LIQUIDITY CONSTRAINTS AND INCENTIVE CONTRACTS

ANDREAS LEHNERT
Board of Governors of the Federal Reserve System

ETHAN LIGON
University of California at Berkeley

ROBERT M. TOWNSEND
University of Chicago
and
Federal Reserve Bank of Chicago

Are firms and households constrained in the use of a productive input? Theoretical approaches to this question range from exogenously imposed credit allocation rules to endogenous market failures stemming from some sort of limited-commitment or moral-hazard problem. However, when testing for constraints, researchers often simply ask firms or households if they would wish to borrow more at the current interest rate and/or test for suboptimal use of inputs in production functions relative to a full-information, full-commitment benchmark. We demonstrate that if credit is part of a much larger information-constrained (or limited-commitment) incentive scheme, then input use may very well be distorted away from the first-best. Further, households and firms, in certain well-defined circumstances, may, at the true interest rate or opportunity cost of credit, desire to borrow more (or less) than the assigned level of credit. In other, more constrained, contractual regimes, firms and households would say that they do not want to borrow more (or less), but these regimes are decidedly suboptimal, although the magnitude of the loss does depend on parameter values. We conclude with empirical methods that, in principle, could allow researchers armed with enough data to estimate parameters and distinguish regimes. Researchers then could see if firms and households are truly constrained and, if so, what the welfare loss might be.

Keywords: Credit Constraints, Moral Hazard, Optimal Contracts

1. INTRODUCTION

Are some firms and households suboptimally constrained in the use of credit to finance productive inputs? Empirically, one might follow the literature in testing this

proposition by estimating a production function and determining if the marginal product of inputs financed via credit was too high (or too low) or influenced by the proprietor's own resources. If designing a survey, one might simply inquire of households (or firms) if they would care to borrow more at the current rate. We demonstrate that optimal information-constrained financial contracts (under a variety of incentive constraints) have the feature that borrowers are assigned, or induced to choose, inputs that would not be optimal by the usual first-best production efficiency standard. We show that in certain circumstances borrowers indeed will claim to be credit constrained.

Examples of the standard approach include Benjamin (1992), who tests for the strong separation of production from a proprietor's resources with data from farm households in Java, though he fails to reject no constraints. Evans and Jovanovic (1989) in the United States specify a credit allocation rule (households may borrow up to a certain multiple of wealth) allowing them to test for constrained entry into business and the sufficiency of credit-financed inputs in business. Feder et al. (1990) estimate a more complex but ad hoc credit-constraint equation from survey data in rural China. Households were asked about their access to credit, and their responses used as dependent variables in a logistic regression to investigate which household characteristics are associated with limited access to credit. This then is used in a switching regression to determine the welfare and production loss from limited credit.

The problem with this literature is that, apart from the benchmark of full efficiency, it lacks an alternative model of how the credit market is supposed to work that is both explicit and compelling. At best, models posit credit allocation rules for formal and informal credit institutions that are crude approximations to reality. At worst, the credit allocation rule is some mongrel empirical relationship, conflating the demand and supply of credit, neither of which is explicitly modeled. Simple palliative suggestions about raising interest rates or reducing collateral requirements ignore the sophisticated resource allocation and incentive problem that the credit supplier must solve and to which the borrower will respond. It is thus difficult to determine empirically if Pareto improvements are really possible, or if the world is efficient.

There are exceptions. The model of Stiglitz and Weiss (1981) and the subsequent literature, including Aghion and Bolton (1992) and Piketty (1994), are explicit about potential difficulties that incentive problems pose for credit markets. A poor household may have little incentive to be diligent and thus may have a large probability of default. Raising the interest rate here only worsens the situation; the borrower is already defaulting in low-output states and works hard only to achieve output net of debt repayment in the high-output states. In equilibrium, households appear to be constrained.

This approach also has its problems. The literature may use special technologies and preferences (for example, two outputs and linear utility) or, a priori, impose the form of the financial instrument. Results derived from models with two states, risk neutrality, or exogenously imposed contracts usually do not hold in more general

settings. The market often is also somewhat stylized, with the supply and demand for credit stemming from an exogenously imposed class difference among agents or some strategic form of presumed competition among lenders.

Our goal is to dispense with special technologies and preferences, an imposed form of the credit contract, and arbitrary notions of the supply and demand for credit. We model credit markets simply as institutions solving an optimal credit-provision problem in the presence of various incentive and information difficulties, with risk-averse borrowers and arbitrary technologies.

The literature may be correct in its implicit assertion that actual contracts are more limited than the information-constrained optimum might suggest. One might verify this assertion by looking at loan contracts directly. Müller and Townsend (1997) find contingencies in supposedly standard sharecropping contracts by asking if there are exceptions in a bad year. Rashid and Townsend (1993) point out that provision for loan rollover and other methods for treating overdue obligations in effect introduce contingencies into apparently standard contracts. This is complementary with the literature on incomplete markets with default [see Dubey et al. (1987), Duffie (1996)]. Our approach here is to look at the allocation of resources directly and turn the above assertion into an empirical question. That is, in addition to the information-constrained optima, we impose a priori the form of the loan contract and calculate the artificially constrained optimal allocations that would result. If these latter allocations are more consistent with the data than those produced by a less-constrained regime, then there is an obvious policy recommendation: See if it is possible to move toward the more complete information-constrained optimum.

To these ends, we calculate optimal dynamic contracts in principal-agent models of the type studied by Phelan and Townsend (1991) under four regimes with productive credit. This productive material input is available from the outside world at a fixed price. These regimes differ by the information and control available to the principal (or lender), ranging from the benchmark full-information, full-commitment regime, in which effort and loan size are costlessly observed and enforced, to one in which effort is not observed but loan size is controlled (the moral-hazard, credit-control regime), to two in which effort is not observed and credit transactions with the outside world are not controlled with varying bankruptcy possibilities (the moral-hazard, no-credit-control regimes), to one in which consumption gross of loan repayment must equal output (the pure debt regime). This last regime is the usual fixed-obligation debt contract.

Default is not possible in the full-information and moral-hazard, credit-control regimes because the principal controls borrowing and net consumption directly. In the moral-hazard, no-credit-control regimes, however, the principal only controls gross consumption. If the gross consumption assignment is too low, the agent may be unable to repay. We solve this regime both with and without default.

We compare consumption, insurance, effort, credit-financed inputs, and internal interest rates (where applicable) among the four regimes. A borrower's response to a survey question that asked "Are you credit constrained?" would depend on the

underlying regime. A survey question that implicitly assumed the household or firm faced a particular interest rate as in a pure-debt contract (and as in the supply-and-demand approach), while in fact it was involved in a dynamic incentive-constrained insurance scheme, might elicit a misleading response. Even if the survey correctly identified the regime, it could mistakenly assume that the firm or household faced the world interest rate when, in fact, it faced a different, purely internal, rate.

In the full-information regime (in which effort and material input are fully observed, controlled, and contractible by the principal), the principal and the agent agree on the level of borrowing and effort, with the agent having unlimited access to credit at the correct opportunity cost. In this regime there is no barrier to the typical equity contract in which, in effect, the principal buys the enterprise (i.e., the equity) and a certain level of effort from the agent for a lump sum, providing full insurance. This regime assumes legal institutions that can force the agent to carry out the promised effort, otherwise the agent would deviate in effort. Over regions of wealth (or promised utility of the agent) in which effort does not vary, loan size is independent of wealth. This is simply the usual neoclassical separation of production decisions and household consumption. Because (as we explain in Section 4) we take the household's (or firm owner's) labor to be tied to the technology, we don't recover this separation across all promised utility ranges.[1]

In the moral-hazard, credit-control regime (where agents are assigned a loan size but must be induced to work hard), the household or firm generally would choose, if asked, a level of borrowing different than the one it would choose if given access to credit at the outside rate (the rate faced by the lender). The lender could ration the household's or firm's access to credit to ease his signal extraction program. Alternatively, the lender could force the household or firm to accept more credit than it desires for the same reason. We thus derive measures of which firms or households, given a set of technology and preference parameters, would appear to be constrained, and whether or not there is a pattern by wealth or utility class.

In the moral-hazard, no-credit-control regimes, we do allow the agent to choose the level of borrowing (at the world rate) while continuing to provide insurance. The principal alters the contract with this in mind, inducing new levels of borrowing and effort that may or may not coincide with the full-information levels. The agent always would appear unconstrained in these regimes, yet they are Pareto inferior to the moral-hazard, credit-control regime. We vary our treatment of default in these regimes, showing that ruling out default is equivalent to an extra set of constraints on the programming problem.[2] Further, the no-default constraints then interact with the incentive-compatibility constraint to make it quite difficult to assign high levels of borrowing. Indeed, it is this effect that can lead borrowers in the moral-hazard, *credit-control* regime to complain that they were being forced to accept too much credit.

In our pure-debt regime, we imagine that the principal controls the agent's access to the world credit market. The principal induces borrowing and effort with an internal on-lending rate (which may be negative) and a schedule of next period's

promised utilities conditional on output. The principal does not smooth contemporaneous consumption risk beyond bankruptcy. At the internal interest rate, the agent can choose any level of borrowing and, as the equity owner of the enterprise, is the residual claimant after debt repayment. Despite these well-defined property rights, the pure-debt regime can be quite suboptimal. To compensate the borrower for the large amount of contemporaneous risk he faces, the principal has to set the internal interest rate to a very low level, distorting the agent's input choices away from both the full-information optimum and the information-constrained optimum. Still, outcomes with this pure-debt regime with default can Pareto dominate outcomes with the moral-hazard, no-credit-control regime without default.

To review, we consider optimal dynamic lending contracts in a principal-agent model with essentially four different sets of constraints, called regimes. In the first regime, the full-information regime, the principal is unconstrained by any incentive problems. In the second regime, the moral-hazard, credit-control regime, the principal must only induce effort but can control borrowing. There are two variants on the third regime, the moral-hazard, no-credit-control regime, depending on how we handle default. In both cases the principal must induce both effort and borrowing. In one case the agent may default on the loan; in the other, he may not. In the fourth regime, the pure-debt regime, the principal is restricted to using the normal debt contract with default. In Proposition 1 we show which regimes are exactly nested in others, and which are not.

So far, we have compared regimes by asking the agents in each regime if they are credit constrained, or by comparing input choices across regimes, but there are empirical implications to each regime that go well beyond this. In all but the full-information regimes, the agent is rewarded or penalized on the basis of output not only in the present but in the future as well. These intertemporal tie-ins create transition probabilities across promised utilities that vary by regime. For a given set of parameters, it is thus possible to determine the invariant distribution over observable variables (such as output, consumption, and credit) and to compare them across regimes. As a step in this direction, we pick a fairly general model and vary its technology and risk-aversion parameters. We then posit an empirical researcher with access to as much output, consumption, and credit data as required, and determine how well this researcher would do in discriminating across regimes and estimating the particular model parameters within a regime. In this way, a researcher could determine if households or firms are truly constrained.

In Section 2, we outline the preferences, endowment, and technology in the model that we study; we develop notation that will be used throughout the paper; and we outline the rules under which the principal and agent trade. In Section 3, we consider in some detail the four main regimes that we study. In Section 4, we clarify the dynamics of the model. In Section 5, we Pareto rank the models. In Section 6, we provide results about information, the effect of default and credit constraints, as well as several numerical examples of interest. We conclude with Section 7.

2. PREFERENCES, ENDOWMENTS, TECHNOLOGY AND TRADING

The canonical principal-agent problem is easily understood as a single principal or lender contracting with a single agent. Yet it could easily generalize to a single lender dealing separately with several agents, each of whom cannot communicate or trade with the others. The set of contracts will still be optimal because we are assuming no aggregate shocks or other physical tie-ins among the agents. A third interpretation, and one that holds great interest for us, is an environment with a continuum of such agents. In this view, what is a probability number for a single agent takes the form of a fraction of agents in the continuum; the principal is used as an artifact of the programming approach. In fact, when the expected profit of the notional principal is set to zero, we have detailed an optimum for the entire set of agents that satisfies present-value budget balance. All of these formulations are computationally identical.

Agents are infinitely lived and have risk-averse preferences over net consumption c_1, effort a, and promised utility next period w' of the form

$$U(c_1, a) = u(c_1) + v(1 - a) + \beta w', \tag{1}$$

where c_1, a, and w' must be in the bounded sets C_1, A, W',[3] and additionally $\max\{A\} \leq 1$. The discount factor β satisfies $0 < \beta < 1$, and the functions u and v satisfy u', $v' > 0$ and u'', $v'' < 0$. Agents will have a satiation or bliss point in consumption of $c_{1_{\max}}$. Net consumption assignments above this level will not increase the agent's utility. In some regimes the principal will only be able to assign the agent a level of gross consumption c_0 in C_0, from which the agent must repay some amount x if he is able. We explicitly define how x must depend on the interest rate and borrowing amount in the appropriate sections. From (1), we see that the agent has preferences over gross consumption assignments c_0 in C_0, effort a in A, and promised utility w in W' of

$$U(c_0, a) = u([c_0 - x]_+) + v(1 - a) + \beta w', \tag{2}$$

where, in effect allowing bankruptcy,

$$[c_0 - x]_+ = \begin{cases} 0 & c_0 < x \\ c_0 - x & c_0 \geq x. \end{cases}$$

The principal is assumed to be risk neutral with respect to consumption, has the same discount factor β as the agent, and has no leisure. The principal has access to a world market for interperiod credit at the constant gross interest rate of $1/\beta$. Although our focus is on within-period credit, we discuss this issue further in Section 4.[4]

Neither the agent nor the principal is endowed with any stocks of the material input at the beginning of each period. Instead, they rely on purchases from the outside world, at a constant price ρ. The principal and the agent together therefore form a small open economy. Depending on the particular regime, the principal acts

as an intermediary between the agent and the outside market, although in certain regimes the agent could just as easily be dealing directly with the outside market. In all cases, however, the outside lender is never defaulted upon. That is, if the agent defaults, the principal pays the remainder, thus keeping the price at ρ.

The agent is endowed with a unit of time each period that he may split between effort a in A and leisure $1 - a$. The credit-financed material input, b, must lie in the set $\boldsymbol{B}$, which contains 0 as an element.

The agent has access to a technology mapping inputs of effort and material (a, b) in $\boldsymbol{A} \times \boldsymbol{B}$ into a probability distribution over the set of allowed outputs $\boldsymbol{Q}$, called $P(q \mid a, b)$. One or both of the inputs may be private and/or controlled by the agent, but output q in $\boldsymbol{Q}$ is public and may be taken at zero cost by the principal. The material input b is completely consumed in the productive process.

Although one might imagine that, in the real world, objects such as effort, output, and material input can take on a continuum of values, for the purposes of the numerical application here, these ranges are approximated by discrete sets. For example, the allowed set of gross consumption allocations $\boldsymbol{C_0}$ is a vector of (nonnegative) real numbers of n_{C_0} elements.

Given sets of values for effort $\boldsymbol{A}$ of dimension n_A, credit-financed material input $\boldsymbol{B}$ of dimension n_B, gross consumption $\boldsymbol{C_0}$ of dimension n_{C_0}, output $\boldsymbol{Q}$ of dimension n_Q, and the internal price of credit $\boldsymbol{R}$ of dimension n_R, the set of net consumptions $\boldsymbol{C_1}$ is determined. The maximum and minimum possible values of promised utility this period, $\boldsymbol{W}$, and next period, $\boldsymbol{W'}$, which may vary across regimes, also are determined. The dimensions of $\boldsymbol{W}$ and $\boldsymbol{W'}$ are on a continuum but we again approximate them with a grid.

Although the discussion that follows is valid for all specifications of sets of effort $\boldsymbol{A}$, material $\boldsymbol{B}$, output $\boldsymbol{Q}$, and so on, in our numerical work we chose a particular formulation and varied technology and preference parameters to investigate various interesting effects.

With this notation in place, we briefly review here the treatment of consumption and default across the four regimes to clarify the commodity space. We can dispense with the distinction between net consumption and gross consumption in the full-information and moral-hazard, credit-control regimes because, in those regimes, the principal controls the borrowing amount b. The principal then can assign net consumption c_1 as a function of output q. It is not necessary (though equivalent) to assign higher gross consumption $c_0 = c_1 + \rho b$ only to take it away with the debt repayment ρb.

In the moral-hazard, no-credit-control regime without default, the principal is, in essence, selling the agent an insurance policy in gross consumption c_0 as a function of output q. This induces an amount of borrowing b. Because the agent's borrowing is private and not controlled, it is as if the agent were dealing directly with an outside loan market at the world rate of ρ. In equilibrium, this credit also could have come from the principal under a contract separate from the insurance arrangement. To prevent the agent from defaulting, in this version of the no-credit-control regime, we require that the gross consumption allocation c_0 never fall below

$\rho\bar{b}$, where $\bar{b}$ is the assigned loan size. The market then is assured of repayment and does not need to charge a risk premium.

