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

Borys Grochulski† Yuzhe Zhang‡

May 19, 2011

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.

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, 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

*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 [3], Albuquerque and Hopenhayn [2], Kehoe and Perri [13].

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 autarky 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 autarky value. Note now that even when the agent cannot commit, the full-insurance contract does not cause the agent to default (revert to autarky) 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 autarky value—being strictly increasing in income—fluctuates below the agent's initial autarky 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 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 autarky 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 autarky 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 it 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 the volatility of his income process.

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 [22], Golosov et al. [9]). 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. Our implementation delivers an optimal borrowing limit derived from the underlying commitment friction.

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 [10], Thomas and Worrall [24], Marcet and Marimon [19], Kehoe and Levine [14], Kocherlakota [15], Alvarez and Jermann [3], Albuquerque and Hopenhayn [2], Ljungqvist and Sargent [18], Krueger and Perri [16], Krueger and Uhlig [17]. 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] Our method for the characterization of the optimal contract, however, is not specific to our continuous-time framework. Zhang [26] shows how our method can be extended to a discrete-time model with a general Markov income process and a general outside option value function.

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 [23]. Our analysis suggests that limited-

[2]Monge-Naranjo [20] 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 [6], Biais et al. [4], Sannikov [23], Piskorski and Tchistyi [21], He [11], Zhang [25].

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 the 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 [7]). 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 [3] 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 [3] 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 and finance. 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 value, we can easily compute the borrowing constraints with no need for the fixed-point iteration procedure used in Alvarez and Jermann [3].

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. 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 provides an application of our method to an optimal contracting problem with two-sided lack of commitment.

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 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 that ensures 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

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

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

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 under full commitment, 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 $V_{aut}(y_0)$. Under limited commitment, however, 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 [5, 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

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

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

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

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.

Figure 1 shows a sample path of income y_t along with the corresponding path of the state variable m_t and the optimal consumption path c_t. Clearly, the path for c_t is non-decreasing and increases when the path for m_t does. In fact, because the utility function is CRRA in this example, c_t is a constant fraction of m_t.

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),

[6] See Karatzas and Shreve [12] for proof.

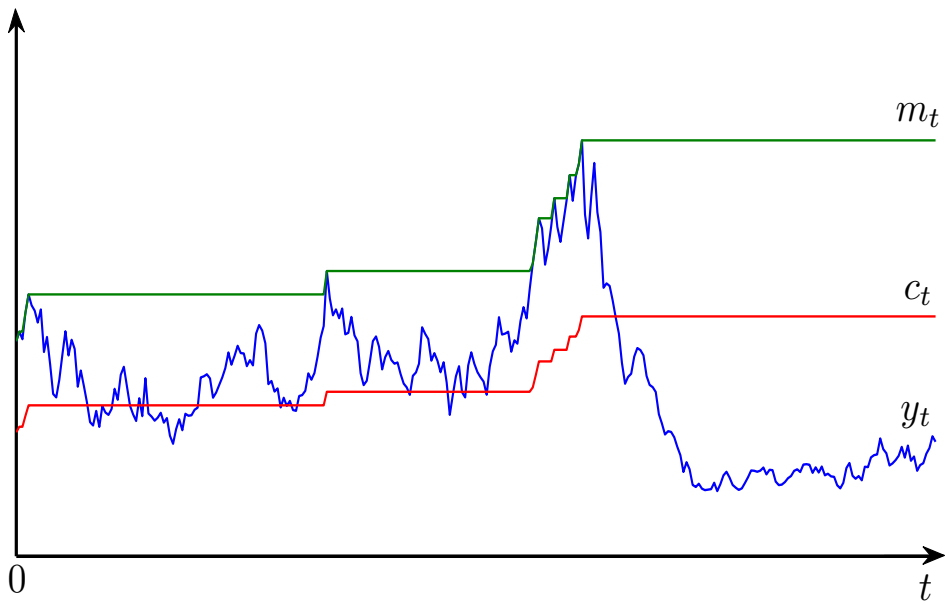

Figure 1: Optimal consumption under CRRA preferences.

