Optimal risk sharing and borrowing constraints in a continuous-time model with limited commitment*

Borys Grochulski† Yuzhe Zhang‡

December 10, 2010

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 a time-invariant, strictly increasing function of a single state variable: the maximal level of the agent's income realized to date. We characterize this function in terms of the agent's outside option value function and the discounted amount of time in which the agent's income process is expected to reach a new to-date maximum. Under constant relative risk aversion we solve the model in closed-form: optimal consumption of the agent equals a constant fraction of his maximal 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 constraints on the amounts they can borrow. There is a large literature exploring the relation between borrowing constraints and limited contract enforcement.[1] When contract enforcement is limited, lenders face the risk of borrower default. The role of borrowing constraints is to mitigate this risk efficiently. In this paper, we contribute to this literature by studying an optimal contracting problem with limited enforcement in a tractable continuous-time framework that allows us to obtain a sharp characterization of the optimal contract as well as of the borrowing constraints that implement it.

Our analysis has two parts. In the first part, we study an optimal long-term contracting problem between a risk-neutral, fully-committed, deep-pocketed principal and a risk-averse,

*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.

[1] Examples of contributions to this literature include Alvarez and Jermann [1], Albuquerque and Hopenhayn [2], Kehoe and Perri [3].

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 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, his optimal consumption is simply given by a constant fraction of the maximal level of his income realized to date.

To see the intuition behind our characterization of the optimal contract, suppose that the principal is to deliver to the agent the level of utility exactly equal to the agent's value of autarchy as of time zero. If the agent could commit to never defaulting, the optimal contract would give the agent a constant consumption flow forever. This is because the principal is risk-neutral, does not face a flow resource constraint, and discounts future payoffs at the same rate as the agent. Under this full-insurance contract, the agent's value of continuing with the contract does not change over time, i.e., remains equal to his initial autarchy value. Note now that even when the agent cannot commit, the full insurance contract does not cause the agent to default (revert to autarchy) for as long as his income fluctuates below its time-zero level, i.e., for as long as the date-zero level remains the to-date maximum level attained by the agent's income process. This is because during any such time interval the agent's autarchy value—being strictly increasing in income—fluctuates below the agent's initial autarchy value, which means that the value of defaulting remains below the value of continuing with the contract (the agent's participation constraint is satisfied). Under the full-insurance contract, however, the agent will default as soon as his income exceeds its time-zero level—i.e., when income attains a new to-date maximum—precisely because the agent's outside option value will at that point exceed the value of continuing with the full-insurance contract. In order to prevent default, the principal has to deviate from the full-insurance contract by increasing the agent's consumption at that moment (as the agent's participation constraint binds), but not before then. So, even when the agent cannot commit, the principal will give the agent a constant consumption level for as long as the agent's income is not at its to-date maximum. The same logic applies after an all-time maximum has been realized and the agent's consumption has been increased: consumption remains constant until income hits its next all-time maximum level. And so on. Optimal consumption, therefore, is always an increasing function of the current to-date maximum level of income.

For a given amount of lifetime utility that the principal provides to the agent, the future consumption increases that are necessary under limited commitment imply that the initial consumption level delivered to the agent is lower than what it would be under the full-insurance contract. The key question is by how much. The answer depends on the magnitude of the future consumption hikes and on how soon they are expected to occur. The advantage of our model is that we can use the properties of the geometric Brownian motion process to give an exact answer to this question. We derive an explicit formula for the mapping from the current to-date maximum income level to the optimal consumption level. At each point in time, the utility flow the agent receives equals the level he would receive under full insurance less the increase in his outside option value that will occur the next time his income reaches a new maximum, divided by the amount of time in which his income is expected to reach it. The increase in the outside option value is measured by the first derivative of the agent's autarchy value function. The amount of time before income reaches its next all-time maximum is an example of the so-called hitting time. When income is a geometric Brownian motion, the expected discounted hitting time needed in our formula is given by a simple, closed-form expression. Our continuous-time framework therefore allows us to express the agent's optimal consumption in terms of the agent's autarchy value function, its derivative, and an expected discounted Brownian hitting time.