In the version of this no-credit-control regime when we do allow default, the outside market would no longer be willing to lend to the agent at the fixed rate ρ because the expected rate of return on loans would be below ρ. We could have used an iterative process to determine a risk premium over ρ charged to agents at which the default rate just guaranteed the lenders an expected rate of return of ρ. But this adds to computational complexity. Instead, we have the principal guarantee the loan, so that, in effect, the agent is getting subsidized credit.

In the pure-debt regime, gross consumption allocations c_0 are required to be equal to output q. The principal acts as an exclusive intermediary between the outside market and the agent, on-lending to the agent at an endogenous internal interest rate of r. Although the principal can observe borrowing b, we imagine that he does not try to control it directly. The agent defaults when $q - rb < 0$, consuming zero, and the principal takes q (therefore, the cost of the default is borne by the principal, not the outsider).

3. REGIMES

We now consider the four regimes in detail.

3.1. Full-Information Regime

In this regime, the principal (lender) observes the agent's (borrower's) effort a in A and can control effort as well as the agent's borrowing choice b in B. Given a set of net consumption allocations, C_1, the relevant principal-agent problem is familiar. Here we write it down as a simple extension of the program considered by Phelan and Townsend (1991) in order to compare it with the nested regimes that follow.

The principal's objective function, $V(w_0)$, given an initial promised utility level to the agent of w_0, is the expected value of output net of materials cost, ρb, and payments to the agent, c_1, plus the discounted value of next period's value function at the continuation utility, $V(w')$. The expectation is taken with respect to the joint probability of outcomes (c_1, q, a, b, w') in $C_1 \times Q \times A \times B \times W'$, namely $\pi(c_1, q, a, b, w')$. It is these joint probabilities π that form the principal's choice objects. So the objective function becomes

$$V(w_0) = \max_{\pi} \left\{ \sum_{C_1 \times Q \times A \times B \times W'} \pi(c_1, q, a, b, w')[q - c_1 - \rho b + \beta V(w')] \right\}, \quad \textbf{(3)}$$

where, again, ρ is the exogenous outside cost to the principal of obtaining the material input b.

The maximization in (3) proceeds subject to a set of constraints on the contract π. First, we have the promise-keeping constraint, which requires the agent's expected

utility to equal w_0. This may be written as

$$\sum_{C_1 \times Q \times A \times B \times W'} \pi(c_1, q, a, b, w')[u(c_1) + v(1 - a) + \beta w'] = w_0, \qquad (4)$$

where $u(\cdot)$, $v(\cdot)$, and β are as defined in equation (1).

Second, the contract π must, for each value of $(\bar{a}, \bar{b})$ in $A \times B$, be Bayes consistent with the exogenously specified probability distribution over outputs $P(q \mid \bar{a}, \bar{b})$. For each $(\bar{q}, \bar{a}, \bar{b})$ in $Q \times A \times B$,

$$\sum_{C_1 \times W'} \pi(c_1, \bar{q}, \bar{a}, \bar{b}, w') = P(\bar{q} \mid \bar{a}, \bar{b}) \sum_{C_1 \times Q \times W'} \pi(c_1, q, \bar{a}, \bar{b}, w'). \qquad (5)$$

Finally, the contract must form a valid probability mass function over the points $C_1 \times Q \times A \times B \times W'$. Thus, for each point, the contract must satisfy $\pi \geq 0$, and over all the points

$$\sum_{C_1 \times Q \times A \times B \times W'} \pi(c_1, q, a, b, w') = 1. \qquad (6)$$

The region of possible promised utilities ranges from the lowest (net) consumption $[c_{1_{\min}} = \min(C_1)]$ and highest effort $[a_{\max} = \max(A)]$ promised with certainty forever at the low end, to the highest consumption $[c_{1_{\max}} = \max(C_1)]$ and lowest effort $[a_{\min} = \min(A)]$ promised with certainty forever at the high end. Thus, the minimum $(w_{\min})$ and maximum $(w_{\max})$ values of promised utility (W') are defined by the equations

$$w_{\min} = \frac{u(c_{1_{\min}}) + v(1 - a_{\max})}{1 - \beta}, \qquad (7)$$

$$w_{\max} = \frac{u(c_{1_{\max}}) + v(1 - a_{\min})}{1 - \beta}. \qquad (8)$$

Note that the only way to achieve these promised utilities is by actually assigning $(c_{1_{\min}}, a_{\max})$ or $(c_{1_{\max}}, a_{\min})$ forever; hence these utilities become absorbing states.

We now are ready to describe the optimization program solved by the principal in this regime.

PROGRAM 1 (Full-information). *Maximize the principal's objective function (3) by choice of contract weights $\pi(c_1, q, a, b, w')$ over points in $(C_1 \times Q \times A \times B \times W')$ satisfying constraints (4), (5), and (6). The range of promised utilities W' is defined by equations (7) and (8).*

Because we approximate the continuous variables C_1 (net consumption) and W' (promised utility) with discrete grids, we cannot solve analytically for the optimal solution in this regime. Because the principal does not have to induce effort or borrowing, we know that the optimal solution will feature full-consumption insurance and no intertemporal tie-ins. It is thus possible to solve a much simpler problem than

(1) to arrive at the solution. However, we solve the program numerically to gauge the distortions introduced by using grids for consumption and promised utility.

3.2. Moral-Hazard, Credit-Control Regime

A natural first step away from the full-information regime is one in which effort is unobserved. In this case, the principal must assign efforts a that satisfy an incentive compatibility constraint. The principal's objective function is unchanged, and the optimal contract in this regime also must satisfy the constraints (4), (5), and (6).

The incentive compatibility constraint in this regime says that, given an assigned effort $\bar{a}$ and known assigned and enforced loan level $\bar{b}$, the agent must not be able to achieve a higher expected utility from any deviation to some other effort $\hat{a}$. In the notation used here, this requires $\forall\,(\bar{a}, \bar{b}, \hat{a}) \in A \times B \times A$,

$$\sum_{C_1 \times Q \times W'} \pi(c_1, q, \bar{a}, \bar{b}, w')[u(c_1) + v(1 - \bar{a}) + \beta w']$$

$$\geq \sum_{C_1 \times Q \times W'} \pi(c_1, q, \bar{a}, \bar{b}, w') \frac{P(q \mid \hat{a}, \bar{b})}{P(q \mid \bar{a}, \bar{b})}[u(c_1) + v(1 - \hat{a}) + \beta w']. \qquad (9)$$

Generally speaking, this equation simply requires the expected utility from undertaking the assigned effort to be weakly greater than the expected utility from any deviation. For an exact derivation, see Prescott and Townsend (1984a, b); for the application in these sorts of programming problems, see Phelan and Townsend (1991). The term $P(q \mid \hat{a}, \bar{b})/P(q \mid \bar{a}, \bar{b})$ appears because if the agent deviates, he induces a new probability distribution over outputs; hence the joint probability π must be corrected. This term is the likelihood ratio associated with that particular output, i.e., the ratio of probability of q given the agent deviates in effort to $\hat{a}$ to the probability of q given the agent supplies the assigned effort $\bar{a}$. If the likelihood ratio is significantly above or below unity, the principal can infer quite a bit ex post from that output, and hence either heavily punish the agent at that output (if the ratio is above unity) or reward the agent at that output (if the ratio is below unity). Note that this ratio is affected by the assigned level of credit, $\bar{b}$. We consider the effect of this in much greater detail below.

Clearly, the contract that defined the lowest allowed promised utility under the full-information regimes, given in equation (7), is not now incentive compatible. Because the principal here cannot directly observe or control the agent's effort, an agent assigned $c_{1_{\min}}$ with certainty forever would always select the lowest effort $a_{\min}$. The new range of W' is given by the equations

$$w_{\min} = \frac{u(c_{1_{\min}}) + v(1 - a_{\min})}{1 - \beta}, \qquad (10)$$

$$w_{\max} = \frac{u(c_{1_{\max}}) + v(1 - a_{\min})}{1 - \beta}, \qquad (11)$$

where the new value for $w_{\min}$ now corresponds to the incentive-compatible contract delivering the least amount of utility. For a proof of this point, see Phelan and Townsend (1991). Loosely speaking, any contract that delivered a promised utility below $w_{\min}$ as defined in equation (10), e.g., one that promised a little extra consumption in return for a large enough increase in effort to drive utility below $w_{\min}$, would not be incentive compatible. The agent considers deviating to $a_{\min}$, realizing that the worst punishment the principal can deliver is $c_{1_{\min}}$ with certainty. But, by construction, the utility from this deviation is greater than the utility of the contract. Therefore, in any period, no contract can be considered incentive compatible if it delivers a one-period utility below $u(c_{1_{\min}}) + v(1 - a_{\min})$. Then, the principal cannot threaten future utilities below $w_{\min}$.

We can now write the program defining optimal contracts in the control regime.

PROGRAM 2 (Moral-hazard, credit-control). *Maximize the principal's objective* (3) *by choice of nonnegative contract weights* $\pi(c_1, q, a, b, w')$ *over points in* $(C_1 \times Q \times A \times B \times W')$ *satisfying constraints* (4), (5), (6), *and* (9). *The range of promised utilities* W' *is defined by equations* (10) *and* (11).

3.3. Two Moral-Hazard, No-Credit-Control Regimes

In this section we discuss the case in which the principal has a further incentive compatibility constraint to satisfy, one on borrowing. It turns out that the effect depends crucially on our treatment of default. When default is prohibited, the principal can be seen as insuring the agent in gross consumption allocations, taking as given the agent's choice problem over borrowing. It is as if the principal assigns an incentive compatible level of borrowing, $\bar{b}$, and the agent never has a gross consumption allocation too low to repay a lender at rate ρ for $\bar{b}$.

When default is allowed, the principal is still selling the agent an insurance policy in gross consumptions, but now the principal is also guaranteeing, to the outside lender, that if the agent defaults he will pay off the entire loan. Because the principal is now paying for the agent's defaults, he presumably can observe, but not control, the agent's borrowing.

Our presentation assumes first that default is allowed. We then show that prohibiting default simply requires a set of extra constraints.

With this in mind, the principal's objective function becomes, for a typical borrower with promised utility w_0,

$$V(w_0) = \max_{\pi} \left\{ \sum_{C_0 \times Q \times A \times B \times W'} \pi(c_0, q, a, b, w') \right.$$

$$\left. \times [q - [c_0 - \rho b]_+ - \rho b + \beta V(w')] \right\}, \quad \textbf{(12)}$$

where C_0 is now a gross consumption allocation, so that net consumption $c_1 = [c_0 - \rho b]_+$. Notice that if the agent does not default, when gross consumption c_0

is greater than the repayment ρb, the principal's objective is $q - c_0$. If the agent does default, when $c_0 < \rho b$, the principal's objective is $q - \rho b$. When the agent defaults, he delivers the entire amount c_0 (which he got from the principal) to the outside lender, and the principal delivers the balance, $\rho b - c_0$. It is exactly as if the principal gave the agent a gross consumption c_0 of 0 and paid the loan off entirely himself.

As usual, we require Bayes consistency and summing to one. These equations, call them (5′) and (6′), do not differ very much from the equivalent equations (5) and (6); one has only to replace C_1 with C_0 (we do not rewrite them here). Given that the principal is assigning gross consumptions c_0 in C_0, for promise keeping, the policies must satisfy

$$\sum_{C_0 \times Q \times A \times B \times W'} \pi(c_0, q, a, b, w')(u([c_0 - \rho b]_+) + v(1 - a) + \beta w') = w_0. \quad \textbf{(13)}$$

The assigned effort and loan size must be jointly incentive compatible so that, for all assigned $(\bar{a}, \bar{b})$, there is no deviation $(\hat{a}, \hat{b})$ that makes the agent better off, given conditional gross consumption and promised utility policies. Hence we require the policies to satisfy, for all $(\bar{a}, \bar{b}, \hat{a}, \hat{b})$ in $A \times B \times A \times B$,

$$\sum_{C_0 \times Q \times W'} \pi(c_0, q, \bar{a}, \bar{b}, w')[u([c_0 - \rho \bar{b}]_+) + v(1 - \bar{a}) + \beta w']$$

$$\geq \sum_{C_0 \times Q \times W'} \pi(c_0, q, \bar{a}, \bar{b}, w') \frac{P(q \mid \hat{a}, \hat{b})}{P(q \mid \bar{a}, \bar{b})} [u([c_0 - \rho \hat{b}]_+) + v(1 - \hat{a}) + \beta w'].$$

$$\textbf{(14)}$$

The ratio $P(q \mid \hat{a}, \hat{b}) / P(q \mid \bar{a}, \bar{b})$ is a measure of the information available from output to the principal about the agent's joint choice of inputs $(\hat{a}, \hat{b})$.

So far, we have allowed the agent to default on his obligation to the outside lender, although the principal makes up the difference through equation (12). We now consider how to restrict the principal's contracts so that, along the equilibrium path, the agent never defaults on his obligations to the outside lenders.

To prevent default, the principal must assign gross consumption that always covers the cost of borrowing. This means that $c_0 \geq \rho \bar{b}$, which in turn implies that $[c_0 - \rho \bar{b}]_+ = c_0 - \rho \bar{b}$. Technically, we model this constraint by requiring that the contract π assign zero-probability weight for all points $\bar{c}_0, \bar{b}$ in $C_0 \times B$ such that $\bar{c}_0 < \rho \bar{b}$. That is, for all $\bar{c}_0, \bar{b}$, in $C_0 \times B$, such that $\bar{c}_0 < \rho \bar{b}$,

$$\sum_{Q \times A \times W'} \pi(\bar{c}_0, q, a, \bar{b}, w') = 0. \quad \textbf{(15)}$$

Off the equilibrium path, when the agent is considering some alternative level of borrowing, $\hat{b}$, the principal will not guarantee the agent's loan and the outside lender may no longer be guaranteed payment. However, because the incentive

compatibility constraints (14) are satisfied, the agent never deviates, and the outside lenders know this.

As in the moral-hazard, credit-control regime above, the lowest incentive-compatible utility occurs when the lowest net consumption $c_{1_{min}}$ and the lowest effort a_{min} are assigned with certainty forever. The highest utility, as before, occurs when the highest net consumption and lowest effort level are assigned with certainty forever. Hence, with the appropriate choice of space C_0 so that net consumptions C_1 match those used in the first two regimes, the utility bounds on the no-control problem match those of the moral-hazard, credit-control program (2) above, as detailed in equations (10) and (11).

PROGRAM 3A (Moral-hazard, no-credit-control, with default). *Maximize $V(w)$, the principal's objective function* (12), *by choice of nonnegative contract weights $\pi(c_0, q, a, b, w')$ over points in $(C_0 \times Q \times A \times B \times W')$ satisfying constraints* (5′), (6′), (13), *and* (14). *The range of promised utilities W' is defined by equations* (10) *and* (11).

PROGRAM 3B (Moral-hazard, no-credit-control, without default). *Maximize the principal's objective function* (12) *by choice of nonnegative contract weights $\pi(c_0, q, a, b, w')$ over points in $(C_0 \times Q \times A \times B \times W')$ satisfying constraints* (5′), (6′), (13), (14), *and* (15). *The range of promised utilities W' is defined by equations* (10) *and* (11).

3.4. Pure Debt Regime

Define a pure-debt contract as one in which the agent borrows an amount b at an internal rate r from the principal, which amount he invests fully in his productive technology. The internal rate r in R is a new choice object and may be viewed as the price of the material input faced by the agent. In this program, the principal assumes the role of an intermediary offering credit by on-lending to the borrowing agent from some larger world market. The principal offers limited insurance via default but not otherwise. Thus we are assuming that the principal has sole access to the outside credit market and knows the amount lent to the agent. The agent cannot go elsewhere for credit. Nevertheless the agent is presumed to decide on the amount borrowed, b, at the rate charged, r, by the principal.

The borrower (that is, the agent) repays the fixed obligation rb if and only if the actual output from his risky technology q allows it, that is, if $q \geq rb$. If there is no default, then the borrower consumes the residual $q - rb$. Default occurs when $q < rb$: The lender (that is, the principal) collects q and the agent consumes zero. The lender, however can promise continuation utilities, w', conditional on output, e.g., higher for high outputs. Because utility is closely (and inversely) tied to the rate charged, r, this conditionality can be viewed as allowing future interest rates to depend on current output realizations in general and on default in particular.

By formulating the pure-debt program in this fashion, we perhaps overstate the power of the lender. Even if loan size b were zero, so that the borrower is

not currently involved in a credit contract, the lender still can make promised continuation utilities conditional on current output. Our interpretation is that a lender can decide on loan terms conditional on the output history of the agent regardless of whether the agent has borrowed or not.

