

INSURANCE AND INEQUALITY WITH PERSISTENT PRIVATE INFORMATION

ALEXANDER W. BLOEDEL
Department of Economics, UCLA

R. VIJAY KRISHNA
Department of Economics, Florida State University

OKSANA LEUKHINA
Research Division, Federal Reserve Bank of St. Louis

We study the implications of optimal insurance provision for long-run welfare and inequality in economies with persistent private information. A principal insures an agent whose private type follows an ergodic, finite-state Markov chain. The optimal contract always induces *immiseration*: the agent's consumption and utility decrease without bound. Under positive serial correlation, it also *backloads high-powered incentives*: the sensitivity of the agent's utility with respect to his reports increases without bound. These results extend—and help elucidate the limits of—the hallmark immiseration results for economies with i.i.d. private information. Numerically, we find that persistence yields faster immiseration, higher inequality, and novel short-run distortions. Our analysis uses recursive methods for contracting with persistent types and allows for binding global incentive constraints.

KEYWORDS: Immiseration, insurance, inequality, backloaded incentives, recursive contracts, persistent private information.

1. INTRODUCTION

The idea that a society's income distribution arises, in large part, from the way it deals with individual risks is a very old and fundamental one, one that is at least implicit in all modern studies of distribution.

[Lucas \(1992, p. 234\)](#)

Many contemporary debates over rising inequality concern institutions designed to facilitate risk-sharing, such as social insurance programs and redistributive tax systems. More generally, there is a fundamental connection between the way that a society provides insurance and its degree of inequality. In particular, because individuals typically have private information about their idiosyncratic risks—such as shocks to income, tastes, or productivity—and must be incentivized to reveal that information, the markets for insuring against these risks are inevitably incomplete. Consequently, the way that a society resolves the tradeoff between providing insurance and incentives determines how these private shocks affect individuals' consumption, and hence how the population distributions of consumption and wealth evolve over time. What, then, is the *optimal* degree of inequality?

Alexander W. Bloedel: abloedel@econ.ucla.edu

R. Vijay Krishna: rvk3570@gmail.com

Oksana Leukhina: oksana.m.leukhina@gmail.com

This paper subsumes the working paper [Bloedel and Krishna \(2015\)](#). We thank the Co-Editor, anonymous referees, Andy Atkeson, Simon Board, Gabriel Carroll, Hector Chade, V.V. Chari, Sebastian Di Tella, Ed Green, Christian Hellwig, Pablo Kurlat, Paul Milgrom, Andy Newman, Shunya Noda, Ilya Segal, Andy Skrzypacz, Bruno Strulovici, Bob Wilson, and various seminar and conference audiences for many helpful comments. Aristotle Magganas provided superb research assistance. Bloedel is grateful to the B.F. Haley and E.S. Shaw Fellowship for Economics for financial support and to Caltech for its hospitality while some of this work was completed. Krishna is grateful to the Economic Department at Duke University, where he was a visitor at the inception of this project.

The classic answer is striking: in private information economies, the social optimum induced by a utilitarian planner with full commitment generates *immiseration*. That is, in the long run, almost all individuals become completely impoverished while a vanishing lucky few consume all of society’s resources (Green, 1987, Thomas and Worrall, 1990). Furthermore, cross-sectional consumption and wealth inequality increase without bound, so the economy does not converge to a well-defined steady-state (Atkeson and Lucas, 1992). In other words, there is no meaningful long-run tradeoff between efficiency and equity: the optimal provision of incentives demands destitution and infinite inequality.

This extreme conclusion has sparked debates over immiseration’s intuitive appeal and inspired the search for alternative models.¹ Nonetheless, immiseration is “often regarded as being the hallmark result of dynamic social contracting in the presence of private information” (Kocherlakota (2010, p. 70)) and constitutes an apparently fundamental feature of the workhorse normative models that economists use to study optimal risk-sharing. It is therefore important to understand its robustness within this canonical class of models.

A key gap in this understanding concerns the dynamic nature of individuals’ private information. At one extreme, the classic literature universally assumes that private information is i.i.d. over time (i.e., shocks are “completely transient”), which is analytically convenient but unrealistic: shocks to individuals’ incomes, tastes, and productivities are often highly persistent.² At the opposite extreme, in an influential paper, Williams (2011) shows that when individuals’ private types follow a random walk (i.e., shocks are “permanent”), immiseration not only disappears, but the optimal contract generates long-run *bliss*: it sends individuals’ consumption and utility to their *upper* bounds. This raises two questions. First, what is driving these diametrically opposing results? Second, what long-run outcomes emerge in the generic and empirically relevant intermediate case in which private information is imperfectly persistent?

In this paper, we study the long-run properties of optimal insurance contracts under general forms of persistent private information. We address the second question by showing that immiseration arises for a broad class of *ergodic* private information processes. While prior work (Zhang (2009)) has considered specific parametric examples, our generality permits us, in essence, to interpolate between the i.i.d. and permanent shock benchmarks. Consequently, we also elucidate the first question by identifying “mean-reversion” of the private information process as a key determinant of long-run outcomes. In particular, our analysis suggests that immiseration fails *only* in the knife-edge case of permanent shocks.

Model. We study a canonical principal-agent formulation of the planning problem.³ A risk-neutral principal (she) interacts with a risk-averse agent (he) over an infinite horizon in discrete

¹Various authors have critiqued immiseration on normative grounds (it seems perverse that *ex ante* efficiency should require *ex post* destitution), descriptive grounds (commitment to one’s own impoverishment may not be enforceable), and practical grounds (without a well-defined steady-state, one cannot meaningfully analyze the long-run tradeoff between equity and efficiency). Correspondingly, a number of papers provide foundations for bounded long-run inequality by appealing to alternative models with limited commitment, different normative criteria, and various other features (see Section 1.1 for references).

²See Storesletten, Telmer, and Yaron (2004), Meghir and Pistaferri (2004) for evidence on income.

³Specifically, we consider the “open economy” planning problem of a utilitarian planner who maximizes the average welfare of a large population of agents subject to an *intertemporal* resource constraint (Green, 1987). By standard duality arguments (Golosov, Tsyvinski, and Werquin, 2016b), such problems can be equivalently “decentralized” into a collection of one-on-one principal-agent problems in which the principal minimizes the expected cost of delivering a particular lifetime utility to each agent. Following Thomas and Worrall (1990) and others, we analyze this principal-agent version of the problem. We revisit the planning interpretation in Sections 5.1 and 6.3 and, in the latter, discuss how our analysis relates to the “closed economy” problem in which the planner faces a *per-period* resource constraint (Atkeson and Lucas, 1992).

time. The agent's preferences in each period are determined by his privately observed *type*, which evolves stochastically over time. The principal provides insurance to the agent via an infinite-horizon insurance contract—which specifies the agent's allocation in each period conditional on the history of reported types—with the goal of achieving constrained Pareto optimality, i.e., minimizing costs while delivering a pre-specified lifetime utility to, and eliciting truthful reports from, the agent. Both parties commit to the contract at the initial date and discount the future at the same rate. For concreteness, we focus on a setting in which the agent's type corresponds to his privately observed endowment (as in [Green, 1987](#), [Thomas and Worrall, 1990](#)), but this is not essential.

We make two main assumptions about the environment. First, the agent's type evolves according to a fully connected, finite-state Markov chain. This implies that the agent's type process is ergodic and bounded but allows for otherwise arbitrary serial correlation, and thus permits forms of persistence studied in the recent theoretical literature and the kinds of asymmetric and skewed shocks documented in the recent empirical literature.⁴ Second, the agent can neither covertly save outside of the contract nor over-report his type within the contract (though he can freely under-report his type). In our favored interpretation, this captures the idea that the agent cannot covertly engage in trade or production.⁵

Main Results. Our main result, [Theorem 1](#), shows that the optimal contract always generates *immiseration*: the agent's consumption and utility converge in probability (and, under certain conditions, almost surely) to their lower bounds. This result encompasses prior immiseration results in i.i.d. settings ([Green, 1987](#), [Thomas and Worrall, 1990](#)) and specific settings with persistence ([Zhang, 2009](#)). Our second result, [Theorem 2](#), shows that when the agent's type is positively serially correlated, the optimal contract additionally features *backloaded high-powered incentives*: in the long-run, the sensitivity of the agent's continuation utility with respect to his reports increases without bound. This can be viewed as a kind of *relative immiseration* wherein the impact of an incrementally higher type realization is magnified over time: the welfare difference between two agents with identical histories before time T , but with different shocks at time T , grows unboundedly as $T \rightarrow \infty$.

A simple intuition for these results is as follows. Incentive compatibility requires that high and low endowment reports be, respectively, rewarded and punished with higher and lower average transfers from the principal in future periods. Thus, the agent's continuation utility must vary with his reported endowment, with larger variation corresponding to higher-powered incentives. *Ceteris paribus*, the *cost of incentive provision* is lower for the principal when the agent's continuation utility is *lower*: due to risk aversion, this is precisely when the agent's *marginal utility of consumption is higher*, so that a given variation in utility can be induced by smaller variations in consumption. Thus, by immiserating the agent, the principal drives her cost of incentive provision to zero and makes it affordable to provide arbitrarily high-powered

⁴Recent theoretical studies of dynamic contracting primarily focus on either binary-state Markov chains or continuous-state ergodic Markov processes (see [Footnote 6](#) and [Section 1.1](#) for references). Our framework includes the former processes and distributional approximations of the latter ones. Empirically, [Arellano, Blundell, and Bonhomme \(2017\)](#) and [Güvenen, Karahan, Ozkan, and Song \(2021\)](#) find that the relative persistence of positive and negative income shocks depends sensitively on the individual's current income level.

⁵The restriction on savings is standard (cf. [Allen \(1985\)](#); [Cole and Kocherlakota \(2001b\)](#)). The restriction on over-reporting is common in the literature and often natural. When the agent's private type is his endowment, it amounts to the assumption that he cannot covertly borrow outside of the contract ([Williams \(2011\)](#); [Bloedel, Krishna, and Strulovici \(2023\)](#)). When his private type is his labor productivity, it amounts to the assumption that he cannot covertly engage in home production ([Golosov and Tsyvinski \(2007\)](#)). Under these interpretations, these restrictions permit us to isolate the incentive constraints corresponding to the agent's exogenous private information from any other constraints arising from imperfect enforcement of contractual terms. See [Section 6.1](#) for further discussion.

incentives in later periods. Such *backloading* of high-powered incentives allows the principal to provide better insurance in early periods while maintaining incentive compatibility, and thus to optimally smooth her costs over time.

While similar intuitions have appeared in the classic i.i.d. setting (Thomas and Worrall, 1990), formalizing them is significantly more challenging in settings with persistence. The key subtlety is that, with persistence, the agent has two sources of information rents. First, there is the classic *i.i.d. information rent* arising from the agent’s private information about his current type, and hence his *static* preferences over allocations in the current period. Second, there is the new *Markov information rent* arising from the agent’s private information about the distribution of his future types, and hence his *intertemporal* preferences over the path of allocations in the “continuation contract” for future periods. Due to the latter, the principal can no longer screen the agent’s static preferences by using continuation utility as a type-independent “numeraire” with which to deliver rewards and punishments. Instead, she must also screen his intertemporal preferences through continuation contracts that carefully spread those rewards and punishments over multiple periods. If the agent’s type is sufficiently persistent, these new intertemporal screening considerations might dominate the classic motive to backload incentives and yield qualitatively different contractual dynamics (as in the permanent shock benchmark of Williams (2011)). In essence, our key lesson is that, given *any* mean-reversion in the agent’s type, the classic i.i.d. forces remain dominant *in the long run* and render immiseration optimal.

Our formal proofs of Theorems 1 and 2 involve three main steps, the latter two of which constitute our primary methodological contribution. First, we formulate the principal’s problem recursively, using as the main state variable a vector of *interim promised utilities* that encode the agent’s continuation utility conditional on each possible current type. Second, we identify a specific directional derivative of the principal’s value function with respect to this state variable—her *marginal cost martingale*—as the appropriate object with which to analyze long-run outcomes. This step builds on Thomas and Worrall’s (1990) observations that, in the i.i.d. special case, (i) the principal’s optimal cost-smoothing implies that her marginal cost of providing continuation utility to the agent follows a martingale process, and (ii) the optimal contract’s long-run behavior is pinned down by convergence properties of this martingale. However, adapting these ideas to the general Markovian setting requires our analysis to overcome several nontrivial challenges. Third, in the final step, we identify a novel “renewal” property of the martingale’s dynamics that lets us characterize long-run outcomes using probabilistic arguments.

Notably, this analysis neither (i) imposes any restrictions (beyond finiteness) on the number of agent types nor (ii) relies on the “first-order approach” in which one studies a relaxed problem that ignores the agent’s global incentive constraints. In contrast, most work on contracting with persistent types either studies binary-type settings, which is restrictive, or studies continuous-type settings using the first-order approach, which can yield solutions that violate the agent’s global incentive constraints.⁶ Moving beyond these two approaches has been identified as an important goal for the contracting literature.⁷

⁶See Section 1.1 for references to recent binary-type models. The first-order approach to dynamic mechanism design is developed in Pavan, Segal, and Toikka (2014) and applied to dynamic insurance problems in Williams (2011), Kapička (2013), Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2016a), Hellwig (2021), Brendon (2022). In a monopolistic screening setting, Battaglini and Lamba (2019) show that, if period lengths are short, the first-order approach fails for a generic class of finite-state Markov type processes; when it fails, the candidate contract it produces can differ qualitatively from the true optimal contract. It is natural to expect that these lessons also apply to dynamic insurance problems.

⁷See Pavan (2016), Garrett, Pavan, and Toikka (2018), Battaglini and Lamba (2019).

Solved Examples. To obtain a more detailed description of the optimal contract, we specialize to the setting with CARA utility and two endowment levels (which facilitates simple graphical depictions and direct comparisons to prior work). We analytically characterize key structural properties of the optimal contract and use numerical analysis to fully solve for it. This exercise yields two main sets of insights.

First, persistence affects the *speed* of immiseration. To illustrate this, we simulate many sample paths of the optimal contract and study how the implied consumption distribution evolves over time. We interpret this as describing the dynamics of the cross-sectional distribution of consumption in an “open economy” that can borrow and lend in the world market, as in [Green \(1987\)](#) (cf. [Footnote 3](#)). As persistence increases, we find that both the mean and variance of consumption diverge *more quickly* (to minus and plus infinity, respectively). That is, the economy’s aggregate consumption decreases without bound and its consumption inequality increases without bound, and both occur at a faster rate when individuals’ endowment shocks are more persistent. Motivated by these findings, we discuss how our analysis speaks to “closed economy” settings in which aggregate consumption is required to be constant, as in [Atkeson and Lucas \(1992\)](#).

Second, persistence gives rise to novel *short-run* dynamics and distortions. For instance, the optimal contract (i) uses history-dependent rewards and punishments that condition on the *order* of the agent’s shocks, in contrast to an order-invariance property from the i.i.d. case ([Thomas and Worrall, 1990](#)), and (ii) delivers these incentives using a fundamentally different mix of *intra*temporal and *inter*temporal distortions. Furthermore, the relative importance of these new features increases as the agent’s type becomes more persistent; by several measures, higher persistence results in worse insurance.

Roadmap. [Section 2](#) develops the model. [Section 3](#) explains the recursive formulation. [Section 4](#) presents our main long-run results. [Section 5](#) presents our solved examples. [Section 6](#) concludes with a discussion of modeling assumptions and various extensions. Appendices A–B present facts about the recursive formulation. Supplemental Appendices C–G contain proofs of our main results. Proofs of auxiliary technical results are in [Bloedel, Krishna, and Leukhina \(2024a\)](#).

1.1. Related Literature

Insurance and Immiseration. Substantively, our work contributes to the literature on long-run properties of optimal insurance contracts initiated by [Green \(1987\)](#), [Thomas and Worrall \(1990\)](#), and [Atkeson and Lucas \(1992\)](#). We build on the connection between immiseration and the Martingale Convergence Theorem identified in [Thomas and Worrall \(1990\)](#), extending it to settings with general forms of persistent private information.

Most directly, our work relates to several papers that study optimal insurance and immiseration under specific forms of persistent private information (and parametric assumptions on preferences).⁸ [Zhang \(2009\)](#) studies the special case of our framework with two types and symmetric transition probabilities; using arguments that rely on this special structure, he shows that the optimal contract induces immiseration.⁹ [Williams \(2011\)](#) studies a model in which the agent’s type follows a Brownian motion (i.e., a continuous-time Gaussian random walk)

⁸[Farinha Luz \(2022\)](#) studies optimal insurance when the agent has a persistent private “risk type,” but his realized income is contractible (as in [Rothschild and Stiglitz \(1976\)](#), [Wilson \(1977\)](#)). The latter feature yields qualitatively different optimal contracts than those in the social contracting literature on which we build.

⁹Formally, [Zhang \(2009\)](#) studies a model with private productivity shocks (an extension we discuss in [Section 6.1](#)) and works in continuous time, but these differences are inessential for the present comparison.

and shows that the optimal contract yields long-run bliss, rather than immiseration.¹⁰ Bloedel, Krishna, and Strulovici (2024b) consider a more general class of Markovian type processes with “permanent shocks” and show that it is effectively impossible to elicit information from the agent, which helps explain why optimal contracts under permanent shocks (as in Williams (2011)) differ from those in settings with mean-reversion. In contrast to these works, we study a general class of Markovian type processes with mean-reversion (and make minimal assumptions on preferences).

