Borrowing Constraint as an Optimal Contract*

Borys Grochulski† Yuzhe Zhang‡

October 20, 2009

Abstract

We study a continuous-time version of the optimal risk-sharing problem with one-sided commitment. In the optimal contract, the agent's consumption is non-decreasing and depends only on the maximal level of the agent's income realized to date. In the complete-markets implementation of the optimal contract, the Alvarez-Jermann solvency constraints take the form of a simple borrowing constraint familiar from the Bewley-Aiyagari incomplete-markets models. Unlike in the incomplete-markets models, however, the asset buffer stock held by the agent is negatively correlated with income.

1 Introduction

Individuals, firms, and sovereigns alike face limits on the amounts they can borrow. In this paper, we show how borrowing constraints (credit limits) emerge as a key element of an optimal contractual arrangement in a risk-sharing problem subject to limited commitment. In our model, a simple credit limit is precisely what differentiates the optimal risk-sharing arrangement with limited commitment from the optimal risk-sharing arrangement with full commitment. In addition, we show that the optimal risk-sharing with limited commitment implies that the financial buffer stock (assets in excess of the credit limit) is negatively correlated with income. This prediction stands in stark contrast to the implications of the incomplete-markets models of self-insurance, in which the financial buffer stock is positively correlated with income.

Our analysis has two parts. In the first part, we study an optimal contracting problem between a risk-neutral, fully-committed, deep-pocketed principal and a risk-averse, non-committed agent whose stochastic income process is a geometric Brownian motion. Autarky represents the agent's outside option. All information is public. In this setting, we show that under the optimal contract the agent's consumption can be represented as a strictly increasing function of the maximal level of the agent's income realized to date. In the optimal contract, therefore, the

*The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

†Federal Reserve Bank of Richmond, borys.grochulski@rich.frb.org.

‡University of Iowa, yuzhe-zhang@uiowa.edu.

consumption path of the agent is weakly increasing and constant whenever current income is strictly below its to-date maximum but strictly increasing when income achieves a new all-time maximum. At all times, the optimal amount of risk-sharing is less than full. If the agent's preferences exhibit constant relative risk-aversion (CRRA), his optimal consumption is simply given by a constant fraction of the maximal level of his income realized to date.

In the second part, we study a simple trading mechanism that implements efficient allocations. This mechanism consists of two trading accounts that work as follows. The principal makes available to the agent a bank account, in which the agent can save or borrow at a riskless interest rate equal to the principal's and agent's common rate of time preference. The principal also gives the agent access to a hedging account, in which the agent can transfer his income risk to the principal with fair-odds pricing. In the hedging account, the agent faces no limits on the size of the hedge he can take out, i.e., he can transfer 100 percent of his income risk to the principal. In the bank account, however, the agent faces a borrowing limit. The borrowing limit is always greater than zero, i.e., the agent has access to credit. The size of the borrowing limit depends only on the agent's current level of income, and has a simple characterization: it is equal to the total value of the relationship between the principal and the agent. In this mechanism, the agent can freely choose his trading strategy and his consumption process. As well, the agent can default (revert to permanent autarky) at any point in time.

We show that under these conditions, the agent's equilibrium (that is, individually-optimal) trading strategy results in an efficient allocation of consumption. This two-account trading mechanism, thus, implements efficient risk sharing. In equilibrium, the agent never defaults and, despite being able to fully hedge his income risk at any point in time, the agent chooses a hedging strategy that less-than-fully insures his income risk at all times. Also, we show that the financial buffer stock that the agent maintains in equilibrium is negatively correlated with his income.

In an environment otherwise identical to ours but in which the agent can fully commit, any efficient allocation of consumption, clearly, would provide the agent with full insurance. Such allocations can be implemented with a combination of a hedging account with no restrictions on hedging and a riskless bank account with no restrictions on borrowing (other than a never-binding no-Ponzi-scheme condition). Furthermore, it is clear that the trading mechanism in which borrowing limits are absent would not implement any efficient allocation of the limited-commitment environment. This is because over the desired no-default equilibrium strategy the agent would prefer to accumulate debt and default. The limited-commitment optimum, therefore, is implementable if and only if the agent faces the borrowing constraint. In our model, thus, a simple borrowing constraint is precisely the difference between an optimal trading mechanism in the limited-commitment environment (in which default risk is present) and an optimal trading mechanism in the full-commitment environment (in which default risk is absent). Our model pinpoints the role of the borrowing constraint. Namely, this role is to efficiently mitigate the default risk.

Relation to the literature Our paper is not the first one to identify a role for restrictions on borrowing in mitigating the risk of default. In the existing literature, this role has been studied in two contexts.

First, it has been studied in equilibrium models of borrowing and default that exogenously restrict the contract structure to debt contracts (e.g., Eaton and Gersovitz (1981)). In these models, the equilibrium credit limits and other costs to access credit are not necessarily optimal. In contrast, our analysis imposes no restrictions on the structure of the contract. The equilibrium credit limits that we obtain are optimal, i.e., a part of a mechanism supporting the optimal level of risk sharing with limited commitment.

Second, Alvarez and Jermann (2000) study a general equilibrium economy with limited commitment and impose no exogenous restrictions on the structure of the contract. They show that optimal allocations can be implemented via decentralized trade in a complete set of state-contingent claims if agents face solvency constraints that prevent default. The solvency constraints of Alvarez and Jermann (2000) take the form of limits on portfolios of state-contingent claims. Our model is essentially a continuous-time, partial-equilibrium version of the Alvarez-Jermann model with one-sided commitment. Our analysis shows that in this setting the state-contingent solvency constraints collapse to a simple borrowing constraint, which, literally, is a limit on the amount the agent can borrow. Thus, the borrowing constraint that emerges in our version of the Alvarez-Jermann model has the same form as the classic borrowing constraints of the Bewley-type models, which have been widely used in macroeconomics and finance.

This simplification in the form of the endogenous restrictions on borrowing allows us to compare our Alvarez-Jermann-type complete-markets model with the Bewley-type incomplete-markets model. Clearly, the key difference between these two models is the availability of hedging. Our analysis shows an important implication of this difference: the role of financial wealth. We show that, although the agent's total wealth (i.e., his financial wealth plus the present value of his future income) is positively correlated with current income, the correlation between the agent's current income and his financial buffer stock (i.e., his bank account balance in excess of the credit limit) is negative. This feature of our model stands in stark contrast to Bewley-type incomplete-markets models, in which the correlation between the agents' income and financial buffer stock is positive.

In addition to Alvarez and Jermann (2000), our paper is closely related to other papers studying optimal contracts and equilibrium outcomes in environments with commitment frictions. Contributions to this literature include Harris and Holmstrom (1982), Thomas and Worral (1988), Kehoe and Levine (1993), Kocherlakota (1996), Albuquerque and Hopenhayn (2004), Ljungqvist and Sargent (2004), Krueger and Perri (2006), Krueger and Uhlig (2006). Our paper extends the analysis to a continuous-time setting with persistent shocks, which allows for closed-form solutions and a detailed characterization of the dynamics of the optimal contract and its implementation. In particular, the continuous time structure allows us to sign the correlation between the financial buffer stock and income. As we show in Appendix C, however, our method for the characterization of the optimal contract is not specific to our

continuous-time framework.

Krueger and Perri (2006) compare the implications of Alvarez-Jermann-type models and Bewley-type models for the relation between income inequality and consumption inequality, as well as confront these implications with U.S. data. In the discussion of their quantitative results, they note that the correlation between assets and income is negative in the Alvarez-Jermann-type model, but they do not provide analytical results. We prove this result analytically. Also, because we characterize the optimal contract in closed form and show that the borrowing constraint in the implementation corresponds to the principal's maximized profit, we can easily compute the borrowing constraints with no need for the fixed-point iteration procedure used in Alvarez and Jermann (2000). In particular, we show that the optimal borrowing constraint is proportional to the agent's current income when the preferences of the agent satisfy CRRA.

Our paper is also related to several recent studies of optimal contracting problems in continuous time with private information.[1] In particular, our proof of the optimality of the contract is based on the techniques developed in Sannikov (2008). Our analysis suggests that limited-commitment environments are more tractable than private information environments, both in the study of the optimal allocation and its implementation. In particular, in our model we can provide closed-form solutions without value function iteration or having to solve a second-order differential equation.

Organization In Section 2, we present the environment and a general class of contracting problems we study. In Section 3, we characterize the solutions to these problems. In Section 4, we study implementation and provide characterization of optimal policies. In Section 5, we discuss extensions. In Section 6, we sum up our conclusions. Appendix A contains proofs of all lemmas and propositions presented in the text. Appendix B contains a formal verification argument for the optimality of the contract characterized in Section 3. Appendix C extends our analysis to a class of discrete-time models with persistence.

2 The contracting problem

Consider the following dynamic contracting problem in continuous time. There is a risk-neutral principal and a risk-averse agent. Let w be a standard Brownian motion $w = \{w_t, \mathcal{F}_t; 0 \leq t < \infty\}$ on a probability space $(\Omega, \mathcal{F}, \mathbf{P})$. The agent's income process $y = \{y_t, \mathcal{F}_t; 0 \leq t < \infty\}$ is a geometric Brownian motion, i.e., for $t \geq 0$

$$y_t = y_0 \exp(\alpha t + \sigma w_t),$$

where $y_0 \in \mathbb{R}_{++}$, $\alpha \in \mathbb{R}$, and $\sigma \in \mathbb{R}_{++}$.

We assume that the principal and the agent discount at a common rate r. Preferences of

[1] E.g., Demarzo and Sannikov (2007), Biais et al. (2007), Sannikov (2008), Piskorski and Tchistyi (forthcoming), He (2009), Zhang (2009).

the agent are represented by the expected utility function

$$E\left[\int_0^\infty re^{-rt}u(c_t)dt\right],$$

where c_t is the agent's consumption at time t, $u:\mathbb{R}_{++}\to\mathbb{R}$ is a strictly increasing and concave smooth period utility function, and E is the expectations operator. The agent's income process y is publicly observable by both the principal and the agent. Since the agent is risk averse and the principal is risk neutral, there are gains from trade to be realized between the principal and the agent. The principal offers the agent a long-term contract in which he provides the agent with a consumption allocation $c=\{c_t;t\geq 0\}$ in return for the agent's income process y. We require that c be progressively measurable with respect to the filtration $\{\mathcal{F}_t;t\geq 0\}$. The principal's discounted cost of a contract with the agent's consumption c is given by

$$E\left[\int_0^\infty re^{-rt}(c_t-y_t)dt\right].$$

To ensure that the value of the agent's income process is finite, we restrict parameters to satisfy

$$r>\alpha+\frac{\sigma^2}{2}, \tag{1}$$

that is, we assume that the common discount rate is larger than the average growth rate of the income process. We will denote $\alpha+\sigma^2/2$ by μ. Also, for any t, the present value of the agent's future income (i.e., the agent's "human capital," or "human wealth") will be denoted by $P(y_t)$. Using the fact that $E[y_{t+s}|\mathcal{F}_t]=y_t\exp(\mu s)$ for any $t,s>0$, we have that

$$\begin{aligned}P(y_t) &= E\left[\int_0^\infty e^{-rs}y_{t+s}ds|\mathcal{F}_t\right]\\ &= \frac{y_t}{r-\mu}.\end{aligned} \tag{2}$$

The principal can commit to a contract, but the agent cannot. In particular, the agent is always free to walk away from the principal and consume his income. If he does, he loses all future insurance possibilities, i.e., he has to remain in autarky forever. Because income is persistent, the value that the autarky option presents to the agent depends on the current income level. Denoting this value by $V_{aut}(y_t)$, we have

$$V_{aut}(y_t)=E\left[\int_0^\infty re^{-rs}u(y_{t+s})ds|\mathcal{F}_t\right].$$

Let v_t denote the conditional expected utility of the agent under allocation c from time t onwards:

$$v_t=E\left[\int_0^\infty re^{-rs}u(c_{t+s})ds|\mathcal{F}_t\right]. \tag{3}$$

The agent will have no incentive to renege on the contract with the principal if the following *participation constraint*,

$$v_t\geq V_{aut}(y_t),$$

holds at each date t and in every state $\omega \in \Omega$. An allocation that satisfies these participation constraints will be called *enforceable*.

We consider a family of contracting problems indexed by y_0 and $\bar{V}$, where $\bar{V} \geq V_{aut}(y_0)$ is the total utility value that the principal must deliver to the agent. For each pair $(y_0, \bar{V}) \in \Theta \equiv \{(y, v) : y > 0, v \geq V_{aut}(y)\}$, the principal's problem is to design an enforceable allocation c that delivers to the agent utility $\bar{V}$ at a minimum cost $C(y_0, \bar{V})$. That is, the principal's problem at $(y_0, \bar{V})$ is

$$C(y_0, \bar{V}) = \min_{c} \quad E\left[\int_0^\infty re^{-rt}(c_t - y_t)dt\right] \tag{4}$$

$$s.t. \quad v_t \geq V_{aut}(y_t), \text{ all } t \text{ and } \omega, \tag{5}$$

$$v_0 = \bar{V}.$$

Any contract that solves this problem will be called *efficient*. Let $c(y_0, \bar{V})$ denote an efficient contract in the planner's problem at $(y_0, \bar{V})$. For each $(y_0, \bar{V}) \in \Theta$, the contract consumption allocation $c(y_0, \bar{V})$ is a process on $(\Omega, \mathcal{F}, \mathbf{P})$ progressively measurable with respect to the filtration $\{\mathcal{F}_t\}$. Let $\Psi = \{c(y_0, \bar{V}); (y_0, \bar{V}) \in \Theta\}$ denote the family of all efficient contracts. Our task is to characterize the contracts in Ψ.

3 Efficient contracts

This section is devoted to the characterization of efficient contracts. In order to provide economic intuition, we first derive the efficient contracts heuristically and give the main properties of these contracts. The formal verification of optimality is done in subsection 3.5. We start out by considering the contracting problems in which all surplus is given to the principal. That is, for a given y_0, let $\bar{V} = V_{aut}(y_0)$. We postpone the analysis of the problems with $\bar{V} > V_{aut}(y_0)$ until subsection 3.3.

Let us first review the case of full commitment. The optimal contract under full commitment provides full insurance to the agent. Since the principal and the agent discount at the same rate, the optimal full-commitment contract provides the agent with constant consumption $u^{-1}(V_{aut}(y_0))$. Under this contract, the agent's continuation value is constant, i.e., $v_t = V_{aut}(y_0)$ at all dates t and in every state $\omega \in \Omega$.

Under one-sided commitment, this full-insurance contract is not feasible because the agent's autarky value $V_{aut}(y_t)$ will exceed $V_{aut}(y_0)$ when y_t exceeds y_0 for the first time. At this time, the full-insurance contract would violate the agent's participation constraint. As long as y_t does not exceed y_0, however, the participation constraint does not bind. Inside the time interval in which y_t fluctuates below the initial level y_0, thus, the principal's profit maximization problem is the same under both one-sided and full commitment. Therefore, the consumption path that the principal optimally provides to the agent during this time must be constant in the one-sided commitment case, as it is in the case of full commitment.