Formally, optimal contracts of this form may be calculated by adding two constraints and a choice object to the moral-hazard, no-credit-control regime with default, Regime 3a presented above. The new restrictions preclude a priori insurance in gross consumption. First, we require that gross consumption allocations c_0 must equal realized output q in each state. Next, we require that the internal price r not vary with output q. In terms of Regime 3a, the principal collects the output, but he now is constrained to return that output as the agent's gross consumption allocation. The debt contract then determines net consumption as a function of the amount borrowed, the internal price, and the output realization (and hence bankruptcy). The new choice object is the internal price r in R charged by the principal.

For technical reasons, the simplest way to implement the new constraint that gross consumption equal output is to require that, for each $(\bar{c}_0, \bar{q})$, such that $\bar{c}_0 \neq \bar{q}$,

$$\sum_{A \times B \times R \times W'} \pi(\bar{c}_0, \bar{q}, a, b, r, w') = 0. \tag{16}$$

Finally, we require, consistent with our interpretation, that the principal announce the value of r *prior* to the realization of uncertainty q. That is, r must be independent of q. So, for all $\bar{r}$ in R and each value $q_i, q_j \in Q$, we require $\Pr(\bar{r} \mid q_i) = \Pr(\bar{r} \mid q_j)$, or

$$\frac{\sum_{A \times B \times W'} \pi(q_1, a, b, \bar{r}, w')}{\sum_{A \times B \times R \times W'} \pi(q_1, a, b, r, w')} = \cdots = \frac{\sum_{A \times B \times W'} \pi(q_{n_Q}, a, b, \bar{r}, w')}{\sum_{A \times B \times R \times W'} \pi(q_{n_Q}, a, b, r, w')}. \tag{17}$$

Note that this constraint does not rule out lotteries over r; it merely requires that these lotteries be resolved before output is realized. If r could depend on the realized output q, then the principal could insure the agent in net consumption.[5]

As usual, the principal must satisfy promise keeping, Bayes consistency, and the adding-up condition. With the control variable r in R, the promise-keeping constraint becomes

$$\sum_{C_0 \times Q \times A \times B \times R \times W'} \pi(c_0, q, a, b, r, w')\{u([c_0 - rb]_+) + v(1 - a) + \beta w'\} = w_0. \tag{18}$$

Bayes consistency still requires the joint probability $\Pr(q, a, b)$ implied by the contract π to be equal to the conditional probability specified by nature $P(q \mid a, b)$ times the marginal probability over inputs specified in the contract $\Pr(a, b)$. That is,

$$\sum_{C_0 \times R \times W'} \pi(c_0, \bar{q}, \bar{a}, \bar{b}, r, w') = P(\bar{q} \mid \bar{a}, \bar{b}) \sum_{C_0 \times Q \times R \times W'} \pi(c_0, q, \bar{a}, \bar{b}, r, w'). \tag{19}$$

the contracts π still must be nonnegative, and we must now require adding-up over one extra variable so that

$$\sum_{C_0 \times Q \times A \times B \times R \times W'} \pi(c_0, q, a, b, r, w') = 1. \tag{20}$$

In this regime, the principal's objective function is rather complicated. In words, the principal (acting as an on-lender) collects the agent's output q, repays the outside lenders an amount ρb, transfers to the agent a net amount $[c_0 - rb]_+$ and then receives in return any amount of net consumption greater than the satiation point $c_{1_{\max}}$. Hence, the principal's objective function is

$$V(w_0) = \sum_{C_0 \times Q \times A \times B \times R \times W'} \pi(c_0, q, a, b, r, w')\{q - \rho b - [c_0 - rb]_+$$

$$+ \left[[c_0 - rb]_+ - c_{1_{\max}}\right]_+ + \beta V(w')\}. \tag{21}$$

Note that if net consumption exceeds the satiation point $c_{1_{\max}}$, the agent refunds the difference to the principal. This can happen at high outputs when r is negative.

We continue to assume not only a moral-hazard problem on effort, but also that the principal cannot directly control the quantity borrowed.[6] Hence, the principal must respect a joint incentive-compatibility constraint on effort and credit similar to (14), but now with the addition of the internal price of credit r in R as a choice variable. As we have maintained in this section, the agent knows r before making any input decisions. Thus, for all assigned effort, credit combinations $(\bar{a}, \bar{b})$ and potential deviations $(\hat{a}, \hat{b})$, given the internal price of credit $\bar{r}$ in R, the policy π must satisfy

$$\sum_{C_0 \times Q} \pi(c_0, q, \bar{a}, \bar{b}, \bar{r}, w')[u([c_0 - \bar{r}\bar{b}]_+) + v(1 - \bar{a}) + \beta w']$$

$$\geq \sum_{C_0 \times Q} \frac{P(q \mid \hat{a}, \hat{b})}{P(q \mid \bar{a}, \bar{b})} \pi(c_0, q, \bar{a}, \bar{b}, \bar{r}, w')[u([c_0 - \bar{r}\hat{b}]_+) + v(1 - \hat{a}) + \beta w'].$$
$$\tag{22}$$

When the agent considers deviating in borrowing, to some amount $\hat{b}$, he uses the known internal price $\bar{r}$ to determine the change in net consumption. If he had direct access to the outside lenders, he would use the world price ρ.

The smallest incentive-compatible promised utility in this regime is the present discounted value of remaining in autarky permanently. In autarky the agent borrows nothing, consumes the output, and chooses effort to maximize expected contemporaneous utility. If the agent were assigned a discounted utility value below autarky, he could always increase his utility by operating the technology at no credit in this period and in all future periods. Thus,

$$w_{\min} = \frac{1}{1 - \beta} \left(\max_a \sum_Q P(q \mid a, 0) \cdot [u(q) + v(1 - a)] \right). \tag{23}$$

Note that, although the lowest promised utility that the agent can be assigned is that associated with autarky, at that lowest point in W, the optimal contract does not necessarily put the agent in autarky forever. That is, unlike the lower bounds on utility calculated in the previous regimes, for the pure-debt regime here, the lowest promised utility is not necessarily an absorbing state (although it can be). That is, an equivalent utility value can be obtained by a nontrivial credit contract offering greater surplus to the principal.

It is worth noting that the agent is driven to autarky by values of the internal price r satisfying

$$r \geq r_{max} \equiv q_{max}/b_2, \tag{24}$$

where b_2 is the smallest nonzero element of B. At $r \geq r_{max}$, even the smallest positive amount of borrowing leads to zero net consumption at the highest output (and hence all lower outputs). Thus the agent would never borrow. In our numerical examples, we specify r_{max} as the maximum value of r in R.

The minimum value of r in R, r_{min}, is determined by the satiation point, the upper bound on net consumption, $c_{1_{max}}$ and, in general, will be negative. If the agent were assigned net consumptions above $c_{1_{max}}$, the difference would be (naturally) refunded to the principal. Thus, r_{min} is the interest rate at which the agent, in principle, is assigned at least the satiation quantity of net consumption, $c_{1_{max}}$, at the lowest output. If the agent gets $c_{1_{min}}$ at the lowest output and highest borrowing level, he will get it at all higher outputs. Thus, r_{min} is implicitly defined by

$$q_{min} - r_{min}b_{max} = c_{1_{max}}, \tag{25}$$

and explicitly by

$$r_{min} = \frac{q_{min} - c_{1_{max}}}{b_{max}}. \tag{26}$$

Note that r_{min} may be negative. The range for R thus is determined by equations (24) and (26).[7]

Therefore, with r_{min} from (26), the highest promised utility in the pure-debt regime occurs when $r = r_{min}$ with certainty permanently. In this case, the agent consumes $c_{1_{max}}$ with certainty each period, and hence sets effort to its lowest level. So, for the pure-debt regime,

$$w_{max} = \frac{1}{1 - \beta} \left[u\left(c_{1_{max}}\right) + v(1 - a_{min}) \right]. \tag{27}$$

Note that the debt regime begins to offer insurance at promised utility values below the maximum, however. At the lowest promised utility value w_i at which the borrower realizes, at the highest output, $q_{max} - r_i b_i > c_{1_{max}}$, net consumption is truncated to $c_{1_{max}}$. At promised utilities above w_i, net consumption is truncated for intermediate outputs, until eventually even at the lowest output, net consumption meets or exceeds $c_{1_{max}}$. This occurs, by construction, at the highest promised utility.

We now can formulate the program to generate optimal contracts in the pure-debt regime.

PROGRAM 4 (Pure debt). *Maximize the principal's objective function* (21) *by choice of nonnegative contract weights* $\pi(c_0, q, a, b, r, w')$ *over points in* $(C_0 \times Q \times A \times B \times R \times W')$ *satisfying constraints* (18), (19), (20), (22), (16), *and* (17). *The range of promised utilities* W' *is defined by equations* (23) *and* (27).

4. DYNAMICS

One key feature not yet emphasized is that in all of the regimes considered, the principal controls the borrowing and lending across periods, denying the agents individual access to across-period world credit markets.

The surplus of a risk-neutral principal, as in (3), (12), and (21), is the expected value of contemporary surplus plus the surplus from the next period discounted by β, with the expectation taken over the optimal policy function $\pi^*(s, w')$.[8] We thus can identify an intertemporal (across-period) interest rate ρ^* by $\beta = 1/(1 + \rho^*)$ or $\rho^* = 1/\beta - 1$, but again, no agent can borrow or lend at this rate. The principal might offer internal across-period credit at an internal interest rate $r^*(w, w')$ indexed by agent type (promised utility) in both periods. Then, from the agent's intertemporal, borrowing-lending first-order (Euler equation) condition, $r^*(w, w')$ may be adjusted to equate both sides of

$$u'(c_t) = \beta(1 + r^*(w, w'))E_{\pi^*(s', w'')}\{u'(c_{t+1}) \mid w'\}.$$

At $r^*(w, w')$, no agent would desire additional across-period credit. In this interpretation, agents are forbidden from trading with each other. Given the differing internal prices faced by each agent, there otherwise would be arbitrage opportunities.[9]

We start off all agents in each regime at $t = 0$ at the promised utility w^*, which gives a surplus of zero, so that $V(w^*) = 0$ for the surplus function as defined for each regime, in equations (3), (12), and (21). This is equivalent, assuming a continuum of agents, to a present discounted surplus of zero for the intermediary. Because the surplus function differs across regimes, the value of w^* will as well.

The dynamics of the model then are completely determined by the policy functions $\pi(s, w' \mid w)$, which themselves are determined as the solutions to the principal's maximization program in each of the regimes. The dynamics of a regime are characterized by considering policies $\pi(s, w' \mid w)$ as Markov transition matrices, and the period-by-period allocations are determined along the way.

Specifically, given promised utility w_i in this period,

$$\Pr(s \mid w_i) = \sum_{w'_j \in W'} \pi(s, w'_j \mid w_i)$$

gives the probability of contemporary outcomes, s, this period. Then, the transition probability matrix $\mathbf{\Psi}$ with entries Ψ_{ij} may be written

$$\Psi_{ij} \equiv \Psi(w_j' \mid w_i) \equiv \sum_{s \in S} \pi(s, w_j' \mid w_i). \tag{28}$$

Each row in $\mathbf{\Psi}$ corresponds to a promised utility w_i this period, whereas each column corresponds to a promised utility w_j' next period, so that entries across a row give the probability distribution over next period's utilities and must sum to unity. Finally, given a particular promised utility w_i this period, the probability distribution over an event (s', w'') in $(S \times W')$ next period is given by

$$\Pr(s', w'' \mid w_i) = \pi(s', w'' \mid w_j')\Psi(w_j' \mid w_i).$$

In a similar fashion, we also can deduce any desired state transition, $\Pr(s', s)$ (for example, the probability of getting no credit tomorrow, given that no credit was assigned today). Thus the utility transition matrix, $\mathbf{\Psi}$, along with the policies $\pi(s, w' \mid w)$, completely characterizes the within-period and across-period dynamics of this model.

When we have a continuum of agents, we simply interpret probabilities over promised utilities as fractions of the agent population actually given those promised utilities. Starting the economy at $t = 0$ in present-value budget balance, let i^* be an $n_W \times 1$ column indicator vector of zeros except where the vector of possible promised utilities is equal to w^*, and put a one in that position.[10] Recall that the transition matrix $\mathbf{\Psi}$ is $n_W \times n_W$. In the usual Markov fashion, the distribution of agents over promised utilities at date $t = 1$, i_1, after one period of uncertainty, is

$$i_1 = \mathbf{\Psi}'i^*,$$

and at date $s \geq 1$ it is

$$i_s = (\mathbf{\Psi}')^s i^*.$$

Hence the invariant distribution reached by starting at i^* is given by

$$i_\infty = (\mathbf{\Psi}')^\infty i^*.$$

The invariant distribution will have associated with it a distribution over within-period states, $H(s)$, and transitions $D(s, s')$, defined by

$$H(s) \equiv \sum_W \Pr(s \mid w)i_\infty(w) \tag{29}$$

$$D(s, s') \equiv \sum_{W, W'} \Pr(s \mid w)\Pr(s' \mid w')\Psi(w' \mid w)i_\infty(w). \tag{30}$$

Assuming a continuum of agents, $H(s)$ can be interpreted as the proportion of agents observed in contemporaneous state s at the invariant distribution. This could be approximated by a large cross section with data on the variables s. Further,

$D(s, s')$ can be interpreted as the proportion of agents moving from state s this period to state s' next period at the invariant distribution. This is turn could be approximated by a large panel with data on s and s'.

Because in the pure-debt regime the principal has no right to output q when the agent does not borrow, w_{min} calculated in equation (23) generally will be higher than the minimum promised utilities in the first three regimes. Any division of property rights that decreases the choice set of the principal (by limiting W') shifts the entire Pareto frontier down. If we begin all agents at the promised utility yielding zero surplus, all agents are, paradoxically, made worse off by such a division of property rights. In other words, the ability to punish the agent is valuable in solving the moral-hazard problem.

Still, the variation in lower promised utility end points across regimes is bothersome. If we wish to Pareto compare regimes, we have to control for the fact that, for example, the full-information regime has many more choice variables (larger range of promised utilities W') than do the others. This variation seems to result from the property-rights division implied by each regime. The pure-debt regime has the highest minimum promised utility simply because, as we defined it, the principal has no legal claim to output under the specified contract form if borrowing is zero. Autarky is a reasonable reservation value for the household or firm in the other regimes. We take this autarky utility to be the uniform minimum utility end point across all regimes, with the idea that agents could drop out of any of the other regimes if pressed lower than this point.

In Section 6.2, we do report Pareto frontiers without imposing autarky as the lower bound, demonstrating the effect of the uniformity assumption.

In the full-information regime, each promised utility point ought to be an absorbing state, because there are no intertemporal tie-ins.[11] In the control and no-control regimes (Regimes 2, 3a, and 3b) there are at least two absorbing promised utilities (the very highest and very lowest promised utilities as specified in Sections 3.2 and 3.3), since they were defined assuming that they were permanently replicated with certainty. Because conditional promised utilities (along with conditional consumptions) are used to reward desired outputs and punish undesired outputs, the observed distributions of promised utilities over time of these models often degenerate to a mass of unity at the lowest promised utility. This makes the steady-state distributions uninteresting from an empirical point of view.[12] Again, in the first three regimes, we discard all promised utilities below autarky. Because autarky generally lies above the minimum promised utility in these regimes, it is not an absorbing state. The principal can achieve the autarky utility point without setting the continuation utilities to autarky with certainty. Therefore, we do not expect to see, and, in the computed solutions, do not observe, steady-state distributions of promised utilities that degenerate to mass points at the lowest promised utility for the first three regimes. Autarky is potentially an absorbing state for the pure-debt regime, so we might observe degenerate probability distributions at autarky. Depending on parameter values, however, there may be nonautarky policies that deliver the same (autarky) promised utility but with a higher surplus. Also, the highest promised

utility is an absorbing state for all regimes, and so, it is technically possible for a steady-state distribution to be degenerate at the highest promised utility.

5. PARETO COMPARISONS

Given uniform ranges for promised utility, we can explicitly compare outcomes under the regimes. Because the regimes are constructed to isolate the effects of specific credit market features, by ranking the regimes we can determine their effects. In particular, we find that access to the outside lenders significantly alters the optimal contracts; that pure-debt contracts may Pareto dominate more complex contracts without default; and that allowing default is Pareto improving. We also can determine, from within this framework, which agents (if any) will report themselves credit constrained.