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 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

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

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 [15] and Ljungqvist and Sargent [18] describe as amnesia: history does not matter, i.e., the continuation contract is the same for all paths of past income $\{y_s; 0 \leq s < t\}$.

Lemma 2 *The function V satisfies*

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

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

[8] In general, a committed principal will provide the agent with permanent full insurance starting at a point when the agent's outside option attains its highest possible value for the first time. 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 permanent full insurance as of time $\tau_K = \min\{t : y_t = K\}$. After τ_K, new all-time maxima that income may attain will not increase the agent's consumption because his outside option is not further improved when these maxima are attained.

(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., the agent's utility flow at t is larger under the income history $\{y^1_s; 0 \leq s \leq t\}$. The same

[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.

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, 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 and total contract surplus

In this subsection, we describe the principal's continuation cost under the optimal contract c given in (11). In particular, we show that the total surplus of the relationship between the principal and the agent is strictly positive. In Section 4, we will show that this surplus represents the size of the optimal borrowing limit for the agent.

Recall first 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 therefore given by

$$C^f(y, v) = u^{-1}(v) - rP(y), \tag{18}$$

[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.

where $u^{-1}(v)$ is the constant consumption level needed to deliver promised utility v. In the limited commitment case, 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

$$Z(y,m) = \left(1 - \left(\frac{y}{m}\right)^{\kappa}\right) u^{-1}(\bar{u}(m)) + \left(\frac{y}{m}\right)^{\kappa} \int_m^{\infty} u^{-1}(\bar{u}(x)) d\left(1 - \left(\frac{m}{x}\right)^{\kappa}\right) - rP(y). \tag{19}$$

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).

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 (19) 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 $P(y) =$

[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 $x \geq m$.

$y/(r-\mu)$. 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: savings and hedging accounts

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 [3].[12] The partial-equilibrium implementation result that we present is a restricted version of the general-equilibrium decentralization result obtained in Alvarez and Jermann [3]. Tractability is an advantage of our continuous-time model. We are able to characterize the solvency constraints in detail. In particular, we show that optimal solvency constraints take in our model a simple form of a borrowing constraint. In addition, by comparing optimal trading arrangements under limited and full commitment, we show that the borrowing constraint is the only difference between the two.

[12] See also Krueger and Perri [16] and Krueger and Uhlig [17]. Albanesi and Sleet [1] consider a similar implementation in an economy with full enforcement, private information, and taxes.

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 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

[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.

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:

$$\begin{aligned} c_t &= u^{-1}(\bar{u}(\bar{m}_t)), \\ A_t &= \frac{Z(y_t, \bar{m}_t)}{r}, \qquad (27) \\ \beta_t &= \frac{Z_y(y_t, \bar{m}_t)\sigma y_t}{r}, \qquad (28) \end{aligned}$$

where $\bar{m}_t = \max\{\max_{0 \leq s \leq t} y_t, \bar{m}_0\}$*,* $\bar{u}$ *is given in (12),* Z *is given in (19), 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 [3]. 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.

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

[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 $B(\cdot) = C(\cdot, V_{aut}(\cdot))/r$.

profit $-C^f(y_0, \bar{V})$, where $C^f(y_0, \bar{V})$ is given in (18). 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 continuity of the agent's income path. We discuss this property in the next section.

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. As already mentioned, Zhang [26] shows that our method can be used to study discrete-time models. In addition, our characterization extends to other models with continuous time. 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 characterizing the optimal contract remains valid.[16]

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 account balance process $\{A_t; t \geq 0\}$. 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 paying off enormous amounts in some states of nature and requiring 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, which 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

[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.

[16] 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 [5, page 524, formula 2.0.1].