We now calculate the level of consumption that the principal will optimally provide to the agent during this time interval. A technical difficulty associated with this calculation stems from

the fact that the length of the time interval in which the principal can provide full insurance is zero, i.e., $\inf_t\{t > 0 : y_t > y_0\} = 0$ almost surely.[2] To deal with this difficulty, we first relax the principal's problem by a small amount and construct an optimal contract in the relaxed problem. Then we take a limit of the optimal contract as the size of the relaxation amount goes to zero. Finally, we check that the limiting contract is feasible in the unrelaxed problem.

We fix $\varepsilon > 0$ and drop the agent's participation constraints $v_t \geq V_{aut}(y_t)$ for all $t < \tau_{y_0+\varepsilon}$, where $\tau_{y_0+\varepsilon} = \min_t \{t > 0 : y_t = y_0 + \varepsilon\}$ is the first time when the agent's income reaches $y_0+\varepsilon$. Because ε is strictly positive, $\tau_{y_0+\varepsilon} > 0$ almost surely, and thus the time interval $[0, \tau_{y_0+\varepsilon})$ has non-zero length. In this relaxed problem, there are no participation constraints inside $[0, \tau_{y_0+\varepsilon})$ and thus the principal provides full insurance to the agent over this time interval. At $\tau_{y_0+\varepsilon}$, the principal provides the agent with continuation value

$$v_{\tau_{y_0+\varepsilon}} = V_{aut}(y_0 + \varepsilon), \tag{6}$$

as this value constitutes the minimal departure from the full-commitment contract. This departure is necessary to ensure that the agent's participation constraint $v_t \geq V_{aut}(y_t)$ is satisfied at $\tau_{y_0+\varepsilon}$.

Under the above contract, the agent's utility flow inside the interval $[0, \tau_{y_0+\varepsilon})$ is constant. We will denote this utility flow level by $\bar{u}^\varepsilon(y_0)$. Using this notation and equation (6), the agent's expected utility from this contact can be split into the part before and after time $\tau_{y_0+\varepsilon}$ as follows:

$$v_0 = E\left[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}\bar{u}^\varepsilon(y_0)dt + e^{-r\tau_{y_0+\varepsilon}}V_{aut}(y_0+\varepsilon)\right].$$

Since the value being provided to the agent is $\bar{V} = V_{aut}(y_0)$, the constant utility flow rate $\bar{u}^\varepsilon(y_0)$ must be chosen at a level at which $v_0 = V_{aut}(y_0)$. Thus, $\bar{u}^\varepsilon(y_0)$ satisfies

$$V_{aut}(y_0) = E\left[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}\bar{u}^\varepsilon(y_0)dt + e^{-r\tau_{y_0+\varepsilon}}V_{aut}(y_0+\varepsilon)\right]. \tag{7}$$

Note also that under autarky, the autarky value $V_{aut}(y_0)$ can also be split into the value of the consumption of income received up to the time $\tau_{y_0+\varepsilon}$ and after:

$$V_{aut}(y_0) = E\left[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}u(y_t)dt + e^{-r\tau_{y_0+\varepsilon}}V_{aut}(y_0+\varepsilon)\right]. \tag{8}$$

Comparing (7) and (8) and canceling common terms, we obtain

$$E\left[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}\bar{u}^\varepsilon(y_0)dt\right] = E\left[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}u(y_t)dt\right].$$

Thus, the utility flow rate $\bar{u}^\varepsilon(y_0)$ is the certainty equivalent of the stochastic utility flow rate that the agent receives under autarky over the time interval $[0, \tau_{y_0+\varepsilon})$. For any $\varepsilon > 0$, the

[2]This is because a typical path of Brownian motion has infinite variation and thus crosses y_0 infinitely many times immediately after $t = 0$.

optimal contract in the relaxed problem simply delivers full insurance until $\tau_{y_0+\varepsilon}$, and the minimal continuation value required to satisfy the participation constraint at time $\tau_{y_0+\varepsilon}$.

By taking ε to zero, we now obtain the formula for the certainty equivalent utility flow rate $\bar{u}(y_0)$ in the unrelaxed planner's problem:

$$\begin{aligned}\bar{u}(y_0) &= \lim_{\varepsilon\to 0}\bar{u}^{\varepsilon}(y_0)\\ &= \lim_{\varepsilon\to 0}\frac{E[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}u(y_t)dt]}{E[\int_0^{\tau_{y_0+\varepsilon}} re^{-rt}dt]}\\ &= \lim_{\varepsilon\to 0}\frac{V_{aut}(y_0)-E[e^{-r\tau_{y_0+\varepsilon}}]V_{aut}(y_0+\varepsilon)}{1-E[e^{-r\tau_{y_0+\varepsilon}}]}.\end{aligned}$$

Denote $1-E[e^{-r\tau_{y_0+\varepsilon}}]$ by $g(\varepsilon)$. Then, applying d'Hospital's rule and using $g(0)=0$, we get

$$\begin{aligned}\bar{u}(y_0) &= \lim_{\varepsilon\to 0}\frac{g'(\varepsilon)V_{aut}(y_0+\varepsilon)-(1-g(\varepsilon))V'_{aut}(y_0+\varepsilon)}{g'(\varepsilon)}\\ &= V_{aut}(y_0)-V'_{aut}(y_0)/g'(0).\end{aligned}$$

This expression for the certainty equivalent utility flow rate is intuitive. Note that $g(\varepsilon)\approx g'(0)\varepsilon$ is the amount of discounted time spent before hitting $y_0+\varepsilon$, the income level at which the participation constraint binds. If the constraint never binds, as is the case in the full-commitment case, then the discount factor at the hitting time is zero (i.e., $E[e^{-r\tau_{y_0+\varepsilon}}]=0$) and $g'(0)\approx\infty$, in which case the formula for $\bar{u}(y_0)$ collapses to the full-commitment level $V_{aut}(y_0)$. In the limited-commitment case, the income level at which the participation constraint binds, $y_0+\varepsilon$, is expected to be reached in finite time. At this time, $\tau_{y_0+\varepsilon}$, the agent expects to receive $V'_{aut}(y_0)\varepsilon$ units of extra continuation utility. Thus, the constant flow rate $\bar{u}(y_0)$ over the interval $[0,\tau_{y_0+\varepsilon})$ is reduced below the full-commitment level $V_{aut}(y_0)$ by the amount of the expected gain $V'_{aut}(y_0)\varepsilon$ divided by the expected discounted waiting time $g'(0)\varepsilon$, which is reflected in the above formula for $\bar{u}$.

Using the structure of the agent's income process y, we can characterize the certainty equivalent utility flow rate more closely. Borodin and Salminen (2002, page 622) show that if $y=\{y_t,\mathcal{F}_t;0\le t<\infty\}$ is the geometric Brownian motion, then for any $y\ge y_0$

$$E[e^{-r\tau_y}]=\left(\frac{y_0}{y}\right)^{\kappa}, \tag{9}$$

where

$$\kappa=\left(\sqrt{\alpha^2+2r\sigma^2}-\alpha\right)\sigma^{-2} \tag{10}$$

is a strictly positive constant.[3] Thus, $g'(0)=\kappa/y_0$ and

$$\bar{u}(y_0)=V_{aut}(y_0)-\kappa^{-1}y_0V'_{aut}(y_0).$$

Having described the contract inside the initial time interval $[0,\tau_{y_0+\varepsilon})$, let us now consider the continuation contract starting at time $\tau_{y_0+\varepsilon}$. As we noted before, since the participation

[3] In fact, (1) implies that $\kappa>1$.

constraint binds at $\tau_{y_0+\varepsilon}$, the agent's continuation value at $\tau_{y_0+\varepsilon}$ equals his autarky value $V_{aut}(y_0+\varepsilon)$. The principal's problem of designing a profit-maximizing contract is thus the same at $t=\tau_{y_0+\varepsilon}$ as it was at $t=0$ but with the new initial value $\bar{V}=V_{aut}(y_0+\varepsilon)$ and the new initial income state $y_0+\varepsilon$. The solution to this problem, therefore, must be the same: Consumption is stabilized until the agent's income exceeds $y_0+\varepsilon$ for the first time. The flow utility provided in the meantime, $\bar{u}(y_0+\varepsilon)$, is at the level necessary to deliver value $V_{aut}(y_0+\varepsilon)$ to the agent given that the autarky value will be delivered to the agent as of the future moment when income first exceeds $y_0+\varepsilon$. The same steps we used earlier to calculate $\bar{u}(y_0)$ let us now calculate $\bar{u}(y_0+\varepsilon)=V_{aut}(y_0+\varepsilon)-\kappa^{-1}(y_0+\varepsilon)V'_{aut}(y_0+\varepsilon)$. And so forth.

Repeating this construction for all dates and possible realizations of income paths, we note that under the resulting contract, current utility flow delivered to the agent at any t is determined by the maximum level the income path attained up to time t. Denote this level by

$$m_t=\max_{0\leq s\leq t} y_s.$$

Whenever income y_t is strictly below m_t, the value of m_t remains constant. As we argued earlier, at these times it is efficient to provide the agent with constant consumption flow. Thus, m_t can be used as a state variable sufficient to determine current consumption flow given to the agent under this contract.

In sum, we have argued (so far heuristically) that the optimal contract delivering the value $\bar{V}=V_{aut}(y_0)$ to the agent is given as follows. At any $t\geq 0$, the agent's consumption is given by

$$c_t=u^{-1}(\bar{u}(m_t)), \tag{11}$$

where $\bar{u}:\mathbb{R}_{++}\to\mathbb{R}$ is

$$\bar{u}(y)=V_{aut}(y)-\kappa^{-1}yV'_{aut}(y), \tag{12}$$

and where the constant $\kappa>1$ is given in (10).

Next, we provide some basic properties of this contract. Our heuristic discussion provides simple intuition why this contract is in fact optimal. We postpone the formal verification of this intuition to subsection 3.5. Also, we still need to check that this contract, which we obtained as a limit of optimal contracts from relaxed problems, does satisfy all participation constraints in the unrelaxed problem. We check this later in this section, after we provide basic properties of the contract.

3.1 Increasing consumption paths

We see in (11) that consumption c_t is constant when y_t fluctuates below m_t. Intuitively, this is optimal because the agent's participation constraint is not binding during these times. Under (11), the agent's consumption changes only when y_t attains a new all-time maximum. Intuitively, this adjustment is necessary because the participation constraint of the agent binds at this time. Consistent with this intuition, consumption c_t increases when a new all-time

maximum is realized. To see that this in fact is the case, note that u^{-1} is strictly increasing, and, by the following lemma, so is $\bar{u}$.

Lemma 1 *$\bar{u}$ is strictly increasing and $\bar{u} < u$.*

Proof In Appendix A. ■

The above lemma verifies that $u^{-1}(\bar{u}(\cdot))$ is a strictly increasing function. Since the process m_t is weakly increasing, (11) implies that the agent's consumption paths are weakly increasing for any ω. In particular, the agent's consumption path is constant when $y_t < m_t$ and it increases whenever $y_t = m_t$. It is a standard result in the mathematics of Brownian motion that $y_t < m_t$ at almost all t, and $y_t = m_t$ occurs on a set of Lebesgue measure zero.[4] Thus, consumption c_t is constant at almost all dates t. Moreover, because $\bar{u} < u$, we have that $c_t < m_t$ at all t. In particular, we have $c_0 < y_0$. This means that the contract begins with net payments from the agent to the principal, which is akin to prepayment of an insurance premium.

Example If utility is logarithmic, $u(c) = \log(c)$, then

$$\begin{aligned} V_{aut}(y_t) &= E\left[\int_t^\infty re^{-r(s-t)}\log(y_s)ds|\mathcal{F}_t\right] \\ &= \int_t^\infty re^{-r(s-t)}\left(\log(y_0)+\alpha s+\sigma E[w_{s+t}|\mathcal{F}_t]\right)ds \\ &= \int_t^\infty re^{-r(s-t)}(\log(y_0)+\alpha t+\alpha(s-t)+\sigma w_t)ds \\ &= \log(y_t)\int_t^\infty re^{-r(s-t)}ds+\alpha\int_t^\infty re^{-r(s-t)}(s-t)ds \\ &= \log(y_t)+\frac{\alpha}{r}. \end{aligned}$$

So

$$\begin{aligned} \bar{u}(y) &= V_{aut}(y)-\kappa^{-1}yV'_{aut}(y) \\ &= \log(y)+\frac{\alpha}{r}-\frac{1}{\kappa} \\ &= \log(y)-\frac{\kappa\sigma^2}{2r}, \end{aligned}$$

where the last line follows from an easy-to-verify equality

$$\frac{\alpha}{r}+\frac{\kappa\sigma^2}{2r}=\frac{1}{\kappa}. \tag{13}$$

Applying the inverse utility function $u^{-1}(u) = \exp(u)$, we thus get

$$\begin{aligned} c_t &= u^{-1}(\bar{u}(m_t)) \\ &= m_t\exp\left(-\frac{\kappa\sigma^2}{2r}\right). \end{aligned}$$

[4] See Karatzas and Shreve (1991) for proof.

Thus, with log preferences, the agent consumes a constant fraction of his to-date maximal income m_t. Similar calculations show that the same is true for any constant relative risk-aversion (CRRA) utility function. ■

To understand the structure of the optimal contract a little better, let us discuss how it delivers the initial utility $V_{aut}(y_0)$ to the agent over time. The monotonicity of the consumption paths allows us to see this structure very clearly. For any ω, the agent's utility flow $u(c_t) = \bar{u}(m_t)$ is weakly increasing in t. The total discounted utility of the agent, thus, depends on how fast the utility flow path $\{u(c_t); 0 \leq t < \infty\}$ attains higher and higher levels. Note now that for any $x > y_0$, we have $u(c_t) \geq \bar{u}(x)$ if and only if $m_t \geq x$. Thus,

$$\min\{t : u(c_t) \geq \bar{u}(x)\} = \min\{t : m_t \geq x\} = \min\{t : y_t = x\} = \tau_x. \tag{14}$$

This means that the utility flow $u(c_t)$ attains the level $\bar{u}(x)$ for the first time precisely at τ_x, i.e., when income y_t hits the level x for the first time. Because the distribution of this hitting time is known, we can compute the expected speed with which the utility flow paths $u(c_t)$ increase. More precisely, as we are interested in agent's discounted expected utility, we can compute the expected amount of discounted time that $u(c_t)$ spends above $\bar{u}(x)$, for any $x \geq y_0$. Using (14), we have

$$\begin{aligned} E\left[\int_0^\infty re^{-rt} 1_{[\bar{u}(x),\infty)}(u(c_t))dt\right] &= E\left[\int_{\tau_x}^\infty re^{-rt}dt\right] \\ &= E[e^{-r\tau_x}] \\ &= \left(\frac{y_0}{x}\right)^\kappa, \end{aligned}$$

where $1_{[a,b)}(\cdot)$ is the indicator function of the interval $[a, b)$, and the last line uses (9). Because the total amount of the discounted time is normalized to unity, $1 - (\frac{y_0}{x})^\kappa$ is the expected discounted amount of time that the agent's utility flow spends below the level $\bar{u}(x)$, for any $x > y_0$. Therefore, $\int_{y_0}^\infty \bar{u}(x)d(1-(\frac{y_0}{x})^\kappa)$ represents the total expected discounted utility delivered to the agent in the contract. By the construction of the contract, we know that this value equals $V_{aut}(y_0)$.[5]

It is also worth pointing out that partial insurance is not a transitory phenomenon in our model. At any t, the probability of a consumption path increase in the future is strictly positive.