PROPOSITION 1 (Regime Pareto rank). *Given the same set of underlying parameters, the principal's surplus* (*considered across the set of all promised utilities*) *will be ranked from highest to lowest across regimes*:

1. *Full-information* (*Regime* 1)
2. *Moral-hazard, credit-control* (*Regime* 2)
3. *Moral-hazard, no-credit-control, with default* (*Regime 3a*)
4. *Moral-hazard, no-credit-control, without default* (*Regime 3b*).

Regime 4, *the pure-debt regime, will certainly lie below Regime* 2, *but cannot be Pareto ranked relative to Regimes* 3a *and* 3b.

Proof. The principal's surplus is the objective function to the programming problem that forms each regime. Regime 2 is simply Regime 1 with an incentive compatibility constraint on effort [equation (9)]. Regime 3a is Regime 2 with incentive compatibility constraints (14) on borrowing as well as effort. Regime 3b is simply Regime 3a with the no-default constraints (15) added. Regime 4, the pure-debt regime, is also a constrained form of Regime 3a, but it has the additional choice variable of r, the internal price of credit, which also applies when the agent considers deviating in borrowing level. The principal in the no-credit-control regimes (3a and 3b) cannot prevent the agent from concluding trading arrangements with the outside lenders, whereas the principal in the pure-debt regime (Regime 4) apparently does have that ability. Hence Regimes 3a and 3b are not directly comparable with Regime 4. The moral-hazard, credit-control regime (Regime 2) certainly does dominate the pure-debt regime (Regime 4) because, with control of borrowing and net consumption, the principal in Regime 2 can replicate any pure-debt contract. ∎

6. RESULTS, DISCUSSION, AND EXAMPLES

6.1. Credit Constraints

We now turn our attention to credit constraints. We begin with a definition, and then move on to consider the incidence of credit constraints.

DEFINITION 1 (Credit constraint). *Agents would report being credit constrained only if they would choose, at a specified price, a level of credit-financed material input different than the one that they are assigned under the contract they face.*

Because the internal price is part of the contract faced by agents in the pure-debt regime, we assume that they are asked if they would like to choose a different amount of the material input at that internal price. Similarly, we assume that agents in all other regimes are asked if they would choose a different level of the credit-financed material input at the world price ρ. With this definition, we can discuss which agents would claim to be credit constrained.

PROPOSITION 2 (Conditions for self-reported credit constraints). *Agents who are in the full-information, moral hazard, no-credit-control and pure debt regimes (regimes 1, 3a, 3b and 4) will never report themselves to be credit constrained under the contract they face. At least one agent type in the moral hazard, credit-control regime (regime 2) will report himself as credit constrained if the optimal contracts for the no-credit-control regime with default and the credit-control regime differ at at least one w in **W**. In addition, the direction of the desired deviation can be determined.*

Proof. Agents in the full-information regime, who are fully insured, will agree with the principal about the proper inputs necessary to maximize the expected output of the project. Agents in the moral-hazard, no-credit-control regimes with and without default (Regimes 3a and 3b), face contracts that must satisfy the joint incentive-compatibility constraints (14), whereas agents in the pure-debt regime (Regime 4) face contracts that must satisfy the joint incentive-compatibility constraints (22). This means that agents in Regimes 3a, 3b, and 4, if offered a deviation in inputs, would always decline. Hence they would never report themselves as credit constrained, although their credit-financed inputs may differ across regimes. The moral-hazard, credit-control regime (Regime 2) and the moral-hazard, no-credit-control regime with bankruptcy (Regime 3a) differ only in the extra incentive-compatibility constraints on the material input b. Therefore, if the solutions in the two regimes differ, it must be because the extra constraints on the material input were binding for at least one level of promised utility.

If offered the contract (that is, the assigned inputs, conditional consumption, and promised utility schedules) from the credit-control regime (Regime 2), at one or more promised utilities, an agent in the no-credit-regime (Regime 3a) would choose some different level of the material input. By checking this deviation, we can determine if the agent in the credit-control regime (Regime 2) would prefer to use more or less of the material input. That is, given that the contract $\pi^*(c_1, q, a, b, w' \mid w)$ is the solution to Program 2, the moral-hazard, credit-control regime, we can determine, for all w in $\boldsymbol{W}$ the agent's choice of effort and material input if he is allowed to choose b (along with a) while facing the contract π^* (which was written assuming that the agent could not choose b). We assume that the principal in the moral-hazard, credit-control regime assigns the agent a gross consumption of $c_0 = c_1 + \rho\bar{b}$ (because the principal controls b directly, this

is merely an accounting measure). But now, the agent will be able to affect net consumption directly by choosing different levels of b at the world price ρ. Define $(\hat{a}^*(w), \hat{b}^*(w))$ as

$$(\hat{a}^*(w), \hat{b}^*(w)) = \arg \max_{(\hat{a}, \hat{b}) \in A \times B} \sum_{C_1 \times Q \times A \times B \times W'} \pi^*(\cdot \mid w) \frac{P(q \mid \hat{a}, \hat{b})}{P(q \mid a, b)}$$

$$\times \{u([c_1 + \rho(b - \hat{b})]_+) + v(1 - \hat{a}) + \beta w'\}. \quad \textbf{(31)}$$

If the agent is assigned a level of borrowing $\bar{b} > 0$, by deviating to a lower level $\hat{b} < \bar{b}$, the agent gains a consumption bonus at each output realization. This effect we call credit diversion and discuss in more detail below. There is, however, another effect manifest in equation (31). Because the agent has chosen a new set of inputs, the probability distribution over outputs (and consumption assignments conditional on output) has changed. This can lead to complaints about credit shortages, in which the optimal deviation $\hat{b}$ is greater than the assigned level $\bar{b}$.   ■

Credit assignments will thus depend on the productivity and information provided by the underlying technology. We define two useful measures, first productivity and next information.

DEFINITION 2 (Productivity). *Productivity is defined in terms of expected net output, which is the natural measure, given that the principal is risk neutral. First, define $\mu(a, b)$ as the expected output at a given input combination (a, b), so that*

$$\mu(a, b) = \sum_Q q P(q \mid a, b).$$

We say that effort is productive for all material input levels when, for all $a' > a$ and b in B,

$$\mu(a', b) > \mu(a, b). \quad \textbf{(32)}$$

In the same way, we say that the material input is productive for an assigned effort $\bar{a}$ if, for all $b' > b$, and $\bar{a}$ in A,

$$\mu(\bar{a}, b') - \rho b' > \mu(\bar{a}, b) - \rho b. \quad \textbf{(33)}$$

For $a' > a$ and $b' > b$, a and b are jointly productive if $\mu(a', b') - \rho b' > \mu(a, b) - \rho b$. We say effort is neutral in production if, for all $a' \neq a$ and b in B, $\mu(a', b) = \mu(a, b)$, and material input is neutral in production if, for all $b' \neq b$ and a in A, $\mu(a, b') - \rho b' = \mu(a, b) - \rho b$.

DEFINITION 3 (Information). *At a particular output q in Q and assigned credit level b in B, let the information about a particular assigned effort $\hat{a}$ and deviation in effort $\hat{a}$, $I(q, b; (\bar{a}, \hat{a}))$ be defined as*

$$I(b, q; (\bar{a}, \hat{a})) = \left[\log\left(\frac{P(q \mid \hat{a}, b)}{P(q \mid \bar{a}, b)} \right) \right]^2. \quad \textbf{(34)}$$

Note that if the probability of a particular output q is the same when the inputs are $(\hat{a}, b)$ as when they are $(\bar{a}, b)$, then $I = 0$. The information measure I increases as the probabilities diverge. For example, if at $\hat{a}$, the probability of a particular output q fell to zero, then at that output, information would be unbounded. The principal would know with certainty, were that output realized, that the agent had not deviated.

Consider a technology in which capital and effort were jointly productive, but in which $I(b', q; (a', a)) \leq I(b, q; (a', a))$ for $b' > b$ and $a' > a$, that is, higher capital levels provide less information about effort at each output realization. If the productivity effect were small enough, then the principal in the moral-hazard regimes (both with and without credit control) would have an incentive to deny the agent credit. Denying credit, in this circumstance, has the flavor of using credit as a monitoring device.

As an example of this effect, we set output $Q = \{0, 1\}$, effort $A = \{0, 1\}$, material $B = \{0, 1\}$, with the world price of material $\rho = 0.0499$, the discount factor $\beta = 0.8$, and the risk-aversion parameter $\gamma = 0.2$. The agent will have linear disutility in effort equal to $\bar{v} \cdot (1-a)$, with $\bar{v} = 0.0375$. Finally, we set the technology $P(q \mid a, b)$ as in Table 1A.

In the technology defined in Table 1A, the credit-financed input is productive when effort is high, because $\mu(a = 1, b = 1) - \rho \cdot 1 > \mu(a = 1, b = 0) - \rho \cdot 0$. However, information about deviations to the low-effort level $\hat{a} = 0$ is worse when b is high than when it is low. By assigning the agent the low-material-input level along with the high-effort level, the principal loses some expected output because of the productivity effect but gains information about the agent's actions. Plugging into the formula for I when assigned effort $\bar{a} = 1$ and the deviation is $\hat{a} = 0$, we

TABLE 1. Technology: Mapping from inputs A and B into conditional probabilities $P(q \mid a, b)$

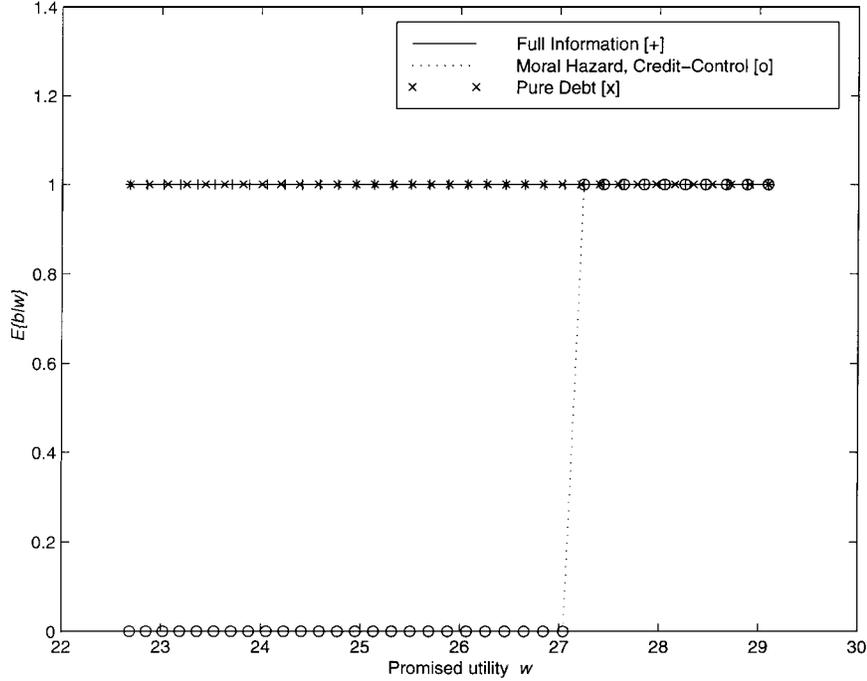| Inputs		Outputs	
B	**A**	$P(q = 0 \mid a, b)$	$P(q = 1 \mid a, b)$
A. Technology with $A = \{0, 1\}$, $B = \{0, 1\}$			
$b = 0$	$a = 0$	0.95	0.05
	$a = 1$	0.10	0.90
$b = 1$	$a = 0$	0.10	0.90
	$a = 1$	0.05	0.95
B. Technology with $A = \{0, 1\}$, $B = \{0, 0.001, 1\}$			
$b = 0$	$a = 0$	0.95	0.05
	$a = 1$	0.10	0.90
$b = 0.001$	$a = 0$	0.50	0.50
	$a = 1$	0.50	0.50
$b = 1$	$a = 0$	0.10	0.90
	$a = 1$	0.05	0.95

FIGURE 1A. Assigned credit when capital is productive but not informative, for Regime 1 (full-information), Regime 2 (moral-hazard, credit-control), and Regime 4 (pure-debt).

see that when output is low, I declines from $(\log(0.95/0.10))^2 = (\log(9.5))^2$ at the low-capital level to $(\log(0.10/0.05))^2 = (\log(2))^2$ at the high-capital level. At the high output, I declines from $(\log(0.90/0.05))^2 = (\log(18))^2$ at the low-capital level to $(\log(0.95/0.90))^2 = (\log(1.0556))^2$ at the high-capital level.

In Figure 1A, we present the assigned capital level under the optimal contract for three regimes: the full-information regime (Regime 1), the moral-hazard, credit-control regime (Regime 2), and the pure-debt regime (Regime 4). We suppress the graphs for the other regimes for clarity. Notice that, in the moral-hazard credit-control regime, agents with low promised utilities are assigned lower credit levels than agents with higher promised utilities. This is because, in the moral-hazard, credit-control regime, agents with low promised utilities are assigned the high-effort level, whereas agents with high promised utilities are assigned the low-effort level. Once the principal is no longer assigning the high-effort level there is no moral-hazard problem and the principal can ignore the effect of the material input on information.

Assigned credit under the moral-hazard, no-credit-control regimes (Regimes 3a and 3b) is almost exactly the same as assigned capital under the moral-hazard, credit-control regime, which is presented here.
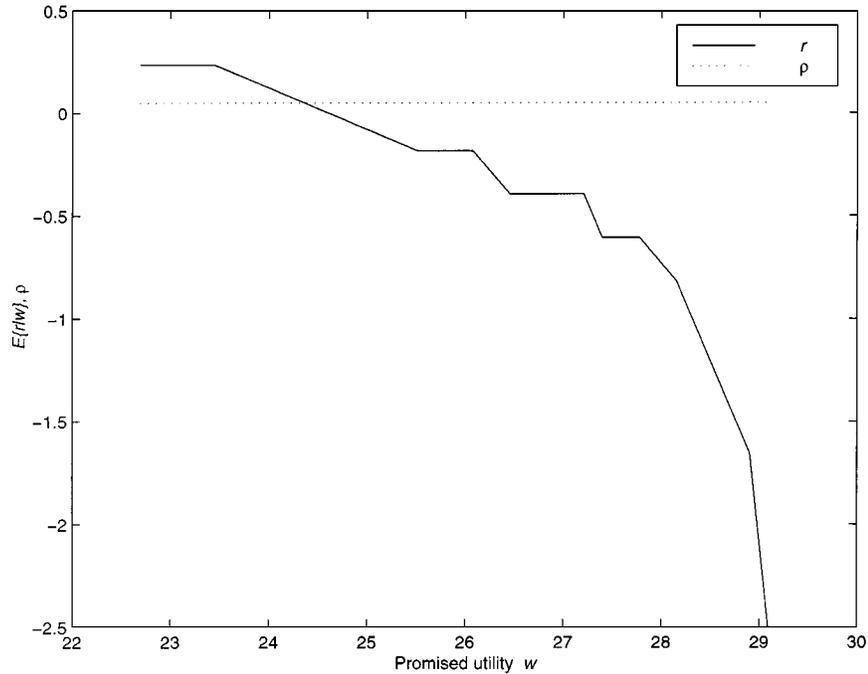
FIGURE 1B. Internal price r assigned in the pure-debt regime (Regime 4): Dotted line is ρ, the external price of credit.

Agents in the pure-debt regime (Regime 4) are assigned the highest credit level over the entire range of promised utilities, including the lowest promised utility. From Figure 1B, one can see that the principal charges agents with low promised utilities internal prices r above ρ, but for agents with higher promised utilities the principal sets $r < \rho$, and eventually $r < 0$. Indeed, the autarky promised utility can be achieved by a contract that pushes the agent's utility this period below what he could realize in autarky, in exchange for future payoffs above autarky.

Although poor agents in the moral-hazard, credit-control regime are assigned low levels of credit, they are *not* credit constrained in the sense of Definition 1, above. Agents with low promised utilities have relatively low consumption assignments and high marginal utilities. They are unwilling to give up own-consumption to finance higher values of the material input b.

With a slight change in the technology presented in Table 1A, however, we can generate poor agents in the moral-hazard, credit-control regime who would report being credit constrained in the sense that they would like to borrow more than the assigned level. The key is to introduce a new, extremely low, level of credit that is not productive, so that the principal would not assign it. The agent is willing to pay a small amount to get credit if it generates a different distribution over outputs, one

that allows him to deviate downward in effort without suffering a large expected utility penalty.

To make this point, we altered the setup slightly and recomputed. Now the credit-financed material input can take on three levels, $B = \{0, 0.001, 1\}$. All other parameters are unchanged, and the augmented technology is presented in Table 1B. Agents with low promised utilities are again assigned no credit and the high effort in the moral-hazard, credit-control regime. The only difference is that now these agents would, if allowed, deviate in borrowing from the assigned level of $b = 0$ to the new level $b = 0.001$.