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. Because the agent cannot generate a discontinuous time path for his bank account balance, he cannot violate his borrowing constraint by a meaningful amount. Continuity of the agent's income path 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 income path discontinuity) and gambles with large upside potential are possible. As a result, asset position 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 contract and its implementation, which we provide in Section 3 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 is more challenging because the pattern of binding participation constraints must be determined for multiple agents at the same time. The continuous-time setup seems particularly useful in the multi-sided case because it greatly simplifies the computation of the hitting times that determine this pattern. To illustrate this point with a concrete example, in Appendix C we solve a continuous-time version of the mutual insurance problem with two-sided lack of commitment similar to Kocherlakota [15].

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.

Appendix A. Proofs

Proof of Lemma 1

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

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

where $f(y_0, y)$ is the density of the expected discounted amount of time that the income process starting from y_0 spends at each level $y \in (0, \infty)$. From Borodin and Salminen [5, page 132], we know that

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

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

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

Then

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

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

Proof of Lemma 2

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

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

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

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

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

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

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

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

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

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

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) We have

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

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 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., $Z_m(y,m) \geq 0$. ■

Appendix B. Verification of optimality

This appendix provides a formal verification of the optimality of the contract (17).

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

By Ito's formula, y_t satisfies

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

where $\mu = \alpha + \sigma^2/2$. In this representation, the income process is decomposed into a drift and a volatility component. The same decomposition can be provided for the agent's continuation value process v_t. In particular, the following proposition of Sannikov [23] 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

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

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

Proof See Sannikov [23]. ■

In this representation, $r(v_t - u(c_t))$ is the drift of the promised utility process v_t and Y_t is the sensitivity of v_t to income shocks dw_t.

In our problem, the Dynamic Principle of Optimality implies that efficient contracts in Ψ are representable by a pair of real-valued policy rules $(c(y_t, v_t), Y(y_t, v_t))$, where $c : \Theta \to \mathbb{R}_{++}$ and $Y : \Theta \to \mathbb{R}$. With these policy rules we can express the law of motion for the state vector (y_t, v_t) as

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

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 (see, for example, Fleming and Soner [8, equation (5.8), page 165]):

$$\begin{aligned} rC(y_t, v_t) &= \min_{c,Y} \Big\{ r(c - y_t) + C_y(y_t, v_t)\mu y_t + C_v(y_t, v_t) r(v_t - u(c)) \\ &\quad + \frac{\sigma^2 y_t^2}{2} C_{yy}(y_t, v_t) + \sigma y_t Y C_{vy}(y_t, v_t) + \frac{Y^2}{2} C_{vv}(y_t, v_t) \Big\}, \end{aligned} \tag{34}$$

where subscripts on C denote partial derivatives, and at the boundary $v_t = V_{aut}(y_t)$ 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 (17), $Z(y, M(y, v))$, by $J(y, v)$. We can now show that $J(y, v)$ satisfies the HJB equation (34).

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

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

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

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

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

To see the first of these equalities, recall from the proof of Lemma 2(iii) that $V_m = -\int_y^m \bar{u}'(m)d(\frac{y}{x})^\kappa$. From (19) we calculate the partial derivative $Z_m = -\int_y^m (u^{-1})'(\bar{u}(m))\bar{u}'(m)d(\frac{y}{x})^\kappa$. 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 (17) is efficient.*

Proof Let $N > 0$ be any positive number and define $\Theta^{(N)} = \{(y,v) \in \Theta : 0 < y \leq N, v \leq V_{aut}(N)\}$. Pick an initial condition $(y,v) \in \Theta^{(N)}$ and consider an auxiliary dynamic programming problem in which we remove the participation constraints 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 removing 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 satisfies

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

where $\bar{m}_t = \max\{m_t, M(y,v)\}$. We define, for $m \in [y, N]$,

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

First, we show that 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 $\Theta^{(N)}$,

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

Pick any contract $\{\tilde{c}_t; t \geq 0\}$ starting from the initial condition $(y,v) \in \Theta^{(N)}$. Denote the volatility term of $\tilde{v}_t$ in Proposition 2 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

