect utility function is zero when the level of debt is zero. The derivative of the borrower's indirect utility function with respect to the level of debt is

$$\phi_{B,d}(s_d, a) = - \int_d^\infty \int_0^\infty u'(\mu(y, a) + \varepsilon + y - d)g(\varepsilon)f(y|a)d\varepsilon dy,$$

and hence is strictly negative, satisfying our monotonicity requirement. The derivative of the

lender's indirect utility function is

$$\phi_{L,d}(s_d, a) = -\beta_L(a)f(d|a)\bar{c} + \beta_L(a) \int_d^\infty f(y|a)dy.$$

We now turn to the question of super-modularity for ϕ_B and sub-modularity for ϕ_L . In the main text, we have assume that a takes on values in a finite set, A . For this discussion, it will be easier to differentiate with respect to a , assuming that $\beta_L(a)$ and $\mu(y, a)$ are differentiable with respect to a . It is also convenient to observe, by the requirement that $\int_0^\infty f(y|a)dy = 1$ for all a , that

$$\psi'(a) = E[\ln(y)|a].$$

We first consider the super-modularity of ϕ_B , which is a necessary condition implied by condition

2. Taking the derivative with respect to a for $\phi_{B,d}$, we have

$$\begin{aligned} \phi_{B,da}(s_d, a) = & - \int_d^\infty \int_0^\infty u'(\mu(y, a) + \varepsilon + y - d)g(\varepsilon)[\ln(y) - E[\ln(\hat{y})|a]]f(y|a)d\varepsilon dy \\ & - \int_d^\infty \int_0^\infty u''(\mu(y, a) + \varepsilon + y - d)g(\varepsilon)\mu_a(y, a)f(y|a)d\varepsilon dy. \end{aligned}$$

Super-modularity requires that $\phi_{B,da} \geq 0$. Sufficient conditions for super-modularity are therefore $E[\ln(\hat{y})|a] \geq 0$ for all $a \in A$ and, for all $d, \varepsilon \in \mathbb{R}_{++}$, all $y > d$, and all $a \in A$,

$$-\frac{u''(\mu(y, a) + \varepsilon + y - d)}{u'(\mu(y, a) + \varepsilon + y - d)}\mu_a(y, a) \geq \ln(y).$$

This condition has an economic interpretation. The external state must impact the non-verifiable income so that the marginal utility of income is lower in good external states, despite the “direct” effect that better states lead to higher expected debt repayments. This condition is satisfied if, for

example, utility is CARA with coefficient α , $\mu_a(y, a) = \xi y$, and $\alpha\xi \geq 1$. One interpretation of the complementary role that the external state and verifiable income play in the non-verifiable income is that better external states imply better investment opportunities, or more persistent income processes (in a multi-period model).

Why do we need to augment the standard CSV model with this non-verifiable income? In the standard CSV model, if the level of debt is sufficiently high, the borrower's marginal utility can be very high, even in the best external states (because she is highly likely to default regardless of the state of the world). As a result, the direct effect of better states making default less likely overwhelms the risk-aversion of the borrower, and the borrower becomes risk-seeking with respect to the external state, contradicting our condition 2. This does not imply the non-existence of the non-contingent equilibrium; it only implies that our proposition 2 cannot be directly applied. In other words, our conditions are sufficient conditions, not necessary conditions.

The sub-modularity of ϕ_L is enforced through the marginal utility $\beta_L(a)$. For sufficiently dispersed $f(y|a)$, it will be the case that, for all d , $\phi_{L,d}(s_d, a) > 0$, and hence that decreasing marginal utility $\beta_L(a)$ promotes sub-modularity. Under this assumption,

$$\begin{aligned} \phi_{L,da}(s_d, a) &= \frac{\beta'_L(a)}{\beta_L(a)} \phi_{L,d}(s_d, a) + \beta_L(a) \int_d^\infty (\ln(y) - E[\ln(\hat{y})|a]) f(y|a) dy \\ &\quad - \beta_L(a) \bar{c}(\ln(d) - E[\ln(\hat{y})|a]) f(d|a). \end{aligned}$$

A sufficient condition for sub-modularity is that $\beta'_L(a)$ be sufficiently negative.