3.2 Continuation value dynamics

Let us now examine the dynamics of the continuation value process v_t delivered to the agent under the contract c in (11). Because consumption c_s is determined by m_s at all dates $s \geq t$, the knowledge of m_t and y_t is sufficient to determine the continuation value v_t delivered to the agent. In fact, at all dates and states under the optimal contract (11) we can decompose v_t as follows

$$v_t = E\left[\int_t^{\tau_{m_t}} re^{-r(s-t)}\bar{u}(m_t)ds + e^{-r(\tau_{m_t}-t)}V_{aut}(m_t)|\mathcal{F}_t\right],$$

[5]Taking the limit $m \to \infty$ in equation (31) in Appendix A, we can confirm that $V_{aut}(y_0) = -\int_{y_0}^\infty \bar{u}(x)d(\frac{y}{x})^\kappa$, which means that the contract indeed delivers $V_{aut}(y_0)$.

where $\tau_{m_t} = \min_s \{s \geq t : y_s = m_t\}$ is the first time when y_t returns to its to-date maximum m_t. From the above we have that

$$v_t = (1 - E[e^{-r(\tau_{m_t}-t)}|\mathcal{F}_t])\bar{u}(m_t) + E[e^{-r(\tau_{m_t}-t)}|\mathcal{F}_t]V_{aut}(m_t), \tag{15}$$

which means that v_t is a weighted average of $\bar{u}(m_t)$ and $V_{aut}(m_t)$. From (9), we know that

$$E\left[e^{-r(\tau_{m_t}-t)}|\mathcal{F}_t\right] = \left(\frac{y_t}{m_t}\right)^{\kappa}.$$

We thus have that $v_t = V(y_t, m_t)$ where

$$V(y,m) = \left(1 - \left(\frac{y}{m}\right)^{\kappa}\right)\bar{u}(m) + \left(\frac{y}{m}\right)^{\kappa} V_{aut}(m), \text{ for any } m \geq y > 0. \tag{16}$$

The sufficiency of the pair (y, m) to determine the continuation allocation (and therefore the value to the agent and the cost to the principal) is a remarkable feature of the optimal contract. In particular, when $y_t = m_t$, the contract shows what Kocherlakota (1996) and Ljungqvist and Sargent (2004) describe as amnesia: history does not matter, i.e., the continuation contract is the same for all paths of past income $\{y_s; 0 \leq s < t\}$.

Lemma 2 *The function V satisfies*

(i) $0 < V_y(y,m) \leq V'_{aut}(y)$ *with equality only if* $y = m$*;*

(ii) $V_y(y,m)$ *is strictly increasing in* y*;*

(iii) $0 \leq V_m(y,m)$ *with equality only if* $y = m$*.*

Proof In Appendix A. ■

The above lemma provides a lot of information about the dynamics of the agent's continuation value process v_t under the optimal contract c.

As we have seen in the previous subsection, the optimal contract (11) provides constant consumption at almost all dates t. However, the continuation value under (11), v_t, fluctuates at all t. This is because the continuation value depends on the distance between y_t and m_t, which fluctuates continuously. The larger this distance, the longer the expected waiting time for the next permanent increase in consumption. Thus, v_t is positively correlated with y_t at all times.

This correlation measures the degree of insurance against innovations in income that the optimal contract provides to the agent. Let us define full insurance against income innovations at time t as $dv_t/dy_t = 0$, no insurance against income innovations at t as $dv_t/dy_t = V'_{aut}(y_t)$, and partial insurance as $0 < dv_t/dy_t < V'_{aut}(y_t)$.[6] Then, the first conclusion in the above lemma tells us that the optimal contract never provides full insurance, and provides no insurance if and

[6]Note that the optimal contract under full commitment provides full insurance against the innovations at all times, while the autarky allocation provides no insurance against innovations at all times.

only when $y_t = m_t$. Thus, at almost all times, the contract provides partial insurance against income innovations.

The partial insurance property is intuitive. When a negative innovation in y_t occurs (i.e., y_t goes down), v_t suffers because the expected waiting time until the next permanent consumption hike (i.e., when y_{t+s} achieves $y_t + \varepsilon$) lengthens. So v_t responds negatively to drops in y_t. But upon any such drop in y_t, $V_{aut}(y_t)$ suffers *even more* because not only the same waiting time lengthens (i.e., when $V_{aut}(y_{t+s})$ climbs up to $V_{aut}(y_t + \varepsilon)$) but also temporary consumption drops, as $c_t = y_t$ under autarky, while it does not drop under the optimal contract allocation c in (11).

This difference between the responses of v_t and $V_{aut}(y_t)$ to the innovations in y_t shrinks as y_t closes on m_t, because the expected duration of smoothed consumption under the optimal contract decreases as y_t approaches m_t. Thus, as the second property in the above lemma demonstrates, the degree of insurance is monotone in the distance between m_t and y_t. The farther away y_t is from its to-date maximum m_t, the smaller the effect of an income innovation on the expected time until the next consumption hike, and so the more stable the continuation value under the optimal contract. Therefore, the farther away from the boundary of consumption adjustment an innovation in income takes place, the more fully it is insured.

The third property in Lemma 2, $V_m \geq 0$, is intuitive. Fix some two paths of past income $\{y_s^1; 0 \leq s \leq t\}$ and $\{y_s^2; 0 \leq s \leq t\}$ such that $y_t^1 = y_t^2$ but $m_t^1 > m_t^2$. Consider the continuation value v_t^i that the optimal contract delivers to the agent under past income history $\{y_s^i; 0 \leq s \leq t\}$ for $i = 1, 2$. Because $\bar{u}$ is strictly increasing, we have $u(c_t^1) = \bar{u}(m_t^1) > \bar{u}(m_t^2) = u(c_t^2)$, i.e., the agent's utility flow at t is larger under the income history $\{y_s^1; 0 \leq s \leq t\}$. The same remains true at all dates $s \in [t, \tau_{m_t^1})$, i.e., as long as the state m_s remains below m_t^1. At date $\tau_{m_t^1}$, however, the continuation value of the agent will be the same, $V_{aut}(m_t^1)$, independently of the past income history (amnesia). Thus, with the income history $\{y_s^1; 0 \leq s \leq t\}$, the agent receives a higher utility flow relative to the income history $\{y_s^2; 0 \leq s \leq t\}$ during the time interval $[t, \tau_{m_t^1})$, and the same continuation value from time $\tau_{m_t^1}$ onward along every income path.[7] Thus, $v_t^1 > v_t^2$, which means that, keeping current income y_t fixed, the continuation value delivered to the agent by the optimal contract is strictly increasing in m_t.

Finally, it follows as a simple corollary of Lemma 2 that the contract defined in (11) is enforceable (sustainable), i.e., that $v_t \geq V_{aut}(y_t)$ at all dates and states. In fact, we have directly from our construction of the contract that if $y_t = m_t$, then $v_t = V(y_t, y_t) = V_{aut}(y_t)$. For $y_t < m_t$, Lemma 2(iii) implies that $V(y_t, m_t) > V(y_t, y_t)$, and so $v_t > V_{aut}(y_t)$.

3.3 Optimal contract when $\bar{V} > V_{aut}(y_0)$

When $\bar{V} > V_{aut}(y_0)$, we can obtain the optimal contract from continuation of the optimal contract that starts at $\bar{V} = V_{aut}(y_0)$, as this continuation must be optimal (for otherwise the contract c would not be optimal in the first place). For this case, it is enough to modify the

[7] Also, the expectation over continuation paths is the same under both past income histories because $y_t^1 = y_t^2$ and income is a Markov process.

initial condition of the state variable. Let $\bar{m}_0$ be defined by

$$V(y_0, \bar{m}_0) = \bar{V}.$$

Because, by Lemma 2, $V(y, m)$ is strictly increasing in m, a unique solution $\bar{m}_0$ to the above equation exists for any $\bar{V} \geq V_{aut}(y_0)$. At any $t \geq 0$, let the agent's consumption be given by

$$c_t = u^{-1}(\bar{u}(\bar{m}_t)), \tag{17}$$

where $\bar{m}_t = \max\{m_t, \bar{m}_0\}$. Note in particular that when $\bar{V} = V_{aut}(y_0)$, we have $\bar{m}_0 = y_0$.

For any y, let us denote the inverse of $V(y, \cdot)$ by $M(y, \cdot)$. In this notation, $\bar{m}_0 = M(y_0, \bar{V})$ and for any pair $(y_0, \bar{V})$ the optimal contract is given by $c_t = u^{-1}(\bar{u}(\max\{m_t, M(y_0, \bar{V})\}))$. Our heuristic derivation makes it clear that this contract is indeed optimal for any pair $(y_0, \bar{V})$. We formally verify this in subsection 3.5.

3.4 Cost to the principal

In this subsection, we study the properties of the principal's continuation cost under the contract c in (11), expressed as a function of the state $(y_t, \bar{m}_t)$. Denoting the principal's continuation cost process by Z_t, we have that, at all t, $Z_t = Z(y_t, \bar{m}_t)$, where

$$Z(y, m) = \left(1 - \left(\frac{y}{m}\right)^{\kappa}\right) u^{-1}(\bar{u}(m)) + \left(\frac{y}{m}\right)^{\kappa} \int_m^{\infty} u^{-1}(\bar{u}(x)) d\left(1 - \left(\frac{m}{x}\right)^{\kappa}\right) - rP(y). \tag{18}$$

The first term on the right-hand side of this expression represents the expected present value of the constant consumption flow the agent receives for as long as his income does not exceed m. The second term is the expected present value of consumption delivered to the agent from the moment his income hits m onward.[8] The third term, $rP(y) = ry/(r - \mu)$, is the present value of the agent's future income (in flow units).

This expression allows us to study the properties of the process Z_t through the properties of the function Z.

Lemma 3 *The function Z satisfies:*

(i) $Z_y(y, m) > -\frac{r}{r-\mu}$ and is strictly increasing in y with $\lim_{y \to 0} Z_y(y, m) = -\frac{r}{r-\mu}$;

(ii) $Z_m(y, m) \geq 0$, with equality only if $y = m$;

(iii) For a given m, if $\frac{dZ(y,y)}{dy}|_{y=m} \leq 0$, then $Z_y(y, m) \leq 0$ for all $y \leq m$.

Proof In Appendix A. ■

Recall that in the case of full commitment, under optimal contract, the agent's consumption is constant. The principal's cost to deliver a continuation value v to an agent with current income y is given by

$$C^f(y, v) = u^{-1}(v) - rP(y), \tag{19}$$

[8] Recall that when $y = m$, then $1 - (\frac{m}{x})^{\kappa}$ is the expected discounted time that the agent's consumption flow spends below the level $u^{-1}(\bar{u}(x))$ for $x \geq m$.

where $u^{-1}(v)$ is the constant consumption level needed to deliver promised utility v. We see that, under full commitment, the present value of the agent's future consumption, $u^{-1}(v)$, is always constant. Because

$$\begin{aligned} C_y^f(y,v) &= -rP'(y) \\ &= -\frac{r}{r-\mu}, \end{aligned}$$

the principal's cost negatively co-varies one-for-one with the present value of the agent's future income.

In the one-sided commitment case, Lemma 3(i) shows that the principal's cost does not respond as strongly to the changes in income as it does under full commitment. This is because the present value of the agent's future consumption is not constant under one-sided commitment. In fact, it is strictly increasing in current income. Thus, when y increases, the drop in the principal's continuation cost that is due to the increase in $P(y)$ is offset by an increase in the present value of the agent's future consumption.

In general, this offsetting effect can be strong enough to cause the overall cost to increase when income increases. Intuitively, this can happen if the agent's utility function approaches risk neutrality at high consumption levels. When income is low, the agent is risk averse, and the principal's profit is high. But when income is high, the agent is almost risk neutral, thus the principal's profit can be lower. Part (iii) of Lemma 3 provides a sufficient condition for this not to be the case (the principal's profit is increasing in agent's income when $dZ(y,y)/dy|_{y=m} \leq 0$). It is easy to check that this sufficient condition is met when the agent's preferences satisfy CRRA (see also the Example below).

Part (ii) of Lemma 3 has a simple intuition. Since higher promised utility to the agent incurs more cost to the principal, $Z_m \geq 0$ follows directly from $V_m \geq 0$.

The total surplus from the relationship between the principal and the agent can be defined as $-C(y, V_{aut}(y))/r$. This quantity represents the amount of profit (measured as a stock) that the principal can generate by efficiently providing to the agent whose income is y the autarky value $V_{aut}(y)$. Under the optimal contract, we have $C(y, V_{aut}(y)) = Z(y,y)$. Since the autarkic contract (i.e., $c_t = y_t$ for all t) generates zero surplus, the surplus from the optimal contract, which is different from autarky under agent risk aversion, is strictly positive. Thus, $-Z(y,y)/r > 0$ for all y.

Example (continued) If utility is logarithmic, $u(c) = \log(c)$, then, after substituting $c_t = m_t \exp\left(-\kappa\sigma^2/(2r)\right)$ in (18) and simplifying, we get

$$Z(y,m) = m \exp\left(-\frac{\kappa\sigma^2}{2r}\right)\left(1 + \frac{1}{\kappa - 1}\left(\frac{y}{m}\right)^{\kappa}\right) - y\frac{r}{r-\mu}. \tag{20}$$

The total contract surplus is given by

$$\begin{aligned} -\frac{Z(y,y)}{r} &= -\left(y\exp\left(-\frac{\kappa\sigma^2}{2r}\right)\left(1+\frac{1}{\kappa-1}\right)\frac{1}{r}-y\frac{1}{r-\mu}\right) \\ &= -\left(\exp\left(-\frac{\kappa\sigma^2}{2r}\right)\left(1+\frac{\kappa\sigma^2}{2r}\right)-1\right)\frac{1}{r-\mu}y, \end{aligned}$$

where the second line uses (13). Let

$$\psi = \exp\left(-\frac{\kappa\sigma^2}{2r}\right)\left(1+\frac{\kappa\sigma^2}{2r}\right). \tag{21}$$

Because $\exp(x) > 1+x$ for any $x > 0$, we have $0 < \psi < 1$. We can now write

$$-\frac{Z(y,y)}{r} = (1-\psi)\frac{1}{r-\mu}y, \tag{22}$$

which shows that the total contract surplus is strictly positive and proportional to y. Equivalently, the total contract surplus is a constant fraction of the agent's human wealth $P(y) = y/(r-\mu)$. Similar calculations show that the same is true for any CRRA utility function. Also, one can show that with CRRA preferences the contract surplus is strictly increasing in the coefficient of relative risk aversion. ■

3.5 Formal verification of optimality

Our heuristic derivation of the optimal contract c in (11) contains the intuition for why it in fact is optimal. Because the principal is risk-neutral, it is efficient to provide the agent with full insurance. Permanent full insurance, however, is not feasible, because of the agent's participation constraints. The contract c in (11) is a minimal deviation from permanent full insurance that satisfies the participation constraints. This heuristic argument must, however, be verified formally. That is, we need to show that the principal's cost under this contract, i.e., $Z(y_0, M(y_0, \bar{V}))$, in fact equals the minimum cost $C(y_0, \bar{V})$ of providing the agent whose initial income level is y_0 with utility $\bar{V}$. We provide this formal verification argument in Appendix B.

