ECONOMIC RESEARCH
FEDERAL RESERVE BANK OF ST. LOUIS
WORKING PAPER SERIES

Insurance and Inequality with Persistent Private Information

Authors	Alex Bloedel, Vijay Krishna, and Oksana Leukhina
Working Paper Number	2018-020E
Revision Date	December 2021
Citable Link	https://doi.org/10.20955/wp.2018.020
Suggested Citation	Bloedel, A., Krishna, V., Leukhina, O., 2021; Insurance and Inequality with Persistent Private Information, Federal Reserve Bank of St. Louis Working Paper 2018-020. URL https://doi.org/10.20955/wp.2018.020

Federal Reserve Bank of St. Louis, Research Division, P.O. Box 442, St. Louis, MO 63166

The views expressed in this paper are those of the author(s) and do not necessarily reflect the views of the Federal Reserve System, the Board of Governors, or the regional Federal Reserve Banks. Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment.

Insurance and Inequality with Persistent Private Information^{*}

ALEXANDER W. BLOEDEL[†]   R. VIJAY KRISHNA[‡]   OKSANA LEUKHINA[§]

12TH 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 workhorse normative models that economists use to study optimal risk-sharing (e.g., in the large literatures on dynamic taxation and unemployment insurance).

---

(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.

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 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

---

(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.

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 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

Main Results. Our main result, Theorem 4.1, 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 4.2, 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 $T \rightarrow \infty$.

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 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

---

Subsection 5.2 for further discussion.

(9) For instance, the specialization of Theorem 4.1 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 4.1 can easily be adapted to generalize related long-run convergence results in some other dynamic contracting environments.

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

(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”

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 4.1 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 4.1 and 4.2) and explicitly solve for the full “recursive domain” of implementable promised utility vectors (Theorem B.1) 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 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

---

(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.

- (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

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 Supplementary Appendix (henceforth SA).

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 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 $U : (\underline{c}, \infty) \rightarrow \mathbb{R}$, where $\underline{c} \geq -\infty$. Let $\mathcal{U} = U((\underline{c}, \infty)) \subset \mathbb{R}$ 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 $\lim_{c \rightarrow \underline{c}} U'(c) = +\infty$ and $\lim_{c \rightarrow \infty} U'(c) = 0$.
- (b) It is bounded above and unbounded below: $\lim_{c \rightarrow \infty} U(c) = 0$ and $\lim_{c \rightarrow \underline{c}} U(c) = -\infty$, so that $\mathcal{U} = (-\infty, 0)$.

---

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.

- (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 $\omega_d > \omega_{d-1} > \dots > \omega_1$. 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 $j \leq i$.

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

(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 $\eta = \underline{c} = 0$.

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

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 $\mathbf{f}_i = (f_{i1}, \dots, f_{id})$, 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 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 $\mathbf{v}^{(0)} := (v_1, \dots, v_d) \in \mathcal{U}^d$ to the agent¹⁶ 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

---

(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 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*, 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 $\mathbf{v}^{(0)} = (v_1, \dots, v_d)$ 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

---

(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.

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* $(u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times \mathcal{U}^d)^d$, 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 $\psi : \mathcal{U} \times S \times S \rightarrow \mathcal{U}$ is defined as

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

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

$$\begin{aligned} [\text{PK}_i] \quad & v_i = u_i + \alpha \mathbf{E}^{f_i} [\mathbf{w}_i] \\ [\text{IC}_{ij}] \quad & v_i \geq \psi(u_j, i, j) + \alpha \mathbf{E}^{f_i} [\mathbf{w}_j] \end{aligned}$$

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

-
- (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 .
  - (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.

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

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

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., $\psi(u_1, 2, 1) = U(\omega_2 + C(u_1, 1)) > u_1$. The Markov information rent is also positive *if* the contingent 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 B.1 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 $\xi : D \times S \times S \rightarrow \mathcal{U} \times D$, written as $\xi = (\xi^f, \xi^c)$, where $\xi^f(\mathbf{v}, s, i) = u_i(\mathbf{v}, s) \in \mathcal{U}$ provides *flow* consumption utilities to i -type reports, and $\xi^c(\mathbf{v}, s, i) = \mathbf{w}_i(\mathbf{v}, s) \in D$ similarly provides *contingent continuation* utilities. We say that ξ is *feasible at* $(\mathbf{v}, s) \in D \times S$ if the menu $(\xi(\mathbf{v}, s, i))_{i \in S} \in \Gamma(\mathbf{v})$, where the principal's *constraint correspondence* $\Gamma : D \rightrightarrows (\mathcal{U} \times D)^d$ is defined by

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

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 $\tilde{u}_\xi := (u_\xi^{(t)})_{t=0}^\infty$, 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:

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

Note that the expectation in $[\mathbf{RP}]$ 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” $\mathbf{f}_s$ over the initial $t = 0$ type. A recursive contract ξ^* is *optimal* if it attains the infimum in $[\mathbf{RP}]$.

Transversality. Since the agent's utility function is unbounded below, *a priori* it is possible that the principal's recursive problem $[\mathbf{RP}]$ is a relaxation of her sequential problem $[\mathbf{SP}]$ because the recursive constraints only impose one-step promise keeping and deter one-step deviations from truthtelling. While standard induction implies that they

---

(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 $[\mathbf{RP}]$ and $[\mathbf{FE}]$ 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).

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.

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

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

where $(\mathbf{v}_\xi^{(t)}(h))_{t=0}^\infty$ denotes the (deterministic) sequence of contingent promises along the path $h \in \mathcal{H}$. 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 $\mathbf{v} \in \mathcal{U}^d$ 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) $v_1^{(t)} = \dots = v_d^{(t)}$. The principal’s value function in the full information problem, which will be referenced below, is denoted $Q^* : \mathcal{U}^d \times S \rightarrow \mathbb{R}$.

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.

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

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 $D \times S$.

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

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

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 [\[IC_{ij}\]](#) ($i > j$) as *strict* inequalities. For each $\mathbf{v} \in D$, $\Gamma_\circ(\mathbf{v}) \neq \emptyset$.

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 $P : D \times S \rightarrow \mathbb{R}$ satisfies the functional equation

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

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 the SA, 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 $\pi \in \Delta(S)$ 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 terminology 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 4.1 (Immiseration). *Suppose the environment is [TVC]-Regular. Under the optimal contract, as $t \rightarrow \infty$:*

- (a) *Promised utilities diverge: $v_i^{(t)} \rightarrow -\infty$ almost surely for all $i \in S$.*
- (b) *Flow utilities diverge: $u_i^{(t)} \rightarrow -\infty$ 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: $c_i^{(t)} + \omega_i \rightarrow \underline{c}$ 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 4.1 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 4.1 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 of Assumption Markov in Section 5.

-
- (26) For each $i \in S$, let $(c_i^{(t)})_{t \geq 0}$ 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.

The basic intuition for Theorem 4.1 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 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 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 4.1, which relies on the Martingale Convergence Theorem.

4.2. Backloaded Incentives and Relative Immiseration

Intuitively, the principal drives the 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 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.

(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.

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

- (a) $v_i^{(t)} - v_{i-1}^{(t)} \rightarrow +\infty$ in probability for all $i = 2, \dots, d$. Moreover, for all $t \in \mathbb{N}$,

$$\text{Var} \left(v_{s^{(t+k)}}^{(t+k-1)} \middle| \mathbf{v}^{(t)}, s^{(t)} \right) \rightarrow +\infty \quad \text{as } k \rightarrow \infty.$$

- (b) *If, in addition, immiseration occurs almost surely (namely, $u_i^{(t)} \rightarrow -\infty$ 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*

$$\text{Var} \left(v_{s^{(t+1)}}^{(t)} \middle| \mathbf{v}^{(t)}, s^{(t)} \right) \rightarrow +\infty \quad \text{as } t \rightarrow \infty$$

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

The proof of Theorem 4.2 is in Online Appendix D. While we have argued that Theorem 4.2 in some sense explains why immiseration arises in the first place, mathematically speaking it is a fairly simple consequence of Theorem 4.1. Intuitively, the principal does not waste the ability to provide high-powered incentives as they become affordable (as they do per Theorem 4.1). To see the gist of the argument, note that, as in Subsection 3.1, substituting the promise keeping constraint $[\mathbf{PK}_i]$ ($i = j$) into the incentive constraint $[\mathbf{IC}_{ij}]$ ($i > j$) gives an alternative way to write the latter:

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

When the type process satisfies FOSD, Theorem B.1 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 4.1 implies that the iid information rent term, which is always non-negative,

grows without bound. It then follows from [IC*_{ij}] (with $j = i - 1$) that the difference $v_i - v_{i-1}$ must also grow without bound.

It follows that the agent's continuation utility becomes increasingly uncertain over time, as in Theorem 4.2'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 $v_{s(t+1)}^{(t)}$ is a random variable with conditional mean given by the agent's date- t *ex ante* promised utility

$$[4.1] \quad \mathbb{E} \left[v_{s(t+1)}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right] = \sum_{i=1}^d f_{s^{(t)}, i} v_i^{(t)}$$

and conditional variance

$$[4.2] \quad \text{Var} \left(v_{s(t+1)}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right) = \sum_{i=1}^d f_{s^{(t)}, i} \left(v_i^{(t)} - \mathbb{E} \left[v_i^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right] \right)^2.$$

When the type process is PPR, Theorem 4.2(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 $t \rightarrow \infty$. This represents one sense in which the quality of risk sharing degrades in the long run.

