Insurance and Inequality with Persistent Private Information*

Alexander W. Bloedel† R. Vijay Krishna‡ Oksana Leukhina§

11th December 2021

Abstract

This paper studies the implications of optimal insurance provision for long-run welfare and inequality in economies with persistent private information. We consider a model in which a principal insures an agent whose privately observed endowment follows an ergodic, finite Markov chain. The optimal contract always induces *immiseration*: the agent's consumption and utility decrease without bound. Under positive serial correlation, the optimal contract also features *backloaded high-powered incentives*: the sensitivity of the agent's utility with respect to his report increases without bound. These results significantly extend — and elucidate the limits of — the hallmark immiseration results for economies with iid private information. Our analysis utilizes recursive techniques for contracting with persistent states, accounts for the possibility of binding global incentive constraints, extends to other canonical insurance settings (e.g., Mirrleesian economies), and has additional implications for the short-run dynamics of optimal contracts.

JEL Codes: C73; D30; D31; D80; D82; E61

Keywords: Immiseration; insurance; inequality; backloaded incentives; recursive contracts; persistent private information

(*) First version: September 2018. This paper subsumes an earlier working paper circulated as "Misery, Persistence, and Growth" (Bloedel and Krishna 2015). We are grateful to Gabriel Carroll, Hector Chade, V.V. Chari, Sebastian Di Tella, Ed Green, Pablo Kurlat, Andy Newman, Shunya Noda, Andy Skrzypacz, Bruno Strulovici, Bob Wilson, and especially Paul Milgrom and Ilya Segal for insightful comments. We also thank audiences at Stanford, Virginia Tech, the 2018 Econometric Society NASM, the 2019 Midwest Macro Meetings, the 2021 SED Annual Meetings, and the 2021 Econometric Society ESM for useful feedback.

(†) Caltech and UC Davis. Email: abloedel@caltech.edu

(‡) Florida State University. Email: rvk3570@gmail.com

(§) Federal Reserve Bank of St. Louis. Email: oksana.m.leukhina@gmail.com

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 income and wealth inequality center on institutions designed to facilitate risk-sharing, such as social insurance programs and redistributive tax systems. This is only natural: because individuals typically have private information about their idiosyncratic risks — such as shocks to income, tastes, productivity, or employment — and must be incentivized to reveal that information, the markets for insuring against such risks are inevitably incomplete. Consequently, the manner in which society resolves the tradeoff between providing incentives and delivering partial insurance determines the extent to which these shocks are propagated, and perhaps even amplified, over time. What, then, is the *optimal* degree of inequality?

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). Consequently, cross-sectional consumption and wealth inequality increase without bound, and 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

(1) For example, various authors have argued that immiseration may be counterintuitive or unappealing on normative grounds (as it seems perverse that *ex ante* efficiency should require *ex post* destitution), descriptive grounds (because commitment to one's own impoverishment may not be enforceable), and practical grounds (because the absence of a well-defined steady-state hampers any meaningful study of the tradeoff between equity and efficiency). Thus, a sizable literature attempts to justify bounded long-run inequality by relaxing the contracting parties' commitment power through participation constraints for the agents (Atkeson and Lucas 1995; Phelan 1995) and credibility constraints for the planner (Sleet and Yeltekin 2006, 2008; Farhi et al. 2012), or by considering alternative normative criteria while maintaining full commitment (Phelan 2006; Farhi and Werning 2007). See Subsection 5.4 for discussion of other model variants.

workhorse normative models that economists use to study optimal risk-sharing (e.g., in the large literatures on dynamic taxation and unemployment insurance). It is therefore important to understand the robustness of the forces that drive it within that 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 iid over time (i.e., shocks are "completely transient"), which is analytically convenient but unrealistic: shocks to individuals' incomes, tastes, productivities, and health or employment statuses are often highly persistent.[2] At the opposite extreme, in an influential paper, Williams (2011) assumes that private information follows a random walk (i.e., shocks are "permanent") and finds, strikingly, that immiseration not only disappears, but the optimal contract generates long-run *bliss* by sending individuals' consumption and utility to their *upper* bounds.[3] This raises two questions: First, what is driving these diametrically opposing results? Second, what kinds of long-run properties arise in the (generic and empirically relevant) intermediate case in which private information is persistent, but only imperfectly so (i.e., shocks are "partially transient")?

In this paper, we investigate the long-run properties of optimal insurance contracts under general forms of persistent private information. We address the second question above by showing that immiseration arises for a broad class of *ergodic* private information processes, allowing us, in essence, to fully interpolate between the iid and random walk benchmarks. Our analysis also sheds light on the first question by identifying mean-reversion of the private information process as a key determinant of the optimal contract's long-run properties and, in particular, suggesting that immiseration fails (as in Williams (2011)) only in the knife-edge case of "permanent" shocks.

Model. We study a canonical principal-agent version of the planning problem.[4] A risk-neutral principal (she) interacts with a risk-averse agent (he) over an infinite horizon in

(2) See Storesletten, Telmer and Yaron (2004) and Meghir and Pistaferri (2004) for evidence on individual-level labor earnings.

(3) See Subsection 5.1 for further discussion of Williams (2011) and its relation to the present paper.

(4) It is well known that the planning problem in which a utilitarian planner maximizes the average welfare of a large population of agents subject to an intertemporal (Green 1987) or per-period (Atkeson and Lucas 1992) resource constraint, 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 (e.g., Golosov, Tsyvinski and Werquin 2016). Single-agent immiseration results in the principal-agent setting, which concern the *level* of the agent's utility and consumption, translate to unbounded cross-sectional *inequality* in the many-agent planning version of the problem.

discrete time. The agent's stage preferences 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 at the same rate. For concreteness, we focus on a baseline model in which the agent's type corresponds to his privately observed endowment (e.g., Green 1987; Thomas and Worrall 1990), but this is not essential.[5]

We make two main assumptions about the environment (in addition to standard assumptions about the agent's risk preferences). First, the agent's type evolves according to a fully connected, finite-state Markov chain. This requires that the agent's type process be ergodic and bounded, but allows for otherwise arbitrary positive and negative serial correlation, including the kinds of asymmetric and skewed shocks documented in the recent empirical literature,[6] and is flexible enough to either nest or approximate most type processes considered in the theoretical literature.[7] Second, we assume that the agent can neither covertly save outside of the contract nor lie within the contract by overstating his type, capturing the idea that the agent interacts exclusively with the principal and is therefore unable to engage in trade or production outside of the contract.[8]

(5) As discussed in Subsection 5.2, our main analysis extends to other canonical insurance settings, such as taste-shock models in which the agent's type corresponds to his privately observed marginal utility of consumption (e.g., Atkeson and Lucas (1992); Farhi and Werning (2007)) and Mirrleesian models in which the agent's type corresponds to his privately observed labor productivity (e.g., Zhang (2009); Farhi and Werning (2013)).

(6) For instance, Arellano, Blundell and Bonhomme (2017) and Guvenen et al. (2021) document that individual labor earnings exhibit significant departures from the standard log-normal specifications, emphasizing that the relative persistence of positive and negative income shocks depends sensitively on the individual's current income level.

(7) This class includes all finite iid processes (as in the classic literature) and binary-state Markov chains (as in much of the recent contract theory literature; see Subsection 5.4), and allows for discrete-state (distributional) approximations of any ergodic Markov process, including those with unbounded and continuous state spaces (as in much of the recent theoretical dynamic taxation literature, e.g., Farhi and Werning (2013); Golosov, Troshkin and Tsyvinski (2016)). While the random walk processes studied in Williams (2011) (see also Bloedel, Krishna and Strulovici 2020, 2021) are not ergodic, their finite-dimensional distributions can be similarly approximated by those of the processes considered here.

(8) The former restriction on savings is standard (cf. Allen (1985); Cole and Kocherlakota (2001b)). The latter restriction is widespread in the literature and natural in many environments: when the agent's private information concerns his endowment, it amounts to the assumption that he cannot covertly borrow outside of the contract (Williams (2011); Bloedel, Krishna and Strulovici (2020)), and when his private information concerns his labor productivity, it amounts to the assumption that he cannot covertly engage in home production (Golosov and Tsyvinski (2007)). Relaxing these assumptions from the outset would

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 sometimes almost surely) to their (possibly infinite) lower bounds. This generalizes most known immiseration results in the contracting literature.[9] Our second result, Theorem 2, shows that — at least 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 report increases without bound. This can be viewed as a kind of *relative immiseration* wherein the welfare 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 $T \to \infty$.

A simple intuition for these results and the relation between them 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 when the agent's *marginal* utility of consumption is *higher*, so that a given variability of utility can be induced by smaller variations in consumption. By immiserating the agent, the principal therefore drives her cost of incentive provision to zero in the long run, making it affordable to provide arbitrarily high-powered incentives in later periods. Such *backloading* of high-powered incentives allows the principal to reduce her costs (by better smoothing the agent's consumption) in early periods while maintaining incentive compatibility. This is optimal for the principal because it allows her to smooth costs over time.

While similar intuitions have been put forth in the classical iid setting (e.g., Thomas and Worrall 1990), converting them into a formal argument is significantly more challenging

conflate constraints arising from the agent's exogenous private information with additional constraints arising from the principal's limited ability to enforce contractual terms. See Subsection 5.2 for further discussion.

(9) For instance, the specialization of Theorem 1 to "pseudo-renewal" type processes subsumes Green's (1987) and Thomas and Worrall's (1990) iid immiseration results and Zhang's (2009) immiseration result for binary-type, symmetric-transition Markovian types. Moreover, as discussed in Subsection 5.4, the *proof* of Theorem 1 can easily be adapted to generalize related long-run convergence results in some other dynamic contracting environments.

when the agent's private information is persistent. The key subtlety is that persistence gives rise to a new source of information rents for the agent: in addition to the traditional *iid information rent* arising from his private information about his current type, there is a *Markov information rent* arising from his private information about the distribution of his future types. Specifically, in the iid case, the principal screens the agent by manipulating (only) his iid information rent, which is determined by the quality of risk-sharing in the current period. But with persistence, the principal may also screen the agent by manipulating his Markov information rent — in particular, by offering inefficient (and hence non-renegotiation-proof) continuation contracts that are assigned different continuation values depending on the agent's private belief about his future types.

Thus, persistence gives rise to a "horse race" between two distinct forms of screening based on, respectively, the agent's privately known intra- and inter-temporal preferences over contracts. If the agent's type is sufficiently persistent, so that his Markov information rents are "sufficiently large," then the latter screening channel may dominate and generate very different contractual dynamics. A key insight emerging from our analysis is that, despite the presence of this new screening-through-continuation-contracts channel, the familiar backloaded incentives channel remains dominant — and renders immiseration optimal — given *any* amount of mean-reversion in the agent's type process.

This insight helps to resolve open questions about the robustness (or fragility) of immiseration. In particular, as discussed above, the classic literature finds that immiseration arises very generally when the agent's type is iid over time (shocks are "perfectly transient"), while more recent work (e.g., Williams 2011) finds that it fails when the agent's type follows a random walk (shocks are "permanent"). By characterizing the generic case in which types are imperfectly, but otherwise arbitrarily, persistent (shocks are "partially transient"), our results demonstrate that the failure of immiseration in the random walk case is a knife-edge phenomenon: as discussed in Subsection 5.1, when shocks are "permanent" — and *only* then — the agent's Markov information rents may be so large that it is impossible to insure him at all, giving rise to a very different set of tradeoffs for the principal and a qualitatively different optimal contract.

Techniques. Our formal analysis is based on a recursive formulation of the contracting problem in which the main state variable for the principal is a vector of *interim* promised utilities, which encode the agent's continuation utility contingent on each possible current type.[10] This allows us to represent the agent's incentive constraints recursively and characterize the optimal contract using dynamic programming techniques.

Our main object of study is a particular directional derivative of the principal's value function with respect to the interim promised utility state, which we call her *marginal cost martingale*.[11] To understand this, recall from Thomas and Worrall (1990) that, in the special case of iid types, the principal's cost-smoothing dictates that, at the optimum, her marginal cost of providing promised utility to the agent follows a martingale process and, moreover, that convergence properties of this martingale can be used to pin down the optimal contract's long-run behavior. The proof of Theorem 1 builds on this idea, but requires new techniques to deal with several complications caused by persistence. For instance, persistence generally gives rise to a non-monotone Pareto frontier: increasing the agent's interim promised utility can decrease his Markov information rents, thereby decreasing the principal's cost. This makes it challenging to formulate the appropriate notion of the principal's marginal cost; we identify the essentially unique way of doing so, which involves varying the agent's promised utility while holding the schedule of information rents fixed. Even then, discerning properties of the optimal contract from the marginal cost martingale's behavior is subtle because the optimal contract is not renegotiation-proof and exhibits rich history-dependence; we identify a "renewal" property of the martingale's dynamics that allows us to establish long-run convergence using novel probabilistic arguments.

Notably, our analysis also overcomes two well-known difficulties that arise more generally in contracting problems with persistent types. The first is the *curse of dimensionality*: because the dimension of the promised utility state in our and related recursive formulations (e.g., Fernandes and Phelan 2000) scales with the number of possible types, prior work based on such techniques has mostly been confined to settings with binary types or relied on numerical simulations. By contrast, we analytically derive general properties of the optimal contract (Theorems 1 and 2) and explicitly solve for the full "recursive domain" of implementable promised utility vectors (Theorem 3) while allowing for any finite number of types and a wide range of serial correlation structures. In doing so, we do not rely on the popular first-order approach in which one first solves a relaxed problem incorporating

(10) See Subsection 5.4 and Appendix B for comparison to the recursive formulation of Fernandes and Phelan (2000).

(11) As described in Subsection 5.2, this martingale reduces to the celebrated "inverse Euler equation" (e.g., Golosov, Kocherlakota and Tsyvinski 2003) when the agent's preferences satisfy particular separability conditions, which, however, are not satisfied in our baseline hidden endowment model.

only the agent's local incentive constraints and subsequently verifies that the candidate solution is actually feasible.[12] While convenient, that approach has limited scope due to the possibility of *binding global incentive constraints*: without stringent assumptions on the agent's type process, global incentive constraints are liable to bind at the true optimum, rendering the first-order approach invalid.[13] We instead account for the possibility of binding global incentive constraints by directly tackling the principal's full problem, which has been identified as an important goal for the dynamic contracting literature (Pavan 2016; Garrett, Pavan and Toikka 2018; Battaglini and Lamba 2019). Many aspects of our analysis are not specific to the insurance problems on which we focus, and therefore some of the techniques utilized here may be of broader interest.

Roadmap. The rest of the paper is organized as follows. Section 2 presents the baseline hidden endowment model. Section 3 develops the recursive formulation of the contracting problem. In Section 4, we present our main long-run results and sketch their proofs. Section 5 discusses implications of our main results, their extension to several variants of the baseline model, short-run properties of the optimal contract, and related literature. In Appendices A–B, we present structural results that are important for the recursive formulation but omitted from the main text. The Online Appendices C–F present self-contained proofs of our main results. Proofs of some auxiliary technical results and details for a solved example are in the working paper (Bloedel, Krishna and Leukhina 2021, henceforth BKL21).

2. Baseline Model

2.1. Environment

A risk-neutral principal (she) provides insurance to a risk-averse agent (he). Time is discrete and runs over an infinite horizon, indexed by $t \geq 0$. We begin by describing the primitives of the environment. The assumptions DARA, NHB, and Markov stated below hold for the

(12) The first-order approach is developed for general dynamic mechanism design problems by Pavan, Segal and Toikka (2014) and applied to problems of dynamic insurance and taxation by, e.g., Williams (2011), Kapička (2013), Farhi and Werning (2013), and Golosov, Troshkin and Tsyvinski (2016).

(13) For instance, Battaglini and Lamba (2019) study a canonical monopolistic screening model and show that the first-order approach is generically invalid when the agent's type follows a fully connected, finite-state Markov chain (as in the present paper) and period lengths are short. They also show by example that, when the first-order approach fails, the candidate optimal contract derived from the relaxed problem may exhibit qualitatively different features from the true optimal contract. While their analysis does not directly apply to the insurance models studied here, it does suggest that binding global incentive constraints are likely to remain an issue.

remainder of paper.

Preferences. Both the principal and agent discount the future at common rate $\alpha \in (0,1)$. The agent has utility function over consumption $U : (\underline{c}, \infty) \to \mathbb{R}$, where $\underline{c} \geq -\infty$. Let $\mathcal{U} = U\big((\underline{c}, \infty)\big) \subset \mathbb{R}$ denote the range of feasible utilities. We make the following assumptions on the utility function.

Assumption 1 (DARA)**.** $U(\cdot)$ satisfies the following properties:

(a) It is strictly increasing, strictly concave, continuously differentiable, and satisfies the Inada conditions $\lim_{c \to \underline{c}} U'(c) = +\infty$ and $\lim_{c \to \infty} U'(c) = 0$.

(b) It is bounded above and unbounded below: $\lim_{c \to \infty} U(c) = 0$ and $\lim_{c \to \underline{c}} U(c) = -\infty$, so that $\mathcal{U} = (-\infty, 0)$.

(c) It has *decreasing absolute risk aversion* (DARA): the mapping $c \mapsto -\log(U'(c))$ is (weakly) concave.

Assumption DARA is fairly weak and common in the literature (e.g., Thomas and Worrall 1990; Cole and Kocherlakota 2001b). For example, it is satisfied by CARA utility, which is the benchmark specification in the literature, and also by many utility functions in the HARA class.[14] Part (a) of Assumption 1 is ubiquitous in models of risk-sharing. The latter two parts serve to simplify aspects of the subsequent analysis: part (b) implies that the range of feasible utility levels $\mathcal{U}$ is an open set, and part (c) ensures that various constraint sets (defined in later sections) are convex. Looking ahead, together they imply that optimal contracts are interior and fully characterized by first-order conditions (see Subsection 5.2).

Information. In every period, the agent receives a random endowment of $\omega_i \in \mathbb{R}$, where $i \in S := \{1, \dots, d\}$ and $\omega_d > \omega_{d-1} > \cdots > \omega_1$. We say that the agent is of *type* $i \in S$ when his current endowment is ω_i. The principal does not observe the agent's endowments and must rely on his reports.

Assumption 2 (NHB)**.** The agent cannot over-state his endowment in any period. That is, an agent of type i can only claim to be types $j \leq i$.

(14) For instance, CARA utility is assumed in Green (1987), Phelan (1995), and Phelan (1998), as well as in the main solved examples in Thomas and Worrall (1990), Wang (1995), Williams (2011), Bloedel, Krishna and Strulovici (2020), and Strulovici (2020). Assumption DARA is also satisfied by all HARA utility functions of the form $U(c) = (c+\eta)^{1-\gamma}/(1-\gamma)$ with $\gamma > 1$ and $\eta \geq -\underline{c}$, which includes CRRA utilities corresponding to $\eta = \underline{c} = 0$.

Assumption NHB, which stands for *No Hidden Borrowing*, is motivated by the ideas that (i) endowments are partially verifiable and (ii) the agent does not have access to a market (or private storage technology) outside of his relationship with the principal. For example, before receiving any transfers from the principal, the agent might be required to deposit his reported endowment in an account that the principal can monitor. If the agent is not able to unilaterally store the consumption good across periods or borrow units of the consumption good without the principal's knowledge, then he can deposit at most his true endowment. NHB is commonly assumed in the literature and natural in the present context (see Subsection 5.2).

Type Process. Aside from finiteness of the type space, we make just one substantive assumption on the *endowment (or type) process*, which is denoted by $(\omega^{(t)})_{t\in\mathbb{N}}$.[15]

Assumption 3 (Markov)**.** The agent's endowment follows a first-order, fully connected Markov process with transition probabilities $\mathbf{P}(\omega^{(t+1)} = j \mid \omega^{(t)} = i) = f_{ij} > 0$, where $\mathbf{P}$ is the probability measure on S^∞ that induces this process.

Note that Assumption Markov implies that the type process is ergodic and bounded, but allows for essentially arbitrary (imperfect) positive and negative serial correlation. The role of ergodicity is discussed further in Subsection 5.1. Going forward, we represent transition probabilities as a $d \times d$ transition matrix with rows $\mathbf{f}_i = (f_{i1}, \ldots, f_{id})$, where $\mathbf{f}_i$ denotes the probability vector over tomorrow's type if today's type is $i \in S$.

Remark 2.1. While, for expositional simplicity, we assume here that the agent's private information concerns his endowment, our analysis extends to other canonical insurance settings in which the agent's private information concerns his tastes, as in Atkeson and Lucas (1992), or his labor productivity, as in Mirrleesian models of optimal dynamic taxation (Kocherlakota 2010). These model variants are discussed in Subsection 5.2.

2.2. Principal's (Sequential) Problem

We study constrained-efficient risk-sharing schemes. At the initial date, $t = 0$, the principal offers the agent a long-term insurance contract that specifies transfers of the consumption good as a function of all reported shocks, past and present. (The restriction to direct truthful mechanisms is, as usual, without loss by a version of the Revelation Principle.) By entering

(15) Throughout, we adopt the convention that $0 \in \mathbb{N}$.

the contract at $t = 0$, both parties fully commit to its terms at all future dates; in particular, neither party is allowed to renege later on. The principal's objective is to minimize her expected lifetime costs, given some (possibly degenerate) prior belief over the agent's initial type, and subject to (i) delivering a pre-specified schedule of promised utilities $\mathbf{v}^{(0)} := (v_1, \ldots, v_d) \in \mathcal{U}^d$ to the agent[16] and (ii) providing appropriate incentives for truthtelling. The promised utility v_i is interpreted as the lifetime utility promised to the agent when his true initial type is $\omega^{(0)} = \omega_i$ and conditional on truthtelling.

We refer to the contracts informally described above as *sequential contracts* because they are defined via the full sequence of history-dependent transfers. Similarly, we refer to the principal's optimization problem over sequential contracts as her *sequential problem*. This is the standard way of defining the dynamic contracting problem but is intractable in our setting, which features both imperfectly transferable utility (due to the agent's risk aversion) and persistent private information. For brevity, the standard but notationally cumbersome details of the sequential formulation are relegated to Appendix A.1, where the principal's sequential problem is formally stated as the program [**SP**]. We now proceed directly to a more tractable recursive formulation of the contracting problem, which will serve as the basis for all subsequent analysis.

3. Recursive Contracts

This section consists of several parts. Subsection 3.1 describes the recursive formulation of the principal's problem. Subsection 3.2 introduces regularity conditions imposed in the subsequent analysis. Subsection 3.3 presents a Bellman equation that generates optimal contracts.

3.1. Principal's Recursive Problem

State variable. When the type process is iid, Green (1987) and Thomas and Worrall (1990) show that the principal's problem can be formulated as a Markov decision process using the agent's *ex ante* promised utility — i.e., his lifetime expected utility starting from a given history *before his current type is realized*, assuming he is truthful in the current and all future periods — as a state variable. However, when types are serially correlated, *ex ante* promised utility is not a sufficient state variable because the agent's *true* type determines both his current marginal utility of consumption *and his beliefs about future endowment realizations*,

(16) This initial condition may be given exogenously, or optimally chosen as in the efficiency problem [**Eff**$_i$] described below in Subsection 4.3.

the latter of which determines his preferences over continuation contracts. Importantly, with persistence, this latter aspect of the agent's preferences is also his private information. To ensure incentive compatibility, the recursive formulation therefore requires keeping track of the agent's private beliefs so that the principal can use continuation contracts as a screening device. Since Assumption Markov implies that the agent's beliefs are determined by his current endowment realization, it suffices to keep track of the latter.

In particular, recall that at time $t = 0$ the principal promises the agent with initial type $i \in S$ exactly $v_i \in \mathcal{U}$ lifetime utiles, as summarized by the vector $\mathbf{v}^{(0)} = (v_1, \ldots, v_d)$ of type-contingent promised utilities. The principal's recursive problem uses the pair $(\mathbf{v}, s)$ of *contingent* (or *interim*) promised utilities and the previously-reported type s as state variables. Importantly, v_i is the lifetime utility promised to a type-i agent *assuming he reports truthfully going forward.*[17] As we shall see below, keeping track of the vector $\mathbf{v}$ allows for a simple recursive formulation of the agent's incentive constraints, while keeping track of the previous report s is needed to compute the principal's expected continuation costs.

Recursive constraints. Given a state $(\mathbf{v}, s)$, the principal offers the agent a *menu* $(u_i, \mathbf{w}_i)_{i \in S} \in \left(\mathcal{U} \times \mathcal{U}^d\right)^d$, where u_i is the agent's flow utility and $\mathbf{w}_i$ is the vector of contingent continuation utilities promised to the agent if he reports his current type to be $i \in S$.[18] When the agent reports to be of type i, the principal's state variable transitions to $(\mathbf{w}_i, i)$ in the next period. An agent of type i who claims to be of type j gets the flow utility $\psi(u_j, i, j)$, where $\psi : \mathcal{U} \times S \times S \to \mathcal{U}$ is defined as

$$\psi(u, i, j) := U\big(\omega_i + C(u, j)\big)$$

where $C(u, j) := U^{-1}(u) - \omega_j$ is the amount of consumption that the principal must transfer to the agent in order to deliver exactly u flow utiles to a type-j agent.