Less directly, our work also relates to the literature that studies the robustness of immiseration under alternative assumptions on preferences, technologies, and institutions (while maintaining the i.i.d. assumption). Several papers show that bounded long-run inequality is optimal in settings with limited commitment, such as when the agent faces interim participation constraints (Atkeson and Lucas, 1995, Phelan, 1995) or the principal faces credibility constraints (Sleet and Yeltekin, 2006, 2008, Farhi, Sleet, Werning, and Yeltekin, 2012), and in settings with full commitment but alternative normative criteria that place Pareto weight directly on “future generations” of agents (Phelan, 2006, Farhi and Werning, 2007). Other papers derive weaker forms of immiseration in economies with production (Khan and Ravikumar (2001), Khan, Popov, and Ravikumar (2020)) or endogenous fertility (Hosseini, Jones, and Shourideh (2013)), and under alternative assumptions on agents’ risk preferences (Phelan (1998); Olszewski and Safronov (2021)). In each case, long-run outcomes can be viewed as deriving from a modified version of the principal’s marginal cost martingale.

Recursive Contracts. Methodologically, our work contributes to the literature on recursive methods for dynamic contracting initiated by Green (1987), Thomas and Worrall (1990), Spear and Srivastava (1987). Fernandes and Phelan (2000) introduce the first recursive formulation for settings with persistence. Their main state variable comprises the agent’s “ex ante” promised utility, which encodes his continuation payoff conditional on having reported truthfully in the previous period, and a vector of *off-path* “threat-point” utilities that encode his continuation payoffs conditional on having misreported in the previous period. Building on Cole and Kocherlakota (2001a), Doepke and Townsend (2006), our main state variable instead comprises a vector of *on-path* “interim” promised utilities that encode the agent’s continuation payoff conditional on his current type, which makes the agent’s incentive constraints simpler to interpret and analyze.¹¹ For instance, we are able to fully characterize the “recursive domain” of implementable promised utility vectors for any number of types (Theorem 3). See Appendix B.1 for further discussion.

Our main methodological contribution is not the recursive formulation itself, but our use of it to characterize optimal contracts in settings with persistence.¹² Several recent papers share this goal, but in different economic environments and under the restriction to binary-type settings. Halac and Yared (2014) use the Fernandes and Phelan (2000) formulation to study optimal discretion in fiscal policymaking. Concurrent to our working paper, Guo and Hörner (2020), Fu and Krishna (2019) apply the same recursive formulation used here to study firm financing and allocation problems, respectively, with a risk-neutral agent and limited transfers. Fu and

¹⁰Williams (2011) also permits the agent’s type to follow a mean-reverting Gaussian AR(1) process, but does not characterize the optimal contract in that case (see Bloedel, Krishna, and Strulovici (2023)).

¹¹Cole and Kocherlakota (2001a) develop their approach in the context of stochastic games with hidden states, building on Abreu, Pearce, and Stacchetti (1990); Athey and Bagwell (2008) apply their methods to study dynamic Bertrand competition. Doepke and Townsend (2006) develop their approach in the context of contracting problems with both hidden states and hidden actions, emphasizing settings with both frictions.

¹²Fernandes and Phelan (2000), Doepke and Townsend (2006) numerically solve binary-type examples (as in Section 5), but do not analytically study optimal contracts (see also Broer, Kapička, and Klein (2017)).

Krishna (2019) establish convergence to long-run efficiency using a special case of the martingale methods that we develop here to establish immiseration, while Guo and Hörner (2020) establish long-run “polarization” (cf. Phelan (1998)) through a detailed construction of the optimal contract. Krasikov and Lamba (2021) use methods similar to Fu and Krishna (2019) to study dynamic procurement problems with limited liability. As our martingale methods for characterizing long-run outcomes (Theorems 1 and 2) and construction of the recursive domain (Theorem 3) apply for a broad class of Markovian type processes, they may prove useful for extending the analysis of optimal contracts in these and other applications beyond the binary-type case.

2. MODEL

Time is discrete and runs over an infinite horizon. At $t = 0$, a principal (she), who is risk-neutral and has discount factor $\alpha \in (0, 1)$, offers an insurance contract to a risk-averse agent (he). Neither party can renege on the contract at a later date.

Agent’s Preferences. The agent has the same discount factor $\alpha \in (0, 1)$. His Bernoulli utility of consumption is $U : (\underline{c}, \infty) \rightarrow \mathbb{R}$, where $\underline{c} \geq -\infty$ is a minimal consumption level.

ASSUMPTION 1—DARA: $U(\cdot)$ satisfies the following properties:

- (a) It is strictly increasing, strictly concave, continuously differentiable, and satisfies the Inada conditions $\lim_{c \rightarrow \underline{c}} U'(c) = +\infty$ and $\lim_{c \rightarrow \infty} U'(c) = 0$.
- (b) It is bounded above and unbounded below: $\lim_{c \rightarrow \infty} U(c) = 0$ and $\lim_{c \rightarrow \underline{c}} U(c) = -\infty$.
- (c) It has decreasing absolute risk aversion: the map $c \mapsto -\log(U'(c))$ is (weakly) concave.

Assumption DARA is common in the literature and includes standard parametric classes of utilities, such as those in the CARA class and many in the HARA class.¹³ We denote the range of U by $\mathcal{U} := U((\underline{c}, \infty)) \subset \mathbb{R}$. Note that part (b) implies $\mathcal{U} = (-\infty, 0)$.

Type Process. In each period t , the agent privately observes his endowment $\omega^{(t)} \in \mathbb{R}$, which can take any of the d values $\omega_d > \omega_{d-1} > \dots > \omega_1$. If $\omega^{(t)} = \omega_i$, we say that the agent’s time- t type is $i \in S := \{1, \dots, d\}$. The endowment (or type) process $(\omega^{(t)})_{t=0}^\infty$ is Markovian and the measure over sample paths is given by $\mathbf{P} \in \Delta(S^\infty)$.¹⁴

ASSUMPTION 2—Markov: The agent’s type follows a fully connected, time-homogeneous, first-order Markov chain with transition probabilities $f_{ij} := \mathbf{P}(\omega^{(t+1)} = \omega_j \mid \omega^{(t)} = \omega_i) > 0$.

Assumption Markov implies that the type process is ergodic and bounded, but otherwise allows for essentially arbitrary serial correlation. Going forward, we represent the transition probabilities as the $d \times d$ transition matrix with rows $\mathbf{f}_i = (f_{i1}, \dots, f_{id})$, where \mathbf{f}_i denotes the conditional distribution of tomorrow’s type if today’s type is $i \in S$.¹⁵

¹³We adopt DARA following Cole and Kocherlakota (2001b), Thomas and Worrall (1990). CARA utility is the benchmark specification in the literature: see the models or main examples in Green (1987), Phelan (1995, 1998), Williams (2011), Bloedel, Krishna, and Strulovici (2023), Thomas and Worrall (1990), Wang (1995), Strulovici (2020), Atkeson and Lucas (1992). DARA is satisfied by all HARA utilities $U(c) = (c + \eta)^{1-\gamma}/(1-\gamma)$ with $\gamma > 1$ and $\eta \geq -\underline{c}$ (where CRRA utility corresponds to $\eta = \underline{c} = 0$).

¹⁴Throughout, we identify endowment realizations with their corresponding type realizations, and identify the space $\{\omega_1, \dots, \omega_d\}^\infty$ of sample paths of endowments with the space S^∞ of sample paths of types.

¹⁵We also implicitly assume that the initial distribution $\mathbf{P}(\omega^{(0)} = \cdot) = \mathbf{f}_i$ for some $i \in S$, which serves to simplify notation in the definition of the principal’s problem below (see (RP) in Section 3.3 and (SP) in Appendix A.1). This assumption is not essential and can easily be relaxed.

Reporting Strategies. In each period, the agent submits a report about his current type. We assume, as is common in the literature, that the agent cannot over-report his type:

ASSUMPTION 3—NHB: *An agent of type i can only report to be of types $j \leq i$.*

Assumption NHB (“No Hidden Borrowing”) captures two standard ideas: (i) endowments are partially verifiable, and (ii) the agent cannot covertly save, borrow, or produce the consumption good (e.g., via a market or private technology). For example, the principal might require the agent to deposit his reported endowment in a jointly monitored account; if he cannot covertly engage in such activities, he can deposit at most his true endowment.

Principal’s Sequential Problem. A contract offered by the principal at $t = 0$ specifies transfers (of the consumption good) to the agent in each period as a function of the history of type reports. The principal aims to minimize her expected lifetime cost of financing the contract, subject to two constraints: (i) *promise keeping* and (ii) *incentive compatibility*. Formally, the contract must (i) deliver a given schedule of *promised utilities* $\mathbf{v}^{(0)} := (v_1, \dots, v_d) \in \mathcal{U}^d$ to the agent, where v_i denotes the agent’s lifetime utility if his initial endowment is $\omega^{(0)} = \omega_i$ and he reports truthfully, and (ii) ensure that truthful reporting in all periods maximizes the agent’s lifetime utility.¹⁶

We call this the principal’s *sequential problem*, as it describes her optimization over full sequences of history-dependent transfers. While this “sequential formulation” is standard, our main analysis uses a more tractable recursive formulation of contracts (see Section 3 below). Thus, for brevity, we defer the details of the sequential formulation to Appendix A.1, where the principal’s sequential problem is stated as program (SP).

REMARK 2.1: Assumptions 1–3 hold throughout the paper. We discuss them further in Section 6.1. There, we also discuss model variants in which the agent’s type represents his tastes or productivity, rather than his endowment.

3. RECURSIVE CONTRACTS

Section 3.1 develops the recursive formulation. Section 3.2 introduces regularity conditions used in our analysis. Section 3.3 presents the principal’s Bellman equation.

3.1. Principal’s Recursive Problem

State Variable. Recall from Green (1987), Thomas and Worrall (1990) that, when types are i.i.d., the principal’s problem can be written recursively with the agent’s *ex ante promised utility* as a state variable (i.e., his expected continuation utility under the contract *before* his type in the current period is realized, and *assuming* truthful reporting in the current and future periods). The fact that this one-dimensional state can encode all payoff-relevant aspects of the history relies on the *symmetry* of information *across* periods implied by the i.i.d. assumption. Intuitively, ex ante promised utility functions as a “numeraire” because the agent’s preferences over “continuation contracts” for future periods are common knowledge, even at *off-path* histories where the agent has lied.¹⁷

¹⁶This initial $\mathbf{v}^{(0)}$ could be given exogenously or arise from the principal’s cost-minimization given some prior belief over the agent’s initial type (e.g., the “efficiency problem” (Eff _{i}) in Section 4.3 below).

¹⁷Formally, at each history (i) the agent’s beliefs about his future types do *not* depend on his current type, and (ii) the principal’s beliefs about the agent’s current type do *not* depend on his previous reports.

Plainly, this logic fails in settings with persistence. Thus, as noted by [Fernandes and Phelan \(2000\)](#), [Doepke and Townsend \(2006\)](#), we require a richer state variable. We use a state variable (\mathbf{v}, s) with two components. First, $\mathbf{v} = (v_1, \dots, v_d) \in \mathcal{U}^d$ denotes a vector of *interim promised utilities*, where $v_i \in \mathcal{U}$ denotes the agent's continuation utility *contingent on* his current type being $i \in S$ and *assuming* truthful reporting in the future. As we will see below, \mathbf{v} encodes all aspects of the history that are payoff-relevant to the *agent* (even off-path). Second, $s \in S$ denotes the agent's *report* in the *previous* period. This is payoff-relevant to the *principal*, as it determines her beliefs about the agent's current type. Thus, the overall state (\mathbf{v}, s) encodes all payoff-relevant aspects of the history.¹⁸

Recursive Constraints. At each state (\mathbf{v}, s) , we view the contract as offering the agent report-contingent consumption transfers and interim continuation utility vectors. Formally, if the agent *reports* to be of type j , two things happen. First, he receives a consumption transfer c_j , so that his overall consumption is $c_j + \omega_i$, where i is his *true* type. Second, he is promised the continuation utilities $\mathbf{w}_j := (w_{j1}, \dots, w_{jd})$, where w_{jk} denotes his interim promised utility in the *next* period, contingent on k being his true type therein. Thus, the state transitions to (\mathbf{w}_j, j) in the next period, and the process begins again.

It will be convenient to express the current period's allocation in terms of the agent's *flow utility*, rather than his consumption. To this end, let $C(u_j, j) := U^{-1}(u_j) - \omega_j$ denote the consumption transfer that delivers exactly u_j utilities to a *truthful* agent of type j . Suppose the principal intends to give u_j utilities to a type j agent, and so transfers $c_j = C(u_j, j)$ units of consumption when the agent *reports* to be of type j . Then, if an agent of true type i reports to be type j , his flow utility is $\psi(u_j, i, j) := U(\omega_i + C(u_j, j))$.

Given this change of variables, at each state (\mathbf{v}, s) , we can view the contract as offering the agent a *menu* $(u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times \mathcal{U}^d)^d$ consisting of *flow utilities* $u_i \in \mathcal{U}$ and *interim continuation utilities* $\mathbf{w}_i \in \mathcal{U}^d$ for each type $i \in S$. Clearly, this menu should (i) deliver the correct promised utility to every type, and (ii) ensure that truthful reporting in the current period (assuming truthful reporting in future periods) is optimal for the agent. Formally, these *recursive constraints* consist of the (interim) *promise keeping* and *incentive compatibility* conditions

$$v_i = u_i + \alpha \mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_i] \quad (\text{PK}_i)$$

$$v_i \geq \psi(u_j, i, j) + \alpha \mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_j] \quad (\text{IC}_{ij})$$

for all $i, j \in S$ with $i > j$ (per Assumption [NHB](#)), where $\mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_j] := \sum_{k=1}^d f_{ik} w_{jk}$ is the *expected* promised utility for an agent of type i who reports to be type j . We call $\mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_i]$ the *ex ante* promised utility for type i .¹⁹

The recursive constraints have two important properties. First, they are independent of s , the agent's previous report. Second, even if the agent lies today, his expectation over tomorrow's type in (IC_{ij}) is still determined by his true current type. Thus, the recursive constraints ensure one-step promise keeping and deter one-shot deviations from truthtelling at *every* history, regardless of the agent's previous true and reported types.²⁰

¹⁸In general, a contract may condition on aspects of the history beyond (\mathbf{v}, s) . However, it is well understood that *without loss of optimality*, we can restrict attention to contracts that are Markovian (or "recursive") with respect to (\mathbf{v}, s) . See, e.g., [Fernandes and Phelan \(2000, Lemma 2.3\)](#) and [Doepke and Townsend \(2006, Proposition 4\)](#). As is standard in the literature, we implicitly restrict attention to such contracts throughout.

¹⁹When types are i.i.d., so that $\mathbf{f}_i = \mathbf{f}_j$ for all $i, j \in S$, the quantity $w_i := \mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_i]$ is precisely the state variable from [Green \(1987\)](#), [Thomas and Worrall \(1990\)](#).

²⁰In principle, requiring that truthful reporting is optimal for the agent at *every* history might be stronger than requiring that it is optimal for him *conditional on truthful reporting in the past* (cf. incentive compatibility in the




FIGURE 1.—The domain $D = V_2$ for $d = 2$ and positive correlation ($f_{11} \geq f_{21}, f_{22} \geq f_{12}$).

Recursive Domain. Even if a menu $(u_i, \mathbf{w}_i)_{i \in S}$ satisfies the recursive constraints (\mathbf{PK}_i) – (\mathbf{IC}_{ij}) in state (\mathbf{v}, s) , there may not exist menus satisfying these constraints in all of the possible continuation states (\mathbf{w}_i, i) . No incentive compatible contract could ever lead to states at which such “dead ends” arise—there must always exist a well-defined continuation contract. Thus, an essential part of the recursive formulation is determining which promised utilities $\mathbf{v} \in \mathcal{U}^d$ are *implementable* via some underlying incentive compatible contract.

Formally, we call a set $D \subseteq \mathcal{U}^d$ *the (recursive) domain* if (i) for every $\mathbf{v} \in D$, there exists a menu $(u_i, \mathbf{w}_i)_{i \in S}$ satisfying the recursive constraints (\mathbf{PK}_i) – (\mathbf{IC}_{ij}) and $\mathbf{w}_i \in D$ for all $i \in S$, and (ii) any set $D' \subseteq \mathcal{U}^d$ with the preceding property satisfies $D' \subseteq D$.²¹ The domain characterizes the implementable promised utilities: all (and only those) $\mathbf{v} \in D$ are viable. [Theorem 3](#) in [Appendix B](#) shows that D exists and is a nonempty, open, convex cone. For type processes satisfying standard forms of positive serial correlation, [Theorem 3](#) also yields the closed-form solution $D = V_d := \{\mathbf{v} \in \mathcal{U}^d : v_d > v_{d-1} > \dots > v_1\}$.

To illustrate the logic of this solution, consider the $d = 2$ case with positive correlation depicted in [Figure 1](#). Substituting the promise keeping constraint (\mathbf{PK}_i) (for $i = 1$) into the incentive constraint (\mathbf{IC}_{ij}) (for $i = 2$ and $j = 1$) yields an equivalent version of the latter:

$$v_2 - v_1 \geq \underbrace{\psi(u_1, 2, 1) - u_1}_{\text{i.i.d. info rent}} + \alpha \underbrace{[\mathbf{E}^{f_2}(\mathbf{w}_1) - \mathbf{E}^{f_1}(\mathbf{w}_1)]}_{\text{Markov info rent}} \quad (\mathbf{IC}_{21}^*)$$

The *i.i.d. information rent* is clearly positive. If $\mathbf{w}_1 \in V_2$, then the *Markov information rent* is also positive, as $j \mapsto w_{1j}$ is increasing and f_2 first-order stochastically dominates f_1 . Thus, implementability of \mathbf{w}_1 implies that $\mathbf{v} \in V_2$. Conversely, it can be shown that each $\mathbf{v} \in V_2$ can be induced by some menu $(u_i, \mathbf{w}_i)_{i=1,2}$ with $\mathbf{w}_1, \mathbf{w}_2 \in V_2$, which in turn implies that both \mathbf{w}_1 and \mathbf{w}_2 can be induced by menus with this property, and so on *ad infinitum*. In sum, all (and only those) $\mathbf{v} \in V_2 = \{\mathbf{v}' \in \mathcal{U}^d : v'_2 > v'_1\}$ are implementable in this example.