Another interpretation of Theorem 4.2 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 $t \rightarrow \infty$. Viewed in this light, Theorem 4.2 expresses the idea that “pathwise welfare inequality” increases without bound in the long run.³⁰

---

(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 $t \rightarrow \infty$. In their model, which considers a society with a continuum of agents and a per-period aggregate resource constraint, this implies that the society's *cross-sectional* inequality explodes due to the *cumulative* effect of all past shocks (in that some agents will have received unboundedly more total endowment in periods $\tau \leq t$ than other agents as $t \rightarrow \infty$). By contrast, Theorem 4.2(b) says that a *single* shock leads to unboundedly large welfare differences between two agents with *otherwise identical histories* as $t \rightarrow \infty$.

4.3. Proof Sketch for Theorem 4.1

This subsection sketches how we translate the cost-smoothing intuition described above into a formal proof of Theorem 4.1 (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 4.1 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 $v = \mathbb{E}^f[\mathbf{v}]$ is the agent's *ex ante* promised utility under the (type-independent) transition probability $f \in \Delta(S)$. 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 $DP(\mathbf{v}, s) = (P_1(\mathbf{v}, s), \dots, P_d(\mathbf{v}, s))$, where $P_i(\mathbf{v}, s)$ is the partial derivative with respect to the component v_i . The directional derivative of $P(\cdot, s)$ in direction $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$ is written as $D_1 P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} [P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)] / \varepsilon$.

Proposition 4.3. If the environment is Regular, then the stochastic process $(D_1 P(\mathbf{v}^{(t)}, s^{(t)}))_{t=0}^{\infty}$ 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 $D_1 P(\mathbf{v}, s)$. 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 $\sum_{i=1}^d f_{si} D_1 P(\mathbf{w}_i, i)$. By an envelope argument, this perturbation is locally optimal, implying that the marginal costs are equal, giving us precisely the martingale property $D_1 P(\mathbf{v}, s) = \sum_{i=1}^d f_{si} D_1 P(\mathbf{w}_i, i)$.
- (b) *Differentiation in Direction 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 [IC_{ij}^{*}] for the agent's incentive constraints, it is clear that a perturbation of $\mathbf{v}$ leaves the left-hand side of the [IC_{ij}^{*}] constraints unchanged if, and only if, it is taken in the direction 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., $E^{f_i}[\hat{\mathbf{w}}_i] - E^{f_i}[\mathbf{w}_i] = \varepsilon$. In general, other types $j \neq i$ value this perturbation differently, as they have different beliefs about future types (i.e., $f_j \neq f_i$). 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 $E^{f_j}[\hat{\mathbf{w}}_i - \mathbf{w}_i] = \varepsilon$ 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 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 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

-
- (31) When the transition probabilities $[f_i]_{i=1}^d$ 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 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 the SA shows that for each $i > 1$, there exists an open set of implementable promised utility vectors $\mathbf{v} \in D$ for which $P_i(\mathbf{v}, s) < 0$.

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 $w_i = \mathbb{E}^f[\mathbf{w}_i]$ vary with his reported type i (here 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 continuation contracts, giving the principal new opportunities to *screen the agent through continuation contracts* that manipulate his Markov information rent (recall $[\mathbf{IC}_{ij}^*]$). 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 $\mathbb{E}^{f_i}[\mathbf{w}_i]$ independent of his reported type i by choosing the interim continuation promises $\{\mathbf{w}_i\}_{i \in S}$ to manipulate the deviation continuation payoffs $\mathbb{E}^{f_j}[\mathbf{w}_i]$ ($j \neq i$) appropriately. Given this additional channel for incentive provision, it is not clear whether (a) the marginal cost martingale splits (i.e., satisfies $D_1 P(\mathbf{w}_1) < D_1 P(\mathbf{v}) < D_1 P(\mathbf{w}_d)$) 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., $[\mathbf{IC}_{ij}^*]$ 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 $w \in \mathcal{U}$ starting in period $t + 1$ (before his period $t + 1$ type is

realized). The cost-minimizing way to do this is to solve the *efficiency problem (at (w, i))*:

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

Solutions to $[\mathbf{Eff}_i]$ 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 $[\mathbf{IC}_{ij}^*]$ with $i > j$. However, Assumption NHB implies that the highest type's promised utility $\mathbf{w}_d$ does not enter any of these incentive constraints, 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 $[\mathbf{Eff}_i]$ in terms of dual variables — namely, the derivative $DP(\mathbf{w}_i, i)$, which is pinned down by the first-order conditions for $[\mathbf{Eff}_i]$. When types are iid (the right-hand panel in Figure 2), there is a single efficiency ray because the principal's value

(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 note that the efficiency problem $[\mathbf{Eff}_i]$ is analogous to what Fernandes and Phelan (2000) call the “planner's problem,” while our recursive problem $[\mathbf{RP}]$ 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).*

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: $\mathbf{v}^{(t)} \rightarrow \mathbf{y} \in D$ 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 extended to (almost) all histories, establishing that $D_1 P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ 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

(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.

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 $DP(\mathbf{v}^{(t)}, s^{(t)})$ 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 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 4.1 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 4.1–4.2 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 $c_i^{(t)} \rightarrow +\infty$ and $u_i^{(t)} \rightarrow 0$ 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 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 $T \rightarrow \infty$), 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 4.1 (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 $T \rightarrow \infty$, 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

---

(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.

(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).

domain D has nonempty interior (Theorem B.1(a)), so that the principal’s marginal cost martingale $D_1P(\mathbf{v}^{(t)}, s^{(t)})$ 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 *deterministic* (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 *submartingale*, 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 4.1: 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 $\lim_{c \rightarrow \infty} U'(c) = 0$ 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

---

(37) If $M := \lim_{c \rightarrow \infty} U'(c) > 0$, 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

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 4.1, 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-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

---

$M \rightarrow 0$. 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 4.2 *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 $c_i^{(t)} + \omega_i \rightarrow \underline{c}$ for all i would imply that $v_i^{(t)}, u_i^{(t)} \rightarrow 0$ 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.

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 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, Troschkin 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 underreporting 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

-
- (40) In the separable Mirrleesian model, the principal’s marginal cost of promised utility admits the representation $DP_1(\mathbf{v}, s) = 1/u'(c(\mathbf{v}, s))$, 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.
- (41) **NHB** is used in the proof of Theorem 4.1 (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 4.2 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).

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

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 the SA, 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 $[\mathbf{Eff}_i]$ (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 $[\mathbf{IC}_{21}^*]$. 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

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

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 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

---

(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 4.1 (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).

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 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 4.1 and B.1).

1.31

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 B.1) 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 $\mathcal{G}$ 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,

(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.

(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).

$\mathcal{H}$ 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 $\sigma_t : G^{t+1} \times H^t \rightarrow S$. The *truthful strategy* σ^* is defined by $\sigma_t^*((g^t, s^t), h^t) = s^{t+1}$ for all $g^t \in G^t$ and $h^t \in H^t$. Strategy σ is *admissible* if $\sigma_t((g^t, s^{t+1}), h^t) \leq s^{t+1}$ 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 $u_i := U(c_i + \omega_i)$. 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 $\tilde{u} := (u^{(t)})_{t=0}^\infty$, 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 $\psi : \mathcal{U} \times S \times S \rightarrow \mathcal{U}$ defined as $\psi(u, i, j) := U(\omega_i + C(u, j))$ 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 $(s^{(t+1)}, \hat{s}^{(t+1)}, u^{(t)})_{t=0}^\infty$ over true types, reported types, and flow utility allocations. Taking $\tilde{u}$ as given, denote the law of this process by $\mathbf{P}^\sigma \in \Delta(S^\infty \times S^\infty \times \mathcal{U}^\infty)$ and its associated expectation operator by $\mathbf{E}^\sigma[\cdot]$. Thus, the agent's preferences over sequential contracts and admissible reporting strategies are represented by the lifetime utility function $\hat{U} : \mathcal{A} \times \Sigma \times S \rightarrow \mathcal{U} \cup \{-\infty\}$ defined by

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

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

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

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

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

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

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

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

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 $i \in S$

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

(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

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

for all $i \in S$. Sending $T \rightarrow \infty$, using the Monotone Convergence Theorem on the first term, and using the Bounded Convergence Theorem on the second term (it applies under [\[TVC\]](#)) yields [\[DP\]](#), as desired. The proof of part (b) is analogous (cf. Green 1987, Lemma 2). $\square$

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

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

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

- (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 $\Gamma : D \rightarrow (\mathcal{U} \times D)^d$ 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 B.1(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 the SA. 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 $\{f_i\}_{i \in S}$ 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 the SA. 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 solved binary-state example described in Subsection 5.3 is Regular. Since AR(1)

(50) For example, consider the case in which $d = 2$ and types are MLRP. In our formulation, Theorem B.1(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 $v^p(s) := E^{f^s}[\mathbf{v}]$ for an agent whose report of s was truthful, and the corresponding *threat point* utility $v^\dagger(s) := E^{f_{3-s}}[\mathbf{v}]$ 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 B.1(e), it is easy to see that $W(1) = \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{bmatrix} V_2$ and $W(2) = \begin{bmatrix} f_{21} & f_{22} \\ f_{11} & f_{12} \end{bmatrix} V_2$. Thus, $W(1)$ and $W(2)$ are 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.

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.