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

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

$$\mathbb{E}[G_{t\wedge T_n}] \geq G_0 = J^{(N)}(y,v).$$

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

$$\mathbb{E}\left[\int_0^\lambda re^{-rs}(\tilde{c}_s - y_s)ds\right] = \lim_{n\to\infty}\mathbb{E}\left[\int_0^{n\wedge T_n} re^{-rs}(\tilde{c}_s - y_s)ds\right]. \tag{35}$$

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 $e^{-r\lambda}J^{(N)}(y_\lambda,\tilde{v}_\lambda) = e^{-r\lambda}J^{(N)}(y_\lambda, V_{aut}(N))$, and the bounded convergence theorem implies

$$\mathbb{E}\left[e^{-r\lambda}J^{(N)}(y_\lambda, V_{aut}(N))\right] = \lim_{n\to\infty}\mathbb{E}\left[e^{-r(n\wedge T_n)}J^{(N)}(y_{n\wedge T_n},\tilde{v}_{n\wedge T_n})\right]. \tag{36}$$

Combining (35) and (36), we get

$$\mathbb{E}\left[\int_0^\lambda re^{-rs}(\tilde{c}_s - y_s)ds + e^{-r\lambda}J^{(N)}(y_\lambda, V_{aut}(N))\right] = \lim_{n\to\infty}\mathbb{E}[G_{n\wedge T_n}] \geq J^{(N)}(y,v).$$

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

Second, since the auxiliary problem has less constraints than the original problem, the cost of the auxiliary problem is below that of the original problem, i.e., for all $N > 0$,

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

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

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

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

Appendix C. Optimal risk sharing without commitment

Consider two agents with identical period utility functions $u(c) = -\exp(-c)$. The income process of agent $i = 1,2$, denoted by $\{y_t^i; t\geq 0\}$, is given by, respectively, $y_t^1 = \sigma w_t$ and $y_t^2 = -\sigma w_t$, where w_t is a standard Brownian motion, and $\sigma > 0$. Hence, there is no aggregate risk and

the two agents are symmetric. We are looking for an optimal risk-sharing contract between these two agents. As in Kocherlakota [15], we assume that neither agent can commit, so the contract (allocation of consumption) must be self-enforcing: at all times both agents' participation constraints (PCs) must be satisfied, which means that for both agents the continuation value under the contract must be at least as large as the value of reverting to permanent autarky. Assuming that the agents' common rate of time preference r is larger than $\sigma^2/2$, the autarky value function, which by symmetry is the same for both agents, can be computed here in closed-form:

$$V_{aut}(y_t^i) = u(y_t^i)F, \tag{37}$$

where $F = \left(1 - \frac{\sigma^2}{2r}\right)^{-1} > 1$.

Because $y_t^1 + y_t^2 = 0$ at all t, any efficient allocation must have $c_t^2 = -c_t^1$. We now state two additional properties that efficient allocations must satisfy in this model: *symmetry* and *scalability*.[17] Suppose at some t agents' incomes are $(y_t^1, y_t^2) = (\bar{y}, -\bar{y})$, and let $\left\{\left(\bar{c}_{t+s}^1, \bar{c}_{t+s}^2\right); s \geq 0\right\}$ be an efficient continuation allocation delivering continuation values $(v_t^1, v_t^2) = (\bar{v}^1, \bar{v}^2)$. The following two properties hold. First, if $(y_{t'}^1, y_{t'}^2) = (-\bar{y}, \bar{y})$ at some t', then the continuation allocation $\left(c_{t'+s}^1, c_{t'+s}^2\right) = \left(\bar{c}_{t+s}^2, \bar{c}_{t+s}^1\right)$ for all $s \geq 0$ is efficient and delivers continuation values $(v_{t'}^1, v_{t'}^2) = (\bar{v}^2, \bar{v}^1)$. Second, if $(y_{t'}^1, y_{t'}^2) = (\bar{y}+k, -\bar{y}-k)$ for some t' and k, then the continuation allocation $\left(c_{t'+s}^1, c_{t'+s}^2\right) = \left(\bar{c}_{t+s}^1 + k, \bar{c}_{t+s}^2 - k\right)$ for all $s \geq 0$ is efficient and delivers continuation values $(v_{t'}^1, v_{t'}^2) = (\exp(-k)\bar{v}^1, \exp(k)\bar{v}^2)$. We will call these two properties, respectively, symmetry and scalability.

