



Research Division
Federal Reserve Bank of St. Louis
Working Paper Series



**Moral Hazard and Lack of Commitment in
Dynamic Economies**

**Alexander K. Karaivanov
and
Fernando M. Martin**

Working Paper 2011-030A
<http://research.stlouisfed.org/wp/2011/2011-030.pdf>

October 2011

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

Moral Hazard and Lack of Commitment in Dynamic Economies*

Alexander K. Karaivanov[†] and Fernando M. Martin[‡]

August 3, 2011

Abstract

We revisit the role of limited commitment in a dynamic risk-sharing setting with private information. We show that a Markov-perfect equilibrium, in which agent and insurer cannot commit beyond the current period, and an infinitely-long contract to which only the insurer can commit, implement identical consumption, effort and welfare outcomes. Unlike contracts with full commitment by the insurer, Markov-perfect contracts feature non-trivial and determinate asset dynamics. Numerically, we show that Markov-perfect contracts provide sizable insurance, especially at low asset levels, and are able to explain a significant part of wealth inequality beyond what can be explained by self-insurance. The welfare gains from resolving the commitment friction are larger than those from resolving the moral hazard problem at low asset levels, while the opposite holds for high asset levels.

Keywords: risk-sharing, optimal insurance, lack of commitment, moral hazard, wealth inequality.

JEL classification: D11, E21.

*Parts of this paper were originally included in our companion paper, “Dynamic optimal insurance and lack of commitment”. We thank Gabriele Camera, Michael Golosov, Edward Green, the Associate Editor and an anonymous referee for helpful comments and suggestions. We also thank seminar and conference participants at the UBC-Kobe-SFU Workshop, Penn State, the Vienna Macro Workshop, the Cleveland Fed, the Midwest Macro Meeting, the Society for Economic Dynamics Meeting, the Far-Eastern Meeting of the Econometric Society, and the Canadian Macroeconomics Study Group. Karaivanov acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada.

[†]Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada V5A 1S6. Email: akaraiva@sfu.ca.

[‡]Department of Economics, Simon Fraser University, 8888 University Drive, Burnaby, B.C., Canada V5A 1S6. Email: fmartin@sfu.ca.

1 Introduction

It is well-understood that in many settings risk-sharing can improve on what agents can achieve through self-insurance. Indeed, this provides justification for a variety of social programs (e.g., unemployment insurance) as well as private arrangements (e.g., corporate compensation, personal insurance). The classic approach to studying optimal contracts in dynamic settings with private information assumes the contracting parties have the ability to bind themselves to a lifetime agreement.¹ These contracts, however, have unappealing properties, such as the agents' ultimate "immiserization", exploding inequality, and/or degenerate long-run wealth and consumption distributions (see Phelan, 1998 for a review).

A natural way to avoid those problems is to study insurance arrangements in which at least one contracting party lacks the ability to commit. For example, allowing the agent to walk away from the agreement puts a limit on how much he can be punished after a sequence of bad outcomes (see, for example, Phelan, 1995 or Krueger and Uhlig, 2006). Another possibility is to allow contracting parties to threaten with perpetual autarky to support a long-term contractual relationship (e.g., Kocherlakota, 1996; Thomas and Worrall, 1988, 1994; or Ligon, Thomas and Worrall, 2002 among others).

In this paper, we revisit the role of limited commitment in a dynamic risk-sharing setting with private information. Specifically, we consider a risk-averse agent endowed with a technology that transforms effort into stochastic output and who can imperfectly self-insure through savings. A profit-maximizing insurer can observe the agent's asset holdings, but not his effort. We assume that agent and insurer cannot commit to an agreement beyond the current period and study the resulting implications for risk-sharing, asset dynamics and wealth inequality. We focus on Markov-perfect insurance contracts, i.e., one-period agreements which only depend on payoff-relevant variables: in our case, beginning-of-period asset holdings and the current output state.

We show that a Markov-perfect equilibrium and an infinitely-long contract where only the insurer can commit ("one-sided commitment") are equivalent in terms of the history-contingent sequences for consumption and effort they implement. However, their implications for asset dynamics are very different. An insurer endowed with commitment power can use utility promises and agent's savings interchangeably to implement desired risk-sharing allocations. This leads to an indeterminacy of asset dynamics. In contrast, we show that the inability of the insurer to commit to a long-term contract, implies that asset holdings by the agent become an integral part of insurance contracts, which results in non-trivial and determinate asset dynamics. In effect, one could interpret a Markov-perfect equilibrium as a specific, empirically usable implementation of a one-sided commitment contract.

We further show that Markov-perfect risk-sharing contracts provide partial insurance and are characterized by "inverse Euler equations", thus preserving standard properties of contracts with full commitment in moral hazard settings familiar from the existing literature. However, unlike in many of those settings, both theoretically and computationally, Markov-perfect equilibria are easy to define and characterize—the optimal contracting problem can be written recursively with a single scalar state variable, the agent's asset holdings. This simplicity holds also if agents can save privately on the side, as long as their asset holdings are observable to the insurer. Curse of dimensionality issues which have been shown to arise in the "hidden savings" literature (e.g., Fernandes and Phelan, 2000; Kocherlakota, 2004; Doepke and Townsend, 2006) are avoided.

Our focus on Markov-perfect contracts therefore allows us to explore the role of asset contractibility in a limited commitment environment in a tractable way. We find that whether

¹Green (1987), Spear and Srivastava (1987), Thomas and Worrall (1990), Phelan and Townsend (1991), Atkeson and Lucas (1992) among many others.

agent’s asset accumulation decisions can be contracted upon or not affects the contract terms as long as insurers make positive expected profits in equilibrium. This result arises because of a disagreement in the marginal value of future assets between the agent and the insurer. Specifically, the agent wants to save more than the amount preferred by the insurer since this enables him to raise his future outside option. The disagreement disappears when all the surplus of the contract goes to the agent, as in the case of perfectly competitive insurance markets.

As an application of the theory, we show that Markov-perfect equilibrium insurance contracts can be calibrated along some dimensions to resemble actual economies. The calibrated model is used to gauge and compare the welfare costs of frictions like lack of commitment and private information. We find that Markov-perfect contracts provide sizable insurance, particularly for agents with low asset holdings. This has consequences for wealth inequality, as wealth-poor agents have lower incentives to self-insure through asset accumulation when they have access to insurance markets. For our parameterization, this effect is quite significant: compared to autarky (self insurance), long-run wealth inequality, as measured by the Gini coefficient, increases by about 30% when Markov-perfect contracts are made available. We find that the welfare gains from resolving the commitment friction are larger than those from resolving the moral hazard problem at low asset levels, while the opposite holds for high asset levels. Finally, we calculate the gains associated with asset contractibility and show that they can be sizable. For example, if the insurer is a monopolist facing the median-wealth agent, net present value profits are almost 30% higher when assets are contractible.

Our paper relates and builds upon previous work on both repeated moral hazard and limited commitment. Unlike here, the early papers on optimal risk-sharing in repeated moral hazard settings by Townsend (1982) and Rogerson (1985b) assume full commitment by both sides and do not allow agents to save privately or any other role for assets. Hidden private saving is allowed by Allen (1985) and Cole and Kocherlakota (2001) in a stochastic income environment. They show that, under certain conditions, no additional consumption smoothing over self-insurance can be provided. In contrast, we obtain that additional smoothing over self-insurance is always possible, even with double-sided lack of commitment. The reason is that in our setting, unlike in the stochastic endowment case, the probability distribution over output is endogenous through agent’s effort and so the insurer’s transfers affect the level of uncertainty the agent faces (see also Abraham and Pavoni, 2008).

We differ from most classic papers on “limited commitment” by agents (e.g., Thomas and Worrall, 1988, 1994; Kocherlakota, 1996; Ligon et al., 2002; Krueger and Uhlig, 2006 among others) in terms of precisely when agents can renege on a contract and, more importantly, what this timing implies for the resulting model dynamics. In those papers, the main issue is the potential inability to support full insurance when agents can opportunistically renege on the contract (and go to autarky or another insurer) within the period, after a high output is realized. Asset accumulation is not studied since, unlike here, insurers are assumed to be able to commit to an infinitely-long contract and thus implement any desired incentive-feasible allocation through promised utility. In contrast, our main focus is on the private asset dynamics that arise when both agents and insurers can commit within a period, but cannot commit across periods.²

The rest of the paper is organized as follows. Section 2 presents the model environment. Section 3 defines and characterizes a Markov-perfect equilibrium (MPE), assuming perfectly competitive insurance markets. Section 4 shows that MPE implements the same consumption and effort sequences as an infinitely-long contract to which the insurer can commit subject to per-period participation constraints by the agent. Section 5 presents a numerical analysis of

²Our timing regarding when agents can leave the contract is similar to that in Phelan (1995) who, however, assumes that: insurers can fully commit (thus, no asset accumulation); agents’ income is unobservable; and agents need to sit out one period if they renege before they can sign-up with another insurer.