Clearly, a menu should (i) deliver the appropriate promised utility to each agent type, and (ii) ensure that reporting truthfully is optimal for the agent at any date, assuming he

(17) We will always use s to denote the *previous period's* type, while indices i, j, k denote the *current period's* type. Thus, $\omega^{(t)} = \omega_i$ if and only if $s^{(t+1)} = i$. With a slight abuse of terminology, we will also refer to the stochastic process $(s^{(t)})_{t=0}^{\infty}$ is the *type process* despite this timing discrepancy.

(18) The promised utilities in the vector $\mathbf{w}_i$ are contingent on the agent's *subsequent* type. For instance, the j^{th} element of $\mathbf{w}_i$, denoted by w_{ij}, represents the agent's promised utility starting tomorrow if his type today is i and his type tomorrow is j.

reports truthfully in the future. These *recursive constraints* are the *promise keeping* and *incentive compatibility* conditions

$$[\mathbf{PK}_i] \qquad v_i = u_i + \alpha\, \mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i]$$

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

where $\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_j] := \sum_{k=1}^{d} f_{ik} w_{jk}$ is the *expected* promised utility for an agent whose current type is i but who reports the type j. (We refer to $\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i]$ as *ex ante* promised utility for type i.) We require that [$\mathbf{PK}_i$] holds for all $i \in S$ and, by Assumption NHB, that [$\mathbf{IC}_{ij}$] holds for all $i, j \in S$ with $i > j$.

Two aspects of the recursive constraints are noteworthy. First, they are independent of s, the agent's previous report. Second, it is evident from [$\mathbf{IC}_{ij}$] that, even if the agent lies today, his expectation over tomorrow's type is still governed by his true current type. In this way, the principal can incentivize truthful revelation in the current period regardless of the agent's previous history of actual and reported types.[19] This recursive formulation therefore solves the issue of the agent's private preferences over continuation contracts.

Implementability. The next step in the recursive formulation is to specify which $\mathbf{w}_i$'s are feasible for the principal to offer, i.e., which promised utility vectors are *implementable*. To illustrate, consider the special case in which $d = 2$ and suppose the agent's type is positively serially correlated (i.e., $\mathbf{f}_2$ first-order stochastically dominates $\mathbf{f}_1$, which entails $f_{11} \geq f_{21}$ and thus $f_{22} \geq f_{12}$). Subtracting both sides of the promise keeping constraint [$\mathbf{PK}_i$] ($i = 1$) from the incentive constraint [$\mathbf{IC}_{ij}$] (setting $j = 1$ and $i = 2$) yields

$$[\mathbf{IC}^*_{21}] \qquad v_2 - v_1 \geq \underbrace{\psi(u_1, 2, 1) - u_1}_{\text{iid info rent}} + \underbrace{\alpha\big[\, \mathsf{E}^{\mathbf{f}_2}(\mathbf{w}_1) - \mathsf{E}^{\mathbf{f}_1}(\mathbf{w}_1)\big]}_{\text{Markov info rent}}$$

The iid information rent is clearly positive, for an agent with endowment ω_2 achieves larger flow utility than one with ω_1 from the given transfer $C(u_1, 1)$, i.e., $\psi(u_1, 2, 1) = U\big(\omega_2 + C(u_1, 1)\big) > u_1$. The Markov information rent is also positive *if* the contingent

(19) More precisely, the recursive formulation requires that truthtelling be optimal for the agent even if he had lied in the past, which is in principle stronger than the requirement that he finds it optimal to truthfully report *conditional on truthtelling in the past*. However, since the agent's type is Markovian and payoffs are time-separable, these two requirements are in fact equivalent (cf. Pavan, Segal and Toikka 2014, Section 3.3).

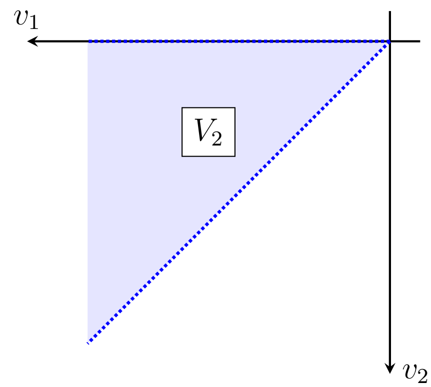

Figure 1: *The largest recursive domain $D = V_2$ in the $d = 2$ and persistent case.*

continuation utility $\mathbf{w}_1 \in V_2 := \{\mathbf{v} \in \mathcal{U}^2 : v_1 < v_2\}$ (see Figure 1). Consequently, if $\mathbf{w}_1 \in V_2$ then we must also have $\mathbf{v} \in V_2$. Iterating on this idea — essentially by noting that $\mathbf{w}_1$ must also have some implementation — it can be shown that only $\mathbf{v} \in V_2$ are implementable. The following definition generalizes this idea:

Definition 3.1. A set $D' \subseteq \mathcal{U}^d$ is a *recursive domain* if, for every $\mathbf{v} \in D'$ there is a menu $(u_i, \mathbf{w}_i)_{i \in S} \in \Gamma(\mathbf{v})$ satisfying the recursive constraints such that $\mathbf{w}_i \in D'$ for all $i \in S$. A set in $\mathcal{U}^d$ is the *largest recursive domain* if it (i) contains every recursive domain, and (ii) is itself a recursive domain. The largest recursive domain is denoted by D if it exists (in which case it is unique).

The largest recursive domain D (henceforth, simply *the domain*) characterizes the implementable promised utility vectors: $\mathbf{v}$ should be considered feasible if and only if $\mathbf{v} \in D$. In Theorem 3 of Appendix B, we show that a largest recursive domain exists and characterize its properties.[20] Consistent with the informal argument given above, one special case of this theorem implies that $D = V_2$ when $d = 2$ and the agent's type is positively serially correlated.

Principal's recursive problem. A *recursive contract* is a map $\xi : D \times S \times S \to \mathcal{U} \times D$, written as $\xi = \left(\xi^f, \xi^c\right)$, where $\xi^f(\mathbf{v}, s, i) = u_i(\mathbf{v}, s) \in \mathcal{U}$ provides *flow* consumption utiles

(20) The domain is an essential piece of the solution to the principal's problem: it determines the constraints imposed by incentive compatibility and hence is implicitly part of the definition of the principal's value function P (defined in [**RP**] and [**FE**] below). Notably, different recursive formulations of the contracting problem would give rise to different domains, which may affect the tractability of the analysis (e.g., see footnote 50 in Appendix B for an explicit comparison to Fernandes and Phelan's (2000) recursive formulation).

to i-type reports, and $\xi^c(\mathbf{v}, s, i) = \mathbf{w}_i(\mathbf{v}, s) \in D$ similarly provides *contingent continuation* utiles. We say that ξ is *feasible at* $(\mathbf{v}, s) \in D \times S$ if the menu $\big(\xi(\mathbf{v}, s, i)\big)_{i \in S} \in \Gamma(\mathbf{v})$, where the principal's *constraint correspondence* $\Gamma : D \rightrightarrows (\mathcal{U} \times D)^d$ is defined by

[**3.1**]
$$\Gamma(\mathbf{v}) := \big\{(u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times D)^d : (u_i, \mathbf{w}_i)_{i \in S} \text{ satisfies [}\mathbf{PK}_i\text{] } \forall\, i \in S \text{ and [}\mathbf{IC}_{ij}\text{] } \forall i, j \in S \text{ with } i > j\big\}$$

Naturally, ξ is *feasible* if it is feasible at all $(\mathbf{v}, s) \in D \times S$. Let $\Xi(\mathbf{v})$ denote the set of feasible recursive contracts that are initialized at $\mathbf{v} \in D$. Note that every $\mathbf{v} \in D$ and $\xi \in \Xi(\mathbf{v})$ together induce stochastic processes $\tilde{u}_\xi := \big(u_\xi^{(t)}\big)_{t=0}^{\infty}$, which we call the *induced allocation*, and $(\mathbf{v}_\xi^{(t)})_{t=1}^{\infty}$, which we call the *induced promises*.

The principal's *recursive problem* is to choose the recursive contract that minimizes the lifetime expected cost of the induced allocation, subject to the recursive constraints at each step:

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

Note that the expectation in [**RP**] is taken with respect to the true probability measure over paths of types, i.e., the measure over reported paths induced by truthful reporting. Conditioning on the event $s^{(0)} = s$ denotes that the principal has the "prior" f_s over the initial $t = 0$ type. A recursive contract ξ^* is *optimal* if it attains the infimum in [**RP**].

Transversality. Since the agent's utility function is unbounded below, *a priori* it is possible that the principal's recursive problem [**RP**] is a relaxation of her sequential problem [**SP**] because the recursive constraints only impose one-step promise keeping and deter one-step deviations from truthtelling. While standard induction implies that they also impose finite-step promise keeping and deter finite-step deviations, we also need to ensure that the induced *promises* do not grow too fast, so that the induced *allocation* actually (i) delivers the appropriate level of promised utility and (ii) also deters deviation strategies involving infinitely-many misreports.[21] Any recursive contract violating these conditions effectively involves the principal running a Ponzi scheme.

(21) Desideratum (i) may be formally stated as $v_j = \mathsf{E}\left[\sum_{t=0}^{\infty} \alpha^t \tilde{u}_\xi^{(t)} \mid s^{(0)} = s_j\right]$ for all $j \in S$. Desideratum (ii) effectively means that the agent's reporting problem is "continuous at infinity."

To that end, let $\mathscr{H} := S^\infty$ denote the space of all infinite sequences, or *paths*, of types with generic element $h \in \mathscr{H}$. We say that ξ satisfies *agent transversality at* $\mathbf{v} \in D$ if, starting from $\mathbf{v}$, the induced discounted promises satisfy

$$\text{[TVC]} \qquad \lim_{t\to\infty} \inf_{h\in\mathscr{H}} \alpha^t \mathbf{v}_\xi^{(t)}(h) = 0$$

where $(\mathbf{v}_\xi^{(t)}(h))_{t=0}^\infty$ denotes the (deterministic) sequence of contingent promises along the path $h \in \mathscr{H}$. Any feasible recursive contract ξ that satisfies [**TVC**] at $\mathbf{v} \in D$ is said to be [**TVC**]-*implementable* at $\mathbf{v} \in D$. Lemma A.1 in Appendix A.2 shows that [**TVC**]-implementable contracts indeed satisfy the two desiderata stated above.[22]

Full-Information Benchmark. For future reference, we briefly describe the first-best contract that arises when there is no private information. In the absence of incentive constraints, every promised utility vector $\mathbf{v} \in \mathscr{U}^d$ is implementable and it is optimal to provide full insurance (see Online Appendix E for details). The first-best contract perfectly smooths the agent's consumption over time and across states so that, conditional on his initial type, the agent's flow utility process $(u^{(t)})_{t=0}^\infty$ is constant. In terms of the recursive variables, this means that the optimal full information contract induces a promised utility process $(\mathbf{v}^{(t)})_{t=0}^\infty$ such that, for $t \geq 1$ and along every path, (i) $\mathbf{v}^{(t)} = \mathbf{v}^{(t+1)}$ and (ii) $v_1^{(t)} = \cdots = v_d^{(t)}$. The principal's value function in the full information problem, which will be referenced below, is denoted $Q^* : \mathscr{U}^d \times S \to \mathbb{R}$.

3.2. Regularity Conditions

To ensure that the optimization problem in [**RP**] is sufficiently well-behaved, we require that the environment satisfy a few regularity conditions. For expositional simplicity, we impose a small set of conditions directly on derived objects, which can be re-stated in terms of model primitives on a case-by-case basis.

Definition 3.2. The environment is *Regular* if Conditions R.1–R.3, stated below, all hold. The environment is [**TVC**]-*Regular* if it is regular and, in addition, satisfies Condition R.4.

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

(22) The standard approach in the literature (e.g., Atkeson and Lucas (1992)) is to build [**TVC**]-implementability into the definition of feasible recursive contracts, even though it is not itself a recursive condition. We find it more convenient to impose [**TVC**]-implementability as a separate requirement; see Condition R.4 and the associated discussion below.

R.2 *(Value Continuity)* For any $\mathbf{v} \in D$ and recursive contract $\xi \in \Xi(\mathbf{v})$, the first-best value function Q^* satisfies

$$\liminf_{t\to\infty} \alpha^t \left[\inf_{h\in\mathscr{H}} Q^*\big(\mathbf{v}_\xi^{(t)}(h), s^{(t)}(h)\big)\right] \geq 0$$

R.3 *(Constraint Qualification)* Let $\Gamma_\circ(\mathbf{v}) \subseteq \Gamma(\mathbf{v})$ denote the set of all menus that are feasible at $\mathbf{v}$ and, in addition, satisfy all of the incentive compatibility constraints [$\mathbf{IC}_{ij}$] $(i > j)$ as *strict* inequalities. For each $\mathbf{v} \in D$, $\Gamma_\circ(\mathbf{v}) \neq \varnothing$.

R.4 *([**TVC**] Existence)* There exists an optimal contract ξ^* that is [**TVC**]-implementable at $\mathbf{v}^{(0)}$, the initial condition for the promised utility process.

Regularity (Conditions R.1–R.3) is a mild technical requirement that allow us to establish basic properties of the principal's recursive problem [**RP**]. Conditions R.1 and R.2 ensure that her value function is well-defined and that an optimal contract exists. Condition R.3 is a standard sufficient condition for the existence of Lagrange multipliers, allowing us to use Lagrangian methods to characterize optimal contracts. In Appendix B.2, we present conditions on the primitives that are sufficient for Regularity. For example, the environment is Regular when the agent has CARA utility and the type process satisfies standard notions of positive serial correlation.

The additional Condition R.4 needed to achieve [**TVC**]-Regularity warrants further comment. Formally, it guarantees that *some* optimal contract in fact delivers promises to, and is incentive compatible for, the agent. Were this condition to fail, the principal's recursive problem [**RP**] could be a strict relaxation of her sequential problem [**SP**]. By assuming Condition R.4, we follow the standard convention in the literature on dynamic contracts (and stochastic control more broadly) to impose (non-recursive) transversality-type constraints on the principal's recursive problem that guarantee the equivalence of [**RP**] and [**SP**].[23]

(23) In the standard approach, one ignores such constraints when solving for a candidate optimum and then verifies that they are satisfied *ex post*. This typically requires solving for the optimal contract in semi-closed form, which we are unable to do in our general setting. Nonetheless, Condition R.4 is known to hold under CARA utility in the two limiting cases where types are either iid or follow a random walk (i.e., are subject to "permanent shocks"), and we have no reason to expect that it will fail in the intermediate cases considered in this paper. For the iid case, see Green (1987, Section 10), Thomas and Worrall (1990, Section 7), and Atkeson and Lucas (1992, Section 6). For the random walk case, see Bloedel, Krishna and Strulovici (2020, 2021).

3.3. Bellman Equation

We conclude this section by observing that the principal's value function P satisfies a familiar Bellman equation that generates optimal contracts, and which will be useful for outlining the proofs of our main results in Section 4 below.

Proposition 3.3. Suppose the environment is Regular. Then the principal's value function $P : D \times S \to \mathbb{R}$ satisfies the functional equation

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

and is the pointwise smallest solution that majorizes Q^*. For each $s \in S$, $P(\cdot, s)$ is convex and continuously differentiable. Moreover:

(a) There exists an optimal contract ξ^*, and any recursive contract generated by a policy function from [**FE**] is optimal.

(b) If the environment is [**TVC**]-Regular, then $P(\cdot, s)$ is *strictly* convex for each $s \in S$ and there is a *unique* optimal contract.

Proposition 3.3 follows from Theorem H.1 in BKL21, which also establishes several additional structural properties of the value function and optimal contract.[24]

4. Main Results

We present our main results on the long-run properties of the optimal contract in Subsections 4.1 and 4.2, and outline the main steps of the proofs in Subsection 4.3.

4.1. Immiseration

Our first main result establishes that the optimal contract always induces *immiseration*: the agent's (promised) utility and consumption tend to their lower bounds, so he becomes impoverished in the long run. To formally state this result, we require the following definition:

Definition 4.1. The type process is *pseudo-renewal* if there exists a probability distribution $\boldsymbol{\pi} \in \Delta(S)$ such that $f_{ij} = \pi_j$ whenever $i \neq j$.[25]

(24) While Bellman equations such as [**FE**] are familiar in this class of contracting problems, in our setting the proofs are somewhat non-standard. For example, even under Regularity, the value function is unbounded and may grow fast enough near the boundaries of D to render standard contraction-mapping methods inapplicable. We therefore rely on first-principles and order-theoretic arguments instead.

(25) This termology is from Hörner, Mu and Vielle (2017).

Pseudo-renewal processes have the distinguishing feature that, conditional on a transition occurring, the probability over new types does not depend on the previous type. This class includes some leading classes of type processes considered in the literature: all iid process are pseudo-renewal, as are all Markov processes when $d = 2$.

Theorem 1 (Immiseration)**.** *Suppose the environment is* [**TVC**]*-Regular. Under the optimal contract, as* $t \to \infty$*:*

(a) Promised utilities diverge: $v_i^{(t)} \to -\infty$ *almost surely for all* $i \in S$*.*

(b) Flow utilities diverge: $u_i^{(t)} \to -\infty$ *in probability for all* $i \in S$*. If the type process is pseudo-renewal, this can be strengthened from "in probability" to "almost surely."*

(c) Net consumption converges to its lower bound: $c_i^{(t)} + \omega_i \to \underline{c}$ *in probability for all* $i \in S$*.*[26] *If the type process is pseudo-renewal, this can be strengthened from "in probability" to "almost surely."*

Thus, none of promised utility, flow utility, or net consumption possess a stationary distribution under the optimal contract.[27]

The proof of Theorem 1 is sketched in Subsection 4.3 and presented in full in Online Appendix C. The special case of this theorem with pseudo-renewal types — for which we are able to establish almost sure divergence/convergence of the contractual variables — generalizes most immiseration results in the literature, which to our knowledge have focused either on iid types (e.g., Thomas and Worrall 1990) or Markovian types with $d = 2$ and specific assumptions on the transition probabilities (e.g., Zhang 2009, who assumes that $f_{22} = f_{11} \geq f_{12} = f_{21}$).[28] However, Theorem 1 further establishes that immiseration arises for *any* type process satisfying Assumption Markov, which allows for essentially arbitrary serial correlation and any number of types. We discuss the technical reasons for adopting the weaker mode of convergence in probability in Subsection 4.3 below, and the importance

(26) For each $i \in S$, let $\left(c_i^{(t)}\right)_{t \geq 0}$ denote the stochastic process (induced by the optimal contract) describing the consumption transferred from the principal to the agent in period t when the agent truthfully reports that $\omega^{(t)} = \omega_i$.

(27) Recall that we assume in Subsection 2.1 that the domain of feasible consumption levels is open, so that the lower bound $\underline{c}$ is not feasible even if $\underline{c} > -\infty$.

(28) Specifically, to our knowledge, the result generalizes all prior immiseration results for the benchmark principal-agent model studied here (i.e., without any of the "non-standard" model features described in Subsection 5.4), except for the non-generic case of "permanent shocks" discussed in Subsection 5.1 below (cf. Williams 2011; Bloedel, Krishna and Strulovici 2020, 2021). Moreover, as described above in Footnote 4, such single-agent immiseration results are known to translate naturally to unbounded inequality results in the many-agent planning version of the model studied by Atkeson and Lucas (1992) and others.

of Assumption Markov in Section 5.

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 (promised) utility is lower (cf. Thomas and Worrall (1990)). In particular, while the principal should give larger transfers to the agent when the reported endowment is lower in order to provide effective insurance, she must also accompany large transfers in the current period with lower continuation utility (i.e., lower transfers, on average, in future periods) in order to induce truthful reporting. Therefore, the agent's flow and continuation utilities must vary with his reported endowment, with 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 is 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

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

Thus, because $C(\cdot, 2)$ is convex (due to the agent's risk aversion), the cost of inducing utility variability increases with the the utility level.[29] In the long run, it is optimal for the principal to drive this cost of incentive provision to zero by driving the level of the agent's utility as low as possible. The next two subsections explain why this is so: Subsection 4.2 first provides economic intuition in terms of the power of the optimal contract's power of incentives, and Subsection 4.3 then sketches the formal proof of Theorem 1, which relies on the Martingale Convergence Theorem.

4.2. Backloaded Incentives and Relative Immiseration

Intuitively, the principal drives to cost of incentives to zero in the long run by immiserating the agent because of her cost-smoothing motive. As the level of (promised) utility decreases, it becomes affordable for the principal to provide *backloaded high-powered incentives*, i.e., make the sensitivity of the agent's continuation utility with respect to his report increases without bound in the long run. By backloading high-powered but cheap incentives in later

(29) See Banerjee and Newman (1991) and Newman (2007) for particularly clear articulations of this relation between wealth (i.e., promised utility) and risk-bearing (i.e., high-powered incentives) in a related class of principal-agent problems with production.

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 by showing that the optimal contract does, in fact, backload high-powered incentives. To state it, we require the additional assumption that types exhibit positive serial correlation:

Definition 4.2. The type process is:
(a) *FOSD-ordered* (FOSD) if $\mathbf{f}_i$ first-order stochastically dominates $\mathbf{f}_j$ whenever $i > j$.
(b) *Positive pseudo-renewal* (PPR) if it is FOSD and pseudo-renewal.

FOSD is the standard notion of positive serial correlation in the literature. It is easy to see that (i) every iid process is PPR, and (ii) every FOSD process is PPR when $d = 2$.

Theorem 2 (Backloaded Incentives)**.** *Suppose the environment is* [**TVC**]*-Regular and the type process is FOSD. Under the optimal contract, as $t \to \infty$:*
(a) $v_i^{(t)} - v_{i-1}^{(t)} \to +\infty$ in probability for all $i = 2, \ldots, d$. Moreover, for all $t \in \mathbb{N}$,

$$\mathrm{Var}\left(v_{s^{(t+k)}}^{(t+k-1)} \middle| \mathbf{v}^{(t)}, s^{(t)}\right) \to +\infty \qquad \textit{as } k \to \infty.$$

(b) If, in addition, immiseration occurs almost surely (namely, $u_i^{(t)} \to -\infty$ almost surely for each $i \in S$), then the convergence in part (a) can be strengthened from "in probability" to "almost surely." Moreover, in this case

$$\mathrm{Var}\left(v_{s^{(t+1)}}^{(t)} \middle| \mathbf{v}^{(t)}, s^{(t)}\right) \to +\infty \qquad \textit{as } t \to \infty$$

almost surely. In particular, these conclusions hold if the type process is PPR.

The proof of Theorem 2 is in Online Appendix D. While we have argued that Theorem 2 in some sense explains why immiseration arises in the first place, mathematically speaking it is a fairly simple consequence of Theorem 1. Intuitively, the principal does not waste the ability to provide high-powered incentives as they become affordable (as they do per Theorem 1). To see the gist of the argument, note that, as in Subsection 3.1, substituting the promise keeping constraint [**PK**$_i$] ($i = j$) into the incentive constraint [**IC**$_{ij}$] ($i > j$) gives an alternative way to write the latter:

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

When the type process satisfies FOSD, Theorem 3 in Appendix B guarantees that the Markov information rent term is non-negative, so that both the iid and Markov information rents work in the same direction (recall Subsection 3.1 for the $d = 2$ case). Part (b) of Theorem 1 implies that the iid information rent term, which is always non-negative, grows without bound. It then follows from [$\mathbf{IC}^*_{ij}$] (with $j = i - 1$) that the difference $v_i - v_{i-1}$ must also grow without bound.