6.2. Credit Diversion

We next consider a circumstance in which agents in the moral-hazard, credit-control regime would report themselves as receiving too much credit: Given the opportunity, agents in the moral-hazard, credit-control regime would use less of the material input than the assigned amount. Because their gross consumption assignments have to cover the cost of borrowing in all states of the world, by borrowing less than their initial assignment, agents can consume more in all states. Of course, the probability distribution over states changes with the material input, which then can blunt or reverse this effect (as in the example of Section 6.1).

PROPOSITION 3 (Credit diversion). *If the underlying technology is productive in material inputs and neutral in effort, the solutions to the moral-hazard, credit-control program will coincide with the solutions to the full-information program over the relevant W ranges, and they both will differ from the moral-hazard, no-credit-control regimes (both with and without default). Agents in the moral-hazard, credit-control regime will report that they are credit constrained, in the sense of being forced to accept too much credit.*

Proof. If the underlying technology is productive in material inputs, in the sense of satisfying equation (33) and neutral in effort, then the principal in the full-information and moral-hazard, credit-control regimes simply assigns the lowest-effort, highest-material-input combination, $(a_{\min}, b_{\max})$, and provides perfect insurance. There is no incentive to deviate in effort because the lowest effort has been assigned; therefore, there is no barrier to perfect insurance in these regimes.

The agent in the moral-hazard, credit-control regime (Regime 2), however, would report being forced to use too much of the material input in the sense of Definition 1, above. Given full insurance, net consumption c_1 does not vary with output and the agent is assigned a constant gross consumption of $c_0 = c_1 + b_{\max}$ (as if he repays the lenders in full). This is incentive compatible in Regime 2 because the principal controls b directly. If offered a deviation to any lower borrowing level $\hat{b} < b_{\max}$, with the same constant gross consumption allocation, the agent could realize a new, higher, constant net consumption level of $\hat{c}_1 = c_1 + \rho(b_{\max} - \hat{b})$. Hence, he would accept this deviation and would report desiring a lower level of the material input than assigned under the contract. To induce levels of borrowing

TABLE 2A. Technology: Mapping from inputs A and B into conditional probabilities $P(q \mid a, b)$

Inputs		Outputs		
B	**A**	$P(q = 0 \mid a, b)$	$P(q = 3 \mid a, b)$	$P(q = 5 \mid a, b)$
$b = 0$	$a = 0$	0.500	0.300	0.200
	$a = 0.1$	0.400	0.550	0.050
$b = 0.5$	$a = 0$	0.400	0.250	0.350
	$a = 0.1$	0.200	0.750	0.050
$b = 1$	$a = 0$	0.350	0.075	0.575
	$a = 0.1$	0.030	0.875	0.095

TABLE 2B. Promised utility ranges[a]

Program	Regime	min(W)	max(W)	w^*
Full-information	1	9.4868	33.1753	23.00
Moral-hazard, credit-control	2	10.0000	33.1753	23.00
Moral-hazard, no-credit-control no-default	3b	10.0000	33.1753	22.91
Pure debt	4	19.4416	34.1753	22.21

[a] Most results are displayed with the utility range truncated at autarky (the minimum promised utility in the pure-debt regime). See Section 3.6 for the details.

greater than zero in the regimes in which the principal does not directly control the material input, the principal must vary gross consumption allocations across outputs, returning a certain amount of uncertainty to the agent. ■

We now present results from a choice of parameter values that illustrate this proposition. In Table 2A we specify a technology in which higher levels of the material input raise the expected value and variance of output. Effort is constrained to be high ($a = 0.1$) or low ($a = 0$), with high levels of effort decreasing output variance but leaving mean output unchanged. Output can take on three values, with $Q = \{0, 3, 5\}$. Material input can take on three values, with $B = \{0, 0.5, 1\}$, and the world price of borrowing to finance this input is $\rho = 1$. Agents discount at the common rate $\beta = 0.8$ and have preference parameters $\gamma = 0.4$ and $\delta = 0.5$. Net consumption is constrained to go from $c_{1\min} = 0$ to $c_{1\max} = 5.5$. For the pure-debt regime, the minimum internal price is –5.5 and the maximum is 10.5.

In the presence of perfect insurance, we would expect agents to borrow enough to achieve the level of material input with the greatest expected surplus, net of the world borrowing cost, ρ, and to have the lowest effort level. Because that surplus is strictly increasing in material input level, the highest value of b in B, the riskiest technology, is clearly the social best. Effort is simply a costly method of self-insurance, and in a first-best world, the principal should be providing all of the insurance.

TABLE 2C. Expected output, variance, expected net output, expected utility of consumption, and expected utility of consumption and effort[a]

Inputs	μ	σ^2	$\mu - \rho b$	$E\{u(c_1)\}$	$E\{u(c_1) + v(1-a)\}$
$b=0, a=0$	1.9	4.09	1.9	1.8421	3.8421
$b=0, a=0.1$	1.9	2.59	1.9	1.9910	3.8883
$b=0.5, a=0$	2.5	4.75	2.0	2.1603	4.1603
$b=0.5, a=0.1$	2.5	1.75	2.0	2.3715	4.2689
$b=1, a=0$	3.1	5.44	2.1	2.3911	4.3911
$b=1, a=0.1$	3.1	0.64	2.1	2.5742	4.4715

[a] $c_1 = [q - \rho b]_+$, $\mu = E\{q \mid a, b\}$ and σ^2 is var$(q \mid a, b)$.

In Table 2C we compute the expected value of output, output net of world loan repayment, utility of consumption alone (to highlight the utility cost to risk), and utility of consumption and leisure (to judge the trade-off between effort and self-insurance) for each of the input levels in the model, assuming that the agent is not otherwise insured, and that he borrows at the world rate ρ. The point of this table is demonstrate how a typical borrower would view the different input choices. In this case, the typical borrower would choose the highest effort and the highest loan size. This is because the material input is productive, as noted earlier, and variances are decreasing in effort, increasing in credit but decreasing again in effort and credit jointly (providing a self-insurance motive to higher levels of borrowing). Note that with $\gamma = 0.4$ the agent isn't particularly averse to consumption risk.

In Figure 2A, we present the expected surplus for four of the regimes over the truncated utility points, as defined in the relevant programs (here and elsewhere for this technology we omit for clarity the results for Regime 3a, the moral-hazard, no-credit-control regime with default). In Figure 2B, we present the Pareto frontiers for each of the four programs without truncating the promised utilities, for comparison. Note that only the surplus for the moral-hazard, no-credit-control regime without default (Regime 3b) should be affected because the solutions to the moral-hazard, credit-control regime will coincide with the solutions to the full-information regime by Proposition 3. These are the only results that we present for the untruncated programs.

As expected from Proposition 3, above, the full-information and moral-hazard, credit-control regimes coincide and dominate the moral-hazard, no-credit-control regime without default. This in turn dominates the pure-debt regime. The pure-debt regime features a low but increasing surplus at the low end of promised utility w because the agent is emerging initially from autarky (where the borrower assumed no credit and the lender has no revenue) into the credit market (where the technology is more productive and the lender does realize revenue). Of course, as promised utility increases forever, the lender (in all regimes) begins subsidizing the agent's consumption (and effort) in the sense of realizing a negative surplus or expected profit.
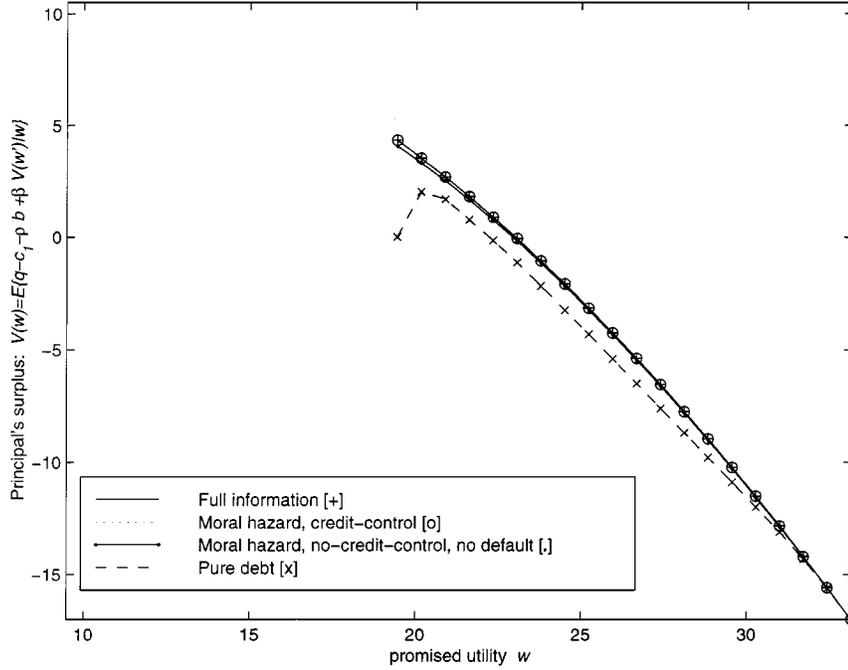
FIGURE 2A. Surplus for each regime when promised utilities are truncated at autarky.

In Figures 2C to 2E, we plot $E\{c_1 \mid q, w\}$, the net consumption allocation conditional on output realization, for the full-information regime (Regime 1), the moral-hazard, no-credit-control, no-default regime (Regime 3b), and the pure-debt regime (Regime 4). Recall that net consumption is constrained to be in [0, 5.5]. The full-information and moral-hazard, credit-control regimes (as expected) provide almost full insurance. The minor deviations from full insurance are the result of the consumption grid. The no-credit-control regime has a significant amount of residual uncertainty. This uncertainty rewards the high output, does not really punish the low output, and severely punishes the middle output. This is because, given the gross consumption allocation, the agent is, on the margin, contemplating deviations to a lower credit level and a higher effort level (the self-insurance motive). Given the technology, these deviations make the middle output more likely; hence it is punished. These consumption allocations provide the most insurance possible, given the possibility of credit diversion.

The most striking of these consumption graphs, though, is Figure 2E. Here, we see the enormous amount of uncertainty present in the pure-debt contract, even relative to the no-credit-control regime.

In the same fashion as above, we plot conditional continuation utility schedules $E\{w' \mid q, w\}$ for the full-information regime in Figure 2F, for the moral-hazard, no-credit-control regime without default in Figure 2G, and for the pure-debt regime
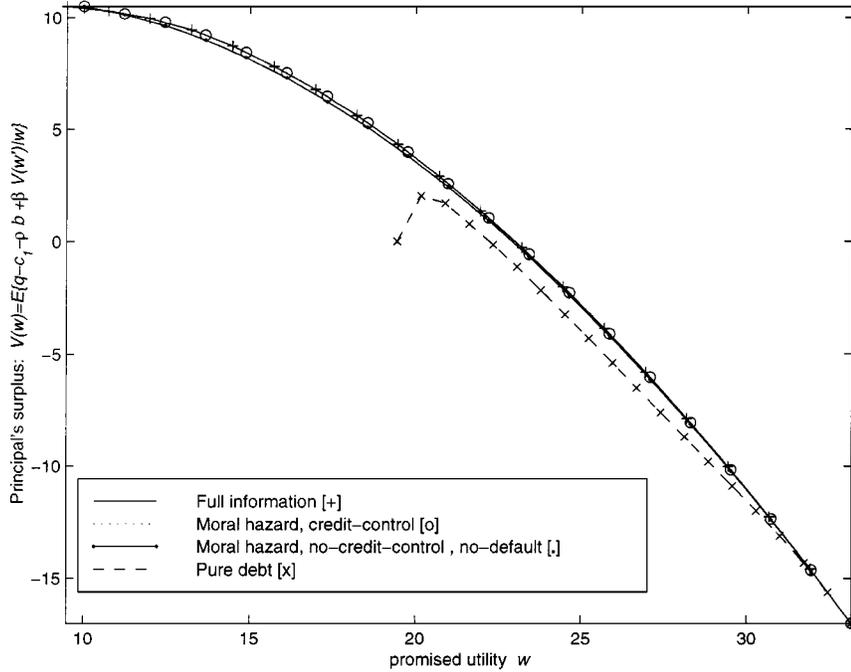
FIGURE 2B. Untruncated surplus for each regime.

in Figure 2H. The full-information and moral-hazard, credit-control regimes have no intertemporal tie-ins, because their conditional continuation schedules lie on the 45-deg line, but where there is a moral-hazard problem and the agent is risk averse, the principal can gain greater surplus by splitting rewards and penalties between within-period and across-period outcomes. We see many of the same features of the conditional-consumption graph echoed here. For example, conditional on observing the middle output, promised utility lies below the 45-deg line. Of course, agents are rewarded with a continuation utility greater than their current level for outcomes with continuation utilities above the 45-deg line and are punished with a continuation utility lower than their current level for outcomes with continuation utilities below the 45-deg line. The conditional continuation utility schedules do not coincide until the highest promised utility, where they must coincide by construction.

In contrast, the moral-hazard, no-credit-control regime displays significant variation across outcomes. In particular, at the autarky promised utility, conditional on high output, the agent is promoted to a higher promised utility. Evidently, autarky is not an absorbing state for the agent in this regime at these parameter values. As with the consumption graphs, the middle output is punished severely, whereas the high output is rewarded.

As demonstrated in Table 2C, the borrower in the pure-debt regime prefers the high-effort, high-material-input choice to all others when $r \leq \rho$. The lender in this
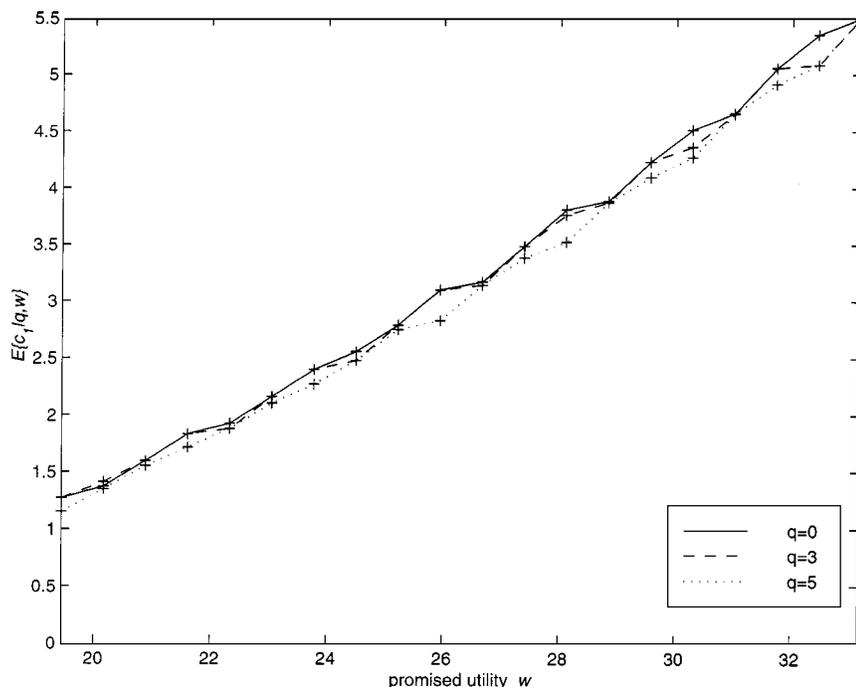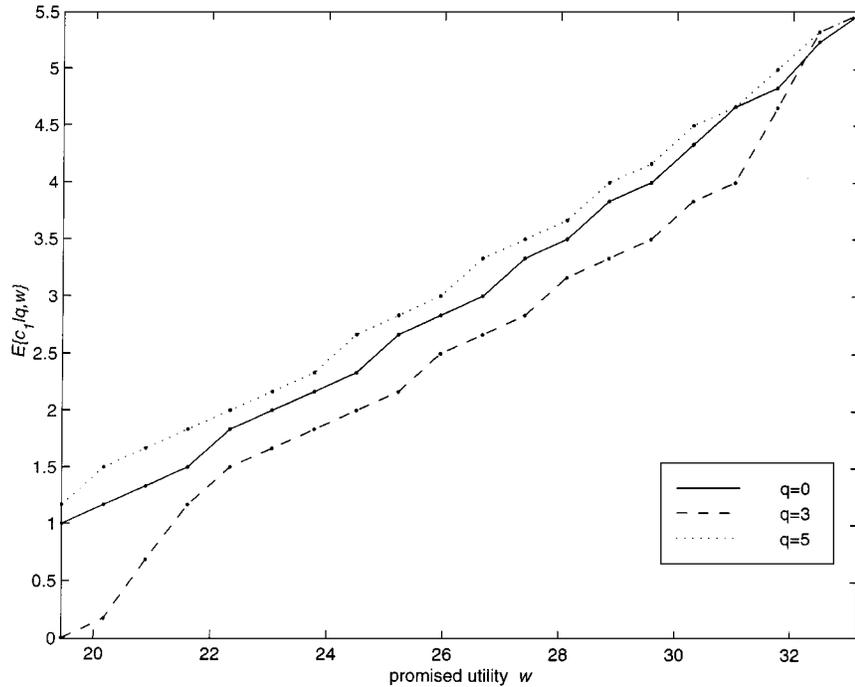
FIGURE 2C. Net consumption allocations for full-information regime.

case assigns exactly those inputs across the first half of promised utilities. At all but the very lowest promised utility points, $r \le \rho$, and so, there is no moral-hazard problem in that region. As a result, in that range, the continuation utilities stick fairly closely to the 45-deg line. In the absence of binding incentive compatibility constraints, there are no intertemporal tie-ins. When, at roughly $w = 27$, assigned inputs change to the low effort, high material input, there is a moral-hazard problem. The binding incentive compatibility constraint is for the deviation to a higher effort. The lender discourages this by heavily punishing the middle output, which is unlikely at the low-effort level, but very likely at the high-effort level. The transitions in this region reflect the intertemporal tie-ins caused by the incentive problem. Similarly, at the very lowest promised utilities, $r > \rho$, and so, input choices have to be induced with conditional promised utilities.