Markov-perfect insurance contracts and welfare analysis of the frictions of lack of commitment and moral hazard in a calibrated version of the model. Section 6 extends the results by allowing insurers to have market power and shows that whether asset are contractible or not affects equilibrium contracts generically. Section 7 concludes.

2 The Environment

Consider an infinitely-lived agent who maximizes expected discounted utility from consumption c and disutility from effort e . Period utility is given by $u(c) - e$, where $u_c(c) > 0$, $u_{cc}(c) < 0$ and u satisfies Inada conditions.³ The agent discounts future utility by factor $\beta \in (0, 1)$.

The agent is endowed with a technology that produces output as a function of effort. There are $n \geq 2$ possible values for output: $0 \leq y^1 < \dots < y^n$. Let $\pi^i(e)$ be the probability of output y^i being realized. Suppose e takes values on the set \mathbb{E} which is a closed interval in \mathbb{R}_+ .

Assumption 1 For all $e \in \mathbb{E}$: (i) $\sum_{i=1}^n \pi^i(e) = 1$; (ii) “full support”, $\pi^i(e) > 0$; (iii) $\pi^i(e)$ is twice continuously differentiable for all $i = 1, \dots, n$; and (iv) the “monotone likelihood ratio property” holds, $\frac{\pi_e^i(e)}{\pi^i(e)}$ is non-decreasing in i .

The agent can save or borrow at the gross rate, r . His asset holdings, a are constrained to the set $\mathbb{A} = [\underline{a}, \bar{a}]$.

Assumption 2 (i) $0 < r < \beta^{-1}$; (ii) $\underline{a} = -\frac{y^1}{r-1}$.

We further assume $\bar{a} \in (0, \infty)$ and sufficiently large so that it never binds. Assumption 2(i) is a standard condition ensuring the existence of a finite upper bound on asset holdings. Assumption 2(ii) sets the minimal asset position, \underline{a} equal to the natural borrowing limit (Aiyagari, 1994) which allows us to focus on interior solutions for asset choice.

3 Markov-Perfect Insurance Contracts

The agent has access to a perfectly competitive insurance market with free entry, populated by risk-neutral, profit-maximizing insurers. Demand for market insurance exists in our setting since the agent cannot span the n -dimensional output uncertainty through just borrowing and saving in the single non-contingent asset a . As discussed in the introduction, we assume that neither the insurers nor the agent can bind themselves to a contract extending beyond the current period. However, within-period contracts are perfectly enforceable. Insurance contracts are offered before effort is exerted and specify the exchange of realized output for state-contingent transfers. Agent’s effort is not observable by insurers. In contrast, the agent’s assets are always observable and their choice—which occurs after output is realized—is assumed to be contractible (relaxing the latter is discussed in Section 6.2).

We study dynamic insurance contracts, the terms of which depend only on fundamentals, i.e., payoff-relevant variables: beginning-of-period assets and current output realization. Due to our lack of commitment assumption insurance contracts cannot depend on variables capturing future promises, such as “promised utility”. Following Maskin and Tirole (2001), we call these contracts Markov-perfect.

³Throughout the paper we use subscripts to denote partial derivatives and primes for next-period values.

Agents starting the period with different asset levels will, in general, be offered different contracts. Insurance contracts, $\{\mathcal{T}^i(a), \mathcal{A}^i(a)\}_{i=1}^n$, consist of transfers and future asset holdings as functions of the agent's beginning-of-period assets, $a \in \mathbb{A}$ and the realized output state, $i = 1, \dots, n$. When entering an agreement with an insurer, the problem of the agent is to choose the level of effort, given the current offered contract and all anticipated future contracts. Call the associated state-contingent agent consumption $\mathcal{C}^i(a) \equiv ra + \mathcal{T}^i(a) - \mathcal{A}^i(a)$. Given anticipated future contracts, $\{\mathcal{T}^i(\cdot), \mathcal{A}^i(\cdot)\}_{i=1}^n$, which induce future value $\mathcal{V}(\mathcal{A}^i(\cdot))$ —to be defined below—the agent's problem is

$$\max_e \sum_{i=1}^n \pi^i(e) [u(\mathcal{C}^i(a)) + \beta \mathcal{V}(\mathcal{A}^i(a))] - e.$$

The first-order condition is

$$\sum_{i=1}^n \pi_e^i [u(\mathcal{C}^i(a)) + \beta \mathcal{V}(\mathcal{A}^i(a))] - 1 = 0. \quad (1)$$

Since the insurer cannot commit beyond the current period, he takes future insurance contracts as given. In other words, the insurer takes $\mathcal{V}(\mathcal{A}^i(a))$ induced by $\{\mathcal{T}^i(a), \mathcal{A}^i(a)\}_{i=1}^n$, as given. Note that the insurer can affect the agent's continuation value through the future-assets component of the contract, $a^i = \mathcal{A}^i(a)$.

Insurers also need to take into account how the agent's effort responds to the offered contract for the period (incentive-compatibility). We use the “first-order approach” (Rogerson, 1985a) and impose the first-order condition of the agent, (1) as a constraint in the insurer's problem. We assume the probability functions $\pi^i(e)$ are such that the first-order approach is valid. For example, if $\sum_{j=1}^i \pi_{ee}^j(e) \geq 0$ for all i on top of Assumption 1, then our probability functions satisfy the sufficient conditions for the validity of the first-order approach: the “monotone likelihood ratio property” (MLRP) and the “convexity of the distribution function condition” (CDFC)—see Rogerson (1985a) for a proof. Of course, many alternative sufficient conditions are possible. For example, in the two-output case ($n = 2$) it is sufficient to assume that $\pi^2(e)$ is strictly increasing and strictly concave in e .

Free entry in the insurance market results in zero expected profits in each period and for any sub-market indexed by agent's assets holdings, a . Cross-subsidization across agents at different asset levels is ruled out—if an insurer makes profits on some agents but makes a loss on others, another insurer could come in and offer a better contract to the former group. This result holds regardless of the rate at which insurers discount profits.

Perfect competition implies all the gains from insurance contracts go to the agent. Thus, the problem of an insurer facing agent with assets a can be written as

$$\max_{\{\tau^i, a^i\}_{i=1}^n, e} \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta \mathcal{V}(a^i)] - e \quad (2)$$

subject to incentive-compatibility and zero per-period profits,

$$\sum_{i=1}^n \pi_e^i(e) [u(c^i) + \beta \mathcal{V}(a^i)] - 1 = 0 \quad (3)$$

$$\sum_{i=1}^n \pi^i(e) [y^i - \tau^i] = 0, \quad (4)$$

and where we used $c^i \equiv ra + \tau^i - a^i$ to simplify notation.

We now formally define a Markov-perfect equilibrium and Markov-perfect contracts in our setting.

Definition 1 A *Markov-perfect equilibrium* (MPE) is a set of functions $\{\{\mathcal{T}^i, \mathcal{A}^i\}_{i=1}^n, \mathcal{E}, \mathcal{V}\} : \mathbb{A} \rightarrow \mathbb{R}^n \times \mathbb{A}^n \times \mathbb{E} \times \mathbb{R}$ such that for all $a \in \mathbb{A}$:

$$\{\{\mathcal{T}^i(a), \mathcal{A}^i(a)\}_{i=1}^n, \mathcal{E}(a)\} = \operatorname{argmax}_{\{\tau^i, a^i\}_{i=1}^n, e} \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta \mathcal{V}(a^i)] - e$$

subject to (3) and (4) and where

$$\mathcal{V}(a) = \sum_{i=1}^n \pi^i(\mathcal{E}(a)) [u(\mathcal{C}^i(a)) + \beta \mathcal{V}(\mathcal{A}^i(a))] - \mathcal{E}(a),$$

where $\mathcal{C}^i(a) = ra + \mathcal{T}^i(a) - \mathcal{A}^i(a)$ and $c^i = ra + \tau^i - a^i$.

A **Markov-perfect contract** for any given asset level $a \in \mathbb{A}$ is the state-contingent transfer and asset choices $\{\tau^i = \mathcal{T}^i(a), a^i = \mathcal{A}^i(a)\}_{i=1}^n$ associated with a MPE.

We next characterize Markov-perfect equilibria in our setting. The constraint set is non-empty for all $a \in \mathbb{A}$ —for example, the full-insurance contract with $e = \min\{\mathbb{E}\}$ or setting $\{e, \tau^i, a^i\}$ to their autarky levels (see Appendix A) are always feasible. Since \mathbb{A} is compact, existence of a fixed point, \mathcal{V} in problem (2)—(4) can be then shown as in Abraham and Pavoni (2006), using standard contraction mapping arguments.

In what follows, to simplify notation, we use $E[x] = \sum_{i=1}^n \pi^i(e) x^i$ for any variable x .

