Journal of Economic Dynamics & Control 94 (2018) 1–23

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Markov-perfect risk sharing, moral hazard and limited commitment☆

Alexander K. Karaivanov[a], Fernando M. Martin[b,*]

[a] Simon Fraser University, Canada
[b] Federal Reserve Bank of St. Louis, United States

ARTICLE INFO

Article history:
Received 11 September 2017
Revised 26 June 2018
Accepted 27 June 2018
Available online 8 July 2018

JEL classification:
D11
E21

Keywords:
Markov-perfect equilibrium
Risk-sharing
Limited commitment
Moral hazard
Consumption smoothing

ABSTRACT

We define, characterize and compute Markov-perfect risk-sharing contracts in a dynamic stochastic economy with endogenous asset accumulation and simultaneous limited commitment and moral hazard frictions. We prove that Markov-perfect insurance contracts preserve standard properties of optimal insurance with private information and are not more restrictive than a long-term contract with one-sided commitment. Markov-perfect contracts imply a determinate asset time-path and a non-degenerate long-run stationary wealth distribution. Quantitatively, we show that Markov-perfect risk-sharing contracts provide sizably more consumption smoothing relative to self-insurance and that the welfare gains from resolving the commitment friction are larger than the gains from resolving the moral hazard friction at low asset levels, while the opposite holds for high asset levels.

1. Introduction

We define, characterize and compute Markov-perfect risk-sharing contracts in a dynamic economy with stochastic income and endogenous asset accumulation by risk-averse agents. We highlight the roles of two market frictions that simultaneously affect the degree of consumption and income smoothing: (i) limited commitment, by which we mean inability to use long-term contracts; and (ii) private information, in the form of moral hazard. Importantly, the agents' wealth interacts with both frictions endogenously, as it affects their demand for insurance and their incentives and intertemporal trade-offs.

We define Markov-perfect insurance as a sequence of one-period incentive-compatible risk-sharing contracts which depend only on payoff-relevant variables. We characterize, both theoretically and quantitatively, the properties of Markov-perfect risk-sharing contracts and compare them to two commonly studied alternatives: self-insurance and long-term contracts. We also quantify the size and distribution of the welfare costs from the commitment and information frictions across agents with different wealth.

☆ We thank the editor, Ravikumar, two anonymous referees, Abraham, Camera, Carroll, Golosov, Green, Pavoni, Quadrini and audiences at CEF, Queen's, SED, Stony Brook and UC Santa Barbara for helpful comments and suggestions. Karaivanov acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada (Grant no. 435-2013-0698). The views expressed in this paper do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

* Corresponding author.
E-mail address: fernando.m.martin@stls.frb.org (F.M. Martin).

https://doi.org/10.1016/j.jedc.2018.06.007

Markov-perfect insurance restricts the contract space to recursive policy functions of observable, payoff-relevant fundamentals: current output and assets. However, we show that in our setting the sequence of one-period Markov-perfect risk-sharing allocations are equivalent to the allocation obtained in a mechanism-design problem with an infinitely-long contract. Specifically, we prove a one-to-one mapping between a Markov-perfect equilibrium (MPE), in which the dynamic state variable is the agent's assets, and a setting with long-term commitment by the insurer, in which the state variable is promised utility. The key assumptions delivering this equivalence result are free entry and non-exclusivity in the insurance market, and the ability of agents to carry assets over time, facing the same rate of return as insurers.

The allocation equivalence obtains regardless of whether the agent's savings or effort are contractible or not. Intuitively, the agent's assets "encode" the history of income shocks and thus completely determine the current and future terms of risk-sharing, just like promised utility in the standard approach. When the agent is subject to a borrowing constraint, as we assume, this constraint maps into a lower bound on promised utility. As we discuss in more detail below, these results are related but also significantly different from recent contributions by Albanesi and Sleet (2006) and Khan et al. (2014).[1]

MPE as a solution concept highlights the main risk-sharing friction on which we focus: the fact that in many situations individuals and firms cannot, or are legally not allowed to, enter binding long-term agreements (e.g., labor, rental, insurance, TV, phone, internet, etc.) Instead, both insurers and the insured can only sign short-term contracts and can costlessly walk away at specified times.