1.2

References

- Abreu, Dilip, David Pearce and Ennio Stacchetti. 1990. ‘Toward a Theory of Discounted Repeated Games with Imperfect Monitoring’. *Econometrica* 58 (5): 1041–1063. (Cit. on p. 70).
- Allen, Franklin. 1985. ‘Repeated principal-agent relationships with lending and borrowing’. *Economics Letters* 17 (1-2): 27–31. (Cit. on pp. 4, 32, 34).
- Arellano, Manuel, Richard Blundell and Stéphane Bonhomme. 2017. ‘Earnings and consumption dynamics: a nonlinear panel data framework’. *Econometrica* 85 (3): 693–734. (Cit. on p. 4).
- 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. (Cit. on p. 2).
- Atkeson, Andrew, and Robert E. Lucas Jr. 1992. ‘On Efficient Distribution with Private Information’. *Review of Economic Studies* 59 (3): 427–453. (Cit. on pp. 2–4, 10, 15, 16, 18, 21, 30, 33).
- 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. (Cit. on p. 19).
- Battaglini, Marco, and Rohit Lamba. 2019. ‘Optimal Dynamic Contracting: The First-Order Approach and Beyond’. *Theoretical Economics* 14 (4): 1435–1482. (Cit. on pp. 7, 8, 28, 32).
- Bloedel, Alexander W., and R. Vijay Krishna. 2015. *Misery, Persistence, and Growth*. Tech. rep. Stanford University. (Cit. on pp. 1, 33, 34).
- 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. (Cit. on pp. 4, 16, 18, 28, 29, 34).
- . 2020. *Persistent Private Information Revisited*. Tech. rep. Stanford University. (Cit. on pp. 4, 9, 16, 18, 28, 29, 34, 40, 111).

- Broer, Tobias, Marek Kapička and Paul Klein. 2017. ‘Consumption risk sharing with private information and limited enforcement’. *Review of Economic Dynamics* 23:170–190. (Cit. on p. 34).
- Cole, Harold L., and Narayana R. Kocherlakota. 2001a. ‘Dynamic Games with Hidden Actions and Hidden States’. *Journal of Economic Theory* 98 (1): 114–126. (Cit. on p. 34).
- 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. (Cit. on pp. 4, 9, 32, 34).
- Di Tella, Sebastian, and Yuliy Sannikov. 2021. ‘Optimal Asset Management Contracts with Hidden Savings’. *Econometrica* 89 (3): 1099–1139. (Cit. on p. 31).
- Doepke, Matthias, and Robert M. Townsend. 2006. ‘Dynamic mechanism design with hidden income and hidden actions’. *Journal of Economic Theory* 126 (1): 235–285. (Cit. on p. 34).
- Echenique, Federico. 2005. ‘A Short and Constructive Proof of Tarski’s Fixed-Point Theorem’. *International Journal of Game Theory* 33 (2): 215–218. (Cit. on pp. 72, 79, 82).
- Farhi, Emmanuel, and Iván Werning. 2007. ‘Inequality and Social Discounting’. *Journal of Political Economy* 115 (3): 365–402. (Cit. on pp. 2, 4).
- . 2013. ‘Insurance and Taxation over the Lifecycle’. *Review of Economic Studies* 80 (2): 596–635. (Cit. on pp. 4, 7, 31).
- Farhi, Emmanuel, et al. 2012. ‘Non-linear capital taxation without commitment’. *Review of Economic Studies* 79 (4): 1469–1493. (Cit. on p. 2).
- Feldman, Naomi E., and Joel Slemrod. 2007. ‘Estimating tax noncompliance with evidence from unaudited tax returns’. *The Economic Journal* 117 (518): 327–352. (Cit. on p. 31).
- Fernandes, Ana, and Christopher Phelan. 2000. ‘A Recursive Formulation for Repeated Agency with History Dependence’. *Journal of Economic Theory* 91:223–247. (Cit. on pp. 6, 7, 14, 25, 31, 34, 39).
- Fu, Shiming, and R. Vijay Krishna. 2019. ‘Dynamic Financial Contracting with Persistence’. *RAND Journal of Economics* 50 (2): 418–452. (Cit. on pp. 31, 35).
- Garrett, Daniel, Alessandro Pavan and Juuso Toikka. 2018. *Robust Predictions in Dynamic Screening*. Tech. rep. Northwestern University. (Cit. on pp. 8, 32).
- Gärtner, Bernd, and Jiri Matousek. 2012. *Approximation Algorithms and Semidefinite Programming*. Springer-Verlag. (Cit. on p. 83).

- Golosov, Mikhail, Narayana R. Kocherlakota and Aleh Tsyvinski. 2003. ‘Optimal Indirect and Capital Taxation’. *Review of Economic Studies* 70 (3): 569–588. (Cit. on pp. 7, 31).
- Golosov, Mikhail, Maxim Troshkin and Aleh Tsyvinski. 2016. ‘Redistribution and Social Insurance’. *American Economic Review* 106 (2): 359–386. (Cit. on pp. 4, 7, 31).
- Golosov, Mikhail, and Aleh Tsyvinski. 2007. ‘Optimal taxation with endogenous insurance markets’. *The Quarterly Journal of Economics* 122 (2): 487–534. (Cit. on pp. 4, 31).
- Golosov, Mikhail, Aleh Tsyvinski and Nicolas Werquin. 2016. ‘Recursive Contracts and Endogenously Incomplete Markets’. In *Handbook of Macroeconomics, Volume 2*. (Cit. on pp. 3, 30, 34).
- 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. (Cit. on pp. 2–5, 9, 11, 16, 32–34, 37).
- Green, Jerry R, and Jean-Jacques Laffont. 1986. ‘Partially verifiable information and mechanism design’. *The Review of Economic Studies* 53 (3): 447–456. (Cit. on p. 35).
- Guo, Yingni, and Johannes Hörner. 2020. *Dynamic Allocation without Money*. Tech. rep. (Cit. on pp. 34, 35).
- Guvenen, Fatih, et al. 2021. ‘What Do Data on Millions of US Workers Reveal About Lifecycle Earnings Dynamics?’ *Econometrica* 89 (5): 2303–2339. (Cit. on p. 4).
- Halac, Marina, and Pierre Yared. 2014. ‘Fiscal rules and discretion under persistent shocks’. *Econometrica* 82 (5): 1557–1614. (Cit. on p. 34).
- Hatcher, Allen. 2001. *Algebraic Topology*. Cambridge University Press. (Cit. on p. 50).
- Hörner, Johannes, Xiaosheng Mu and Nicolas Vieille. 2017. *Keeping Your Story Straight: Truthtelling and Liespotting*. Tech. rep. Yale University. (Cit. on p. 17).
- 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. (Cit. on p. 33).
- Kamihigashi, Takashi. 2014. ‘Elementary results on solutions to the Bellman equation of dynamic programming: existence, uniqueness, and convergence’. *Economic Theory* 56 (2): 251–273. (Cit. on p. 101).
- Kapička, Marek. 2013. ‘Efficient allocations in dynamic private information economies with persistent shocks: A first-order approach’. *The Review of Economic Studies*: rds045. (Cit. on p. 7).

- Khan, Aubhik, and Bashkar Ravikumar. 2001. ‘Growth and risk-sharing with private information’. *Journal of Monetary Economics* 47 (3): 499–521. (Cit. on p. 33).
- Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press. (Cit. on pp. 2, 10, 31).
- Krasikov, Ilia, and Rohit Lamba. 2021. ‘A Theory of Dynamic Contracting with Financial Constraints’. *Journal of Economic Theory* 193. (Cit. on p. 35).
- Kreps, David M. 1977. ‘Decision Problems with Expected Utility Criteria, I: Upper and Lower Convergent Utility’. *Mathematics of Operations Research* 2 (1): 45–53. (Cit. on p. 103).
- Ljungqvist, Lars, and Thomas J. Sargent. 2012. *Recursive Macroeconomic Theory*. 3rd ed. Cambridge, Ma: MIT Press. (Cit. on p. 34).
- Lucas, Robert E., Jr. 1992. ‘On Efficiency and Distribution’. *Economic Journal* 102 (411): 233–247. (Cit. on p. 2).
- Luenberger, David G. 1969. *Optimization by Vector Space Methods*. Wiley. (Cit. on pp. 45, 54, 87).
- Matoušek, Jiří, and Bernd Gärtner. 2007. *Understanding and Using Linear Programming*. New York: Springer. (Cit. on p. 77).
- Meghir, Costas, and Luigi Pistaferri. 2004. ‘Income variance dynamics and heterogeneity’. *Econometrica* 72 (1): 1–32. (Cit. on p. 3).
- Newman, Andrew F. 2007. ‘Risk-bearing and entrepreneurship’. *Journal of Economic Theory* 137 (1): 11–26. (Cit. on p. 19).
- Olszewski, Wojciech, and Mikhail Safronov. 2021. *Efficient dynamic allocations in endowment economies with linear utilities and private information*. Tech. rep. (Cit. on p. 33).
- Pavan, Alessandro. 2016. ‘Dynamic mechanism design: Robustness and endogenous types’. In *2015 World Congress of the Econometric Society*. (Cit. on pp. 8, 32).
- Pavan, Alessandro, Ilya Segal and Juuso Toikka. 2014. ‘Dynamic Mechanism Design: A Myersonian Approach’. *Econometrica* 82 (2): 601–653. (Cit. on pp. 7, 12, 28).
- Phelan, Christopher. 1998. ‘On the Long Run Implications of Repeated Moral Hazard’. *Journal of Economic Theory* 79:174–191. (Cit. on pp. 9, 29–31, 33, 34).
- . 2006. ‘Opportunity and social mobility’. *Review of Economic Studies* 73 (2): 487–504. (Cit. on p. 2).