The steady-state utility distributions, displayed in Figure 2I, differ across regimes because of the different transitions in each regime. The full-information and moral-hazard, credit-control regimes coincide with spikes at w^* (the initial promised utility), because their transitions lie mainly along the 45-deg line. The initial distribution over promised utilities will not necessarily be a spike at w^* if w^* lies between grid points in W. The initial distribution instead reflects the weights placed on each of the neighboring points to put the expected initial value of w to

FIGURE 2D. Net consumption allocations for the moral-hazard, no-credit-control, no-default regime.

w^*. In Figure 2J, we display the initial and invariant distributions for the moral-hazard, no-credit-control regime. The relatively complicated transitions of this regime, which are used to discourage deviations in credit, induce a steadily rising distribution over promised utilities.

The pure-debt regime also shows a spike at w^*, although w^* for the pure-debt regime is quite a bit lower than w^* for the other regimes, because they Pareto dominate the pure-debt regime. There is a small amount of mass at the autarky promised utility. This is because there are multiple ergodic sets in the pure-debt transition matrix. The spike in the pure-debt steady-state distribution at the absorbing state (autarky) is because, at the starting point w^*, agents divide between ergodic sets. A certain percentage are demoted to autarky forever and are barred permanently from the credit market. Others are kept near the initial level (although not, of course, exactly at the initial level) of promised utility.

6.3. Default

Default enters our analysis in two ways. First, we explicitly allow and then prohibit default in the moral-hazard, no-credit-control regimes (Regimes 3a and 3b).
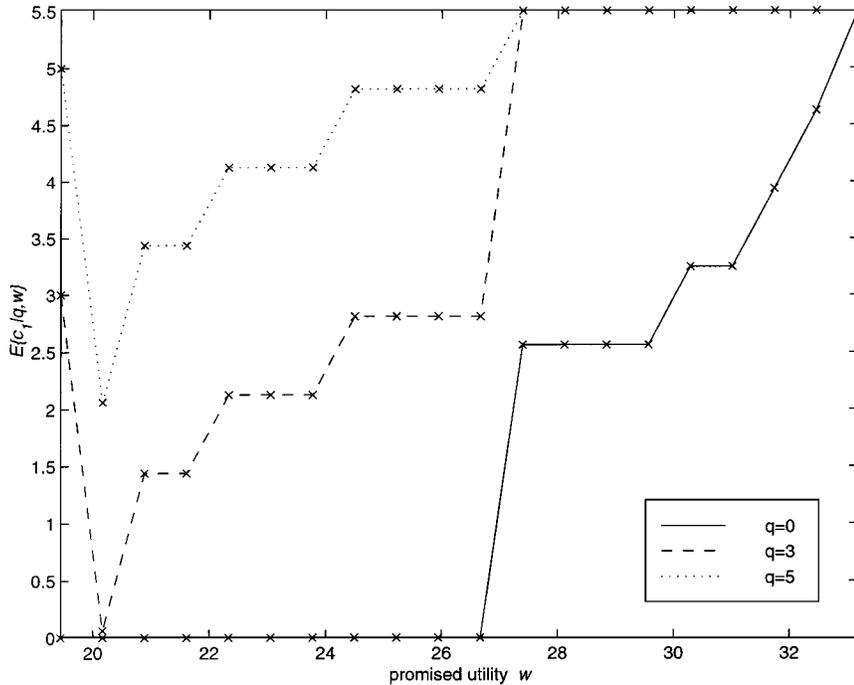
FIGURE 2E. Net consumption allocations for the pure-debt regime.

Second, default is the only mechanism for insurance in the pure-debt regime (Regime 4). In both senses, default can have an important impact.

The effect of prohibiting default in the moral-hazard, no-credit-control regime (i.e., of moving from Regime 3a to Regime 3b) is to force the principal to provide the agent with gross consumption allocations large enough to cover the borrowing assignment, so that $c_0 \geq \rho \bar{b}$. If the principal wishes to assign zero net consumption, he has to set $c_0 = \rho \bar{b}$. When default is allowed, the principal can achieve zero net consumption by setting $c_0 = 0$ and having the agent default. Along the equilibrium path, net consumptions are the same in the two cases. When the agent considers deviating downward in assigned capital, however, his consumption when default is prohibited is larger than when default is allowed. Put differently, when default is prohibited, the agent has to pay back the entire loan out of his gross consumption allocation no matter what the output level. This is a force for borrowing less. Technically, the right-hand side of the incentive compatibility constraints (14) will be larger at high levels of the assigned material input $\bar{b}$. The principal will find it more difficult to assign the high credit level. This effect will be larger at lower levels of promised utility because this is exactly the range in which the principal will assign net consumptions at or close to zero. Again, allowing default eliminates this effect and makes it easier to provide credit to poor agents.

TABLE 3. Technology: Mapping from inputs A and B into conditional probabilities $P(q \mid a, b)$

Inputs		Outputs	
B	A	$P(q = 0 \mid a, b)$	$P(q = 2 \mid a, b)$
$b = 0$	$a = 0$	0.50	0.50
$b = 1$	$a = 0$	0.05	0.95
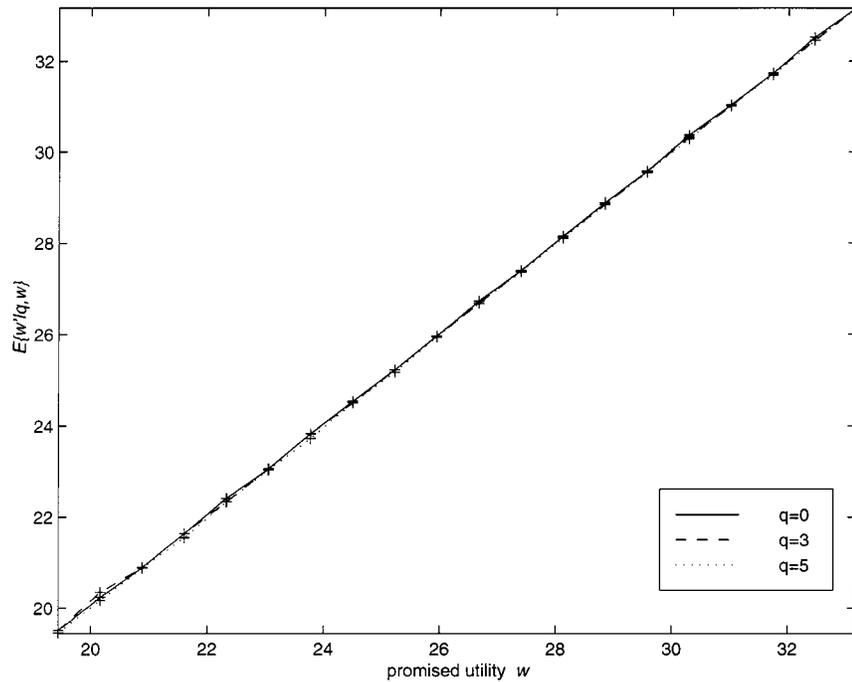
FIGURE 2F. Continuation utility schedule for full-information regime.

We now present an example that highlights this effect. We isolate the effect of the moral-hazard problem on borrowing by eliminating deviations in effort, so that $A = \{0\}$, and allowing only two levels of credit, so that $B = \{0, 2\}$. We set output to $Q = \{0, 2\}$ and $\rho = 0.0499$, with the technology presented in Table 3.

In Figure 3A, we present the assigned credit levels in the moral-hazard, no-credit-control regimes both with and without default. When default is prohibited, the agent is assigned, at low promised utilities, a credit level less than half the level assigned when default is allowed. In Figure 3B, we present the expected surplus in both regimes. Because input choice has been distorted by the no-default constraint, at the lowest promised utility the Pareto frontier is more than 10% lower.

The effect of allowing default as in the pure-debt regime is to introduce some contingencies. Sometimes this is enough so that the value functions actually cross,
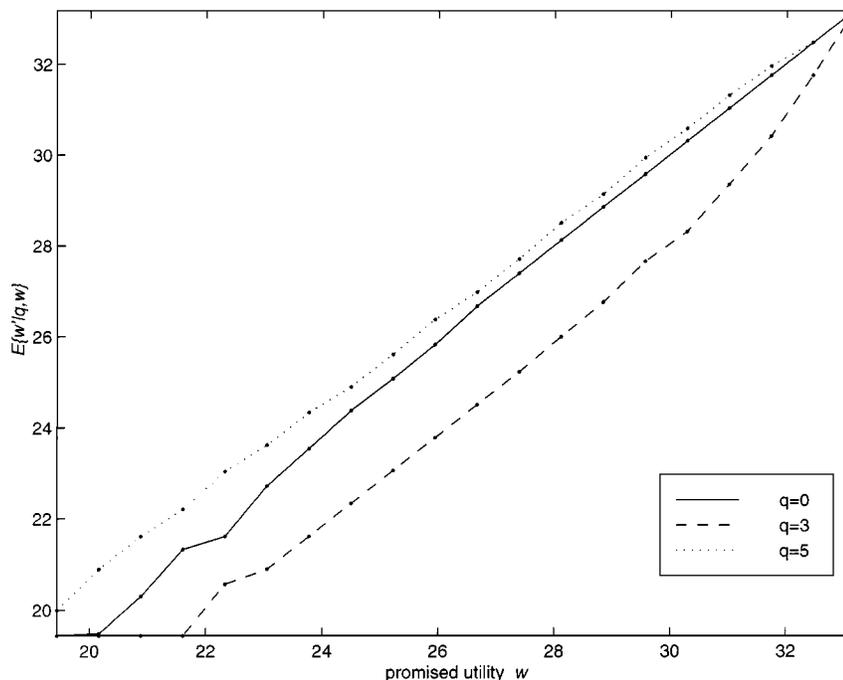
FIGURE 2G. Continuation utility schedule for moral-hazard, no-credit-control, no-default regime.

with the surplus from the pure-debt regime greater than the surplus from the moral-hazard, no-credit-control regime. In particular, consider the case in which $Q = \{0, 2\}$, $A = \{0, 0.24\}$, $B = \{0, 0.6\}$, $\rho = 0.5$, $\gamma = 0.1$, $\delta = 0.5$, and the technology is as presented in Table 4. We present the Pareto frontiers for the two moral-hazard, no-credit-control regimes (Regimes 3a and 3b) and the pure-debt regime in Figure 4. Despite the fact that, in the pure-debt regime, insurance is limited only to default, whereas in the moral-hazard, no-credit-control regimes the principal may explicitly condition gross consumption allocations on outputs, the principal's surplus is actually quite a bit greater in the pure-debt regime than in the moral-hazard, no-credit-control regime without default. If agents were allowed to choose between the pure-debt regime (which allows default) and the moral-hazard, no-credit-control regime without default, there would be an interesting difference across agents' choices by promised utility.

6.4. Estimating Regimes

A researcher armed with sufficient data and computing power, in principal, could estimate underlying parameters and differentiate among the regimes of the model. In this section we describe how to obtain a distribution over observable variables,

TABLE 4. Technology: Mapping from inputs A and B into conditional probabilities $P(q \mid a, b)$

Inputs		Outputs	
B	**A**	$P(q=0 \mid a, b)$	$P(q=2 \mid a, b)$
$b=0$	$a=0$	0.95	0.05
$b=0$	$a=0.25$	0.95	0.05
$b=0.6$	$a=0$	0.05	0.50
$b=0.6$	$a=0.25$	0.05	0.95
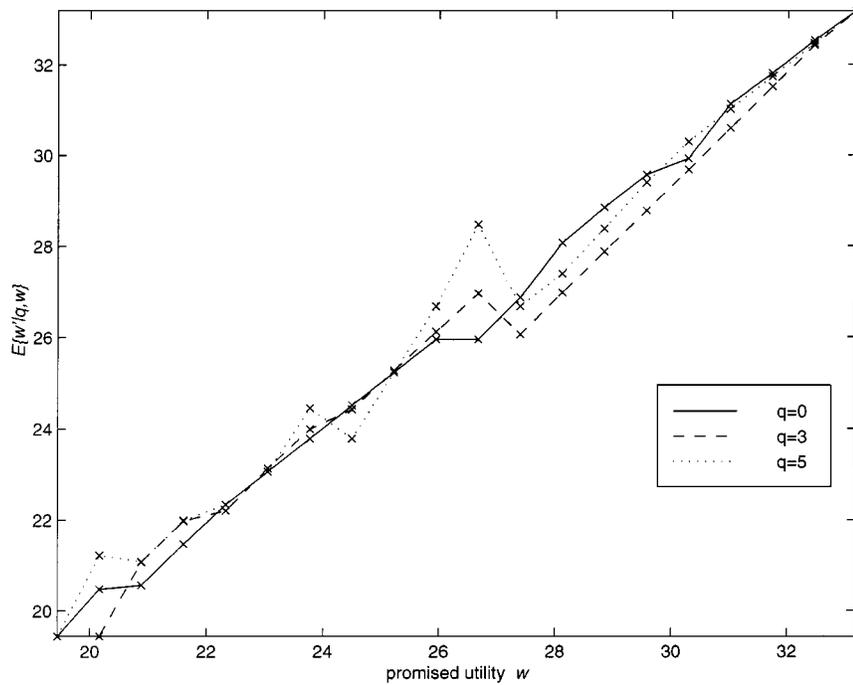
FIGURE 2H. Continuation utility schedule for the pure-debt regime.

such as output, credit, and consumption, from the optimal policies generated by a particular parameter set in a particular regime. This distribution then can be compared against data by computing a goodness-of-fit statistic, which then can be used to choose a parameter set and regime that best describe the data. We computed optimal policies for several combinations of technology and preference parameters in each regime. We then compared the implied distributions over observables to determine if parameters and regimes could be distinguished, given an infinite quantity of data. We find that certain regimes can be rejected. But we also find that certain technology parameters cannot be identified and that, in a few
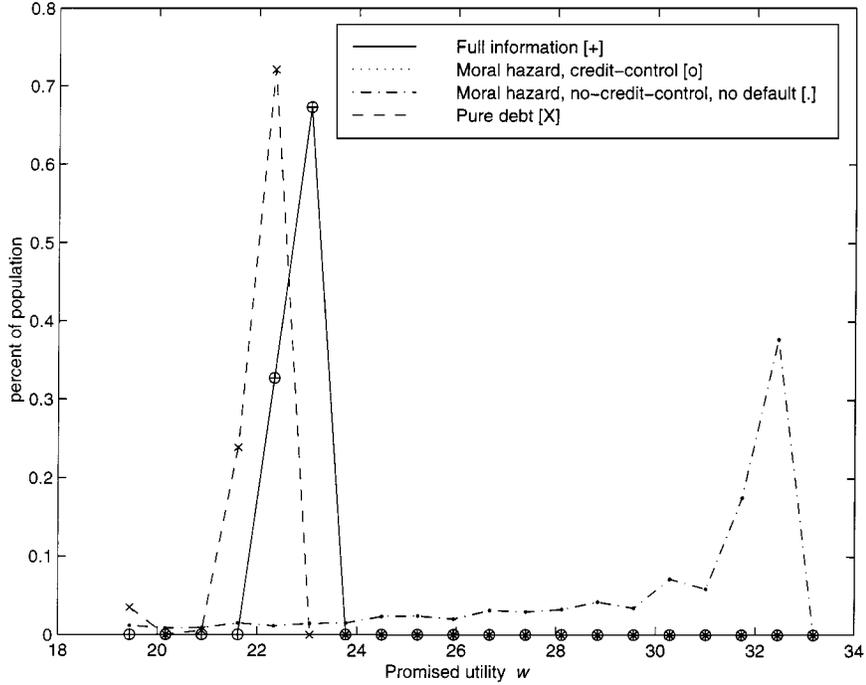
FIGURE 2I. Steady-state (invariant) distribution over W for all regimes.

instances, regimes and preference parameters also are not identified. That two different technologies, with different degrees of risk aversion, when computed under two different regimes, produce the same invariant distribution over observables is in itself an interesting result.