Proposition 1 A *Markov-perfect equilibrium* is characterized by:

- (i) $\mathcal{C}^i(a)$ non-decreasing in i , with $\mathcal{C}^1(a) < \mathcal{C}^n(a)$, for all $a \in \mathbb{A}$;
- (ii) the inverse Euler equations

$$u_c(\mathcal{C}^i(a)) = \frac{\beta r}{E\left[\frac{1}{u_c(\mathcal{A}^i(a))}\right]},$$

for all $a \in \mathbb{A}$ and $i = 1, \dots, n$.

Proof. With Lagrange multipliers λ and μ on (3) and (4), respectively, the first-order conditions with respect to transfers and assets are

$$\begin{aligned} u_c(c^i) [\pi^i(e) + \lambda \pi_e^i(e)] - \mu \pi^i(e) &= 0 \\ [-u_c(c^i) + \beta \mathcal{V}_a(a^i)] [\pi^i(e) + \lambda \pi_e^i(e)] &= 0, \end{aligned} \tag{5}$$

for all $i = 1, \dots, n$. Re-arrange (5) as,

$$\frac{\mu \pi^i(e)}{u_c(c^i)} = \pi^i(e) + \lambda \pi_e^i(e).$$

Since this expression holds for any i , sum over $i = 1, \dots, n$ and obtain

$$\mu \sum_{i=1}^n \frac{\pi^i(e)}{u_c(c^i)} = \sum_{i=1}^n \pi^i(e) + \lambda \sum_{i=1}^n \pi_e^i(e).$$

Given $\sum_{i=1}^n \pi^i(e) = 1$ for all e , we have $\sum_{i=1}^n \pi_e^i(e) = 0$, which implies

$$\mu = \frac{1}{E\left[\frac{1}{u_c(c)}\right]} > 0. \tag{7}$$

Given Assumption 1 (the monotone likelihood ratio property) and $\mu > 0$, part (i) follows from the fact that the incentive-compatibility constraint (3) binds at optimum (i.e., $\lambda > 0$)—see Rogerson (1985a) or Bolton and Dewatripont (2004) for discussion and proofs. The monotonicity of period consumption in output is a static property—see the first-order condition (5)—so the standard proofs carry over.

To show part (ii), note that given $u_c(c^i) > 0$ and $\mu > 0$, (5) implies $\pi^i(e) + \lambda\pi_e^i(e) > 0$ for all $i = 1, \dots, n$. Thus, from (6) we get

$$-u_c(c^i) + \beta\mathcal{V}_a(a^i) = 0, \quad \forall i \quad (8)$$

The envelope condition implies

$$\mathcal{V}_a(a) = r \sum_{i=1}^n u_c(c^i)[\pi^i(e) + \lambda\pi_e^i], \quad (9)$$

which, using (5), implies $\mathcal{V}_a = r\mu$. Replace μ from (7) and update one period. Plugging the resulting expression into (8), we obtain the inverse Euler equations as stated in the proposition. ■

Proposition 1 shows that Markov-perfect contracts preserve standard features of insurance under asymmetric information. Part (i) states that these contracts do not provide full insurance, as is generally the case for insurance arrangements in the presence of private information. Part (ii) shows that consumption paths are characterized by the familiar inverse Euler equations from the full commitment literature (e.g., Rogerson, 1985b; Golosov et al., 2006).

4 Equivalence of Markov-perfect and one-sided commitment contracts

In this section, we investigate further the role of commitment by the insurer. Specifically, we show that in our setting the insurer's inability to commit to a long-term contract is immaterial, in the sense that endowing him with commitment would not alter the consumption and effort allocation the agent receives when signing a contract, nor social welfare. However, we argue that the Markov-perfect equilibrium analyzed in the previous section provides a justification for specific asset dynamics, which are indeterminate in the setting with full commitment by the insurer.

Consider a “one-sided commitment” contract, such that insurers are able to commit to an infinite sequence of state-contingent transfers, but agents are allowed to walk away at the beginning of each period, *before* output is realized. Although the agent is assumed to be unable to commit to stay in the contract beyond the current period, we assume that he cannot renege on the terms of the contract for the period if he decides to stay on. Thus, the agent's participation constraint must be satisfied in every period. This timing of events is similar to Phelan (1995), except that we allow the agent to sign up with another insurer right away, instead of sitting-out for one period. Crucially, our timing is different from that in much of the “limited commitment” literature (Thomas and Worrall, 1988, 1994; Ligon, Thomas and Worrall, 2002; Kocherlakota, 1996; Krueger and Uhlig, 2006 among many others) where agents can walk away *after* observing the output realization. The lack of commitment friction that we stress is thus not about short-term opportunistic behavior of agents when output is high, but about their ability to walk away at the beginning of any time period, should the terms of the contract leave them worse-off than their best alternative.

Let $s_t \in \{1, \dots, n\}$ be the output state in period t and $s^t \equiv \{s_0, \dots, s_t\}$ denote the history of output states from period 0 up to period t . Given initial asset holdings, $a_0 \in \mathbb{A}$, a one-

sided commitment contract consists of history-dependent sequences for consumption, assets and effort. Let $\{c(a_0, s^t), a(a_0, s^t), e(a_0, s^t)\}_{t=0}^{\infty}$ denote those sequences. For any $t = 0, \dots, \infty$ and any history s^t , let $\alpha(s^t)$ denote the beginning-of-period asset holdings by the agent, obtained using the policy rule in a MPE, $\mathcal{A}^i(a)$, i.e., $\alpha(s^t) = \mathcal{A}^{s^t}(\alpha(s^{t-1}))$ with $\alpha(s^{-1}) = a_0$.

Proposition 2 *Given $\alpha(s^{-1}) = a_0 \in \mathbb{A}$ and any history s^∞ , a one-sided commitment contract implements history-contingent consumption and effort sequences $\{c(a_0, s^t), e(a_0, s^t)\}_{t=0}^{\infty}$ identical to the sequences $\{C^{s^t}(\alpha(s^{t-1}), \mathcal{E}(\alpha(s^{t-1})))\}_{t=0}^{\infty}$ implemented in a Markov-perfect equilibrium solving problem (2)–(4).*

Proof. See Appendix B. ■

Whether the insurer can or cannot commit in our setting does not affect the insurance received by the agent—our Markov-perfect contracts in which insurers can only commit for the current period implement the same outcome as when insurers can commit to infinitely-long contracts. As shown in Karaivanov and Martin (2011), the key to this equivalence result is that agent and insurer face the same rate of return r . Instead, if the insurer had superior return on assets, then his ability to commit allows him to support higher promised utility in the future and offer a contract with a higher net present value than achieved in a Markov-perfect equilibrium. On the other hand, the commitment ability of the agent matters regardless of the parties’ returns on assets. If both sides could commit to an infinitely-long contract, promises lower than any outside option can be used at some point of time, after some histories, which results in higher ex-ante welfare—see Karaivanov and Martin (2011) for more discussion.

Proposition 2 shows that consumption and effort allocations in a Markov-perfect equilibrium are equivalent to those in one-sided commitment contracts, which have been studied often in the literature. Markov-perfect insurance contracts are thus not as restrictive as they may initially appear. There is, however, an important difference between MPE and one-sided commitment contracts. One-sided commitment contracts feature an indeterminate path for assets since assets and promised-utility are interchangeable in terms of implementing future allocations (see the proof of Proposition 2 for details). Infinitely many asset paths are thus consistent with the same insurance contract. In contrast, in a Markov-perfect equilibrium insurers cannot use promised utility to implement dynamic insurance arrangements since they lack the ability to commit to any promises involving events beyond the present period. In this case, the agent’s asset holdings are the only instrument insurers can use to manipulate the contract’s future value. Importantly, this leads to non-trivial, determinate asset dynamics which we can use to evaluate the model empirically.

5 Numerical Analysis

In this section we calibrate a version of our model and solve for the resulting Markov-perfect equilibrium. In addition to demonstrating the relative numerical simplicity of the solution, we show that Markov-perfect contracts possess characteristics that can be calibrated to match certain features of US macro data. We also use our calibration to perform a simple welfare analysis comparing the gains from resolving private information vs. commitment problems in our setting. All references to a Markov-perfect equilibrium in this section correspond to the *computed* MPE.

5.1 Parametrization and computation

Assume a generalized CRRA parameterization for the utility function,

$$u(c) = \frac{\alpha(c^{1-\sigma} - 1)}{1 - \sigma}.$$

where $\alpha > 0$ and $\sigma > 0$ with $u(c) = \ln c$ for $\sigma = 1$. We consider an economy with $n = 3$, “low”, “medium” and “high” output levels, labeled y^L , y^M and y^H , respectively. The probability functions, $\pi^i(e)$ are given by

$$\begin{aligned}\pi^L(e) &= 1 - \pi^M(e) - \pi^H(e) \\ \pi^M(e) &= \frac{\varphi e^\nu}{1 + e^\nu} \\ \pi^H(e) &= \frac{(1 - \varphi) e^\nu}{\gamma + e^\nu},\end{aligned}$$

where $\gamma, \nu > 0$ and $\varphi \in (0, 1)$. It can be directly verified that the chosen probability functions satisfy the sufficient conditions for the validity of the first-order approach (Rogerson, 1985a). If $\gamma \neq 1$ the likelihood ratios $\frac{\pi_e^i(e)}{\pi^i(e)}$ are different for medium and high output and so transfers differ for all levels of output.