Recursive Contracts. With the domain D in hand, we can define contracts. A *(recursive) contract* is a map $\xi : D \times S \rightarrow (\mathcal{U} \times D)^d$ that assigns to each state $(\mathbf{v}, s) \in D \times S$ a menu

sequential problem, (\mathbf{SP})). However, since the type process is Markovian and payoffs are time-separable, these two requirements are, in fact, equivalent (e.g., [Pavan, Segal, and Toikka \(2014, Section 3.3\)](#)).

²¹In the terminology of [Abreu, Pearce, and Stacchetti \(1990\)](#), the domain D is (i) *self-generating* and (ii) the *largest* self-generating set. Note that D is clearly unique (if it exists).

$(u_i(\mathbf{v}, s), \mathbf{w}_i(\mathbf{v}, s))_{i \in S}$ of report-contingent flow utilities $u_i(\mathbf{v}, s) \in \mathcal{U}$ and continuation utilities $\mathbf{w}_i(\mathbf{v}, s) \in D$. We are interested only in those ξ that satisfy the recursive constraints at each step. To this end, define the *constraint correspondence* $\Gamma : D \rightrightarrows (\mathcal{U} \times D)^d$ as²²

$$\Gamma(\mathbf{v}) := \{(u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times D)^d : (\mathbf{PK}_i) \text{ and } (\mathbf{IC}_{ij}) \text{ hold } \forall i > j \in S \text{ at } \mathbf{v} \in D\}. \quad (3.1)$$

We say that a contract ξ is *feasible* if $\xi(\mathbf{v}, s) \in \Gamma(\mathbf{v})$ for all $(\mathbf{v}, s) \in D \times S$. Let Ξ denote the set of feasible contracts.

We can describe a contract's dynamics as follows. Each reporting strategy for the agent generates an S -valued stochastic process $(s^{(t+1)})_{t=0}^\infty$, where $s^{(t+1)}$ denotes his *current* report in period t (and thus his *previous* report from the perspective of period $t+1$). Let $\mathcal{H} := S^\infty$ denote the set of *paths* of type reports. Given any $\xi \in \Xi$, initial $\mathbf{v} \in D$, and distribution over \mathcal{H} (determined by the agent's reporting strategy), iterating on ξ in the natural way yields two stochastic processes: (i) the \mathcal{U} -valued process $\tilde{u}_\xi := (u_\xi^{(t)})_{t=0}^\infty$ of *induced allocations*, and (ii) the D -valued process $(\mathbf{v}_\xi^{(t)})_{t=0}^\infty$ of *induced promises* (where $\mathbf{v}_\xi^{(0)} := \mathbf{v}$).²³ In this notation, the recursive state variable at time t is $(\mathbf{v}_\xi^{(t)}, s^{(t)})$.

Optimality. Given an initial state $(\mathbf{v}, s) \in D \times S$, the principal minimizes her lifetime cost across all feasible recursive contracts $\xi \in \Xi$. Formally, her *recursive problem* is

$$P(\mathbf{v}, s) := \inf_{\xi \in \Xi} \mathbf{E} \left[\sum_{t=0}^{\infty} \alpha^t C(u_\xi^{(t)}, s^{(t+1)}) \mid (\mathbf{v}_\xi^{(0)}, s^{(0)}) = (\mathbf{v}, s) \right], \quad (\text{RP})$$

where the expectation is taken with respect to the distribution over paths induced by truthful reporting and the prior belief \mathbf{f}_s over initial types $\omega^{(0)} = \omega_{s(1)}$.²⁴ As this notation suggests, under truthful reporting, $s^{(t+1)}$ is the agent's true type in period t (i.e., $\omega^{(t)} = \omega_{s^{(t+1)}}$). For this reason, going forward, we will often refer to $(s^{(t+1)})_{t=0}^\infty$ itself as the *type process*.

A contract ξ^* is *optimal* if it attains the infimum in (RP) at every $(\mathbf{v}, s) \in D \times S$. We represent the infimal cost in (RP) by the principal's *value function* $P : D \times S \rightarrow \overline{\mathbb{R}}$.

Transversality. The recursive constraints only impose one-step promise keeping and deter one-step deviations. For a solution ξ^* in (RP) to be feasible in the sequential problem (SP), the induced *allocation* \tilde{u}_{ξ^*} must actually (i) deliver the appropriate level of promised utility and (ii) deter infinite-length deviations from truthful reporting.²⁵ Intuitively, any contract violating (i) or (ii) corresponds to a Ponzi scheme in which the principal drives the induced *promises* to minus infinity too quickly. We wish to rule out such schemes.

To this end, we say that a contract $\xi \in \Xi$ is (TVC)-*implementable* at $\mathbf{v} \in D$ if, starting from the initial condition $\mathbf{v}_\xi^{(0)} := \mathbf{v}$, it satisfies

$$\lim_{t \rightarrow \infty} \inf_{h \in \mathcal{H}} \alpha^t \mathbf{v}_\xi^{(t)}(h) = \mathbf{0} \quad (\text{TVC})$$

²²Note that $\Gamma(\cdot)$ depends only on the \mathbf{v} component of the state (\mathbf{v}, s) because, as discussed above, the recursive constraints are independent of s .

²³Implicit in this description are the usual restrictions that (a) the agent's report process is adapted to his type process, and (b) the induced allocation and promised utility processes are adapted to the report process. See [Appendix A](#) for these details in the context of the sequential formulation of the contracting problem.

²⁴This distribution coincides with the measure \mathbf{P} over paths of types defined in [Section 2](#), except for one difference: the prior belief over the initial $\omega^{(0)}$ may differ. We formulate (RP) at all conceivable initial states, following the standard convention in dynamic programming.

²⁵That is, (i) $v_j = \mathbf{E}[\sum_{t=0}^{\infty} \alpha^t u_{\xi^*}^{(t)} \mid s^{(1)} = j]$ for all $j \in S$, and (ii) the agent's reporting problem is suitably "continuous at infinity." Both (i)–(ii) would trivially hold if the agent's utility function were bounded.

where $(\mathbf{v}_\xi^{(t)}(h))_{t=0}^\infty$ is the sequence of promises along path $h \in \mathcal{H}$. All (TVC)-implementable contracts satisfy properties (i) and (ii) stated above (Lemma A.1 in Appendix A.2).

Full-Information Benchmark. For future reference, we briefly describe the *first-best* contract that is optimal when there is no private information, i.e., there are no incentive constraints (see Supplemental Appendix F for details). In this setting, every promised utility vector $\mathbf{v} \in \mathcal{U}^d$ is implementable, and a *feasible full-information contract* is a map $\zeta : \mathcal{U}^d \times S \rightarrow (\mathcal{U} \times \mathcal{U}^d)^d$ such that $\zeta(\mathbf{v}, s)$ satisfies (PK_{*i*}) for all $i \in S$ at every $(\mathbf{v}, s) \in \mathcal{U}^d \times S$. We let Ξ^{FB} denote the set of all such contracts. The first-best contract ζ^* perfectly smooths the agent's consumption conditional on the initial type $s^{(1)}$: given any initial condition $(\mathbf{v}_{\zeta^*}^{(0)}, s^{(0)}) \in \mathcal{U}^d \times S$ and realization of $s^{(1)} \in S$, (i) the induced allocations $(u_{\zeta^*}^{(t)})_{t=0}^\infty$ are constant, and (ii) the induced promises $(\mathbf{v}_{\zeta^*}^{(t)})_{t=1}^\infty$ are constant and satisfy $v_{\zeta^*,1}^{(t)} = \dots = v_{\zeta^*,d}^{(t)}$ for all $t \geq 1$. The principal's *first-best value function* is denoted by $Q^* : \mathcal{U}^d \times S \rightarrow \mathbb{R}$.²⁶ Naturally, the first-best contract is not incentive compatible: in every period $t \geq 1$, the agent's consumption is maximized by reporting to be of type $i = 1$ (cf. Lemma C.16 in Supplemental Appendix C.3).

3.2. Regularity Conditions

To ensure that problem (RP) is mathematically well-behaved, we impose some regularity conditions on the environment. For expositional simplicity, we impose them directly on derived objects; sufficient conditions on model primitives can be obtained case-by-case.

DEFINITION 3.1: The environment is *Regular* if Conditions R.1–R.3 below hold, and is (TVC)-*Regular* if Conditions R.4–R.5 below also hold.

R.1 (Finite Value) The value function P for (RP) is well-defined and finite-valued on $D \times S$.

R.2 (FB Regularity) The first-best value function Q^* satisfies, for all $\zeta \in \Xi^{\text{FB}}$ and $\mathbf{v}_\zeta^{(0)} \in \mathcal{U}^d$,

$$\liminf_{t \rightarrow \infty} \alpha^t \left[\inf_{h \in \mathcal{H}} Q^*(\mathbf{v}_\zeta^{(t)}(h), s^{(t)}(h)) \right] \geq 0.$$

R.3 (Constraint Qualification) For all $\mathbf{v} \in D$, we have $\Gamma_\circ(\mathbf{v}) \neq \emptyset$, where $\Gamma_\circ(\mathbf{v}) \subseteq \Gamma(\mathbf{v})$ is the subset of feasible menus at \mathbf{v} that satisfy all of the incentive compatibility constraints (IC_{*ij*}) (for $i > j \in S$) as *strict* inequalities.

R.4 (TVC) Existence) There exists an optimal contract ξ^* that is (TVC)-implementable at every $\mathbf{v} \in D$.

R.5 (Smooth Value) The value function $P(\cdot, s)$ for (RP) is differentiable on D for each $s \in S$.

Regularity (R.1–R.3) is a mild technical assumption that facilitates analysis of (RP). Conditions R.1 and R.2 ensure that the principal's value function is well-defined and that an optimal contract exists. Condition R.3 ensures the existence of Lagrange multipliers, allowing us to use Lagrangian methods. One can obtain sufficient conditions for Regularity in terms of model primitives. For instance, Regularity holds for CARA utility and type processes with positive serial correlation (Lemma B.2 in Appendix B.2).

The assumption of (TVC)-Regularity comprises two additional technical conditions (R.4–R.5) that can be verified in special cases but are harder to check in general. First, Condition R.4

²⁶Formally, this characterization of the first-best holds under Regularity Condition R.2 (stated below in Definition 3.1), which ensures that the full-information problem (FB) in Supplemental Appendix F is well-behaved.

ensures that *some* optimal contract in the recursive problem (RP) is feasible in the sequential problem (SP). Were this conclusion to fail, (RP) would be a strict relaxation of (SP); we view this possibility as pathological. Second, Condition R.5 ensures that $P(\cdot, s)$ is continuously differentiable *everywhere*, rather than merely *almost everywhere* (see Proposition 3.2 below).²⁷ In Supplemental Appendix H, we outline an approach to verifying Condition R.5 based on the envelope theorem of Rincón-Zapatero and Santos (2009). Both of Conditions R.4 and R.5 (or suitable versions thereof) can be verified under CARA utility in the two limiting cases of i.i.d. and random walk type processes.²⁸

3.3. Bellman Equation

To conclude the recursive formulation, we note that the principal's value function P satisfies a familiar Bellman equation. This equation is useful for generating optimal contracts and sketching the proofs of our main results in Section 4 below.

PROPOSITION 3.2: *Suppose the environment is Regular. Then the principal's value function $P : D \times S \rightarrow \mathbb{R}$ satisfies the functional equation*

$$P(\mathbf{v}, s) = \min_{(u_i, \mathbf{w}_i)_{i \in S} \in \Gamma(\mathbf{v})} \sum_{i \in S} f_{si} [C(u_i, i) + \alpha P(\mathbf{w}_i, i)], \quad (\text{FE})$$

is the pointwise smallest solution to (FE) that is pointwise greater than Q^* , and is convex in its first argument. Moreover:

- (a) There exists an optimal contract ξ^* , and any contract ξ generated by P via a policy function from (FE) is optimal.
- (b) If the environment is (TVC)-Regular, then $P(\cdot, s)$ is continuously differentiable and strictly convex for each $s \in S$, and there exists a unique optimal contract ξ^* , which is continuous on $D \times S$.

The routine but lengthy proof of Proposition 3.2 is in Section J of Bloedel, Krishna, and Leukhina (2024a), where we also establish some further properties of P and ξ^* .²⁹

4. LONG-RUN OUTCOMES

Sections 4.1 and 4.2 present our main results on long-run properties of optimal contracts. Section 4.3 provides a sketch of the proof.

4.1. Immiseration

Our first main result establishes that the optimal contract always induces *immiseration*: the agent's utility and consumption tend to their lower bounds, so he becomes impoverished in the long run. To state it, let $c^{(t)} := C(u^{(t)}, s^{(t+1)})$ denote the process of consumption transfers under the optimal contract. (For convenience, we henceforth omit ξ^* subscripts on all processes induced by the optimal contract.)

²⁷Under Regularity, Proposition 3.2 shows that $P(\cdot, s)$ is convex, which implies that it is (i) differentiable almost everywhere and (ii) continuously differentiable at every point of differentiability.

²⁸In the i.i.d. case, see Atkeson and Lucas (1992), Green (1987), Thomas and Worrall (1990) for Condition R.4 and Supplemental Appendix H for Condition R.5. In the permanent shock case, see Bloedel, Krishna, and Strulovici (2023, 2024b) for suitable versions of Conditions R.4 and R.5.

²⁹The value function P is unbounded and diverges near the boundaries of the domain D , so we cannot use standard contraction-mapping results. Instead, we appeal to first-principles and order-theoretic arguments.

THEOREM 1—Immiseration: *Suppose the environment is (TVC)-Regular. Under the optimal contract, as $t \rightarrow \infty$:*

- (a) *Promised utilities decrease without bound: $v_i^{(t)} \rightarrow -\infty$ in probability for all $i \in S$.*
- (b) *Flow utilities decrease without bound: $u^{(t)} \rightarrow -\infty$ in probability.*
- (c) *Consumption converges to its lower bound: $c^{(t)} + \omega^{(t)} \rightarrow \underline{c}$ in probability.*

Consequently, none of the promised utility, flow utility, or consumption processes have limiting distributions, and the state process $(\mathbf{v}^{(t)}, s^{(t)})$ does not have a stationary distribution.

The proof of [Theorem 1](#) is in Supplemental Appendix C. We sketch the argument in [Section 4.3](#) and discuss the importance of Assumption [Markov](#) for this result in [Section 6.2](#). As a corollary of the proof of [Theorem 1](#), we also obtain a stronger mode of convergence in the two special cases considered in prior work.³⁰

COROLLARY 4.1: *The convergence in [Theorem 1](#) can be strengthened from “in probability” to “almost surely” if types are either (i) i.i.d. ($\mathbf{f}_i = \mathbf{f}_j$ for all $i, j \in S$) or (ii) binary ($d = 2$).*

The proof of [Corollary 4.1](#) is in Supplemental Appendix D. The i.i.d. case essentially restates the main result of [Thomas and Worrall \(1990\)](#). The binary-state case can be viewed as a generalization of the immiseration result in [Zhang \(2009\)](#), which further assumes positive correlation and symmetric transition probabilities (i.e., $f_{11} = f_{22} \geq f_{12} = f_{21}$).

The basic intuition for [Theorem 1](#) is that the agent’s risk aversion implies that providing incentives is cheaper for the principal when the level of the agent’s (promised) utility is lower ([Thomas and Worrall, 1990](#)). In particular, while effective insurance provision requires the principal to give larger transfers to the agent when his reported endowment is lower, incentive compatibility requires these larger transfers in the current period to be accompanied by lower continuation utility (i.e., lower transfers, on average, in future periods). Thus, the agent’s flow and continuation utilities must vary with his reported endowment, with larger variation corresponding to higher-powered incentives. All else equal, the cost of incentive provision is lower for the principal when the level of the agent’s utility is lower, as this is when the agent’s *marginal* utility of consumption is higher. For a simple illustration, suppose that $d = 2$ and the principal wants to induce a spread of $\varepsilon > 0$ between the two types’ flow utilities (i.e., $u_2 - u_1 = \varepsilon$). If the previous report was $s \in S$ and $\varepsilon > 0$ is small, this costs the principal approximately

$$\underbrace{f_{s1}C(u_1, 1) + f_{s2}C(u_1, 2)}_{\text{level}} + \varepsilon \cdot \underbrace{f_{s2}C'(u_1, 2)}_{\text{variability}}$$

Thus, because $C(\cdot, 2)$ is convex (due to risk aversion), the cost of inducing utility variability increases with the utility level.³¹ In the long run, it is optimal for the principal to drive this cost of incentive provision to zero by immiserating the agent. The next two sections explain why this is so: [Section 4.2](#) provides economic intuition in terms of the power of the optimal contract’s incentives, and [Section 4.3](#) sketches the formal proof of [Theorem 1](#).

4.2. Backloaded Incentives and Relative Immiseration

Intuitively, immiserating the agent allows the principal to optimally smooth her costs over time. As the level of (promised) utility decreases, it becomes affordable for the principal to

³⁰We conjecture that, in general, the convergence in parts (a)–(c) of [Theorem 1](#) occurs almost surely.

³¹See [Banerjee and Newman \(1991\)](#), [Newman \(2007\)](#) for a clear articulation of this relation between wealth (or utility) levels and risk-bearing in a related class of principal-agent problems with production.

provide *high-powered incentives*, i.e., make the sensitivity of the agent's continuation utility with respect to his report increase without bound. By *backloading* high-powered but cheap incentives in later periods, the principal can make utility and consumption less variable in earlier periods while maintaining incentive compatibility, thereby reducing the cost of incentive provision in those early periods. This pattern of incentive provision allows her to intertemporally smooth, and thus minimize, her costs.