We will now use a heuristic argument similar to that in Section 3 to qualitatively characterize efficient allocations of consumption in this environment. In our informal exposition, we will invoke the symmetry and scalability properties described above. After that, as in Section 3 again, we will use a relaxed version of this contracting problem to compute the solution. The closed-form expression for the distribution of a Brownian motion hitting time will be very useful in this computation.

Qualitative properties

Clearly, as there is no aggregate risk in this model, full insurance for both agents is feasible as long as both agents' participation constraints are slack. Any efficient allocation, therefore, will give constant consumption flows to both agents at all times at which no PC binds.

Let us start out from an initial condition in which the PC of agent 1 holds as an equality. This means that agent 1's continuation value at $t = 0$ is equal to $V_{aut}(y_0^1) = V_{aut}(0)$, i.e., agent 1 is indifferent between staying in the contract and defaulting to autarky. Thus, agent 1 receives no surplus from this insurance relationship as of time zero, and agent 2 receives the whole surplus. At this point, we do not know the size of the surplus. We will denote the value that

[17] These properties follow from the assumption that both agents have identical exponential (i.e., CARA) utility function and their income processes are random walks with the same distribution. Due to space constraints, we do not provide formal proofs in this appendix, but we can make them available upon request.

agent 2 obtains under this initial condition by $V_{max}(y_0^2) = V_{max}(0)$.[18] If $V_{max}(0) = V_{aut}(0)$, the surplus from the relationship is zero. This means that each agent can get at most his autarky value, i.e., no insurance can be sustained. If $V_{max}(0) > V_{aut}(0)$, the surplus is positive and the PC of agent 2 is slack, which makes some insurance sustainable.

In particular, starting from this initial condition, insurance can be sustained along those sample paths in which w_t decreases (goes negative) immediately after $t = 0$. When w_t decreases, y_t^1 decreases and y_t^2 increases. The autarky value of agent 1 thus decreases and his PC is relaxed, i.e., it becomes non-binding. At the same time, the autarky value of agent 2 increases and thus his PC becomes tighter. But because the PC of agent 2 was slack at $t = 0$, it will remain slack for some time as w_t becomes more and more negative. As both agents' participation constraints are slack, efficiency requires full risk sharing during that time, so both agents' consumption is constant.

Constant consumption, however, will not be sustainable forever. Consumption will have to be adjusted when one agent's PC becomes binding. Starting from the initial condition described above, the PC of agent 1 will bind as soon as y_t^1 increases above its initial value of zero. Alternatively, if y_t^1 stays below zero, the PC of agent 2 will become binding when w_t becomes sufficiently negative, y_t^2 sufficiently positive, and the autarky value of agent 2 sufficiently high. Let us denote the level of y_t^2 at which this will take place by $2a$. Full risk sharing can be sustained for as long as y_t^1 remains below zero and above $-2a$. If y_t^1 crosses zero, consumption of agent 1 must be increased in order to satisfy his participation constraint. Because aggregate income remains constant, this means that consumption of agent 2 must at that point decrease. If y_t^1 crosses $-2a$, which means that y_t^2 crosses $2a$, consumption of agent 2 must be increased in order to satisfy his participation constraint. At that point, consumption of agent 1 must decrease.

Let us now consider the agents' continuation values when y_t^1 hits $-2a$. Because the PC of agent 2 binds, clearly, his continuation value equals $V_{aut}(2a)$. Symmetry implies that agent 1's continuation value is $V_{max}(-2a)$, as his PC is maximally slack at this point and thus he gets the whole surplus from the relationship. Note that the transfer of the surplus from agent 2 to agent 1 that take place while y_t^1 drops from zero to $-2a$ provides insurance to agent 1.