For this discussion we have to augment slightly the notation developed in Section 4, equations (29) and (30). There we defined the proportion of agents at the invariant distribution who are in state s, $H(s)$, and the proportion of agents who, while still in the invariant distribution, move from state s to s', $D(s, s')$. The empirical version of H can be formed from a large cross section, whereas the empirical version of D can be formed from a large panel. For the purposes of this section, we assume that the researcher has access to enough data to be able to compute the population versions of H and D directly, with no error. This simplifying assumption allows us to abstract from the problems of inference and concentrate on identification issues.

We augment this notation in two ways: First, from now on, we restrict the states s to be over observable variables only, which we take to be net consumption c_1, output q, and credit b. That is, we integrate out the unobserved variables effort a and, in the pure-debt regime, the internal interest rate r. Second, we now recognize that H and D are contingent on the choice of regime and parameters. From now on,
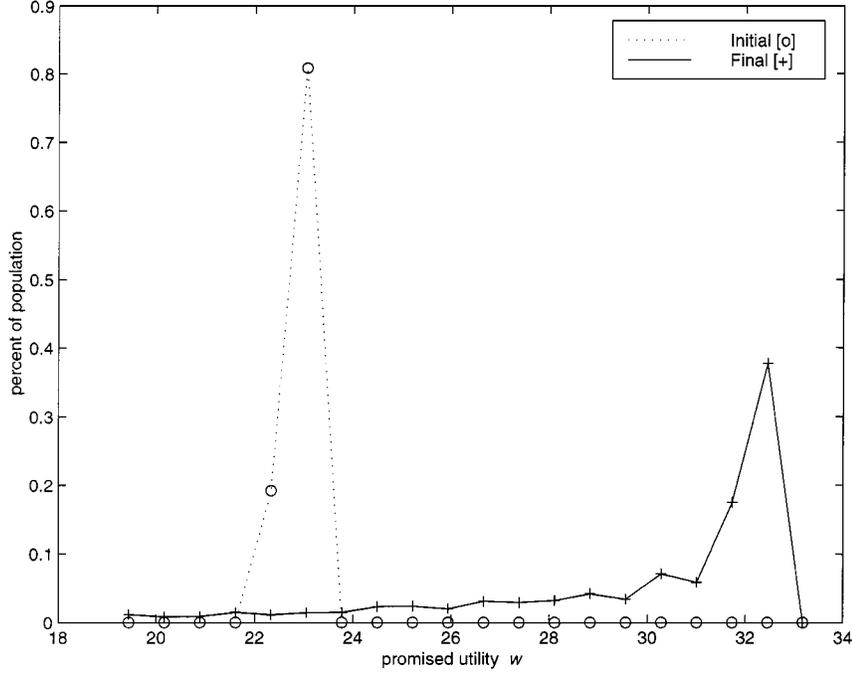
FIGURE 2J. Initial and steady-state distributions for the moral-hazard, no-credit-control, no-default regime.

we write $H(s \mid \boldsymbol{\theta}, \text{reg})$ and $D(s, s' \mid \boldsymbol{\theta}, \text{reg})$, where s is an element of $\boldsymbol{C}_1 \times \boldsymbol{Q} \times \boldsymbol{B}$, $\boldsymbol{\theta}$ is a vector of parameters, and reg is the choice of regime under which the optimal policy was computed.

Now we define a scalar metric that measures how different one distribution is from another. There is a large literature on nonparametric density estimation, which a researcher with a finite sample would have to undertake to find the empirical counterparts to H and D. In keeping with that literature, we use the mean value of the squared differences across the distributions (know as the mean integrated squared error), or what we call the error score.[13] The error scores when comparing two regimes, reg and reg′, and two sets of parameters, θ and θ', are defined, for the cross section and the panel, respectively, as

$$\sigma_1(\boldsymbol{\theta}, \text{reg}, \boldsymbol{\theta}', \text{reg}') \equiv (1/N) \sum_s [H(s \mid \boldsymbol{\theta}, \text{reg}) - H(s \mid \boldsymbol{\theta}', s')]^2, \quad \textbf{(35)}$$

$$\sigma_2(\boldsymbol{\theta}, \text{reg}, \boldsymbol{\theta}', \text{reg}') \equiv (1/N^2) \sum_{s,s'} [D(s, s' \mid \boldsymbol{\theta}, \text{reg}) - D(s, s' \mid \boldsymbol{\theta}', \text{reg}')]^2, \quad \textbf{(36)}$$

where N is the number of values that s can take on, $N = n_{C_1} n_{Q} n_{B}$.
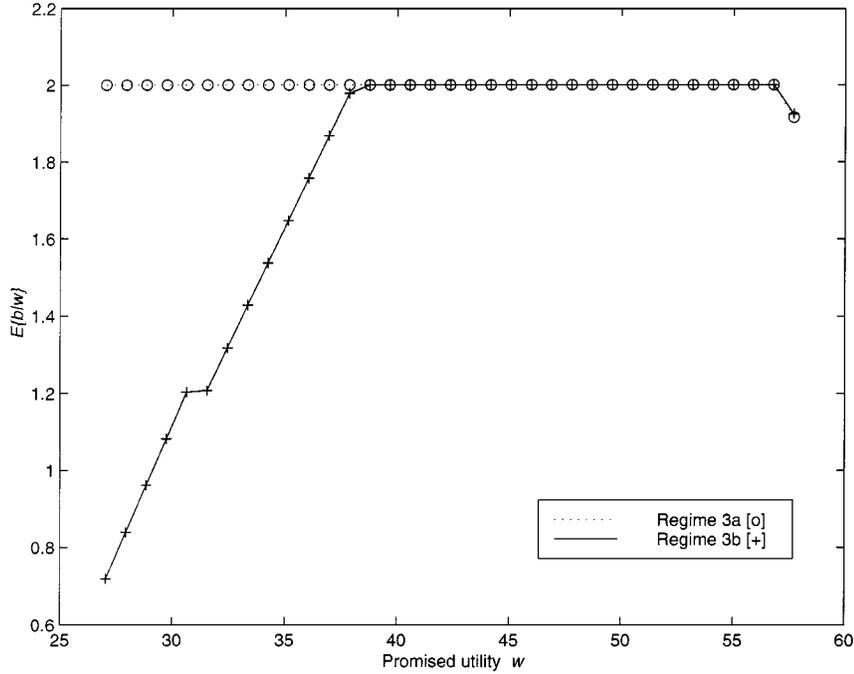
FIGURE 3A. Effect of prohibiting default on assigned credit levels: Solid line is assigned credit under the moral-hazard, no-credit-control regime with default (Regime 3a); the dotted line is the same regime without default (Regime 3b).

We computed results for a fairly general setup with many different technology and preference parameter combinations. In this setup, effort $A = \{0, 0.25\}$, output $Q = \{0, 2\}$, material $B = \{0, 0.6\}$, the discount factor $\beta = 0.8$, and the world price of credit $\rho = 0.5$. We varied the risk-aversion parameter γ across the points $\{0.1, 0.3, 0.5, 0.7, 0.9\}$. The technology $P(q \mid a, b)$ is described in Table 5. We assumed that the probability of the high output at the high effort and material inputs was 0.95, and that the probability of the low output at the low effort and material inputs was also 0.95. The parameter T_1 indexes the probability of the high output when effort is high but material is low, whereas T_2 indexes the probability of the high output when effort is low but material is high.

With this technology, the information measure $I(b, q; (\bar{a} = 0.25, a = 0))$ is increasing (at both outputs) in T_1 when $b = 0$ and decreasing (at both outputs) in T_2 when $b = 0.6$. When T_1 and T_2 are large, there is an information cost to assigning the high material input level (as in the example in Section 6.1).

We varied the parameters (T_1, T_2) across the points $\{0.05, 0.275, 0.5, 0.725, 0.95\}$ each. Thus we computed the optimal solutions to all regimes at 125 different combinations of the parameters (γ, T_1, T_2).

TABLE 5. Technology: Mapping from inputs A and B into conditional probabilities $P(q \mid a, b)$

Inputs		Outputs	
B	A	$P(q = 0 \mid a, b)$	$P(q = 2 \mid a, b)$
$b = 0$	$a = 0$	0.95	0.05
	$a = 1$	$1 - T_1$	T_1
$b = 0.6$	$a = 0$	$1 - T_2$	T_2
	$a = 1$	0.05	0.95
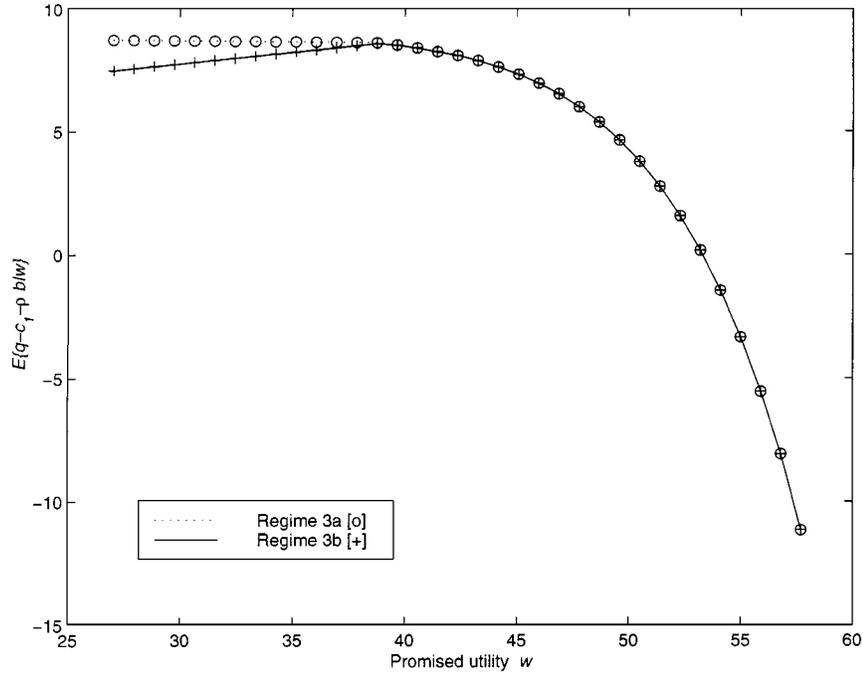
FIGURE 3B. Effect of prohibiting default on the principal's surplus $V(w)$: Solid line is assigned credit under the moral-hazard, no-credit-control regime with default (Regime 3a); the dotted line is the same regime without default (Regime 3b).

Assume that the true values of (γ, T_1, T_2) are (0.5, 0.5, 0.5) (we simply need a baseline for comparison). Now assume further that the researcher has chosen γ and the regime properly. It turns out that T_1 is not identified in several regimes. In Figure 5A, we plot the σ_2 error score (assuming that the researcher has access to data from a large panel) over all values of (T_1, T_2) at the true regime (here, full information) and the true γ. Notice that, along the true value of T_2, the surface is zero for the smaller values of T_1 [the x's mark the true value of (T_1, T_2)]. When
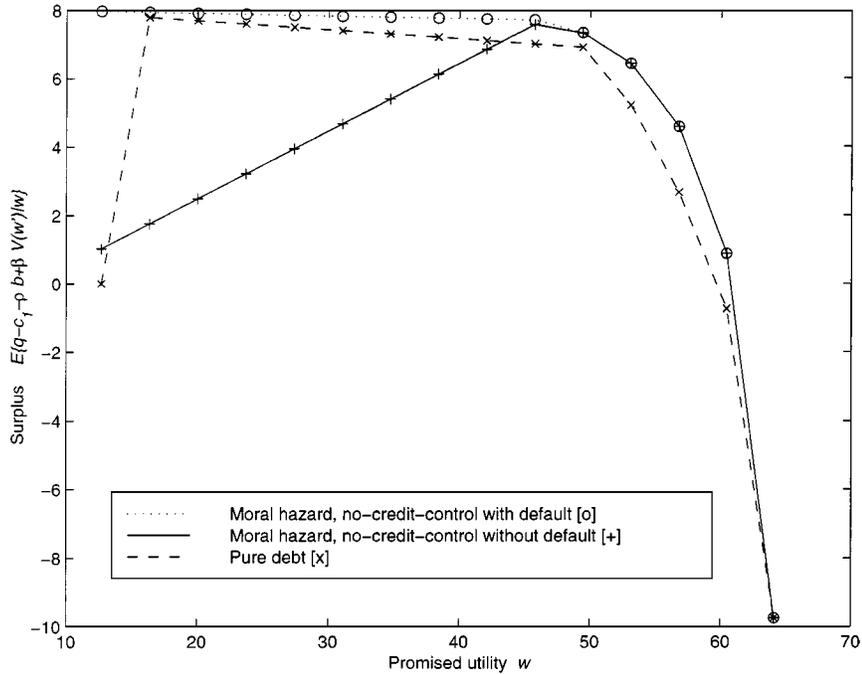
FIGURE 4. Effect of prohibiting default on the Pareto frontiers. Three regimes are presented: the moral-hazard, no-credit-control regimes with and without default, and the pure-debt regime. Notice that the no-credit-control regime with default dominates both the regime without default and the pure-debt regime, but that the pure-debt regime and the no-credit-control regime without default cross.

T_1 is small relative to T_2, the principal assigns the high material input level and then either the high or low effort level (depending on the promised utility of the agent). Because there is no moral-hazard problem in the full-information regime, the principal does not have to worry about deviations in effort or material, and T_1 does not enter the solution at all. This is a common feature, even in regimes that do feature moral hazard. The reason there is that agents contemplate deviating to the low-effort, low-material combination, which has a fixed probability distribution over outputs. When $T_1 = 0.95$, the high material input level is no longer productive, and the principal switches to the low material input. This is precisely the point at which the σ_2 surface is no longer zero—the researcher is only able to pin T_1 down into a range of values.

In Figure 5B, we again plot the error score σ_2 over all values of the technology parameters, but this time the researcher has guessed that the true regime is moral-hazard, credit-control when it is still full-information. Again, the lowest value of σ_2 is quite low, but the values of T_1 and T_2 that produce this minimum are quite wrong. This is unsurprising because the observed distribution D features full insurance in
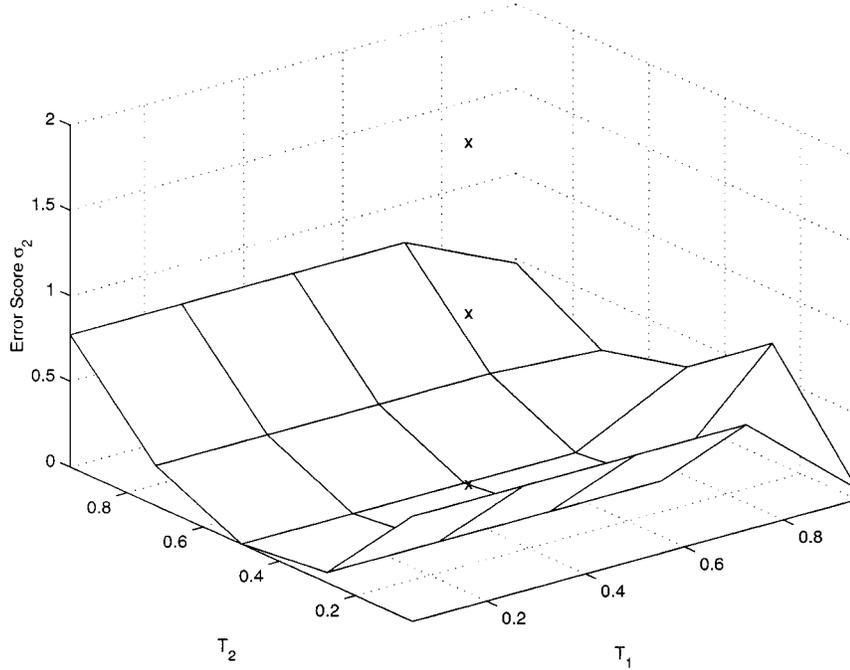
Figure 5A. Values of σ_2 norm at all values of (T_1, T_2) when the researcher knows that $\gamma = 0.5$ and that the true regime is the full-information regime.

consumption (because the true regime is full information), whereas any regime with a binding moral-hazard constraint will feature only partial insurance. In fact, the error score is minimized at the point $(T_1 = 0.05, T_2 = 0.95)$. At this technology, effort is not productive at either level of the material input, but the material input is productive at both levels of effort. Thus the principal assigns the low-effort, high-material-input combination, there is no binding incentive constraint, and the principal provides full insurance. This mimics the full insurance in the observed distribution D.