Our formula for the optimal consumption process allows us to provide a detailed characterization of the dynamics of the agent's continuation value in the contract and the principal's profit from the relationship. The agent's continuation value is always positively correlated with his income. This correlation, however, is almost always strictly less than what it would be in autarky, except on a measure-zero set of times at which the agent's participation constraint binds, when the two are equal. This correlation decreases with the distance between the agent's current income level and its to-date maximum. Thus, for a given to-date maximum, the agent is more fully insured at lower income levels. As the agent's income approaches its current to-date maximum, the degree of insurance provided to the agent decreases, i.e., the agent becomes progressively more exposed to risk. We also show that, due to the fact that the present value of all future consumption given to the agent is increasing in the agent's current income level, the principal's profit can be decreasing in the current level of the agent's income, which would not take place under full commitment.

In the second part of the paper, we study a simple trading mechanism that implements the optimal long-term contract. 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 replicates the optimal long-term contract. 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 only partially insures his income.

As already mentioned, in an environment otherwise identical to ours but in which the agent can fully commit, any efficient allocation of consumption 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, 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, therefore, shows clearly the role of borrowing constraints in mitigating the risk of borrower default.

The implementation exercise with the two-account trading mechanism provides two additional insights. First, it gives us a better understanding of the optimal long-term risk-sharing contract by identifying a set of restrictions on trading consistent with optimal risk sharing that are weaker than the strong restrictions implicit in the optimal long-term contract itself, where no retrading is allowed. For example, the implementation exercise lets us see that the optimal contract with limited commitment does not place any restrictions on how much the agent is allowed to save. In dynamic risk-sharing problems with private information, in contrast, optimal contracts typically do restrict agents' savings (Rogerson [4], Golosov et al. [5]).

Second, the implementation exercise delivers a theory of optimal borrowing constraints. The standard Bewley-Aiyagari incomplete-markets model does not endogenously determine what agents' borrowing limits should be. Moreover, in the implementing mechanism we can examine the dynamics of the agent's financial wealth and compare it with the dynamics of financial wealth in the standard Bewley-Aiyagari model. Although the correlation between the agent's total wealth (i.e., financial wealth plus the present value of his future income) and his current income is positive in the limited-commitment model, the correlation between the agent's buffer stock of financial assets (i.e., financial assets held in excess of the borrowing limit) and the agent's current income level is negative. This prediction stands in stark contrast to the implications of the Bewley-Aiyagari incomplete-markets models of self-insurance, in which the financial buffer stock is positively correlated with income. The implementation thus clarifies the difference in the role that financial wealth fulfills in the incomplete- and complete-markets models. With incomplete markets, where agents can only self-insure, the role of financial wealth is to buffer off the negative shocks to income. With complete markets, where agents can explicitly hedge their income shocks, the role of financial wealth is to buffer off the losses from hedging, which occur precisely when the shocks to income are positive. For that reason, the correlation between income and financial wealth takes the opposite sign in incomplete- and complete-markets models.

Relation to the literature Our paper is closely related to the literature studying optimal contracts and equilibrium outcomes in environments with commitment frictions. Contributions to this literature include Harris and Holmstrom [6], Thomas and Worrall [7], Marcet and Marimon [8], Kehoe and Levine [9], Kocherlakota [10], Alvarez and Jermann [1], Albuquerque and Hopenhayn [2], Ljungqvist and Sargent [11], Krueger and Perri [12], Krueger and Uhlig [13]. 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.[2] As we show in Appendix C, however, our method for the characterization of the optimal contract is not specific to our continuous-time framework.

Our paper is also related to several recent studies of optimal contracting problems in continuous time with private information.[3] In particular, our proof of the optimality of the contract is based on the techniques developed in Sannikov [17]. 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.

In addition to the optimal contracting papers, our paper is related to papers studying the 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 [21]). 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 [1] 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 [1] 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-Aiyagari-type models, which have been widely used in macroeconomics

[2] Monge-Naranjo [14] studies an optimal contracting problem with limited enforcement in continuous time. In the model studied in that paper, there are no shocks (deterministic dynamics) and agents have no preference for intertemporal smoothing (linear utility). In this paper, we study a stochastic model with a risk-averse agent.

[3] E.g., Demarzo and Sannikov [15], Biais et al. [16], Sannikov [17], Piskorski and Tchistyi [18], He [19], Zhang [20].

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-Aiyagari-type incomplete-markets model. Clearly, the key difference between these two models is the availability of hedging. Our analysis shows an important testable implication of this difference: the correlation between the agent's current income and his financial buffer stock has the opposite sign in the two classes of models.