We can now compute the levels of the constant consumption flow the two agents receive during the time interval in which y_t^1 stays between zero and $-2a$. Agent 1 gets some level x, and agent 2 gets $-x$. If we switched the two agents' positions so that the PC of agent 2 were binding at $t = 0$, then, by symmetry, agent 2 would be getting consumption x. But the PC of agent 2 in fact does bind when y_t^1 hits $-2a$. By scalability, therefore, agent 2's consumption at this point must be $x + 2a$, because y_t^2 equals $2a$ when y_t^1 hits $-2a$. Since consumption is constant over this time interval, agents 2's consumption is the same at time zero and at the time when y_t^1 hits $-2a$, which means $-x = x + 2a$. Solving for x we get that agent 1 consumes

[18] In general, we will define $V_{max}(y_t^2)$ as the maximum value that agent 2 can get out of the insurance relationship with agent 1 when income of agent 2 is y_t^2. By symmetry, $V_{max}(y_t^1)$ equals the maximum value that agent 1 can get out of this insurance relationship when his income is y_t^1. Thus, the same functional form, $V_{max}(\cdot)$, applies to both agents.

$-a$ and agent 2 consumes a over this time interval.

After consumption is updated because of a binding PC for agent 1 or 2, the agents face the same contracting problem as at $t = 0$, with possibly two differences. First, the roles of agent 2 and 1 may be reversed depending on whose PC is currently binding. Second, the income levels of the two agents, although adding up to zero, may be non-zero. By symmetry and scalability, however, optimal risk-sharing starting from this new initial condition takes the same form as it did starting from the original initial condition: consumption will be stabilized until the next time the PC of one of the agents binds.

Following Thomas and Worrall [24] and Ljungqvist and Sargent [18, chap. 20], we can express the efficient allocation we have just descried in terms of a "participation-ensuring" interval $[y_t^1 - a, y_t^1 + a]$, and the following updating rule for consumption c_t^1: Keep consumption c_t^1 constant unless doing so would cause it to fall outside of the participation-ensuring interval $[y_t^1 - a, y_t^1 + a]$. When updating c_t^1, change it by the smallest amount necessary to keep it inside the interval $[y_t^1 - a, y_t^1 + a]$.

In the initial condition described above, the initial participation-ensuring interval is $[-a, a]$. With $c_0^1 = -a$, consumption of agent 1 is at the very bottom of the participation-ensuring interval as of $t = 0$. As soon as y_t^1 exceeds 0, the interval shifts up and the bottom end of it bumps agent 1's consumption up. Consumption of agent 2 at this point decreases. On the other extreme, if y_t^1 does not exceed zero but rather falls to $-2a$, the participation-ensuring interval shifts down to $[-3a, -a]$ and its upper end bumps c_t^1 down, which means that c_t^2 is bumped up, consistent with agent 2's PC binding at this point. As long as y_t^1 does not exceed 0 or fall below $-2a$, however, c_t^1 remains inside $[y_t^1 - a, y_t^1 + a]$, which means that both c_t^1 and c_t^2 stay constant.

To compute the solution, we need to jointly determine the function $V_{max}(\cdot)$ and the number a. This task is simplified in our model by the fact that $V_{max}(\cdot)$ is proportional to $V_{aut}(\cdot)$. To see this, note that the scalability property of this model implies that

$$V_{max}(y_t^i) = V_{max}(0)\exp(-y_t^i) \tag{38}$$

for both i and any y_t^i. From (37) and (38), the ratio

$$\frac{V_{max}(y_t^i)}{V_{aut}(y_t^i)} = -V_{max}(0)F^{-1}$$

does not depend on the level of income y_t^i. Denote this ratio by $g \in (0, 1]$. For both i and any y_t^i, we can now write (38) as $V_{max}(y_t^i) = gV_{aut}(y_t^i)$. Solving for the optimal allocation, thus, boils down to finding two constants: $a \geq 0$ and $0 < g \leq 1$.