In picking parameter values, we follow Castañeda, Díaz-Giménez, and Ríos-Rull (2003) who match earnings, the wealth distribution and other aggregates for the US economy in a model with a mix of dynastic and life cycle features. Table 1 displays our selection of parameters. The values for β and σ are taken straight from Castañeda et al. We also set r equal to the value which their calibration implies for the annual interest rate net of depreciation. For the three output levels, we take their values for the endowments of labor efficiency units.⁴ We set the value for ν arbitrarily to 0.5 and pick φ and γ so that the long-run distribution of realized output levels in a Markov-perfect equilibrium approximates the proportions for the US reported in Castañeda et al. (2003)—see Table 2 below. The remaining parameter, α affects only the scale of effort levels in equilibrium and is set large enough for effort to be significantly different (in a numerical sense) than zero for high asset levels. Finally, following Castañeda et al. (2003) and most of the related literature, we assume agents are subject to a non-borrowing constraint, i.e., we set $\underline{a} = 0$. For our parameterization, this constraint only binds for the lowest output state.

Table 1: Parameter values

α	β	σ	r	ν	φ	γ	y^L	y^M	y^H
4.000	0.924	1.500	1.061	0.500	0.450	2.000	0.10	0.315	0.978

We start by computing the autarky problem (see Appendix A for its formulation and solution properties). We use a discrete grid of 1,000 points for the state space but allow all choice variables to take any admissible value.⁵ Cubic splines are used to interpolate between grid points.

⁴See Table 5 in Castañeda et al. (2003). Note that they parameterize four endowment levels; the fourth type is about 1,000 times more productive than the first type and comprises 0.04% of working-age households. Since our economy already makes important simplifications with respect to theirs—no life-cycle features, taxes, etc.—we omit this fourth type to simplify the numerical analysis and exposition of results.

⁵We did not find significant gains from further increasing the size of the grid. For example, the value of autarky at zero assets computed with 1,000 and with 10,000 grid points differs by only 0.02%.

The upper bound for assets \bar{a} is set to 60 which ensures that all three asset accumulation functions, $\mathcal{A}^i(a)$ cross the 45-degree line. We use the same asset grid for all computations performed below. For clarity of exposition, graphs only display asset holdings up to $a = 5$, which includes 99.95% of agents in a stationary equilibrium in autarky and virtually all agents in a MPE.

To compute a Markov-perfect equilibrium, we use the following iterative algorithm to find fixed-point in the value function $\mathcal{V}(a)$ and its corresponding policy functions: (i) start with the agent's value in autarky as initial guess for \mathcal{V} ; (ii) solve the insurer's maximization problem (2)—(4) which outputs a new value function; (iii) update and continue iterating until convergence. Subsequently, we use the first-order conditions to the insurer's problem to improve the precision of the solution.

In addition, we compute the long-run distribution of assets by assuming continuum of agents. This is done using standard techniques: the decisions rules derived from the numerical solution imply a transition matrix on which we iterate until obtaining a distribution that maps into itself.

5.2 Autarky vs. Markov-perfect equilibrium

Figure 1 displays transfers from the insurer to the agent in the computed MPE, as functions of the agent's current assets a . Clearly, the principal is able to provide insurance over and above the self-insurance allocation achieved by the agent in autarky. Specifically, for our calibration we obtain $y^L < \mathcal{T}^L(a) < y^M < \mathcal{T}^M(a) < \mathcal{T}^H(a) < y^H$ for all $a \in \mathbb{A}$. Thus, if realized output is low, the principal provides the agent with higher consumption than in autarky while, if realized output is medium or high, he provides the agent with lower consumption than in autarky. Although not shown on the graph, the transfer functions for the low and medium states become almost flat for high asset levels, whereas the transfer function for the high state converges towards that for the medium state.


We next compare the state-contingent consumption levels in autarky and the computed Markov-perfect equilibrium. For all asset levels, consumption for the low and medium output states is higher in a MPE, whereas consumption in the high output state is always higher in autarky. The differences for the low and medium states are significant at low asset levels. For example, at zero assets, computed MPE levels of c^L and c^M are 99% and 38% higher than those in autarky, respectively.

Figure 2 provides a measure of the degree of consumption insurance achieved in a MPE as opposed to in autarky as function of assets a . Specifically, we display c^M/c^L and c^H/c^L in autarky and in the computed Markov-perfect equilibrium. We see that for high asset holdings, the agent is self-insuring adequately and thus, both ratios approach one in either environment. For low asset levels, however, there is significant demand for market insurance. For example, at zero assets, the agent contracting with a competitive insurer obtains about 1.5 times of his consumption in the low state if output is medium or high. In autarky, the same agent would instead consume about 2.2 and 3.4 times his consumption in the low state, respectively, indicating much lower ability to smooth consumption.

Naturally, given the additional insurance provided in a Markov-perfect equilibrium over autarky, agent's effort decreases when contracting with an insurer (the graph is omitted to save space). At zero assets, effort in autarky is about 3 times higher than in the computed Markov-perfect equilibrium. This difference in effort levels becomes smaller at higher asset levels: for instance, at $a = 5$ the agent only exerts about 15% more effort in autarky.

Despite the above differences, the long-run measures of agents with specific output realiza-

Figure 1: Transfers in a MPE



tions do not differ dramatically between autarky and Markov-perfect equilibrium—see Table 2. However, self-insurance and MPE insurance yield significantly different long-run distribution of assets (wealth). The Gini coefficient of the wealth distribution is 0.35 in autarky and 0.45 in the Markov-perfect equilibrium. As a reference, the canonical model by Aiyagari (1994) which is basically our autarky model from Appendix A, but with exogenous output probabilities, features a Gini coefficient of 0.38, whereas the model by Castañeda et al. (2003) matches the US Gini coefficient of 0.78.⁶ Figure 3 displays the Lorenz curves for autarky and the computed Markov-perfect equilibrium. Most of the difference between the Lorenz curves in the two settings is explained by the fact that wealth-poor agents contracting with an insurer have lower incentives to self-insure through asset accumulation.


Table 2: Long-run measure of agents, according to output realizations

	y^L	y^M	y^H
Autarky	0.585	0.228	0.187
MPE	0.605	0.218	0.176

The parameterization adopted here allows our autarky economy with endogenous effort to

⁶This number is calculated using the 1992 Survey of Consumer Finances. Budría-Rodríguez et al. (2002) update it to 0.803 using the 1998 Survey.

Figure 2: Consumption insurance in Autarky and MPE



Note: dashed line corresponds to autarky, solid line corresponds to MPE.

perform similarly to Aiyagari’s model with exogenous earnings. Taking this economy as a benchmark, our MPE calibration results indicate that augmenting the Aiyagari self-insurance framework to an insurance market environment with lack of commitment and moral hazard frictions, as assumed here, goes in the right direction, in the sense that we show it can help explain a non-trivial additional part of long-run wealth inequality.

We take one final look at the differences between autarky and Markov-perfect equilibrium by focusing on the median-wealth agent in each environment. Table 3 shows his assets, effort, expected income (post-transfers and including asset income) and assets over expected income in both environments. The median-wealth agent is significantly poorer in the computed Markov-perfect equilibrium. In autarky, he holds assets three times as large as his period income; while in a MPE his asset holdings are only about 80% of his period income. Coupled with lower effort level, the median-wealth agent is also income-poorer in a MPE, although the income differential with autarky is much smaller.

Table 3: Statistics for median-wealth agent


	a	e	Ey	a/Ey
Autarky	1.119	1.031	0.380	2.948
MPE	0.247	0.942	0.321	0.772

Note: expected income, $Ey = \sum_{i=1}^n \pi^i(e)y^i + (r-1)a$ for autarky and $Ey = \sum_{i=1}^n \pi^i(e)\tau^i + (r-1)a$ for MPE.