Krueger and Perri [12] compare the implications of Alvarez-Jermann-type models and Bewley-Aiyagari-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 [1].

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 a characterization of optimal borrowing constraints and the agent's optimal trading strategy. 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 the agent are represented by the expected utility function

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

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 $\mathbb{E}$ 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

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

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 $\mathbb{E}[y_{t+s}|\mathcal{F}_t] = y_t \exp(\mu s)$ for any $t, s > 0$, we have that

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

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

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

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

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

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

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

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

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

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.[4] To deal with this difficulty, we first relax

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

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:

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

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

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

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:

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

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

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

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 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 a formula for the certainty equivalent utility flow rate $\bar{u}(y_0)$ in the unrelaxed planner's problem:

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

Denote $1-\mathbb{E}[e^{-r\tau_{y_0+\varepsilon}}]$ 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., $\mathbb{E}[e^{-r\tau_{y_0+\varepsilon}}]=0$) 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 [22, page 622] show that if $y=\{y_t,\mathcal{F}_t; 0\le t<\infty\}$ is the geometric Brownian motion, then for any $y\ge y_0$

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

where

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

is a strictly positive constant.[5] 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 constraint binds at $\tau_{y_0+\varepsilon}$, the agent's continuation value at $\tau_{y_0+\varepsilon}$ equals his autarky value

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

$V_{aut}(y_0 + \varepsilon)$. 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

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

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 $t \ge 0$, 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

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

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

If the utility function u is given by a closed-form expression, the optimal contract can be characterized more closely. The following example obtains a closed-form expression for the class of utility functions satisfying constant relative risk aversion (CRRA).

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

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

So

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

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

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

Thus, with log preferences, the agent consumes a constant fraction of his to-date maximal income m_t. Similar calculations show that the optimal consumption process has the same structure under any CRRA utility function. In particular, if $u(c) = (1-\gamma)^{-1}c^{1-\gamma}$ with $\gamma > 0, \gamma \neq 1$, then the agent's optimal consumption is given by

$$c_t = m_t \left(\frac{\kappa - (1-\gamma)}{\kappa - (1-\gamma)\alpha}\right)^{\frac{1}{1-\gamma}} \exp\left((1-\gamma)\frac{\sigma^2}{2}\right)$$

at all dates and states. ■

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.[6] 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.

To better understand the structure of the optimal contract, let us discuss how the optimal contract 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

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

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 $V_{aut}(y_0)$.[7]

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. This property of the optimal contract is due to the fact that the agent's autarky value function

[6] See Karatzas and Shreve [23] for proof.

[7] Taking the limit $m \to \infty$ in equation (32) 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 $V_{aut}(y_0)$.

does not have a maximum on the support of the agent's income process in our model. As we have seen, the optimal consumption path in our model must increase whenever income and (hence) the autarky value reach a new all-time maximum. For any m_t, y_t and $s > 0$, the probability of $y_{t+s} > m_t$ is strictly positive, so the consumption path never settles permanently. If the support of agents' income process were bounded from above in our model, the agent's consumption path would be permanently stabilized after income hits its upper bound for the first time.[8]

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

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

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

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

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

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

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 [10] and Ljungqvist and Sargent [11] 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* $y = m$*;*

[8] In general, a committed principal will provide the agent with full insurance if and when the agent's outside option attains its highest possible value. For example, if the agent's outside option value equals 1 for all $y_t < K$ and equals 2 for all $y_t \geq K$ with some $K > y_0$, then the agent obtains full insurance at time $\tau_K = \min\{t : y_t = K\}$. After τ_K, new all-time maxima that income will attain will not increase the agent's consumption because his outside option is not further improved when these maxima are attained.

(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 $0 < dv_t/dy_t < V'_{aut}(y_t)$.[9] Then, the first conclusion in the above lemma tells us that the optimal contract never provides full insurance, and provides no insurance if and 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 $\{y^1_s; 0 \leq s \leq t\}$ and $\{y^2_s; 0 \leq s \leq t\}$ such that $y^1_t = y^2_t$ but $m^1_t > m^2_t$. Consider the continuation value v^i_t that the optimal contract delivers to the agent under past income history $\{y^i_s; 0 \leq s \leq t\}$ for $i = 1, 2$. Because $\bar{u}$ is strictly increasing, we have $u(c^1_t) = \bar{u}(m^1_t) > \bar{u}(m^2_t) = u(c^2_t)$, i.e.,

[9] 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.