It follows that the agent's continuation utility becomes increasingly uncertain over time, as in Theorem 2's statements about conditional variances. Specifically, at the beginning of period t, before the current endowment shock $\omega_{s(t+1)}$ is realized, the agent's type-contingent continuation utility $v^{(t)}_{s(t+1)}$ is a random variable with conditional mean given by the agent's date-t *ex ante* promised utility

$$[\mathbf{4.1}] \qquad \mathsf{E}\left[v^{(t)}_{s(t+1)} \middle| \mathbf{v}^{(t)}, s^{(t)}\right] = \sum_{i=1}^{d} f_{s^{(t)},i} v^{(t)}_i$$

and conditional variance

$$[\mathbf{4.2}] \qquad \mathrm{Var}\left(v^{(t)}_{s(t+1)} \middle| \mathbf{v}^{(t)}, s^{(t)}\right) = \sum_{i=1}^{d} f_{s^{(t)},i} \left(v^{(t)}_i - \mathsf{E}\left[v^{(t)}_i \middle| \mathbf{v}^{(t)}, s^{(t)}\right]\right)^2.$$

When the type process is PPR, Theorem 2(b) states that the conditional variance in [**4.2**] diverges, meaning that the agent's uncertainty at the start of period t (before observing his period-t type) about his future prospects (his continuation utility after observing and reporting his period-t type) increases without bound as $t \to \infty$. This represents one sense in which 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. 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 agent Bob. The Theorem states that this transient difference in Alice and Bob's endowments translates to an arbitrarily large permanent difference in welfare as $t \to \infty$. Viewed in this light, Theorem 2 expresses the idea that "pathwise welfare inequality" increases without bound in the long run.[30]

4.3. Proof Sketch for Theorem 1

This subsection sketches how we translate the cost-smoothing intuition described above into a formal proof of Theorem 1 (the full proof is in Online Appendix C). The proof consists of several steps, which are presented in sequence below. Building on Thomas and Worrall (1990), the basic idea is to identify a martingale that represents the principal's optimal cost-smoothing and to show that convergence of this martingale corresponds to the contract inducing immiseration. However, relative to the iid 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. The key to proving Theorem 1 is to formalize the principal's optimal cost-smoothing in terms of a martingale process (in analogy to a standard "Euler equation"). In the special case of iid types, Thomas and Worrall (1990) show that the principal's marginal cost of increasing the agent's *ex ante* promised utility defines a martingale under the optimal contract. In particular, when types are iid, we can write the principal's value function as $\hat{P}(v)$, where $v = \mathsf{E}^{\mathbf{f}}[\mathbf{v}]$ is the agent's *ex ante* promised utility under the (type-independent) transition probability $\mathbf{f} \in \Delta(S)$. The derivative $\hat{P}'(v)$ of this value function then represents the principal's marginal cost of promised utility.

In the general Markovian setting that we study, the appropriate martingale is a somewhat subtle, but natural, generalization. Denote the derivative of P at $(\mathbf{v}, s)$ by $\mathsf{D}P(\mathbf{v}, s) = (P_1(\mathbf{v}, s), \dots, P_d(\mathbf{v}, s))$, where $P_i(\mathbf{v}, s)$ is the partial derivative with respect to the component v_i. The directional derivative of $P(\cdot, s)$ in direction $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$ is written as $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} [P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)] / \varepsilon$.

(30) This contrasts with 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 (as in [**4.1**]) grows without bound as $t \to \infty$. 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 $t \to \infty$). By contrast, Theorem 2(b) says that a *single* shock leads to unboundedly large welfare differences between two agents with *otherwise identical histories* as $t \to \infty$.

Proposition 4.3. If the environment is Regular, then the stochastic process $\left(\mathsf{D}_{\mathbf{1}} P(\mathbf{v}^{(t)}, s^{(t)})\right)_{t=0}^{\infty}$ induced by the optimal contract is a non-negative martingale. If the environment is [**TVC**]-Regular, then this process is also strictly positive.

The proof of Proposition 4.3 is in Online Appendix C.2. We will refer to this stochastic process as the principal's *marginal cost martingale*. The statement of the proposition includes three important pieces:

(a) *The Martingale Property:* Intuitively, the marginal cost process defines a martingale because the principal *optimally smooths costs* over time and across states. More precisely, fix a vector $\mathbf{v} \in D$ 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, the cost of this increase in promises is approximated by the marginal cost $\mathsf{D}_{\mathbf{1}} P(\mathbf{v}, s)$. One specific way to deliver the additional utility in an incentive-compatible manner is to increase each of the continuation promises $\mathbf{w}_i$ to $\mathbf{w}_i' := \mathbf{w}_i + (\varepsilon/\alpha)\mathbf{1}$. Using the Bellman equation [**FE**], the cost of this perturbation is approximately $\sum_{i=1}^{d} f_{si} \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i, i)$. By an envelope argument, this perturbation is locally optimal, implying that the marginal costs are equal, giving us precisely the martingale property $\mathsf{D}_{\mathbf{1}} P(\mathbf{v}, s) = \sum_{i=1}^{d} f_{si} \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i, i)$.

(b) *Differentiation in Direction* $\mathbf{1}$*:* Notice that 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*. Specifically, recalling the form [$\mathbf{IC}^*_{ij}$] for the agent's incentive constraints, it is clear that a perturbation of $\mathbf{v}$ leaves the left-hand side of the [$\mathbf{IC}^*_{ij}$] constraints unchanged if, and only if, it is taken in the direction $\mathbf{1}$. Alternatively, we may consider perturbations of the $\mathbf{w}_i$: 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., $\mathsf{E}^{\mathbf{f}_i}[\hat{\mathbf{w}}_i] - \mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i] = \varepsilon$. In general, other types $j \neq i$ value this perturbation differently, as they have different beliefs about future types (i.e., $\mathbf{f}_j \neq \mathbf{f}_i$). 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 $\mathsf{E}^{\mathbf{f}_j}[\hat{\mathbf{w}}_i - \mathbf{w}_i] = \varepsilon$ for all $j \in S$), ensuring that incentive compatibility is preserved.[31]

(c) *(Strict) Positivity of the Marginal Cost:* Somewhat subtly, it is *not* true that the partial derivatives $P_i(\mathbf{v}, s)$ are always non-negative, as increasing a single component v_i of $\mathbf{v}$ has two countervailing effects.[32] First, it increases the promise to type i, which

(31) When the transition probabilities $[\mathbf{f}_i]_{i=1}^{d}$ are affinely independent, $\mathbf{1}$ is the *unique* direction with this property.

(32) Thus, the "interim" Pareto frontier that traces the principal's costs as a function of the promised utility

mechanically increases costs. Second, it tightens incentive constraints for higher types $j > i$ that can misreport as type i, but also adds slack to incentive constraints for type i when misreporting as a lower type $j < i$. Whenever it adds "enough" slack to the latter incentive constraints, increasing v_i can lead to an overall decrease in costs. However, this is not an issue for the marginal cost martingale because differentiation in the direction $\mathbf{1}$ holds information rents for all types fixed, leaving only the mechanical increase in costs due to higher promised utilities.

Step 2: The Cost of Incentives, Martingale Splitting, and Efficiency. To understand the marginal cost martingale's dynamics, it is useful to recall two benchmarks. First, when the agent does not have any private information, the first-best contract completely stabilizes the agent's consumption and utility (recall Subsection 3.1). The first-best contract thus perfectly smooths the principal's costs, resulting in a *constant* marginal cost martingale. Second, when the agent's type is private but iid over time (as in Thomas and Worrall (1990)), the principal provides incentives by making the agent's *ex ante* continuation utility $w_i = \mathsf{E}^{\mathbf{f}}[\mathbf{w}_i]$ vary with his reported type i (here $\mathbf{f}$ denotes the type-independent transition probability). Indeed, this is the *only* way to provide incentives through continuation contracts (as opposed to current transfers) when types are iid, because the agent's valuation of any given continuation contract is independent of his current type. The optimal contract therefore always sets $w_d > v > w_1$ in order to (a) reward the highest-type agents for not under-reporting and (b) punish the lowest-type agents so as to deter other types from under-reporting. Consequently, the principal's marginal cost (of *ex ante* promised utility) 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, and the "size" of martingale splitting quantifies the *cost of incentives provision* relative to the first-best benchmark: the difference $\hat{P}'(w_d) - \hat{P}'(v) > 0$ (for example) is precisely determined by the magnitude of the Lagrange multipliers on the agent's incentive constraints. Convergence of the marginal cost martingale — which must occur, by the Martingale Convergence Theorem — therefore corresponds to the cost of incentives vanishing in the long run.

Persistent private information significantly complicates the above logic. Most importantly, persistence means that the agent's current type determines his preferences over

vector $\mathbf{v}$ is *not* downward-sloping, which is closely related to the fact that the optimal contract is generally not renegotiation-proof (see Step 2 below). More formally, Theorem H.1 in BKL21 shows that for each $i > 1$, there exists an open set of implementable promised utility vectors $\mathbf{v} \in D$ for which $P_i(\mathbf{v}, s) < 0$.

continuation contracts, giving the principal new opportunities to *screen the agent through continuation contracts* that manipulate his Markov information rent (recall [**IC**$^*_{ij}$]). For example, because different agent types value any given continuation contract differently, it may be possible for the principal to incentivize truthtelling while also making the agent's *ex ante* continuation utility $\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i]$ independent of his reported type i by choosing the interim continuation promises $\{\mathbf{w}_i\}_{i\in S}$ to manipulate the deviation continuation payoffs $\mathsf{E}^{\mathbf{f}_j}[\mathbf{w}_i]$ $(j \neq i)$ appropriately. Given this additional channel for incentive provision, it is not clear whether (a) the marginal cost martingale splits (i.e., satisfies $\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_1) < \mathsf{D}_{\mathbf{1}}P(\mathbf{v}) < \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d)$) or (b) even if so, what relation this has to the cost of incentives, as measured by the Lagrange multipliers on the agent's incentive constraints. An added complication is that global incentive constraints (i.e., [**IC**$^*_{ij}$] with $j < i-1$) may bind when types are persistent, and it is unclear which of them actually do.

The key to overcoming these challenges is to identify a special class 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 $w \in \mathscr{U}$ 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 (at (w,i))*:

$$[\mathbf{Eff}_i] \qquad K(w,i) := \min_{\mathbf{w}_i \in D} P(\mathbf{w}_i, i) \quad \text{s.t.} \quad \mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i] \geq w.$$

Solutions to [**Eff**$_i$] are said to be *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 case of iid types is simple precisely because the optimal contract is renegotiation-proof in this sense (see Remark C.1).[33] However, when types are persistent, 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

(33) This is the same notion of renegotiation-proofness that is informally discussed in Thomas and Worrall (1990, p. 369), and we use the term "renegotiation-proof" in the same informal sense here. We do not delve into the subtleties of formally defining renegotiation-proof contracts; see Strulovici (2017) and Strulovici (2020) for more stringent definitions in closely related settings. Also 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."

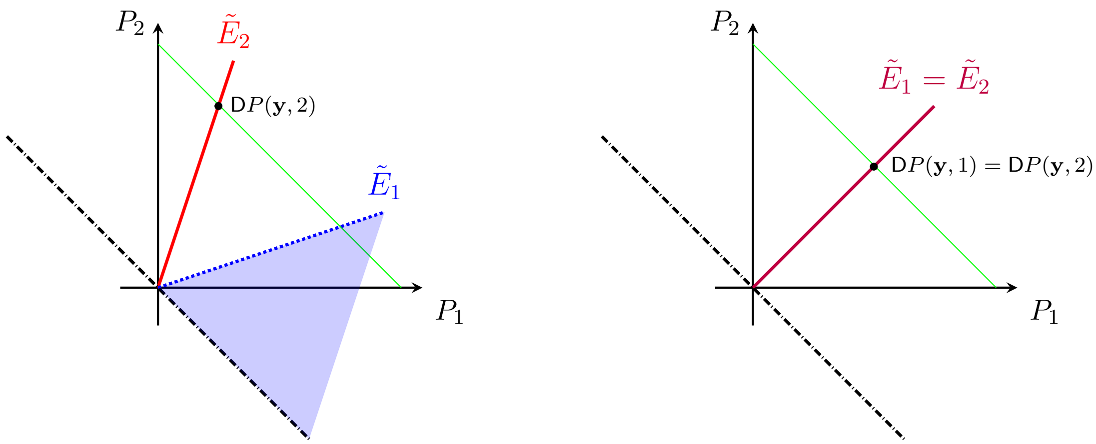

Figure 2: *Martingale convergence argument when $d = 2$ and the type process satisfies FOSD (left) and is iid (right).*

the highest type's promised utility $\mathbf{w}_d$ does not enter any of these incentive constraints, so that the optimal contract is always efficient at those histories where the agent had reported to have the highest endowment level ω_d in the previous period (see Lemma C.4). Moreover, at such histories, the marginal cost martingale splits and the size of its splitting pins down the cost of incentives, in analogy to the iid case (see [**C.10**] in Online Appendix C.3.2). 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.15).

Step 3: Martingale Convergence. The Martingale Convergence Theorem guarantees that the marginal cost martingale converges to a non-negative, integrable random variable. The final two steps of the proof establish, first, that the limit random variable is almost surely zero and, second, that this implies that the optimal contract induces immiseration.

The argument that the martingale converges to zero is illustrated in Figure 2 for 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 — namely, the derivative $\mathsf{D}P(\mathbf{w}_i, i)$, which is pinned down by the first-order conditions for [**Eff**$_i$]. When types are iid (the right-hand panel in Figure 2), there is a single efficiency ray because the principal's value function can be written as a function of *ex ante* promised utility alone, and the derivative process traverses this efficiency ray because the optimal contract is renegotiation-proof. Consequently, it is easy to see 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: $\mathbf{v}^{(t)} \to \mathbf{y} \in D$

almost surely.[34] Per Step 2 above, this would imply that the optimal contract perfectly stabilizes the agent's consumption in the limit (i.e., converges to the first-best), which would violate incentive compatibility. Thus, the martingale must converge to zero along almost every path. When types are persistent (the left-hand panel in Figure 2), we show that a similar argument applies along the subsequence of histories at which the agent (truthfully) reported to have the highest endowment ω_d in the previous period because, as described above, the optimal contract is efficient at such histories. Being that the principal's marginal cost is both a martingale and positive, this convergence can be extend to (almost) all histories, establishing that $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}^{(t)}, s^{(t)}) \to 0$ almost surely (Lemma C.16).

Step 4: Convergence of Contractual Variables. When types are iid, there is a one-to-one relationship between *ex ante* promised utility v and the marginal cost $\hat{P}'(v)$ (again see the right-hand panel of Figure 2), so immiseration follows quickly from martingale convergence. However, the situation is once again complicated by persistence: even though the martingale converges to zero, it is possible that the agent's promised utility vector fails to converge (i.e., remains transient or converges to a non-degenerate stationary distribution) because the principal provides incentives (only) by screening through continuation contracts at histories where the optimal contract is not efficient. (Geometrically, in the left-hand panel of Figure 2, the derivative process $\mathsf{D}P(\mathbf{v}^{(t)}, s^{(t)})$ may traverse the blue region below the efficiency ray $\tilde{E}_1$ at such histories, even as the directional derivative represented by the green line converges to zero.)

To rule this out, the key step is to 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 which it can be deduced that the contract either (a) converges to the first-best or (b) immiserates the agent, only the latter of which is consistent with incentive compatibility. The arguments that establish convergence of the multipliers are fairly involved, but the main ideas are as follows. When types are pseudo-renewal (which includes the iid case), we show that 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 (Lemma C.21). For general type process, we are unable to rule out that the

(34) This relies on strict concavity of $\hat{P}$, so that the mapping $v \mapsto \hat{P}'(v)$ is bijective. The facts that (i) $\mathbf{v}^{(t)}$ converges and (ii) its limit point is in D (rather than boundary point) rely on the Inada conditions in Assumption DARA(a). See Subsection 5.2 for further discussion.

vector of multipliers take infinitely-many excursions away from zero on any given path, but Assumption Markov implies that the agent's endowment almost surely returns to the highest level ω_d infinitely often. We show that this causes the multiplier process to exhibit a "renewal property" whereby it returns to a neighborhood of zero infinitely often and, moreover, mixes quickly enough to guarantee that excursions away from zero are very rare in the long run, allowing us to establish that the multipliers converge to zero in probability (Lemma C.19). Once convergence of the Lagrange multipliers is established, convergence of the contractual variables themselves follows quickly (Lemmas C.22–C.23).

5. Discussion

This section consists of several parts. Subsection 5.1 discusses the importance of Assumption Markov and some implications of Theorem 1 for open questions raised by the recent literature. Subsection 5.2 discusses the roles played by our (other) baseline assumptions and how our results extend to other insurance settings. Subsection 5.3 summarizes the short-run properties of the optimal contract. Finally, Subsection 5.4 surveys strands of the related literature that have not already been discussed.

5.1. The Role of Mean-Reversion

An important counterpoint to our Theorems 1–2 is provided by Williams (2011) (henceforth W11), which 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 an ergodic Markov chain. In stark contrast to our results, W11 finds that the optimal contract generates long-run *bliss* for the agent (i.e., sends $c_i^{(t)} \to +\infty$ and $u_i^{(t)} \to 0$ almost surely),[35] casting doubt on the robustness of immiseration outside of the iid benchmark studied by the classic literature.

By contrast, our results demonstrate that immiseration is robust to a broad class of ergodic type processes. This suggests that ergodicity (or, more informally, "mean-reversion") of the type process is a key determinant of the optimal contract's long-run properties and, in

(35) More precisely, W11 studies a continuous-time variant of the model in which the agent's type follows a Brownian motion and attributes the failure of immiseration, in part, to the difference between discrete- and continuous-time models. However, Bloedel, Krishna and Strulovici (2020, 2021) clarify that W11's results extend to the discrete-time version of the model, facilitating a more direct comparison to our discrete-time analysis. Moreover, W11 finds that long-run bliss also arises when the agent's type follows a mean-reverting AR(1) process, but Bloedel, Krishna and Strulovici (2020) show that the contract derived in W11 is strictly suboptimal in that case, so that the only relevant comparison here is to the random walk case.

particular, that the failure of immiseration in W11 hinges on the knife-edge assumption of *zero* mean-reversion. A simple intuition for this is that, with mean-reversion, the current type in period t is "approximately independent" of distant-future types in period $t+T$ (as $T \to \infty$), so that the mechanics behind immiseration in the classic iid benchmark "should" kick in over long time horizons. Indeed, certain aspects of this informal reasoning are evident in our formal proof of Theorem 1 (e.g., the "renewal" dynamics described in Step 4 of Subsection 4.3).

More formally, the key property of our model is that the agent's type process has *impulse response functions* (in the language of Pavan, Segal and Toikka 2014) that vanish asymptotically: conditional on period-t information, the effect of a marginal increase in the agent's (given) period t type on his (stochastic) period $t+T$ type vanishes as $T \to \infty$, so that the principal and agent have "approximately symmetric" information about the agent's distant-future types and, therefore, the agent's Markov information rents are "sufficiently small."[36] This allows the principal to separately manipulate the *level* of the agent's *ex ante* promised utility and the *slope* of his *interim* promised utility schedule (i.e, the power of incentives). Formally, this is expressed by the fact that the recursive domain D has nonempty interior (Theorem 3(a)), so that the principal's marginal cost martingale $\mathsf{D}_1 P(\mathbf{v}^{(t)}, s^{(t)})$ is a well-defined object (Proposition 4.3).

Using different techniques, Bloedel, Krishna and Strulovici (2020, 2021) study the complementary case of *permanent shocks* in which the impulse response functions are constant and equal to one (i.e., a unit increase in the agent's current type induces a "permanent" unit increase in all future types). This property is equivalent to the agent's type following a "generalized random walk," a class that includes the Gaussian random walks studied in W11. In that setting, the agent's Markov information rents are "sufficiently large" that (i) the agent is indifferent among *all* reporting strategies under *any* incentive compatible contract, and (ii) the schedule of interim promised utilities is uniquely pinned down by the level of ex ante promised utility. Property (i) means that all of the agent's incentive constraints — even global ones — always hold with equality, so that the main economic force underlying immiseration — the backloading of high-powered incentives — is automatically shut off. Consequently, the principal faces a very different set of tradeoffs: there is zero risk-sharing under any incentive compatible contract, and the optimal contract is implemented with *determinstic*

(36) See 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).

(i.e., report-independent) transfers that merely change the drift of the agent's consumption. Property (ii) implies that the principal's marginal-cost martingale is not well-defined: the principal's marginal cost defines a non-negative but unbounded *sub*martingale, for which the Martingale Convergence Theorem has no implications.

5.2. Baseline Assumptions and Extensions

In this subsection, we discuss the baseline model assumptions from Section 2 and how our analysis can be extended to several model variants.

Shape of Utility Function. Assumption DARA imposes both substantive and technical restrictions on the agent's risk preferences. Part (a) embeds the economically substantive conditions needed to establish that immiseration occurs with probability one in Theorem 1: strict monotonicity and concavity of $U(\cdot)$ ensure that first-best insurance is not implementable (i.e., the incentive problem is nontrivial), while the upper Inada condition $\lim_{c\to\infty} U'(c) = 0$ ensures that there is no upper bound for the principal's cost of incentive provision (i.e., maintaining truthtelling requires unboundedly large variability in consumption as the level of utility increases).[37] The remaining pieces of Assumption DARA play purely technical roles (and could likely be relaxed): smoothness of $U(\cdot)$ in part (a) implies that the principal's value function [**RP**] is sufficiently smooth for the martingale convergence proof of Theorem 1, unboundedness below of $U(\cdot)$ in part (b) implies that the optimal contract is interior,[38, 39] and DARA in part (c) implies that the principal's problem is convex, so that report-contingent randomization is unnecessary and the optimal contract can be characterized through first-

(37) If $M := \lim_{c\to\infty} U'(c) > 0$, then the optimal contract would induce long-run *polarization*: the agent's consumption and utility would converge *either* to their lower or upper bounds, both with positive probability (Phelan 1998). However, even then the probability of immiseration would tend to one as $M \to 0$. It is conceptually straightforward to extend our analysis to the $M > 0$ case.

(38) As discussed in Phelan (1998, p. 176), interiority allows us to bypass the complications caused by corner solutions, which are orthogonal to the basic incentive problem of interest. Specifically, when the range of feasible utilities $\mathcal{U} = \{U(c) : c \in \mathcal{C}\}$ contains its lower boundary point (e.g., if the consumption domain is both bounded below and closed), we must deal separately with the contract's behavior at the lower boundary to ensure that it is not "reflecting," which would be a force against immiseration. This is easy to address when types are iid, in which case the domain of *ex ante* promised utilities is unidimensional and hence has a singleton lower boundary (e.g., Proposition 6 and Lemma 2 of Golosov, Tsyvinski and Werquin 2016), but requires more involved analysis when types are persistent, in which case the domain's boundary is multi-dimensional.

(39) Notably, 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 $U(\cdot) \geq 0$, then $c_i^{(t)} + \omega_i \to \underline{c}$ for all i would imply that $v_i^{(t)}, u_i^{(t)} \to 0$ for all i. However, the basic intuition is robust, for risk-aversion alone implies that it is cheapest for the principal to induce variations in utility when its level is lowest.

order conditions.

Source of Private Information. For expositional simplicity, we have assumed that the agent's private information concerns his endowment. However, as noted in Remark 2.1, this aspect of the model is not essential. Indeed, our analysis would extend in a straightforward manner to several other insurance settings, including the following canonical classes of models:

- *Taste shock models:* 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, which coincides with our baseline hidden endowment model when the agent has CARA utility, has been a workhorse specification in the macro literature since Atkeson and Lucas (1992). Our analysis would extend almost verbatim, assuming that $u(\cdot)$ is strictly increasing, strictly concave, continuously differentiable, and unbounded below.
- *Separable Mirrleesian models:* 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 labor 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 Mirrleesian model is the workhorse specification in much of the optimal dynamic taxation literature (e.g., Zhang 2009; Kocherlakota 2010). Our analysis would extend almost verbatim, assuming that (i) $u(\cdot)$ and $-v(\omega, \cdot)$ are strictly increasing, strictly concave, continuously differentiable, and unbounded below, and that (ii) v exhibits strictly decreasing differences in (ω, ℓ) and is such that $x \mapsto v(\omega, v^{-1}(\omega', x))$ is concave for $\omega > \omega'$, conditions which are satisfied by most commonly used parametric specifications.[40]
- *Non-separable Mirrleesian models:* The agent has utility $U(\omega, c, \ell)$ over consumption c and labor effort ℓ, with the same interpretation as above, but the utility function

(40) In the separable Mirrleesian model, the principal's marginal cost of promised utility admits the representation $\mathsf{D}P_{\mathbf{1}}(\mathbf{v}, s) = 1/u'(c(\mathbf{v}, s))$, so that her marginal cost martingale reduces to the celebrated *Inverse Euler Equation* (IEE) (Golosov, Kocherlakota and Tsyvinski 2003). However, this correspondence relies on additive separability, so that the agent's consumption utility $u(c)$ serves as a "numeraire" that the principal can use to deliver promised utility in a type-independent manner (cf. monetary transfers in mechanism design model with quasi-linear utility). In our baseline model and the other model variants described here, such separability is absent and the IEE does not arise. This implies that the IEE is not necessary for immiseration to occur. It is also not sufficient, for even in the separable Mirrleesian model it has no implications for the labor component of the allocation.

need not be additively separable with respect to consumption. This extension of the workhorse optimal taxation model has received significant recent interest (e.g., Farhi and Werning 2013; Golosov, Troshkin and Tsyvinski 2016). Our analysis would extend almost verbatim under assumptions analogous to those described above.