- . 1995. ‘Repeated Moral Hazard and One-Sided Commitment’. *Journal of Economic Theory* 66:488–506. (Cit. on pp. 2, 9).
- Shiryaev, Albert N. 1995. *Probability*. Springer. (Cit. on pp. 58, 68).
- Sleet, Christopher, and Şevin Yeltekin. 2006. ‘Credibility and endogenous societal discounting’. *Review of Economic Dynamics* 9 (3): 410–437. (Cit. on p. 2).
- . 2008. ‘Politically credible social insurance’. *Journal of monetary Economics* 55 (1): 129–151. (Cit. on p. 2).
- Spear, Stephen E, and Sanjay Srivastava. 1987. ‘On Repeated Moral Hazard with Discounting’. *Review of Economic Studies* 54 (4): 599–617. (Cit. on p. 34).
- Stokey, Nancy L., Robert E. Lucas Jr and Edward C. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Cambridge, Ma: Harvard University Press. (Cit. on pp. 101, 104, 105).
- 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. (Cit. on p. 3).
- Strulovici, Bruno. 2017. *Contract Negotiation and Screening with Persistent Information*. Tech. rep. Northwestern University. (Cit. on p. 25).
- . 2020. ‘Renegotiation-proof contracts with persistent private information’. Working paper. (Cit. on pp. 9, 25, 111).
- 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. (Cit. on pp. 2, 4–6, 9, 11, 16, 18, 19, 22, 24, 25, 32–34, 107, 110, 111, 118).
- Vohra, Rakesh V. 2005. *Advanced Mathematical Economics*. New York, NY: Routledge. (Cit. on p. 77).
- Wang, Cheng. 1995. ‘Dynamic insurance with private information and balanced budgets’. *The Review of Economic Studies* 62 (4): 577–595. (Cit. on p. 9).
- Williams, Noah. 2011. ‘Persistent Private Information’. *Econometrica* 79 (4): 1233–1275. (Cit. on pp. 3, 4, 6, 7, 9, 18, 28, 31, 34, 40).
- Zhang, Yuzhe. 2009. ‘Dynamic Contracting with Persistent Shocks’. *Journal of Economic Theory* 144:635–675. (Cit. on pp. 4, 5, 18, 31–34, 107).

Online Appendix

This Online Appendix (henceforth OA) provides proofs omitted from the main text. OA-C first presents the proof of Theorem 4.1, essential steps of which include the proofs of Proposition 4.3 and Theorem B.1(d). OA-D then presents the proof of Theorem 4.2. 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 Supplementary Appendix (SA) contains proofs of some auxiliary technical results referenced herein.

C. Proof of Theorem 4.1

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 4.1. OA-C.4 presents the main proof of convergence for Theorem 4.1.

C.1. Optimality Conditions

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

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

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

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

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

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

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

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

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

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

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

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

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 $D_1 P(\mathbf{v}, s) = \sum_{i \in S} P_i(\mathbf{v}, s)$, 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 $D_1 P(\mathbf{v}^{(t)}, s^{(t)})$ 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

$$\text{[C.1]} \quad D_1 P(\mathbf{v}, s) = \sum_{i=1}^d \lambda_i$$

For fixed $i \in S$, summing the **[FOC_{w_{ij}]}** over $j \in S$ gives

$$\text{[C.2]} \quad f_{si} \cdot D_1 P(\mathbf{w}_i, i) = \lambda_i + \sum_{k=1}^{i-1} \mu_{ik} - \sum_{k=i+1}^d \mu_{ki}$$

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 $\sum_{i \in S} f_{si} D_1 P(\mathbf{w}_i, i) = \sum_{i=1}^d \lambda_i$ which, combined with **[C.1]**, gives the martingale property $D_1 P(\mathbf{v}, s) = \sum_{i=1}^d f_{si} D_1 P(\mathbf{w}_i, i)$

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 $D_1 P(\cdot, s)$ 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 $D_1 P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} \frac{P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)}{\varepsilon}$ 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 $Q : D \times S \rightarrow \mathbb{R}$ 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 C.1. 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 B.1.) 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. $\square$

Claim C.2. The lattice Φ is complete.

Proof of Claim. Let $F \subseteq \Phi$ be nonempty and define $\bar{f}(\mathbf{v}, s) := \sup_{f \in F} f(\mathbf{v}, s)$ 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 $\bar{f} \in F$; 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 $\bar{f}(\mathbf{v}, s) > \bar{f}(\mathbf{v}', s)$. 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 $\bar{f}(\mathbf{v}, s) - \delta \geq \bar{f}(\mathbf{v}', s)$. By definition of the supremum, there exists some $f \in F$ such that $f(\mathbf{v}, s) > \bar{f}(\mathbf{v}, s) - \delta$. Combined with the earlier inequality and the definition of Φ , this implies that $f(\mathbf{v}, s) > \bar{f}(\mathbf{v}', s) \geq f(\mathbf{v}', s)$ which contradicts the fact that f is non-decreasing in the direction $\mathbf{1}$ by virtue of $f \in F \subset \Phi$. $\square$

(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.

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

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

Claim C.3. $T : \Phi \rightarrow \Phi$ 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, $\delta + TQ(\mathbf{v}, s) \geq \sum_{i \in S} f_{si} [C(u_i, i) + \alpha Q(\mathbf{w}_i, i)] \geq \sum_{i \in S} f_{si} [C(u_i, i) + \alpha Q(\mathbf{w}_i - \frac{\varepsilon}{\alpha} \mathbf{1}, i)] \geq TQ(\mathbf{v} - \varepsilon \mathbf{1}, s)$ 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. $\square$

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 $D_1 P(\mathbf{v}, s) = 0$. For this fixed state, define the function $f : Y \rightarrow \mathbb{R}$ 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 $D_1 P(\mathbf{v}, s) = 0$ 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 $D_1 P$ is strictly positive on $D \times S$. $\square$

C.3. Intermediate Steps Towards the Proof of Theorem 4.1

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 4.1. 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:

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

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 $\mathbf{w}^\dagger(\cdot, i) : \mathcal{U} \rightarrow \mathcal{U}^d$ 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 $\lim_{w \rightarrow -\infty} K'(w, i) = 0$ and $\lim_{w \rightarrow 0} K'(w, i) = +\infty$.

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 the SA. 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} ^{FB}] in OA-E. That value function is called $K^*(w, i)$, and clearly satisfies $K^* \leq K$ on $\mathcal{U} \times S$. Lemma E.3 states that $\lim_{w \rightarrow 0} K^*(w, i) = +\infty$. If $K(\cdot, i)$ did not satisfy $\lim_{w \rightarrow 0} K'(w, i) = +\infty$, then there would exist some $v \in \mathcal{U}$ such that $K^*(v, i) > K(v, i)$, a contradiction. Similarly, Lemma E.3 states that $\lim_{w \rightarrow -\infty} K^*(w, i) = 0$. If $K(\cdot, i)$ did not satisfy $\lim_{w \rightarrow -\infty} K'(w, i) = 0$, then there would exist some $v \in \mathcal{U}$ such that $K^*(v, i) > K(v, i)$, again a contradiction. This completes the proof. $\square$

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

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

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

and the envelope condition

$$[\mathbf{Env}_j\text{-}\mathbf{Eff}_i] \quad K'(w, i) = \zeta(w, i)$$

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

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

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

$$[\mathbf{C.3}] \quad K'(w, i) = D_1 P(\mathbf{w}^\dagger(w, i), i)$$

C.3.2. Facts Concerning the Marginal Cost Martingale

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

Proof. We first show that $DP(\cdot, s)$ is injective. To see this, notice first that because P is strictly convex, $DP(\cdot, s)$ is *strictly monotone* in the sense that for all $\mathbf{v}, \mathbf{v}' \in D$ such that $\mathbf{v} \neq \mathbf{v}'$, $\langle DP(\mathbf{v}, s) - DP(\mathbf{v}', s), \mathbf{v} - \mathbf{v}' \rangle > 0$. But now suppose $DP(\cdot, s)$ is not injective, so that there are $\mathbf{v}, \mathbf{v}' \in D$ distinct such that $DP(\mathbf{v}, s) = DP(\mathbf{v}', s)$. But this would imply that $0 = \langle \mathbf{0}, \mathbf{v} - \mathbf{v}' \rangle = \langle DP(\mathbf{v}, s) - DP(\mathbf{v}', s), \mathbf{v} - \mathbf{v}' \rangle > 0$, which is a contradiction. As D is open by Theorem B.1 and the derivative $DP(\cdot, s)$ 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 $DP(D, s) = Y_s$. $\square$

Lemma C.4. Suppose the environment is $[\mathbf{TVC}]$ -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 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. $\square$

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 $D_1 P(\cdot, s) : E_s \rightarrow \mathbb{R}_{++}$ is strictly increasing and is a homeomorphism.