To establish condition 2, we require that the social welfare function $U(s_d, a; \lambda^*)$ be super modular. Note, however, that λ^* is endogenous to the problem, and is strictly increasing in K , with $\lambda^* = 0$ if $K = 0$. As a result, it will generally be the case that $U(s_d, a; \lambda^*)$ is super modular, for sufficiently small K (and hence λ^*), as long as ϕ_B is super modular. The exception to this is if the

tail behavior (as $d \rightarrow 0$ or $d \rightarrow \infty$) for ϕ_B has the cross partial converging to zero faster than the cross partial for ϕ_L . Using an exponentially-decaying distribution for $q(y)$ is sufficient to eliminate this issue.

We establish that our conditions hold in a numerical example, using Mathematica. The table below describes the parameters and functional forms. In this example, we assume that $\varepsilon = 0$.

Table 1: Parameters of CSV Example

Function/Parameter	Functional Form	Value
Utility $u(\cdot)$	CARA	3
Marginal Utility $\beta_L(a)$	$\beta_L(a) = \bar{\beta} a^{-\frac{1}{2}}$	$\bar{\beta} = \frac{\sqrt{2}}{\sqrt{2}+1}$
PDF $q(y)$	Exponential	$\frac{2}{3}$
Verification cost \bar{c}		$\frac{1}{10}$
N.V. Income Function $\mu(y, a)$	$\mu(y, a) = \xi ya$	$\xi = 1$
Required Funds K		1
External States A		$\{\frac{1}{4}, \frac{1}{2}\}$
Probabilities $p(a)$		$(\frac{1}{2}, \frac{1}{2})$

Notes: For these parameters and the model described in this section, we have verified using Mathematica that condition 2 is satisfied.

B Proofs

B.1 Proof of lemma 1

First, note that, for any values of θ for the which the lenders do not offer a security, profits are zero.

Proof by contradiction: suppose that there exists a symmetric pure-strategy equilibrium such that, for some values of $\theta \in \Theta$, the security $s(\theta)$ is offered and equilibrium lender profits are strictly positive.

Let θ' and $s' = s(\theta')$ denote the equilibrium type and security for which lender profits are

positive. In this equilibrium, each lender earns

$$|L|^{-1} \left(\sum_{a \in A, z \in Z} \theta'(a, z) \phi_L(s'_z, a) - K \right) > 0.$$

Let $d'(z)$ be the function satisfying $s'_z = s_{d'(z)}$ for all $z \in Z$. Consider a deviation by some lender to the security $s''_z = s_{d''(z)}$, where $d''(z) = \alpha d'(z)$ from some $\alpha \in (0, 1)$. By assumption, $s'' \in S$. By the monotonicity property of the borrower's indirect utility function, $\phi_B(s_d, a)$, in d , we have

$$\sum_{a \in A, z \in Z} \theta(a, z) \phi_B(s''_z, a) > \sum_{a \in A, z \in Z} \theta(a, z) \phi_B(s'_z, a)$$

for all $\theta \in \Theta$. It follows that, regardless of the beliefs the borrower forms off-equilibrium, she will accept security s'' if offered, for any value of $\alpha \in [0, 1)$.

The change in profits for the deviating lender is

$$\sum_{a \in A, z \in Z} \theta'(a, z) (\phi_L(s''_z, a) - |L|^{-1} \phi_L(s'_z, a)).$$

By the continuity of ϕ_L in d and the fact that $|L| > 1$, there exists an $\alpha \in (0, 1)$ such that this quantity is positive. It follows that an equilibrium with lender profits cannot exist.

B.2 Proof of proposition 1

By 1, this equilibrium delivers weakly positive utility for the borrower. Therefore, the borrower is willing to participate, and lenders earn zero profits (by the construction of $\bar{s}(\theta)$) and therefore are also willing to participate.