4 Implementation

In this section, we show that the optimal contract can be implemented in an arrangement in which the principal, instead of offering a long-term contract that swaps the income process y for a consumption process c, offers to the agent a pair of trading accounts: a simple bank account with a credit line and a hedging account in which the agent can take out insurance against his income risk. The final allocation is then determined by the agent through his trading activity in the two accounts. This mechanism is significantly less restrictive than the "direct" mechanism in which the principal controls the agent's consumption directly. The agent has much more control over his consumption than under the direct long-term swap contract. Yet, we show that

under an appropriate choice of the initial bank account balance and the credit line process, the final allocation is the same as the optimum c.

The trading mechanism we consider here is closely related to the one that agents face in the complete-markets economy with solvency constraints of Alvarez and Jermann (2000).[9] The partial-equilibrium implementation result that we present is a restricted version of the general-equilibrium decentralization result obtained in Alvarez and Jermann (2000). Tractability is an advantage of our continuous-time model. We are able to characterize the solvency constraints in detail. In particular, we show that they take in our model a simple form of a borrowing constraint. Also, we show in our model that although the agent's total (that is, financial and human) wealth is positively correlated with income, the correlation between the agent's financial wealth and his income is negative.

We start this section by studying the implementation of the optimal allocation under the assumption of full commitment. The borrowing constraint turns out to be the only difference between the implementing mechanisms in the full-commitment environment and the one-sided commitment model. In this implementation, therefore, the borrowing constraint is the implication of the limited-commitment friction.

4.1 The agent's problem

The principal offers the agent two accounts: a simple bank account with a credit line and a hedging account in which the agent can hedge his income risk at fair odds. The interest rate in the bank account is equal to the common rate of time preference. We will show that under an appropriate choice of the credit line, this trading mechanism is optimal. By optimality we mean that the agent trading freely in these two accounts will choose individually the same consumption process as that provided by the optimal contract, and thus will achieve the maximum utility at the minimum cost to the principal.

Let A_t denote the agent's bank account balance process. The asset A_t is risk-free and pays a net interest r. The principal imposes a lower bound process $B_t \leq 0$ on the agent's bank account balance, i.e., A_t must satisfy

$$A_t \geq B_t, \text{ at all } t. \tag{23}$$

Because $B_t \leq 0$, the quantity B_t represents the size of the credit line that the principal makes available to the agent within the bank account.

The fair-odds hedging account works as follows. The agent chooses a hedging position at all t. If the agent's hedging position is β_t at t, then at time $t + dt$, the hedging account pays off $\beta_t(w_{t+dt} - w_t)$ to the agent. Thus, the agent can use this account to hedge (bet against) the innovations dw_t to his income process. The payoff flow to the agent can be positive or negative, but its expected value is zero for any choice of the hedging position process β_t

[9] See also Krueger and Perri (2006) and Krueger and Uhlig (2006). Albanesi and Sleet (2006) consider a similar implementation in an economy with full enforcement, private information, and taxes.

because $E[\beta_t dw_t] = E[\beta_t(w_{t+dt} - w_t)] = 0$. Thus, the fair-odds price of the hedging asset is zero.[10]

The agent chooses his consumption process c_t, his bank account balance process A_t, and his hedging position process β_t subject to the credit limit (23) and the flow budget constraint

$$dA_t = (rA_t + y_t - c_t)dt + \beta_t dw_t, \text{ at all } t. \tag{24}$$

The agent's objective is to maximize the utility of consumption. We will refer to any utility-maximizing trading strategy as an equilibrium of the two-account problem.

4.2 Implementation of the full-commitment optimum

In this subsection, we discuss how this two-account trading mechanism could, under the conditions of full commitment, be used to implement the optimal allocation from the full-commitment long-term contracting problem.

Recall that the optimal allocation under full commitment provides full insurance to the agent with a constant consumption process $c_t = u^{-1}(\bar{V})$ at all dates and states. Thus, the agent's continuation value process is constant $v_t = \bar{V}$, and the principal's cost function is $C^f(y_0, \bar{V}) = u^{-1}(\bar{V}) - rP(y_0)$.

Suppose now that the principal offers the two trading accounts to the agent with some initial bank account balance A_0 and some credit limit process B_t. The objective of the principal is to choose these conditions in such a way as to deliver to the agent lifetime utility $\bar{V}$ at the minimum cost $C^f(y_0, \bar{V})$.

Proposition 1 *Suppose the principal offers the initial balance*

$$A_0 = \frac{C^f(y_0, \bar{V})}{r}$$

and no borrowing limit except for the requirement that $\lim_{t\to\infty} E[e^{-rt}A_t] \geq 0$. *Then,*

$$\begin{aligned} c_t &= u^{-1}(\bar{V}), \\ A_t &= \frac{C^f(y_t, \bar{V})}{r}, \\ \beta_t &= -\frac{\sigma y_t}{r - \mu} \end{aligned} \tag{25}$$

are an optimal consumption and trading strategy for the agent.

[10] We could alternatively formulate the hedging account in terms of payoffs contingent on the innovations dy_t, instead of dw_t. Because the income process y is not a martingale (unless $\mu = 0$), in the alternative formulation the principal would have to charge the agent a premium flow of $E[\beta_t dy_t] = \beta_t \mu y_t dt$ so as to break even. The formulation we adopt is simpler because $E[\beta_t dw_t] = 0$ for any β_t, and so the fair-odds premium is zero. These two formulations are otherwise equivalent: the properties of the optimal credit limit and agent's equilibrium consumption, wealth, and hedging ratio processes are the same in both cases.

Proof In Appendix A. ■

Note that the borrowing constraint here is as loose as possible. All that the principal requires is that the agent does not play a Ponzi scheme on the principal. The borrowing constraint could be tighter, as long as it is never binding. For example, it could be $B_t = C^f(y_t, \bar{V})/r$, in which case $A_t = B_t$ at all t but the constraint never binds.

Recall that the cost function C^f given in (19) can be written as $C^f(y, v) = u^{-1}(v) - rP(y)$, where $P(y)$ is agent's human wealth given in (2). In the above implementation of the full-insurance allocation, the agent's bank account balance process is $A_t = C^f(y_t, \bar{V})/r$. We thus have that

$$A_t + P(y_t) = \frac{u^{-1}(\bar{V})}{r}, \text{ at all } t.$$

This means that the agent's total wealth—the sum of his financial wealth A_t and his human wealth $P(y_t)$—is constant at all dates and states. All shocks to the agent's human wealth are perfectly absorbed by his financial wealth, i.e., $dP(y_t) = -dA_t$ at all dates and states. In particular, the agent takes out in equilibrium a hedging position β_t that perfectly offsets the innovations to his human wealth. In fact, by Ito's lemma,

$$dP(y_t) = \frac{1}{r-\mu}\mu y_t dt + \frac{1}{r-\mu}\sigma y_t dw_t,$$

and so the innovations to $P(y_t)$ are represented by $\frac{1}{r-\mu}\sigma y_t dw_t$. By (25), these innovations are equal to $-\beta_t dw_t$. Thus, the agent is fully hedged at all dates and states.

We can see here that access to hedging is necessary for implementation of the optimum. If instead the agent had access to the bank account only, similar to the standard incomplete-markets model, then his flow budget constraint would be

$$dA_t = (rA_t - c_t + y_t)dt. \tag{26}$$

In this case, the total wealth $A_t + P(y_t)$ could not be constant because $dP(y_t)$ contains a volatility term while dA_t does not. Thus, since the present value of future consumption must equal the agent's total wealth, consumption cannot be perfectly smoothed without the hedging account.

We also see that the role of the bank account here, where the agent has access to full hedging opportunities, is much different from what it would be in the incomplete-markets setting, in which no hedging is available to the agent, i.e., the agent uses his bank account to self-insure his income shocks. In particular, with hedging available to him, the agent has no precautionary motive for saving. As a result, his financial wealth A_t co-varies negatively with human wealth $P(y_t)$, and thus also with the current income y_t.

4.3 Implementation of the one-sided optimum

In this section, we return to the case of one-sided commitment, in which the agent can stop participating (default) at any time. At any point in time, thus, the agent can permanently exit the contract and stay in autarky forever. If he does, he loses the credit line and access to

hedging with the principal, but can consume his own income $\{y_{t+s}; s \geq 0\}$ without having to repay his debt $-A_t$, if any, to the principal.

We show that the optimal consumption process c_t given in (11), combined with some trading strategy $\{\beta_t; t \geq 0\}$ and asset level process $\{A_t; t \geq 0\}$ solve the agent's utility maximization problem.

Proposition 2 *Suppose the borrowing constraint is given by*

$$B_t = \frac{C(y_t, V_{aut}(y_t))}{r}, \tag{27}$$

and the agent's initial assets are

$$A_0 = \frac{C(y_0, \bar{V})}{r}. \tag{28}$$

Then, under the above trading mechanism, the agent's optimal consumption and trading strategy are as follows:

$$\begin{aligned} c_t &= u^{-1}(\bar{u}(\bar{m}_t)), \\ A_t &= \frac{Z(y_t, \bar{m}_t)}{r}, \\ \beta_t &= \frac{Z_y(y_t, \bar{m}_t)\sigma y_t}{r}, \end{aligned}$$

where $\bar{m}_t = \max\{\max_{0 \leq s \leq t} y_t, \bar{m}_0\}$, $\bar{u}$ *is given in (12),* Z *is given in (18), and* $\bar{m}_0 = M(y_0, \bar{V})$.

Proof In Appendix A. ■

The credit limit in (27) is not just a no-Ponzi condition that we had in the case of the full-commitment implementation. It is now a binding constraint analogous to the solvency constraints of Alvarez and Jermann (2000). In our model, these constraints can be succinctly characterized by a single function of current income alone. (Note that $B_t = B(y_t)$, where $B(y) = C(y, V_{aut}(y))/r$.) Thus, in our continuous-time model, the state-contingent solvency constraints of Alvarez and Jermann (2000) take the simple form of a borrowing constraint that depends only on current income. Our framework allows for a clear characterization of this borrowing constraint. In fact, we see from (27) that, at any t, the agent's credit limit (the negative of the borrowing constraint value) equals the total surplus from the relationship between the principal and the agent. The initial asset level (28) determines how this surplus is divided between the principal and the agent. If $A_0 = B_0$, the whole surplus goes to the principal. If $A_0 = 0$, the whole surplus goes to the agent.

We now discuss the main properties of the agent's optimal trading strategy in this implementation.

Total wealth When income y_t fluctuates below $\bar{m}_t$, the agent, similar to the case of full-insurance allocation under full commitment, stabilizes his current consumption. However, unlike in the full insurance case, the present value of the agent's future consumption is not stabilized.

Under both full- and one-sided commitment, the present value of future consumption is identically equal to the agent's total wealth $W_t = A_t + P(y_t)$. Stabilization of the present value of future consumption, thus, is equivalent to stabilization of the agent's total wealth, which was the case in the full-commitment optimum, and is not the case in the limited-commitment optimum. Formally, we have the following lemma.

Lemma 4 *At all t, the agent's total wealth W_t is strictly positively correlated with y_t.*

Proof In Appendix A. ■

Financial wealth and the buffer stock When income y_t fluctuates below $\bar{m}_t$, the agent's bank account balance A_t remains strictly above the credit limit B_t. When y_t hits $\bar{m}_t$, the bank account balance A_t reaches the credit limit B_t. This can happen in one of two ways. If $Z_y(y, m) \leq 0$ for all $y \leq m$, which is true under, e.g., CRRA preferences, the correlation between income y_t and bank account balance A_t is always negative. As y_t approaches $\bar{m}_t$ from below, the bank account balance A_t approaches B_t from above, and hits it when y_t hits $\bar{m}_t$. If $Z_y(y, m) > 0$ for y close to m, the correlation between y_t and the bank account balance A_t is positive at y_t close to $\bar{m}_t$. When y_t hits $\bar{m}_t$, the credit limit is reached, i.e., $A_t = B_t$. This, however, is achieved not by A_t dropping down to B_t, because A_t increases when y_t increases as it reaches $\bar{m}_t$. Rather, B_t increases with y_t faster than A_t does and, as a result, B_t hits A_t from below to achieve $A_t = B_t$. In either case, the buffer stock of assets held by the agent, $S_t = A_t - B_t$, is negatively correlated with income y_t, as shown in the following lemma.

Lemma 5 *At all t, the agent's buffer stock $S_t = A_t - B_t$ is negatively correlated with income y_t. This correlation is zero if and only if $S_t = 0$.*

Proof In Appendix A. ■

The negative correlation between financial wealth and income is intuitive in a setting in which shocks to the present value of lifetime income are at least partially insured. Take disability as an example of such a shock. When an agent becomes disabled, his human wealth drops drastically. The value of her disability insurance policy soars, however, as now this policy is "in the money." With disability insurance in place, therefore, the occurrence of disability decreases human wealth and increases financial wealth.

The above lemma clearly shows the difference between the role that financial wealth fulfills in our model, in which the agent can hedge, and the role it fulfills in the self-insurance models, in which agents do not have access to hedging. Under self-insurance, which takes place, e.g., in Bewley-type incomplete-markets models, the financial wealth's function is to buffer off the income shocks. Financial wealth is accumulated when income increases and decumulated when income decreases. The buffer stock of financial assets is thus positively correlated with income in these models. In our model, which is a version of the complete-markets model of Alvarez and Jermann (2000), the agent can obtain insurance via hedging. The role of financial wealth is not to buffer off the income shocks but rather to buffer off the losses from hedging the

shocks to current income (and thus also the human wealth). Hedging generates losses precisely when current income increases. Thus, the financial buffer stock decreases when current income increases. In effect, the correlation between current income and the financial buffer stock is negative in our model.