5.3 Welfare analysis

The dynamic insurance contracts we study above feature two important frictions: private information and limited commitment. In this section, we compute the welfare costs associated

Figure 3: Lorenz curves



with these frictions for our calibrated economy. We start by computing the welfare gains for an agent from moving from autarky to the computed MPE insurance contract. Specifically, for all $a \in \mathbb{A}$ we solve for the consumption equivalent compensation function $\Delta(a)$ defined in


$$\sum_{i=1}^n \pi^i(\hat{\mathcal{E}}(a)) \left\{ u(\hat{\mathcal{C}}^i(a)[1 + \Delta(a)]) + \beta \Omega(\hat{\mathcal{A}}^i(a)) \right\} - \hat{\mathcal{E}}(a) = \mathcal{V}(a), \quad (10)$$

where $\{\hat{\mathcal{C}}^i(a), \hat{\mathcal{A}}^i(a), \hat{\mathcal{E}}(a)\}_{i=1}^n$ are the optimal policy functions in autarky, $\Omega(a)$ is the value function in autarky (see Appendix A) and $\mathcal{V}(a)$ is the agent's value in a Markov-perfect equilibrium (see Definition 1).

Figure 4 displays $\Delta(a)$, i.e., the welfare gains associated with moving from autarky to Markov-perfect insurance as function of agent's current assets, a . Clearly, these gains can be quite sizable. For example, at zero assets, the welfare gains are equivalent to a one-time payment of 326% of period consumption in autarky. Naturally, the gains are decreasing in assets as the demand for market insurance decreases with wealth. Still, the welfare gains are significant even for asset-rich agents; at $a = 5$ (where an agent in autarky is richer than 99.95% of the population), the value of Δ is about 10%.

The relatively large welfare gains we obtain at zero assets are in part due to the non-borrowing constraint, which severely limits the ability of the agent to self-insure when very poor. To see this, consider decreasing the lower bound on asset holdings to $\underline{a} = -1$ (as a reference, in our calibration the “natural borrowing limit”, Aiyagari, 1994 is about -1.64).

Figure 4: Welfare gains from moving from autarky to MPE




The welfare gains at zero assets drop to 127% of consumption in autarky. This number is significantly lower than with a non-borrowing constraint, but still very sizeable (the equivalent compensation at the new asset lower bound $\underline{a} = -1$ climbs to 4,244%)

Next, starting from our computed MPE, we consider the welfare gains that result from resolving the information and commitment frictions. A Markov-perfect equilibrium with full information, i.e., when effort is observable and contractible and so the moral hazard problem is resolved, can be defined analogously to Definition 1 but without the incentive compatibility constraint (see also Karaivanov and Martin, 2011 for more discussion on the full information case). To study the gains from resolving the commitment friction we solve the contracting problem assuming full (double-sided) commitment, i.e., assuming both sides can bind themselves to a lifetime agreement at time zero. This full-commitment problem is formulated exactly as the one-sided commitment problem from Section 4 and Appendix B, but without the lower bound on promised utility.

We compute the equivalent compensation function $\Delta(a)$ for the two scenarios outlined above. This involves replacing $\mathcal{V}(a)$ in (10) with the corresponding value function from either the full-information (with lack of commitment) or full-commitment (with moral hazard) solutions. Figure 5 displays the welfare gains from resolving the information and commitment frictions. We measure the relative magnitudes of these gains by taking the ratio of the corresponding consumption equivalent compensations over all asset levels.

The left panel in Figure 5 displays the gains from resolving the information friction in a

Figure 5: Welfare gains from resolving information and commitment frictions



Note: the dashed line indicates the ratio of consumption equivalent compensations is equal to 1.

Markov-perfect equilibrium. As we see, the welfare gains can be large: at zero assets, going from autarky to a Markov-perfect equilibrium with full information is 19 times better than going to a Markov-perfect equilibrium with private information. Although not visibly apparent in the chart, the value from resolving the information friction is non-monotonic—first decreasing and then increasing in assets (a minimum is achieved at about $a = 4.9$). The reason for this non-monotonicity is that there are two opposing forces at play. On the one hand, demand for insurance is decreasing in assets. With full information we obtain full rather than partial insurance, which is especially valued at low asset levels. On the other hand, as agent’s assets increase, it becomes more costly for the insurer to induce the desired amount of effort when it is private information. Thus, the gains from going to a Markov-perfect equilibrium fall much more rapidly with private information than with full information. For our parameterization, the latter force only dominates for asset levels that are held by virtually no agent in the long-run, but we have computed alternative examples where it dominates for all asset levels.

Endowing the insurer and agent with commitment power results in significant welfare gains for the agent. Recall from Section 4 that simply endowing the insurer with commitment does not affect the resulting contract. Therefore, what is key for the gains from commitment is the inability of the agent to walk away from the agreement, which allows the insurer to front-load agent’s consumption and extract agent’s assets providing better incentives to supply effort. The gains from full commitment at low asset levels can be huge: in our parametrization going from autarky to a setting with full commitment at $a = 0$ is 69 times better in terms of consumption equivalence compensation than going to a Markov-perfect equilibrium. The gains erode rapidly as asset holdings increase (not displayed in the Figure).

The right panel in Figure 5 compares head-to-head the welfare gains from resolving the information and commitment frictions. For our parameterization, the gains from commitment are higher than those from resolving the moral hazard problem at low asset levels, while the opposite holds for high asset levels. This result follows from the analysis above. The gains from agent’s commitment are associated with the insurer’s ability to immiserate the agent in the

long-run and, as such, depend crucially on the initial demand for insurance, i.e., the initial asset level. For high asset levels, there is simply too little scope for market insurance in the presence of the moral hazard problem. On the other hand, the welfare improvement from resolving the private information problem comes from the insurer’s ability to provide full insurance, which is still valued at high asset levels.

6 The role of asset contractibility

In the preceding sections, we assumed that the agent’s assets decisions can be contracted upon. Here, we show that our results for the case of a perfectly competitive insurance market with free entry are unaffected by this assumption—even if insurers cannot control (but can still observe) agent’s savings, the same MPE results. However, we also show that, more generally, when insurers have at least some market power asset (non-)contractibility does affect Markov-perfect insurance contracts.

6.1 Generalized insurance problem

Consider an agent with a general outside option as a function of assets, $B(a)$. To capture the full spectrum from a monopolistic insurer to partial market power to perfect competition, assume that the agent’s outside option is at least as high as his value in autarky—denoted $\Omega(a)$ (see Appendix A)—and up to the value obtained in perfectly competitive Markov-perfect equilibrium, $\mathcal{V}(a)$, as analyzed in Section 3. We make the following assumptions about the agent’s outside option.

Assumption 3 *Let $B(a)$ be differentiable, strictly increasing, strictly concave, and such that for all $a \in \mathbb{A}$: $B(a) \in [\Omega(a), \mathcal{V}(a)]$.*

Insurers discount future profits at rate r . This assumption can be interpreted as the insurer being able to carry resources intertemporally using the same technology as the agent.⁷ The problem of the insurer is

$$\Pi(a) = \max_{\{\tau^i, a^i\}_{i=1}^n, e} \sum_{i=1}^n \pi^i(e) \left[y^i - \tau^i + \frac{\Pi(a^i)}{r} \right] \quad (11)$$

subject to

$$\sum_{i=1}^n \pi_e^i(e) [u(c^i) + \beta B(a^i)] - 1 = 0 \quad (12)$$

$$\sum_{i=1}^n \pi^i(e) [u(c^i) + \beta B(a^i)] - e - B(a) = 0. \quad (13)$$

Since profits strictly decrease in transfers, the profit-maximizing insurer will always drive the continuation value of the agent to the agent’s outside option. Thus, in equilibrium, the agent’s continuation value is $B(a)$. In the special case $B(a) = \mathcal{V}(a)$ for all $a \in \mathbb{A}$ Lemma 2 in Appendix B proves that the above problem is equivalent to the perfectly competitive MPE problem (2)—(4).

⁷In Karaivanov and Martin (2011), for the case with full information, we allow the insurer’s discount rate to be anything in between r and β^{-1} and analyze how differences in technology interact with lack of commitment.

Proposition 3 For any given $B(a)$ satisfying Assumption 3, a Markov-perfect equilibrium is characterized by:

- (i) $\mathcal{C}^i(a)$ non-decreasing in i , with $\mathcal{C}^1(a) < \mathcal{C}^n(a)$, for all $a \in \mathbb{A}$;
- (ii) the inverse Euler equations,

$$u_c(\mathcal{C}^i(a)) = \frac{\beta r}{E\left[\frac{1}{u_c(\mathcal{C}(\mathcal{A}^i(a)))}\right]},$$

for all $a \in \mathbb{A}$ and for all $i = 1, \dots, n$.

Proof. With Lagrange multipliers λ and μ on the incentive and participation constraints, respectively, the first-order conditions with respect to transfers and assets are

$$-\pi^i(e) + u_c(c^i) [\lambda \pi_e^i(e) + \mu \pi^i(e)] = 0 \quad (14)$$

$$\frac{\pi^i(e) \Pi_a(a^i)}{r} + [-u_c(c^i) + \beta B_a(a^i)] [\lambda \pi_e^i(e) + \mu \pi^i(e)] = 0, \quad (15)$$

for all $i = 1, \dots, n$. Adding up (14), we obtain,

$$\mu = E\left[\frac{1}{u_c(c)}\right] > 0. \quad (16)$$

Part (i) then follows from standard arguments as in the proof of Proposition 1.