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.[10] 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). To obtain the optimal contract in this case, it is enough to modify the 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

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

[10] 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.

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.[11] 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, the agent's consumption is constant under the optimal contract. The principal's cost to deliver a continuation value v to an agent with current income y is given by

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

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

[11] 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 $x \geq m$.

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: financial buffer stock and hedging

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. Under the two-account mechanism the agent has much more control over his consumption than he has under the direct long-term swap contract. Yet, we show that for an appropriate choice of the initial bank account balance and the credit line process, the final allocation is the same as the optimal allocation given in (17).

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 [1].[12] The partial-equilibrium implementation result that we present is a restricted version of the general-equilibrium decentralization result obtained in Alvarez and Jermann [1]. 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.

The implementation exercise we conduct in this section identifies a part of the optimal long-term contract with a natural notion of financial wealth. This decomposition of the optimal contract into an insurance component and a savings component allows us to compare our model with the standard incomplete-markets Bewley-Aiyagari model. We show that the sign of the correlation between financial wealth and income is opposite in these two models. In addition, comparing optimal trading arrangements under limited and full commitment we show that the

[12]See also Krueger and Perri [12] and Krueger and Uhlig [13]. Albanesi and Sleet [24] consider a similar implementation in an economy with full enforcement, private information, and taxes.

borrowing constraint is the only difference between the two.

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 absolute value of 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 because $\mathbb{E}[\beta_t dw_t] = \mathbb{E}[\beta_t(w_{t+dt} - w_t)] = 0$. Thus, the fair-odds price of the hedging asset is zero.[13]

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

We now show how this two-account trading mechanism can be used to implement the consumption process obtained in the optimal long-term contracting problem with one-sided

[13] 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 $\mathbb{E}[\beta_t dy_t] = \beta_t \mu y_t dt$ so as to break even. The formulation we adopt is simpler because $\mathbb{E}[\beta_t dw_t] = 0$ 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.

commitment. In that problem, the agent had an option to stop participating (default) at any time. Here, likewise, at any point in time the agent has the option to exit, i.e., to stop trading with the principal 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, if any, to the principal.

We now show that the optimal consumption process c_t given in (17), combined with some trading strategy $\{\beta_t; t \geq 0\}$ and asset level process $\{A_t; t \geq 0\}$, solves the agent's utility maximization problem.

Proposition 1 *Suppose the borrowing constraint is given by*

$$B_t = \frac{C(y_t, V_{aut}(y_t))}{r}, \tag{25}$$

and the agent's initial assets are

$$A_0 = \frac{C(y_0, \bar{V})}{r}. \tag{26}$$

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

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

$$A_t = \frac{Z(y_t, \bar{m}_t)}{r}, \tag{27}$$

$$\beta_t = \frac{Z_y(y_t, \bar{m}_t)\sigma y_t}{r}, \tag{28}$$

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 (25) is our model's version of the solvency constraints of Alvarez and Jermann [1]. In the discrete-time model of Alvarez and Jermann, these solvency constraints are complicated state-contingent restrictions on portfolios of Arrow securities. In our continuous-time model, these constraints take the simple form of a credit limit.

Our framework allows for a clear characterization of the optimal credit limit. The expression in (25) succinctly expresses it in terms of current income alone: $B_t = B(y_t)$ with $B(\cdot) = C(\cdot, V_{aut}(\cdot))/r$.[14] In addition, (25) shows 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 (26) 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.

[14] In particular, the size of the credit limit does not depend on the agent's current asset position or his history of past income. The function $C(\cdot, V_{aut}(\cdot))/r$ is a unique representation of the optimal credit limit process as a continuous function of current income alone. That is, one can show that if B_t is an optimal credit limit process and $B_t = B(y_t)$ for some continuous function $B(\cdot)$, then $B(\cdot) = C(\cdot, V_{aut}(\cdot))/r$.