The negative correlation between the financial buffer stock (or, more directly, assets in the natural case of $Z_y(y,m) \leq 0$) implies that the degree of consumption insurance provided to the agent in the complete-markets model with borrowing constraints we consider here is larger than that provided to the agent in the standard incomplete-markets model, in the following sense. In both models, the expected present value of the agent's future consumption is equal to his total wealth W_t, and total wealth is the sum of human wealth and financial wealth: $W_t = P(y_t) + A_t$, at all t. For a small $dt > 0$, denote the change in the total wealth process W_t by $dW_t = W_{t+dt} - W_t$, with similar notation for the corresponding changes in A_t and $P(y_t)$. Because of the identity between total wealth and the sum of human and financial wealth, we have that

$$cov_t(dW_t, dP(y_t)) = var_t(dP(y_t)) + cov_t(dA_t, dP(y_t)),$$

i.e., the conditional covariance between the change of the present value of future consumption and the change of human wealth equals the sum of the conditional variance of the change in human wealth and the conditional covariance between the changes in financial and human wealth. In our complete-markets model, the covariance between the change in present value of future consumption and the change in human wealth, $cov_t(dW_t, dP(y_t))$, is smaller than the variance of the change of human wealth precisely because $cov_t(dA_t, dP(y_t)) < 0$ at all t.[11] With incomplete markets, the covariance between the change of the present value of consumption and the change in human wealth is equal to the variance of the change in human wealth because $cov_t(dA_t, dP(y_t))$ is zero in that model.[12]

In the full-commitment case, as we have seen in the previous subsection, the agent did not face a limit on how low his financial wealth could become along any realization of the income path. In the one-sided commitment case, the agent faces the constraint $S_t \geq 0$. The next lemma describes the dynamics of the buffer stock as it hits its lower bound.

Lemma 6 *When $S_t = 0$, the volatility of S_t is zero and the drift of S_t is strictly positive.*

Proof In Appendix A. ■

Because the buffer stock S_t remains non-negative at all t, it is clear that its volatility must be zero when $S_t = 0$, for otherwise S_t would become strictly negative with probability one.

[11] In particular, as we have seen in the previous subsection, in the case of full-commitment there are no borrowing constraints and, in effect, the covariance between the changes in the present value of consumption and human wealth is brought down all the way to zero, i.e., $cov_t(dW_t, dP(y_t)) = 0$.

[12] Clearly, $cov_t(dA_t, dP(y_t))$ is zero also under the autarkic allocation, as financial wealth is identically zero in autarky. This covariance is thus the same in the incomplete-markets model and in autarky. The incomplete-markets model, however, delivers some consumption smoothing through self-insurance. To see this, note that $E_t[dA_t dP(y_t)]$ is negative in a typical incomplete-markets model, while zero in autarky.

That the drift of S_t at zero is strictly positive means that zero is not an absorbing barrier for S_t, but rather is reflective.

Outside of the special case in which $Z_y(y, \bar{m})|_{y=\bar{m}} = 0$, the volatility of the agent's financial wealth is non-zero even when the agent reaches his credit limit. This fact highlights some important properties of the optimal allocation. Recall from Section 3 that the agent's continuation value v_t is equal to his autarky value $V_{aut}(y_t)$ when $y_t = \bar{m}_t$, which is a consequence of the binding participation constraint at times t such that $y_t = \bar{m}_t$. As we see in Lemma 2(i), locally at these time points, the principal provides no insurance against the agent's income innovations, so the agent's continuation value locally behaves as if the agent were in autarky. But the optimal contract supports these local properties of the continuation value process with an allocation that is very different from autarky (and more efficient). This fact is easy to see through the properties of the implementing mechanism. In autarky, agent's financial wealth is identically equal to zero, and thus so is its volatility. In the above implementation of the optimal allocation, this volatility is generically non-zero even when $y_t = \bar{m}_t$, which makes it clear that the optimal allocation is not similar to autarky even locally at times when the agent's continuation value is. Also, the non-zero volatility of financial wealth at $y_t = \bar{m}_t$ shows that the optimal allocation is not locally the same as the optimal self-insurance allocation, because under self-insurance the agent's financial wealth's volatility would be zero (cf. (26)).

Hedging Recall that the hedging position β_t represents the sensitivity of the agent's financial wealth A_t to shocks dw_t, and the corresponding sensitivity of the human wealth $P(y_t)$ is represented by $y_t\sigma/(r-\mu)$. Let us thus define the agent's hedging ratio h_t^A as

$$h_t^A = \frac{-\beta_t}{\frac{y_t\sigma}{r-\mu}}.$$

As shown in the previous subsection, under full commitment the agent's hedging ratio equals one at all dates and states. In this way, agent's total wealth W_t is stabilized. Thus, $h_t^A = 1$ defines full hedging. In the limited-commitment optimum, the hedging ratio h_t^A is strictly less than one at all t. This reflects the result of Lemma 4, i.e., that financial wealth A_t does not respond to shocks dw_t as strongly as human wealth $P(y_t)$ does. In effect, the agent's total wealth W_t is positively correlated with current income y_t. For a fixed $\bar{m}_t$, the agent's hedging ratio is decreasing in y_t. When income is very low (y_t is close to zero), the hedging ratio is close to one, i.e., the agent is nearly fully hedged. When y_t increases, the hedging ratio drops. The closer current income y_t approaches $\bar{m}_t$, the less hedged the agent becomes. Let us also define the (implicit) hedging ratio h_t^B that characterizes the borrowing limit process B_t as the negative of the volatility of B_t divided by the volatility of human wealth $P(y_t)$. The agent's hedging ratio h_t^A is bounded below by h_t^B. These results are summed up in the following lemma.

Lemma 7 *At all t, the hedging ratio h_t^A satisfies*

$$h_t^B \leq h_t^A < 1,$$

with equality if and only if $y_t = \bar{m}_t$. *Also,* h_t^A *is decreasing in* y_t *and approaches one as* y_t *approaches zero.*

Proof In Appendix A. ■

Example (continued) With log utility, using (20) we obtain

$$\beta_t = \left(\frac{1}{r} \exp\left(-\frac{\kappa\sigma^2}{2r} \right) \frac{\kappa}{\kappa - 1} \left(\frac{y_t}{\bar{m}_t} \right)^{\kappa - 1} - \frac{1}{r - \mu} \right) \sigma y_t,$$

and so the agent's hedging ratio is given by

$$\begin{aligned} h_t^A &= 1 - \frac{r - \mu}{r} \exp\left(-\frac{\kappa\sigma^2}{2r} \right) \frac{\kappa}{\kappa - 1} \left(\frac{y_t}{\bar{m}_t} \right)^{\kappa - 1} \\ &= 1 - \psi \left(\frac{y_t}{\bar{m}_t} \right)^{\kappa - 1}, \end{aligned}$$

where $\psi \in (0, 1)$ is given in (21). For a fixed m, this hedging ratio is strictly decreasing in y and approaching one (full hedging) from below as $y \to 0$. When $m = y$, the hedging ratio reduces to $1 - \psi$, i.e., is a strictly positive constant, the same for all y. This constant is a lower bound on the hedging ratio. Whenever $y < m$, the hedging ratio is strictly larger than this lower bound. Similar calculations show that the same is true for any CRRA utility function. When the agent's preferences do not exhibit CRRA, the lower bound on the hedging ratio will generally depend on y.

Turning to the credit limit process B_t, using (22), we have $B_t = B(y_t)$ where $B(y)$ is given by

$$\begin{aligned} B(y) &= \frac{Z(y, y)}{r} \\ &= -\frac{1 - \psi}{r - \mu} y. \end{aligned}$$

The dynamics of B_t, therefore, are

$$dB_t = -\frac{1 - \psi}{r - \mu} dy_t = -\frac{1 - \psi}{r - \mu} \mu y_t dt - \frac{1 - \psi}{r - \mu} \sigma y_t dw_t.$$

Comparing this to the dynamics of A_t we have

$$\begin{aligned} -\frac{1 - \psi}{r - \mu} \sigma y_t &= \left(\exp\left(-\frac{\kappa\sigma^2}{2r} \right) \frac{\kappa}{\kappa - 1} \frac{1}{r} - \frac{1}{r - \mu} \right) \sigma y_t \\ &> \left(\exp\left(-\frac{\kappa\sigma^2}{2r} \right) \frac{\kappa}{\kappa - 1} \frac{1}{r} \left(\frac{y_t}{\bar{m}_t} \right)^{\kappa - 1} - \frac{1}{r - \mu} \right) \sigma y_t \\ &= \beta_t, \end{aligned}$$

i.e., the volatility of B_t is less negative than the volatility of A_t, for $y_t < \bar{m}_t$ strictly. The hedging ratio of the credit limit h_t^B is equal to the lower bound on the agent's hedging ratio,

i.e., $h_t^B = 1 - \psi$. We can also directly see in this example that the volatility of the buffer stock $S_t = A_t - B_t$, given by

$$\left(1 - \left(\frac{y_t}{\bar{m}_t}\right)^{\kappa-1}\right) \exp\left(-\frac{\kappa\sigma^2}{2r}\right) \frac{\kappa}{\kappa-1} \frac{1}{r} \sigma y_t,$$

is zero when $A_t = B_t$, i.e., when $y_t = \bar{m}_t$.

The drift of A_t is

$$rA_t + y_t - c_t = Z_t + y_t - \exp\left(-\frac{\kappa\sigma^2}{2r}\right) m_t,$$

which, using (20), is equal to

$$\bar{m}_t \exp\left(-\frac{\kappa\sigma^2}{2r}\right) \frac{1}{\kappa-1} \left(\frac{y_t}{\bar{m}_t}\right)^{\kappa} - y_t \frac{\mu}{r-\mu}.$$

Subtracting the drift of B_t and simplifying, we get that the drift of the buffer stock S_t is

$$y_t \exp\left(-\frac{\kappa\sigma^2}{2r}\right) \frac{\kappa}{\kappa-1} \left[\frac{1}{\kappa}\left(\frac{y_t}{\bar{m}_t}\right)^{\kappa-1} - \frac{\mu}{r}\right].$$

This formula shows explicitly that when $S_t = 0$, the drift of the buffer stock is strictly positive, as $\frac{1}{\kappa} - \frac{\mu}{r} = \frac{(\kappa-1)\sigma^2}{2r} > 0$, which follows from (13). ■

4.4 Other credit limit processes

In general, the credit limit process B_t given in (27) is not unique. In the class of fully history-dependent processes, it could be replaced with a tighter or looser credit limit process. For instance, any process $\tilde{B}_t$ such that $C(y_t, V_{aut}(\bar{m}_t))/r \geq \tilde{B}_t \geq B_t$ would preserve the implementation result. In particular, one could take $\tilde{B}_t = C(y_t, V_{aut}(\bar{m}_t))/r$. In this case, the agent's assets would always sit on the credit limit. Essentially, however, the impact of the limit process $\tilde{B}_t$ would be the same as that of the process B_t, as the agent's optimal trading and consumption strategy is the same.

The credit limit process B_t in (27) depends only on the agent's current income, i.e., $B_t = B(y_t)$, where $B(y_t) = C(y_t, V_{aut}(y_t))/r$. In particular, B_t is independent of the current asset position A_t, or the the history of past income. It turns out that this credit limit is the unique optimal credit limit process in the class of process B_t that can be given as a continuous function of y_t. This means that any relaxation of the borrowing constraint function $B(\cdot)$ would allow the agent to modify his strategy, default, and obtain higher utility while leaving the principal with an inefficiently high cost, which would invalidate the implementation result; while any tightening of the borrowing constraint function $B(\cdot)$ would restrict the budget set of the agent more than necessary to satisfy the participation constraints, which would lead to a loss of efficiency.

Lemma 8 *Let B_t be an optimal credit limit process. If $B_t = B(y_t)$ for some continuous function $B(\cdot)$, then $B(y) = C(y, V_{aut}(y))/r$.*

Proof In Appendix A. ■

5 Extensions

As we show in (11), the optimal consumption process in our model is given as a fixed, increasing function of the to-date maximal income. This property of the optimal contract is not specific to our continuous-time model with geometric Brownian motion income process. In Appendix C, we show how our analytical characterization in (11) can be extended to a class of discrete-time models in which the agent's income process is a first-order Markov chain whose transition matrix satisfies a weak first-order stochastic dominance condition. As well, this characterization extends to other continuous-time models. In particular, it holds for any continuous-path income process under which the derivative of $E[e^{-r\tau_\varepsilon}]$ is continuous at zero. As long as this condition holds, the certainty equivalent utility flow rate, $\bar{u}(y_0)$, can be approximated by the certainty equivalents from relaxed problems, $\bar{u}^\varepsilon(y_0)$, and our method of characterizing the optimal contract remains valid.[13]

In those more general models, even though the optimal contract can be characterized as in (11) and implemented in a trading mechanism with some form of Alvarez-Jermann solvency constraints, the properties of the optimal allocation and the implementing equilibrium analogous to those we discuss in lemmas 2 to 7 would be much more difficult to obtain and present. For this reason, we study in this paper the optimal risk-sharing problem with one-sided commitment in a continuous-time model with a geometric Brownian motion structure for the income process.

In our implementation, as long as the borrowing constraints are enforced, there is no restriction on hedging, i.e., the agent can choose the process $\{\beta_t; t \geq 0\}$ with no size restrictions. This critically depends on the continuity of the time paths of the bank balance process $\{A_t; t \geq 0\}$. In contrast, in a discrete-time model, state-contingent solvency constraints necessarily imply a restriction on the agent's hedging position at all times. Without such a restriction, the agent could take out a hedging position that would pay off enormous amounts in some states of nature and require delivery of enormous amounts in other states. The agent could use this extreme gambling strategy to obtain a profitable deviation from the desired equilibrium strategy, thus invalidating the implementation result. In this deviation, which is often called a double-deviation strategy, the agent combines the extreme gamble against a subset of the possible states of nature with default in the states in which his gamble does not pay off. The upside value of this plan can be made very large while the downside risk is bounded by the value of autarky that the agent obtains when he defaults. This makes the double-deviation strategy profitable. In our model, double deviations cannot provide a large upside potential to the agent because income sample paths are continuous. Intuitively, this means that in our model, in which the income shocks are small (and frequent), the agent cannot take a hedging position large enough to obtain a large gamble, which is necessary to make the double-deviation plan profitable. Equivalently, the agent cannot generate a discontinuous time path for his bank account balance, which means he cannot violate his borrowing constraint by a meaningful amount. The continuous time path property of the income process is important here. In a continuous-time

[13] For example, if the log of the income process is an Ornstein-Unlenbeck mean-reverting process, the formula for the derivative of $E[e^{-r\tau_\varepsilon}]$ can be obtained from Borodin and Salminen (2002, page 524, formula 2.0.1).

model with discontinuous income paths (for example, with discrete income shocks arriving as a Poisson process), individual shocks could be large (at points of time path discontinuity) and gambles with large upside potential are possible. As a result, asset paths could have discrete jumps. In such environments, restrictions on the size of hedging would again become necessary.

In addition, our results can be easily extended to the case of unequal time preference rates between the principal and the agent. If the principal is more patient than the agent, the agent's consumption path drifts down deterministically when participation constraints are not binding and increases when participation constraints bind. Thus the optimal consumption path is non-monotonic, and the stationary distribution of consumption may be non-degenerate. Non-monotonic consumption paths also arise in optimal risk-sharing problems with multi-sided commitment frictions. We conjecture that our method of characterizing the optimal contract and its implementation, which we provide in this paper for a continuous-time model with one-sided commitment, can be extended to study optimal contracts in continuous-time models of optimal risk-sharing with multi-sided commitment frictions.