Now consider a deviation by a single lender: suppose some lender of type θ offers security s' instead of $\bar{s}(\theta)$, and would weakly profit from doing so if the security was accepted. Because the

lender can weakly profit from offering this deviation, the borrower is free to place the full support of her beliefs on the lender's true type. Because the security $\bar{s}(\theta)$ is on the Pareto-frontier, and offers zero profit to lenders, it follows that the borrower must be weakly worse off using security s' , and therefore would prefer the security $\bar{s}(\theta)$. Because there is more than one lender ($|L| > 1$), the borrower can choose the non-deviating lender and reject the deviating security. Given that the security will be rejected, the lender does not profit from offering it, and therefore $s(\theta) = \bar{s}(\theta)$ is an equilibrium.

B.3 Proof of proposition 2

The non-contingent security $s^* = s_{d^*}$ has payoffs that do not depend on the index. As a result, it offers zero profits for the lender, regardless of the lender's type, by the assumption that all $\theta \in \Theta$ have the same marginal distribution with respect to the external state. By assumption 1, s_{d^*} can deliver positive utility to the borrower, and therefore the participation constraints are satisfied in this equilibrium. It is sufficient to rule out deviations in which a single lender offers security s' instead of s_{d^*} , when the common type is θ' , to demonstrate that this is an equilibrium.

The security $s' \in S$, $s'_z = s_{d'(z)}$ with must offer profits for the lender of type θ' (if accepted) to break the equilibrium:

$$\sum_{a \in A, z \in Z} \theta'(a, z) \phi_L(s_{d'(z)}, a) > K.$$

Define $\bar{\theta}$ as a type that would profit the most from offering the security s' :

$$\bar{\theta} \in \arg \max_{\theta \in \Theta: \mu_0(\theta) > 0} \sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_{d'(z)}, a)$$

Consider the set of “elementary transformations” $t \in \mathcal{T}$ defined by Meyer and Strulovici [2015], and suppose that for some $z', z'' \in Z$ that are adjacent in the order on Z (with $z'' \succ z'$),

$d(z'') > d(z')$. By Meyer and Strulovici [2015], we can write

$$\bar{\theta} = \theta_0 + \sum_{t \in \mathcal{T}} \alpha_t t,$$

for some constants $\alpha_t \geq 0$. If there exists a $t \in \mathcal{T}$ with support on z' and z'' such that $\alpha_t > 0$, then by the richness of the type space (condition 3), there exists a type $\tilde{\theta} = \bar{\theta} - \beta t$, for some $\beta > 0$, such that $\tilde{\theta}$ is in the support of μ_0 . By the sub-modularity of ϕ_L (condition 2),

$$\sum_{a \in A, z \in Z} \tilde{\theta}(a, z) \phi_L(s_{d'(z)}, a) \geq \sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_L(s_{d'(z)}, a).$$

Therefore, it is without loss of generality to assume that the security $d'(z)$ is weakly decreasing between adjacent pairs z', z'' such that $\alpha_t > 0$ for some elementary transformation with support on those pairs.

By this result and the super-modularity of the social welfare function with Pareto-weight λ^* (condition 2), we must have

$$\sum_{a \in A, z \in Z} \bar{\theta}(a, z) U(s_{d'(z)}, a; \lambda^*) \leq \sum_{a \in A, z \in Z} \theta_0(a, z) U(s_{d'(z)}, a; \lambda^*).$$

By the Pareto-optimality of the non-contingent security s^* under θ_0 ,

$$\sum_{a \in A, z \in Z} \theta_0(a, z) U(s_{d'(z)}, a; \lambda^*) \leq \sum_{a \in A, z \in Z} \theta_0(a, z) U(s^*, a; \lambda^*).$$

Therefore,

$$\begin{aligned} \sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_B(s_{d'(z)}, a) + \lambda^* \sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_L(s_{d'(z)}, a) \leq \\ \sum_{a \in A, z \in Z} \theta_0(a, z) \phi_B(s^*, a) + \lambda^* \sum_{a \in A, z \in Z} \theta_0(a, z) \phi_L(s^*, a). \end{aligned}$$

It follows that

$$\sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_B(s_{d'(z)}, a) < \sum_{a \in A, z \in Z} \theta_0(a, z) \phi_B(s^*, a),$$

and by the non-contingency of s^* ,

$$\sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_B(s_{d'(z)}, a) < \sum_{a \in A, z \in Z} \bar{\theta}(a, z) \phi_B(s^*, a).$$

By the $D1$ refinement, the borrower can place the support of her beliefs entirely on the type $\bar{\theta}$. Consequently, if there exists a θ' for which the deviation is profitable, the borrower can believe she is worse off and reject the deviation.