Feasible Reporting Strategies. Assumption NHB limits the agent's set of feasible reporting strategies by preventing him from overreporting his type.[41] We view this as a mild assumption. First, in our baseline hidden endowment model, constraints on under-reporting are the empirically relevant concern (e.g., see Feldman and Slemrod (2007) for evidence from US tax data). Second, whenever the agent's private type corresponds to something that is verifiable in principle, such as his endowment or labor productivity, NHB is without loss of generality if the agent cannot covertly save, borrow, or engage in production outside of the contract.[42] In this light, relaxing NHB would correspond to imposing additional constraints on the principal corresponding to her limited ability to enforce the terms of her contract, as in the literature on optimal contracts with hidden saving (e.g., Allen 1985; Cole and Kocherlakota 2001b). Such constraints are conceptually distinct from the basic incentive constraints arising from the hidden information problem on which we focus.

However, a different interpretation is that NHB relaxes the "full" problem in which overreporting is feasible for the agent, which may be more natural in settings where the agent's private type corresponds to something purely subjective, as in the taste-shock model variant described above. Even then, NHB would correspond to a significantly *less* relaxed problem than that from the popular first-order approach (FOA), in which only "local downward" incentive constraints (of the form [$\mathbf{IC}_{ij}$] for $j = i - 1$) are included, and would therefore be "valid" (i.e., yields a solution satisfying the full set of constraints) in a larger set of environments. This additional robustness is important because the FOA is typically valid only under strong assumptions on the agent's type process.[43]

(41) NHB is used in the proof of Theorem 1 (e.g., Lemma C.15 and several inductive arguments in Online Appendix C.4), but we conjecture that it could be dispensed with. Conditional on immiseration occurring, Theorem 2 would go through verbatim.

(42) This interpretation is common in the literature. In the context of hidden endowments, see Phelan (1998), Fernandes and Phelan (2000, footnote 4) (who argue that NHB is without loss of generality when $\underline{c} > -\infty$), and Williams (2011). In the context of hidden productivity, see Golosov and Tsyvinski (2007, p. 492). Notably, NHB is also built by fiat into the closely related family of cash-flow diversion models from corporate finance (e.g., Fu and Krishna 2019 and references therein) and macroeconomics (e.g., Di Tella and Sannikov 2021).

(43) The FOA is known to be valid in the present insurance setting when types are iid (e.g., Thomas and

5.3. Short-Run Properties

While our main interest is in the optimal contract's long-run properties, our analysis also has implications for its short-run dynamics. To isolate the novel effects of persistence, we study a simple example with CARA utility (as in Green 1987; Thomas and Worrall 1990) and binary, positively correlated types with symmetric transitions (i.e., $d = 2$ and $f_{11} = f_{22} \geq f_{12} = f_{21}$, as in Zhang 2009). In BKL21, we provide a detailed description of the optimal contract using both analytical and numerical characterizations. Our main findings include the following:

- Outside the iid case, the optimal contract is *not* renegotiation-proof. In particular, it is never efficient after low-endowment reports: relative to solutions of the efficiency problem [**Eff**$_i$] (with $i = 1$), the agent's continuation utility w_1 is pushed closer to the diagonal (the lower boundary of V_2 in Figure 1) so as to *compress* the difference $w_{12} - w_{11} > 0$, which serves to reduce the agent's Markov information rent $\alpha(f_{22} - f_{12}) \cdot (w_{12} - w_{11})$ in [**IC**$^*_{21}$]. This gives rise to rich "cyclical" dynamics for the agent's promised utility vector, somewhat reminiscent of those in Zhang (2009).
- The optimal contract exhibits a novel *order-dependence* property, whereby the agent's continuation utility is *lower* if a fixed number of low endowment realizations occurs *earlier* in the sequence. This stands in stark contrast to a key finding of Thomas and Worrall (1990) that, in the iid case (with CARA utility), the optimal contract depends on the history of endowment realizations only via the number, but not the order, of high and low shocks.
- With persistence, the optimal contract may *over-insure* the agent after certain histories by allocating more net consumption to low-type agents (i.e., set $c_2 + \omega_2 < c_1 + \omega_1$). Intuitively, this "punishes" high-type agents so as to reduce the low-types' Markov information rents (i.e., compress $w_{12} - w_{11} > 0$) in earlier periods. This contrasts with the iid benchmark, in which the optimal contract always features *under-insurance* (i.e., $c_2 + \omega_2 > c_1 + \omega_1$) so as to reduce the agent's iid information rents. Zhang's (2009) separable Mirrleesian model also always features under-insurance, even when types are persistent, suggesting that allocative distortions behave very differently in separable and non-separable insurance models.
- Numerically, we find that greater persistence gives rise to quantitatively larger distortions away from first-best risk-sharing. These differences are substantial: while our simulations

Worrall 1990). However, the results of Battaglini and Lamba (2019) suggest that it is likely to be invalid when types are highly persistent (see also Pavan (2016) and Garrett, Pavan and Toikka (2018)).

are not meant to be realistically calibrated, we find that high persistence gives rise to *insurance* and *intertemporal wedges* (standard measures of allocative distortions) that are several times greater than in the iid benchmark.

5.4. Related Literature

Dynamic Insurance and Immiseration. The classic immiseration results were established under the assumption of iid private information in "partial equilibrium" by Green (1987) and Thomas and Worrall (1990) and in "general equilibrium" by Atkeson and Lucas (1992). Our analysis builds on Thomas and Worrall's (1990), which first related immiseration to the Martingale Convergence Theorem.

Our paper is related to the sizable literature that studies the robustness or fragility of immiseration under alternative assumptions on preferences, technology, or institutions (while maintaining the iid assumption). First, several papers derive bounded long-run inequality by relaxing the contracting parties' commitment power or, equivalently, by adopting a different normative criterion that places Pareto weight directly on all "future generations," which generates a force towards mean-reversion that destroys the martingale property of the principal's marginal cost process (cf. Footnote 1). Second, several papers derive weaker forms of immiseration in growing economies (Khan and Ravikumar 2001; Bloedel and Krishna 2015), intergenerational economies with endogenous fertility (Hosseini, Jones and Shourideh 2013), and under alternative assumptions on agents' risk preferences (Phelan 1998; Olszewski and Safronov 2021) by studying convergence properties of suitably modified versions of the principal's marginal cost martingale. Finally, immiseration fails when the agent can covertly save outside of the contract, which prevents the principal from providing any insurance (Allen 1985; Cole and Kocherlakota 2001b) and leads to contractual dynamics governed by a different martingale, the agent's Euler equation, which typically generates a backloaded consumption profile that sends the agent to long-run bliss (e.g., Ljunqvist and Sargent 2012, Ch. 17).

In contrast to these studies, we move beyond the restrictive assumption of iid types but otherwise maintain baseline assumptions on preferences and technology. The most related papers in this respect are Zhang (2009), Williams (2011), and Bloedel, Krishna and Strulovici (2020, 2021), which were discussed in Subsection 5.1 above.

Recursive Contracts with Persistent Types. The recursive approach to dynamic screening was introduced by Green (1987) and Thomas and Worrall (1990) for iid types, extended by Fernandes and Phelan (2000) to accommodate Markovian types, and further extended to settings with hidden actions (Doepke and Townsend 2006) and continuous time (Zhang 2009).[44] Our recursive formulation is essentially isomorphic to Fernandes and Phelan's (2000). However, their approach requires tracking off-path "threat-point" promised utilities that represent the agent's continuation payoffs conditional on having lied in the previous period, which never occurs on-path. Instead, our approach uses on-path type-contingent promised utilities, which results in incentive constraints that are easier to interpret and is technically more convenient.[45]

Methodologically, the closest work to ours is a set of three recent papers — either concurrent or subsequent to our working paper (Bloedel and Krishna 2015) — that use the same recursive formulation to study different contracting problems with imperfectly transferable utility and persistent types. Importantly, all three focus on the special case of binary types with positive serial correlation (as in our solved example described in Subsection 5.3).[46] In concurrent work, Guo and Hörner (2020) study a dynamic allocation problem without transfers and establish a version of Phelan's (1998) polarization result through a detailed construction of the optimal contract, rather than through the martingale convergence arguments used in this paper.[47] Fu and Krishna (2019) (concurrently) and Krasikov and Lamba (2021) (subsequently) use the same techniques as this paper to study closely related models of firm financing and repeated procurement in which the agent is risk-neutral, but monetary transfers are subject to a limited liability constraint. In their models, the optimal contract converges to the first-best in finite time, as it is eventually optimal for

(44) See also Spear and Srivastava (1987) for an early treatment with hidden actions, Cole and Kocherlakota (2001a) for an extension to stochastic games without commitment, and Ljunqvist and Sargent (2012, Ch. 20-21) and Golosov, Tsyvinski and Werquin (2016) for surveys of applications in macroeconomics.

(45) See Appendix B.1 for further discussion. As described there, while in our formulation the domain is a set whose shape is (under mild assumptions) independent of the type process, in Fernandes and Phelan's (2000) formulation the domain is a set-valued *function* of past reports, the shape of which changes depending on the persistence of the agent's type. Moreover, in an important intermediate step to proving Theorem 1 (see Online Appendix C.3.3), we rely on a decomposition of the principal's Bellman equation [**FE**] into a collection of type-contingent "interim" problems, which has no natural analogue in the Fernandes and Phelan (2000) approach.

(46) See also Halac and Yared (2014) and Broer, Kapička and Klein (2017).

(47) Guo and Hörner's (2020) model is equivalent to a variant of the taste-shock version of ours (cf. Subsection 5.2) in which the agent is risk-neutral and has a bounded domain of (randomized) consumption levels. These differences mute the consumption-smoothing motive that is central to our model and lead to distinct formal analyses.

the principal to sell the firm to the agent. Notably, the proofs of long-run convergence and recursive domain constructions in Fu and Krishna (2019) and Krasikov and Lamba (2021) are essentially special cases of ours for pseudo-renewal type processes (Theorems 1 and 3).

Appendix

This appendix contains details concerning the recursive formulation of the contracting problem that were omitted from the main text. Appendix A first provides details for the sequential version of the contracting problem described informally in Subsection 2.2. Appendix B then presents several structural results, including a general characterization of the recursive domain (Theorem 3) and sufficient conditions for the Regularity conditions from Subsection 3.2 to hold.

A. Sequential Contracts, Recursive Contracts, and their Equivalence

A.1. Sequential Contracts

This appendix provides details for the sequential formulation of the contracting problem outlined in Subsection 2.2. By the Revelation Principle, we may restrict attention to direct revelation mechanisms.[48] Let $\mathscr{G}$ denote the space of *private* histories, which are sequences of *realized* endowment types of the form $g = (s^1, \dots) \in S^\infty$.[49] Similarly, $\mathscr{H}$ is the space of *public* histories, which are sequences of *reported* endowment types $h = (\hat{s}^0, \hat{s}^1, \dots) \in S^\infty$. Let G^t and H^t denote the spaces of length-t private and public histories, respectively. Thus, $h^t \in H^t$ is of the form $h^t = (s^1, \dots, s^t)$, so that h^t records the type realizations in periods $0, 1, \dots, t-1$. A *(pure) reporting strategy* $\sigma := (\sigma_t)_{t=0}^\infty$ for the agent is a sequence of functions $\sigma_t : G^{t+1} \times H^t \to S$. The *truthful strategy* σ^* is defined by $\sigma_t^*\big((g^t, s^t), h^t)\big) = s^{t+1}$ for all $g^t \in G^t$ and $h^t \in H^t$. Strategy σ is *admissible* if $\sigma_t\left((g^t, s^{t+1}), h^t\right) \leq s^{t+1}$ for all $g^t \in G^t$ and $h^t \in H^t$, i.e., if the agent never over-reports his endowment. The set of admissible strategies is denoted Σ. Under Assumption NHB, the agent only has access to $\sigma \in \Sigma$.

While sequential contracts are naturally described in terms of transfers of the consumption good, it is most convenient to formulate them in terms of flow utilities. A transfer of c_i from the principal to an agent with endowment ω_i delivers to the agent flow utility

(48) Assumption NHB implies that the agent can only misreport his type in one direction, so we formally require a version of the Revelation Principle suited for environments with partial verification. Since the reporting constraints in each period satisfy the "nested range condition" of Green and Laffont (1986), straightforward adaptations of their arguments to our dynamic setting establish the appropriate version.

(49) $s^{(t)}$ is a random variable and s^t is a realization. Recall from Footnote 17 the timing convention whereby $\omega^{(t)} = \omega_i$ is equivalent to $s^{(t+1)} = i$, so that $s^{(t+1)}$ is the period t type (*not* the period $t + 1$ type).

$u_i := U(c_i + \omega_i)$. Thus, any such transfer is equivalent to a flow utility allocation of u_i to an i-type agent at cost $C(u_i, i) := U^{-1}(u_i) - \omega_i$. Thus, a *sequential contract*, denoted $\tilde{u} := \left(u^{(t)}\right)_{t=0}^{\infty}$, is a $\mathcal{U}$-valued stochastic process adapted to the filtration implied by the public histories. Let $u_i^{(t)}$ denote the (random) date t flow allocation when the period t report $\hat{s}^{(t+1)} = i$. The set of all sequential contracts is $\mathcal{A}$. Recall that $\psi : \mathcal{U} \times S \times S \to \mathcal{U}$ defined as $\psi(u, i, j) := U\left(\omega_i + C(u, j)\right)$ specifies how an agent of type i values the flow utility allocation intended for type j. In particular, if an agent of type i lies and claims to be of type $j \neq i$, he receives flow utility $\psi(u_j, i, j)$. If he truthfully reports his type to be i, he receives flow utility $\psi(u_i, i, i) = u_i$.

Every sequential contract $\tilde{u}$ and reporting strategy σ together induce a stochastic process $\left(s^{(t+1)}, \hat{s}^{(t+1)}, u^{(t)}\right)_{t=0}^{\infty}$ over true types, reported types, and flow utility allocations. Taking $\tilde{u}$ as given, denote the law of this process by $\mathbf{P}^{\sigma} \in \Delta\left(S^{\infty} \times S^{\infty} \times \mathcal{U}^{\infty}\right)$ and its associated expectation operator by $\mathsf{E}^{\sigma}[\cdot]$. Thus, the agent's preferences over sequential contracts and admissible reporting strategies are represented by the lifetime utility function $\hat{U} : \mathcal{A} \times \Sigma \times S \to \mathcal{U} \cup \{-\infty\}$ defined by

$$\hat{U}\left(\tilde{u}, \sigma, s\right) := \mathsf{E}^{\sigma}\left[\sum_{t=0}^{\infty} \alpha^t \psi\left(u_{\hat{s}^{(t+1)}}^{(t)}, s^{(t+1)}, \hat{s}^{(t+1)}\right) \big| s^{(0)} = s\right]$$

We say that a sequential contract $\tilde{u}$ *implements* $\mathbf{v} \in \mathcal{U}^d$ if it satisfies

[S-PK$_i$] $$v_i = \hat{U}(\tilde{u}, \sigma^*, i)$$

[S-IC] $$\hat{U}(\tilde{u}, \sigma^*, i) \geq \hat{U}(\tilde{u}, \sigma, i) \qquad \forall\, \sigma \in \Sigma$$

for all $i \in S$. The [**S-PK$_i$**] constraints are the familiar *promise-keeping* conditions and the [**S-IC**] are *incentive compatibility* constraints. The set of sequential contracts that implement $\mathbf{v}$ is $\Pi(\mathbf{v})$.

When the agent follows the strategy σ and the principal has prior belief $\mu \in \Delta(S)$ over the agent's initial type, the cost to the principal of a sequential contract $\tilde{u}$ is

$$R(\tilde{u}, \sigma, \mu) := \mathsf{E}^{\sigma}_{s^{(0)} \sim \mu}\left[\sum_{t=0}^{\infty} \alpha^t C(u_{\hat{s}^{(t+1)}}^{(t)}, \hat{s}^{(t+1)})\right]$$

The principal aims to minimize the cost of delivering $\mathbf{v}$ to the agent:

$$\text{[SP]} \qquad P^*(\mathbf{v},\mu) := \inf_{\tilde{u}\in\Pi(\mathbf{v})} R(\tilde{u},\sigma^*,\mu)$$

We refer to this as the principal's *sequential problem*. A sequential contract is *sequentially optimal* if it attains the infimum in [SP].

A.2. Agent Transversality in the Recursive Problem

The following lemma shows that [TVC]-implementable recursive contracts in fact deliver promises, are fully incentive compatible (i.e., deter infinite-length deviations), and induce allocations that are feasible in the principal's sequential problem [SP]. Consequently, this mild "continuity at infinity" condition renders the sequential and recursive problems, [SP] and [RP], equivalent.

Lemma A.1. If a recursive contract ξ is [TVC]-implementable at $\mathbf{v} \in D$, then:

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

$$\text{[DP]} \qquad v_i = \mathsf{E}\left[\sum_{t=0}^{\infty} \alpha^t \tilde{u}_\xi^{(t)} \mid s^{(0)} = s_i\right]$$

(b) Truthtelling after *every* history is an optimal strategy for the agent.

(c) The induced allocation $\tilde{u}_\xi$ is feasible in [SP].

Conversely, if the induced allocation $\tilde{u}_\xi$ is feasible in [SP], then the recursive contract ξ delivers promises (i.e., satisfies [DP]).

Proof of Lemma A.1. The converse statement at the end of the lemma is obvious. Part (c) follows from parts (a) and (b). For part (a), let $\xi \in \Xi^*(\mathbf{v})$. Iterating forward T times on the recursive promise keeping constraints [PK$_i$] gives

$$v_i = \mathsf{E}\left[\sum_{t=0}^{T-1} \alpha^t \tilde{u}_\xi^{(t)} \mid s^{(0)} = i\right] + \mathsf{E}\left[\alpha^T \, \mathsf{E}^{\mathbf{f}}{}_{s^{(T+1)}}\left[\mathbf{v}_\xi^{(T+1)}\right] \mid s^{(0)} = i\right]$$

for all $i \in S$. Sending $T \to \infty$, 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], as desired. The proof of part (b) is analogous (cf. Green 1987, Lemma 2). □

B. Structural Results Omitted from Section 3

B.1. Recursive Domain

This appendix presents a formal characterization of the principal's (largest) recursive domain D, as defined in Subsection 3.1. We require the following definitions. First, let $\Xi^*(\mathbf{v}) \subseteq \Xi(\mathbf{v})$ denote the set of feasible recursive contracts that are initialized and [**TVC**]-implementable at $\mathbf{v} \in D$, and let

[**B.1**] $$D^* := \{\mathbf{v} \in D : \ \Xi^*(\mathbf{v}) \neq \varnothing \}$$

denote the set of contingent promises that can be generated by some [**TVC**]-implementable contract. Second, consider the following class of type processes:

Definition B.1. The type process is *MLRP-ordered* (or simply MLRP) if the transition probabilities are non-decreasing in the monotone likelihood ratio order, i.e., if the ratio f_{ki}/f_{kj} is non-decreasing in k whenever $i > j$.

When $d \geq 3$, MLRP, like PPR, is stronger than FOSD but satisfied by many type processes considered in applications, such as discretized AR(1) processes. In general, MLRP and PPR have nontrivial intersection. When $d = 2$, all FOSD processes are both MLRP and PPR.

Theorem 3. *Fix $d > 1$ and define the set $V_d := \left\{\mathbf{v} \in \mathcal{U}^d : v_d > v_{d-1} > \cdots > v_1\right\}$.*

(a) There exists a largest recursive domain D. It is a non-empty, convex, and open cone in $\mathcal{U}^d$ that satisfies $V_d \subseteq D$. For fixed type process, D is independent of the discount factor $\alpha \in (0,1)$ and the utility function U (within the class allowed for by Assumption DARA).

(b) The constraint correspondence $\Gamma : D \to (\mathcal{U} \times D)^d$ is nonempty-valued and has a convex graph.

(c) $D^* \subseteq D$ is nonempty, convex, has decreasing returns (i.e., if $\mathbf{v} \in D$, then $a\mathbf{v} \in D$ for all $a \in (0,1]$), and is unbounded below (i.e., for all $k < 0$ there exists some $\mathbf{v} \in D$ such that $\mathbf{v} \leq k\mathbf{1}$).

(d) If the type process satisfies FOSD, then $D^* \subseteq V_d$.

(e) If the type process satisfies either MLRP or PPR, then $D = V_d$.

(f) If the type process satisfies either MLRP or PPR and, in addition, the agent has CARA utility, then $D = V_d = D^*$.

The proof of Theorem 3(d) is in Online Appendix D. The proofs of all other parts are lengthy, but otherwise standard, and can be found in Appendix G of BKL21. Parts (b)–(c) and most of part (a) are established by characterizing fixed points of a set-valued operator via Tarski's theorem. The cone property stated in part (a), as well as parts (e)–(f), are established by characterizing solutions to a system of linear programs.

It is worth noting that different recursive formulations of the contracting problem would yield different recursive domains. For instance, in the well-known Fernandes and Phelan (2000) recursive formulation based on *ex ante* promised and threat-point utilities, the recursive domain is actually a set-valued *function* of past reports, the shape of which changes depending on the transition probabilities of the agent's type. By contrast, in our formulation based on type-contingent promised utilities, the domain is independent of the agent's reports and, at least for type processes in the MLRP and PPR classes, is also independent of the transition probabilities.[50] These features are important for the tractability of our analysis.

B.2. Regularity

The following lemma shows that the Regularity conditions in Definition 3.2 can be verified in terms of model primitives in some cases of special interest.

Lemma B.2. The following hold:

(a) If either (i) the type process satisfies FOSD or (ii) the transition probabilities $\{\mathbf{f}_i\}_{i\in S}$ are affinely independent, then Regularity Condition R.3 holds.

(b) If the agent has CARA utility and the type process satisfies MLRP or PPR, then the environment is Regular.

The proof of Lemma B.2 is in Appendix H of BKL21. Point (i) of part (a) includes all MLRP, PPR, and iid type processes, as well as most others considered in applications. While the affine independence condition in point (ii) of part (a) is violated in the iid case, for any fixed $d \geq 2$ it is satisfied by a generic set of transition matrices. Part (b) implies that the

(50) For example, consider the case in which $d = 2$ and types are MLRP. In our formulation, Theorem 3(e) establishes that $D = V_2$, as illustrated in Figure 1. In Fernandes and Phelan (2000), the recursive state variable consists of the previous report s, the corresponding *ex ante* promised utility $v^p(s) := \mathsf{E}^{\mathbf{f}_s}[\mathbf{v}]$ for an agent whose report of s was truthful, and the corresponding *threat point* utility $v^\dagger(s) := \mathsf{E}^{\mathbf{f}_{3-s}}[\mathbf{v}]$ for an agent whose report of s was a lie. In this formulation, the domain is a correspondence $W : S \rightrightarrows \mathcal{U}^2$, where $W(s)$ consists of implementable $(v^p(s), v^\dagger(s))$ pairs, and the collection of sets $\{W(s)\}_{s\in S}$ must be solved for jointly. From the definition of $(v^p(s), v^\dagger(s))$ and our Theorem 3(e), it is easy to see that $W(1) = \left[\begin{smallmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{smallmatrix}\right] V_2$ and $W(2) = \left[\begin{smallmatrix} f_{21} & f_{22} \\ f_{11} & f_{12} \end{smallmatrix}\right] V_2$. Thus, $W(1)$ and $W(2)$ are cones and are symmetric about the diagonal in $\mathbb{R}^2_{--}$, and their shapes depend on the transition probabilities. For instance, if types are iid, then $W(1) = W(2) = \{(t, t) : t < 0\}$. With positive serial correlation, they are non-degenerate cones.

solved binary-state example described in Subsection 5.3 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 (2020) is Regular.

References

Allen, Franklin. 1985. 'Repeated principal-agent relationships with lending and borrowing'. *Economics Letters* 17 (1-2): 27–31.

Arellano, Manuel, Richard Blundell and Stéphane Bonhomme. 2017. 'Earnings and consumption dynamics: a nonlinear panel data framework'. *Econometrica* 85 (3): 693–734.

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.

Atkeson, Andrew, and Robert E. Lucas Jr. 1992. 'On Efficient Distribution with Private Information'. *Review of Economic Studies* 59 (3): 427–453.

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.

Battaglini, Marco, and Rohit Lamba. 2019. 'Optimal Dynamic Contracting: The First-Order Approach and Beyond'. *Theoretical Economics* 14 (4): 1435–1482.