6 Conclusion

It has long been recognized in the literature that borrowing constraints are an important tool to mitigate the risk of borrower default. Existing models, however, deliver optimal borrowing constraints in the form of complicated restrictions on portfolios of state-contingent assets. Our model shows how simple borrowing constraints—literally, limits on the amount that an agent can borrow—emerge as *the* implication of limited borrower commitment in a continuous-time model of optimal risk sharing.

Our model is highly tractable. We show how optimal allocation can be expressed as a fixed function of a single state variable. In the implementation, the optimal credit limit is simply equal to the total value of the surplus generated by the relationship between the principal and the agent. We closely characterize the dynamics of wealth and the hedging position held by the agent. Financial wealth is negatively correlated with current income and with the degree of hedging. Because of persistence in the income process, a negative shock to current income decreases the agent's human wealth. Because of hedging, it increases the agent's financial wealth. With more financial wealth, the risk of the agent's default decreases, so the agent can better hedge subsequent shocks to his income and human wealth. Our model shows clearly that the role of financial wealth is drastically different in complete- and incomplete-markets models. With incomplete markets, financial assets buffer off income shocks. With complete markets, financial assets buffer off losses generated by the agent's optimal hedging position, although not completely, due to the risk of default.

Appendix A

Proof of Lemma 1

We begin by noting that the autarky value function V_{aut} can be expressed as

$$V_{aut}(y_0) = \int_0^\infty u(y) f(y_0, y) dy, \tag{29}$$

where $f(y_0, y)$ is the density of the expected discounted amount of time that the income process starting from y_0 spends at each level $y \in (0, \infty)$. From Borodin and Salminen (2002, page 132), we know that

$$f(y_0, y) = \begin{cases} \frac{r}{\sigma^2\kappa + \alpha} \frac{1}{y} \left(\frac{y_0}{y}\right)^\kappa & \text{for} \quad y \geq y_0, \\ \frac{r}{\sigma^2\kappa + \alpha} \frac{1}{y} \left(\frac{y}{y_0}\right)^{\kappa + 2\alpha\sigma^{-2}} & \text{for} \quad y \leq y_0, \end{cases}$$

where κ is the constant given in (10). Differentiating (29) yields

$$V'_{aut}(y_0) = \frac{r}{\alpha + \kappa\sigma^2} \left[\kappa y_0^{\kappa-1} \int_{y_0}^\infty u(y) y^{-\kappa-1} dy + (-\kappa - 2\alpha\sigma^{-2}) y_0^{-\kappa-2\alpha\sigma^{-2}-1} \int_0^{y_0} u(y) y^{\kappa+2\alpha\sigma^{-2}-1} dy \right].$$

Then

$$\begin{aligned} \bar{u}(y_0) &= V_{aut}(y_0) - \frac{y_0}{\kappa} V'_{aut}(y_0) \\ &= \frac{r}{\alpha + \kappa\sigma^2} \Big[y_0^\kappa \int_{y_0}^\infty u(y) y^{-\kappa-1} dy + y_0^{-\kappa-2\alpha\sigma^{-2}} \int_0^{y_0} u(y) y^{\kappa+2\alpha\sigma^{-2}-1} dy \\ &\quad - y_0^\kappa \int_{y_0}^\infty u(y) y^{-\kappa-1} dy + \frac{\kappa + 2\alpha\sigma^{-2}}{\kappa} y_0^{-\kappa-2\alpha\sigma^{-2}} \int_0^{y_0} u(y) y^{\kappa+2\alpha\sigma^{-2}-1} dy \Big] \\ &= \frac{2r}{\kappa\sigma^2} y_0^{-\kappa-2\alpha\sigma^{-2}} \int_0^{y_0} u(y) y^{\kappa+2\alpha\sigma^{-2}-1} dy \\ &= \frac{1}{\kappa + 2\alpha/\sigma^2} y_0^{-\kappa-2\alpha\sigma^{-2}} \int_0^{y_0} u(y) y^{\kappa+2\alpha\sigma^{-2}-1} dy \\ &= \int_0^{y_0} u(y) d \left(\frac{y}{y_0}\right)^{\kappa+2\alpha\sigma^{-2}}. \end{aligned}$$

Because u is strictly increasing, it follows that $\bar{u}$ is a strictly increasing function and that $\bar{u}(y_0) < u(y_0)$ for all y_0. ■

Proof of Lemma 2

(i) Directly from (12), we have that $\bar{u}(y) < V_{aut}(y)$ at all y because $\kappa > 0$. We can thus see in (16) that V is strictly increasing in y because the weight on the larger value $V_{aut}(m)$ is strictly increasing in y. Indeed, taking the partial derivative in (16), we have

$$V_y(y, m) = \kappa y^{\kappa-1} m^{-\kappa} (V_{aut}(m) - \bar{u}(m)) > 0.$$

To see that $V_y(y,m) \le V'_{aut}(y)$, first note (16) can be written as

$$V(y,m) = -\int_y^m \bar{u}(m) d\left(\frac{y}{x}\right)^\kappa + \left(\frac{y}{m}\right)^\kappa V_{aut}(m), \tag{30}$$

because $1 - (\frac{y}{m})^\kappa = -\int_y^m d(\frac{y}{x})^\kappa$. Note also that definition of $\bar{u}(\cdot)$ allows us to express $V_{aut}(y)$ as

$$V_{aut}(y) = -\int_y^m \bar{u}(x) d\left(\frac{y}{x}\right)^\kappa + \left(\frac{y}{m}\right)^\kappa V_{aut}(m), \text{ for any } m \ge y > 0. \tag{31}$$

To see this, note that this equation holds trivially for $m = y$ and the derivative of the right-hand side with respect to m

$$\kappa y^\kappa m^{-\kappa-1} \bar{u}(m) - \kappa y^\kappa m^{-\kappa-1} V_{aut}(m) + \left(\frac{y}{m}\right)^\kappa V'_{aut}(m)$$

is zero because $\bar{u}(m) = V_{aut}(m) - \kappa^{-1} m V'_{aut}(m)$. Thus, the right-hand side is constant in m. From (30) and (31) we have

$$V(y,m) - V_{aut}(y) = -\int_y^m \left(\bar{u}(m) - \bar{u}(x)\right) d\left(\frac{y}{x}\right)^\kappa.$$

Introducing a new variable $s = \frac{x}{y}$, we rewrite the above as

$$\begin{aligned} V(y,m) - V_{aut}(y) &= -\int_1^{m/y} (\bar{u}(m) - \bar{u}(sy)) d\left(\frac{1}{s}\right)^\kappa \\ &= \kappa \int_1^{m/y} (\bar{u}(m) - \bar{u}(sy)) s^{-\kappa-1} ds. \end{aligned}$$

Thus $V_y(y,m) - V'_{aut}(y) \le 0$ and equality holds only if $y = m$.

(ii) Since $\kappa > 1$,

$$V_y(y,m) = \kappa y^{\kappa-1} m^{-\kappa} (V_{aut}(m) - \bar{u}(m))$$

is strictly increasing in y.

(iii)

$$\begin{aligned} V(y,m) &= \left(1 - \left(\frac{y}{m}\right)^\kappa\right) \bar{u}(m) + \left(\frac{y}{m}\right)^\kappa V_{aut}(m) \\ &= -\int_y^m \bar{u}(m) d\left(\frac{y}{x}\right)^\kappa - \left(\frac{y}{m}\right)^\kappa \int_m^\infty \bar{u}(x) d\left(\frac{m}{x}\right)^\kappa \\ &= -\int_y^\infty \bar{u}(\max\{m,x\}) d\left(\frac{y}{x}\right)^\kappa. \end{aligned}$$

Thus $V_m(y,m) = -\int_y^m \bar{u}'(m) d\left(\frac{y}{x}\right)^\kappa \ge 0$ with equality only if $y = m$.

■

Proof of Lemma 3

(i) Differentiating (18) with respect to y we have

$$\begin{aligned} Z_y(y,m) &= \kappa y^{\kappa-1}m^{-\kappa}\left[-\int_m^\infty u^{-1}(\bar{u}(x))d\left(\frac{m}{x}\right)^\kappa - u^{-1}(\bar{u}(m))\right] - \frac{r}{r-\mu} \\ &> -\frac{r}{r-\mu}. \end{aligned}$$

Also, because $\kappa > 1$, $Z_y(y,m)$ increases with y and $\lim_{y\to 0} Z_y(y,m) = -r/(r-\mu)$.

(ii) We can write (18) as

$$Z(y,m) = -\int_y^\infty u^{-1}(\bar{u}(\max\{x,m\}))d\left(\frac{y}{x}\right)^\kappa - \frac{r}{r-\mu}y.$$

From here we get that $Z_m(y,m) = -\int_y^m (u^{-1})'(\bar{u}(m))\bar{u}'(m)d\left(\frac{y}{x}\right)^\kappa \geq 0$ with equality only if $y = m$.

(iii) By part (i), $Z_y(y,m)$ is monotonically increasing in y. Thus $Z_y(y,m)|_{y=m} \leq 0$ implies $Z_y(y,m) \leq 0$ for all $y \leq m$. Otherwise, if $Z_y(y,m)|_{y=m} > 0$, then, by continuity, $Z_y(y,m) > 0$ for y sufficiently close to m. The sign of $Z_y(y,m)|_{y=m}$ is the same as that of $\frac{dZ(y,y)}{dy}|_{y=m}$ because $Z_m(y,m)|_{y=m} = 0$ by part (ii).

■

Proof of Proposition 1

By Ito's lemma

$$dA_t = \frac{-1}{r-\mu}\mu y_t dt + \frac{-1}{r-\mu}\sigma y_t dw_t.$$

The flow constraint (24) will be satisfied with $\beta_t = -\frac{\sigma y_t}{r-\mu}$ if we show that

$$\frac{-1}{r-\mu}\mu y_t = rA_t - c_t + y_t,$$

which is easy to check after substituting $-y_t\frac{r}{r-\mu}$ for $rA_t - c_t$. Thus, the proposed strategy is budget-feasible for the agent because $\lim_{t\to\infty} E[e^{-rt}(u^{-1}(\bar{V})/r - y_t\frac{1}{r-\mu})] = 0$. The optimality of this trading strategy can be proved by a duality argument. Suppose not. Then, there is an alternative trading strategy for the agent such that the consumption plan associated with it, $\{\tilde{c}_t; t \geq 0\}$, delivers to the agent utility $\tilde{V} > \bar{V}$. (This alternative consumption plan cannot be constant because the constant consumption plan that delivers $\tilde{V}$, i.e., $c_t = u^{-1}(\tilde{V})$, is not budget-feasible for the agent.) But this means that the consumption allocation $\{\tilde{c}_t; t \geq 0\}$, if used by the principal in the contracting problem $(y_0, \tilde{V})$, could deliver to the agent the utility $\tilde{V} > \bar{V}$ at the cost $C^f(y_t, \bar{V}) < C^f(y_t, \tilde{V})$, which contradicts the optimality of the constant consumption allocation $c_t = u^{-1}(\tilde{V})$ in the contracting problem $(y_0, \tilde{V})$. ■

Proof of Proposition 2

We first show that the strategy $\{c_t, A_t, \beta_t; t \geq 0\}$ described in the statement of the proposition is feasible, then prove that it is optimal. Note that $A_t = Z(y_t, \bar{m}_t)/r = C(y_t, V(y_t, \bar{m}_t))/r \geq C(y_t, V(y_t, y_t))/r = B_t$, thus the borrowing constraint is satisfied. Applying Ito's lemma to the martingale

$$\int_0^t re^{-rs}(c_s - y_s)ds + e^{-rt}Z_t(y_t, \bar{m}_t),$$

we have that the drift of Z_t is $r(Z_t + y_t - c_t)dt$. Applying Ito's lemma to Z_t and noting that $\bar{m}_t$ is monotonically increasing (i.e., no volatility), we have

$$dZ_t = r(Z_t + y_t - c_t)dt + Z_y(y_t, \bar{m}_t)\sigma y_t dw_t.$$

Therefore,

$$\begin{aligned} dA_t &= (rA_t + y_t - c_t)dt + r^{-1}Z_y(y_t, \bar{m}_t)\sigma y_t dw_t \\ &= (rA_t + y_t - c_t)dt + \beta_t dw_t, \end{aligned}$$

which shows that the policy $\{c_t, A_t, \beta_t; t \geq 0\}$ is budget-feasible to the agent.

To see that $\{c_t, A_t, \beta_t; t \geq 0\}$ is optimal, we must argue that the agent cannot do better than $\bar{V}$. By contradiction, suppose the agent's optimal plan is $\{\tilde{c}_t, \tilde{A}_t, \tilde{\beta}_t; t \geq 0\}$ and $E\left[\int_0^\infty re^{-rs}u(\tilde{c}_t)dt\right] > \bar{V}$. Then the consumption allocation $\{\tilde{c}_t; t \geq 0\}$ must satisfy the participation constraints at every time and under all states because $\tilde{A}_t \geq B(y_t)$ for all t and the continuation utility $E\left[\int_0^\infty re^{-rs}u(\tilde{c}_{t+s})ds|\mathcal{F}_t\right]$ is at least as large as $V_{aut}(y_t)$, due to the optimality of $\{\tilde{c}_t; t \geq 0\}$. If the agent follows $\{\tilde{c}_t, \tilde{A}_t, \tilde{\beta}_t; t \geq 0\}$, the bank's cost is still A_0, because the bank's expected return on the fair-odds hedging asset is zero no matter what $\tilde{\beta}_t$ is. Thus, we find an enforceable contract $\{\tilde{c}_t; t \geq 0\}$ that incurs the same cost $rA_0 = C(y_0, v_0)$ to the principal as $\{c_t; t \geq 0\}$ but delivers a utility larger than $\bar{V}$. This contradicts the fact that higher promised utility incurs higher cost, i.e., $Z_m(y, m) \geq 0$. ■

Proof of Lemma 4

We have $W_t = W(y_t, \bar{m}_t)$ where

$$W(y, m) = \frac{Z(y, m)}{r} + \frac{y}{r - \mu}.$$

Thus,

$$W_y(y, m) = \frac{Z_y(y, m)}{r} + \frac{1}{r - \mu} > 0,$$

where the strict inequality follows from Lemma 3(i). ■

Proof of Lemma 5

We have that $S_t = S(y_t, \bar{m}_t)$ where

$$S(y,m) = \frac{Z(y,m) - Z(y,y)}{r}.$$

We need to show that $S_y(y,m) \leq 0$ with equality if and only if $y = m$. It thus suffices to show that $Z_{ym}(y,m) < 0$. Differentiating (18) with respect to m we have

$$Z_m(y,m) = \left(1 - \left(\frac{y}{m}\right)^{\kappa}\right)(u^{-1})'(\bar{u}(m))\bar{u}'(m).$$

Thus

$$Z_{ym}(y,m) = -\kappa y^{\kappa-1} m^{-\kappa}(u^{-1})'(\bar{u}(m))\bar{u}'(m) < 0.$$