To show part (ii), note that for all $i = 1, \dots, n$, (14) implies $\lambda \pi_e^i(e) + \mu \pi^i(e) = \frac{\pi^i(e)}{u_c(c^i)} > 0$ given $u_c(c^i) > 0$ and $\pi^i(e) \in (0, 1)$. Thus, from (15) we get

$$\Pi_a(a^i) - r + \frac{\beta r B_a(a^i)}{u_c(c^i)} = 0, \quad (17)$$

for all $i = 1, \dots, n$.

The envelope condition implies

$$\Pi_a(a) = \lambda r \sum_{i=1}^n \pi_e^i(e) u_c(c^i) + \mu r \sum_{i=1}^n \pi^i(e) u_c(c^i) - \mu B_a(a).$$

Solving for λ from (14) we can rearrange the above expression as follows

$$\Pi_a(a) = r \sum_{i=1}^n \pi^i(e) \left\{ 1 - u_c(c^i) E\left[\frac{1}{u_c(c)}\right] \right\} + \mu r \sum_{i=1}^n \pi^i(e) u_c(c^i) - \mu B_a(a).$$

which, using (16), simplifies to

$$\Pi_a(a) = r - B_a(a) E\left[\frac{1}{u_c(c)}\right]. \quad (18)$$

Update one period and plug into (17) to obtain the inverse Euler equations from the proposition statement. ■

As in Section 2, Markov-perfect equilibria preserve standard properties of dynamic insurance under private information. The “partial insurance” and “inverse Euler equations” properties hold regardless of the agent’s outside option or, equivalently, the insurer’s market power.

6.2 Asset contractibility

So far we have assumed that insurers, whether competitive or possessing some market power can perfectly control agent's assets. Proposition 3 shows that market power does not affect the basic properties of Markov-perfect contracts. So does market power matter generally speaking? As we show next, the answer is affirmative in the case when agent's assets remain observable but cannot be directly controlled by the insurer. Realistically, one can imagine many situations falling in this category—insurers or governments can know agents' asset positions (via tax records, required disclosure, etc.) but may not be able to directly command what agents do with their assets. We show that asset non-contractibility can alter the terms of Markov-perfect insurance contracts when insurers have market power (even a small deviation from free entry is sufficient).

Suppose assets are non-contractible, that is, the agent cannot be contractually bound to a specific asset accumulation choice. In this case, Markov-perfect contracts (defined analogously to Definition 1) consist only of output-contingent transfers, $\{\mathcal{T}^i(a)\}_{i=1}^n$. The problem of the agent is

$$\max_{\{a^i\}_{i=1}^n, e} \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta B(a^i)] - e, \quad (19)$$

where $c^i \equiv ra + \mathcal{T}^i(a) - a^i$. Unlike in the contractible case, the agent chooses both effort, e and future asset holdings, a^i .

Proposition 4 *For any given outside option $B(a)$ satisfying Assumption 3, a Markov-perfect equilibrium does not depend on asset contractibility if $\Pi_a(a) = 0$ for all $a \in \mathbb{A}$ and only if $\Pi_a(\mathcal{A}^i(a)) = 0$ for all $a \in \mathbb{A}$ and all $i = 1, \dots, n$.*

Proof. If assets are non-contractible then, for any given contract, $\{\mathcal{T}^i(a)\}_{i=1}^n$ the agent's savings decision is characterized by the following optimality condition:

$$-u_c(c^i) + \beta B_a(a^i) = 0, \quad (20)$$

for all $i = 1, \dots, n$. At optimum condition (20) must be satisfied by the insurer as an additional constraint in the contracting problem (11). Note that de facto we are using a 'first-order approach' in the actions a^i and e replacing the agent's optimization problem (19) with its first-order conditions. In practice, the validity of this approach can be verified using methods suggested by Abraham and Pavoni (2006) which is what we report on in the following numerical section.

If $\Pi_a(a) = 0$ for all $a \in \mathbb{A}$ then (17) implies (20) and thus, MPE with and without contractible assets are identical. Conversely, suppose Markov-perfect equilibria with and without contractible assets are identical. Together, (17) and (20) then imply $\Pi_a(a^i) = 0$ for all $i = 1, \dots, n$. ■

In general, expected profits depend non-trivially on assets, since the latter affect the agent's demand for insurance. Intuitively, the agent wants to save more than the amount preferred by the insurer since this enables the agent to raise his outside option, $B(a^i)$ and secure higher future utility. This misalignment of incentives disappears in the perfect competition case where all the surplus goes to the agent. In that case expected profits $\Pi(a)$ are identically zero for all $a \in \mathbb{A}$ and so $\Pi_a(a) = 0$ for all $a \in \mathbb{A}$. Therefore, Markov-perfect contracts do not depend on whether agent's assets are contractible or not.

6.3 Welfare

We now evaluate numerically how Markov-perfect insurance contracts are affected by asset (non-)contractibility as a function of the insurer’s market power for our parameterization from Section 4. Define the outside option of the agent as follows,

$$B(a) = (1 - \theta)\Omega(a) + \theta\mathcal{V}(a),$$

where $\theta \in [0, 1]$, $\Omega(a)$ is the agent’s value in autarky, as defined in Appendix A and $\mathcal{V}(a)$ is the value obtained by the agent in a Markov-perfect equilibrium with perfect competition, as defined and characterized in Section 3. Note that, from Proposition 4, the function $\mathcal{V}(\cdot)$ is the same for the cases with contractible and non-contractible assets. The value of θ can be interpreted to reflect the insurer’s market power. When $\theta = 0$, the insurer is a monopolist who maximizes profits subject to the agent receiving at least his autarky value; when $\theta = 1$ we obtain the environment with perfect competition among insurers.

The numerical procedure we use when $\theta < 1$ is similar to what we did in Section 5. The main difference is that here, we search for a fixed-point in profits, $\Pi(a)$ for any given θ . Specifically, we compute Markov-perfect equilibrium for θ on $[0, 1]$, in 0.05 increments. For all θ values less than one we solve both the problem with contractible and non-contractible assets.

As shown above, asset non-contractibility affects the terms of Markov-perfect insurance contracts except in the perfectly competitive case. To quantify this effect, Figure 6 measures the gains (in terms of profits for the insurer) from being able to contract on asset accumulation, as a function of θ for select asset levels—note that, for a given θ , $B(a)$ is the same in the contractible and non-contractible cases. The gains from asset contractibility can be quite large: e.g., for the median-wealth agent in a MPE from Table 3, a monopolist insurer ($\theta = 0$) can make almost 30% more profits if assets are contractible. Note that even though the profit ratios displayed in Figure 6 are not monotone in assets, the differences in profits between the contractible and non-contractible cases are.


Remember, our analysis relies on using the agent’s first-order condition with respect to assets as a constraint in the insurer’s problem. To confirm the validity of this approach we use the verification technique proposed by Abraham and Pavoni (2008). That is, after computing Markov-perfect equilibrium, we solve the agent’s maximization problem (19) globally (without using first-order conditions) and verify that the difference between the resulting net present value for the agent and $B(a)$ is below some desired tolerance. Table 4 reports maximum and median errors for selected values of θ . In all cases, the validity of the approach is verified.

Table 4: Verification of first-order approach when assets are non-contractible

θ	Maximum error	Median error
0.00	1×10^{-6}	2×10^{-11}
0.25	2×10^{-7}	2×10^{-11}
0.50	2×10^{-7}	2×10^{-11}
0.75	2×10^{-7}	2×10^{-11}
1.00	3×10^{-6}	9×10^{-14}

Note: “error” for any asset level a is defined as $[v(a) - B(a)]/B(a)$, where $v(a)$ is the value that results from solving the agent’s maximization problem (19) globally (i.e., without using first-order conditions). See Abraham and Pavoni (2008).

Figure 6: The value of asset contractibility



Note: Each line displays the ratio between profits in a MPE with contractible assets and profits in a MPE with non-contractible assets, as a function of market power, θ , at different asset levels. The median asset holdings corresponds to the case with contractible assets.

7 Concluding remarks

The Markov-perfect dynamic insurance contracts studied in this paper offer several characteristics that make them attractive for empirical analysis. Specifically, we have: simple recursive representation with single, scalar state without the curse of dimensionality even if agents can (observably) save on the side; determinacy in asset dynamics; non-degenerate long-run wealth distribution; and sizable (hence, potentially testable) effects when varying key environment assumptions such as market power and asset contractibility. We plan on pursuing research along these lines using structural estimation techniques, for example, by testing consumption and assets implications of our Markov-perfect insurance contracts against those of alternative models from the literature (permanent income, limited commitment, etc.)