Next, we compare across regimes and values of γ in a more systematic fashion. For each guess at γ and the true regime, we can find the choice of T_1 and T_2 that minimizes the error score σ_2. Because this choice often will not be unique, we plot the largest value of each that minimizes the error score. By plotting σ_2 across choices of γ and the regime, we can determine, for each choice of regime, which value of (γ, T_1, T_2) minimizes the error score σ_2. At the true value of (γ, T_1, T_2) and in the true regime, $\sigma_2 = 0$. If at any other combination of (γ, T_1, T_2) and regime, $\sigma_2 = 0$, we know that the distribution of observables must coincide exactly, and the true value of (γ, T_1, T_2) and regime are not identified. If σ_2 is nonzero but small at a particular value of (γ, T_1, T_2), with an infinite sample that combination would be rejected, but with a finite sample, a researcher would find it harder to differentiate between the two cases.
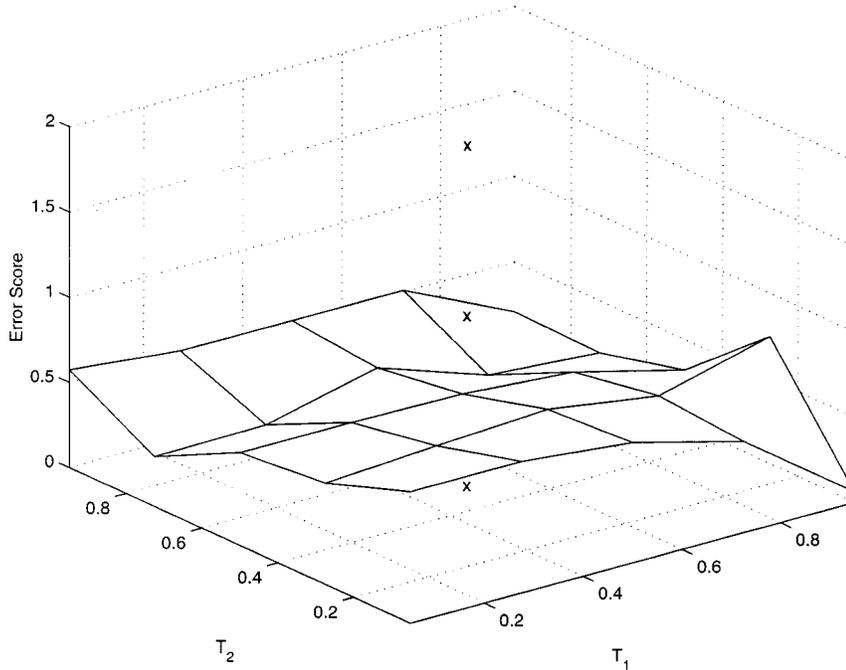
FIGURE 5B. Values of the σ_2 norm at all values of (T_1, T_2) when the researcher knows that $\gamma = 0.5$ but guesses incorrectly that the true regime is moral-hazard, credit-control when in fact it is full-information.

In Figures 5C and 5D, we present the values of σ_2 across different guesses at γ and different guesses at the true regime, when the true regime is, in Figure 5C, the full-information regime, and, in Figure 5D, the moral-hazard, credit-control regime. We also plot the values of T_1 and T_2 that minimize the error score at each guess for γ. (Recall that T_1 often is not identified, and so, we plot the largest value of T_1 that still minimizes the error score.)

Note first that, if the true regime is full-information but the researcher guesses that it is moral-hazard, credit-control, he will choose the correct value of γ but wildly incorrect values of T_1 and T_2. If he guesses that the true regime is the moral-hazard, no-credit-control regime, then he will select $\gamma = 0.7$, and get a very low error score. Indeed, in a finite sample, the researcher would have trouble distinguishing between the true parameters and regime and this alternative.

If the true regime is moral-hazard, credit-control and the researcher guesses that it is full-information, then the researcher will choose the lowest possible value of γ (0.1) and wrong values for T_1 and T_2. If the researcher guesses that the true regime is moral-hazard, no-credit-control, than he will choose the largest value of γ (0.9).
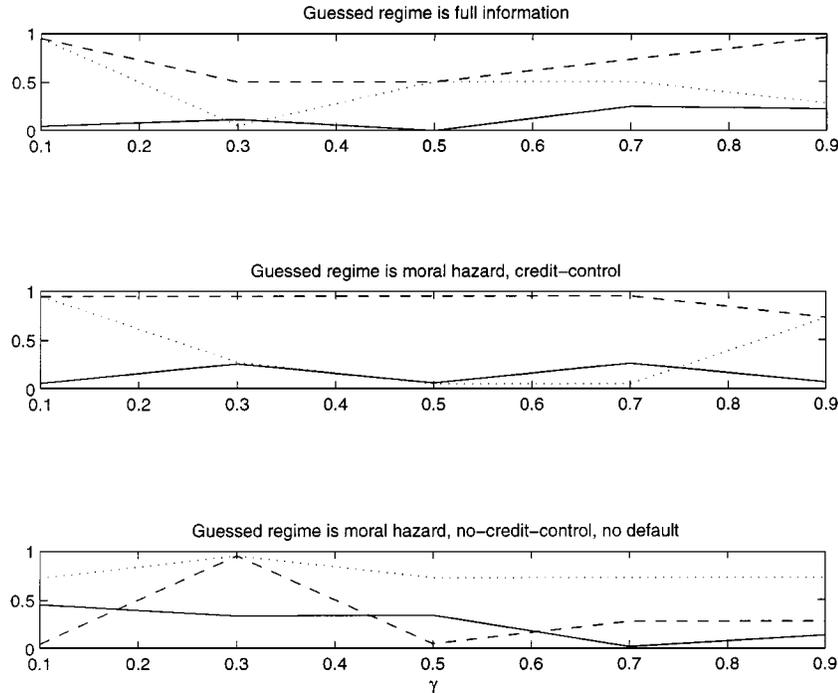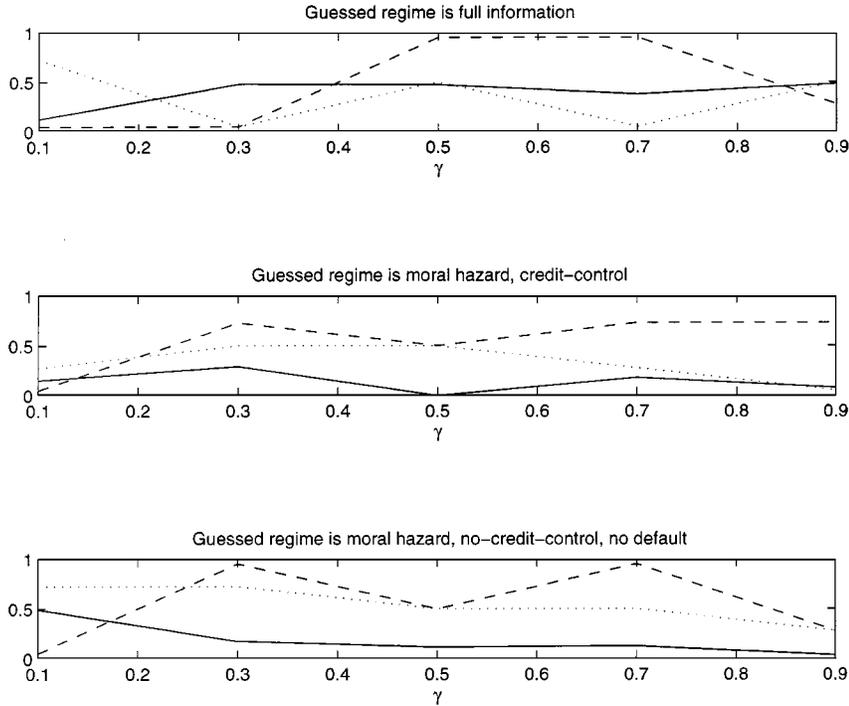
Guessed regime is full information

Guessed regime is moral hazard, credit–control

Guessed regime is moral hazard, no–credit–control, no default

γ

FIGURE 5C. The σ_2 error score across values of γ and regimes when the true regime is full-information: The dotted line gives the error-minimizing value of T_1 at each guess for γ; the dashed line gives the same for T_2.

7. CONCLUSIONS

By casting the credit-provision problem in a principal-agent framework, we are able to consider explicitly how optimal contracts vary with information, control of access to the world lending market, default, and exogenously imposed forms of the credit contract.

We demonstrate that agents would report themselves liquidity-constrained in only one combination of information and control (Regime 2), although in the other regimes credit inputs may be distorted away from their optimal level because of incentive conditions. Agents might assure survey researchers that they were receiving the proper amount of credit, despite the fact that they were in fact assigned a lower-than-optimal level. Agents are also likely to report that they are assigned too much credit, through the credit-diversion effect. Here, they would like to consume their entire gross consumption allocation, even that portion of it which is earmarked for loan repayment.

We find that allowing the agent to default is Pareto improving. Prohibiting default increases the temptation for the agent to divert credit to consumption, and so lower credit levels are assigned. So much so, in fact, that the pure-debt regime (with its

Guessed regime is full information

Guessed regime is moral hazard, credit-control

Guessed regime is moral hazard, no-credit-control, no default

FIGURE 5D. The σ_2 error score across values of γ and regimes when the true regime is moral-hazard, credit-control: The dotted line gives the error-minimizing value of T_1 at each guess for γ; the dashed line gives the same for T_2.

very limited insurance) may Pareto dominate the moral-hazard, no-credit-control regime without default, despite the latter regime's much larger scope for insurance.

Finally, for a general range of technology and preferences, these models produce distinct distributions over observable variables. A researcher with adequate computer time and data should be able to discriminate among parameters and regimes.

NOTES

1. We could also, in principle, decentralize all of these planner's problems, following Prescott and Townsend (1984a, b). Although the last regime, the pure-debt regime, has the flavor of a decentralized scheme, we show that it is nested within the other regimes, which are more clearly planner's problems.

2. As we explain in Section 3.3, although we sometimes allow the agent to default, there is never a free lunch for the agent and principal combined, because the principal always must pay the remainder of the loan.

3. Because the set W contains all of the allowed promised utility points, both promised utility this period, w, and next period, w', must lie in W. With the assumption of an infinite horizon, the set W' must be identical to W. We treat them as two different objects only to clarify whether we mean contemporaneous or future promised utility.

4. We recognize that the distinction between within-period and across-period credit is somewhat artificial, but it does greatly simplify the analytics of this paper. For a strong discussion of the role of intertemporal credit in a two-period principal-agent model that is not inconsistent with the one discussed here, see Bizer and DeMarzo (1996).

5. In practice, we imposed this constraint by using a two-stage algorithm. In the first stage, we solved a subprogram over the choice objects (q, a, b), given an internal price and promised utility (r, w), obtaining the $n_R n_W$ optimal contracts $\pi^1(q, a, b, w' \mid r, w^1)$. These contracts each deliver the promised utility w^1 in W with the interest rate fixed at r in R, and satisfy the Bayes consistency and incentive-compatibility constraints. In the second stage, the choice objects were contracts $\pi^2(r, w^1 \mid w^2)$ over first-stage contracts $\pi^1(q, a, b, w^1 \mid r, w^2)$, given a promised utility of w^2. That is, we fixed the promised utility and interest rate in the first subprogram and then calculated lotteries over those contracts in the second subprogram. From the point of view of an agent, the second subprogram comes first, in the form of a lottery π^2 over contracts π^1.

6. If we imagined that the lender *could* directly control the amount borrowed by the agent, in effect rationing credit, the principal's surplus for a given w would be weakly higher, but the solution still would be different from the moral-hazard, credit-control regime, since contemporaneous net consumption still would be uninsured.

7. The principal thus can assign net consumptions of $[q - rb]_+$ for all Q, R, B combinations. To keep this program computationally comparable with the previous regimes, we set net consumption to exactly this grid. Thus, in all regimes, the set of net consumptions available to the principal is the same.

8. We use s in S from now on to denote all of the contemporary variables in each regime, i.e., everything except promised utilities. For example, in the full-information regime, $S = (C_1 \times Q \times A \times B)$, whereas, in the pure-debt regime, $S = (C_0 \times Q \times A \times B \times R)$. With this notation, all of the features discussed in Section 4 are the same for each regime.

9. Ligon (1997) argues following Rogerson that agents are actually savings constrained intertemporally along an information-constrained optimal path. That is, they would carry consumption forward at rate ρ^* if they were allowed to do so. In this sense, incentive constraints lead to savings constraints and intertemporal distortions much in the spirit of the intratemporal input-financing constraints addressed throughout the paper. Again see Bizer and DiMarzo (1996) for a treatment that suggests that the information-constrained optimum can be supported in an apparently free banking regime with bankruptcy.

10. In practice, because w^* often falls between two of the grid points in W, we weight the two grid points w^*_{-1} and w^*_{+1} that fall on either side of w^* by λ and $1 - \lambda$ so that $\lambda w^*_{-1} + (1 - \lambda)w^*_{+1} = w^*$. The indicator vector i^* then becomes $[0 \ \ 0 \ \ \cdots \ \ 0 \ \ \lambda(1 - \lambda) \ \ \cdots \ \ 0]'$.

11. Numerically, there will be lotteries over promised utility points in the full-information regime if the grid over net consumptions is too coarse.

12. See also Thomas and Worrall (1990), Atkeson and Lucas (1992), and Phelan (1995) for discussions on the distributional dynamics of private information economies.

13. See Scott (1992) for a discussion. Because we are assuming that a researcher has either an infinitely large cross-sectional data set (to compute H directly) or an infinitely large panel data set (to compute D directly), the sample properties of the statistic do not matter.

REFERENCES

Aghion, P. & P. Bolton (1992) Distribution and growth in models of imperfect capital markets. *European Economic Review* 36, 603–611.

Atkeson, A. & R.E. Lucas, Jr. (1992) On efficient distribution with private information. *Review of Economic Studies* 59, 427–453.

Benjamin, D. (1992) Household composition, labor markets, and labor demand: Testing for separation in agricultural household models. *Econometrica* 60, 287–322.

Bizer, D.S. & P.M. DeMarzo (1996) Optimal Incentive Contracts When Agents Can Save, Borrow and Default. Manuscript, Kellog Graduate School of Management, Northwestern University.

Dubey, P., J. Geanokopoulos & M. Shubik (1987) The revelation of information in strategic market games: A critique of rational expectations equilibrium. *Journal of Mathematical Economics* 16, 105–137.

Duffie, D. (1996) Incomplete security markets with infinitely many states: An introduction. *Journal of Mathematical Economics* 26, 1–8.

Evans, D.S. & B. Jovanovic (1989). An estimated model of entrepreneurial choice under liquidity constraints. *Journal of Political Economy* 97, 808–827.

Feder, G., L.J. Lau, J.Y. Lin & X. Luo (1990) The relationship between credit and productivity in Chinese agriculture. *American Journal of Agricultural Economics* 72(4), 1151–1157.

Ligon, E. (1997) Implementation in Village Economies. Giannini Foundation working paper 828, University of California at Berkeley.

Müller, R.A.E. & R.M. Townsend (1997) Mechanism Design and Village Economies: From Credit, to Tenancy, to Cropping Groups. Manuscript, University of Chicago.

Phelan, C. (1995) Repeated moral hazard and one-sided commitment. *Journal of Economic Theory* 66, 488–506.

Phelan, C. & R.M. Townsend (1991) Computing multi-period, information-constrained optima. *Review of Economic Studies* 58, 853–881.

Piketty, T. (1994) Existence of fair allocations in economies with production. *Journal of Public Economics* 55, 391–405.

Prescott, E.C. & R.M. Townsend (1984a) Pareto optima and competitive equilibria with adverse selection and moral hazard. *Econometrica* 52, 21–45.

Prescott, E.C. & R.M. Townsend (1984b) General competitive analysis in an economy with private information. *International Economic Review* 25, 1–20.

Rashid, M. & R.M. Townsend (1993) Targeting Credit and Insurance: Efficiency, Mechanism Design and Program Evaluation. Manuscript, University of Chicago.

Scott, D.W. (1992) *Multivariate Density Estimation*. New York: John Wiley & Sons.

Stiglitz, J.E. & A. Weiss (1981) Credit rationing in markets with imperfect information. *American Economic Review* 71, 393–410.

Thomas, J. & T. Worrall (1990) Income fluctuation and asymmetric information: An example of a repeated principal-agent problem. *Journal of Economic Theory* 51, 367–390.