The two-account trading mechanism could also be used in a full-commitment environment to implement the optimal consumption process $c_t = u^{-1}(\bar{V})$ giving the principal the maximum profit $-C^f(y_0, \bar{V})$, where $C^f(y_0, \bar{V})$ is given in (19). It is easy to check that in that case, similar to (27) and (28), the agent's equilibrium asset holdings would be given by $A_t = C^f(y_t, \bar{V})/r$ and the hedging process would be $\beta_t = -\sigma y_t/(r - \mu)$. However, no borrowing constraint would be necessary in the full-commitment environment.[15] The borrowing constraint, therefore, is the only difference between the implementing mechanisms in the full-commitment environment and the one-sided commitment model.

Proposition 1 lets us better understand the structure of the optimal long-term risk-sharing contract by decomposing it into a saving/borrowing component and an insurance/hedging component. Perhaps surprisingly, it shows that the limited commitment friction does not necessitate in our model any restrictions on the size of the agent's hedging position. This property depends on the path-continuity of the agent's income process. We discuss this property in Section 5.

4.3 Financial wealth, savings buffer stock, and hedging

Total wealth and financial wealth We begin by defining agent's total wealth W_t as the sum of his financial wealth and his human wealth: $W_t = A_t + P(y_t)$ at all t. In equilibrium, agent's total wealth fluctuates with his current income level.

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

Proof In Appendix A. ■

This lemma shows that, despite the fact that consumption is constant when income y_t fluctuates below $\bar{m}_t$, the expected present value of the agent's future consumption responds to changes in y_t at all times. The agent's budget constraints imply that, at all t, the expected present value of his future consumption is equal to his total wealth. Thus, since total wealth positively co-varies with current income, so does the expected future consumption.

In the full-commitment environment, where the optimal consumption process is permanently constant (full insurance), the agent's total wealth in the implementing mechanism would be constant at all t, i.e., would not respond to changes in income. This means that under full insurance, the changes in human wealth brought about by fluctuations in y_t would be perfectly offset by the changes in the financial wealth of the agent. With limited commitment, they are offset only partially.

The buffer stock of savings 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 CRRA preferences, for example), 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

[15] In order to eliminate Ponzi schemes, it would be necessary to require that $\lim_{t\to\infty} \mathbb{E}[e^{-rt}A_t] \geq 0$. That constraint, however, would never bind in equilibrium.

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 his 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.

Lemma 5 shows a fundamental difference in the role of financial wealth in complete- and incomplete-markets models. We discuss this difference in the next subsection.

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. 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.

Hedging As evident in (24), the hedging position β_t represents the sensitivity of the agent's financial wealth A_t to shocks dw_t. Applying Ito's lemma to $P(y_t)$ given in (2), we have

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

which shows that the sensitivity of the human wealth $P(y_t)$ to shocks dw_t is represented by $y_t\sigma/(r-\mu)$. Let us thus define the agent's hedging ratio h_t^A as

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

Under full insurance, the agent's hedging ratio would equal one at all dates and states. In this way, agent's total wealth W_t would be stabilized in the implementation of the full-insurance allocation. 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: 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. The next lemma tell us more about the structure of this response.

Lemma 7 *At all t, h_t^A is decreasing in y_t and approaches one as y_t approaches zero. Also, define the (implicit) hedging ratio h_t^B characterizing 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 satisfies*

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

with equality if and only if $y_t = \bar{m}_t$.

Proof In Appendix A. ■

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. The agent's hedging ratio h_t^A however never falls below a lower bound given by h_t^B.

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 Comparing with self-insurance

If the agent had access to the bank account only, as is assumed in the standard incomplete-markets Bewley-Aiyagari model, then his flow budget constraint would be given by

$$dA_t = (rA_t - c_t + y_t)dt. \tag{29}$$

In this subsection, we make two points. First, financial wealth fulfills a very different role in complete- and incomplete-markets models. Second, the optimal mechanism in the limited-commitment model provides the agent with better insurance than what the agent could obtain through self-insurance, i.e., self-insurance is suboptimal.

Role of financial wealth Lemma 5 lets us understand 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, e.g., in Bewley-Aiyagari incomplete-markets models, the role of financial wealth 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, in contrast, the agent can insure shocks to his income via hedging. The role of financial wealth is not to buffer off the income shocks but rather to buffer off the losses that the agent's hedging activity may generate. 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.