■

Proof of Lemma 6

The volatility of S_t is $\left(Z_y(y_t, \bar{m}_t) - \frac{dZ(y_t,y_t)}{dy_t}\right)\sigma y_t dw_t/r$ and the drift of S_t is

$$Z(y_t, \bar{m}_t) + y_t - u^{-1}(\bar{u}(\bar{m}_t)) - \frac{1}{r}\frac{dZ(y_t,y_t)}{dy_t}\mu y_t - \frac{1}{2}\frac{1}{r}\frac{d^2Z(y_t,y_t)}{d(y_t)^2}(\sigma y_t)^2.$$

It is easy to see that $\left(Z_y(y_t, \bar{m}_t) - \frac{dZ(y_t,y_t)}{dy_t}\right)|_{y_t=\bar{m}_t} = 0$, thus S_t has zero volatility when $y_t = m_t$ (i.e., when $S_t = 0$). To show that the drift when $y_t = \bar{m}_t$ is strictly positive, note that we have

$$\begin{aligned}
\frac{dZ(y,y)}{dy} &= \kappa y^{-1}\left[-\int_y^{\infty} u^{-1}(\bar{u}(x))d\left(\frac{y}{x}\right)^{\kappa} - u^{-1}(\bar{u}(y))\right] - \frac{r}{r-\mu}, \\
\frac{d^2Z(y,y)}{dy^2} &= (\kappa^2-\kappa)y^{-2}\left[-\int_y^{\infty} u^{-1}(\bar{u}(x))d\left(\frac{y}{x}\right)^{\kappa} - u^{-1}(\bar{u}(y))\right] - \kappa y^{-1}(u^{-1})'(\bar{u}(y))\bar{u}'(y).
\end{aligned}$$

Thus

$$\begin{aligned}
& \frac{1}{r}\frac{dZ(y,y)}{dy}\mu y + \frac{1}{2}\frac{1}{r}\frac{d^2Z(y,y)}{dy^2}(\sigma y)^2 \\
= \; & \frac{\kappa\mu + \frac{1}{2}(\kappa^2-\kappa)\sigma^2}{r}\left[-\int_y^{\infty} u^{-1}(\bar{u}(x))d\left(\frac{y}{x}\right)^{\kappa} - u^{-1}(\bar{u}(y))\right] - \frac{\mu}{r-\mu}y \\
& -\frac{\kappa\sigma^2 y}{2r}(u^{-1})'(\bar{u}(y))\bar{u}'(y) \\
= \; & \left[-\int_y^{\infty} u^{-1}(\bar{u}(x))d\left(\frac{y}{x}\right)^{\kappa} - u^{-1}(\bar{u}(y))\right] - \frac{\mu}{r-\mu}y - \frac{\kappa\sigma^2 y}{2r}(u^{-1})'(\bar{u}(y))\bar{u}'(y),
\end{aligned}$$

where the last equality uses (13). Thus, drift of S_t when $y_t = \bar{m}_t$ is

$$\frac{\kappa\sigma^2 y_t}{2r}(u^{-1})'(\bar{u}(y_t))\bar{u}'(y_t) > 0.$$

■

Proof of Lemma 7

Follows directly from Lemma 3(i). ■

Proof of Lemma 8

Suppose there exists another borrowing constraint $\tilde{B}(y_t)$ and initial asset level $\tilde{A}_0$ that also implement the optimal allocation. Since in any implementation, the agent's present value of consumption must equal the present value of total wealth, the agent's assets in his optimal trading strategy must be $A_t = Z(y_t, \bar{m}_t)/r$. Feasibility $Z(y_t, \bar{m}_t)/r \geq \tilde{B}(y_t)$ for all $y_t \leq \bar{m}_t$ requires that $C(y_t, V_{aut}(y_t))/r = Z(y_t, y_t)/r \geq \tilde{B}(y_t)$, i.e., $B(y_t) \geq \tilde{B}(y_t)$ for all y_t. To show that $B(y_t) = \tilde{B}(y_t)$ for all y_t, suppose $B(y^*) > \tilde{B}(y^*)$ for some y^*. Then there is a small $\epsilon > 0$, such that $B(y) > \tilde{B}(y)$ for all $y \in (y^* - \epsilon, y^* + \epsilon)$. Now, an agent with initial income y^* and asset level $A_0 = C(y^*, V_{aut}(y^*))/r$, instead of following the proposed equilibrium strategy which delivers the value $V_{aut}(y^*)$, can achieve utility higher than $V_{aut}(y^*)$. To do this, he can consume $y_t + 1$ (and choose zero hedging, i.e., set $\beta_t = 0$) for a short period of time, and default immediately after his borrowing constraint binds for the first time. The agent's utility from this strategy is higher than the autarky value $V_{aut}(y^*)$ because before he defaults he consumes $y_t + 1$, i.e., more than what he would consume in autarky. He is able to do this for a strictly positive amount of time because the borrowing constraint is initially non-binding and $\tilde{B}(\cdot)$ is continuous. ■

Appendix B

This appendix provides a formal verification of the optimality of the contract (11).

First, we express the principal's cost minimization problem as a dynamic programming problem in a two-dimensional state vector (y, v), where y is the agent's current level of income and v is the current level of the continuation utility that the principal must provide to the agent.

By Ito's formula, y_t satisfies

$$dy_t = \mu y_t dt + \sigma y_t dw_t, \tag{32}$$

where $\mu = \alpha + \sigma^2/2$. In this representation, the income process is decomposed into a drift and a volatility component. The same decomposition can be provided for the agent's continuation value process v_t. In particular, the following proposition of Sannikov (2008) demonstrates how the promised utility process $v = \{v_t; t \geq 0\}$ defined in (3) can be decomposed into the sum of a drift term and a volatility term.

Proposition 3 *Let c be an allocation and v the promised utility process as defined in (3). There exists a progressively measurable process $Y = \{Y_t, \mathcal{F}_t; 0 \leq t < \infty\}$ such that*

$$v_t = v_0 + \int_0^t r(v_s - u(c_s))ds + \int_0^t Y_s dw_s.$$

Put differently, the evolution of the promised utility process v implied by c can be decomposed as

$$dv_t = r(v_t - u(c_t))dt + Y_t dw_t. \tag{33}$$

This decomposition pins down the process Y uniquely up to a subset of measure zero.

Proof See Sannikov (2008). ■

In this representation, $r(v_t - u(c_t))$ is the drift of the promised utility process v_t and Y_t as the sensitivity of v_t to income shocks dw_t. Among useful properties of this representation is $E_t[Y_t dw_t] = 0$.

In our problem, the Dynamic Principle of Optimality (DPO) implies that continuation of any efficient contract is itself efficient. Indeed, let $c(y_0, \bar{V}) \in \Psi$ be an efficient contract in the problem $(y_0, \bar{V})$ and suppose that c is applied over some time interval $[0, t)$. At t, the agents's income is y_t and his continuation utility is v_t. The continuation allocation $\{c_{t+s}(y_0, \bar{V}); s \geq 0\}$ of the efficient contract $c(y_0, \bar{V})$ has to be the same as the efficient allocation starting at (y_t, v_t) for otherwise $c(y_0, \bar{V})$ would not be efficient to begin with. Thus, for any date t and state $\omega \in \Omega$,[14] we have

$$c_t(y_0, \bar{V}) = c_0(y_t, v_t), \tag{34}$$

where contracts on both sides are processes on $(\Omega, \mathcal{F}, \mathbf{P})$. By Proposition 1 both of these contracts are a.e.-uniquely representable by drift and sensitivity components. Thus, (34) implies that the sensitivity components of these representations are the same a.e.:

$$Y_t(y_0, \bar{V}) = Y_0(y_t, v_t), \tag{35}$$

where $Y(y, v)$ denotes the sensitivity process of the efficient contract $c(y, v)$ for all $(y, v) \in \Theta$.

In sum, the DPO implies that the efficient contracts in Ψ are representable by a pair of real-valued functions $(c_0(y_t, v_t), Y_0(y_t, v_t))$, where $c_0 : \Theta \to \mathbb{R}_+$ and $Y_0 : \Theta \to \mathbb{R}$. Because of (34) and (35), these two functions (the so-called policy rules) can be used in (33) to express the law of motion for the state variable (y_t, v_t) as

$$\begin{aligned} dy_t &= \mu y_t dt + \sigma y_t dw_t, \\ dv_t &= r(v_t - u(c_0(y_t, v_t)))dt + Y_0(y_t, v_t)dw_t. \end{aligned}$$

This law of motion and the policy rules can be repeatedly applied to generate the sensitivity process $Y(y_0, \bar{V}) = \{Y_t(y_0, \bar{V}); t \geq 0\}$ and the contract allocation $c(y_0, \bar{V}) = \{c_t(y_0, \bar{V}); t \geq 0\}$ for any initial $(y_0, \bar{V}) \in \Theta$.

The cost function $C(y_t, v_t)$, i.e., the cost of an optimal contract starting from the state (y_t, v_t), must satisfy the necessary *Hamilton–Jacobi–Bellman (HJB)* equation given as follows. For the interior values of the state variable, i.e., for $v_t > V_{aut}(y_t)$, the HJB equation is standard (see, for example, Fleming and Soner (2006, equation (5.8), page 165)):

$$\begin{aligned} rC(y_t, v_t) &= \min_{c,Y} \Big\{ r(c - y_t) + C_y(y_t, v_t)\mu y_t + C_v(y_t, v_t) r(v_t - u(c)) \\ &\quad + \frac{\sigma^2 y_t^2}{2} C_{yy}(y_t, v_t) + \sigma y_t Y C_{vy}(y_t, v_t) + \frac{Y^2}{2} C_{vv}(y_t, v_t) \Big\}, \end{aligned} \tag{36}$$

[14] Actually, the same is true for any stopping time T on $(\Omega, \mathcal{F}, \mathbf{P})$.

where subscripts on C denote partial derivatives. At the boundary $v_t = V_{aut}(y_t)$, the HJB is the same except that the controls (c, Y) must be such that $v_{t+dt} \geq V_{aut}(y_{t+dt})$ with probability one. Otherwise, the agent would revert to permanent autarky with positive probability, which would be inefficient.

Denote the cost under the contract (11) $Z(y, M(y, v))$ by $J(y, v)$. We can now show that $J(y, v)$ satisfies the HJB equation (36).

Proposition 4 *$J(y, v)$ satisfies the HJB equation.*

Proof Consider a contract starting at $(y_0, \bar{V}) = (y, v) \in \Theta$. Recall in the contract $u(c_t) = \bar{u}(m_t) = \bar{u}(M(y_t, v_t))$. Define

$$G_t = \int_0^t re^{-rs}(c_s - y_s)ds + e^{-rt}J(y_t, v_t).$$

Because

$$G_t = E\left[\int_0^\infty re^{-rs}(c_s - y_s)ds|F_t\right],$$

we have that G_t is a martingale, and thus its drift is zero. Calculating this drift by applying Ito's lemma and the fact that the volatility of $V(y, m)$ is $V_y\sigma y$, and setting time equal to zero, we get

$$\begin{aligned} &r(u^{-1}(\bar{u}(m)) - y) - rJ(y, v) + J_y\mu y + J_v r(v - \bar{u}(m)) \\ &+\frac{1}{2}J_{yy}(\sigma y)^2 + J_{yv}(\sigma y)^2 V_y + \frac{1}{2}J_{vv}(\sigma y)^2 V_y^2 = 0, \end{aligned}$$

which is the HJB equation, except for the minimization operator. To verify that in fact

$$\begin{aligned} &r(u^{-1}(\bar{u}(m)) - y) - rJ(y, v) + J_y\mu y + J_v r(v - \bar{u}(m)) + \frac{1}{2}J_{yy}(\sigma y)^2 + J_{yv}(\sigma y)^2 V_y + \frac{1}{2}J_{vv}(\sigma y)^2 V_y^2 \\ = \quad &\min_{u,Y}\left\{r(u^{-1}(u) - y) + J_y\mu y + J_v r(v - u) + \frac{1}{2}J_{yy}(\sigma y)^2 + J_{yv}\sigma y Y + \frac{1}{2}J_{vv}Y^2\right\}, \end{aligned}$$

it suffices to show that $J_v = (u^{-1})'(\bar{u}(m))$ and $V_y = -J_{vy}/J_{vv}$.

To see the first of these equalities, recall from the proof of Lemma 2(iii) that $V_m = -\int_y^m \bar{u}'(m)d(\frac{y}{x})^\kappa$. Recall from the proof of Lemma 3(ii) that $Z_m = -\int_y^m (u^{-1})'(\bar{u}(m))\bar{u}'(m)d(\frac{y}{x})^\kappa$. Since $J(y, v) \equiv Z(y, M(y, v))$, we have

$$J_v = Z_m M_v = \frac{Z_m}{V_m} = (u^{-1})'(\bar{u}(m)).$$

To see the second equality, note $J_v(y, V(y, m)) = (u^{-1})'(\bar{u}(m))$ is independent of y when J_v is interpreted as a function of (y, m). Thus, we have that $J_{vy} + J_{vv}V_y = 0$. Thus $V_y = -J_{vy}/J_{vv}$. Therefore the HJB is verified. ■

We have thus verified a necessary condition for optimality. The next proposition shows sufficiency.