Our second main result formalizes this idea. To state it, we require two definitions. First, note that at the beginning of period t , when the state $(\mathbf{v}^{(t)}, s^{(t)})$ is given but before the current type $s^{(t+1)}$ is realized, the agent's type-contingent continuation utility $v_{s^{(t+1)}}^{(t)}$ is a random variable with support $\{v_1^{(t)}, \dots, v_d^{(t)}\}$. We denote by

$$\mathbf{V}(v_{s^{(t+1)}}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)}) := \sum_{i=1}^d f_{s^{(t)}, i} \left(v_i^{(t)} - \sum_{i=1}^d f_{s^{(t)}, i} v_i^{(t)} \right)^2 \quad (4.1)$$

the variance of this random variable, conditional on information available at the beginning of period t . Second, we focus on type processes that exhibit positive serial correlation:

DEFINITION 4.2: The type process is *FOSD-ordered* (or simply *FOSD*) if $\mathbf{f}_i \in \Delta(S)$ first-order stochastically dominates $\mathbf{f}_j \in \Delta(S)$ whenever $i > j$.

FOSD is the standard notion of positive serial correlation in the literature. Every i.i.d. process is FOSD. When $d = 2$, FOSD is equivalent to $f_{22} \geq f_{12}, f_{11} \geq f_{21}$.

THEOREM 2—Backloaded Incentives: *Suppose the environment is (TVC)-Regular and the type process is FOSD. Under the optimal contract, as $t \rightarrow \infty$:*

- (a) $v_i^{(t)} - v_{i-1}^{(t)} \rightarrow +\infty$ in probability for all $i = 2, \dots, d$.
- (b) $\mathbf{V}(v_{s^{(t+1)}}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)}) \rightarrow +\infty$ in probability.

The proof of [Theorem 2](#) is in Supplemental Appendix [E](#). Again, we obtain a stronger mode of convergence in the two special cases from [Corollary 4.1](#).

COROLLARY 4.3: *The convergence in [Theorem 2](#) can be strengthened from “in probability” to “almost surely” if types are either (i) i.i.d. ($\mathbf{f}_i = \mathbf{f}_j$ for all $i, j \in S$) or (ii) binary ($d = 2$).*

While we have argued that [Theorem 2](#) in some sense explains why immiseration arises in the first place, we derive it as a simple consequence of [Theorem 1](#). We sketch the argument here. As in [Section 3.1](#), substituting the promise keeping constraint (PK_i) for type j into the incentive constraint (IC_{ij}) for type $i > j$ gives an alternative way to write the latter:

$$v_i - v_j \geq \underbrace{\psi(u_j, i, j) - u_j}_{\text{i.i.d. info rent}} + \alpha \left(\mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_j] - \mathbf{E}^{\mathbf{f}_j} [\mathbf{w}_j] \right). \quad (\text{IC}_{ij}^*)$$

For FOSD type processes, [Theorem 3](#) in [Appendix B](#) guarantees that the Markov information rent term is non-negative, so that both the i.i.d. and Markov information rents work in the same direction (recall [Section 3.1](#) for the $d = 2$ case). [Theorem 1](#)(b) implies that the i.i.d. information rent term, which is always non-negative, grows without bound. It then follows from (IC_{ij}^*) (with $j = i - 1$) that the difference $v_i - v_{i-1}$ must also explode.

Consequently, the agent’s continuation utility becomes increasingly uncertain over time: [Theorem 2\(b\)](#) shows that the conditional variance (4.1) diverges. That is, the agent’s uncertainty at the start of period t (before observing his period- t type) about his future prospects (his continuation utility after reporting his period- t type) increases without bound as $t \rightarrow \infty$. In this sense, the quality of risk sharing degrades in the long run.

Another interpretation of [Theorem 2](#) is that the optimal contract induces *relative immiseration*: the difference in promised utilities across different types of agents increases without bound, so that low-type agents become impoverished (in utility terms) *relative to* high-type agents (who also become impoverished). More concretely, imagine two agents, Alice and Bob, who have received the same sequence of realized endowments up through period $t - 1$. In period t , Alice receives a higher endowment than Bob. [Theorem 2\(a\)](#) states that this *transient* difference in Alice and Bob’s endowments translates to an arbitrarily large difference in *lifetime* continuation utility as $t \rightarrow \infty$. Viewed in this light, [Theorem 2](#) expresses the idea that “pathwise welfare inequality” increases without bound in the long run.³²

4.3. Proof Sketch for [Theorem 1](#)

The proof consists of several steps, which we sketch in sequence below. Building on [Thomas and Worrall \(1990\)](#), the basic idea is to identify a martingale that embodies the principal’s optimal cost-smoothing and to show that convergence of this martingale corresponds to the contract inducing immiseration. However, relative to the i.i.d. benchmark, persistence gives rise to both substantively new economic features of the optimal contract and technical complications that must be overcome in the proof.

Step 1: Marginal Cost Martingale. We first formalize the principal’s optimal cost-smoothing in terms of a martingale process (cf. an “Euler equation”). In the special case of i.i.d. types, [Thomas and Worrall \(1990\)](#) show that the principal’s marginal cost of increasing the agent’s *ex ante* promised utility is such a martingale. In particular, when types are i.i.d., we can write the principal’s value function as $\hat{P}(v)$, where $v = \mathbf{E}^f[\mathbf{v}]$ is the agent’s *ex ante* promised utility under the (type-independent) transition probability $\mathbf{f} \in \Delta(S)$. The derivative $\hat{P}'(v)$ represents the principal’s marginal cost of promised utility.

In our general Markovian setting, the appropriate martingale is a subtle but natural generalization of this object. Denote the derivative of P at (\mathbf{v}, s) by $DP(\mathbf{v}, s) = (P_1(\mathbf{v}, s), \dots, P_d(\mathbf{v}, s))$, where $P_i(\mathbf{v}, s)$ is the partial derivative with respect to the component v_i . Let $D_{\mathbf{1}}P(\mathbf{v}, s) := \lim_{\varepsilon \rightarrow 0} [P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)] / \varepsilon = \sum_{i \in S} P_i(\mathbf{v}, s)$ denote the directional derivative of $P(\cdot, s)$ in direction $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$.

PROPOSITION 4.4: *If the environment is (TVC)-Regular, then the process $(D_{\mathbf{1}}P(\mathbf{v}^{(t)}, s^{(t)}))_{t=0}^{\infty}$ induced by the optimal contract is a strictly positive martingale.*

The proof of [Proposition 4.4](#) is in Supplemental Appendix C.2. We henceforth refer to this process as the principal’s *marginal cost martingale*. [Proposition 4.4](#) has three parts:

³²This complements [Atkeson and Lucas’s \(1992\)](#) main finding that the *unconditional* (i.e., from the period 0 perspective) variance of period- t *ex ante* promised utility grows without bound as $t \rightarrow \infty$. In their model, which considers a society with a continuum of agents and a per-period aggregate resource constraint, this implies that the society’s *cross-sectional* inequality explodes due to the *cumulative* effect of all past shocks (in that some agents will have received unboundedly more total endowment in periods $\tau \leq t$ than other agents as $t \rightarrow \infty$). [Theorem 2](#) identifies the complementary property that a *single* shock leads to unboundedly large welfare differences between two agents with *otherwise identical histories* as $t \rightarrow \infty$.

- (a) *Martingale Property*: The marginal cost process defines a martingale because the principal *optimally smooths costs* over time and across states. Fix some $\mathbf{v} \in D$ and consider the cost to the principal of increasing this promise to $\mathbf{v}' := \mathbf{v} + \varepsilon \mathbf{1}$ for some $\varepsilon > 0$. When $\varepsilon > 0$ is small, this cost is approximated by the marginal cost $D_1 P(\mathbf{v}, s)$. One way to deliver the additional utility in an incentive-compatible manner is to increase the continuation promises from \mathbf{w}_i to $\mathbf{w}'_i := \mathbf{w}_i + (\varepsilon/\alpha)\mathbf{1}$ for all $i \in S$. Using the Bellman equation (FE), the cost of this perturbation is approximately $\sum_{i=1}^d f_{s_i} D_1 P(\mathbf{w}_i, i)$. By an envelope argument, this perturbation is locally optimal, yielding equality of these two expressions and thus the martingale property $D_1 P(\mathbf{v}, s) = \sum_{i=1}^d f_{s_i} D_1 P(\mathbf{w}_i, i)$.
- (b) *Differentiation in Direction 1*: This is the *unique* direction of change for \mathbf{v} that increases the agent's *ex ante* continuation utility *while leaving his information rent unchanged*. This can be seen from the incentive constraint (\mathbf{IC}_{ij}^*) in two ways. First, a perturbation of \mathbf{v} leaves the left-hand side of (\mathbf{IC}_{ij}^*) unchanged if and only if it is taken in the direction $\mathbf{1}$. Second, suppose the principal wants to perturb \mathbf{w}_i to some $\hat{\mathbf{w}}_i$ such that type i 's *ex ante* continuation utility increases by $\varepsilon > 0$, i.e., $\mathbf{E}^{\mathbf{f}_i}[\hat{\mathbf{w}}_i] - \mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_i] = \varepsilon$. Other types $j \neq i$ may value this perturbation differently if $\mathbf{f}_j \neq \mathbf{f}_i$. By setting $\hat{\mathbf{w}}_i - \mathbf{w}_i = \varepsilon \mathbf{1}$, the principal guarantees that *all* types value this perturbation in the same way, i.e., that $\mathbf{E}^{\mathbf{f}_j}[\hat{\mathbf{w}}_i - \mathbf{w}_i] = \varepsilon$ for *all* $j \in S$, ensuring that the perturbation preserves (\mathbf{IC}_{ij}^*) .³³
- (c) *Strict Positivity*: Somewhat subtly, the partial derivatives $P_i(\mathbf{v}, s)$ may be *negative*, as increasing a single component v_i of \mathbf{v} has two countervailing effects.³⁴ First, it increases the promise to type i , which mechanically increases costs. Second, it tightens incentive constraints for higher types $j > i$ that can misreport as type i , but also adds slack to type i 's incentive constraints. Whenever it adds “enough” slack to the latter incentive constraints, increasing v_i can lead to an overall decrease in costs. Nevertheless, the marginal cost martingale is always positive because, as described above, differentiation in the direction $\mathbf{1}$ does not affect any incentive constraints. This will be important for applying the Martingale Convergence Theorem below.

To preview how [Proposition 4.4](#) leads to [Theorem 1](#), we recall two benchmarks. First, when the agent does not have private information, the first-best contract fully stabilizes the agent's consumption and utility ([Section 3.1](#)), yielding a *constant* marginal cost martingale.

Second, when the agent's type is private but i.i.d. over time ([Thomas and Worrall, 1990](#)), the principal provides incentives by making the agent's *ex ante* continuation utility $w_i = \mathbf{E}^{\mathbf{f}}[\mathbf{w}_i]$ vary with his reported type i (where \mathbf{f} denotes the type-independent transition probability). The optimal contract sets $w_d > v > w_1$ in order to (a) reward the highest-type agents (who get the lowest current transfers) for not under-reporting and (b) punish the lowest-type agents, so as to deter other types from under-reporting. Consequently, the [Thomas and Worrall \(1990\)](#) martingale always *splits*: $\hat{P}'(w_1) < \hat{P}'(v) < \hat{P}'(w_d)$. This is costly to the principal because her value function is convex, where the “size” of martingale splitting quantifies the *cost of incentives*, e.g., the difference $\hat{P}'(w_d) - \hat{P}'(v) > 0$ is determined by the magnitude of the Lagrange multipliers on the agent's incentive constraints. [Thomas and Worrall \(1990\)](#) show that $\hat{P}'(v^{(t)}) \rightarrow 0$ almost surely, i.e., the cost of incentives *vanishes* in the long run. We sketch their argument as follows. First, since their value function \hat{P} is increasing, the Martingale Convergence Theorem guarantees that $\hat{P}'(v^{(t)})$ converges to *some* non-negative random variable. Second, because

³³In particular, for generic transition matrices $[\mathbf{f}_i]_{i=1}^d$, the direction $\mathbf{1}$ is the *unique* direction with this property.

³⁴See Theorem 6 in Section J of [Bloedel, Krishna, and Leukhina \(2024a\)](#) for a formal statement. Economically, the “interim” Pareto frontier that traces the principal's costs as a function of \mathbf{v} is generally *not* monotonic. This is closely related to the fact that the optimal contract is generally not renegotiation-proof (see Step 2 below).

their state variable v is one-dimensional and their value function \hat{P} is strictly convex and satisfies Inada conditions, the derivative $\hat{P}' : \mathcal{U} \rightarrow \mathbb{R}_{++}$ is a *bijection*. These two facts imply that $\hat{P}'(v^{(t)})$ must converge to zero along almost every path. To see why, note that for any path along which $\hat{P}'(v^{(t)})$ converges to a strictly positive number, bijectivity implies that $v^{(t)}$ must converge to some $v^* \in \mathcal{U}$. But then, letting $(w_i^*)_{i \in S}$ denote the optimal (ex ante) continuation utilities at v^* , martingale convergence requires that $\hat{P}'(w_1^*) = \hat{P}'(v^*) = \hat{P}'(w_d^*)$, which contradicts martingale splitting at v^* . Thus, such paths must arise with zero probability. It follows that $\hat{P}'(v^{(t)}) \rightarrow 0$ and therefore (by bijectivity) that $v^{(t)} \rightarrow -\infty$ almost surely.

Step 2: Martingale Splitting and Convergence. Persistence greatly complicates the above logic. Economically, it implies that the agent's current type determines his preferences over continuation contracts, giving the agent new Markov information rents and the principal new opportunities to *screen the agent through continuation contracts* that manipulate those rents (recall (IC_{ij}^*)). For example, with persistence, it may be possible for the principal to incentivize truthtelling while also making the agent's *ex ante* continuation utility $\mathbf{E}^{f_i}[\mathbf{w}_i]$ constant across all types $i \in S$ by suitably choosing the interim continuation promises $(\mathbf{w}_i)_{i \in S}$ to distort the continuation payoffs $\mathbf{E}^{f_j}[\mathbf{w}_i]$ for non-truthful types $j \neq i$. Given this new channel for incentive provision, it is no longer clear whether the marginal cost martingale splits (i.e., $D_1 P(\mathbf{w}_1, 1) < D_1 P(\mathbf{v}, s) < D_1 P(\mathbf{w}_d, d)$) and, if it does, what relation this has to the cost of incentives. Furthermore, because $D_1 P(\cdot, s) : D \rightarrow \mathbb{R}_{++}$ is *not* injective in our setting due to the multidimensionality of $\mathbf{v} \in D$, it is unclear how to connect convergence of the martingale to limiting properties of the $\mathbf{v}^{(t)}$ process.

To overcome these challenges, we first identify a special class of histories at which the principal does not need to screen through continuation contracts. Suppose that the agent has revealed himself to be of type i in period t , and the principal wants to give him *ex ante* promised utility $w \in \mathcal{U}$ starting in period $t + 1$ (before his period $t + 1$ type is realized). The cost-minimizing way to do this is to solve the *efficiency problem*

$$\begin{aligned} K(w, i) &:= \min_{\mathbf{w}_i \in D} P(\mathbf{w}_i, i) \\ \text{s.t.} \quad &\mathbf{E}^{f_i}[\mathbf{w}_i] \geq w. \end{aligned} \tag{Eff}_i$$

We call solutions to (Eff_i) *efficient*. Efficiency corresponds to a kind of *renegotiation-proofness*: even *after* the agent reveals himself to be of type i , so that incentive compatibility in the current period is no longer a concern, there is no way for the principal to reduce her continuation costs while also improving the agent's (expected) continuation payoff. The i.i.d. case is simple precisely because the optimal contract is always renegotiation-proof in this sense (see Remark C.7).³⁵ But with persistence, the optimal contract is typically *not* renegotiation-proof because the principal's choice of \mathbf{w}_j affects the agent's Markov information rent in each (IC_{ij}^*) with $i > j$. However, Assumption NHB implies that the highest type's promised utility \mathbf{w}_d does not enter any of these incentive constraints. Thus, the optimal contract *is* efficient at those histories where the agent had reported the highest endowment ω_d in the previous period (Lemma C.8). Moreover, at such histories, the marginal cost martingale splits and the size of its splitting pins

³⁵This is the same notion of renegotiation-proofness that is informally discussed in Thomas and Worrall (1990, p. 369), and which corresponds to having a monotonic Pareto frontier. We do not delve into the subtleties of formally defining renegotiation-proof contracts (cf. Strulovici (2017, 2020)). Also note that the efficiency problem (Eff_i) is analogous to what Fernandes and Phelan (2000) call the "planner's problem," while our recursive problem (RP) is analogous to what they call the "auxiliary planner's problem."




FIGURE 2.—Martingale convergence argument for $d = 2$ and FOSD (left) and i.i.d. (right) types.

down the cost of incentives, in analogy to the i.i.d. case (Lemmas C.9 and C.10). In particular, at such histories, the martingale does *not* split if and only if the contract perfectly stabilizes the agent's consumption, which is inconsistent with incentive compatibility (Lemma C.18).

Next, we use these observations to show that the marginal cost martingale converges to 0 almost surely. By Proposition 4.4, the Martingale Convergence Theorem implies that the marginal cost martingale converges to a non-negative, integrable random variable. To show that this limit random variable equals 0, we argue as in Figure 2, which illustrates the special case of binary types ($d = 2$). The efficiency rays \tilde{E}_i characterize the solutions of the efficiency problem (Eff_i) in terms of dual variables, viz., the derivative $DP(\mathbf{w}_i, i)$, which is pinned down by the first-order conditions for (Eff_i). When types are i.i.d., there is a single efficiency ray (corresponding to the one-dimensional state space in Thomas and Worrall (1990)), and the derivative process traverses it because the optimal contract is renegotiation-proof. Consequently, a similar argument to that in Thomas and Worrall (1990) shows that if the marginal cost martingale were to converge to a positive number along some path, then the agent's promised utility vector would also converge: $\mathbf{v}^{(t)} \rightarrow \mathbf{y} \in D$ almost surely.³⁶ This is impossible by the martingale splitting results described above, so the martingale must converge to zero along almost every path. When types are persistent, we show that a similar argument applies *along the subsequence of histories* at which $s^{(t)} = d$ (i.e., the agent truthfully reported to have the highest endowment ω_d in the previous period), where the optimal contract is efficient. Finally, we show that this convergence *extends to (almost) all histories*, establishing that $D_1P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ almost surely (Lemma C.19).