Degree of consumption insurance The negative correlation between income and 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, at all t, the expected present value of the agent's future consumption flow is equal to his current total wealth stock W_t. Recall from the previous subsection that total wealth is defined as the sum of human wealth and financial wealth: $W_t = P(y_t) + A_t$. 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 negative correlation between income and assets (Lemma 5) means that the covariance term $cov_t(dA_t, dP(y_t))$ is negative. The conditional 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 conditional variance of the change of human wealth, $var_t(dP(y_t))$.[16] Because dA_t is deterministic (has zero volatility) in the incomplete-markets structure represented by (29), $cov_t(dA_t, dP(y_t))$ is identically equal zero, and, thus, $cov_t(dW_t, dP(y_t))$ is in that model the same as $var_t(dP(y_t))$, not smaller.[17]

Intuitively, because the volatility term of dA_t is identically zero in the incomplete-markets model, the agent cannot reduce the volatility of his total wealth by just adjusting his financial wealth in response to shocks to his human wealth. Therefore, no wealth process the agent might choose under self-insurance can be consistent with the optimal consumption process (17).[18]

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 $\mathbb{E}[e^{-r\tau_\varepsilon}]$ 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

[16]In the model with full commitment, there are no borrowing limits constraining the dynamics of financial wealth. Therefore, the process A_t can be chosen so that $cov_t(dA_t, dP(y_t)) = -var_t(dP(y_t))$ at all t to obtain full insurance: $cov_t(dW_t, dP(y_t)) = 0$.

[17]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 $\mathbb{E}_t[dA_t dP(y_t)]$ is negative in a typical incomplete-markets model, while zero in autarky.

[18]In particular, Lemma 6 shows that, in the implementation, financial wealth has non-zero volatility at $y_t = \bar{m}_t$, i.e., the optimal allocation is not the same as the optimal self-insurance allocation even locally at times when the borrowing constraint binds.

characterizing the optimal contract remains valid.[19]

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 property 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 model with discontinuous income paths (for example, with discrete income shocks arriving as a Poisson process), individual shocks can be large (at points of time path discontinuity) and gambles with large upside potential are possible. As a result, asset paths can 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

[19] For example, if the log of the income process is an Ornstein-Unlenbeck mean-reverting process, the formula for the derivative of $\mathbb{E}[e^{-r\tau_\varepsilon}]$ can be obtained from Borodin and Salminen [22, page 524, formula 2.0.1].

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 risk sharing with multi-sided commitment frictions. Analysis of the multi-sided case, however, is more challenging because in that case the pattern of binding participation constraints must be determined for multiple agents at the same time.

To see this, note that in the one-sided case studied in this paper we proceed in three steps. We first compute how soon the agent's participation constraint binds and by how much the agent's continuation value must be increased at that time. Second, we compute the agent's initial level of consumption given the timing and the magnitude of the future consumption increases. Third, we use this information to compute the principal's profit from the relationship. If the principal is unable to commit, we can no longer use this three-step procedure. In the contracting problem in which the principal's expected continuation profit must be non-negative at all times (i.e., adding the principal's participation constraints to our model), there are two possible events that necessitate a change in the level of the agent's consumption. One, as in the one-sided case, is the event in which the agent's income process hits a new all-time high and so his consumption must be increased to ensure his continuing participation. The other is the event in which the agent's income hits a level so low that the principal's expected continuation profit drops to zero and so the agent's consumption must be decreased in order to ensure the principal's continuing participation. The likelihood, timing and magnitude of these future consumption increases and decreases determine the agent's initial consumption level, for any fixed amount of ex ante utility that the contract is to deliver to the agent. Computation of the likelihood, the timing, and the magnitude of a consumption decrease, however, is more challenging than that of an increase because it is not a priori known how low income must fall in order to trigger a consumption cut, whereas we know that a consumption hike will be necessary each time income attains a new all-time maximum. The larger the principal's profit from the contract, the lower the agent's income can fall without driving the profit to zero and triggering a consumption cut. Thus, in the two-sided case the agent's initial consumption level cannot be determined without simultaneously determining the principal's profit from the contract.

Given these challenges, the continuous-time setup seems particularly useful to study the multi-sided case because it simplifies the computation of the hitting times that determine the optimal contract. Further, it typically is easier to study the differential equation that in continuous time characterizes the principal's profit function than it is to perform value function iteration on the principal's profit function in discrete time.