Bloedel, Alexander W., and R. Vijay Krishna. 2015. *Misery, Persistence, and Growth*. Tech. rep. Stanford University.

Bloedel, Alexander W., R. Vijay Krishna and Oksana Leukhina. 2021. 'Insurance and Inequality with Persistent Private Information'. Federal Reserve Bank of St. Louis Working Paper, https://research.stlouisfed.org/wp/more/2018-020.

Bloedel, Alexander W., R. Vijay Krishna and Bruno Strulovici. 2021. *On the Limits of Dynamic Screening under Permanent Shocks*. Tech. rep.

– . 2020. *Persistent Private Information Revisited*. Tech. rep. Stanford University.

Broer, Tobias, Marek Kapička and Paul Klein. 2017. 'Consumption risk sharing with private information and limited enforcement'. *Review of Economic Dynamics* 23:170–190.

Cole, Harold L., and Narayana R. Kocherlakota. 2001a. 'Dynamic Games with Hidden Actions and Hidden States'. *Journal of Economic Theory* 98 (1): 114–126.

Cole, Harold L, and Narayana R Kocherlakota. 2001b. 'Efficient allocations with hidden income and hidden storage'. *The Review of Economic Studies* 68 (3): 523–542.

Di Tella, Sebastian, and Yuliy Sannikov. 2021. 'Optimal Asset Managment Contracts with Hidden Savings'. *Econometrica* 89 (3): 1099–1139.

Doepke, Matthias, and Robert M Townsend. 2006. 'Dynamic mechanism design with hidden income and hidden actions'. *Journal of Economic Theory* 126 (1): 235–285.

Farhi, Emmanuel, and Iván Werning. 2007. 'Inequality and Social Discounting'. *Journal of Political Economy* 115 (3): 365–402.

– . 2013. 'Insurance and Taxation over the Lifecycle'. *Review of Economic Studies* 80 (2): 596–635.

Farhi, Emmanuel, et al. 2012. 'Non-linear capital taxation without commitment'. *Review of Economic Studies* 79 (4): 1469–1493.

Feldman, Naomi E, and Joel Slemrod. 2007. 'Estimating tax noncompliance with evidence from unaudited tax returns'. *The Economic Journal* 117 (518): 327–352.

Fernandes, Ana, and Christopher Phelan. 2000. 'A Recursive Formulation for Repeated Agency with History Dependence'. *Journal of Economic Theory* 91:223–247.

Fu, Shiming, and R. Vijay Krishna. 2019. 'Dynamic Financial Contracting with Persistence'. *RAND Journal of Economics* 50 (2): 418–452.

Garrett, Daniel, Alessandro Pavan and Juuso Toikka. 2018. *Robust Predictions in Dynamic Screening*. Tech. rep. Northwestern University.

Golosov, Mikhail, Narayana R. Kocherlakota and Aleh Tsyvinski. 2003. 'Optimal Indirect and Capital Taxation'. *Review of Economic Studies* 70 (3): 569–588.

Golosov, Mikhail, Maxim Troshkin and Aleh Tsyvinski. 2016. 'Redistribution and Social Insurance'. *American Economic Review* 106 (2): 359–386.

Golosov, Mikhail, and Aleh Tsyvinski. 2007. 'Optimal taxation with endogenous insurance markets'. *The Quarterly Journal of Economics* 122 (2): 487–534.

Golosov, Mikhail, Aleh Tsyvinski and Nicolas Werquin. 2016. 'Recursive Contracts and Endogenously Incomplete Markets'. In *Handbook of Macroeconomics, Volume 2*.

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.

Green, Jerry R, and Jean-Jacques Laffont. 1986. 'Partially verifiable information and mechanism design'. *The Review of Economic Studies* 53 (3): 447–456.

Guo, Yingni, and Johannes Hörner. 2020. *Dynamic Allocation without Money*. Tech. rep.

Guvenen, Fatih, et al. 2021. 'What Do Data on Millions of US Workers Reveal About Lifecycle Earnings Dynamics?' *Econometrica* 89 (5): 2303–2339.

Halac, Marina, and Pierre Yared. 2014. ‘Fiscal rules and discretion under persistent shocks’. *Econometrica* 82 (5): 1557–1614.

Hatcher, Allen. 2001. *Algebraic Topology*. Cambridge University Press.

Hörner, Johannes, Xiaosheng Mu and Nicolas Vielle. 2017. *Keeping Your Story Straight: Truthtelling and Liespotting*. Tech. rep. Yale University.

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.

Kapička, Marek. 2013. ‘Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach’. *The Review of Economic Studies*: rds045.

Khan, Aubhik, and Bashkar Ravikumar. 2001. ‘Growth and risk-sharing with private information’. *Journal of Monetary Economics* 47 (3): 499–521.

Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press.

Krasikov, Ilia, and Rohit Lamba. 2021. ‘A Theory of Dynamic Contracting with Financial Constraints’. *Journal of Economic Theory* 193.

Ljunqvist, Lars, and Thomas J. Sargent. 2012. *Recursive Macroeconomic Theory*. 3rd ed. Cambridge, Ma: MIT Press.

Lucas, Robert E., Jr. 1992. ‘On Efficiency and Distribution’. *Economic Journal* 102 (411): 233–247.

Luenberger, David G. 1969. *Optimization by Vector Space Methods*. Wiley.

Meghir, Costas, and Luigi Pistaferri. 2004. ‘Income variance dynamics and heterogeneity’. *Econometrica* 72 (1): 1–32.

Newman, Andrew F. 2007. ‘Risk-bearing and entrepreneurship’. *Journal of Economic Theory* 137 (1): 11–26.

Olszewski, Wojciech, and Mikhail Safronov. 2021. *Efficient dynamic allocations in endowment economies with linear utilities and private information*. Tech. rep.

Pavan, Alessandro. 2016. ‘Dynamic mechanism design: Robustness and endogenous types’. In *2015 World Congress of the Econometric Society*.

Pavan, Alessandro, Ilya Segal and Juuso Toikka. 2014. ‘Dynamic Mechanism Design: A Myersonian Approach’. *Econometrica* 82 (2): 601–653.

Phelan, Christopher. 1998. ‘On the Long Run Implications of Repeated Moral Hazard’. *Journal of Economic Theory* 79:174–191.

– . 2006. ‘Opportunity and social mobility’. *Review of Economic Studies* 73 (2): 487–504.

– . 1995. ‘Repeated Moral Hazard and One-Sided Commitment’. *Journal of Economic Theory* 66:488–506.

Shiryaev, Albert N. 1995. *Probability*. Springer.

Sleet, Christopher, and Şevin Yeltekin. 2006. ‘Credibility and endogenous societal discounting’. *Review of Economic Dynamics* 9 (3): 410–437.

– . 2008. ‘Politically credible social insurance’. *Journal of monetary Economics* 55 (1): 129–151.

Spear, Stephen E, and Sanjay Srivastava. 1987. ‘On Repeated Moral Hazard with Discounting’. *Review of Economic Studies* 54 (4): 599–617.

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.

Strulovici, Bruno. 2017. *Contract Negotiation and Screening with Persistent Information*. Tech. rep. Northwestern University.

– . 2020. ‘Renegotiation-proof contracts with persistent private information’. Working paper.

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.

Wang, Cheng. 1995. ‘Dynamic insurance with private information and balanced budgets’. *The Review of Economic Studies* 62 (4): 577–595.

Williams, Noah. 2011. ‘Persistent Private Information’. *Econometrica* 79 (4): 1233–1275.

Zhang, Yuzhe. 2009. ‘Dynamic Contracting with Persistent Shocks’. *Journal of Economic Theory* 144:635–675.

Online Appendix

This Online Appendix (henceforth OA) provides proofs omitted from the main text. OA-C first presents the proof of Theorem 1, essential steps of which include the proofs of Proposition 4.3 and Theorem 3(d). OA-D then presents the proof of Theorem 2. Finally, OA-E and OA-F collect facts — about the first-best optimal contract and pathwise properties of Markov chains, respectively — that are used in the preceding proofs. While the arguments in this OA are mostly self-contained, the working paper (Bloedel, Krishna and Leukhina 2021, henceforth BKL21) contains proofs of some auxiliary technical results referenced herein.

C. Proof of Theorem 1

This OA is divided into several parts. OA-C.1 presents the Lagrangian and first-order optimality conditions derived from the Bellman equation [**FE**]. OA-C.2 proves Proposition 4.3 regarding the marginal cost martingale. OA-C.3 consists of several intermediate steps towards the proof of Theorem 1. OA-C.4 presents the main proof of convergence for Theorem 1.

C.1. Optimality Conditions

Recall that the set of *recursive constraints* consists of the of the *promise keeping* conditions

$$[\mathbf{PK}_i] \qquad v_i = u_i + \alpha\, \mathsf{E}^{\mathbf{f}_i}\left[\mathbf{w}_i\right]$$

for all $i \in S$, and the *incentive compatibility* conditions

$$[\mathbf{IC}_{ij}] \qquad u_i + \alpha\, \mathsf{E}^{\mathbf{f}_i}\left[\mathbf{w}_i\right] \geq \psi(u_j, i, j) + \alpha\, \mathsf{E}^{\mathbf{f}_i}\left[\mathbf{w}_j\right]$$

for all $i, j \in S$ with $i > j$. (The incentive constraints are written here in a slightly different, but equivalent, form than in Section 3.)

Proposition 3.3 reduces the principal's problem to a smooth, convex, finite-dimensional minimization problem. Thus, under Condition R.3, standard results imply that optimal menus in [**FE**] can be characterized via saddle points of a Lagrangian function (see, e.g., Exercise 7 on p. 236 and Theorem 2 on p. 221 of Luenberger (1969)). The Lagrangian for this problem is

$$\mathscr{L}(\mathbf{v}, s, \mathbf{u}, \mathbf{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_{i=1}^{d} f_{si}\left[C(u_i, i) - \omega_i + \alpha P(\mathbf{w}_i, i)\right] + \sum_{i=1}^{d}\left[\lambda_i\left(v_i - u_i - \alpha\,\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i]\right)\right]$$
$$- \sum_{i=2}^{d}\sum_{j=1}^{i-1}\left[\mu_{ij}\left(u_i + \alpha\,\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i] - \psi(u_j, i, j) - \alpha\,\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_j]\right)\right]$$

where $\lambda_i \in \mathbb{R}$ is the multiplier on the promise keeping constraint [**PK**$_i$] and $\mu_{ij} \geq 0$ is the multiplier on the incentive constraint [**IC**$_{ij}$]. For notational ease, we extend μ_{ij} to all pairs $i, j \in \mathbb{N}$, with the understanding that $\mu_{ij} = 0$ if $j \geq i$, $i \notin S$, or $j \notin S$.

The necessary and sufficient first-order optimality equations consist of the envelope conditions

[**Env**$_i$] $$P_i(\mathbf{v}, s) = \lambda_i$$

for $i \in S$, the first-order conditions for flow utilities

[**FOC**u_i] $$f_{si}C'(u_i, i) = \lambda_i + \sum_{k=1}^{i-1}\mu_{ik} - \sum_{k=i+1}^{d}\psi'(u_i, k, i)\mu_{ki}$$

for $i \in S$, and the first-order conditions for contingent continuation utilities

[**FOC**$\mathbf{w}_{ij}$] $$f_{si}P_j(\mathbf{w}_i, i) = f_{ij}\left(\lambda_i + \sum_{k=1}^{i-1}\mu_{ik}\right) - \sum_{k=i+1}^{d} f_{kj}\mu_{ki}$$

for $i, j \in S$ with $i > j$, and the usual complementary slackness conditions (which we omit).

C.2. Proof of Proposition 4.3

Suppose the environment is Regular. Because P is continuously differentiable by Proposition 3.3, the directional derivative is linear, implying that $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}, s) = \sum_{i \in S} P_i(\mathbf{v}, s)$, and $P_i(\cdot, \cdot)$ is real-valued on $D \times S$ for each $i \in S$. For each $t \in \mathbb{N}$, integrability of the random variable $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}^{(t)}, s^{(t)})$ then follow from non-negativity and finiteness of the directional derivative and finiteness of S.

As for the martingale property, summing the [**Env**$_i$] over $i \in S$ delivers

[**C.1**] $$\mathsf{D}_{\mathbf{1}}P(\mathbf{v}, s) = \sum_{i=1}^{d}\lambda_i$$

For fixed $i \in S$, summing the [FOC$\mathbf{w}_{ij}$] over $j \in S$ gives

$$f_{si} \cdot \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i, i) = \lambda_i + \sum_{k=1}^{i-1} \mu_{ik} - \sum_{k=i+1}^{d} \mu_{ki} \tag{C.2}$$

Now, summing the above display over $i \in S$ and noting that $\sum_{i=1}^{d} \sum_{k=1}^{i-1} \mu_{ik} = \sum_{i=1}^{d} \sum_{k=i+1}^{d} \mu_{ki}$ delivers $\sum_{i \in S} f_{si} \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i, i) = \sum_{i=1}^{d} \lambda_i$ which, combined with [C.1], gives the martingale property $\mathsf{D}_{\mathbf{1}} P(\mathbf{v}, s) = \sum_{i=1}^{d} f_{si} \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i, i)$

Thus, the directional derivative process defines a martingale. It remains to prove that the marginal cost martingale is non-negative/strictly positive. This is implied by the following lemma.

Lemma C.1. Suppose that the environment is Regular. Then the directional derivative $\mathsf{D}_{\mathbf{1}} P(\cdot, s)$ is non-negative for each $s \in S$. If the environment is [TVC]-Regular, then it is strictly positive.

Proof. If $P(\cdot, s)$ were non-decreasing in the direction $\mathbf{1} \in \mathbb{R}^d$, non-negativity of the directional derivative $\mathsf{D}_{\mathbf{1}} P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} \frac{P(\mathbf{v}+\varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)}{\varepsilon}$ would follow directly from the definition. Hence, it suffices to show that $P(\cdot, s)$ is non-decreasing in direction $\mathbf{1}$. Lemma E.2 shows that the first-best value function Q^* is non-decreasing in this direction. We will show that P inherits this property from Q^*. The proof is order-theoretic.[51]

Let $[Q^*, P]$ denote the order interval (in the pointwise order) of functions $Q : D \times S \to \mathbb{R}$ that lie weakly above Q^* and weakly below P. (P is real-valued under Condition R.1, so this order interval is well-defined.) Let $\Phi := \{Q \in [Q^*, P] : Q(\mathbf{v}, s) \geq Q(\mathbf{v} - \varepsilon \mathbf{1}, s) \, \forall \varepsilon > 0\}$. That is, Φ consists of all functions in the order interval $[Q^*, P]$ with the property that they are non-decreasing in the direction $\mathbf{1}$.

Claim 4. Φ is a lattice in the pointwise order.

Proof of Claim. It is easy to see that if $F, G \in \Phi$, then $F \vee G, F \wedge G \in [Q^*, P]$. Now, fix $(\mathbf{v}, s) \in D \times S$ and $\varepsilon > 0$. (We may take $\varepsilon > 0$ sufficiently small that all perturbed vectors defined below are in D, as D is open by part (a) of Theorem 3.) If F and G are ordered the same way at $(\mathbf{v}, s)$ and $(\mathbf{v} + \varepsilon \mathbf{1}, s)$, there is nothing left to prove. So suppose, without loss of generality, that $F(\mathbf{v}, s) \geq G(\mathbf{v}, s)$ and $G(\mathbf{v} + \varepsilon \mathbf{1}, s) \geq F(\mathbf{v} + \varepsilon \mathbf{1}, s)$. Then, $(F \wedge G)(\mathbf{v} + \varepsilon \mathbf{1}, s) = F(\mathbf{v} + \varepsilon \mathbf{1}, s) \geq F(\mathbf{v}, s) \geq (F \wedge G)(\mathbf{v}, s)$. Similarly, $(F \vee G)(\mathbf{v} + (\varepsilon \mathbf{1}, s) \geq F(\mathbf{v} + \varepsilon \mathbf{1}, s) \geq F(\mathbf{v}, s) = (F \vee G)(\mathbf{v}, s)$ which concludes the proof. □

(51) In particular, it does *not* rely on convergence of $T^n Q^*$, the n-fold iterate of the Bellman operator [T] on Q^*, to P in countably-many steps. We are unable to show that such convergence takes place in general.

Claim 5. The lattice Φ is complete.

Proof of Claim. Let $F \subseteq \Phi$ be nonempty and define $\overline{f}(\mathbf{v}, s) := \sup_{f \in F} f(\mathbf{v}, s)$ and $\underline{f}(\mathbf{v}, s) := \inf_{f \in F} f(\mathbf{v}, s)$ for each $(\mathbf{v}, s) \in V_d \times S$. We show that $\overline{f} \in F$; the proof for $\underline{f}$ is symmetric. Suppose towards a contradiction that there exists $(\mathbf{v}, s) \in D \times S$ and some $\varepsilon > 0$ such that $(\mathbf{v}', s) \in D \times S$, where $\mathbf{v}' = \mathbf{v} + \varepsilon\mathbf{1}$, and $\overline{f}(\mathbf{v}, s) > \overline{f}(\mathbf{v}', s)$. Because every function in Φ is bounded by $[Q^*, P]$, both of which are finite, this implies that there exists some $\delta > 0$ such that $\overline{f}(\mathbf{v}, s) - \delta \geq \overline{f}(\mathbf{v}', s)$. By definition of the supremum, there exists some $f \in F$ such that $f(\mathbf{v}, s) > \overline{f}(\mathbf{v}, s) - \delta$. Combined with the earlier inequality and the definition of Φ, this implies that $f(\mathbf{v}, s) > \overline{f}(\mathbf{v}'s,) \geq f(\mathbf{v}', s)$ which contradicts the fact that f is non-decreasing in the direction $\mathbf{1}$ by virtue of $f \in F \subset \Phi$. □

Let $\overline{\mathbb{R}}$ denote the extended reals, and let $\overline{\mathbb{R}}^{V_d \times S}$ denote the space of functions $f : D \times S \to \overline{\mathbb{R}}$. Define the Bellman operator $T : \overline{\mathbb{R}}^{D\times S} \to \overline{\mathbb{R}}^{D\times S}$ by

$$\text{[T]} \qquad TQ(\mathbf{v}, s) := \inf_{(u_i, \mathbf{w}_i)_{i\in S} \in \Gamma(\mathbf{v})} \sum_{i=1}^{d} f_{si} \left[C(u_i, i) + \alpha Q(\mathbf{w}_i, i)\right]$$

Claim 6. $T : \Phi \to \Phi$ is well-defined and monotone.

Proof of Claim. Monotonicity is standard. It is easy to see that for $Q \in \Phi$, $TQ \in [Q^*, P]$. All that remains is to show that $TQ \in \Phi$.

To see this, fix $(\mathbf{v}, s) \in D \times S$ and $\delta > 0$, and let $(u_i, \mathbf{w}_i)_{i\in S}$ be a δ-optimal pair for the Bellman operator. Then, $\delta + TQ(\mathbf{v}, s) \geq \sum_{i\in S} f_{si}\big[C(u_i, i) + \alpha Q(\mathbf{w}_i, i)\big] \geq \sum_{i\in S} f_{si}\big[C(u_i, i) + \alpha Q(\mathbf{w}_i - \frac{\varepsilon}{\alpha}\mathbf{1}, i)\big] \geq TQ(\mathbf{v} - \varepsilon\mathbf{1}, s)$ where the first inequality uses the fact that $Q \in \Phi$ and the second inequality follows because $(u_i, \mathbf{w}_i - \varepsilon/\alpha \cdot \mathbf{1})_{i\in S} \in \Gamma(\mathbf{v} - \varepsilon\mathbf{1})$. It follows that $TQ \in \Phi$, since $\delta > 0$ was arbitrary. This proves that T is well defined. □

Now, because Φ is a complete lattice by Claims 1 and 2 and T is well-defined and monotone on Φ by Claim 3, it follows from Tarski's Fixed Point Theorem that the Bellman operator T has a fixed point in Φ. Let $\hat{P}$ be the smallest fixed point of T in Φ. As $\Phi \subseteq [Q^*, P]$, it follows from Proposition 3.3 that P is this smallest fixed point.

This establishes that the directional derivative is non-negative when the environment is Regular. To see that it is strictly positive under Condition R.4, suppose there exists $(\mathbf{v}, s) \in D \times S$ such that $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}, s) = 0$. For this fixed state, define the function $f : Y \to \mathbb{R}$ by $f(y) := P(\mathbf{v} - y\mathbf{1}, s)$ where $Y \subset \mathbb{R}_+$ has nonempty interior and is small enough that $\mathbf{v} - y\mathbf{1} \in D$ for all $y \in Y$. The function $f(\cdot)$ is strictly concave because P is strictly convex under Condition R.4 by Proposition 3.3, and is non-increasing by Lemma C.1. Hence, the hypothesis that $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}, s) = 0$ implies that $f(y) \equiv f(0)$ for all $y \in Y$. But this contradicts strict convexity of P and is therefore impossible. It follows that $\mathsf{D}_{\mathbf{1}}P$ is strictly positive on $D \times S$. □

C.3. Intermediate Steps Towards the Proof of Theorem 1

This OA consists of several parts, culminating in Lemmas C.13, C.14, and C.15, which establish that (a) the marginal cost martingale necessarily "splits" after consecutive high-endowment shocks and that (b) if this were not the case, the optimal contract would in fact implement the first-best solution from Subsection 3.1 (which is impossible). These lemmas are key inputs into the main proof of Theorem 1. To that end, OA-C.3.1 first establishes some preliminary facts about the efficiency problem [**Eff**$_i$] from Subsection 4.3. OA-C.3.2 then establishes that the optimal contract is efficient (i.e., solves [**Eff**$_i$]) after consecutive high-endowment shocks, and records important properties of the marginal cost martingale after such shocks. OA-C.3.3 introduces a reformulation of the principal's recursive that allows us to relate policy functions and optimal Lagrange multipliers across different values of the previous report s. Finally, OA-C.3.4 uses the results of OA-C.3.2 and OA-C.3.3 to establish Lemmas C.13, C.14, and C.15.

C.3.1. *The Efficiency Problem*

Recall the *efficiency problem* [**Eff**$_i$] from Subsection 4.3, re-stated here for convenience:

[**Eff**$_i$]
$$K(w, i) := \min_{\mathbf{w}_i \in D} P(\mathbf{w}_i, i)$$
$$\text{s.t.} \qquad \mathsf{E}^{\mathbf{f}_i}[\mathbf{w}_i] \geq w$$

Lemma C.2. Suppose that the environment is [**TVC**]-Regular. Then the following hold:

(a) For each $i \in S$ and $w \in \mathcal{U}$, the efficiency problem has a unique solution $\mathbf{w}^{\dagger}(w)$.

(b) For each $i \in S$, the policy function $\mathbf{w}^{\dagger}(\cdot, i) : \mathcal{U} \to \mathcal{U}^d$ is continuous.

(c) For each $i \in S$, the value function $K(\cdot, i)$ is well-defined, finite-valued, strictly increasing, strictly convex, continuously differentiable, and satisfies the Inada conditions $\lim_{w \to -\infty} K'(w, i) = 0$ and $\lim_{w \to 0} K'(w, i) = +\infty$.

Proof. Parts (a), (b), and part (c) — aside from the claimed strict monotonicity and Inada conditions — follow from the same arguments used to establish Theorem H.1 of BKL21. Proofs of these properties are thus omitted.

For the strict monotonicity in part (c), let $w \in \mathcal{U}$ and $w' := w - \varepsilon$ for some $\varepsilon > 0$. Clearly $\mathbf{w}' := \mathbf{w}^{\dagger}(w, i) - \varepsilon\mathbf{1}$ is feasible in [**Eff**$_i$] at (w', i) and, by Proposition 4.3 (or Lemma C.1), $P(\mathbf{w}', i) < K(w, i)$. Thus, by revealed preference, $K(w', i) < K(w, i)$, which proves strict monotonicity.

Next, the Inada conditions in part (c) follow from properties of the value function for the analogue of [**Eff**$_i$] in the full-information problem, which is defined as [**Eff**$_i^{\mathbf{FB}}$] in OA-E. That value function is called $K^*(w,i)$, and clearly satisfies $K^* \leq K$ on $\mathscr{U} \times S$. Lemma E.3 states that $\lim_{w\to 0} K^{*\prime}(w,i) = +\infty$. If $K(\cdot,i)$ did not satisfy $\lim_{w\to 0} K'(w,i) = +\infty$, then there would exist some $v \in \mathscr{U}$ such that $K^*(v,i) > K(v,i)$, a contradiction. Similarly, Lemma E.3 states that $\lim_{w\to -\infty} K^{*\prime}(w,i) = 0$. If $K(\cdot,i)$ did not satisfy $\lim_{w\to -\infty} K'(w,i) = 0$, then there would exist some $v \in \mathscr{U}$ such that $K^*(v,i) > K(v,i)$, again a contradiction. This completes the proof. □

For each $i \in S$, define the set $E_i := \{\mathbf{v} \in D : \mathbf{v} = \mathbf{w}^\dagger(w,i) \text{ for some } w \in \mathscr{U}\}$. Note that $E_i = \mathbf{w}^\dagger(\mathscr{U}, i)$, be the range of efficient solutions given past report $i \in S$.