Step 3: Convergence of Multipliers and Allocations. As noted above and depicted in Figure 2, the principal's marginal cost $D_1P(\cdot, s)$ is *not* injective when types are persistent. Thus, even though the marginal cost martingale converges to 0, the agent's promised utility $\mathbf{v}^{(t)}$ may remain transient along histories where the optimal contract is not efficient.³⁷

To rule this out, we show that martingale convergence implies that the cost of incentives (i.e., the vector of Lagrange multipliers on the agent's incentive constraints) also converges to zero. Given this, it can be deduced that the contract either (a) converges to the first-best or (b)

³⁶The facts that (i) $\mathbf{v}^{(t)}$ converges and (ii) its limit point is in D (rather than being a boundary point) rely on the Inada conditions in Assumption DARA(a). See Section 6.1 for further discussion.

³⁷For instance, in the left panel of Figure 2, the derivative $DP(\mathbf{v}^{(t)}, s^{(t)})$ may traverse the blue region below \tilde{E}_1 at such histories even as $DP_1(\mathbf{v}^{(t)}, s^{(t)})$ (depicted by the green line) converges to zero.

immiserates the agent, only the latter of which is consistent with incentive compatibility. These arguments are involved, but the main ideas are as follows. When types are i.i.d. or binary (as in [Corollary 4.1](#)), we can show that, in essence, the marginal cost martingale provides a uniform bound for the multipliers, so that almost sure convergence of the former implies almost sure convergence of the latter. For general type processes, while we cannot rule out that the vector of multipliers take infinitely-many excursions away from zero, [Assumption Markov](#) implies that the agent’s endowment almost surely returns to the highest level ω_d infinitely often. Using the fact that the contract is efficient at such histories, we show that the multiplier process exhibits a “renewal property” whereby it returns to a neighborhood of zero infinitely often; moreover, it mixes quickly enough to ensure that long excursions away from zero are very rare. This lets us show that the multipliers converge to zero in probability ([Lemma C.22](#)). Given convergence of the multipliers, convergence of consumption and utility follows quickly ([Lemma C.23](#)).

5. SOLVED EXAMPLES

We now turn to the optimal contract’s short- and medium-run properties. [Section 5.1](#) investigates the *speed* of immiseration. [Section 5.2](#) studies short-run dynamics and distortions.

For these purposes, we focus on the special case of CARA utility and $d = 2$ types with FOSD transitions (and assume [\(TVC\)](#)-Regularity). While restrictive, these assumptions yield two key benefits. First, they permit a direct comparison to the i.i.d. benchmark, in which full solutions are known only under CARA utility ([Thomas and Worrall, 1990](#), [Green, 1987](#)). Second, they allow for a simple graphical depiction of the optimal contract’s dynamics. We can analytically derive structural properties of the optimal contract in this setting (Properties (a)–(e) in [Section 5.2](#)), but a full solution requires numerical simulations.

Parameters. For the numerical simulations, we use the following parameters. The two endowment levels are $\omega_1 = -\log(5)$ and $\omega_2 = \log(10)$. The agent’s utility function is $U(c + \omega_i) = -e^{-(\omega_i + c)}$ and his discount factor is $\alpha = 0.5$. For simplicity, the type process has *symmetric* transitions: $q := f_{11} = f_{22}$ and $q \geq 1/2$.³⁸ We focus on three persistence levels: $q = 0.5$ (i.i.d.), $q = 0.65$ (low persistence), and $q = 0.8$ (high persistence). In [Section 5.1](#), we also consider $q = 0.95$ (very high persistence).

5.1. How Fast is the Ride on the Highway to Hell?

To begin, we investigate how persistence affects the *speed* of immiseration. Our main finding is that, in the medium- and long-run, greater persistence yields *faster* immiseration.

For this exercise, it is useful to slightly reinterpret our model. Consider, as in [Green \(1987\)](#), the problem of a utilitarian planner presiding over an economy populated by a continuum of homogeneous agents who face idiosyncratic endowment shocks. The economy is *open*: the planner can finance consumption in each period by borrowing and saving in the “world market” at the risk-free rate $R = 1/\alpha$, and therefore faces a single intertemporal resource constraint. It is well known that our principal-agent problem is equivalent to the *dual* of the open economy planning problem (by viewing the open-economy resource constraint as our principal’s objective function).³⁹ Given this correspondence, we find it most insightful to focus on

³⁸Neither our analytical results nor the qualitative findings from our numerical simulations rely on this symmetry assumption. It merely simplifies the exposition by (i) letting us parameterize persistence with a *single* number and (ii) ensuring that the stationary distribution of types is *fixed* (at $(0.5, 0.5)$) as we vary persistence.

³⁹See, e.g., [Golosov, Tsyvinski, and Werquin \(2016b\)](#), [Green \(1987\)](#).

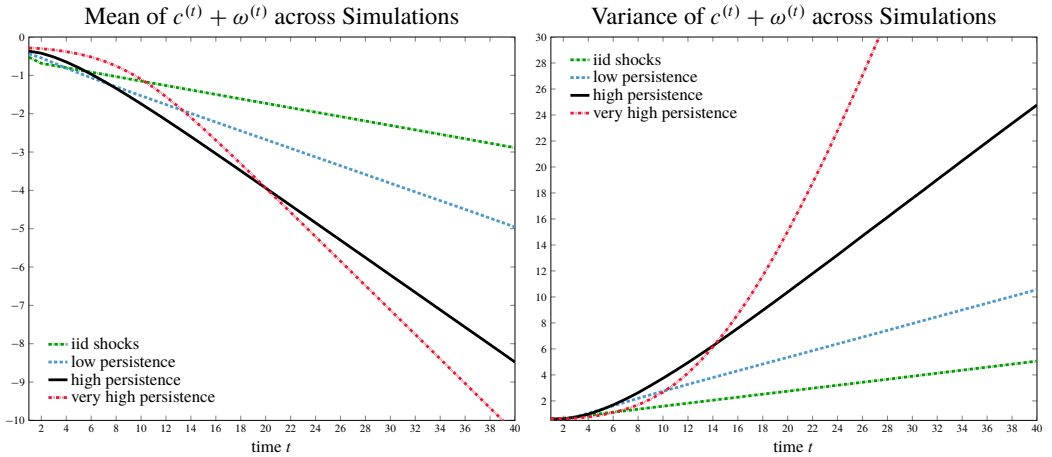


FIGURE 3.—Mean consumption (left) and variance of consumption (right) across simulations.

the time evolution of two statistics: (i) *mean consumption* $\mu_{C,t} := \mathbf{E}[c^{(t)} + \omega^{(t)}]$ at time t and (ii) the *variance of consumption* $\sigma_{C,t}^2 := \mathbf{V}[c^{(t)} + \omega^{(t)}]$ at time t , where the expectations are over paths of the type process from the perspective of time 0. We can then interpret (i) $\mu_{C,t}$ as describing the open economy's *aggregate consumption* at time t , and (ii) $\sigma_{C,t}^2$ as describing its *cross-sectional consumption inequality* at time t . We explore how persistence shapes their evolution.

To implement this numerically, for each of the persistence levels $q \in \{.5, .65, .8, .95\}$, we simulate 420,000 sample paths of the optimal contract for 40 time periods.⁴⁰ Figure 3 plots the corresponding sample means and variances of consumption (which serve as proxies for $\mu_{C,t}$ and $\sigma_{C,t}^2$, respectively). The following key patterns emerge:

1. *The economy tends towards aggregate misery and infinite inequality:* at all persistence levels, $\mu_{C,t} \rightarrow -\infty$ and $\sigma_{C,t}^2 \rightarrow +\infty$. These features are consistent with Theorem 1 and the fact that, in the i.i.d. case, consumption follows a random walk with negative drift (Green, 1987, Thomas and Worrall, 1990).⁴¹ As in Green (1987), the fact that aggregate consumption $\mu_{C,t} \rightarrow -\infty$ effectively means that the planner finances consumption in early periods by going into unbounded debt in later periods.
2. *In the medium-run, greater persistence yields faster immiseration and higher inequality:* after the first few periods, greater persistence causes $\mu_{C,t}$ to decrease more quickly and $\sigma_{C,t}^2$ to increase more quickly.
3. *In the short-run, the patterns in Point 2 are reversed:* for the first few periods, greater persistence causes $\mu_{C,t}$ to decrease more slowly and $\sigma_{C,t}^2$ to increase more slowly. Furthermore, the initial aggregate consumption level $\mu_{C,0}$ increases with persistence.

We view Point 2 as being the primary lesson of this exercise, and interpret it as suggesting that private information has an even more pernicious effect on welfare and equity outside the classic i.i.d. benchmark. In Section 5.2 below, we attribute this comparative static result to

⁴⁰We draw 21 initial $\mathbf{v}^{(0)}$ points uniformly from a grid in the domain $D = V_2$ (recall Figure 1), with the $v_1^{(0)}$ and $v_2^{(0)}$ values evenly spaced between -10 and -1 . For each such $\mathbf{v}^{(0)}$, we run 20,000 simulations, half of which start from $\omega^{(0)} = \omega_1$. Thus, we effectively initialize the type process at its stationary distribution.

⁴¹Two clarifications: (i) as consumption is unbounded, Theorem 1 does not *directly* imply that $\mu_{C,t} \rightarrow -\infty$, and (ii) beyond the i.i.d. case, consumption does *not* follow a random walk (see Section 5.2 below).

novel intertemporal distortions that arise from the agent’s Markov information rents, and which become more severe as persistence increases.

To reconcile Points 2 and 3, note that the $\mu_{C,t}$ and $\sigma_{C,t}^2$ curves in Figure 3 are *eventually* linear, but (except in the i.i.d. case) are *initially* concave and convex, respectively.⁴² Moreover, increasing persistence makes these curves “flatter” in the initial periods. Intuitively, this arises from the mechanical fact that greater persistence increases the “mixing time” that it takes for the promised utility process to become “stationary” in an appropriate sense (see Section 5.2 for a more precise discussion). For reference, it can be shown that the optimal contract under *perfect* persistence ($q = 1$) is “static,” i.e., it specifies a constant path of consumption conditional on the initial endowment report (which it imperfectly insures against). Thus, as we approach the perfect persistence limit ($q \rightarrow 1$), it is natural for the optimal contract to specify “almost constant” consumption in early periods (see the $q = 0.95$ case in Figure 3).⁴³ Finally, note that because we are holding the initial promised utility fixed, $\mu_{C,0}$ *must* increase with persistence to compensate for the faster negative trend in later periods.

Next, in Section 5.2, we study the optimal contract’s short-run properties in detail. Later, in Section 6.3, we revisit Point 1 when discussing a *closed economy* variant of the planning problem described here, as in Atkeson and Lucas (1992).

5.2. Short-Run Properties

We begin by describing the optimal contract’s basic structural properties. We then turn to the qualitative and quantitative implications of persistence.

Basic Structure. The joint assumptions of CARA utility and binary types ($d = 2$) yield two simplifications. First, CARA implies that optimal contracts are *homogeneous of degree 1 (HDI)* with respect to interim promises: $\xi^*(a\mathbf{v}, s) = a \cdot \xi^*(\mathbf{v}, s)$ for all states $(\mathbf{v}, s) \in D \times S$ and positive scalars $a > 0$.⁴⁴ Second, $d = 2$ implies that the domain $D = V_2 := \{\mathbf{v} \in \mathcal{U} : v_2 > v_1\}$ (by Theorem 3 in Appendix B), as illustrated in Figure 1. Thus, on the equilibrium path, the optimal contract has a simple structure determined by a countable infinity of rays in V_2 . Figure 4 illustrates the following properties:

- (a) There are a countable infinity of *rays* in V_2 , denoted E_1 , E_2 , and B_k for $k \in \mathbb{N}$. Each E_i ray traces out solutions to the corresponding efficiency problem (Eff_i).
- (b) Starting from *any* $\mathbf{v} \in V_2$, promised utility jumps to E_2 after a high shock ($i = 2$).
- (c) Starting from $\mathbf{v}^{(0)} \in E_2$, a high shock sends promised utility to the “northeast” along E_2 , while k *consecutive* low shocks send it to B_k . If there is a high shock after k consecutive low shocks, promised utility returns to E_2 , and the cycle starts again. Due to HD1, the proportional scaling $\|\mathbf{v}^{(t+1)}\|/\|\mathbf{v}^{(t)}\|$ between pre- and post-shock promised utilities at each step does not depend on $\|\mathbf{v}^{(t)}\|$, only on the ray that contains $\mathbf{v}^{(t)}$.

⁴²In the i.i.d. case, these curves are linear after an initial kink at $t = 1$. This kink arises because, in the i.i.d. case, the $\mathbf{v}^{(t)}$ process “immediately jumps” at $t = 1$ from its exogenous initial condition $\mathbf{v}^{(0)}$ to a specific subset of the domain that it remains in thereafter (see properties (c) and (d) in Section 5.2 for details).

⁴³Specifically, it is natural for the optimal contract to converge *pointwise* in the $q \rightarrow 1$ limit to the optimal “static” contract at $q = 1$ *along the constant histories that arise with positive probability at $q = 1$* . Point 2 indicates that this convergence is not uniform or does not occur along non-constant histories, so that even when $q \approx 1$, the compound effect of multiple type transitions eventually becomes dominant. See Battaglini and Lamba (2019, Section 4.3) and Krasikov and Lamba (2021, Proposition 2) for related points.

⁴⁴This fact and Properties (a)–(e) below follow directly from the CARA, $d = 2$, and FOSD assumptions and the optimality conditions stated in Supplemental Appendices C.1 and C.3.3. For completeness, derivations are in Section K of Bloedel, Krishna, and Leukhina (2024a).

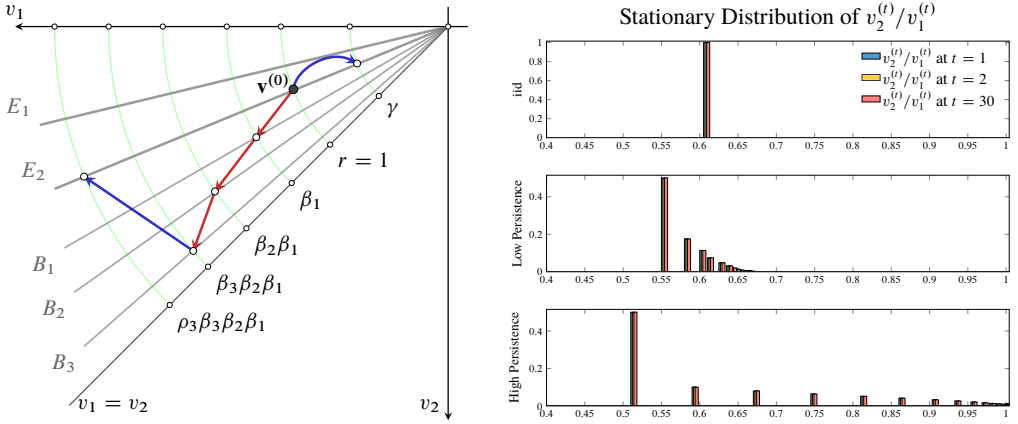


FIGURE 4.—**LEFT: QUALITATIVE DEPICTION OF OPTIMAL CONTRACT.** Initial $\mathbf{v}^{(0)} \in E_2$ normalized so that $\|\mathbf{v}^{(0)}\| = 1$. Proportional scalings are $\gamma \in (0, 1)$ if first shock is *good*, $\beta_k > 1$ if k th consecutive shock is *bad*, and $\rho_k > 1$ when *returning* to E_2 with the first good shock after k consecutive bad shocks. **RIGHT: STATIONARY DISTRIBUTION OVER E_2 AND B_k RAYS.** Simulated distribution of $v_2^{(t)}/v_1^{(t)}$ given initial $\mathbf{v}^{(0)}$ drawn from the stationary distribution over rays, which is $\mathbf{P}(\mathbf{v}^{(0)} \in E_2) = 0.5$ and $\mathbf{P}(\mathbf{v}^{(0)} \in B_k) = 0.5(1 - q)q^{k-1}$ if all rays are distinct.

- (d) If types are i.i.d., all the rays coincide ($E_1 = E_2 = B_k$ for all k). But if types are persistent (FOSD but not i.i.d.), E_2 and the B_k rays lie strictly “below” E_1 .
- (e) The cyclicity and HD1 in Property (c) imply that the optimal contract induces a stationary distribution over the E_2 and B_k rays.⁴⁵ The right panel of Figure 4 plots the simulated stationary distributions in our examples, indicating both (i) the slopes of E_2 and the B_k rays and (ii) the fraction of time spent on each. (Both panels of Figure 4 depict that B_1 lies strictly below E_2 and that, for all k , B_k lies strictly below B_{k-1} . This is true in all of our numerical examples, but we did not prove these properties.)

Next, we unpack the economic meaning and implications of these Properties (a)–(e).

Managing Markov Information Rents. As noted in Section 4.3, a fundamental implication of persistence is that the optimal contract is *not* renegotiation-proof: it need *not* solve the efficiency problem (Eff_i), except after the highest type realizations ($i = d$). Intuitively, these inefficiencies *optimally reduce the agent’s Markov information rents*.

Property (d) and our simulations help clarify this intuition. The fact that the B_k rays lie below E_1 means that, after a low shock, the agent’s continuation utility \mathbf{w}_1 is pushed *inefficiently close* to the lower boundary of V_2 . For any fixed expected continuation utility $\mathbf{E}^{\mathbf{f}^1}[\mathbf{w}_1]$ for a low-type agent, this distortion *compresses* the difference $w_{12} - w_{11} > 0$ and thus reduces a high-type agent’s Markov information rent $\alpha(2q - 1)(w_{12} - w_{11}) > 0$, adding slack to (IC_{21}^*) . However, it also yields a *tighter* incentive constraint (i.e., left-hand side of (IC_{21}^*)) in the *next* period, when the new promised utility state is \mathbf{w}_1 . This makes further compression *even more* attractive after a *second* consecutive low shock, and so on.