Proposition 5 *$J = C$, i.e., that the contract c constructed in (11) is efficient.*

Proof Let $N > 0$ be any positive number and consider an initial state $(y_0, \bar{V}) \in \Theta^{(N)} = \{(y, v) \in \Theta : 0 < y \leq N, v \leq V_{aut}(N)\}$. Consider an auxiliary dynamic programming problem in which the participation constraints are deleted (i.e., not required to hold) after the hitting time $\lambda = \min_t \{t : v_t = V_{aut}(N)\}$. Note that, since $v_t \geq V_{aut}(y_t)$ when $t \leq \lambda$, we have $\lambda \leq \tau_N$. An implication of deleting participation constraints is that the optimal consumption is perfectly smoothed after λ, i.e., $c_t = u^{-1}(V_{aut}(N))$ for $t \geq \lambda$, even as income y_t continues to fluctuate. To study the auxiliary problem, we can restrict attention to the interior of $\Theta^{(N)}$, where the law of motion of the state variable is the same as before. The cost function on the boundary $\partial\Theta^{(N)} = \{(y, v) \in \Theta : v = V_{aut}(N)\}$ is the full-commitment cost, i.e., $C^{(N)}(y, V_{aut}(N)) = u^{-1}(V_{aut}(N)) - \frac{ry}{r-\mu}$, because consumption is perfectly smoothed from the date λ on. The cost function $C^{(N)}(y, v)$ in the interior is by definition the cost of the optimal policies in the auxiliary dynamic programming problem. To solve the auxiliary problem, we make the same guess as before, i.e., consumption is defined as in equation (9) in section 3 before τ_N. That is, for any $t < \tau_N$ and $m_t < N$,

$$c_t = u^{-1}(\bar{u}(m_t)), \tag{37}$$

where $\bar{u} : \mathbb{R}_{++} \to \mathbb{R}$ is

$$\bar{u}(y) = V_{aut}(y) - \kappa^{-1} y V'_{aut}(y). \tag{38}$$

Any $\bar{V} \in [V_{aut}(y), V_{aut}(N)]$ is uniquely associated with an $m \in [y, N]$, because $V(y, m) = (1 - (\frac{y}{m})^\kappa)\bar{u}(m) + (\frac{y}{m})^\kappa V_{aut}(m)$ is strictly increasing in m. We define, for $m \in [y, N]$,

$$Z^{(N)}(y, m) = -\int_y^N u^{-1}(\bar{u}(\max\{x, m\}))d\left(\frac{y}{x}\right)^\kappa + u^{-1}(V_{aut}(N))\left(\frac{y}{N}\right)^\kappa - \frac{r}{r-\mu}y.$$

We first claim that for any $(y, v) \in \Theta^{(N)}$, the function $J^{(N)}$ defined as $J^{(N)}(y, v) = Z^{(N)}(y, M(y, v))$ is the optimal cost function $C^{(N)}(y, v)$. To see this, note that $J^{(N)}$ satisfies the HJB on the state space $\Theta^{(N)}$,

$$\begin{aligned} rJ^{(N)}(y_t, v_t) &= \min_{c,Y}\Big\{r(c - y_t) + J_y^{(N)}(y_t, v_t)\mu y_t + J_v^{(N)}(y_t, v_t)r(v_t - u(c)) \\ &\quad + \frac{\sigma^2 y_t^2}{2}J_{yy}^{(N)}(y_t, v_t) + \sigma y_t Y J_{vy}^{(N)}(y_t, v_t) + \frac{Y^2}{2}J_{vv}^{(N)}(y_t, v_t)\Big\}. \end{aligned}$$

Pick any contract $\{\tilde{c}_t; t \geq 0\}$ and denote the volatility process of $\tilde{v}_t$ in Proposition 1 by $\{\tilde{Y}_t; t \geq 0\}$. We introduce, for each $n \geq 1$, the stopping time

$$T_n = \inf_t \left\{t \geq 0 : \int_0^t \tilde{Y}_s^2 ds \geq n \text{ or } \tilde{v}_t \geq V_{aut}(N)\right\}.$$

We define

$$G_t = \int_0^t re^{-rs}(\tilde{c}_s - y_s)ds + e^{-rt}J^{(N)}(y_t, \tilde{v}_t).$$

We apply the Ito's lemma to G_t and obtain

$$\begin{aligned} G_{t\wedge T_n} &= G_0 + \int_0^{t\wedge T_n} e^{-rs}\Big[r(\tilde{c}_s - y_s) - rJ^{(N)}(y_s,\tilde{v}_s) + J_y^{(N)}(y_s,\tilde{v}_s)\mu y_s + J_v^{(N)}(y_s,\tilde{v}_s)r(\tilde{v}_s - u(\tilde{c}_s)) \\ &\quad + \frac{\sigma^2 y_s^2}{2}J_{yy}^{(N)}(y_s,\tilde{v}_s) + \sigma y_s \tilde{Y}_s J_{vy}^{(N)}(y_s,\tilde{v}_s) + \frac{\tilde{Y}_s^2}{2}J_{vv}^{(N)}(y_s,\tilde{v}_s)\Big]ds \\ &\quad + \int_0^{t\wedge T_n} e^{-rs}\left[J_y^{(N)}(y_s,\tilde{v}_s)\sigma y_s + J_v^{(N)}(y_s,\tilde{v}_s)\tilde{Y}_s\right]dw_s. \end{aligned}$$

Since $\int_0^{t\wedge T_n} e^{-rs}[J_y^{(N)}(y_s,\tilde{v}_s)\sigma y_s + J_v^{(N)}(y_s,\tilde{v}_s)\tilde{Y}_s]dw_s$ has zero mean and the drift is non-negative, taking expectation, we see that

$$E(G_{t\wedge T_n}) \geq G_0 = J^{(N)}(y_0,\bar{V}).$$

In particular $E(G_{n\wedge T_n}) \geq J^{(N)}(y_0,\bar{V})$. Since $\lim_{n\to\infty} n\wedge T_n = \lambda$, $E[\int_0^\infty (\tilde{c}_s)re^{-rs}ds] < \infty$ and $E[\int_0^\infty (y_s)re^{-rs}ds] < \infty$, the dominated convergence theorem yields

$$E\left[\int_0^\lambda re^{-rs}(\tilde{c}_s - y_s)ds\right] = \lim_{n\to\infty} E\left[\int_0^{n\wedge T_n} re^{-rs}(\tilde{c}_s - y_s)ds\right].$$

Furthermore, since $J^{(N)}$ is bounded, $\lim_{n\to\infty} e^{-r(n\wedge T_n)}J^{(N)}(y_{n\wedge T_n},\tilde{v}_{n\wedge T_n})$ equals $e^{-r\lambda}J^{(N)}(y_\lambda,\tilde{v}_\lambda) = e^{-r\lambda}J^{(N)}(y_\lambda,V_{aut}(N))$ if $\lambda < \infty$, and equals 0 if $\lambda = \infty$. Thus

$$E\left[e^{-r\lambda}J^{(N)}(y_\lambda,V_{aut}(N))\right] = \lim_{n\to\infty} E\left[e^{-r(n\wedge T_n)}J^{(N)}(y_{n\wedge T_n},\tilde{v}_{n\wedge T_n})\right].$$

We get

$$E\left[\int_0^\lambda re^{-rs}(\tilde{c}_s - y_s)ds + e^{-r\lambda}J^{(N)}(y_\lambda,V_{aut}(N))\right] \geq J^{(N)}(y_0,\bar{V}).$$

This means that $J^{(N)}(y_0,\bar{V})$ is (weakly) less than the cost of any other contract $\{\tilde{c}_t; t \geq 0\}$, i.e., $J^{(N)} = C^{(N)}$.

Since the auxiliary problem has less constraints than those in the original problem, we know that the cost of the auxiliary problem is below that of the original problem, i.e., for all $N > 0$,

$$J^{(N)}(y,v) \leq C(y,v), \text{ for } (y,v) \in \Theta^{(N)}.$$

Taking limit $N \to \infty$, we have

$$\begin{aligned} J(y,v) &= -\int_y^\infty u^{-1}(\bar{u}(\max\{x,m\}))d\left(\frac{y}{x}\right)^\kappa - \frac{r}{r-\mu}y \\ &= \lim_{N\to\infty}\left(-\int_y^N u^{-1}(\bar{u}(\max\{x,m\}))d\left(\frac{y}{x}\right)^\kappa + u^{-1}(V_{aut}(N))\left(\frac{y}{N}\right)^\kappa\right) - \frac{r}{r-\mu}y \\ &= \lim_{N\to\infty} J^{(N)}(y,v) \\ &\leq C(y,v), \end{aligned}$$

where $m = M(y,v)$. Thus we have $J(y,v) = C(y,v)$ for all $(y,v) \in \Theta$. ■

Appendix C

Here we consider a discrete-time, one-sided commitment model in which the agent's preferences are represented by the expected utility function $E[\sum_{t=0}^{\infty}(1-\beta)\beta^t u(c_t)]$, with the discount factor $\beta \in (0,1)$ and u strictly increasing and strictly concave. His income process $y_t \in \{\bar{y}_1, \bar{y}_2, ..., \bar{y}_n\}$ is a Markov chain, where $\bar{y}_1 < \bar{y}_2 < ... < \bar{y}_n$. The transition probability π satisfies first-order stochastic dominance, i.e., $\pi_i(\cdot) = \Pr(\cdot|y_t = \bar{y}_i)$ first-order stochastic dominates $\pi_j(\cdot) = \Pr(\cdot|y_t = \bar{y}_j)$, when $i > j$. If $y_0 = \bar{y}_i$, let τ_i be the stopping time when income exceeds $\bar{y}_i$ for the first time, i.e., $\tau_i = \min\{t \geq 0 : y_t > \bar{y}_i\}$. Define

$$\bar{u}(\bar{y}_i) = \frac{E\left[\sum_{t=0}^{\tau_i - 1} \beta^t u(y_t) | y_0 = \bar{y}_i\right]}{E\left[\sum_{t=0}^{\tau_i - 1} \beta^t | y_0 = \bar{y}_i\right]}$$

to be the average utility when income does not exceed $\bar{y}_i$. Using first-order stochastic dominance, we can verify that $\bar{u}(\bar{y}_i)$ is strictly increasing in i.[15] Now consider a contracting problem, where $y_0 = \bar{y}_i$ and the agent's promised utility is $\bar{V} = V_{aut}(\bar{y}_i)$. Let $m_t = \max_{0 \leq s \leq t} y_s$ be the to-date maximal realized income. We construct a contract as

$$c_t = u^{-1}(\bar{u}(m_t)). \tag{39}$$

Note that the sequence $\{c_t, t \geq 0\}$ is weakly increasing, because $\bar{u}$ is increasing. Denote the continuation value in the above contract by $V(y_t, m_t)$. It is easily seen that the function $V(y, m)$ is increasing in m. To show that this contract satisfies all the participation constraints, note that $Z(\bar{y}_i, \bar{y}_i) = V_{aut}(\bar{y}_i)$ by the definition of $\bar{u}$. Thus $Z(y_t, m_t) \geq Z(y_t, y_t) = V_{aut}(y_t)$ for all (y_t, m_t). To show that the contract is optimal, we verify that the sufficient and necessary Lagrange conditions are satisfied. Borrowing notation from Ljungqvist and Sargent (2004, page 660), let $\beta^t \alpha_t$ be the multiplier on

$$E_t\left[\sum_{s=t}^{\infty}(1-\beta)\beta^{s-t}u(c_s)\right] - V_{aut}(y_t) \geq 0, \text{ for } t \geq 1,$$

and let ϕ be the multiplier on

$$E\left[\sum_{s=0}^{\infty}(1-\beta)\beta^s u(c_s)\right] - V_{aut}(y_0) = 0.$$

We construct $\alpha_t = (u^{-1})'(u(c_t)) - (u^{-1})'(u(c_{t-1}))$ and $\phi = (u^{-1})'(u(c_0))$, which satisfies the non-negativity of the multipliers since consumption is non-decreasing. It is easy to verify that the analogs of the first-order conditions in Ljungqvist and Sargent (2004, equation 19.4.6a) are satisfied by this construction. The analogs of the complementary slackness conditions in Ljungqvist and Sargent (2004, equation 19.4.6b) are satisfied as well. This is because when

[15] Proof is available upon request.

$\alpha_t > 0$, then $c_t > c_{t-1}$ and $y_t = m_t > m_{t-1}$. Then participation constraint holds with equality when $\alpha_t > 0$ because $Z(\bar{y}_i, \bar{y}_i) = V_{aut}(\bar{y}_i)$ for all $\bar{y}_i$.

First-order stochastic dominance is necessary for the characterization in (39). The following example shows that when no structure on the income process is imposed, then consumption in the optimal contract can increase when income decreases. Take $\epsilon \in (0,1)$ and $\delta \in (0,1)$. Suppose $y_0 = 1$ with probability 1 and $\Pr(y_1 = 1-\epsilon | y_0 = 1) = \delta$ and $\Pr(y_1 = \frac{1}{2} | y_0 = 1) = 1-\delta$. Assume also that $\Pr(y_{t+1} = 1-\epsilon | y_t = 1-\epsilon) = 1$ and $\Pr(y_{t+1} = \frac{1}{2^{t+1}} | y_t = \frac{1}{2^t}) = 1, t \geq 1$. That is, there is uncertainty only at the beginning of period 1 and income afterwards is either constant or monotonically decreasing, depending on the realization of income in the first period. When ϵ is sufficiently small, the optimal contract to deliver to the agent the ex ante autarky value $V_{aut}(y_0)$ involves only two consumption values. If income is $y_1 = 1-\epsilon$, the agent consumes $1-\epsilon$ in every period $t \geq 1$ because of the binding participation constraint. If $y_1 = 1/2$, the participation constraint does not bind and it is efficient to smooth the consumption path at the constant level, denoted here by $\bar{c}$, at all $t \geq 0$. In order to deliver $V_{aut}(y_0)$, consumption $\bar{c}$ must satisfy

$$(1-\beta)u(\bar{c}) + (1-\delta)\beta u(\bar{c}) = (1-\beta)u(1) + (1-\delta)\sum_{t=1}^{\infty}(1-\beta)\beta^t u\left(\frac{1}{2^t}\right).$$

When ϵ is sufficiently small, $\bar{c} < 1-\epsilon$. Thus, the optimal contract starts with consumption $\bar{c}$ and jumps up to $1-\epsilon$ when income decreases from 1 to $1-\epsilon$.

References

Albanesi, S., Sleet, C., 2006. Dynamic optimal taxation with private information. Review of Economic Studies 73, 1–30.

Albuquerque, R., Hopenhayn, H., 2004. Optimal lending contracts and firm dynamics. Review of Economic Studies 71, 285–315.

Alvarez, F., Jermann, U., 2000. Efficiency, equilibrium, and asset pricing with risk of default. Econometrica 68(4), 775–797.

Biais, B., Mariotti, T., Plantin, G., Rochet, J., 2007. Dynamic security design: Convergence to continuous time and asset pricing implications. Review of Economic Studies 74, 345–390.

Borodin, A., Salminen, P., 2002. Handbook of Brownian Motion - Facts and Formulae, 2nd ed. Birkhauser.

Demarzo, P., Sannikov, Y., 2007. Optimal security design and dynamic capital structure in a continuous-time agency model. The Journal of Finance 61, 2681–2724.

Eaton, J., Gersovitz, M., 1981. Debt with potential repudiation: Theoretical and empirical analysis. Review of Economic Studies 48(2), 289–309.

Fleming, W., Soner, H., 2006. Controlled Markov Processes and Viscosity Solutions, Second Edition. Springer Science+ Business Media, Inc., New York, NY, USA.

Harris, M., Holmstrom, B., 1982. A theory of wage dynamics. Review of Economic Studies 49, 315–333.

He, Z., 2009. Optimal executive compensation when firm size follows geometric brownian motion. Review of Financial Studies 22, 859–892.

Karatzas, I., Shreve, S., 1991. Brownian Motion and Stochastic Calculus, 2nd ed. Springer-Verlag.

Kehoe, T., Levine, D., 1993. Debt-constrained asset markets. Review of Economic Studies 60(4), 865–888.

Kocherlakota, N., 1996. Implications of efficient risk sharing without commmitment. Review of Economic Studies 63(4), 595–609.

Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality? evidence and theory. Review of Economic Studies 73, 163–193.

Krueger, D., Uhlig, H., 2006. Competitive risk sharing contracts with one-sided commitment. Journal of Monetary Economics 53, 1661–1691.

Ljungqvist, L., Sargent, T., 2004. Recursive Macroeconomic Theory, Second Edition. The MIT Press, Cambridge, Massachusetts, USA.

Piskorski, T., Tchistyi, A., forthcoming. Optimal mortgage design. Review of Financial Studies.

Sannikov, Y., 2008. A continuous-time version of the principal-agent problem. Review of Economic Studies 75, 957–984.

Thomas, J., Worral, T., 1988. Self-enforcing wage contracts. Review of Economic Studies 55, 541–554.

Zhang, Y., 2009. Dynamic contracting with persistent shocks. Journal of Economic Theory 144, 635–675.