The efficiency problem [**Eff**$_i$] admits a Lagrangian $\mathscr{L}^E(w,i,\zeta,\mathbf{w}) = P(\mathbf{w},i) - \zeta \cdot (\mathsf{E}^{\mathbf{f}_i}[\mathbf{w}] - w)$ where $\zeta \geq 0$. By Lemma C.2, the unique solution to [**Eff**$_i$] is characterized by the first-order conditions

[**FOC**$_j$**-Eff**$_i$] $$P_j(\mathbf{w}^\dagger(w,i),i) = z(w,i) f_{ij}$$

and the envelope condition

[**Env**$_j$**-Eff**$_i$] $$K'(w,i) = \zeta(w,i)$$

where $z(w,i) \geq 0$ denotes the optimal multiplier. It is easy to see from this and part (c) of Lemma C.2 (namely, continuous differentiability and the Inada conditions) that

[$\tilde{\mathbf{E}}_i$] $$\tilde{E}_i := \left\{ (P_1, \ldots, P_d) \in \mathbb{R}^d_{++} : \frac{P_1}{f_{i1}} = \cdots = \frac{P_d}{f_{id}} \right\}$$

is the image of E_i under the derivative mapping $\mathsf{D}P(\cdot,i)$. Moreover, by summing the first-order conditions [**FOC**$_j$**-Eff**$_i$] over $j \in S$ and combining with the envelope condition [**Env**$_j$**-Eff**$_i$], we get

[**C.3**] $$K'(w,i) = \mathsf{D}_\mathbf{1} P(\mathbf{w}^\dagger(w,i),i)$$

C.3.2. *Facts Concerning the Marginal Cost Martingale*

Lemma C.3. Suppose the environment is [**TVC**]-Regular. Let $s \in S$ be given, and define $Y_s := \mathsf{D}P(D,s) \subseteq \mathbb{R}^d$ to be the image of D under the derivative mapping $\mathsf{D}P(\cdot,s)$. Then, the mapping $\mathsf{D}P(\cdot,s) : D \to Y_s$ is a homeomorphism.

Proof. We first show that $\mathsf{D}\,P(\cdot,s)$ is injective. To see this, notice first that because P is strictly convex, $\mathsf{D}\,P(\cdot,s)$ is *strictly monotone* in the sense that for all $\mathbf{v},\mathbf{v}'\in D$ such that $\mathbf{v}\neq\mathbf{v}'$, $\langle \mathsf{D}\,P(\mathbf{v},s)-\mathsf{D}\,P(\mathbf{v}',s),\mathbf{v}-\mathbf{v}'\rangle>0$. But now suppose $\mathsf{D}\,P(\cdot,s)$ is not injective, so that there are $\mathbf{v},\mathbf{v}'\in D$ distinct such that $\mathsf{D}\,P(\mathbf{v},s)=\mathsf{D}\,P(\mathbf{v}',s)$. But this would imply that $0=\langle \mathbf{0},\mathbf{v}-\mathbf{v}'\rangle=\langle \mathsf{D}\,P(\mathbf{v},s)-\mathsf{D}\,P(\mathbf{v}',s),\mathbf{v}-\mathbf{v}'\rangle>0$, which is a contradiction. As D is open by Theorem 3 and the derivative $\mathsf{D}P(\cdot,s)$ is continuous on D by Proposition 3.3, it follows from Brouwer's Invariance of Domain Theorem (e.g., Hatcher 2001, Theorem 2B.3) that D is homeomorphic to $\mathsf{D}\,P(D,s)=Y_s$. □

Lemma C.4. Suppose the environment is [**TVC**]-Regular. Let $(\mathbf{v},s)\in D\times S$ be given. Then $\mathbf{w}_d(\mathbf{v},s)\in E_d$.

Proof. Follows immediately from the the optimality conditions laid out in OA-C.1, the definition of the set $\tilde{E}_d$ in [$\tilde{\mathbf{E}}_i$], and Lemma C.3. □

Remark C.1. By comparing the optimality conditions for [**RP**] (enumerated in OA-C.1) with the optimality conditions for [**Eff**$_i$] stated above, it is easy to see — using the same logic as in the proof of Lemma C.4 above — that *when types are iid* $\mathbf{w}_i(\mathbf{v},s)\in E_i$ for all $i,s\in S$ and $\mathbf{v}\in D$. That is, the optimal contract is always *efficient* in the iid case.

Lemma C.5. Suppose the environment is [**TVC**]-Regular. Let $s\in S$ be given. Then, the function $\mathsf{D}_{\mathbf{1}}P(\cdot,s):E_s\to\mathbb{R}_{++}$ is strictly increasing and is a homeomorphism.

Proof. This follows immediately from [**C.3**] and part (c) of Lemma C.2. □

Lemma C.6. Suppose the environment is [**TVC**]-Regular. Let $(\mathbf{v},s)\in D\times S$ be given, and define $\mathbf{w}_d:=\xi^c((\mathbf{v},s),d)$. For all $i\in S$, define $\tilde{\mathbf{w}}_i:=\xi^c((\mathbf{w}_d,d),i)$. Then, for all $i\in S$ we have

[**MS**$_i$]
$$\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i,i)=\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d,d)+\frac{\sum_{k=1}^{i-1}\mu_{ik}(\mathbf{w}_d,d)}{f_{di}}-\frac{\sum_{k=i+1}^{d}\mu_{ki}(\mathbf{w}_d,d)}{f_{di}}$$

Proof. Begin with the case $i=d$, which corresponds to consecutive ω_d realizations. From the optimality conditions, we have $f_{dd}\left(\lambda_d(\mathbf{v},s)+\sum_{k=1}^{d-1}\mu_{dk}(\mathbf{v},s)\right)=f_{sd}P_d(\mathbf{w}_d,d)=f_{sd}\lambda_d(\tilde{\mathbf{w}}_d,d)$ where the first equality is the FOC for w_{dd} at state $(\mathbf{v},s)$ and the second equality follows from the dth envelope condition at state $(\mathbf{w}_d,d)$. It follows that

[**C.4**]
$$f_{sd}\left[\lambda_d(\mathbf{w}_d,d)+\sum_{k=1}^{d-1}\mu_{dk}(\mathbf{w}_d,d)\right]=f_{dd}\left[\lambda_d(\mathbf{v},s)+\sum_{k=1}^{d-1}\mu_{dk}(\mathbf{v},s)\right]+f_{sd}\sum_{k=1}^{d-1}\mu_{dk}(\mathbf{w}_d,d)$$

Summing over the FOCs for $\{\mathbf{w}_{dj}\}_{j=1}^{d}$ at state $(\mathbf{v}, s)$, we obtain

$$f_{sd}\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d) = \lambda_d(\mathbf{v}, s) + \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{v}, s) \tag{C.5}$$

Similarly, summing over the FOCs for $\{\tilde{\mathbf{w}}_{dj}\}_{j=1}^{d}$ at state $(\mathbf{w}_d, d)$, we obtain

$$f_{dd}\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_d, d) = \lambda_d(\mathbf{w}_d, d) + \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d) \tag{C.6}$$

Substituting [C.5] and [C.6] into [C.4] and dividing through by $f_{sd} \cdot f_{dd}$ delivers

$$\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_d, d) = \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d) + \frac{\sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d)}{f_{dd}} \tag{C.7}$$

which is precisely [MS$_i$] for $i = d$.

Now, consider any $i < d$. Summing over the FOCs for $\{\tilde{\mathbf{w}}_{ij}\}_{j=1}^{d}$ at state $(\mathbf{w}_d, d)$, we obtain

$$f_{di}\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i, i) = \lambda_i(\mathbf{w}_d, d) + \sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d) - \sum_{k=i+1}^{d} \mu_{ki}(\mathbf{w}_d, d) \tag{C.8}$$

Now, combining [C.6] and [C.8] gives

$$\begin{aligned}\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i, i) = \mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_d, d) &- \frac{\sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d)}{f_{dd}} - \frac{\sum_{k=i+1}^{d} \mu_{ki}(\mathbf{w}_d, d)}{f_{di}} \\ &+ \frac{\sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d)}{f_{di}} - \left[\frac{\lambda_d(\mathbf{w}_d, d)}{f_{dd}} - \frac{\lambda_i(\mathbf{w}_d, d)}{f_{di}}\right]\end{aligned} \tag{C.9}$$

From the envelope conditions [Env$_i$], the bracketed term in the second line is equal to $\frac{P_d(\mathbf{w}_d,d)}{f_{dd}} - \frac{P_i(\mathbf{w}_d,d)}{f_{di}}$ and, by Lemma C.4 and [$\tilde{\text{E}}_i$], this term vanishes. Thus, plugging the first line of [C.9] into [C.7] and rearranging delivers [MS$_i$]. □

Lemma C.7. Suppose the environment is [TVC]-Regular. Let $(\mathbf{v}, s) \in D \times S$ be given, and define $\mathbf{w}_d := \xi^c((\mathbf{v}, s), d)$. For all $i \in S$, define $\tilde{\mathbf{w}}_i := \xi^c((\mathbf{w}_d, d), i)$. Then the following are equivalent:

(1) $\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i, i) = \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d)$ for all $i \in S$;

(2) $\mu_{ij}(\mathbf{w}_d, d) = 0$ for all $i, j \in S$.

Proof. From [MS$_i$] in Lemma C.6, it is easy to see that $(2) \implies (1)$. To show the converse, we proceed by induction through the type space. For the base step, let $i = d$. From [**MS**$_i$] with $i = d$ and dual feasibility (i.e., $\mu_{dk}(\mathbf{v}, s) \geq 0$ for all $k \in S$ and $(\mathbf{v}, s) \in D \times S$), it is easy to see that $\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_d, d) = \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d)$ if and only if $\mu_{dk}(\mathbf{w}_d, d) = 0$ for all $k \in S$. For the inductive step, let $i < d$ be given and suppose we have shown that $\mu_{jk}(\mathbf{w}_d, d) = 0$ for all $(j, k) \in S \times S$ such that $k < j$ and $j \geq i + 1$. Then [**MS**$_i$] reduces to

[C.10]
$$\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i, i) = \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d) + \frac{\sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d)}{f_{di}}$$

It follows from [**C.10**] and dual feasibility that $\mathsf{D}_{\mathbf{1}}P(\tilde{\mathbf{w}}_i, i) = \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_d, d)$ only if $\mu_{ik}(\mathbf{w}_d, d) = 0$ for all $k < i$. The type space S is finite, so this process terminates, establishing the converse as desired. □

C.3.3. *An "Interim" Formulation*

Consider the following *interim* formulation of the principal's recursive problem, in which she optimizes over contractual variables *contingent on* the current period's report. Given $\mathbf{v} \in D$ and for each $i \in S$, the Principal solves the i^{th} *interim problem*:

[FE-Q^i]
$$Q^i(v_i, \ldots, v_d) := \inf_{(u_i, \mathbf{w}_i) \in \mathcal{U} \times D} [C(u_i, i) + \alpha P(\mathbf{w}_i, i)]$$

subject to

[PK$_i$]
$$v_i = u_i + \alpha \, \mathsf{E}^{\mathbf{f}_i} [\mathbf{w}_i]$$

[IC$^*_{ji}$]
$$v_j - v_i \geq \psi(u_i, j, i) - u_i + \alpha \left(\mathsf{E}^{\mathbf{f}_j} [\mathbf{w}_i] - \mathsf{E}^{\mathbf{f}_j} [\mathbf{w}_j] \right)$$

for all $j \in S$ with $j > i$. That is, suppose the agent reports that he is of type $i \in S$ in the current period. Given this report, the principal optimizes over flow and continuation utilities for type i, namely $(u_i, \mathbf{w}_i) \in \mathcal{U} \times D$, subject to promise keeping [**PK**$_i$] for type i and incentive compatibility [**IC**$^*_{ji}$] for all *higher* types $j > i$. As the notation suggests, the function $Q^i(\cdot)$ depends on $\mathbf{v}$ only through the components $(v_i, v_{i+1}, \ldots, v_d)$, as these are the only components that enter the constraints. Notably, $Q^d(\cdot)$ is a function of v_d alone, and is subject only to the promise keeping constraint [**PK**$_i$] ($i = d$).

For each $i \in S$, define

[C.11]
$$\Gamma_i(\mathbf{v}) := \left\{ (u_i, \mathbf{w}_i) \in \mathcal{U} \times D : (u_i, \mathbf{w}_i) \text{ satisfies [PK}_i\text{] and [IC}^*_{ji}\text{]} \ \forall j \in S, \ j > i \right\}$$

It is easy to see that, for any $\mathbf{v} \in D$, the constraint set $\Gamma(\mathbf{v})$ (defined in [3.1]) is the Cartesian product of the $\Gamma_i(\mathbf{v})$, i.e., $\Gamma(\mathbf{v}) = \Gamma_1(\mathbf{v}) \times \cdots \times \Gamma_d(\mathbf{v})$.

Lemma C.8. Suppose the environment is Regular. The collection of functions $Q^i : D \to \mathbb{R}$ satisfy

$$P(\mathbf{v}, s) = \sum_{i=1}^{d} f_{si} Q^i(v_i, \ldots, v_d) \tag{C.12}$$

for all $(\mathbf{v}, s) \in D \times S$. Moreover, a menu $(u_i, \mathbf{w}_i)_{i \in S}$ is a minimizer in [FE] at $(\mathbf{v}, s)$ if and only if, for all $i \in S$, $(u_i, \mathbf{w}_i)$ is a minimizer in [FE-Q^i] at $\mathbf{v}$.

Proof. It follows from the observation that $\Gamma(\mathbf{v}) = \Gamma_1(\mathbf{v}) \times \cdots \times \Gamma_d(\mathbf{v})$ that we may rewrite [FE] as $P(\mathbf{v}, s) = \sum_{i=1}^{d} \inf_{(u_i, \mathbf{w}_i) \in \Gamma_i(\mathbf{v})} [C(u_i, i) + \alpha P(\mathbf{w}_i, i)]$ from which the lemma immediately follows. □

Lemma C.9. Suppose the environment is Regular. Then:

(a) There exists an optimal contract ξ^* such that, for each $i \in S$, the functions $\xi^{*f}(\cdot, \cdot, i) : D \times S \to \mathcal{U}$ and $\xi^{*c}(\cdot, \cdot, i) : D \times S \to D$ depend on $(\mathbf{v}, s)$ only through the components $(v_i, \ldots, v_d)$.[52]

(b) If the environment is [TVC]-Regular, the unique optimal contract ξ^* satisfies the independence property in part (a).

Proof. Consider first part (a). Existence of an optimal contract is established in Proposition 3.3. The existence of an optimal contract with the desired properties then follows immediately from Lemma C.8. Part (b) then follows from part (a) of the present lemma and part (b) of Proposition 3.3. □

Lemma C.10. Suppose the environment is Regular. For each $i \in S$, the interim value function $Q^i : D \to \mathbb{R}$ satisfies the following properties:

(a) It is convex and continuously differentiable.

(b) For every $\mathbf{v} \in D$ and $i \in S$, there exists some $(u_i, \mathbf{w}_i) \in \Gamma_i(\mathbf{v})$ such that all of the [IC^*_{ji}] $(j > i)$ hold as strict inequalities.

Proof. Part (a) follows from the definition of [FE-Q^i], convexity and continuous differentiability of $P(\cdot, i)$ (Proposition 3.3). Part (b) follows from the observation that $\Gamma(\mathbf{v}) = \Gamma_1(\mathbf{v}) \times \cdots \times \Gamma_d(\mathbf{v})$ and Condition R.3. □

(52) Formally, $\xi^{*f}(\mathbf{v}, s, i) = \xi^{*f}(\mathbf{v}', s', i)$ for all $(\mathbf{v}, s), (\mathbf{v}', s') \in D \times S$ such that $v_j = v'_j$ for all $j \geq i$.

By Lemma C.10, the solutions of problem [**FE-**Q^i] are characterized by saddle points of the Lagrangian (see, e.g., Exercise 7 on p. 236 and Theorem 2 on p. 221 of Luenberger (1969))

$$\mathscr{L}^i(\mathbf{v}, \boldsymbol{\eta}, \boldsymbol{\sigma}, \mathbf{u}, \mathbf{w}) = C(u_i, i) + \alpha P(\mathbf{w}_i, i) + \eta_i \Big[v_i - u_i + \alpha \, \mathsf{E}^{\mathbf{f}_i} [\mathbf{w}_i] \Big]$$

[$\mathbf{L}_i$]
$$- \sum_{j=i+1}^{d} \sigma_{ji} \Big[v_j - v_i - \psi(u_i, j, i) + u_i - \alpha \left(\mathsf{E}^{\mathbf{f}_j} [\mathbf{w}_i] - \mathsf{E}^{\mathbf{f}_j} [\mathbf{w}_j] \right) \Big]$$

and, in particular, by the appropriate envelope, first-order, and complementary slackness conditions. Here, $\eta_i(\mathbf{v}) \in \mathbb{R}$ is the multiplier on [$\mathbf{PK}_i$] and $\sigma_{ji}(\mathbf{v}) \in \mathbb{R}_+$ is the multiplier on [$\mathbf{IC}^*_{ji}$].

Lemma C.11. Suppose the environment is Regular. At the optimum:

(a) For every $(\mathbf{v}, s) \in D \times S$, the multipliers satisfy

[**C.13**]
$$\frac{\lambda_i(\mathbf{v}, s)}{f_{si}} = \eta_i(\mathbf{v}) + \sum_{k=i+1}^{d} \sigma_{ki}(\mathbf{v}) - \sum_{k=1}^{i-1} \frac{f_{sk}}{f_{si}} \sigma_{ik}(\mathbf{v}) \qquad \text{for all } i \in S$$

(b) For every $(\mathbf{v}, s) \in D \times S$, the multipliers satisfy, for all $i \in S$,

[**C.14**]
$$0 = \sum_{k=1}^{i-1} \Big[\mu_{ik}(\mathbf{v}, s) - f_{sk} \sigma_{ik}(\mathbf{v}) \Big] + \sum_{k=i+1}^{d} \psi'(u_i, k, i) \Big[f_{si} \sigma_{ki}(\mathbf{v}) - \mu_{ki}(\mathbf{v}, s) \Big]$$

(c) For every $\mathbf{v} \in D$, the following are equivalent: (i) $\sigma_{ij}(\mathbf{v}) = 0$ for all $i > j$, (ii) for *some* $s \in S$, $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i > j$, (iii) for *all* $s \in S$, $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i > j$.

Proof. The lemma follows from Lemma C.8 and comparison of optimality conditions from the "ex ante" problem in OA-C.1 and those derived from the "interim" Lagrangians [$\mathbf{L}_i$], $i \in S$. Begin with part (a). The envelope conditions from [$\mathbf{L}_i$] read

[**C.15**]
$$Q^i_j(\mathbf{v}) = \mathbf{1}(i = j) \cdot \left[\eta_i(\mathbf{v}) - \sum_{j=i+1}^{d} \sigma_{ji}(\mathbf{v}) \right] - (1 - \mathbf{1}(j > i)) \cdot \sigma_{ji}(\mathbf{v})$$

for all $i, j \in S$. It follows from Lemma C.8 that $P_j(\mathbf{v}, s) = \sum_{i=1}^{d} f_{si} Q^i_j(\mathbf{v})$ and thus, substituting in the interim envelope conditions [**C.15**], that

[**C.16**]
$$P_j(\mathbf{v}, s) = f_{sj} \left[\eta_j(\mathbf{v}) - \sum_{k=j+1}^{d} \sigma_{kj}(\mathbf{v}) \right] - \sum_{k=1}^{j-1} f_{sk} \sigma_{jk}(\mathbf{v})$$

Substituting the ex ante envelope condition [**Env**$_i$] ($i = j$) into [**C.16**] delivers $\lambda_j(\mathbf{v}, s) = f_{sj}\big[\eta_j(\mathbf{v}) - \sum_{k=j+1}^{d} \sigma_{kj}(\mathbf{v})\big] - \sum_{k=1}^{j-1} f_{sk}\sigma_{jk}(\mathbf{v})$. Dividing both sides through by f_{sj} and replacing the dummy index j with i delivers [**C.13**].

Now, consider part (b). The first-order condition with respect to u_i in [**L**$_i$] is

[**C.17**] $$C'(u_i, i) = \eta_i(\mathbf{v}) + \sum_{k=i+1}^{d} \sigma_{ki}(\mathbf{v}) - \sum_{k=i+1}^{d} \psi'(u_i, k, i)\sigma_{ki}(\mathbf{v})$$

and the first-order condition for u_i in the ex ante problem, [**FOC**u_i], is

[**C.18**] $$C'(u_i, i) = \frac{\lambda_i(\mathbf{v}, s)}{f_{si}} + \sum_{k=1}^{i-1} \frac{\mu_{ik}(\mathbf{v}, s)}{f_{si}} - \sum_{k=i+1}^{n} \psi'(u_i, k, i)\frac{\mu_{ki}(\mathbf{v}, s)}{f_{si}}$$

Setting the RHS of [**C.17**] equal to the RHS of [**C.18**] and substituting in [**C.13**] delivers

$$\begin{aligned}
&\frac{\lambda_i(\mathbf{v}, s)}{f_{si}} + \sum_{k=1}^{i-1} \frac{\mu_{ik}(\mathbf{v}, s)}{f_{si}} - \sum_{k=i+1}^{n} \psi'(u_i, k, i)\frac{\mu_{ki}(\mathbf{v}, s)}{f_{si}} \\
&\quad = \sum_{k=i+1}^{d} \sigma_{ki}(\mathbf{v})\big(1 - \psi'(u_i, k, i)\big) + \underbrace{\left[\frac{\lambda_i(\mathbf{v}, s)}{f_{si}} - \sum_{k=i+1}^{d} \sigma_{ki}(\mathbf{v}) + \sum_{k=1}^{i-1} \frac{f_{sk}}{f_{si}}\sigma_{ik}(\mathbf{v})\right]}_{= \eta_i(\mathbf{v})}
\end{aligned}$$

Simplifying the above display yields [**C.14**].

Finally, consider part (c). Let $\mathbf{v} \in D$ be given. We show that (i) implies (iii) by induction. So suppose that (i) holds, and let $s \in S$ be given. For the base step, note that [**C.14**] with $i = d$ becomes $0 = \sum_{k=1}^{d-1} \Big[\mu_{dk}(\mathbf{v}, s) - f_{sk}\sigma_{dk}(\mathbf{v})\Big] = \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{v}, s)$ where the second equality follows because (i) holds at $\mathbf{v}$. Because $\mu_{dk}(\cdot) \geq 0$ on $D \times S$ for all $k < d$, it follows that $\mu_{dk}(\mathbf{v}, s) = 0$ for all $k < d$. For the inductive step, suppose we have shown, for all $i > \ell$, that $\mu_{ij}(\mathbf{v}, s) = 0$ for all $j < i$. Then [**C.14**] with $i = \ell$ reads

$$\begin{aligned}
0 &= \sum_{k=1}^{\ell-1} \Big[\mu_{\ell k}(\mathbf{v}, s) - f_{sk}\sigma_{\ell k}(\mathbf{v})\Big] + \sum_{k=\ell+1}^{d} \psi'(u_\ell, k, \ell)\Big[f_{s\ell}\sigma_{k\ell}(\mathbf{v}) - \mu_{k\ell}(\mathbf{v}, s)\Big] \\
&= \sum_{k=1}^{\ell-1} \mu_{\ell k}(\mathbf{v}, s) - \sum_{k=\ell+1}^{d} \psi'(u_\ell, k, \ell)\mu_{k\ell}(\mathbf{v}, s) = \sum_{k=1}^{\ell-1} \mu_{\ell k}(\mathbf{v}, s)
\end{aligned}$$

where the second equality follows because (i) holds at $\mathbf{v}$ and the third equality follows from the induction hypothesis. As before, it follows that $\mu_{\ell k}(\mathbf{v}, s) = 0$ for all $k < \ell$. Thus, by induction, we see that $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i, j \in S$ with $i > j$. The given $s \in S$ was arbitrary, so this holds for all $s \in S$. Thus, we have shown that (i) implies (iii).

Given Assumption Markov, the proof that (ii) implies (i) is completely analogous. It is obvious that (iii) implies (ii). Thus, we have shown the desired equivalence. □

Lemma C.12. Suppose the environment is [**TVC**]-Regular. Let $\mathbf{v} \in D$ be given.