⁴⁵Formally, the “slope” process $v_2^{(t)}/v_1^{(t)}$ has a stationary distribution on $(0, 1)$. The support characterizes $\{E_2, B_1, \dots\}$, while the probabilities are determined by the stationary distribution of the “counting process” over the number of consecutive low shocks, i.e., the time elapsed since the last high shock. This process is defined as $n^{(t)} := \inf\{k \in \mathbb{N} \cup \{0\} : s^{(t-k)} = 2\}$ and its stationary distribution is given by $\lim_{t \rightarrow \infty} \mathbf{P}(n^{(t)} = 0) = 0.5$ and $\lim_{t \rightarrow \infty} \mathbf{P}(n^{(t)} = k) = 0.5(1 - q)q^{k-1}$ for all $k \in \mathbb{N}$. If all the rays are distinct, $n^{(t)} = 0$ corresponds to $\mathbf{v}^{(t)} \in E_2$ and $n^{(t)} = k \geq 1$ corresponds to $\mathbf{v}^{(t)} \in B_k$, as described in Figure 4.

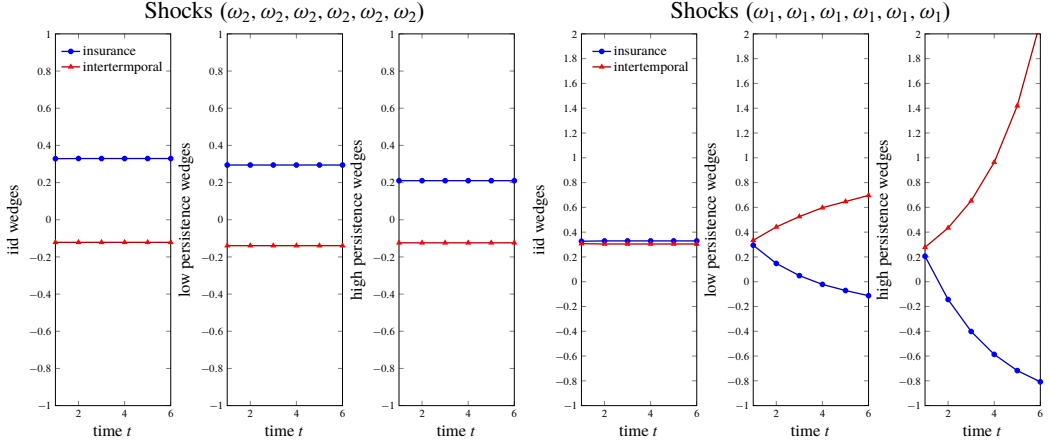


FIGURE 5.—Wedges along strings of 6 consecutive high (left) and low (right) shocks in periods $t \in \{1, \dots, 6\}$, starting from $\mathbf{v}^{(0)} \in E_2$. *Insurance wedge* $\equiv U'(c_1^{(t)} + \omega_1)/U'(c_2^{(t)} + \omega_2) - 1$ measures departures from within-period perfect insurance. *Intertemporal wedge* $\equiv \mathbf{E}[U'(c^{(t+1)} + \omega^{(t+1)}) | \omega^{(t)}]/U'(c^{(t)} + \omega^{(t)}) - 1$ measures departures from Euler equation. Both are 0 in the first-best.

This is borne out in Figure 4. As depicted in the left panel, except in the i.i.d. case, B_k lies below B_{k-1} for all k , indicating greater compression as the number of consecutive low shocks increases. Furthermore, as depicted in the right panel, the B_k rays become steeper and more “spread out” as persistence increases, indicating greater compression for each fixed number of consecutive low shocks.

Implications. These forces shape economically important features of optimal contracts:

1. *Speed of Immiseration:* As noted in Section 5.1, greater persistence yields (i) *faster* immiseration in the medium-run, but (ii) *slower* immiseration in the short-run. We can now explain these findings as follows: greater persistence simultaneously (i) yields steeper B_k rays, so that $\mathbf{v}^{(t)}$ decreases (to the “southwest” in Figure 4) *more quickly* after long strings of consecutive low shocks (which also become more likely), and (ii) makes initial strings of high shocks more likely, increasing the “mixing time” for $\mathbf{v}^{(t)}$ to settle into its cyclical dynamics on the E_2 and B_k rays (Properties (c) and (e)).
2. *Mix of Incentives:* Persistence causes the optimal contract to use a different mix of *intra-* and *inter-*temporal incentives. To illustrate, Figure 5 plots the standard *insurance* and *intertemporal wedges* along sequences of consecutive high shocks (left panel) and consecutive low shocks (right panel). In the i.i.d. case, (i) the insurance wedge is *always* positive because the contract *under-*insures the agent (Thomas and Worrall, 1990) and (ii) the intertemporal wedge is positive after low shocks (which cause consumption to drift down). By contrast, with persistence, *after consecutive low shocks* (i) the insurance wedge *becomes negative* (*over-*insurance), while (ii) the intertemporal wedge remains positive but *becomes an order of magnitude larger*. Both features arise from the fact that persistence yields steeper B_k rays: (i) over-insurance ($u_1 > u_2$) after low shocks serves to “compress” promised utilities (by increasing w_{11} relative to w_{12}), while (ii) promised (and hence flow) utilities decrease more quickly after low shocks.
3. *Order-Dependence:* With i.i.d. types and CARA utility, the optimal contract is order-independent: the agent’s consumption and continuation utility depend only on the *number* of high and low shocks to date, but *not* the order in which they arrived. Thomas and Worrall (1990) attribute this to the absence of wealth effects under CARA. We find that it also

depends critically on the i.i.d. assumption. In our examples with persistence, for a fixed number of high and low shocks, the agent is *worse off* if the low shocks occurred *earlier*—initial bad luck is punished more than later bad luck.⁴⁶ Intuitively, this is driven by the following property (depicted in [Figure 4](#)): following a high shock, the low-type continuation utility $v_1^{(t)}$ *increases* if the preceding shock was also high ($\mathbf{v}^{(t-1)} \in E_2$) but may *decrease* if the preceding shock was low ($\mathbf{v}^{(t-1)} \in B_k$). The latter feature, which does not arise in the i.i.d. case, serves to *re-introduce slack* into the (\mathbf{IC}_{21}^*) constraint after it was successively tightened along a string of low shocks.

6. DISCUSSION

[Section 6.1](#) revisits our baseline model assumptions. [Section 6.2](#) discusses the role of mean-reversion. [Section 6.3](#) considers an extension to a closed-economy planning problem.

6.1. The Role of Baseline Assumptions

We discuss (i) Assumptions [DARA](#) and [NHB](#) and (ii) other sources of private information.

Shape of Utility Function. Assumption [DARA](#) plays both substantive and technical roles in our analysis. Part (a) embeds the economically substantive conditions needed for [Theorem 1](#): strict monotonicity and concavity of U ensure that first-best insurance is not implementable (i.e., the incentive problem is nontrivial), while the upper Inada condition $\lim_{c \rightarrow \infty} U'(c) = 0$ ensures that there is no upper bound for the principal's cost of incentive provision (i.e., maintaining truth-telling requires unboundedly large variability in consumption as the level of utility increases).⁴⁷ The remaining pieces of Assumption [DARA](#) play technical roles: smoothness of U in part (a) lets us take well-behaved first-order conditions and is related to smoothness of the value function P , unboundedness below of U in part (b) implies that the optimal contract is interior, and [DARA](#) in part (c) implies that the principal's problem is convex, so that report-contingent randomization is not needed. It is an open question whether these technical assumptions can be relaxed.⁴⁸

Feasible Reporting Strategies. Assumption [NHB](#) prevents the agent from over-reporting his type. We view this as a relatively mild assumption. First, under-reporting constraints appear to be the empirically relevant concern in hidden endowment settings (see [Feldman and Slemrod \(2007\)](#) for evidence from US tax data). Second, it is often argued that [NHB](#) is without loss of generality if (i) the agent's private type represents a partially verifiable variable, like his endowment or labor productivity, and (ii) he cannot covertly trade or produce outside of the

⁴⁶For example, consider period $t = 6$ after the agent has had three high shocks and three low shocks. In the high persistence ($q = 0.8$) case, the agent's flow and continuation utilities contingent on a low shock at $t = 6$ ($u_1^{(6)}$ and $v_1^{(6)}$) are approximately *two times lower* if the three prior low shocks occurred in periods $\{0, 1, 2\}$ than if they occurred in periods $\{3, 4, 5\}$.

⁴⁷In the i.i.d. setting, [Phelan \(1998\)](#) observes that if $M := \lim_{c \rightarrow \infty} U'(c) > 0$, then the principal's marginal cost martingale is bounded above and thus converges to either 0 or its upper bound, both with positive probability. Consequently, the optimal contract generates *polarization*: both immiseration and long-run bliss arise with positive probability. The probability of immiseration nonetheless tends to one as $M \rightarrow 0$.

⁴⁸We note that [Theorem 2](#) *does* rely on the agent's utility function being unbounded below, for otherwise it is impossible to simultaneously immiserate the agent and provide unboundedly high-powered incentives. For instance, if $\inf_{c \in \mathbb{C}} U(c) = 0$, then immiseration corresponds to $v_i^{(t)} \rightarrow 0$ for all $i \in S$. However, the basic intuition—that risk-aversion makes incentive provision cheaper at lower utility levels—is robust.

contract.⁴⁹ While covert trade and production are surely important in some settings (Allen, 1985, Cole and Kocherlakota, 2001b), they are conceptually distinct from the pure reporting problems studied herein and in much of the prior literature.

Alternatively, one can interpret NHB as a *relaxation* of the “full” problem in which over-reporting is feasible for the agent, which is natural when his type is purely subjective (e.g., a taste shock). Even then, NHB corresponds to a significantly *less* relaxed problem than that considered in the popular first-order approach (FOA), which in our context would include only the “local downward” incentive constraints (\mathbf{IC}_{ij}) with $j = i - 1$. This additional robustness is notable because the FOA may be “invalid” (i.e., yield solutions that violate global incentive constraints) without strong assumptions on the type process.⁵⁰

Source of Private Information. For concreteness, we have assumed that the agent’s private information concerns his endowment, but this is not essential. Our main long-run analysis (Sections 3 and 4) can be adapted to other canonical insurance settings, including:

- *Taste shocks:* The agent has utility $U(\omega, c) = \omega u(c)$ over consumption c , where $\omega \in \mathbb{R}_+$ is privately observed taste shock, and the principal minimizes the lifetime cost of providing consumption to the agent. This model has been a workhorse specification since Atkeson and Lucas (1992), and coincides with our baseline hidden endowment model under CARA utility. Our analysis extends to this model almost verbatim, provided that $u(\cdot)$ satisfies parts (a) and (b) of Assumption DARA.
- *Productivity shocks:* The agent has utility $U(\omega, c, \ell) = u(c) - v(\omega, \ell)$ over consumption c and labor effort ℓ , where $\omega \in \mathbb{R}$ is a privately observed productivity shock. In each period, the principal offers a menu of report-contingent (c, ℓ) bundles, transfers consumption c to the agent, and collects his labor output $\omega\ell$, aiming to minimize the lifetime cost of the contract. This is the workhorse specification in the dynamic Mirrleesian taxation literature (e.g., Zhang (2009), Farhi and Werning (2013)). Many aspects of our analysis apply to this model with minor modification; we leave a full analysis to future work.⁵¹

6.2. The Role of Mean-Reversion

To clarify the role of Assumption Markov, it is useful to recall the analysis of Williams (2011) (henceforth W11), which provides an important counterpoint to our Theorems 1 and 2. W11 studies the special case of our baseline model in which the agent has CARA utility, but with one key difference: the agent’s endowment evolves as a Gaussian random walk, rather than as an

⁴⁹See Phelan (1998), Williams (2011), Golosov and Tsyvinski (2007) for versions of this interpretation. For hidden endowment settings, Fernandes and Phelan (2000, fn. 4) also argue that NHB is without loss of generality if the agent’s minimal consumption $\underline{c} > -\infty$, as the principal can then use stochastic contracts that punish an over-report with infinite disutility (with vanishing probability and cost). NHB also appears in the related family of “cash-flow diversion” models (e.g., Fu and Krishna (2019) and references therein).

⁵⁰In the i.i.d. case, Thomas and Worrall (1990) argue that only local downward constraints bind. Battaglini and Lamba (2019), Pavan (2016) discuss potential issues with the FOA when types are highly persistent.

⁵¹With separable preferences $U(\omega, c, \ell) = u(c) - v(\omega, \ell)$, consumption utility $u(c)$ serves as a “numeraire” for delivering value to the agent in a type-independent manner. Thus, the principal’s marginal cost martingale reduces to $D_1 P(\mathbf{v}, s) = 1/u'(c(\mathbf{v}, s))$, yielding the classic *Inverse Euler Equation* (IEE) which states that the inverse marginal utility of consumption is a martingale (e.g., Golosov and Tsyvinski (2007)). When preferences are not separable between consumption and type—as in hidden endowment, taste shock, and non-separable Mirrleesian models—this correspondence and the classic IEE break down. In such cases, the martingale $D_1 P(\mathbf{v}, s)$ remains the relevant object for characterizing long-run outcomes, while short-run distortions can be characterized with augmented versions of the IEE that include additional Lagrange multiplier terms (e.g., Farhi and Werning (2013), Hellwig (2021), Golosov, Troshkin, and Tsyvinski (2016a)).

ergodic finite-state Markov chain.⁵² This seemingly minor difference yields markedly different results: the optimal contract generates long-run *bliss* ($c^{(t)} + \omega^{(t)} \rightarrow +\infty$ and $u^{(t)} \rightarrow 0$ almost surely) and *vanishing incentives* ($v_i^{(t)} - v_{i-1}^{(t)} \rightarrow 0$ almost surely).

Taken together, our and W11’s results indicate that ergodicity (or “mean-reversion”) of the type process is a key determinant of the optimal contract’s long-run properties. In particular, we interpret our results as suggesting that the failure of immiseration in W11 hinges on the knife-edge assumption of *zero* mean-reversion (“permanent shocks”).⁵³

Intuitively, with *any* nonzero mean-reversion, the agent’s current type (in period t) is approximately independent of his distant-future types (in period $t + T$ as $T \rightarrow \infty$), implying that the principal and agent have approximately symmetric information about the distant future. This suggests that the classic forces underlying immiseration in the i.i.d. benchmark “should” kick in over long time horizons. Indeed, this logic is borne out in parts of our proof of [Theorem 1](#) (e.g., the “renewal” dynamics in Step 3 of [Section 4.3](#)).

Formally, Assumption [Markov](#) ensures that the type process has *impulse response functions* ([Pavan, Segal, and Toikka, 2014](#)) that vanish over long time horizons, and thus that the agent’s Markov information rents are “not too large.”⁵⁴ In contrast, when the agent’s type follows a random walk as in W11, the impulse response functions are *constant* over time because a change in the current type has a permanent additive effect on all future types. In a generalization of W11’s model, [Bloedel, Krishna, and Strulovici \(2024b\)](#) show that this implies that the agent’s Markov information rents are “so large” that (i) he is globally indifferent among *all* reporting strategies under *any* incentive compatible contract and (ii) the principal optimally elicits *no* information from the agent and provides *no* insurance. In other words, permanent shocks effectively leave the principal no latitude over the power of incentives, shutting off the classic rationale for immiseration by construction.

This discussion invites two conjectures: (i) immiseration arises for all utility functions satisfying Assumption [DARA](#) if and only if the agent’s type process has asymptotically vanishing impulse response functions, and (ii) immiseration occurs “more slowly” as we approach the permanent shock limit.⁵⁵ While our results offer a significant step towards (i), definitive answers require extending our analysis to unbounded, continuous, and non-Markovian type processes. We leave this important and challenging task to future work.

6.3. Closing the Economy

Our analysis has focused on a principal-agent version of the insurance problem ([Thomas and Worrall, 1990](#)). As noted in [Section 5.1](#), this framework is dual to the utilitarian planning

⁵²More precisely, W11 studies a continuous-time model in which the agent’s type follows a Brownian motion. [Bloedel, Krishna, and Strulovici \(2023, 2024b\)](#) clarify that W11’s results carry over to the discrete-time version of the model, permitting a more direct comparison to our analysis. W11 and [Bloedel, Krishna, and Strulovici \(2023\)](#) also allow for the agent’s type to follow a mean-reverting AR(1) process (which is closer to the Markov chains studied here) but do not characterize the optimal contract in that case.

⁵³Note that this notion of “permanent shocks” is different from the “perfect persistence” limit in which types become absorbing (cf. [Section 5.1](#)). Immiseration also fails when there are absorbing states for the mechanical reason that the optimal contract becomes “static” when such a state is reached. See [Ravikumar and Zhang \(2012\)](#), [Fuller, Ravikumar, and Zhang \(2015\)](#), [Golosov and Tsyvinski \(2006\)](#) for models of dynamic insurance with absorbing states.

⁵⁴See [Pavan, Segal, and Toikka \(2014, Theorem 1\)](#) for the definition of impulse response functions and their relation to information rents (see also [Battaglini and Lamba \(2019, Lemma 1\)](#)).

⁵⁵In our model, the permanent shock limit corresponds to increasing the number of types d without bound while suitably adjusting the transition probabilities $[\mathbf{f}_i]_{i=1}^d$ so that the type process approximates a random walk as $d \rightarrow \infty$. [Theorem 1](#) suggests that the long-run outcome “discontinuously” switches from immiseration to bliss at the permanent shock limit (i.e., the order in which we send $d \rightarrow \infty$ and $t \rightarrow \infty$ matters), which is consistent with *pointwise* but *non-uniform* convergence of optimal contracts in this limit.

problem for a continuum-agent *open economy* in which the planner can borrow and save at the risk-free rate $R = 1/\alpha$ on the world market (Green, 1987). Both of these formulations are “partial equilibrium” because they endow the principal/planner with deep pockets. As such, we view them as applying most readily to “small” economies (e.g., individual communities, private insurance exchanges, or social insurance programs at the state or national level) and over “short” horizons (e.g., the lifespan of a single individual).