B.4 Proof of (2)

The text demonstrates that it is without loss of generality to consider contracts with a constant promise, and with default only due to an inability to pay. The utility function is

$$U(s_d, a; \lambda) = \frac{\bar{e}}{2} - (\lambda - \beta_L(a)^{-1}) \phi_L(s_d, a) - \frac{d}{\bar{e}} (C(a) - L),$$

with

$$\phi_L(s_d, a) = \beta_L(a) \left(1 - \frac{d}{\bar{e}}\right) d + \beta_L(a) \frac{d}{\bar{e}} \min\{d, L\}.$$

The borrower's utility function is maximized at $d = 0$. The lender's utility function is differentiable almost everywhere with respect to d , with

$$\frac{\partial}{\partial d} \phi_L(s_d, a) = \begin{cases} \beta_L(a) & d < L \\ \beta_L(a) \left(\frac{\bar{e} + L - 2d}{\bar{e}} \right) & d > L. \end{cases}$$

It follows that the function $\phi_L(s_d, a)$ is maximized at

$$\bar{d} = \frac{\bar{e} + L}{2} > L,$$

irrespective of the distribution of states.

At the maximum debt value \bar{d} ,

$$\begin{aligned} \sum_{a \in A} p(a) \phi_L(s_{\bar{d}}, a) &= \sum_{a \in A} p(a) \beta_L(a) \left\{ \left(1 - \frac{\bar{d}}{\bar{e}}\right) \bar{d} + \frac{\bar{d}}{\bar{e}} L \right\} \\ &= \frac{\bar{e} + L}{2} \left(1 - \frac{\bar{e} + L}{2\bar{e}}\right) + \frac{\bar{e} + L}{2\bar{e}} L \\ &= \frac{\frac{1}{2}(\bar{e} + L)(\bar{e} - L) + L(\bar{e} + L)}{2\bar{e}} \\ &= \frac{(\bar{e} + L)^2}{4\bar{e}}, \end{aligned}$$

and therefore trade is feasible.

B.5 Proof of lemma 3

First, we show that $\phi_L(s_d, a)$ is sub-modular. This function is differentiable with respect to $d \in D$ almost everywhere, with

$$\frac{\partial}{\partial d} \phi_L(s_d, a) = \begin{cases} \beta_L(a) & d < L \\ \beta_L(a) \left(\frac{\bar{e} + L - 2d}{\bar{e}} \right) & d > L. \end{cases}$$

By assumption 4, $\beta_L(a)$ is weakly decreasing, by the definition of D , $2d \leq \bar{e} + L$, and therefore ϕ_L is sub-modular.

Differentiating $U(s_d, a; \lambda^*)$ with respect to d ,

$$\frac{\partial}{\partial d} U(s_d, a; \lambda^*) = \frac{\lambda^* \beta_L(a) - 1}{\beta_L(a)} \frac{\partial}{\partial d} \phi_L(s_d, a) - \frac{C(a) - L}{\bar{e}}.$$

For any $d < L$, for $U(s_d, a; \lambda^*)$ to be super modular, we must have

$$\lambda^* \beta_L(a) \bar{e} + L - C(a) - \bar{e}$$

increasing in a . For $d > L$, the condition is

$$\lambda^* \beta_L(a) (\bar{e} + L - 2d) - C(a) - \bar{e} + 2d$$

increasing in a . By the definition of D , $\bar{e} > (\bar{e} + L - 2d) > 0$, and by the fact that $\beta_L(a)$ is decreasing, it follows in both cases that

$$\lambda^* \beta_L(a) \bar{e} - C(a)$$

increasing is sufficient.