References

- [1] Abraham, A. and N. Pavoni (2008), “Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation”, *Review of Economic Dynamics* 11(4): 781-803.
- [2] Aiyagari, S.R. (1994), “Uninsured Idiosyncratic Risk and Aggregate Saving”, *Quarterly Journal of Economics* 109(3): 659-84.
- [3] Allen, F. (1985), “Repeated principal-agent relationships with lending and borrowing”, *Economics Letters* 17(1-2): 27-31.
- [4] Atkeson, A. and R. Lucas (1992), “On efficient distribution with private information”, *Review of Economic Studies* 59: 427-53
- [5] Bolton, P and M. Dewatripont (2004), *Contract Theory*, The MIT Press, Cambridge, MA.
- [6] Budría-Rodríguez, S., J. Díaz-Giménez, V. Quadrini and J.-V. Ríos-Rull (2002), “Updated Facts on the U.S. Distributions of Earnings, Income, and Wealth”, *Federal Reserve Bank of Minneapolis Quarterly Review* 26(3): 2-35.
- [7] Castañeda, A., J. Díaz-Giménez and J.-V. Ríos-Rull (2003), “Accounting for the U.S. Earnings and Wealth Inequality”, *Journal of Political Economy* 111(4): 818-57.
- [8] Cole, H. and N. Kocherlakota (2001), “Efficient Allocations with Hidden Income and Hidden Storage”, *Review of Economic Studies* 68(3): 523-42.
- [9] Doepke, M. and R. Townsend (2006), “Dynamic mechanism design with hidden income and hidden actions”, *Journal of Economic Theory* 126(1): 235-85.
- [10] Fernandes, A. and C. Phelan (2000), “A Recursive Formulation for Repeated Agency with History Dependence”, *Journal of Economic Theory* 91, pp. 223-47.
- [11] Golosov, M., A. Tsyvinski and I. Werning (2006), “New Dynamic Public Finance: A User’s Guide”, NBER Chapters, in: NBER Macroeconomics Annual 2006, Volume 21, 317-388, National Bureau of Economic Research, Inc.
- [12] Green, E. (1987), “Lending and Smoothing of Uninsurable Income”, in E.C. Prescott and N. Wallace (eds), *Contractual Arrangements for Intertemporal Trade* pp. 3-25. University of Minnesota Press.
- [13] Karaivanov, A. and F. Martin (2011), “Dynamic Optimal Insurance and Lack of Commitment”, manuscript, Simon Fraser University.
- [14] Kocherlakota, N. (1996), “Implications of Efficient Risk Sharing without Commitment”, *Review of Economic Studies* 63(4): 595-609.
- [15] Kocherlakota, N. (2004), “Figuring out the Impact of Hidden Savings on Optimal Unemployment Insurance”, *Review of Economic Dynamics* 7(3): 541-54.
- [16] Krueger, D. and H. Uhlig (2006), “Competitive risk-sharing contracts with one-sided commitment”, *Journal of Monetary Economics* 53, pp. 1661-91.
- [17] Ligon, E., J. Thomas and T. Worrall, (2002) “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies”, *Review of Economic Studies* 69(1): 209-44.

- [18] Maskin, E. and J. Tirole (2001), “Markov Perfect Equilibrium”, *Journal of Economic Theory* 100(2): 191-219.
- [19] Phelan, C. (1995), “Repeated Moral Hazard and One-Sided Commitment”, *Journal of Economic Theory* 66(2): 488-506.
- [20] Phelan, C. (1998), “On the Long Run Implications of Repeated Moral Hazard”, *Journal of Economic Theory* 79: 174-91.
- [21] Phelan, C. and R. Townsend (1991), “Computing Multi-Period, Information-Constrained Equilibria”, *Review of Economic Studies* 58: 853-81.
- [22] Rogerson, W. (1985a), “The First Order Approach to Principal-Agent Problems”, *Econometrica* 53: 1357-68.
- [23] Rogerson, W. (1985b) “Repeated Moral Hazard”, *Econometrica* 53: 69–76.
- [24] Schechtman, J. and V. Escudero (1977), “Some Results on an Infinite Fluctuation Problem”, *Journal of Economic Theory* 16: 151-66.
- [25] Spear, S. and S. Srivastava (1987), “On Repeated Moral Hazard with Discounting”, *Review of Economic Studies* 54(4): 599-617.
- [26] Thomas, J. and T. Worrall (1988), ”Self-Enforcing Wage Contracts”, *Review of Economic Studies* 55(4): 541-553.
- [27] Thomas, J. and T. Worrall (1990), “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem”, *Journal of Economic Theory* 51(2): 367-90.
- [28] Thomas, J. and T. Worrall (1994), “Foreign Direct Investment and the Risk of Expropriation”, *Review of Economic Studies* 61(1): 81-108.
- [29] Townsend, R. (1982), “Optimal Multi-period Contracts and the Gain from Enduring Relationships Under Private Information”, *Journal of Political Economy* 61: 166-86.

A Autarky

This appendix characterizes the agent’s autarky, or self-insurance, problem which is one possible way to determine his outside option when contracting with the insurer. The timing in each period is as follows. First, the agent decides how much effort e to put in, then output y^i is realized, and finally the agent decides how much to consume, c^i and to save, a^i in each state $i = 1, \dots, n$. The agent faces the state-by-state budget constraint, $c^i + a^i = ra + y^i$. Applying standard arguments, write the agent’s self-insurance problem recursively as

$$\Omega(a) = \max_{\{a^i\}_{i=1}^n, e} \sum_{i=1}^n \pi_i(e) [u(ra + y^i - a^i) + \beta\Omega(a^i)] - e,$$

where $\Omega(a)$ is the agent’s value function. Our assumptions on u , the condition $r < \beta^{-1}$ and the upper bound on assets are sufficient for the autarky problem to be well-defined. Alternatively, assuming $r < \beta^{-1}$ and DARA utility it is easy to adapt the proof of Schechtman and Escudero (1977) to our environment with endogenous income distribution showing that assets stay bounded (details available upon request).

The autarky value function $\Omega(a)$ is strictly increasing in agent’s assets. Intuitively, since $r < \beta^{-1}$, the agent saves only to insure against current and future consumption volatility. An agent with more assets can do everything a poorer agent can, but is in a better position to self-insure against a long sequence of low outputs.

As is standard in self-insurance models, consumption smoothing is imperfect (c^i differ across states with different y^i). Other features of the autarky solution are that optimal consumption and savings in each state are increasing in assets, a . Assets are reduced if the agent is in the lowest income state(s) and increased (for some asset range) if the agent is in the highest state(s).⁸

B Equivalence between MPE and one-sided commitment contract

B.1 Recursive one-sided commitment problem

As we show in Karaivanov and Martin (2011), without loss of generality, the one-sided commitment contract described in section 4 can be written as a two-stage recursive problem where agent’s assets are taken away by the insurer in the initial period and replaced by future utility promises.⁹ The reason is that future assets and utility promises are completely interchangeable to an insurer who has full commitment—any allocation can be implemented via infinitely many appropriately chosen combinations of these two instruments. In the first stage of the one-sided commitment problem, insurers solve a static problem offering the agent the maximum possible utility subject to breaking-even and incentive-compatibility. Without loss of generality an insurer sets the agent’s asset future holdings to their minimum, \underline{a} and, in exchange, promises the agent lifetime utility w_0^i . Since the agent can walk away at the beginning of each period, the promised utility needs to be at least as high as the agent’s outside option: the value of contracting with another insurer starting with \underline{a} assets (since ω_0^i enters the objective with positive sign this constraint will not bind typically). Let \underline{w} be the value of this outside option, which is endogenous and will be determined below.

⁸The proofs of these statements (available upon request) are relatively standard and hence, omitted.

⁹In that paper, we show this result in the full-information setting, but it is straightforward to adapt the proof for the moral hazard case here.

The first-stage ($t = 0$) problem of the insurer is

$$\mathcal{V}^C(a_0) \equiv \max_{\{c_0^i, \omega_0^i\}_{i=1}^n, e_0} \sum_{i=1}^n \pi^i(e_0) [u(c_0^i) + \beta \omega_0^i] - e_0 \quad (21)$$

subject to

$$\begin{aligned} \sum_{i=1}^n \pi_e^i(e_0) [u(c_0^i) + \beta \omega_0^i] - 1 &= 0 \\ \sum_{i=1}^n \pi^i(e_0) [y^i + r a_0 - c_0^i - \underline{a} + r^{-1} \Pi^C(\omega_0^i)] &= 0 \\ \omega_0^i - \underline{w} &\geq 0, \forall i. \end{aligned}$$

The first constraint ensures incentive-compatibility. The second constraint is the “zero ex-ante profits” condition for the insurer. The third constraint ensures the agent remains in the contract from tomorrow on, for any realization of current output.