Equally important is the “general equilibrium” perspective pioneered by Atkeson and Lucas (1992), which considers the more constrained *closed economy* problem in which the utilitarian planner faces a sequence of date-by-date resource constraints. Specifically, in the notation of Section 5.1, aggregate consumption $\mu_{C,t} := \mathbf{E} [c^{(t)} + \omega^{(t)}]$ must satisfy $\mu_{C,t} \leq e$ at all dates t , where $e \in \mathbb{R}$ is a per-period aggregate endowment. (The corresponding open economy problem is subject to the weaker intertemporal constraint $\sum_{t=0}^{\infty} \alpha^t \mu_{C,t} \leq \frac{e}{1-\alpha}$ on the present values of consumption and the endowment.) This approach is well suited for modeling “large” economies over “long” horizons, such as when each period corresponds to a new generation in a dynasty (Farhi and Werning, 2007, Phelan, 2006).

It is important to understand the extent to which Theorems 1 and 2 translate to such settings. Notably, our optimal contracts may be infeasible in the closed economy problem: if the open economy constraint binds and aggregate consumption is declining over time (as in Section 5.1), then the closed economy constraints must be violated in the initial periods. While it is beyond the scope of this paper to solve the closed economy problem, our analysis does have some notable implications for it. We outline these implications and associated challenges below, with the hope of stimulating further work on this problem.

To begin, note that the most tractable way to tackle the closed economy problem is to study the *decentralized (dual) problem* from Atkeson and Lucas (1992), in which the planner minimizes the expected lifetime cost of consumption transfers $\mathbf{E} [\sum_{t=0}^{\infty} q_t c^{(t)}]$ under a given sequence $\tilde{q} := \{q_t\}_{t=0}^{\infty}$ of intertemporal prices, subject to promise keeping and incentive compatibility for each agent.⁵⁶ The analysis then consists of two steps:

- *Step 1:* For each fixed \tilde{q} , the decentralized problem separates agent-by-agent into a collection of principal-agent problems of the sort we have studied, *except that the principal’s discount factor between periods t and $(t + 1)$ is now q_{t+1}/q_t , which may differ from the agent’s discount factor α* . Solve the principal-agent problem for each \tilde{q} .
- *Step 2:* Find a \tilde{q}^* at which the Step 1 solution satisfies the closed economy constraints with equality (“market clearing”). This yields a solution to the closed economy problem.

A General Immiseration Result? Our recursive approach extends, with minor modifications and technical qualifications, to the Step 1 subproblem. Let $P_{\tilde{q}}(\cdot, \cdot, t) : D \times S \rightarrow \mathbb{R}$ denote the principal’s value function in the time- t continuation of this subproblem, which is analogous to (RP) but may now depend on calendar time if $t \mapsto q_{t+1}/q_t$ is non-constant. Adapting Proposition 4.4, we see that the principal’s marginal cost process now satisfies

$$D_1 P_{\tilde{q}}(\mathbf{v}^{(t)}, s^{(t)}, t) = \frac{q_{t+1}/q_t}{\alpha} \cdot \mathbf{E} [D_1 P_{\tilde{q}}(\mathbf{v}^{(t+1)}, s^{(t+1)}, t+1) \mid \mathbf{v}^{(t)}, s^{(t)}] \quad (6.1)$$

and thus is a martingale only if $q_{t+1}/q_t = \alpha$ for all t . More generally, however, it is a non-negative *supermartingale* and thus guaranteed to converge whenever $q_{t+1}/q_t \geq \alpha$ for all t . Consequently, Theorem 1 and its proof can be extended to \tilde{q} satisfying this inequality. If such \tilde{q} could be shown to satisfy market clearing in Step 2, this would yield a general proof of immiseration for the closed economy problem.

⁵⁶See, e.g., Atkeson and Lucas (1995), Farhi and Werning (2007), Phelan (1998, 1994).

Unfortunately, market clearing (at least) sometimes requires the *opposite* inequality $q_{t+1}/q_t < \alpha$, which makes it *cheaper* than in our baseline model to backload consumption. Intuitively, this is needed to counteract the downward drift in aggregate consumption that arises under “equal discounting” (Section 5.1).⁵⁷ For such \tilde{q} , the principal’s marginal cost (6.1) is a non-negative (and unbounded) *submartingale*, and hence not guaranteed to converge in general. Thus, our proof method for Theorem 1 does not *directly* extend to the closed economy model, at least without further structure on market-clearing prices.

Towards a Special Case. As it is hard to imagine a general immiseration proof that does *not* use martingale convergence, further progress apparently requires parametric assumptions. This is not surprising: even in the i.i.d. case, immiseration-type results have only been obtained for CARA utility (or CRRA utility in taste-shock settings).⁵⁸

In particular, Atkeson and Lucas (1992) and most of the subsequent literature rely on the wealth (or scale) invariance of these preferences to ensure that optimal contracts are “scale invariant,” i.e., characterized by d numbers that (i) determine how promised utility is (proportionally or additively) rescaled following each of the d shocks and (ii) crucially, are *independent* of the current promised utility state. Without such scale-invariance, the market clearing prices in Step 2 will generally depend on the (entire path of the) economy’s cross-sectional promised utility distribution, which in turn depends on the path of prices; this requires one to solve an infinite-dimensional fixed-point problem before determining whether or not (6.1) is a supermartingale.⁵⁹

With this in mind, the analysis in Section 5 suggests a roadmap for how to generalize the CARA (or CRRA) example from Atkeson and Lucas (1992) to persistent types. Even with such preferences, the scale-invariance from the i.i.d. case *fails* with persistence due to the distortions described in Section 5.2. Nonetheless, because optimal contracts are HD1 and induce a stationary distribution over the E_2 and B_k rays, we expect that analogous scale-invariance holds *at the population level*. Thus, we conjecture that the long-run evolution of the *cross-sectional promised utility distribution* can be characterized through an eigendistribution/eigenvalue problem. This is an exciting direction for future research.

APPENDIX

These appendices contain additional details concerning the recursive formulation. Appendix A presents the sequential formulation of the contracting problem described informally in Section 2. Appendix B then presents a characterization of the recursive domain (Theorem 3) and sufficient conditions for Regularity from Section 3.

APPENDIX A: SEQUENTIAL CONTRACTS, RECURSIVE CONTRACTS, AND THEIR EQUIVALENCE

A.1. Sequential Formulation

By the Revelation Principle, we may restrict attention to truthful direct revelation mechanisms.⁶⁰ Let $\mathcal{H} := S^\infty$ denote the set of paths of type *reports*, and let $\mathcal{G} := S^\infty$ denote the set

⁵⁷Atkeson and Lucas (1992), Phelan (1998) formalize this for the i.i.d. case with CARA utility.

⁵⁸See Atkeson and Lucas (1992), Phelan (1998, 1994). Atkeson and Lucas (1995), Farhi and Werning (2007), Phelan (2006) study long-run outcomes for more general preferences in alternative models where non-degenerate stationary distributions exist, using arguments that rely on this different structure.

⁵⁹See, e.g., Atkeson and Lucas (1992, p. 446) for further discussion of the role of scale-invariance.

⁶⁰Because the agent can only under-report his type (Assumption NHB), we technically require a version of the Revelation Principle suited for environments with partial verification. The reporting constraints in each period satisfy

of paths of type *realizations*. For notational consistency with our formulation of the recursive problem (RP) in Section 3.1, we adopt the following timing convention. Denote a generic path of reports by $h = (s^{t+1})_{t=0}^{\infty}$, where $s^{t+1} \in S$ represents the agent's reported type in period t . Analogously, denote a generic path of type realizations by $g = (\delta^{t+1})_{t=0}^{\infty}$, where $\delta^{t+1} \in S$ represents the agent's realized type in period t . That is, along path $g = (\delta^{t+1})_{t=0}^{\infty}$, the agent's realized endowment in period t is $\omega^t = \omega_{\delta^{t+1}}$.

Let G^t and H^t denote the sets of length- t private and public histories, respectively. That is, $g^t := (\delta^1, \dots, \delta^t) \in G^t$ and $h^t := (s^1, \dots, s^t) \in H^t$ encode the sequences of realized and reported types, respectively, in periods $\{0, \dots, t-1\}$. Together, g^t and h^t specify the full history at the start of period t , before the agent's period- t type or report are realized.

A (pure) reporting strategy $\sigma := (\sigma_t)_{t=0}^{\infty}$ for the agent is a sequence of functions $\sigma_t : G^{t+1} \times H^t \rightarrow S$. A strategy σ is *admissible* if it never specifies over-reporting the current type, i.e., $\sigma_t((g^t, \delta^{t+1}), h^t) \leq \delta^{t+1}$ for all $g^t \in G^t$, $h^t \in H^t$, and $\delta^{t+1} \in S$. Under Assumption NHB, the agent can only use admissible strategies. The set of admissible strategies is denoted by Σ . Every $\sigma \in \Sigma$, together with the measure $\mathbf{P} \in \Delta(\mathcal{G})$ over type paths, generates a stochastic process $(\delta^{(t+1)}, s^{(t+1)})_{t=0}^{\infty}$ on $S \times S$. A strategy $\sigma^* \in \Sigma$ is *on-path truthful* if $\sigma^*((g^t, \delta^{t+1}), h^t) = \delta^{t+1}$ for all $\delta^{t+1} \in S$ and $g^t \in G^t, h^t \in H^t$ such that $g^t = h^t$ (where we define $g^0 \equiv h^0$). All on-path truthful strategies generate the same joint distribution on $\mathcal{G} \times \mathcal{H}$, under which $\delta^{(t+1)} = s^{(t+1)}$ for all $t \geq 0$ almost surely and the marginal distribution on \mathcal{H} equals \mathbf{P} . The *globally truthful* strategy σ^* is defined as $\sigma^*((g^t, \delta^{t+1}), h^t) = \delta^{t+1}$ for all $\delta^{t+1} \in S$ and $g^t \in G^t, h^t \in H^t$ (i.e., including $g^t \neq h^t$).

Following Section 3.1, we define contracts in terms of flow utilities, rather than consumption transfers. A *sequential contract* is a controlled stochastic process $\tilde{u} := (u^{(t)})_{t=0}^{\infty}$, where (i) $u^{(t)} : H^{t+1} \rightarrow \mathcal{U}$ maps histories of type reports in periods $\{0, \dots, t\}$ to flow utilities for period t , and (ii) the law of \tilde{u} is controlled by the agent's reporting strategy σ . In particular, given a type history $(g^t, \delta^{t+1}) \in G^{t+1}$ and report history $(h^t, s^{t+1}) \in H^{t+1}$, the agent's true flow utility in period t is $\psi(u^{(t)}(h^t, s^{t+1}), \delta^{t+1}, s^{t+1}) := U(C(u^{(t)}(h^t, s^{t+1}), s^{t+1}) + \omega_{\delta^{t+1}})$. We denote the set of sequential contracts by Π .

The lifetime utility function $V : \Pi \times \Sigma \times S \rightarrow \mathcal{U} \cup \{-\infty\}$, defined as

$$V(\tilde{u}, \sigma, i) := \mathbf{E}^{\sigma} \left[\sum_{t=0}^{\infty} \alpha^t \psi(u^{(t)}, \delta^{(t+1)}, s^{(t+1)}) \mid \delta^{(1)} = i \right],$$

represents the agent's time-0 preferences over sequential contracts and reporting strategies, conditional on observing the realization of his time-0 type $\delta^{(1)}$. A sequential contract $\tilde{u} \in \Pi$ *implements* $\mathbf{v} \in \mathcal{U}^d$ if it satisfies promise keeping and incentive compatibility:

$$v_i = V(\tilde{u}, \sigma^*, i) \quad \forall i \in S, \quad (\text{S-PK})$$

$$V(\tilde{u}, \sigma^*, i) \geq V(\tilde{u}, \sigma, i) \quad \forall i \in S, \sigma \in \Sigma, \quad (\text{S-IC})$$

where $\sigma^* \in \Sigma$ denotes any on-path truthful strategy. For each $\mathbf{v} \in \mathcal{U}^d$, we denote by $\Pi(\mathbf{v}) \subseteq \Pi$ the subset of sequential contracts that implement \mathbf{v} . Let $D_S := \{\mathbf{v} \in \mathcal{U}^d : \Pi(\mathbf{v}) \neq \emptyset\}$ denote the set of *sequentially implementable* promised utility vectors.

the ‘‘nested range condition’’ of Green and Laffont (1986), so straightforward adaptations of their arguments to our dynamic setting establish the appropriate version.

Given any promised utility vector $\mathbf{v} \in D_S$ and type $s \in S$, the principal's *sequential problem* is to implement \mathbf{v} at minimal cost:

$$P_S(\mathbf{v}, s) := \inf_{\tilde{u} \in \Pi(\mathbf{v})} \mathbf{E} \left[\sum_{t=0}^{\infty} \alpha^t C(u^{(t)}, s^{(t+1)}) \mid s^{(0)} = s \right], \quad (\text{SP})$$

where the expectation is taken with respect to the measure $\mathbf{P} \in \Delta(\mathcal{H})$ over paths of reports induced by (on-path) truthful reporting and the prior belief \mathbf{f}_s over initial types $\omega^{(0)} = \omega_{s^{(1)}}$. A sequential contract is *sequentially optimal* at initial state (\mathbf{v}, s) if it attains the infimum in (SP) at (\mathbf{v}, s) . For every $\mathbf{v} \in D_S$ and $s \in S$, we assume that (a) $P_S(\mathbf{v}, s)$ is well-defined and finite, and (b) a sequentially optimal contract $\tilde{u}^* \in \Pi(\mathbf{v})$ exists.⁶¹ We refer to the map $P_S : D_S \times S \rightarrow \mathbb{R}$ as the principal's *sequential value function*.

A.2. Transversality and the Equivalence of Formulations

For every recursive contract $\xi \in \Xi$ and initial $\mathbf{v}_\xi^{(0)} := \mathbf{v} \in D$, the induced allocation \tilde{u}_ξ defines a sequential contract, although it is possible that $\tilde{u}_\xi \notin \Pi(\mathbf{v})$ (recall the discussion of (TVC) in Section 3.1). We say that $\tilde{u} \in \Pi(\mathbf{v})$ is *recursively generated* if $\tilde{u} = \tilde{u}_\xi$ for some $\xi \in \Xi$ initialized at $\mathbf{v}_\xi^{(0)} = \mathbf{v}$. As is well understood, for every $(\mathbf{v}, s) \in D_S \times S$, there exists a sequentially optimal contract $\tilde{u} \in \Pi(\mathbf{v})$ that is recursively generated.⁶² It follows that $D_S \subseteq D$ and $P_S \geq P$. Furthermore, as claimed in Sections 3.1 and 3.2, the sequential problem (SP) and the recursive problem (RP) are equivalent under (TVC)-Regularity:

LEMMA A.1: *If a recursive contract $\xi \in \Xi$ is (TVC)-implementable at $\mathbf{v} \in D$, then the induced allocation \tilde{u}_ξ is feasible in (SP), i.e., $\tilde{u}_\xi \in \Pi(\mathbf{v})$. In particular:*

(a) \tilde{u}_ξ delivers promises at \mathbf{v} , i.e., for all $i \in S$,

$$v_i = \mathbf{E} \left[\sum_{t=0}^{\infty} \alpha^t \tilde{u}_\xi^{(t)} \mid s^{(1)} = i \right]. \quad (\text{DP})$$

(b) Under \tilde{u}_ξ , the globally truthful reporting strategy σ^* is optimal for the agent.

(c) Moreover, if ξ is optimal in (RP) at initial state (\mathbf{v}, s) , then \tilde{u}_ξ is optimal in (SP) at (\mathbf{v}, s) . Consequently, if the environment is (TVC)-Regular, then $D = D_S$ and $P = P_S$.

PROOF: For part (a), let ξ be (TVC)-implementable at \mathbf{v} . Iterating (PK_i) forward T times gives $v_i = \mathbf{E} \left[\sum_{t=0}^T \alpha^t \tilde{u}_\xi^{(t)} \mid s^{(1)} = i \right] + \mathbf{E} \left[\alpha^{T+1} \mathbf{E}^{f_{s^{(T+1)}}} [\mathbf{v}_\xi^{(T+1)}] \mid s^{(1)} = i \right]$ for all $i \in S$. Sending $T \rightarrow \infty$, using the Monotone Convergence Theorem on the first term, and using the Bounded Convergence Theorem on the second term (it applies under (TVC)) yields (DP). The proof of part (b) is analogous (see, e.g., (Green, 1987, Lemma 2)). Part (c) follows from parts (a) and (b) and the fact that $P_S \geq P$. Under (TVC)-Regularity, parts (a)–(c) then imply that $D_S \supseteq D$ and $P_S = P$ on $D \times S$; the fact that $D_S \subseteq D$ completes the proof. *Q.E.D.*

APPENDIX B: STRUCTURAL RESULTS OMITTED FROM SECTION 3

B.1. Recursive Domain

Herein, we characterize the recursive domain D . We need two definitions. First, for every $\mathbf{v} \in D$, we denote by $\Xi^*(\mathbf{v}) \subseteq \Xi$ the subset of feasible contracts that are (TVC)-implementable

⁶¹These two assumptions mirror Conditions R.1 and R.4, respectively, from the recursive problem (RP).