Proof. This follows immediately from [C.3] and part (c) of Lemma C.2. $\square$

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

$$[\text{MS}_i] \quad D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\mathbf{w}_d, d) + \frac{\sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d)}{f_{di}} - \frac{\sum_{k=i+1}^d \mu_{ki}(\mathbf{w}_d, d)}{f_{di}}$$

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 d th envelope condition at state $(\mathbf{w}_d, d)$. It follows that

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

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

$$[\text{C.5}] \quad f_{sd} D_1 P(\mathbf{w}_d, d) = \lambda_d(\mathbf{v}, s) + \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{v}, s)$$

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

$$[\text{C.6}] \quad f_{dd} D_1 P(\tilde{\mathbf{w}}_d, d) = \lambda_d(\mathbf{w}_d, d) + \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d)$$

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

$$[\text{C.7}] \quad D_1 P(\tilde{\mathbf{w}}_d, d) = D_1 P(\mathbf{w}_d, d) + \frac{\sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d)}{f_{dd}}$$

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

$$[C.8] \quad f_{di} D_1 P(\tilde{\mathbf{w}}_i, i) = \lambda_i(\mathbf{w}_d, d) + \sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d) - \sum_{k=i+1}^d \mu_{ki}(\mathbf{w}_d, d)$$

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

$$[C.9] \quad \begin{aligned} D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\tilde{\mathbf{w}}_d, d) &- \frac{\sum_{k=1}^{d-1} \mu_{dk}(\mathbf{w}_d, d)}{f_{dd}} - \frac{\sum_{k=i+1}^d \mu_{ki}(\mathbf{w}_d, d)}{f_{di}} \\ &+ \frac{\sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d)}{f_{di}} - \left[\frac{\lambda_d(\mathbf{w}_d, d)}{f_{dd}} - \frac{\lambda_i(\mathbf{w}_d, d)}{f_{di}} \right] \end{aligned}$$

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 [E_{*i*}], this term vanishes. Thus, plugging the first line of [C.9] into [C.7] and rearranging delivers [MS_{*i*}]. $\square$

Lemma C.7. Suppose the environment is [TV_C]-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) $D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\mathbf{w}_d, d)$ 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 $D_1 P(\tilde{\mathbf{w}}_d, d) = D_1 P(\mathbf{w}_d, d)$ 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

$$[C.10] \quad D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\mathbf{w}_d, d) + \frac{\sum_{k=1}^{i-1} \mu_{ik}(\mathbf{w}_d, d)}{f_{di}}$$

It follows from [C.10] and dual feasibility that $D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\mathbf{w}_d, d)$ 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. $\square$

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*:

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

subject to

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

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

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 $[\text{PK}_i]$ for type i and incentive compatibility $[\text{IC}_{ji}^*]$ for all *higher* types $j > i$. As the notation suggests, the function $Q^i(\cdot)$ depends on $\mathbf{v}$ only through the components $(v_i, v_{i+1}, \dots, v_d)$, 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 $[\text{PK}_i]$ ($i = d$).

For each $i \in S$, define

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

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., $\Gamma(\mathbf{v}) = \Gamma_1(\mathbf{v}) \times \dots \times \Gamma_d(\mathbf{v})$.

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

$$[\text{C.12}] \quad P(\mathbf{v}, s) = \sum_{i=1}^d f_{si} Q^i(v_i, \dots, v_d)$$

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

Proof. It follows from the observation that $\Gamma(\mathbf{v}) = \Gamma_1(\mathbf{v}) \times \dots \times \Gamma_d(\mathbf{v})$ that we may rewrite $[\text{FE}]$ 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. $\square$

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 $\xi^{*f}(\cdot, \cdot, i) : D \times S \rightarrow \mathcal{U}$ and $\xi^{*c}(\cdot, \cdot, i) : D \times S \rightarrow D$ depend on $(\mathbf{v}, s)$ only through the components $(v_i, \dots, v_d)$.⁵²

---

(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 $j \geq i$.

- (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. $\square$

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

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))

$$\begin{aligned} \mathcal{L}^i(\mathbf{v}, \boldsymbol{\eta}, \boldsymbol{\sigma}, \mathbf{u}, \mathbf{w}) = & C(u_i, i) + \alpha P(\mathbf{w}_i, i) + \eta_i \left[v_i - u_i + \alpha \mathbf{E}^{\mathbf{f}_i} [\mathbf{w}_i] \right] \\ \text{[L}_i\text{]} \quad & - \sum_{j=i+1}^d \sigma_{ji} \left[v_j - v_i - \psi(u_i, j, i) + u_i - \alpha (\mathbf{E}^{\mathbf{f}_j} [\mathbf{w}_i] - \mathbf{E}^{\mathbf{f}_j} [\mathbf{w}_j]) \right] \end{aligned}$$

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 [PK_i] and $\sigma_{ji}(\mathbf{v}) \in \mathbb{R}_+$ is the multiplier on [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

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

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

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

- (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 [L_i], $i \in S$. Begin with part (a). The envelope conditions from [L_i] read

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

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

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

Substituting the ex ante envelope condition [Env_i] ($i = j$) into [C.16] delivers $\lambda_j(\mathbf{v}, s) = f_{sj} [\eta_j(\mathbf{v}) - \sum_{k=j+1}^d \sigma_{kj}(\mathbf{v})] - \sum_{k=1}^{j-1} f_{sk} \sigma_{jk}(\mathbf{v})$. 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

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

and the first-order condition for u_i in the ex ante problem, [FOC_{u_i}], is

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

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

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

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 $0 = \sum_{k=1}^{d-1} [\mu_{dk}(\mathbf{v}, s) - f_{sk}\sigma_{dk}(\mathbf{v})] = \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{v}, s)$ 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

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

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. $\square$

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

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

(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 $[\tilde{E}_i]$ and the envelope conditions $[\text{Env}_i]$, 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 $[\text{Env}_i]$ 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. $\square$

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 $c_1 > \dots > c_d$ satisfying $\omega_1 + c_1 = \dots = \omega_d + c_d$ 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 . $\square$

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 $\mathbf{v}_i$], [FOC u_i], [FOC $\mathbf{w}_{ij}$], respectively) reduce to

$$[\text{C.20}] \quad P_i(\mathbf{v}, s) = \lambda_i(\mathbf{v}, s) \quad \forall i \in S$$

$$[\text{C.21}] \quad f_{si}C'(u_i) = \lambda_i(\mathbf{v}, s) \quad \forall i \in S$$

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

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

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

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

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 $D_1P(\mathbf{v}, s') = \sum_{i=1}^d f_{s'i}\hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$. Similarly, set $i = s'$ in [C.24] and sum over $j \in S$ to get $D_1P(\mathbf{w}_{s'}, s') = \sum_{j=1}^d f_{s'j}\hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$. 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 $(\mathbf{v}, s)$. $\square$

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 $D_1 P(\tilde{\mathbf{w}}_{i^*(\mathbf{w}_d)}, i^*(\mathbf{w}_d)) \neq D_1 P(\mathbf{w}_d, d)$.

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. $\square$

C.4. Main Proof of Theorem 4.1

The following pieces of notation will be used extensively during the proof. Recall that we denote the space of (*infinite*) *histories*, or *paths*, by $\mathcal{H} := S^\infty$ 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 $\tau^{(t)}(h) := \sup \{T \leq t : s^T = d\}$. 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 $(\tau^{(t)})_{t=0}^\infty$ is $\mathbf{P}$ -a.s. non-decreasing.

Martingale Convergence. Define the event

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

It is easy to see that $\lim_{t \rightarrow \infty} \tau^{(t)}(h) = +\infty$ for all $h \in \mathcal{F}$. We note here that $\mathbf{P}(\mathcal{F}) = 1$ 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 $D_1(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ 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 $(D_1 P(\mathbf{v}^{(t)}, s^{(t)}))_{t=0}^\infty$ 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 $\mathcal{F} \subseteq \mathcal{H}$. So fix an arbitrary path $h := (s^t)_{t=0}^\infty \in \mathcal{F}$. 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 $\tau_i^t := \sup \{T \leq t : (s^T, s^{T+1}) = (d, i)\}$. It follows from the definition of the event $\mathcal{F}$ that $\lim_{t \rightarrow \infty} \tau_i^t = +\infty$ for all $i \in S$.

Suppose, towards contradiction, that $D_1 P(\mathbf{v}^t, s^t) \rightarrow C > 0$. It then follows from Lemmas C.4 and C.5 that $\mathbf{v}^{\tau_i^t} \rightarrow \mathbf{w}_d^* \in E_d$. The policy functions are continuous under Condition R.4, so $\xi^c(\mathbf{v}^{\tau_i^t}, d, i) \rightarrow \xi^c(\mathbf{w}_d^*, d, i) =: \tilde{\mathbf{w}}_i^*$. Fix $i \in S$. Because $P(\cdot, i)$ is continuously differentiable by Proposition 3.3, it follows that $\lim_{t \rightarrow \infty} D_1 P(\mathbf{v}^{\tau_i^t+1}, s^{\tau_i^t+1}) = D_1 P(\tilde{\mathbf{w}}_i^*, i)$, as the sequence $\{s^{\tau_i^t+1}\}$ 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 $D_1 P(\mathbf{w}_d^*, d) = D_1 P(\tilde{\mathbf{w}}_i^*, i)$ 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, $D_1 P(\mathbf{v}^t, s^t) \rightarrow 0$, 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 $(\boldsymbol{\nu}^{(t)})_{t=0}^\infty$, where $\boldsymbol{\nu}^{(t)} := \boldsymbol{\nu}(\mathbf{v}^{(t)}, s^{(t)})$ for each $t \in \mathbb{N}$. Similarly define the processes $(\boldsymbol{\lambda}^{(t)})_{t=0}^\infty$ and, for each $i \in S$, $(\boldsymbol{\mu}_{*,i}^{(t)})_{t=0}^\infty$, where $\boldsymbol{\lambda}(\mathbf{v}, s) := (\lambda_1(\mathbf{v}, s), \dots, \lambda_d(\mathbf{v}, s)) \in \mathbb{R}^d$ and $\boldsymbol{\mu}_{*,i}(\mathbf{v}, s) := (\mu_{i+1,i}(\mathbf{v}, s), \dots, \mu_{d,i}(\mathbf{v}, s)) \in \mathbb{R}_+^{d-i}$. 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, $\boldsymbol{\nu}^{(\tau^{(t)})} \rightarrow \mathbf{0}$ and $\boldsymbol{\lambda}^{(\tau^{(t)}+1)} \rightarrow \mathbf{0}$ almost surely.