The function $\Pi^C(\omega)$ solves the following, second-stage problem for any $\omega \geq \underline{w}$:

$$\Pi^C(\omega) = \max_{\{c^i, \omega^i\}_{i=1}^n, e} \sum_{i=1}^n \pi^i(e) [y^i + (r-1)\underline{a} - c^i + r^{-1} \Pi^C(\omega^i)] \quad (22)$$

subject to

$$\begin{aligned} \sum_{i=1}^n \pi_e^i(e) [u(c^i) + \beta \omega^i] - 1 &= 0 \\ \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta \omega^i] - e &= \omega \\ \omega^i - \underline{w} &\geq 0, \forall i. \end{aligned}$$

The $(r-1)\underline{a}$ term in the objective reflects that agent’s assets are fixed at \underline{a} each period. The first constraint ensures incentive compatibility with the agent’s unobserved effort choice. The second constraint, referred to as “promise keeping” embodies the commitment ability of insurers to always deliver on their utility promise ω . The last constraint reflects the bound on future promises implied by the agent’s inability to commit for more than one period.

In a competitive equilibrium, by free entry we must have that $\underline{w} = \mathcal{V}^C(\underline{a})$.

B.2 Auxiliary results

To prove Proposition 2 we will use two auxiliary Lemmas.

Lemma 1 *For any $\Delta\tau > 0$ there exist $\varepsilon^i > 0$, $i = 1, \dots, n$ such that $\sum_{i=1}^n \varepsilon^i = \Delta\tau$ and $\sum_{i=1}^n \pi_e^i [u(c^i + \varepsilon^i) - u(c^i)] = 0$.*

Proof. Since $\sum_i \pi_e^i = 0$, it is enough to show that we can choose ε^i so that $u(c^i + \varepsilon^i) - u(c^i) = b$ where b is some appropriately chosen positive constant. From the strict monotonicity of u a unique solution to this equation in ε^i on $[0, \infty)$ exists and is strictly increasing in b . Call this solution $\phi^i(b, c^i)$. Since $\sum_i \phi^i(b, c^i)$ is also strictly increasing in b and since $\phi^i(0, c^i) = 0$ for all c^i , $\exists b > 0$ that solves $\sum_i \phi^i(b, c^i) = \Delta\tau$ for any $\Delta\tau > 0$. ■

Let $P1$ be the problem of a competitive insurer as stated in (2)—(4). Let $P2$ be the problem of a profit-maximizing insurer facing an agent with general outside option $B(a)$, which we write as follows:

$$\Pi(a) = \max_{\{\tau^i, a^i, e\}_{i=1}^n} \sum_{i=1}^n \pi^i(e) \left[y^i - \tau^i + \frac{\Pi(a^i)}{r} \right]$$

subject to

$$\begin{aligned} \sum_{i=1}^n \pi_e^i(e) [u(c^i) + \beta B(a^i)] - 1 &= 0 \\ \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta B(a^i)] - e - B(a) &= 0. \end{aligned}$$

Lemma 2 *Equivalence of problems P1 and P2 for $B(a) = \mathcal{V}(a)$.*

- (i) For $B(a) = \mathcal{V}(a)$ in P2 any MPE solving P2 features $\Pi(a) = 0$ for all $a \in \mathbb{A}$ and is also a solution to P1.
- (ii) Any MPE solving P1 is a solution to P2 at $B(a) = \mathcal{V}(a)$.

Proof. For part (i), call S1 a MPE solving P1 and S2 a MPE solving P2. Any P2 solution at $B(a) = \mathcal{V}(a)$ satisfies $\mathcal{V}(a) = \sum_{i=1}^n \pi^i(e_2) [u(c_2^i) + \beta \mathcal{V}(a_2^i)] - e_2$, i.e., it achieves the same value, $\mathcal{V}(a)$ as S1. Clearly, at $B(a) = \mathcal{V}(a)$ S2 also satisfies (3). However, could it violate (4)? Suppose yes. Suppose first (4) did not bind at S2 so that its expected transfers are too low relative to expected output, i.e., $\sum_{i=1}^n \pi^i(e_2) [y^i - \tau_2^i] > 0$. But then, in P1 we can increase transfers starting from S2 while satisfying (3) (see Lemma 1) until (4) binds and get higher ex-ante value for the agent—recall that S2 achieves $\mathcal{V}(a)$. This is a contradiction with the optimality of S1. Next, suppose (4) did not bind in the opposite direction (transfers, τ_2^i are too high relative to expected output at S2). This implies that S1 has lower expected transfers than S2 since it satisfies (4) at equality. Thus, at $B(a) = \mathcal{V}(a)$ S1 satisfies all constraints of P2 but yields higher profits than S2 (due to its lower expected transfers) which contradicts the optimality of S2. Therefore, it must be that S2 at $B(a) = \mathcal{V}(a)$ satisfies (4), which in turn, implies that its associated profits, $\Pi(a)$ are identically zero for any $a \in \mathbb{A}$.

For part (ii), note that S1 is feasible for P2 at $B(a) = \mathcal{V}(a)$ and yields $\Pi(a) = 0$ for all $a \in \mathbb{A}$. Suppose, however, S1 is not optimal for P2, i.e., profits, $\Pi(a)$ at S2 are actually positive. But then, since S2 satisfies (3) and achieves $\mathcal{V}(a)$, as does S1, then back in P1 we can increase expected transfers starting from S2 until (4) (zero profits) is satisfied while keeping (3) satisfied (see Lemma 1) which would generate a higher value $\mathcal{V}(a)$ —a contradiction with the optimality of S1. ■

B.3 Proof of Proposition 2

Proof. Define $\hat{\Pi}(a) = \Pi(a) - ra$ where Π is the profits function in problem P2 above. Since B is assumed strictly increasing, we can perform a change of variables by calling $w = B(a)$ and call $\bar{\Pi}(w) \equiv \hat{\Pi}(B^{-1}(w))$. This yields the following mathematically equivalent formulation of P2:

$$\bar{\Pi}(w) = \max_{\{\tau^i, w^i, e\}_{i=1}^n} \sum_{i=1}^n \pi^i(e) \left[y^i - c^i + \frac{\bar{\Pi}(w^i)}{r} \right] \quad (23)$$

subject to

$$\begin{aligned} \sum_{i=1}^n \pi_e^i(e) [u(c^i) + \beta w^i] - 1 &= 0 \\ \sum_{i=1}^n \pi^i(e) [u(c^i) + \beta w^i] - e - w &= 0 \\ w^i - B(\underline{a}) &\geq 0. \end{aligned}$$

Note that in the second-stage of the one-sided commitment problem (22), calling $\bar{\Pi}^C(w) \equiv \Pi^C(w) - r\underline{a}$ and $\underline{w} = B(\underline{a})$ yields an equivalent problem to (23) for any function $B(a)$ so their solutions must coincide. In terms of value functions, we must have

$$\Pi^C(w) - r\underline{a} = \Pi(a) - ra. \quad (24)$$

Also, at $B(a) = \mathcal{V}(a)$ we know by Lemma 2 that problem (23), which is equivalent to P2 for any $B(a)$, will have the same solution(s) as problem P1 defined in (2)—(4), starting from the same a .

Changing variables from w to $B(a)$, the first-stage of the one-sided commitment problem (21) is equivalent to:

$$\max_{\{c_0^i, a_0^i, e_0\}_{i=1}^n} \sum_{i=1}^n \pi^i(e_0) [u(c_0^i) + \beta B(a_0^i)] - e_0 \quad (25)$$

subject to

$$\begin{aligned} \sum_{i=1}^n \pi_e^i(e_0) [u(c_0^i) + \beta B(a_0^i)] - 1 &= 0 \\ \sum_{i=1}^n \pi^i(e_0) [y^i + ra_0 - c_0^i - \underline{a} + r^{-1}\Pi^C(B(a_0^i))] &= 0 \\ B(a_0^i) - B(\underline{a}) &\geq 0, \forall i. \end{aligned}$$

Plugging in for Π^C in terms of Π from (24), the second constraint can be rewritten as

$$\sum_{i=1}^n \pi^i(e_0) [y^i - \tau_0^i + r^{-1}\Pi(a_0^i)] = 0, \quad (26)$$

where $\tau_0^i = c_0^i + \underline{a} - ra_0$. From the second-stage problem (22) (which was shown to be equivalent to (23) and hence to P2 earlier on) we have $\Pi(a_0^i) = 0$ at $B(a) = \mathcal{V}(a)$ —see Lemma 2, part (i). Thus, constraint (26) becomes $\sum_{i=1}^n \pi^i(e_0) [y^i - \tau_0^i] = 0$ and so the first-stage problem (25) at $B(a) = \mathcal{V}(a)$ is equivalent to problem P1, (2)—(4) starting from the same a_0 . Overall, we have shown that the one-sided commitment problem with free entry by insurers (21)—(22) is equivalent to P1—the Markov-perfect insurance problem with free entry. Thus, starting at the same initial asset level, $a_0 \in \mathbb{A}$, we obtain equivalence of consumption and effort allocations, for any given sequence of output realizations, as written in the proposition statement. ■