6 Conclusion

We view our analysis in this paper as making two contributions. First, we provide a closed-form characterization of the optimal long-term risk-sharing contract in a dynamic environment in which the insured agent has a limited ability to commit. We build our construction of the optimal contract on a simple observation that it is efficient for the principal to provide the agent with a constant level of consumption whenever the agent's income process is not at its all-time high. The maximum level of income attained to-date, therefore, is the only state variable needed to determine the agent's current consumption. The geometric Brownian motion structure of the agent's income process allows us to give a simple formula, (12), mapping this state variable into the optimal level of consumption. This formula lets us characterize precisely the dynamics of the agent's continuation value and the principal's profit under the optimal long-term contract.

Second, we relate our results to the literature studying borrowing constraints as a tool to mitigate the risk of borrower default. Existing models deliver optimal borrowing constraints in the form of complicated restrictions on portfolios of state-contingent assets. Our model shows that simple borrowing constraints—literally, limits on the amount that agents can borrow—emerge as the implication of limited borrower commitment in a continuous-time model of optimal risk sharing. In our model, we show that the optimal credit limit equals the total value of the surplus generated by the relationship between the principal and the agent. We closely characterize the dynamics of the financial wealth and the hedging position held by the agent. Both are negatively correlated with current income. 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

$$V_{aut}(y_0) = \int_0^\infty u(y) f(y_0, y) dy, \tag{30}$$

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 [22, 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 (30) 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

$$
V(y, m) = -\int_y^m \bar{u}(m) d\left(\frac{y}{x}\right)^{\kappa} + \left(\frac{y}{m}\right)^{\kappa} V_{aut}(m), \tag{31}
$$

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

$$
V_{aut}(y) = -\int_y^m \bar{u}(x) d\left(\frac{y}{x}\right)^{\kappa} + \left(\frac{y}{m}\right)^{\kappa} V_{aut}(m), \text{ for any } m \ge y > 0. \tag{32}
$$

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 (31) and (32) we have

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

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 $V_y(y,m) - V'_{aut}(y) \leq 0$ 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 $V_m(y,m) = -\int_y^m \bar{u}'(m) d\left(\frac{y}{x}\right)^\kappa \geq 0$ 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

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 $\mathbb{E}\left[\int_0^\infty re^{-rs}u(\tilde{c}_t)dt\right] > \bar{V}$. 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 $\mathbb{E}\left[\int_0^\infty re^{-rs}u(\tilde{c}_{t+s})ds|\mathcal{F}_t\right]$ 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 principal's cost is still A_0 because the principal'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

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

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

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

Thus

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

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). ■

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

$$dy_t = \mu y_t dt + \sigma y_t dw_t, \tag{33}$$

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 [17] 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 2 *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

$$dv_t = r(v_t - u(c_t))dt + Y_t dw_t. \tag{34}$$

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

Proof See Sannikov [17]. ■

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 $\mathbb{E}_t[Y_t dw_t] = 0$.

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 $\omega \in \Omega$,[20] we have

$$c_t(y_0, \bar{V}) = c_0(y_t, v_t), \tag{35}$$

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

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, (35) implies that the sensitivity components of these representations are the same a.e.:

$$Y_t(y_0, \bar{V}) = Y_0(y_t, v_t), \tag{36}$$

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 (35) and (36), these two functions (the so-called policy rules) can be used in (34) 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 [25, equation (5.8), page 165]):

$$\begin{aligned} rC(y_t, v_t) &= \min_{c,Y} \Big\{ r(c - y_t) + C_y(y_t, v_t)\mu y_t + C_v(y_t, v_t) r(v_t - u(c)) \\ &\quad + \frac{\sigma^2 y_t^2}{2} C_{yy}(y_t, v_t) + \sigma y_t Y C_{vy}(y_t, v_t) + \frac{Y^2}{2} C_{vv}(y_t, v_t) \Big\}, \end{aligned} \tag{37}$$

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 (37).

Proposition 3 *$J(y, v)$ 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

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

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

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

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

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

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 4 *$J = C$, 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 $m_t < N$,

$$c_t = u^{-1}(\bar{u}(m_t)), \tag{38}$$

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

$$\bar{u}(y) = V_{aut}(y) - \kappa^{-1} y V'_{aut}(y). \tag{39}$$

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

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

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

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

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

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

We get

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

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 $\mathbb{E}[\sum_{t=0}^\infty (1-\beta)\beta^t u(c_t)]$, 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 $i > j$.[21] 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

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

[21] Zhang [26] generalizes the results of this appendix by relaxing the first-order stochastic dominance assumption and allowing for a general outside option value function.