⁶²See, e.g., Fernandes and Phelan (2000), Atkeson and Lucas (1992), Doepke and Townsend (2006).

at \mathbf{v} . Thus, the set of initial $\mathbf{v} \in D$ at which (TVC) can be satisfied is

$$D^* := \{\mathbf{v} \in D : \Xi^*(\mathbf{v}) \neq \emptyset\}. \quad (\text{B.1})$$

Second, consider the following classes of type processes:

DEFINITION B.1: The type process is:

- *MLRP-ordered* (or simply *MLRP*) if f_{ki}/f_{kj} is non-decreasing in k whenever $i > j$.
- *Persistent pseudo-renewal* (or simply *PPR*) if it is FOSD and, moreover, there exists a vector $\boldsymbol{\pi} \in \mathbb{R}_+^S$ such that $f_{ij} = \pi_j$ whenever $i \neq j$.⁶³

When $d = 2$, the three properties FOSD, MLRP, and PPR are equivalent. When $d \geq 3$, (i) MLRP is stronger than FOSD but satisfied by many type processes considered in applications, such as discretized AR(1) processes, and (ii) PPR processes are distinguished within the FOSD class by the property that, conditional on a transition occurring, the probability over new types does not depend on the previous type. The MLRP and PPR classes always include all i.i.d. processes and, when $d \geq 3$, have nontrivial intersection.

THEOREM 3: Fix $d \geq 2$ and define the set $V_d := \{\mathbf{v} \in \mathcal{U}^d : v_d > v_{d-1} > \dots > v_1\}$.

- The domain D exists. It is an open, convex cone satisfying $V_d \subseteq D$. It is also independent of the discount factor $\alpha \in (0, 1)$ and utility function U (within the *DARA* class).
- $\Gamma : D \rightrightarrows (\mathcal{U} \times D)^d$ is nonempty-valued and has a convex graph.
- $D^* \subseteq D$ is nonempty, convex, unbounded below, and has decreasing returns.⁶⁴
- If FOSD holds, then $D^* \subseteq V_d$.
- If either MLRP or PPR holds, then $D = V_d$.
- If either MLRP or PPR holds and the agent has CARA utility, then $D = D^* = V_d$.

We prove Theorem 3(d) in Supplemental Appendix E during the proof of Theorem 2. The proofs of all other parts of Theorem 3 are in Section I of Bloedel, Krishna, and Leukhina (2024a); while the details are nontrivial and lengthy, the core logic builds on standard iterative arguments (Abreu, Pearce, and Stacchetti, 1990). We prove parts (b)–(c) and most of part (a) by characterizing fixed points of a set-valued operator via Tarski’s theorem. We prove parts (e)–(f) and the cone property in part (a) by characterizing solutions to systems of linear and conic programs.

Notably, alternative recursive formulations would yield different domains. For instance, in the well-known Fernandes and Phelan (2000) formulation based on *ex ante* promised and threat-point utilities, the domain is a set-valued *function* of past reports, the shape of which *changes* with the type process.⁶⁵ In contrast, our domain is independent of the agent’s reports (by construction) and type process (at least within the MLRP and PPR classes). These features are important for the tractability of our analysis.

⁶³The latter property is from Renault, Solan, and Vielle (2013), Hörner, Mu, and Vielle (2017).

⁶⁴We say that D^* is *unbounded below* if, for every $r < 0$, there exists some $\mathbf{v} \in D^*$ such that $\mathbf{v} \leq r\mathbf{1}$. We say that D^* has *decreasing returns* if $\mathbf{v} \in D^*$ implies that $a\mathbf{v} \in D^*$ for all $a \in (0, 1]$.

⁶⁵To illustrate, consider the case in which $d = 2$ and types are MLRP. The Fernandes and Phelan (2000) recursive state variable consists of the previous report $s \in S$, the *ex ante* promised utility $v^p(s) := \mathbf{E}^s[v]$ for an agent whose report of s was truthful, and the *threat point* utility $v^\dagger(s) := \mathbf{E}^{s-s}[v]$ for an agent whose report of s was a lie. The domain is then a correspondence $W : S \rightrightarrows \mathcal{U}^2$, where $W(s)$ consists of implementable $(v^p(s), v^\dagger(s))$ pairs. The sets $W(1)$ and $W(2)$ must be solved for jointly. Theorem 3(e) implies that $W(1) = \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{bmatrix} V_2$ and $W(2) = \begin{bmatrix} f_{21} & f_{22} \\ f_{11} & f_{12} \end{bmatrix} V_2$. Thus, $W(1)$ and $W(2)$ are reflections of each other about the diagonal in $\mathcal{U}^2 = \mathbb{R}_+^2$, and each $W(s)$ is a convex cone whose shape depends on the transition probabilities. If types are i.i.d., both sets collapse to the diagonal, i.e., $W(1) = W(2) = \{(t, t) : t < 0\}$.

B.2. Regularity

Regularity Conditions [R.1–R.3](#) can be verified in some cases of special interest.

LEMMA B.2: *The following hold:*

- (a) *If either (i) the type process is FOSD or (ii) the transition probabilities $\{f_i\}_{i \in S}$ are affinely independent, then Condition [R.3](#) holds.*
- (b) *If the agent has CARA utility and the type process is MLRP or PPR, then the environment is Regular (i.e., Conditions [R.1–R.3](#) hold).*

The proof of [Lemma B.2](#) is in Section J of [Bloedel, Krishna, and Leukhina \(2024a\)](#). In part (a), the FOSD case covers all instances of (weakly) positive correlation and the affine independence case covers many instances of negative correlation. Part (b) implies that the example in [Section 5](#) is Regular. Since AR(1) processes satisfy MLRP, it also implies that a suitably discretized version of the setup from [Williams \(2011\)](#) and [Bloedel, Krishna, and Strulovici \(2023\)](#) is Regular.

REFERENCES

- ABREU, DILIP, DAVID PEARCE, AND ENNIO STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58 (5), 1041–1063. [[6](#), [10](#), [32](#)]
- ALLEN, FRANKLIN (1985): “Repeated principal-agent relationships with lending and borrowing,” *Economics Letters*, 17, 27–31. [[3](#), [26](#)]
- ARELLANO, MANUEL, RICHARD BLUNDELL, AND STÉPHANE BONHOMME (2017): “Earnings and consumption dynamics: a nonlinear panel data framework,” *Econometrica*, 85 (3), 693–734. [[3](#)]
- ATHEY, SUSAN AND KYLE BAGWELL (2008): “Collusion with persistent cost shocks,” *Econometrica*, 76 (3), 493–540. [[6](#)]
- ATKESON, ANDREW AND ROBERT E. LUCAS, JR. (1992): “On Efficient Distribution with Private Information,” *Review of Economic Studies*, 59 (3), 427–453. [[2](#), [5](#), [7](#), [13](#), [16](#), [22](#), [26](#), [28](#), [29](#), [31](#)]
- ATKESON, ANDREW AND ROBERT E. LUCAS, JR (1995): “Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance,” *Journal of Economic Theory*, 66, 64–88. [[6](#), [28](#), [29](#)]
- BANERJEE, ABHIJIT V. AND ANDREW F. NEWMAN (1991): “Risk-bearing and the Theory of Income Distribution,” *The Review of Economic Studies*, 58 (2), 211–235. [[14](#)]
- BATTAGLINI, MARCO AND ROHIT LAMBA (2019): “Optimal Dynamic Contracting: The First-Order Approach and Beyond,” *Theoretical Economics*, 14 (4), 1435–1482. [[4](#), [22](#), [26](#), [27](#)]
- BLOEDEL, ALEXANDER W. AND R. VIJAY KRISHNA (2015): “Misery, Persistence, and Growth,” Working paper, Stanford University. [[1](#)]
- BLOEDEL, ALEXANDER W., R. VIJAY KRISHNA, AND OKSANA LEUKHINA (2024a): “Extended Supplement to ‘Insurance and Inequality with Persistent Private Information,’” . [[5](#), [13](#), [17](#), [22](#), [32](#), [33](#)]
- BLOEDEL, ALEXANDER W., R. VIJAY KRISHNA, AND BRUNO STRULOVICI (2023): “Persistent Private Information Revisited,” Working paper, UCLA. [[3](#), [6](#), [7](#), [13](#), [27](#), [33](#)]
- (2024b): “On the Limits of Dynamic Screening under Permanent Shocks,” Working paper, UCLA. [[6](#), [13](#), [27](#)]
- BRENDON, CHARLES (2022): “Optimal Nonlinear Savings Taxation,” Working paper. [[4](#)]
- BROER, TOBIAS, MAREK KAPIČKA, AND PAUL KLEIN (2017): “Consumption risk sharing with private information and limited enforcement,” *Review of Economic Dynamics*, 23, 170–190. [[6](#)]
- COLE, HAROLD L. AND NARAYANA R. KOCHERLAKOTA (2001a): “Dynamic Games with Hidden Actions and Hidden States,” *Journal of Economic Theory*, 98 (1), 114–126. [[6](#)]
- (2001b): “Efficient allocations with hidden income and hidden storage,” *The Review of Economic Studies*, 68 (3), 523–542. [[3](#), [7](#), [26](#)]
- DOEPKE, MATTHIAS AND ROBERT M. TOWNSEND (2006): “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory*, 126 (1), 235 – 285. [[6](#), [9](#), [31](#)]
- FARHI, EMMANUEL, CHRISTOPHER SLEET, IVÁN WERNING, AND SEVIN YELTEKIN (2012): “Non-linear capital taxation without commitment,” *Review of Economic Studies*, 79 (4), 1469–1493. [[6](#)]
- FARHI, EMMANUEL AND IVÁN WERNING (2007): “Inequality and Social Discounting,” *Journal of Political Economy*, 115 (3), 365–402. [[6](#), [28](#), [29](#)]

- (2013): “Insurance and Taxation over the Lifecycle,” *Review of Economic Studies*, 80 (2), 596–635. [4, 26]
- FARINHA LUZ, VITOR (2022): “Optimal dynamic insurance contracts,” Working paper. [5]
- FELDMAN, NAOMI E. AND JOEL SLEMROD (2007): “Estimating tax noncompliance with evidence from unaudited tax returns,” *The Economic Journal*, 117 (518), 327–352. [25]
- FERNANDES, ANA AND CHRISTOPHER PHELAN (2000): “A Recursive Formulation for Repeated Agency with History Dependence,” *Journal of Economic Theory*, 91, 223–247. [6, 9, 18, 26, 31, 32]
- FU, SHIMING AND R. VIJAY KRISHNA (2019): “Dynamic Financial Contracting with Persistence,” *RAND Journal of Economics*, 50 (2), 418–452. [6, 7, 26]
- FULLER, DAVID L., B. RAVIKUMAR, AND YUZHE ZHANG (2015): “Unemployment insurance fraud and optimal monitoring,” *American Economic Journal: Macroeconomics*, 7 (2), 249–290. [27]
- GARRETT, DANIEL, ALESSANDRO PAVAN, AND JUUSO TOIKKA (2018): “Robust Predictions in Dynamic Screening,” Tech. rep., Northwestern University. [4]
- GOLOSOV, MIKHAIL, MAXIM TROSHKIN, AND ALEH TSYVINSKI (2016a): “Redistribution and Social Insurance,” *American Economic Review*, 106 (2), 359–386. [4, 26]
- GOLOSOV, MIKHAIL AND ALEH TSYVINSKI (2006): “Designing optimal disability insurance: A case for asset testing,” *Journal of Political Economy*, 114 (2), 257–279. [27]
- (2007): “Optimal taxation with endogenous insurance markets,” *The Quarterly Journal of Economics*, 122 (2), 487–534. [3, 26]
- GOLOSOV, MIKHAIL, ALEH TSYVINSKI, AND NICOLAS WERQUIN (2016b): “Recursive Contracts and Endogenously Incomplete Markets,” in *Handbook of Macroeconomics, Volume 2*. [2, 20]
- GREEN, EDWARD J. (1987): “Lending and the Smoothing of Uninsurable Income,” in *Contractual Agreements for Intertemporal Trade*, ed. by Edward C. Prescott and Neil Wallace, University of Minnesota Press. [2, 3, 5, 6, 7, 8, 9, 13, 20, 21, 28, 31]
- GREEN, JERRY R. AND JEAN-JACQUES LAFFONT (1986): “Partially verifiable information and mechanism design,” *Review of Economic Studies*, 53 (3), 447–456. [30]
- GUO, YINGNI AND JOHANNES HÖRNER (2020): “Dynamic Allocation without Money,” Tech. rep. [6, 7]
- GUVENEN, FATIH, FATIH KARAHAN, SERDAR OZKAN, AND JAE SONG (2021): “What Do Data on Millions of US Workers Reveal About Lifecycle Earnings Dynamics?” *Econometrica*, 89 (5), 2303–2339. [3]
- HALAC, MARINA AND PIERRE YARED (2014): “Fiscal rules and discretion under persistent shocks,” *Econometrica*, 82 (5), 1557–1614. [6]
- HELLWIG, CHRISTIAN (2021): “Static and Dynamic Mirrleesian Taxation with Non-separable Preferences: A Unified Approach,” Working paper. [4, 26]
- HÖRNER, JOHANNES, XIAOSHENG MU, AND NICOLAS VIELLE (2017): “Keeping Your Story Straight: Truth-telling and Liespotting,” Tech. rep., Yale University. [32]
- HOSSEINI, ROOZBEH, LARRY E. JONES, AND ALI SHOURIDEH (2013): “Optimal contracting with dynastic altruism: Family size and per capita consumption,” *Journal of Economic Theory*, 148, 1806–1840. [6]
- KAPIČKA, MAREK (2013): “Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach,” *The Review of Economic Studies*, 80 (3), 1027–1054. [4]
- KHAN, AUBHIK, LATCHEZAR POPOV, AND B. RAVIKUMAR (2020): “Enduring Relationships in an Economy with Capital and Private Information,” *FRB St. Louis Working Paper*. [6]
- KHAN, AUBHIK AND BASHKAR RAVIKUMAR (2001): “Growth and risk-sharing with private information,” *Journal of Monetary Economics*, 47 (3), 499–521. [6]
- KOCHERLAKOTA, NARAYANA R. (2010): *The New Dynamic Public Finance*, Princeton, NJ: Princeton University Press. [2]
- KRASIKOV, ILIA AND ROHIT LAMBA (2021): “A Theory of Dynamic Contracting with Financial Constraints,” *Journal of Economic Theory*, 193. [7, 22]
- LUCAS, JR., ROBERT E. (1992): “On Efficiency and Distribution,” *Economic Journal*, 102 (411), 233–247. [1]
- MEGHIR, COSTAS AND LUIGI PISTAFERRI (2004): “Income variance dynamics and heterogeneity,” *Econometrica*, 72 (1), 1–32. [2]
- NEWMAN, ANDREW F. (2007): “Risk-bearing and entrepreneurship,” *Journal of Economic Theory*, 137 (1), 11–26. [14]
- OLSZEWSKI, WOJCIECH AND MIKHAIL SAFRONOV (2021): “Efficient dynamic allocations in endowment economies with linear utilities and private information,” Tech. rep. [6]
- PAVAN, ALESSANDRO (2016): “Dynamic mechanism design: Robustness and endogenous types,” in *2015 World Congress of the Econometric Society*. [4, 26]
- PAVAN, ALESSANDRO, ILYA SEGAL, AND JUUSO TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82 (2), 601–653. [4, 10, 27]
- PHELAN, CHRISTOPHER (1994): “Incentives and aggregate shocks,” *The Review of Economic Studies*, 61 (4), 681–700. [28, 29]

- (1995): “Repeated Moral Hazard and One-Sided Commitment,” *Journal of Economic Theory*, 66, 488–506. [6, 7]
- (1998): “On the Long Run Implications of Repeated Moral Hazard,” *Journal of Economic Theory*, 79, 174–191. [6, 7, 25, 26, 28, 29]
- (2006): “Opportunity and social mobility,” *Review of Economic Studies*, 73 (2), 487–504. [6, 28, 29]
- RAVIKUMAR, B. AND YUZHE ZHANG (2012): “Optimal auditing and insurance in a dynamic model of tax compliance,” *Theoretical Economics*, 7 (2), 241–282. [27]
- RENAULT, JEROME, EILON SOLAN, AND NICOLAS VIELLE (2013): “Dynamic Sender-Receiver Games,” *Journal of Economic Theory*, 148 (2), 502–534. [32]
- RINCÓN-ZAPATERO, JUAN PABLO AND MANUEL S. SANTOS (2009): “Differentiability of the value function without interiority assumptions,” *Journal of Economic Theory*, 144 (5), 1948–1964. [13]
- ROTHSCHILD, MICHAEL AND JOSEPH STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *The Quarterly Journal of Economics*, 90 (4), 629–649. [5]
- SLEET, CHRISTOPHER AND ŞEVİN YELTEKIN (2006): “Credibility and endogenous societal discounting,” *Review of Economic Dynamics*, 9 (3), 410–437. [6]
- (2008): “Politically credible social insurance,” *Journal of monetary Economics*, 55 (1), 129–151. [6]
- SPEAR, STEPHEN E. AND SANJAY SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54 (4), 599–617. [6]
- STORESLETTEN, KJETIL, CHRISTOPHER I. TELMER, AND AMIR YARON (2004): “Consumption and risk sharing over the life cycle,” *Journal of Monetary Economics*, 51 (3), 609–633. [2]
- STRULOVICI, BRUNO (2017): “Contract Negotiation and Screening with Persistent Information,” Tech. rep., Northwestern University. [18]
- (2020): “Renegotiation-proof contracts with persistent private information,” Working paper. [7, 18]
- THOMAS, JONATHAN P. AND TIM WORRALL (1990): “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem,” *Journal of Economic Theory*, 51, 367–390. [2, 3, 4, 5, 6, 7, 8, 9, 13, 14, 16, 17, 18, 19, 20, 21, 24, 26, 27]
- WANG, CHENG (1995): “Dynamic insurance with private information and balanced budgets,” *The Review of Economic Studies*, 62 (4), 577–595. [7]
- WILLIAMS, NOAH (2011): “Persistent Private Information,” *Econometrica*, 79 (4), 1233–1275. [2, 3, 4, 5, 6, 7, 26, 33]
- WILSON, CHARLES (1977): “A Model of Insurance Markets with Incomplete Information,” *Journal of Economic theory*, 16 (2), 167–207. [5]
- ZHANG, YUZHE (2009): “Dynamic Contracting with Persistent Shocks,” *Journal of Economic Theory*, 144, 635–675. [2, 3, 5, 14, 26]