Computation

We are now ready to follow the approach from Section 3 of the paper, which lets us characterize a and g. We fix a small $\epsilon > 0$ and consider a relaxed problem. Let τ be the stopping time when income y_t^1 first reaches one of the boundaries of the relaxed participation-ensuring

interval $[-2a, \epsilon]$. Agent 1's expected utility from this contract (equal to his autarky value) can be split into the part before and after time τ as follows:

$$V_{aut}(0) = \mathbb{E}\left[\int_0^\tau re^{-rt}u(-a)dt + e^{-r\tau}1_{\{y_\tau^1=-2a\}}V_{max}(-2a) + e^{-r\tau}1_{\{y_\tau^1=\epsilon\}}V_{aut}(\epsilon)\right],$$

where $1_{\{y_\tau^1=x\}}$ is the indicator set of all paths of Brownian motion $\{w_t; t \geq 0\}$ such that $y_\tau^1 = x$. In this formula, as we have discussed informally earlier, along the sample paths in which y_{1t} reaches $-2a$ before it reaches ε agent 1's continuation utility at $t = \tau$ is $V_{max}(-2a)$ because agent 2's participation constraint binds and the whole surplus is at that point given to agent 1. Using $V_{max}(-2a) = gV_{aut}(-2a)$, we can write the above as

$$V_{aut}(0) = \mathbb{E}\left[\int_0^\tau re^{-rt}u(-a)dt + e^{-r\tau}1_{\{y_{1\tau}=-2a\}}gV_{aut}(-2a) + e^{-r\tau}1_{\{y_{1\tau}=\epsilon\}}V_{aut}(\epsilon)\right]. \tag{39}$$

Because g is not known, we cannot use this equation alone to compute a. This contrasts with the one-sided commitment model studied in Section 3, where we could compute the agent's consumption level before time τ without knowing the size of the total contract surplus $V_{max}(\cdot)$. This is because in the one-sided limited commitment model there is only one possible reason for adjusting consumption: the single uncommitted agent's participation constraint binds. We know that when this happens, the agent's continuation utility equals his autarky value. In Section 3, we used this fact to compute the initial optimal consumption level using just the autarky value function $V_{aut}(\cdot)$ and the expected discounted hitting time $\mathbb{E}\left[e^{-r\tau}\right]$. In the two-sided case, in order to compute an agent's consumption we need to know the continuation value that will be delivered to an agent in the event in which the other agent's participation constraint binds. This value, however, depends on the size of the surplus from the relationship, which a priori is unknown.

We can resolve this, however, by obtaining for agent 2 a condition analogous to (39). Because we assumed that the whole contract surplus as of $t = 0$ goes to agent 2, his total expected utility as of time zero is $V_{max}(0) = gV_{aut}(0)$. In the relaxed problem, this total expected utility value can be split into the part before and after time τ as follows

$$gV_{aut}(0) = \mathbb{E}\left[\int_0^\tau re^{-rt}u(a)dt + e^{-r\tau}1_{\{y_\tau^1=-2a\}}V_{aut}(2a) + e^{-r\tau}1_{\{y_\tau^1=\epsilon\}}gV_{aut}(-\epsilon)\right], \tag{40}$$

which gives us the second equation we need to solve simultaneously for a and g. Note that we have substituted in this formula $gV_{aut}(-\epsilon)$ for $V_{max}(-\epsilon)$, the value of agent 2's continuation utility when y_t^1 reaches ϵ.