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.[22] 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 $m_t = \max_{0 \le s \le t} y_s$ be the to-date maximal realized income. We construct a contract as

$$c_t = u^{-1}(\bar{u}(m_t)). \tag{40}$$

Note that the sequence $\{c_t, t \ge 0\}$ 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 $Z(y_t, m_t) \ge Z(y_t, y_t) = V_{aut}(y_t)$ 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 [11, page 660], let $\beta^t \alpha_t$ be the multiplier on

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

and let ϕ be the multiplier on

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

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 [11, equation 19.4.6a] are satisfied by this construction. The analogs of the complementary slackness conditions in Ljungqvist and Sargent [11, equation 19.4.6b] are satisfied as well. This is because when $\alpha_t > 0$, 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 (40). 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 $\Pr(y_{t+1} = \frac{1}{2^{t+1}} | y_t = \frac{1}{2^t}) = 1, t \ge 1$. 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 $t \ge 1$ 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

[22] Proof is available upon request.

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

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

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

[1] F. Alvarez, U. Jermann, Efficiency, Equilibrium, and Asset Pricing with Risk of Default, Econometrica 68 (4) (2000) 775–797.

[2] R. Albuquerque, H. Hopenhayn, Optimal Lending Contracts and Firm Dynamics, Review of Economic Studies 71 (2004) 285–315.

[3] P. J. Kehoe, F. Perri, Competitive equilibria with limited enforcement, Journal of Economic Theory 119 (1) (2004) 184 – 206.

[4] W. Rogerson, Repeated Moral Hazard, Econometrica 53 (1) (1985) 69–76.

[5] M. Golosov, N. Kocherlakota, A. Tsyvinski, Optimal Indirect and Capital Taxation, Review of Economic Studies (2003) 569–587.

[6] M. Harris, B. Holmstrom, A theory of wage dynamics, Review of Economic Studies 49 (1982) 315–333.

[7] J. Thomas, T. Worrall, Self-Enforcing Wage Contracts, Review of Economic Studies 55 (1988) 541–554.

[8] A. Marcet, R. Marimon, Communication, commitment, and growth, Journal of Economic Theory 58 (2) (1992) 219 – 249.

[9] T. Kehoe, D. Levine, Debt-Constrained Asset Markets, Review of Economic Studies 60 (4) (1993) 865–888.

[10] N. Kocherlakota, Implications of Efficient Risk Sharing without Commmitment, Review of Economic Studies 63 (4) (1996) 595–609.

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

[12] D. Krueger, F. Perri, Does Income Inequality Lead to Consumption Inequality? Evidence and Theory, Review of Economic Studies 73 (2006) 163–193.

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

[14] A. Monge-Naranjo, Entrepreneurship and firm heterogeneity with limited enforcement, Annals of Finance 5 (2009) 465–494.

[15] P. Demarzo, Y. Sannikov, Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model, The Journal of Finance 61 (2007) 2681–2724.

[16] B. Biais, T. Mariotti, G. Plantin, J. Rochet, Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications, Review of Economic Studies 74 (2007) 345–390.

[17] Y. Sannikov, A Continuous-Time Version of the Principal-Agent Problem, Review of Economic Studies 75 (2008) 957–984.

[18] T. Piskorski, A. Tchistyi, Optimal Mortgage Design, Review of Financial Studies 23 (3) (2010) 3098–3140.

[19] Z. He, Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion, Review of Financial Studies 22 (2009) 859–892.

[20] Y. Zhang, Dynamic Contracting with Persistent Shocks, Journal of Economic Theory 144 (2009) 635–675.

[21] J. Eaton, M. Gersovitz, Debt with Potential Repudiation: Theoretical and Empirical Analysis, Review of Economic Studies 48 (2) (1981) 289–309.

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

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

[24] S. Albanesi, C. Sleet, Dynamic Optimal Taxation with Private Information, Review of Economic Studies 73 (2006) 1–30.

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

[26] Y. Zhang, Characterization of a Risk Sharing Contract with One-Sided Commitment, Working Paper Univeristy of Iowa .