(a) Suppose that, for *some* $s \in S$, $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i, j \in S$. Then, for *all* $r \in S$ we have

$$\frac{\lambda_i(\mathbf{v}, r)}{f_{ri}} = \eta_i(\mathbf{v}) \qquad \text{for all } i \in S \tag{C.19}$$

(b) Suppose that, for *some* $s \in S$, $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i, j \in S$. Suppose, in addition, that $\mathbf{v} \in E_s$. Then, $\mathbf{v} \in E_r$ for *all* $r \in S$.

Proof. Begin with part (a). The hypothesis together with part (c) of Lemma C.11 implies that $\sigma_{ij}(\mathbf{v}) = 0$ for all $i, j \in S$ with $i > j$. Plugging this into [**C.13**] yields [**C.19**] for every $r, i \in S$.

Now consider part (b). From [$\tilde{\mathbf{E}}_i$] and the envelope conditions [$\mathbf{Env}_i$], part (a) implies that there exists a single number, call it $\hat{\eta}(\mathbf{v})$, such that $\eta_i(\mathbf{v}) = \hat{\eta}(\mathbf{v})$ for all $i \in S$. Applying part (a) again delivers $\frac{\lambda_i(\mathbf{v}, s')}{f_{s'i}} = \hat{\eta}(\mathbf{v})$ for all $s', i \in S$. Using the envelope conditions [$\mathbf{Env}_i$] to replace $\lambda_i(\mathbf{v}, s')$ with $P_i(\mathbf{v}, s')$ delivers that $\mathbf{v} \in E_{s'}$ for all $s' \in S$, as desired. □

C.3.4. *Self-Generation and the First-Best*

Say that a recursive contract ξ *self-generates at* $\mathbf{v} \in V$ if $\xi^c((\mathbf{v}, s), i) = \mathbf{v}$ for all $s, i \in S$. Say that a recursive contract ξ *implements the first-best at* $(\mathbf{v}, s) \in D \times S$ if the induced allocation $\tilde{u}_\xi$ solves the first-best problem [**FB**] given initial condition $(\mathbf{v}, s)$.

Lemma C.13. The first-best contract (see Section 3.1 and OA-E) is not feasible at any $\mathbf{v} \in D$.

Proof. By Lemma E.2 in OA-E, the first-best contract is characterized by perfect consumption smoothing. That is, there exist consumption levels $c_1 > \cdots > c_d$ satisfying $\omega_1 + c_1 = \cdots = \omega_d + c_d$ such that the principal gives the agent c_i units of consumption good whenever the agent reports that his endowment is ω_i. This clearly violates the incentive constraints, as the agent's unique best reply is to always report the lowest endowment, ω_1. □

Lemma C.14. Suppose the environment is [**TVC**]-Regular. Let $(\mathbf{v}, s) \in D \times S$ be given such that $\mathbf{v} \in E_s$. If $\mu_{ij}(\mathbf{v}, s) = 0$ for all $i, j \in S$, then:

(i) The optimal contract self-generates at $\mathbf{v}$;

(ii) The optimal contract implements the first-best at $(\mathbf{v}, s)$.

Proof. We begin with part (i). As usual, define $u_i := \xi^f((\mathbf{v}, s), i)$ and $\mathbf{w}_i := \xi^c((\mathbf{v}, s), i)$. Under the hypothesis of the lemma, the optimality conditions ([**Env**$_i$], [**FOC**u_i], [**FOCw**$_{ij}$], respectively) reduce to

$$P_i(\mathbf{v}, s) = \lambda_i(\mathbf{v}, s) \qquad \forall\, i \in S \tag{C.20}$$

$$f_{si} C'(u_i) = \lambda_i(\mathbf{v}, s) \qquad \forall\, i \in S \tag{C.21}$$

$$f_{si} P_j(\mathbf{w}_i, i) = f_{ij} \lambda_i(\mathbf{v}, s) \qquad \forall\, i, j \in S \tag{C.22}$$

It is easy to see that [**C.22**] implies that $\mathbf{w}_i \in E_i$ for all $i \in S$. Moreover, the hypothesis of the lemma and part (iii) of Lemma C.12 imply that $\mathbf{v} \in E_{s'}$ for all $s' \in S$. Now, plugging [**C.19**] from Lemma C.12 into [**C.20**] and [**C.22**], and invoking part (iii) of Lemma C.12, delivers

$$\frac{P_i(\mathbf{v}, s')}{f_{s'i}} = \hat{\eta}(\mathbf{v}) \qquad \forall\, i, s' \in S \tag{C.23}$$

$$\frac{P_j(\mathbf{w}_i, i)}{f_{ij}} = \hat{\eta}(\mathbf{v}) \qquad \forall\, i, j \in S \tag{C.24}$$

where, as in the proof of Lemma C.12, we are denoting by $\hat{\eta}(\mathbf{v})$ the common value taken by each of the $\{\eta_i(\mathbf{v})\}_{i \in S}$. Recall, from Lemma C.9, that $\mathbf{w}_i$ is independent of s, so [**C.23**] and [**C.24**] do not depend on the given s at all. Thus, take an arbitrary $s' \in S$. Summing over $i \in S$ in [**C.23**] delivers $\mathsf{D}_{\mathbf{1}} P(\mathbf{v}, s') = \sum_{i=1}^{d} f_{s'i} \hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$. Similarly, set $i = s'$ in [**C.24**] and sum over $j \in S$ to get $\mathsf{D}_{\mathbf{1}} P(\mathbf{w}_{s'}, s') = \sum_{i=j}^{d} f_{s'j} \hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$. Now, we have established that $\mathbf{v}, \mathbf{w}_{s'} \in E_{s'}$. By Lemma C.5, the above display implies that $\mathbf{w}_{s'} = \mathbf{v}$. But $s' \in S$ was arbitrary, so we have $\mathbf{w}_i = \mathbf{v}$ for all $i \in S$, which establishes part (i) of the lemma.

Now consider part (ii). Combining [**C.20**], [**C.21**], and [**C.23**], we see that there exists some $\hat{u}(\mathbf{v}, s) \in \mathbb{R}_{--}$ such that $u_i = \hat{u}(\mathbf{v}, s)$ for all $i \in S$. Because the policy functions are independent of s (part (b) of Lemma C.9) and, by part (i) of the present lemma, the optimal contract ξ^* self-generates at $\mathbf{v}$, it follows that the induced allocation $\tilde{u}_{\xi^*}$ is constant and equal to $\hat{u}$ when initialized at $(\mathbf{v}, s)$. It follows from Lemma E.2 (see OA-E) that ξ^* implements the first-best at $(\mathbf{v}, s)$. □

Lemma C.15. Suppose the environment is [**TVC**]-Regular. Let $(\mathbf{v}, s) \in D \times S$ be given, and define $\mathbf{w}_d := \xi^c((\mathbf{v}, s), d)$. For all $i \in S$, define $\tilde{\mathbf{w}}_i := \xi^c((\mathbf{w}_d, d), i)$. There exists some $i \in S$, call it $i^*(\mathbf{w}_d)$, such that $\mathsf{D}_{\mathbf{1}} P(\tilde{\mathbf{w}}_{i^*(\mathbf{w}_d)}, i^*(\mathbf{w}_d)) \neq \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_d, d)$.

Proof. Suppose not. Then Lemmas C.14 and C.7 imply that the optimal contract implements the first best at $(\mathbf{w}_d, d)$. But this is impossible by C.13, a contradiction. □

C.4. Main Proof of Theorem 1

The following pieces of notation will be used extensively during the proof. Recall that we denote the space of *(infinite) histories*, or *paths*, by $\mathscr{H} := S^\infty$ with generic element $h := (s^t)_{t=0}^\infty$, where s^t denotes the realized type in period $t-1$. (Recall that s denotes the *previous period's* realized type.) Let $\tau^{(t)}$ denote the *random* time defined pathwise by $\tau^{(t)}(h) := \sup\left\{T \leq t : \; s^T = d\right\}$. That is, given path h, $\tau^{(t)}(h)$ is the last date (i) that precedes t and (ii) that was immediately preceded by a realized endowment ω_d. It is easy to see that $\tau^{(t)}$ is well-defined *stopping time*, and that the stochastic process $\left(\tau^{(t)}\right)_{t=0}^\infty$ is $\mathbf{P}$-a.s. non-decreasing.

Martingale Convergence. Define the event

$$\mathscr{I} := \left\{h \in \mathscr{H} : \forall i \in S, \, (s^t, s^{t+1}) = (d, i) \text{ occurs for infinitely-many } t \right\}$$

It is easy to see that $\lim_{t\to\infty} \tau^{(t)}(h) = +\infty$ for all $h \in \mathscr{I}$. We note here that $\mathbf{P}(\mathscr{I}) = 1$ by Corollary F.3 in OA-F.

Lemma C.16. Suppose the environment is TVC-Regular. Under the optimal contract, the marginal cost martingale satisfies $\mathsf{D}_\mathbf{1}(\mathbf{v}^{(t)}, s^{(t)}) \to 0$ almost surely.

Proof of Lemma C.16. Suppose that the environment is [TVC]-Regular, as hypothesized by the theorem. Proposition 4.3 shows that, under the optimal contract, the process $\left(\mathsf{D}_\mathbf{1}P(\mathbf{v}^{(t)}, s^{(t)})\right)_{t=0}^\infty$ defines a strictly positive martingale. By Doob's Martingale Convergence Theorem (see Theorem 2 in Shiryaev (1995, p. 517)), it must converge $\mathbf{P}$-a.s. to a non-negative, $\mathbf{P}$-integrable random variable. For purposes of establishing almost sure convergence, we may restrict attention to the event $\mathscr{I} \subseteq \mathscr{H}$. So fix an arbitrary path $h := (s^t)_{t=0}^\infty \in \mathscr{I}$. Since the path is fixed, let $\tau^t := \tau^{(t)}(h)$ and $\mathbf{v}^t := \mathbf{v}^{(t)}(h)$ for all $t \in \mathbb{N}$. Similarly, for each $i \in S$ define $\tau_i^t := \sup\left\{T \leq t : (s^T, s^{T+1}) = (d, i)\right\}$. It follows from the definition of the event $\mathscr{I}$ that $\lim_{t\to\infty} \tau_i^t = +\infty$ for all $i \in S$.

Suppose, towards contradiction, that $\mathsf{D}_\mathbf{1}P(\mathbf{v}^t, s^t) \to C > 0$. It then follows from Lemmas C.4 and C.5 that $\mathbf{v}^{\tau_t} \to \mathbf{w}_d^* \in E_d$. The policy functions are continuous under Condition R.4, so $\xi^c(\mathbf{v}^{\tau_t}, d, i) \to \xi^c(\mathbf{w}_d^*, d, i) =: \tilde{\mathbf{w}}_i^*$. Fix $i \in S$. Because $P(\cdot, i)$ is continuously differentiable by Proposition 3.3, it follows that $\lim_{t\to\infty} \mathsf{D}_\mathbf{1}P(\mathbf{v}^{\tau_t^i+1}, s^{\tau_t^i+1}) = \mathsf{D}_\mathbf{1}P(\tilde{\mathbf{w}}_i^*, i)$, as the sequence $\{s^{\tau_t^i+1}\}$ is constant and equal to i, by definition of the sequence $\{\tau_i^t\}$. This holds for all $i \in S$, so by the supposition it follows that $\mathsf{D}_\mathbf{1}P(\mathbf{w}_d^*, d) = \mathsf{D}_\mathbf{1}P(\tilde{\mathbf{w}}_i^*, i)$ for all $i \in S$. But this is impossible, as it implies that the optimal contract self-generates and implements the first-best at $\mathbf{w}_d^*$ by Lemmas C.13, C.14, and C.15. Thus, $\mathsf{D}_\mathbf{1}P(\mathbf{v}^t, s^t) \to 0$, as desired. $\square$

Convergence of Multipliers. We next characterize the limit properties of the Lagrange multipliers. This is done in a series of lemmas presented below. Let $\boldsymbol{\nu}(\mathbf{v}, s) \in \mathbb{R}^{d(d+1)2}$ denote the vector of Lagrange multipliers induced by the optimal contract at state $(\mathbf{v}, s) \in D \times S$, obtained by stacking the d multipliers $\lambda_i(\mathbf{v}, s) \in \mathbb{R}$ on the promise keeping constraints [**PK**$_i$] and the $d(d-1)/2$ multipliers $\mu_{ij}(\mathbf{v}, s) \in \mathbb{R}_+$ on the incentive constraints [**IC**$_{ij}$] ($j > i$). The optimal contract induces a process $\left(\boldsymbol{\nu}^{(t)}\right)_{t=0}^{\infty}$, where $\boldsymbol{\nu}^{(t)} := \boldsymbol{\nu}(\mathbf{v}^{(t)}, s^{(t)})$ for each $t \in \mathbb{N}$. Similarly define the processes $\left(\boldsymbol{\lambda}^{(t)}\right)_{t=0}^{\infty}$ and, for each $i \in S$, $\left(\boldsymbol{\mu}_{*,i}^{(t)}\right)_{t=0}^{\infty}$, where $\boldsymbol{\lambda}(\mathbf{v}, s) := \left(\lambda_1(\mathbf{v}, s), \ldots, \lambda_d(\mathbf{v}, s)\right) \in \mathbb{R}^d$ and $\boldsymbol{\mu}_{*,i}(\mathbf{v}, s) := (\mu_{i+1,i}(\mathbf{v}, s), \ldots, \mu_{d,i}(\mathbf{v}, s)) \in \mathbb{R}_+^{d-i}$. Finally, let $\boldsymbol{\mu}(\mathbf{v}, s) \in \mathbb{R}^{d(d-1)/2}$ denote the vector that stacks each of the $\boldsymbol{\mu}_{*,i}(\mathbf{v}, s)$.

Lemma C.17. Suppose the environment is [**TVC**]-Regular. Under the optimal contract, $\boldsymbol{\nu}^{(\tau^{(t)})} \to \mathbf{0}$ and $\boldsymbol{\lambda}^{(\tau^{(t)}+1)} \to \mathbf{0}$ almost surely.

Proof of Lemma C.17. Fix some path $h \in \mathcal{J}$ along which the marginal cost martingale converges to zero. By part (a) of Theorem 1, the set of such paths has full measure, and is thus sufficient for establishing almost sure convergence. By [$\tilde{\mathbf{E}}_i$] and Lemmas C.4 and C.5, convergence of the marginal cost martingale implies that the derivative process converges along the subsequence $\{\tau_t\}_{t=0}^{\infty}$, and part (a) of Theorem 1 (proved above) requires that $\mathsf{D}P(\mathbf{v}^{\tau_t}, s^{\tau_t}) \to \mathbf{0}$. By the envelope conditions [**Env**$_i$], this translates to

$$\text{[C.25]} \qquad \boldsymbol{\lambda}^{\tau_t} \to \mathbf{0}$$

It remains to show that $\boldsymbol{\mu}_{*,i}^{\tau_t} \to 0$ for all $i \in S$. To do so, we will use the optimality conditions and induct through the type space, starting from the bottom. Define $\tilde{\mathbf{w}}_i^{\tau_t} := \xi^c(\mathbf{v}^{\tau_t}, d, i)$.

Base step: The first-order condition [**FOCw**$_{ij}$] with $i = 1$ at state $(\mathbf{v}^{\tau_t}, d)$ is $f_{d1} P_j(\tilde{\mathbf{w}}_1^{\tau_t}, 1) = f_{1j}\left(\lambda_1(\mathbf{v}^{\tau_t}, d) + 0\right) - \sum_{k=2}^{d} f_{kj}\mu_{k1}(\mathbf{v}^{\tau_t}, d)$. Because $\boldsymbol{\mu}(\cdot) \geq \mathbf{0}$ on $D \times S$, it follows from [**C.25**] that $f_{d1} P_j(\tilde{\mathbf{w}}_1^{\tau_t}, 1) \to 0$ and thus also that $\sum_{k=2}^{d} f_{kj}\mu_{k1}(\mathbf{v}^{\tau_t}, d) \to 0$. This is true for all $j \in S$. Because the type process is fully connected (Assumption Markov), it follows that

$$\text{[C.26]} \qquad \mathsf{D}P(\tilde{\mathbf{w}}_1^{\tau_t}, 1) \to \mathbf{0}$$

$$\text{[C.27]} \qquad \boldsymbol{\mu}_{*,1}(\mathbf{v}^{\tau_t}, d) \to \mathbf{0}$$

Inductive step: Let $m \in S$. Suppose we have shown that $\mathsf{D}P(\tilde{\mathbf{w}}_\ell^{\tau_t}, \ell) \to \mathbf{0}$ and $\boldsymbol{\mu}_{*,\ell}(\mathbf{v}^{\tau_t}, d) \to \mathbf{0}$ for all $\ell < m$. The first order condition [**FOCw**$_{ij}$] with $i = m$ at state $(\mathbf{v}^{\tau_t}, d)$ is $f_{dm} P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) = f_{mj}\left(\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d)\right) - \sum_{k=m+1}^{d} f_{kj}\mu_{km}(\mathbf{v}^{\tau_t}, d)$.

By [C.25] and the supposition, it follows that $\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d) \to \mathbf{0}$. As $\boldsymbol{\mu}(\cdot) \geq \mathbf{0}$ on $D \times S$, it follows from the previous two displays that $f_{dm} P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) \to 0$ and $\sum_{k=m+1}^{d} f_{kj}\mu_{km}(\mathbf{v}^{\tau_t}, d) \to 0$. This is true for all $j \in S$. Because the type process is fully connected (Assumption Markov), it follows that

[C.28] $$\mathsf{D}P(\tilde{\mathbf{w}}_m^{\tau_t}, m) \to \mathbf{0}$$

[C.29] $$\boldsymbol{\mu}_{*,m}(\mathbf{v}^{\tau_t}, d) \to \mathbf{0}$$

Thus, by induction, we have shown (i) that $\mathsf{D}P(\tilde{\mathbf{w}}_i^{\tau_t}, i) \to \mathbf{0}$ for all $i \in S$ and (ii) that $\boldsymbol{\mu}^{\tau_t} \to \mathbf{0}$. It follows from (i) and the envelope conditions [$\mathbf{Env}_i$] that $\boldsymbol{\lambda}^{\tau_t+1} \to \mathbf{0}$. It follows from (ii) and [C.25] that $\boldsymbol{\nu}^{\tau_t} \to \mathbf{0}$. □

Lemma C.18. Suppose the environment is [TVC]-Regular. Under the optimal contract and for all $k \in \mathbb{N}$, $\boldsymbol{\nu}^{(\tau^{(t)}+k)} \to \mathbf{0}$ and $\boldsymbol{\lambda}^{(\tau^{(t)}+k+1)} \to \mathbf{0}$ almost surely.

Proof of Lemma C.18. Suppose we have shown the desired convergence for all $k = 0, \ldots, m-1$. Replicating the proof of Lemma C.17 with "τ_t" replaced everywhere by "$\tau_t + m - 1$" shows that we obtain the desired convergence for $k = m$, as well. The lemma then follows from induction, with Lemma C.17 serving as the base step. □

Lemma C.19. Suppose the environment is [TVC]-Regular. Under the optimal contract, $\boldsymbol{\nu}^{(t)} \to \mathbf{0}$ in probability.

Proof of Lemma C.19. Define the stochastic processes $\left(\delta^{(t)}\right)_{t=0}^{\infty}$ and $\left(L^{(t)}\right)_{t=0}^{\infty}$ by $\delta^{(t)} := ||\boldsymbol{\nu}^{(t)}||$, where $||\cdot||$ denotes the Euclidean norm on $\mathbb{R}^{d(d+1)/2}$, and $L^{(t)} := t - \tau^{(t)}$. Thus, $\delta^{(t)}(h) \geq 0$ denotes the *d*istance of $\boldsymbol{\nu}^{(t)}(h)$ from the zero vector and $L^{(t)}(h) \in \mathbb{N}$ denotes the *l*ag since the last ω_d realization at date t, along path $h \in \mathscr{H}$.

To show convergence in probability, we must show that $\limsup_{t\to\infty} \mathbf{P}(\delta^{(t)} > \varepsilon) = 0$ for all $\varepsilon > 0$. To that end, let $\varepsilon > 0$ and $k \in \mathbb{N}$ be given. For each $t \in \mathbb{N}$, define the following events: $A_{\varepsilon,t} := \left\{h \in \mathscr{H} : \delta^{(t)}(h) > \varepsilon\right\}$, $B_{k,t} := \left\{h \in \mathscr{H} : L^{(t)}(h) > k\right\}$, and $C_{\varepsilon,k,t} := \cup_{T \geq t}\left[A_{\varepsilon,T} \cap B_{k,T}^c\right]$. Note that $C_{\varepsilon,k,t+1} \subseteq C_{\varepsilon,k,t}$ and $A_{\varepsilon,t} \cap B_{k,t}^c \subseteq C_{\varepsilon,k,t}$ for each $t \in \mathbb{N}$. We have $\mathbf{P}\left(A_{\varepsilon,t}\right) = \mathbf{P}\left(A_{\varepsilon,t} \cap B_{k,t}^c\right) + \mathbf{P}\left(A_{\varepsilon,t} \cap B_{k,t}\right) \leq \mathbf{P}\left(C_{\varepsilon,k,t}\right) + \mathbf{P}\left(B_{k,t}\right)$, where the inequality follows from monotonicity of probability. Observe that $\lim_{t\to 0} \mathbf{P}\left(C_{\varepsilon,k,t} \cap \mathscr{I}\right) = 0$, as the sequence of sets $\{C_{\varepsilon,k,t}\}_{t=0}^{\infty}$ is non-increasing (in the set inclusion order) and because Lemmas C.17 and C.18 imply that $\delta^{(\tau^{(t)}+m)} \to 0$ almost surely for all $m = 0, \ldots, k$. Thus, it follows that

[C.30] $$\limsup_{t\to\infty} \mathbf{P}\left(A_{\varepsilon,t}\right) \leq \limsup_{t\to\infty} \mathbf{P}\left(B_{k,t}\right)$$

The remaining step of the proof is to show that there exists a function $H : \mathbb{N} \to [0,1]$ that satisfies $\lim_{k\to\infty} H(k) = 1$, and such that $\lim_{t\to\infty} \mathbf{P}(B_{k,t}) = 1 - H(k)$ for all k. If such H exists, then since [C.30] is valid for all $k \in \mathbb{N}$ and only the RHS depends on k, it must be that $\limsup_{t\to\infty} \mathbf{P}(A_{\varepsilon,t}) \le \inf_{k\in\mathbb{N}} (1 - H(k)) = 0$ and, since $\varepsilon > 0$ was arbitrary, this establishes the desired convergence in probability.

We now show that such H exists. Under Assumption Markov, the type process $\left(s^{(t)}\right)$ is ergodic. Thus, there exists a unique stationary distribution $\boldsymbol{\pi} \in \Delta(S)$, $\pi_i > 0$ for all $i \in S$, and $\lim_{t\to\infty} \mathbf{P}(s^{(t)} = i) = \pi_i$ for all $i \in S$. Define the Markov process $\left(r^{(t)}\right)$ via the transition probabilities $\mathbf{Q}\left(r^{(t+1)} = j \mid r^{(t)} = i\right) = g_{ij} := \frac{\pi_j}{\pi_i} \cdot f_{ji}$ and let $\mathbf{Q} \in \Delta(S^\infty)$ denote the induced measure over paths. The Markov process $\left(r^{(t)}\right)$ is the *time-reversed* version of $\left(s^{(t)}\right)$. (Note that the backward transition probabilities g_{ij} are defined from the forward transition probabilities f_{ij} via Bayes' Rule, with the stationary distribution of the forward chain, $\boldsymbol{\pi}$, acting as the prior.)

Let T_d^R denote the hitting time of state d for the time-reversed chain, i.e., it is the $\mathbb{N} \cup \{+\infty\}$-valued random variable defined by $T_d^R := \inf\left\{t \in \mathbb{N} : r^{(t)} = d\right\}$. For each $i \in S$, define the function $H_i : \mathbb{N} \to [0,1]$ by $H_i(k) := \mathbf{Q}\left(T_d^R \le k \mid r^{(0)} = i\right)$. Thus, $H_i(\cdot)$ is the CDF of T_d^R given that the time-reversed chain starts in state $i \in S$.

Claim 7. Let $\mathbf{P}$ and $\mathbf{Q}$ be as defined above. Then, for every $i \in S$:

(a) $\lim_{k\to\infty} H_i(k) = 1$.