Proof of Lemma C.17. Fix some path $h \in \mathcal{F}$ along which the marginal cost martingale converges to zero. By part (a) of Theorem 4.1, the set of such paths has full measure, and is thus sufficient for establishing almost sure convergence. By [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 4.1 (proved above) requires that $DP(\mathbf{v}^{\tau_t}, s^{\tau_t}) \rightarrow \mathbf{0}$. By the envelope conditions [Env_{*i*}], this translates to

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

It remains to show that $\boldsymbol{\mu}_{*,i}^{\tau_t} \rightarrow \mathbf{0}$ 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 [FOC_{w_{*ij*}}] with $i = 1$ at state $(\mathbf{v}^{\tau_t}, d)$ is $f_{d1} P_j(\tilde{\mathbf{w}}_1^{\tau_t}, 1) = f_{1j} (\lambda_1(\mathbf{v}^{\tau_t}, d) + 0) - \sum_{k=2}^d f_{kj} \mu_{k1}(\mathbf{v}^{\tau_t}, d)$. Because $\boldsymbol{\mu}(\cdot) \geq \mathbf{0}$ on $D \times S$, it follows from [C.25] that $f_{d1} P_j(\tilde{\mathbf{w}}_1^{\tau_t}, 1) \rightarrow 0$ and thus also that $\sum_{k=2}^d f_{kj} \mu_{k1}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$. This is true for

all $j \in S$. Because the type process is fully connected (Assumption [Markov](#)), it follows that

$$[C.26] \quad DP(\tilde{\mathbf{w}}_1^{\tau_t}, 1) \rightarrow 0$$

$$[C.27] \quad \mu_{*,1}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$$

Inductive step: Let $m \in S$. Suppose we have shown that $DP(\tilde{\mathbf{w}}_\ell^{\tau_t}, \ell) \rightarrow 0$ and $\mu_{*,\ell}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$ for all $\ell < m$. The first order condition [\[FOC \$w_{ij}\$ \]](#) with $i = m$ at state $(\mathbf{v}^{\tau_t}, d)$ is $f_{dm}P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) = f_{mj}(\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d)) - \sum_{k=m+1}^d f_{kj}\mu_{km}(\mathbf{v}^{\tau_t}, d)$. By [\[C.25\]](#) and the supposition, it follows that $\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$. As $\mu(\cdot) \geq 0$ on $D \times S$, it follows from the previous two displays that $f_{dm}P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) \rightarrow 0$ and $\sum_{k=m+1}^d f_{kj}\mu_{km}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$. This is true for all $j \in S$. Because the type process is fully connected (Assumption [Markov](#)), it follows that

$$[C.28] \quad DP(\tilde{\mathbf{w}}_m^{\tau_t}, m) \rightarrow 0$$

$$[C.29] \quad \mu_{*,m}(\mathbf{v}^{\tau_t}, d) \rightarrow 0$$

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

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

Proof of Lemma C.18. Suppose we have shown the desired convergence for all $k = 0, \dots, m-1$. 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. $\square$

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

Proof of Lemma C.19. Define the stochastic processes $(\delta^{(t)})_{t=0}^\infty$ and $(L^{(t)})_{t=0}^\infty$ by $\delta^{(t)} := \|\nu^{(t)}\|$, 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 distance of $\nu^{(t)}(h)$ from the zero vector and $L^{(t)}(h) \in \mathbb{N}$ denotes the lag since the last ω_d realization at date t , along path $h \in \mathcal{H}$.

To show convergence in probability, we must show that $\limsup_{t \rightarrow \infty} \mathbf{P}(\delta^{(t)} > \varepsilon) = 0$ 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: $A_{\varepsilon,t} := \{h \in \mathcal{H} : \delta^{(t)}(h) > \varepsilon\}$, $B_{k,t} := \{h \in \mathcal{H} : L^{(t)}(h) > k\}$, and $C_{\varepsilon,k,t} := \cup_{T \geq t} [A_{\varepsilon,T} \cap B_{k,T}^c]$. 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 $\mathbf{P}(A_{\varepsilon,t}) = \mathbf{P}(A_{\varepsilon,t} \cap B_{k,t}^c) + \mathbf{P}(A_{\varepsilon,t} \cap B_{k,t}) \leq \mathbf{P}(C_{\varepsilon,k,t}) + \mathbf{P}(B_{k,t})$, where the inequality follows from monotonicity of probability. Observe that $\lim_{t \rightarrow 0} \mathbf{P}(C_{\varepsilon,k,t} \cap \mathcal{F}) = 0$, 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 $\delta^{(\tau^{(t)}+m)} \rightarrow 0$ almost surely for all $m = 0, \dots, k$. Thus, it follows that

$$[\text{C.30}] \quad \limsup_{t \rightarrow \infty} \mathbf{P}(A_{\varepsilon,t}) \leq \limsup_{t \rightarrow \infty} \mathbf{P}(B_{k,t})$$

The remaining step of the proof is to show that there exists a function $H : \mathbb{N} \rightarrow [0, 1]$ that satisfies $\lim_{k \rightarrow \infty} H(k) = 1$, and such that $\lim_{t \rightarrow \infty} \mathbf{P}(B_{k,t}) = 1 - H(k)$ 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 $\limsup_{t \rightarrow \infty} \mathbf{P}(A_{\varepsilon,t}) \leq \inf_{k \in \mathbb{N}} (1 - H(k)) = 0$ 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 $(s^{(t)})$ is ergodic. Thus, there exists a unique stationary distribution $\pi \in \Delta(S)$, $\pi_i > 0$ for all $i \in S$, and $\lim_{t \rightarrow \infty} \mathbf{P}(s^{(t)} = i) = \pi_i$ for all $i \in S$. Define the Markov process $(r^{(t)})$ via the transition probabilities $\mathbf{Q}(r^{(t+1)} = j \mid r^{(t)} = i) = g_{ij} := \frac{\pi_j}{\pi_i} \cdot f_{ji}$ and let $\mathbf{Q} \in \Delta(S^{\infty})$ denote the induced measure over paths. The Markov process $(r^{(t)})$ is the *time-reversed* version of $(s^{(t)})$. (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, π , 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 $T_d^R := \inf \{t \in \mathbb{N} : r^{(t)} = d\}$. For each $i \in S$, define the function $H_i : \mathbb{N} \rightarrow [0, 1]$ by $H_i(k) := \mathbf{Q}(T_d^R \leq k \mid r^{(0)} = i)$. Thus, $H_i(\cdot)$ is the CDF of T_d^R given that the time-reversed chain starts in state $i \in S$.

Claim C.4. Let $\mathbf{P}$ and $\mathbf{Q}$ be as defined above. Then, for every $i \in S$:

- (a) $\lim_{k \rightarrow \infty} H_i(k) = 1$.
- (b) For all $k \in \mathbb{N}$, $\lim_{t \rightarrow \infty} \mathbf{P}(s^{(t-m)} \neq d \forall m = 0, \dots, k \mid s^{(t)} = i) = 1 - H_i(k)$

Proof of Claim C.4. It is clear that the time-reversed process is fully connected, and thus each state is recurrent, i.e., $\mathbf{Q}(T_d^R < \infty \mid r^{(0)} = i) = 1$ for all $i \in S$. It follows from the Bounded Convergence Theorem that $\lim_{k \rightarrow \infty} H_i(k) = 1$ 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 $\mathbf{P}(s^{(T)} = i) > 0$ for all $T \geq t$. (This is possible because $\pi_i > 0$.) Then,

$$\mathbf{P}(s^{(t-m)} \neq d \forall m = 0, \dots, k \mid s^{(t)} = i) = \frac{\mathbf{P}(s^{(t-m)} \neq d \forall m = 0, \dots, k \text{ and } s^{(t)} = i)}{\mathbf{P}(s^{(t)} = i)}$$

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

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

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 $\lim_{t \rightarrow \infty} \mathbf{P}(s^{(t-k)} = j_k) = \pi_{j_k}$ and $\lim_{t \rightarrow \infty} \mathbf{P}(s^{(t)} = i) = \pi_i$ completes the proof of the claim. $\square$

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 $\lim_{k \rightarrow \infty} H(k) = 1$ by part (a) of Claim C.4. Notice that we may write $\mathbf{P}(B_{k,t}) = \sum_{i=1}^d \mathbf{P}(s^{(t)} = i) \cdot \mathbf{P}(s^{(t-m)} \neq d \forall m = 0, \dots, k \mid s^{(t)} = i)$. By part (b) of Claim C.4 and ergodicity of the type process, it follows that $\lim_{t \rightarrow \infty} \mathbf{P}(B_{k,t}) = 1 - H(k)$, as desired. $\square$