Brownian motion allows us to compute the discounted stopping time in closed-form:

$$\begin{aligned}
\mathbb{E}\left[\int_0^\tau re^{-rt}dt\right] &= \frac{1 - e^{(-2a)\sqrt{2r}/\sigma}}{1 + e^{(-2a-\epsilon)\sqrt{2r}/\sigma}}(1 - e^{-\epsilon\sqrt{2r}/\sigma}), \\
\mathbb{E}[e^{-r\tau}1_{\{y_\tau^1=-2a\}}] &= \frac{e^{-2a\sqrt{2r}/\sigma} - e^{(-2a-2\epsilon)\sqrt{2r}/\sigma}}{1 - e^{2(-2a-\epsilon)\sqrt{2r}/\sigma}}, \\
\mathbb{E}[e^{-r\tau}1_{\{y_\tau^1=\epsilon\}}] &= \frac{e^{-\epsilon\sqrt{2r}/\sigma} - e^{(-4a-\epsilon)\sqrt{2r}/\sigma}}{1 - e^{2(-2a-\epsilon)\sqrt{2r}/\sigma}}.
\end{aligned}$$

Following the steps in Section 3, i.e., taking ϵ to zero and simplifying terms, we express (39) and (40) as a system of two non-linear equations with two unknowns (a, g):

$$\begin{aligned} V_{aut}(0) &= \frac{(1-e^{-2a\sqrt{2r}/\sigma})^2}{1+e^{-4a\sqrt{2r}/\sigma}}u(-a) + \frac{2e^{-2a\sqrt{2r}/\sigma}}{1+e^{-4a\sqrt{2r}/\sigma}}gV_{aut}(-2a) \\ &\quad + \frac{1-e^{-4a\sqrt{2r}/\sigma}}{1+e^{-4a\sqrt{2r}/\sigma}}V'_{aut}(0)\frac{1}{\sqrt{2r}/\sigma}, \end{aligned} \tag{41}$$

$$\begin{aligned} gV_{aut}(0) &= \frac{(1-e^{-2a\sqrt{2r}/\sigma})^2}{1+e^{-4a\sqrt{2r}/\sigma}}u(a) + \frac{2e^{-2a\sqrt{2r}/\sigma}}{1+e^{-4a\sqrt{2r}/\sigma}}V_{aut}(2a) \\ &\quad - \frac{1-e^{-4a\sqrt{2r}/\sigma}}{1+e^{-4a\sqrt{2r}/\sigma}}gV'_{aut}(0)\frac{1}{\sqrt{2r}/\sigma}. \end{aligned} \tag{42}$$

This system always has autarky (i.e., $a = 0$ and $g = 1$) as a solution. Although hard to solve analytically, this system is inexpensive to study numerically. Our numerical explorations (details of which we do not include here) show that autarky is the only solution if the volatility parameter σ is low, but solutions that sustain insurance (i.e., such that $a > 0$ and $g < 1$) exist if σ is sufficiently high. This result is very intuitive when we note that the agents' value of autarky is decreasing in σ. With low volatility, the autarky value is high, so the participation constraints are tight to the point that no insurance can be sustained, i.e., autarky is efficient. When we increase the volatility parameter σ, the agents' value of autarky decreases. This makes autarky less desirable and thus relaxes the participation constraints. For sufficiently high σ, autarky becomes inefficient, i.e., agents are able to sustain mutual insurance despite high (full, in fact) persistence of the income shock in this model.

References

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

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

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

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

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

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

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

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

[9] Golosov, M., Kocherlakota, N., Tsyvinski, A., 2003. Optimal indirect and capital taxation. Review of Economic Studies, 569–587.

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

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

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

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

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

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

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

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

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

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

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

[21] Piskorski, T., Tchistyi, A., 2010. Optimal mortgage design. Review of Financial Studies 23 (3), 3098–3140.

[22] Rogerson, W., 1985. Repeated moral hazard. Econometrica 53 (1), 69–76.

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

[24] Thomas, J., Worrall, T., 1988. Self-enforcing wage contracts. Review of Economic Studies 55, 541–554.

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

[26] Zhang, Y., 2010. Characterization of a risk sharing contract with one-sided commitment. Working Paper, Univeristy of Iowa.