(b) For all $k \in \mathbb{N}$, $\lim_{t\to\infty} \mathbf{P}\left(s^{(t-m)} \neq d \; \forall\, m = 0, \ldots, k \;\middle|\; s^{(t)} = i\right) = 1 - H_i(k)$

Proof of Claim 7. It is clear that the time-reversed process is fully connected, and thus each state is recurrent, i.e., $\mathbf{Q}\left(T_d^R < \infty \mid r^{(0)} = i\right) = 1$ for all $i \in S$. It follows from the Bounded Convergence Theorem that $\lim_{k\to\infty} H_i(k) = 1$ for all $i \in S$. This establishes part (a).

For part (b), let $k \in \mathbb{N}$ be given and consider only $t > k$ large enough that $\mathbf{P}\left(s^{(T)} = i\right) > 0$ for all $T \ge t$. (This is possible because $\pi_i > 0$.) Then,

$$\mathbf{P}\left(s^{(t-m)} \neq d \; \forall\, m = 0, \ldots, k \;\middle|\; s^{(t)} = i\right) = \frac{\mathbf{P}\left(s^{(t-m)} \neq d \; \forall\, m = 0, \ldots, k \text{ and } s^{(t)} = i\right)}{\mathbf{P}\left(s^{(t)} = i\right)}$$

If $i = d$, we are done, so suppose $i \neq d$. Then we have

$$\begin{aligned}
&\mathbf{P}\left(s^{(t-m)} \neq d \;\forall\, m = 0, \dots, k \text{ and } s^{(t)} = i\right) \\
&\quad = \sum_{(j_k,\dots,j_1)\in\{1,\dots,d-1\}^k} \mathbf{P}\left(s^{(t-k)} = j_k\right) \cdot f_{j_k,j_{k-1}} \cdots f_{j_2,j_1} \cdot f_{j_1,i} \\
&\quad = \sum_{(j_k,\dots,j_1)\in\{1,\dots,d-1\}^k} \pi_i \cdot \left(g_{i,j_1} \cdots g_{j_{k-1},j_k}\right) \cdot \frac{\mathbf{P}\left(s^{(t-k)} = j_k\right)}{\pi_{j_k}}
\end{aligned}$$

where the first line follows from the Markov property (for the forward chain) and the second line follows from the definition of the time-reversed transition probabilities. Combining the two displays above and noting that $\lim_{t\to\infty} \mathbf{P}\left(s^{(t-k)} = j_k\right) = \pi_{j_k}$ and $\lim_{t\to\infty} \mathbf{P}(s^{(t)} = i) = \pi_i$ completes the proof of the claim. □

To conclude the proof of the lemma, we claim that the function $H(k) := \sum_{i=1}^d \pi_i H_i(k)$ satisfies the desired properties. It clearly satisfies $\lim_{k\to\infty} H(k) = 1$ by part (a) of Claim 7. Notice that we may write $\mathbf{P}(B_{k,t}) = \sum_{i=1}^d \mathbf{P}\left(s^{(t)} = i\right) \cdot \mathbf{P}\left(s^{(t-m)} \neq d \;\forall\, m = 0, \dots, k \;\middle|\; s^{(t)} = i\right)$. By part (b) of Claim 7 and ergodicity of the type process, it follows that $\lim_{t\to\infty} \mathbf{P}(B_{k,t}) = 1 - H(k)$, as desired. □

Lemma C.20. Suppose the type process is pseudo-renewal, with $\boldsymbol{\pi} \in \Delta(S)$ such that $f_{ij} = \pi_j$ for all $i \neq j$. Then it satisfies $f_{ii} - \pi_i = f_{jj} - \pi_j$ for all $i, j \in S$.

Proof of Lemma C.20. Let $i, j \in S$ be given. By assumption, we have $1 = f_{ii} + \pi_j + \sum_{k\neq i,j} \pi_k$ and $1 = f_{jj} + \pi_i + \sum_{k\neq i,j} \pi_k$. Combining these two equations yields $f_{ii} - \pi_i = f_{jj} - \pi_j$, as desired. □

Lemma C.21. Suppose the environment is [**TVC**]-Regular. Suppose that the type process is pseudo-renewal. Under the optimal contract, $\boldsymbol{\nu}^{(t)} \to \mathbf{0}$ almost surely.

Proof of Lemma C.21. Fix some path $h = (s^t)_{t=0}^\infty \in \mathcal{I}$ along which the marginal cost martingale converges to zero. By part (a) of Theorem 1, the set of such paths has full measure, and is thus sufficient for establishing almost sure convergence. Recall that the first-order condition [**FOC**$\mathbf{w}_{ij}$] at state $(\mathbf{v}, s)$ reads

$$f_{si} P_j(\mathbf{w}_i(\mathbf{v}, s), i) = f_{ij}\left(\lambda_i(\mathbf{v}, s) + \sum_{k=1}^{i-1} \mu_{ik}(\mathbf{v}, s)\right) - \sum_{k=i+1}^{n} f_{kj}\mu_{ki}(\mathbf{v}, s) \tag{C.31}$$

and that, at the optimum, the directional derivative $\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)$ satisfies

[**C.32**] $$f_{si}\cdot \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i) = \lambda_i(\mathbf{v},s) + \sum_{k=1}^{i-1}\mu_{ik}(\mathbf{v},s) - \sum_{k=i+1}^{d}\mu_{ki}(\mathbf{v},s)$$

(This is [**C.2**], from the proof of Proposition 4.3 in OA-C.2.) Substituting [**C.32**] into [**C.31**] and rearranging delivers

[**C.33**] $$f_{si}\left(\frac{P_j(\mathbf{w}_i(\mathbf{v},s),i)}{f_{ij}} - \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)\right) = \sum_{k=i+1}^{d}\left(1-\frac{f_{kj}}{f_{ij}}\right)\mu_{ki}(\mathbf{v},s)$$

Now, since the type process is pseudo-renewal, there exists some $\boldsymbol{\pi}\in\Delta(S)$ such that $f_{ij}=\pi_j$ whenever $i\neq j$. It is then easy to see that [**C.33**] reduces to

[**C.34**] $$f_{si}\left(\frac{P_j(\mathbf{w}_i(\mathbf{v},s),i)}{f_{ij}} - \mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)\right) = \begin{cases} 0, & \text{for } j\leq i \\ \left(1-\frac{f_{jj}}{\pi_j}\right)\mu_{ji}(\mathbf{v},s), & \text{for } j>i \end{cases}$$

By Lemma C.20, there are two cases to consider.

Case 1: First, suppose that $f_{ii}\geq\pi_i$ for all $i\in S$. Because $\boldsymbol{\mu}(\cdot)\geq 0$ on $D\times S$ and since the type process is fully connected (Assumption Markov), it follows that

[**C.35**] $$P_j(\mathbf{w}_i(\mathbf{v},s),i)\leq f_{ij}\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)$$

for all $i,j\in S$ (with equality when $j\leq i$). Now, part (a) of Theorem 1 states that $\mathsf{D}_{\mathbf{1}}P(\mathbf{v}^{(t)},s^{(t)})\to 0$ almost surely. It follows from this, the martingale property (Proposition 4.3), and non-negativity of the directional derivative $\mathsf{D}_{\mathbf{1}}P(\cdot,\cdot)$ on $D\times S$ (Lemma C.1) that $\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v}^{(t)},s^{(t)}),i)\to 0$ almost surely. Thus, [**C.35**] implies that

$$\mathsf{P}\left(\limsup_{t\to\infty}P_j(\mathbf{w}_i(\mathbf{v}^{(t)},s^{(t)}),i)\leq 0\;\forall\, i,j\in S\right)=1$$

Since $\sum_{j=1}^{d}P_j(\mathbf{w}_i(\mathbf{v},s),i)=\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)\geq 0$ this implies that $\mathsf{D}P(\mathbf{w}_i(\mathbf{v}^{(t)},s^{(t)}),i)\to\mathbf{0}$ $\mathbf{P}$-a.s. for all $i\in S$.

Case 2: Second, suppose that $f_{ii}\leq\pi_i$ for all $i\in S$. The appropriate analogue of [**C.35**] is

[**C.36**] $$P_j(\mathbf{w}_i(\mathbf{v},s),i)\geq f_{ij}\mathsf{D}_{\mathbf{1}}P(\mathbf{w}_i(\mathbf{v},s),i)$$

which, by the same argument, implies that $\mathbf{P}\left(\liminf_{t\to\infty} P_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \geq 0 \; \forall \, i, j \in S\right) = 1$. Because $\sum_{j=1}^{d} P_j(\mathbf{w}_i(\mathbf{v}, s), i) = \mathsf{D}_{\mathbf{1}} P(\mathbf{w}_i(\mathbf{v}, s), i) \to 0$, this implies that $\mathsf{D}P(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \to \mathbf{0}$ $\mathbf{P}$-a.s. for all $i \in S$.

We may now complete the proof. In either case, note that because the process $(\mathbf{v}^{(t)})_{t=0}^{\infty}$ is itself generated by the policy functions at the optimum — so that along each path $h \in \mathscr{H}$, $\mathbf{v}^{(t)}(h) = \mathbf{w}_i(\mathbf{v}^{(t-1)}(h), s^{(t-1)}(h))$ for some $i \in S$ — it must be that $\mathsf{D}P(\mathbf{v}^{(t)}, s^{(t)}) \to \mathbf{0}$ almost surely. Replicating the proof of Lemma C.17 with "τ_t" replaced everywhere by "t" establishes that $\boldsymbol{\nu}^{(t)} \to \mathbf{0}$ almost surely, as desired. □

Convergence of Allocations. We next translate convergence properties of the Lagrange multipliers into convergence properties of the allocation itself.

Lemma C.22. Let $\left(n^{(t)}\right)_{t=0}^{\infty}$ be a non-decreasing sequence of random times. If $\boldsymbol{\nu}^{(n^{(t)})} \to \mathbf{0}$ almost surely (in probability), then $u_i^{(n^{(t)})} \to -\infty$ almost surely (in probability) for all $i \in S$.

Proof of Lemma C.22. Let $i \in S$ be given. The first-order condition [**FOC**u_i] at state $(\mathbf{v}, s)$ is

$$f_{si} C'\left(u_i(\mathbf{v}, s), i\right) = \underbrace{\lambda_i(\mathbf{v}, s) + \sum_{k=1}^{i-1} \mu_{ij}(\mathbf{v}, s)}_{A(\mathbf{v}, s)} - \underbrace{\sum_{k=i+1}^{d} \psi'\left(u_i(\mathbf{v}, s), k, i\right) \mu_{ki}(\mathbf{v}, s)}_{B(\mathbf{v}, s)}$$

Now, $B(\cdot, \cdot) \geq 0$ on $D \times S$ and the hypothesis of the lemma implies that $A\left(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}\right) \to 0$ almost surely (in probability), so it follows that the process $f_{s^{(n^{(t)})}, i} C'\left(u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}, i\right) \to 0$ almost surely (in probability). Because the type process is fully connected (Assumption Markov), it follows that $C'\left(u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}), i\right) \to 0$ almost surely (in probability). Finally, $C'(\cdot, i) : \mathscr{U} \to \mathbb{R}_{++}$ is a homeomorphism by Assumption DARA. Thus, an application of the Continuous Mapping Theorem for almost sure convergence (convergence in probability) to the inverse of $C'(\cdot, i)$ implies that $u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}) \to -\infty$ almost surely (in probability). □

Lemma C.23. Consider any feasible recursive contract such that the induced process $u_d(\mathbf{v}^{(t)}, s^{(t)}) \to -\infty$ almost surely. Then the induced process $v_i^{(t)} \to -\infty$ almost surely for all $i \in S$.

Proof of Lemma C.23. Denote the recursive contract by ξ, and its policy functions by $\xi^f(\mathbf{v}, s, i) := u_i(\mathbf{v}, s)$ and $\xi^c(\mathbf{v}, s, i) = \mathbf{w}_i(\mathbf{v}, s)$. Let $i \in S$ be fixed. Iterating the promise keeping constraints [**PK**$_i$] one step ahead delivers

$$v_i^{(t)} = u_i(\mathbf{v}^{(t)}, s^{(t)}) + \alpha \sum_{j=1}^{d} f_{ij} \left(u_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) + \alpha\, \mathsf{E}^{\mathbf{f}_j} \left[\mathbf{w}_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \right] \right)$$
$$\leq \alpha f_{id} \cdot u_d(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i)$$

where the second line follows from Assumption DARA. It follows from the hypothesis of the lemma that $u_d(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \to -\infty$ almost surely, which completes the proof. □

Wrapping Up. We can now consolidate the preceding lemmas into the proof of Theorem 1.

Proof of Theorem 1. Part (a) follows from Lemma C.17, Lemma C.22 with the process $(n^{(t)})_{t=0}^{\infty}$ defined by $n^{(t)} := \tau^{(t)}$, and Lemma C.23. The portion of part (b) concerning convergence in probability follows from Lemmas C.19 and C.22. The portion of part (b) concerning almost sure convergence follows from Lemmas C.21 and C.22. Each portion of part (c) follows from the corresponding portion of part (b) and the appropriate version of Continuous Mapping Theorem. □

D. Proof of Theorem 2

We begin by proving part (d) of Theorem 3 (from Appendix B.1), which is then invoked in the main proof of Theorem 2.

D.1. Proof of part (d) of Theorem 3

Recall the set D^* defined in [**B.1**] and suppose that the type process satisfies FOSD. Consider the IC constraint [**IC**$^*_{ij}$]. Let $F^*(i \mid j) := \sum_{k \geq i} f_{kj}$ for each j. By FOSD, it follows that $F^*(\cdot \mid j) \leq F^*(\cdot \mid k)$ whenever $j < k$. It is easy to see that we can rewrite [**IC**$^*_{ij}$] (for $i > j$) as $v_i - v_j \geq \big[U(\omega_i + c_j) - U(\omega_j + c_j)\big] + \alpha \sum_{k=2}^{d} \big[F^*(k \mid i) - F^*(k \mid j)\big](w_{j,k} - w_{j,k-1})$. Notice that $U(\omega_i + c_j) > U(\omega_j + c_j)$ because $\omega_i > \omega_j$. More generally, we have, $w_{j,k}^{(n)} - w_{j,k-1}^{(n)} \geq \big[U(\omega_k + c_{k-1}) - U(\omega_{k-1} + c_{k-1})\big] + \alpha \sum_{\ell=2}^{d} \big[F^*(\ell \mid k) - F^*(\ell \mid k - 1)\big](w_{k-1,\ell}^{(n+1)} - w_{k-1,\ell}^{(n+1)})$. Iterating the inequalities above, for any sequence of types, we have $v_i - v_j \geq (\text{strictly positive terms}) + \alpha^n \chi_n \left(w_{\ell,k}^{(n)} - w_{\ell,k-1}^{(n)} \right)$ where ℓ is the type in period $n-1$ and χ_n is the product of terms that lie in $[0, 1]$. By assumption, $\mathbf{v} \in D^*$ if, and only if, there exists a recursive contract that satisfies [**TVC**]. Therefore, if $\mathbf{v} \in D^*$ and we are considering a [**TVC**]-implementable contract, then we must have $\lim_{n \to \infty} \alpha^n \chi_n \left(w_{j,k}^{(n)} - w_{j,k-1}^{(n)} \right) \to 0$. But this implies $v_i > v_j$ whenever $i > j$, i.e., $D^* \subset V$.

D.2. Main Proof of Theorem 2

Proof of part (a). Recall that [$\mathbf{IC}^*_{ij}$] (with $i = j + 1$) can be written as

$$v_{j+1} - v_j \geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) + \alpha \left(\mathsf{E}^{\mathbf{f}_{j+1}}[\mathbf{w}_j] - \mathsf{E}^{\mathbf{f}_j}[\mathbf{w}_j]\right)$$

[**D.1**] $$\geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j))$$

where the second inequality follows from the FOSD assumption and part (d) of Theorem 3. It is thus sufficient to show that the difference of flow utilities in [**D.1**] grows without bound.

By part (b) of Theorem 1, $u_j^{(t)} \to -\infty$ in probability, and by definition, $U\big(\omega_j + C(u_j^{(t)}, j)\big) = u_j^{(t)}$. By the Continuous Mapping Theorem, this implies that $\omega_j + c_j^{(t)} \to \underline{c}$ in probability, and furthermore, that $\lim_{t\to\infty}\left[\omega_{j+1} + c_j^{(t)}\right] = \underline{c} + (\omega_{j+1} - \omega_j)$ in probability. There are two cases to consider, depending on whether or not the consumption domain is bounded below.

Case 1: Suppose first that $\underline{c} > -\infty$. Then, by the Continuous Mapping Theorem, $U\big(\omega_{j+1} + C(u_j^{(t)}, j)\big) - U\big(\omega_j + C(u_j^{(t)}, j)\big) \to \infty$ in probability because $U\big(\underline{c} + (\omega_{j+1} - \omega_j)\big) > -\infty$.

Case 2: Suppose now that $\underline{c} = -\infty$. Observe that $U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) = \int_{\omega_j + C(u_j,j)}^{\omega_{j+1}+C(u_j,j)} U'(y)\,\mathrm{d}y \geq (\omega_{j+1} - \omega_j)U'(\omega_{j+1} + C(u_j, j))$ where the inequality follows from the concavity of U. But because $\underline{c} = -\infty$, we have $\omega_{j+1} + c_j^{(t)} \to -\infty$ in probability. It now follows from continuous differentiability of $U(\cdot)$ and the Inada conditions in part (a) of Assumption DARA and the Continuous Mapping Theorem that $U'\big(\omega_{j+1} + C(u_j^{(t)}, j)\big) \to +\infty$ in probability.

Together, these two cases prove part (a). The claim concerning conditional variances then follows as a trivial consequence of the first piece and the Markov property of the process $(\mathbf{v}^{(t)}, s^{(t)})_{t=0}^{\infty}$. □

Proof of part (b). The strengthening to almost sure convergence follows from exactly the same argument used to prove part (a), with "in probability" replaced everywhere by "almost surely." By part (b) of Theorem 1, a sufficient condition for $u_j^{(t)} \to -\infty$ (each $j \in S$) almost surely is that the type process be pseudo-renewal; PPR processes are exactly those that are both FOSD and pseudo-renewal. The claim concerning conditional variances then follows immediately. □

E. First-Best Optimal Contract

This OA describes the first-best contract that arises under symmetric information. For each $\mathbf{v} \in \mathcal{U}^d$, let $\Gamma^{\mathrm{FB}}(\mathbf{v}) := \left\{(u_i, \mathbf{w}_i)_{i\in S} \in (\mathcal{U} \times \mathcal{U}^d)^d : [\mathbf{PK}_i] \text{ holds for all } i \in S\right\}$. The principal's first-best value function is then

$$\text{[FB]} \qquad Q^*(\mathbf{v}, s) := \inf_{\xi \in \Xi^{FB}(\mathbf{v})} \mathsf{E}\left[\sum_{t=0}^{\infty} \alpha^t C\big(u_\xi^{(t)}, s^{(t+1)}\big) \Big| s^{(0)} = s\right]$$

Lemma E.1. The first-best value function $Q^* : \mathcal{U}^d \times S \to \mathbb{R}$ satisfies the functional equation

$$\text{[E.1]} \qquad Q^*(\mathbf{v}, s) = \inf_{(u_i, \mathbf{w}_i)_{i \in S} \in \Gamma^{FB}(\mathbf{v})} \sum_{i \in S} f_{si} \left[C(u_i, i) + \alpha Q^*(\mathbf{w}_i, i)\right]$$

and each $Q^*(\cdot, s)$ is convex and continuously differentiable. Moreover, the infimum in [**E.1**] is attained at each $(\mathbf{v}, s) \in \mathcal{U}^d \times S$.

The proof of Lemma E.1 is standard, and hence omitted. Given Lemma E.1, solving for the optimal policy in [**E.1**] reduces to a smooth, convex, finite-dimensional minimization problem. From the first-order and envelope conditions, it is easy to deduce the following characterization of the first-best contract:

Lemma E.2. Fix any initial promise $\mathbf{v} \in \mathcal{U}^d$. There exists a unique optimal recursive contract. If the first type is $s^{(0)} = s \in S$, the optimal contract satisfies

$$\text{[E.2]} \qquad u_i^{(t)}(\mathbf{v}, s^{(0)} = s) = (1 - \alpha) \cdot v_s \qquad \text{for all } i \in S$$

$$\text{[E.3]} \qquad \mathbf{w}_i^{(t)}(\mathbf{v}, s^{(0)} = s) = v_s \mathbf{1} \qquad \text{for all } i \in S$$

The value function $Q^*(\cdot, s)$ is strictly convex for each $s \in S$, the derivative is strictly positive $\mathsf{D}Q^*(\mathbf{v}, s) \gg \mathbf{0}$, and thus $\mathsf{D}_1\, Q^*(\mathbf{v}, s) > 0$ for all $(\mathbf{v}, s) \in \mathcal{U}^d \times S$.

Consider also the *full-information efficiency problem*

$$\text{[}\mathbf{Eff}_i^{\mathbf{FB}}\text{]} \qquad \begin{aligned} K^*(v, s) := \min_{\mathbf{v} \in \mathcal{U}^d} &\; Q^*(\mathbf{v}, s) \\ \text{s.t.} \qquad &\mathsf{E}^{\mathbf{f}_s}[\mathbf{v}] \geq v \end{aligned}$$

for each $i \in S$. This is the full-information analogue of the efficiency problem [$\mathbf{Eff}_i$] defined in Subsection 4.3. The following characterization then follows immediately:

Lemma E.3. For each $s \in S$ and $v \in \mathcal{U}$, the full-information efficiency problem [$\mathbf{Eff}_i^{\mathbf{FB}}$] has the unique solution

$$\text{[E.4]} \qquad \mathbf{v}^*(v, s) = v \cdot \mathbf{1}$$

and its value function is given by

$$K^*(v,s) = \frac{U^{-1}((1-\alpha)v)}{1-\alpha} - \mathsf{E}\left[\sum_{t=0}^{\infty} \alpha^t \omega^{(t)} \middle| \, s^{(0)} = s\right] \tag{E.5}$$

which is strictly increasing, strictly convex, and continuously differentiable, and satisfies the Inada conditions $\lim_{v\to-\infty} K^{*\prime}(v,s) = 0$ and $\lim_{v\to 0} K^{*\prime}(v,s) = +\infty$.

F. Pathwise Properties of Markov Chains

We collect here some miscellaneous facts about paths of Markov chains, which are used in the proof of Theorem 1 (see OA-C.4). Let $\mathfrak{X}$ be the (countable) state space for a Markov process with transition probabilities $P(x,B)$ denoting the probability of transitioning from x to $B \subseteq \mathfrak{X}$. Let $\mathbf{P}$ denote the induced probability measure on the path space $\mathfrak{X}^\infty$.

Lemma F.1. Let (X_n) be an $\mathfrak{X}$-valued Markov process with transitions given by the kernel P, and suppose $x \in \mathfrak{X}$ is *recurrent*. Then, $\mathbf{P}(X_n = x$ for infinitely many $n \mid X_0 = x) = 1$.

An elementary proof can be found in Shiryaev (1995, p. 577). To apply this result to our setting, let $\mathfrak{X} := S$ and let $\mathbf{P}$ denote the measure on $\mathscr{H} = S^\infty$ induced by the type process defined in Section 2.

Proposition F.2. The event $\{h \in \mathscr{H} : (s^{t-1}, s^t) = (i,j)$ for infinitely many $t\}$ occurs $\mathbf{P}$-a.s. for all $i, j \in S$.

Proof of Proposition F.2. Recall from Section 2 that $S = \{1,\dots,d\}$ with transition probabilities $P(i,j) := f_{ij} > 0$. It is useful to consider the *bivariate* Markov chain with states $B_S := \{(i,j) : i,j \in S\}$ and transition probabilities Q given by $Q\big((i,j),(k,\ell)\big) = \mathbf{1}\{j = k\}P(k,\ell)$, where the indicator $\mathbf{1}\{j=k\} = 1$ if $j = k$ and 0 otherwise. Then, the two-step transition probabilities are given by $Q^{(2)}\big((i,j),(k,\ell)\big) = P(j,k)P(k,\ell) > 0$. Therefore, all states communicate with each other, which implies that the Markov chain is indecomposable. But because the state space B_S is finite, by Theorem 1 (and the subsequent discussion) on p. 580 of Shiryaev (1995), at least one of the states must be recurrent. The indecomposability of the process then implies that all states are recurrent. An application of Lemma F.1 to the bivariate chain completes the proof. □

Corollary F.3. Let $\mathscr{I}_j := \{h \in \mathscr{H} : (s^{t-1}, s^t) = (d,j)$ infinitely often$\}$, and $\mathscr{I} := \bigcap_{j=1}^d \mathscr{I}_j$. Then, $\mathbf{P}(\mathscr{I}_j) = 1$ and hence $\mathbf{P}(\mathscr{I}) = 1$.