Lemma C.20. Suppose the type process is pseudo-renewal, with $\pi \in \Delta(S)$ 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. $\square$

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

Proof of Lemma C.21. Fix some path $h = (s^t)_{t=0}^\infty \in \mathcal{J}$ along which the marginal cost martingale converges to zero. By part (a) of Theorem 4.1, 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

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

and that, at the optimum, the directional derivative $D_1 P(\mathbf{w}_i(\mathbf{v}, s), i)$ satisfies

$$[\text{C.32}] \quad f_{si} \cdot D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) = \lambda_i(\mathbf{v}, s) + \sum_{k=1}^{i-1} \mu_{ik}(\mathbf{v}, s) - \sum_{k=i+1}^d \mu_{ki}(\mathbf{v}, s)$$

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

$$[\text{C.33}] \quad f_{si} \left(\frac{P_j(\mathbf{w}_i(\mathbf{v}, s), i)}{f_{ij}} - D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) \right) = \sum_{k=i+1}^d \left(1 - \frac{f_{kj}}{f_{ij}} \right) \mu_{ki}(\mathbf{v}, s)$$

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

$$[\text{C.34}] \quad f_{si} \left(\frac{P_j(\mathbf{w}_i(\mathbf{v}, s), i)}{f_{ij}} - D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) \right) = \begin{cases} 0, & \text{for } j \leq i \\ \left(1 - \frac{f_{jj}}{\pi_j} \right) \mu_{ji}(\mathbf{v}, s), & \text{for } j > i \end{cases}$$

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 $\mu(\cdot) \geq 0$ on $D \times S$ and since the type process is fully connected (Assumption Markov), it follows that

$$[\text{C.35}] \quad P_j(\mathbf{w}_i(\mathbf{v}, s), i) \leq f_{ij} D_1 P(\mathbf{w}_i(\mathbf{v}, s), i)$$

for all $i, j \in S$ (with equality when $j \leq i$). Now, part (a) of Theorem 4.1 states that $D_1 P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ almost surely. It follows from this, the martingale property (Proposition 4.3), and non-negativity of the directional derivative $D_1 P(\cdot, \cdot)$ on $D \times S$ (Lemma C.1) that $D_1 P(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow 0$ almost surely. Thus, [C.35] implies that

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

Since $\sum_{j=1}^d P_j(\mathbf{w}_i(\mathbf{v}, s), i) = D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) \geq 0$ this implies that $DP(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow 0$ $\mathbf{P}$ -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

$$[\text{C.36}] \quad P_j(\mathbf{w}_i(\mathbf{v}, s), i) \geq f_{ij} D_1 P(\mathbf{w}_i(\mathbf{v}, s), i)$$

which, by the same argument, implies that $\mathbf{P} \left(\liminf_{t \rightarrow \infty} P_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \geq 0 \forall i, j \in S \right) = 1$. Because $\sum_{j=1}^d P_j(\mathbf{w}_i(\mathbf{v}, s), i) = D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) \rightarrow 0$, this implies that $DP(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow 0$ $\mathbf{P}$ -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 $h \in \mathcal{H}$, $\mathbf{v}^{(t)}(h) = \mathbf{w}_i(\mathbf{v}^{(t-1)}(h), s^{(t-1)}(h))$ for some $i \in S$ — it must be that $DP(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ almost surely. Replicating the proof of Lemma C.17 with “ τ_t ” replaced everywhere by “ t ” establishes that $\nu^{(t)} \rightarrow 0$ almost surely, as desired. $\square$

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

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

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

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

Now, $B(\cdot, \cdot) \geq 0$ on $D \times S$ and the hypothesis of the lemma implies that $A(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}) \rightarrow 0$ almost surely (in probability), so it follows that the process $f_{s^{(n^{(t)})}, i} C'(u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}), i) \rightarrow 0$ almost surely (in probability). Because the type process is fully connected (Assumption **Markov**), it follows that $C'(u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}), i) \rightarrow 0$ almost surely (in probability). Finally, $C'(\cdot, i) : \mathcal{U} \rightarrow \mathbb{R}_{++}$ 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 $u_i(\mathbf{v}^{(n^{(t)})}, s^{(n^{(t)})}) \rightarrow -\infty$ almost surely (in probability). $\square$

Lemma C.23. Consider any feasible recursive contract such that the induced process $u_d(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow -\infty$ almost surely. Then the induced process $v_i^{(t)} \rightarrow -\infty$ 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 $[\text{PK}_i]$ one step ahead delivers

$$\begin{aligned} v_i^{(t)} &= u_i(\mathbf{v}^{(t)}, s^{(t)}) + \alpha \sum_{j=1}^d f_{ij} (u_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) + \alpha \mathbb{E}^{f_j} [\mathbf{w}_j(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i)]) \\ &\leq \alpha f_{id} \cdot u_d(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \end{aligned}$$

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

Wrapping Up. We can now consolidate the preceding lemmas into the proof of Theorem 4.1.

Proof of Theorem 4.1. 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. $\square$

D. Proof of Theorem 4.2

We begin by proving part (d) of Theorem B.1 (from Appendix B.1), which is then invoked in the main proof of Theorem 4.2.

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

Recall the set D^* defined in [B.1] and suppose that the type process satisfies FOSD. Consider the IC constraint [IC $_{ij}^*$]. Let $F^*(i | j) := \sum_{k \geq i} f_{kj}$ for each j . By FOSD, it follows that $F^*(\cdot | j) \leq F^*(\cdot | k)$ whenever $j < k$. It is easy to see that we can rewrite [IC $_{ij}^*$] (for $i > j$) as $v_i - v_j \geq [U(\omega_i + c_j) - U(\omega_j + c_j)] + \alpha \sum_{k=2}^d [F^*(k | i) - F^*(k | j)](w_{j,k} - w_{j,k-1})$. Notice that $U(\omega_i + c_j) > U(\omega_j + c_j)$ because $\omega_i > \omega_j$. More generally, we have, $w_{j,k}^{(n)} - w_{j,k-1}^{(n)} \geq [U(\omega_k + c_{k-1}) - U(\omega_{k-1} + c_{k-1})] + \alpha \sum_{\ell=2}^d [F^*(\ell | k) - F^*(\ell | k-1)](w_{k-1,\ell}^{(n+1)} - w_{k-1,\ell}^{(n)})$. Iterating the inequalities above, for any sequence of types, we have $v_i - v_j \geq (\text{strictly positive terms}) + \alpha^n \chi_n (w_{\ell,k}^{(n)} - w_{\ell,k-1}^{(n)})$ 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 $\lim_{n \rightarrow \infty} \alpha^n \chi_n (w_{j,k}^{(n)} - w_{j,k-1}^{(n)}) \rightarrow 0$. But this implies $v_i > v_j$ whenever $i > j$, i.e., $D^* \subset V$.

D.2. Main Proof of Theorem 4.2

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

$$\begin{aligned} v_{j+1} - v_j &\geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) + \alpha (\mathbf{E}^{f_{j+1}} [\mathbf{w}_j] - \mathbf{E}^{f_j} [\mathbf{w}_j]) \\ \text{[D.1]} \quad &\geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) \end{aligned}$$

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

By part (b) of Theorem 4.1, $u_j^{(t)} \rightarrow -\infty$ in probability, and by definition, $U(\omega_j + C(u_j^{(t)}, j)) = u_j^{(t)}$. By the Continuous Mapping Theorem, this implies that $\omega_j + c_j^{(t)} \rightarrow \underline{c}$

in probability, and furthermore, that $\lim_{t \rightarrow \infty} [\omega_{j+1} + c_j^{(t)}] = \underline{c} + (\omega_{j+1} - \omega_j)$ 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, $U(\omega_{j+1} + C(u_j^{(t)}, j)) - U(\omega_j + C(u_j^{(t)}, j)) \rightarrow \infty$ in probability because $U(\underline{c} + (\omega_{j+1} - \omega_j)) > -\infty$.

Case 2: Suppose now that $\underline{c} = -\infty$. Observe that $U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) = \int_{\omega_j + C(u_j, j)}^{\omega_{j+1} + C(u_j, j)} U'(y) dy \geq (\omega_{j+1} - \omega_j) U'(\omega_{j+1} + C(u_j, j))$ where the inequality follows from the concavity of U . But because $\underline{c} = -\infty$, we have $\omega_{j+1} + c_j^{(t)} \rightarrow -\infty$ 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 $U'(\omega_{j+1} + C(u_j^{(t)}, j)) \rightarrow +\infty$ 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 $(\mathbf{v}^{(t)}, s^{(t)})_{t=0}^\infty$. $\square$

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 [4.1](#), a sufficient condition for $u_j^{(t)} \rightarrow -\infty$ (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. $\square$

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 $\Gamma^{\text{FB}}(\mathbf{v}) := \{(u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times \mathcal{U}^d)^d : [\mathbf{PK}_i] \text{ holds for all } i \in S\}$. The principal’s first-best value function is then

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

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

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

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

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

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

The value function $Q^*(\cdot, s)$ is strictly convex for each $s \in S$, the derivative is strictly positive $DQ^*(\mathbf{v}, s) \gg \mathbf{0}$, and thus $D_1 Q^*(\mathbf{v}, s) > 0$ for all $(\mathbf{v}, s) \in \mathcal{U}^d \times S$.