While in our model the risk-sharing outcome does not depend on whether the insurers can commit to long-term contracts or not, Markov-perfect contracts lend themselves more directly to quantitative analysis and empirical work as the agents' asset holdings are an integral determinant of the contract, with non-trivial endogenous dynamics and non-degenerate limiting distributions that can be taken to data, as we demonstrate. This differs from the mechanism-design problem with one-sided commitment by the insurer in which the contracts and allocations are expressed in terms of the mathematical abstraction of promised utility as the state variable and asset dynamics are indeterminate.

Our equivalence result can also be interpreted as a *decomposition* of a long-term insurance contract with limited commitment and private information into a sequence of short-term contracts that are only a function of agents' assets and current income. Such decomposition does not emerge in many dynamic contracting settings in which there are gains from enduring relationships (e.g., Townsend, 1982).[2]

We model risk-averse agents endowed with a stochastic technology that transforms labor effort into output. An agent can imperfectly self-insure through accumulating or drawing down a risk-free asset. A perfectly competitive risk-neutral insurer offers a risk-sharing contract. The agent's assets and output are observable, but the agent's effort is not observable, which gives rise to a moral hazard problem. Though the insurer observes the agent's assets, he cannot control them. That is, given the contract terms (insurance premia/transfers), the agent makes his own effort, consumption and savings decisions.

We first show that Markov-perfect risk-sharing contracts with moral hazard provide partial insurance and are characterized by "inverse Euler equations" relating the reciprocals of current and future marginal utility of consumption, therefore preserving the standard properties of optimal insurance with private information from the literature. Similarly, Markov-perfect insurance contracts preserve the standard properties of optimal insurance with full information: consumption is equalized across states of the world and the relationship between present and future consumption is described by an Euler equation equalizing current marginal utility of consumption with appropriately discounted future marginal utility. Assuming free entry in the insurance market, we prove and numerically verify that Markov-perfect insurance contracts yield history-contingent time paths for consumption and effort which are equivalent to the time paths implied by an infinitely-long contract with commitment by the insurer.

We use numerical methods to further characterize the properties of Markov-perfect insurance contracts. The MPE problem is tractable, low-dimensional and relatively easy to compute, both with and without private information. We show that an MPE can be parameterized to match several broad dimensions of US data. We find that Markov-perfect insurance contracts provide sizeable additional consumption smoothing relative to self-insurance, particularly for agents with low asset holdings. Even in the presence of very volatile output-shock realizations, access to Markov-perfect insurance allows agents to smooth consumption to a considerably higher degree compared to when only relying on their own savings. In addition, the transfers implemented by the Markov-perfect insurance contracts imply a much smoother income process for the agent.

The properties of Markov-perfect insurance described above have important consequences for wealth inequality. In an MPE a large fraction of agents have zero assets in the limiting distribution, since poorer agents have weaker incentive to supply effort and accumulate assets when they have access to outside insurance. Our model also delivers a long-term consumption distribution that is broadly in line with the data, avoiding the counterfactual degree of left skewness in models with limited commitment alone. In addition, our approach allows us to compute the welfare costs of moral hazard and limited commitment for any level of wealth. For our parameterization, we find sizable gains from resolving either friction which are the highest at low-wealth levels, near the borrowing constraint. We also find that the gains from resolving the commitment friction are larger than the gains from resolving the moral hazard friction at low asset levels, while the opposite

[1] See also DeMarzo and Sannikov (2006) who use continuous-time methods and show how a dynamic contracting problem with hidden cash flow and unobserved effort can be decentralized via the firm's capital structure (credit line, debt and equity).

[2] In a broader sense, rewriting a dynamic problem recursively in terms of promised utility reduces it to a sequence of one-period problems that can be decentralized via "component planners" who trade contracts (see Atkeson and Lucas, 1992). However, a decentralization in terms of other observable variables such as wealth, debt, etc. is not always available (see Golosov et al., 2016).