At the non-contingent optimal security,

$$\sum_{a \in A} p(a) \frac{\lambda^* \beta_L(a) - 1}{\beta_L(a)} \frac{\partial}{\partial d} \phi_L(s_d, a) \Big|_{d=d^*} = \sum_{a \in A} p(a) \frac{C(a) - L}{\bar{e}}.$$

By assumption,

$$\sum_{a \in A} p(a) \phi_L(s_{d^*}, a) = K,$$

and $\phi_L(s_{d^*}, a) \leq \beta_L(a) d^*$, and therefore $d^* > K > L$. It follows that

$$\lambda^* (\bar{e} + L - 2d^*) = \bar{e} - 2d^* + \sum_{a \in A} p(a) C(a)$$

and that

$$K = \left(1 - \frac{d^*}{\bar{e}}\right) d^* + \frac{d^*}{\bar{e}} L,$$

or

$$(d^*)^2 - (\bar{e} + L) d^* + K \bar{e} = 0.$$

Solving,

$$d^* = \frac{(\bar{e} + L) - \sqrt{(\bar{e} + L)^2 - 4K\bar{e}}}{2},$$

noting that the sign is pinned down by $d^* \leq \bar{d}$. Therefore,

$$\lambda^* = 1 + \frac{\sum_{a \in A} p(a) (C(a) - L)}{\sqrt{(\bar{e} + L)^2 - 4K\bar{e}}},$$

and therefore

$$\lambda^* \leq \frac{\bar{e} + \sum_{a \in A} p(a) C(a)}{\bar{e} + L}.$$

It follows that assumption 4 is sufficient.

B.6 Proof of lemma 4

By definition, the set Θ is bounded, by assumption it is closed, and therefore it is compact. It follows that $\Theta^*(\theta)$ is non-empty. By the linearity of

$$\sum_{a \in A, z \in Z} \theta''(a, z) \phi_L(\bar{s}_z(\theta), a)$$

in θ'' , $\Theta^*(\theta)$ is convex.

By condition 1, $\bar{s}_z(\theta) = s_{d(z, \theta)}$ for some $d(z, \theta) \in D$. By definition,

$$\bar{s}(\theta) = \max_{\{d(z) \in D\}_{z \in Z}} \sum_{a \in A, z \in Z} \theta(a, z) \phi_B(s_{d(z)}, a)$$

subject to

$$\sum_{a \in A, z \in Z} \theta(a, z) \phi_L(s_{d(z)}, a) \geq K.$$

By assumption, $\phi_L(s_d, a)$ and $\phi_B(s_d, a)$ are continuous in d , and hence it follows that $d(z, \theta)$ is continuous in θ and that

$$\sum_{a \in A, z \in Z} \theta''(a, z) \phi_L(\bar{s}_z(\theta), a)$$

is jointly continuous in (θ, θ'') . Therefore by Berge's theorem (the theorem of the maximum), $\Theta^*(\theta)$ is upper semi-continuous.

It follows that Kakutani's fixed point theorem holds, and therefore that there exists a θ^* such that

$$\theta^* \in \Theta^*(\theta^*),$$

as claimed.

Now suppose there is another $\theta' \in \Theta^*(\theta^*)$. We must have

$$\sum_{a \in A, z \in Z} \theta'(a, z) \phi_L(s_{d(z, \theta)}, a) = K.$$

Because the index and the external state are perfectly correlated, we can rewrite this as

$$\sum_{a \in A, z \in Z} q(z; \theta') \delta(a, z) \phi_L(s_{d(z, \theta)}, a) = K,$$

where $q(z; \theta')$ is the marginal distribution associated with θ' .

By the definition of $\bar{s}(\theta)$, for all $z \in Z$ (by the full support assumption, $q(z; \theta^*) > 0$),

$$\phi_B(s_{d(z; \theta^*)}, a) + \lambda \phi_B(s_{d(z; \theta')}, a) \geq \phi_B(s_{d'}, a) + \lambda \phi_B(s_{d'}, a)$$

for all $d' \in D$ and some multiplier $\lambda > 0$. It follows that the optimality would also hold if $d(z; \theta') = d(z; \theta^*)$, and feasibility is satisfied, and therefore

$$\bar{s}(\theta') = \bar{s}(\theta^*),$$

as required.