Consider also the *full-information efficiency problem*

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

for each $i \in S$. This is the full-information analogue of the efficiency problem [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 [Eff_i^{FB}] has the unique solution

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

and its value function is given by

$$[\text{E.5}] \quad K^*(v, s) = \frac{U^{-1}((1 - \alpha)v)}{1 - \alpha} - \mathbf{E} \left[\sum_{t=0}^{\infty} \alpha^t \omega^{(t)} \mid s^{(0)} = s \right]$$

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

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 4.1 (see OA-C.4). Let $\mathcal{X}$ be the (countable) state space for a Markov process with transition probabilities $P(x, B)$ denoting the probability of transitioning from x to $B \subseteq \mathcal{X}$. Let $\mathbf{P}$ denote the induced probability measure on the path space $\mathcal{X}^\infty$.

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

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

Proposition F.2. The event $\{h \in \mathcal{H} : (s^{t-1}, s^t) = (i, j) \text{ for infinitely many } t\}$ 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 $Q((i, j), (k, \ell)) = \mathbf{1}\{j = k\}P(k, \ell)$, where the indicator $\mathbf{1}\{j = k\} = 1$ if $j = k$ and 0 otherwise. Then, the two-step transition probabilities are given by $Q^{(2)}((i, j), (k, \ell)) = P(j, k)P(k, \ell) > 0$. 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. $\square$

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

Supplementary Appendix

G. Proof of Theorem B.1

This SA, which is divided into several parts, provides the proof of Theorem B.1. SA-G.1 introduces the operator $\mathcal{B}_\circ$, of which D is the largest fixed point, and establishes basic existence and characterization results concerning D , D^* , and $\Gamma(\cdot)$. The heart of the proof is contained in the various parts of SA-G.2, which is focused mainly on two auxiliary operators, $\mathcal{B}$ and $\mathcal{B}^\dagger$. First, SA-G.2.1 establishes several results relating fixed points of the three operators to each other and to the set V_d . Second, SA-G.2.2 and SA-G.2.3 focus on the special case of CARA utility. SA-G.2.2 characterizes iterates of $\mathcal{B}^\dagger$ via solutions to an appropriate linear program and its dual, and culminates in closed-form expressions for these iterates when the type process satisfies MLRP or PPR. SA-G.2.3 characterizes iterates of $\mathcal{B}$ via solutions to an appropriate conic program and its dual (which are not characterized in closed form). Third, SA-G.2.4 returns to the case of general DARA utility functions and, using a Lagrangian argument, shows that this actually reduces to the CARA case studied in SA-G.2.3. Next, SA-G.3 constructs a useful suboptimal contract that is [TVC]-implementable and stays bounded away from the boundaries of D . Finally, SA-G.4 collects these various pieces into a proof of Theorem B.1.

G.1. Preliminary Observations

We begin with an elementary lemma concerning the $\psi(\cdot, i, j)$ functions.

Lemma G.1. Let $i, j \in S$ with $i > j$. Under Assumption DARA:

- (a) $\psi(u, i, j)$ is convex, $\psi'(u, i, j) \in (0, 1)$, and $\psi(u, i, j) > u$ for all $u \in \mathcal{U}$.
- (b) The extension of $\psi(\cdot, i, j)$ to $\mathbb{R}_-$ given by $\psi(0, i, j) := 0$ and $\psi'(0, i, j) := 1$ preserves continuity, convexity, and continuous differentiability.
- (c) $a\psi(u, i, j) \geq \psi(au, i, j)$ for all $u \in \mathcal{U}$ and $a \in [0, 1]$.
- (d) Under CARA utility, $\psi(u, i, j) = (\theta_i/\theta_j) \cdot u$ where $\theta_k := e^{-\omega_k}$ for all $k \in S$. Thus, $a\psi(u, i, j) = \psi(au, i, j)$ for all $u \in \mathcal{U}$ and $a \in \mathbb{R}_+$.

Proof. By the Inverse Function Theorem, we get

$$[\text{G.1}] \quad \psi'(u, i, j) = \frac{U'(\omega_i + C(u, j))}{U'(\omega_j + C(u, j))} > 0$$

It follows immediately from part (a) of Assumption DARA that $\psi(u, i, j) - u > 0$ and $\psi'(u, i, j) \in (0, 1)$. Moreover, we have

$$\log(\psi'(u, i, j)) = \log(U'(\omega_i + C(u, j))) - \log(U'(\omega_j + C(u, j)))$$

and because $C(\cdot, j)$ is strictly increasing, by part 3 of Assumption [DARA](#), $\log(\psi'(u, i, j))$ is increasing. It follows that $\psi'(u, i, j)$ is increasing, ie, $\psi(\cdot, i, j)$ is convex.

For part (b), note that, since the consumption domain is assumed unbounded above, $\lim_{u \rightarrow 0} \psi(u, i, j) = \lim_{u \rightarrow 0} U(\omega_i + C(u, j)) = \lim_{c \rightarrow \infty} U(\omega_i + c) = 0$. Thus this extension preserves continuity and convexity of $\psi(\cdot, i, j)$. Moreover, taking the limit $u \rightarrow 0$ in [\[G.1\]](#) delivers $\psi'(0, i, j) = 1$.

For part (c), let $u \in \mathcal{U}$ and $a \in [0, 1]$. We have $\psi(au, i, j) \leq a\psi(u, i, j) + (1 - a)\psi(0, i, j) = a\psi(u, i, j)$, where the inequality is from convexity and the equality is from $\psi(0, i, j) = 0$, shown above.

Part (d) is immediate from the definitions. $\square$

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

$$[\mathbf{PK}_j] \quad u_j + \alpha \mathbf{E}^{f_j} [\mathbf{w}_j] = v_j$$

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

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

for all $i, j \in S$ with $i > j$.

For any $\mathbf{v} \in \mathcal{U}^d$, we say that the menu $(\mathbf{u}, \mathbf{w}) := (u_i, \mathbf{w}_i)_{i \in S} \in (\mathbb{R}_- \times \mathbb{R}_-^d)^d$ (where $\mathbf{u} := (u_1, \dots, u_d)$ and $\mathbf{w} := (\mathbf{w}_1, \dots, \mathbf{w}_d)$) *implements* $\mathbf{v}$ if it satisfies all of the recursive constraints at $\mathbf{v}$. (Note that implementation does not require even that $u_i < 0$ or $\mathbf{w}_i \ll \mathbf{0}$. Such feasibility requirements will be included as separate constraints.) Consider — in the manner of Abreu, Pearce and Stacchetti ([1990](#)) — the *operator* $\mathcal{B}_\circ : 2^{\mathcal{U}^d} \rightarrow 2^{\mathcal{U}^d}$, defined as follows for any $K \subset \mathcal{U}^d$:

[\[G.2\]](#)

$$\mathcal{B}_\circ(K) := \{\mathbf{v} \in \mathcal{U}^d : \exists (\mathbf{u}, \mathbf{w}) \text{ that implements } \mathbf{v} \text{ with } (u_i, \mathbf{w}_i) \in \mathcal{U} \times K \text{ for all } i \in S\}$$

Recalling Definition [3.1](#), it is easy to see that any fixed point of $\mathcal{B}_\circ$ is a recursive domain, and that the *largest recursive domain* D is precisely the *largest* fixed point of the operator $\mathcal{B}_\circ$.

Define the correspondence $\Gamma : \mathcal{U}^d \times \mathcal{K}_d \rightrightarrows (\mathcal{U} \times \mathcal{U}^d)^d$ by

$$[\mathbf{G.3}] \quad \Gamma(\mathbf{v}, K) := \left\{ (\mathbf{u}, \mathbf{w}) \in (\mathcal{U} \times K)^d : (\mathbf{u}, \mathbf{w}) \text{ implements } \mathbf{v} \right\}$$

Thus $\Gamma(\mathbf{v})$, defined in [\[3.1\]](#) in the main text, is equal to $\Gamma(\mathbf{v}, D)$. Clearly, an alternative way of defining the operator in [\[G.2\]](#) is $\mathcal{B}_\circ(K) = \{\mathbf{v} \in \mathcal{U}^d : \Gamma(\mathbf{v}, K) \neq \emptyset\}$.

Say that set $K \subseteq \mathbb{R}^d$ has *decreasing returns* if $\mathbf{v} \in K$ implies that $a\mathbf{v} \in K$ for all $a \in (0, 1]$. Say that a set $K \subseteq \mathbb{R}^d$ is *unbounded below* if, for all $r < 0$, there exists $\mathbf{v} \in K$ such that $\mathbf{v} \leq r\mathbf{1}$.

Lemma G.2. Let $K \subseteq \mathcal{U}^d$. Then:

- (a) If K is convex, then $\Gamma(\cdot, K)$ has convex graph.
- (b) If K has decreasing returns, then the graph of $\Gamma(\cdot, K)$ has decreasing returns.
- (c) If K is a cone and utility is CARA, then the graph of $\Gamma(\cdot, K)$ is a cone.

Proof. For part (a), suppose that K is convex and let $\gamma \in \Gamma(\mathbf{v}, K)$ and $\gamma' \in \Gamma(\mathbf{v}', K)$. Let $a \in (0, 1)$ and define $\mathbf{v}'' := a\mathbf{v} + (1 - a)\mathbf{v}' \in \mathcal{U}^d$ and $\gamma'' := a\gamma + (1 - a)\gamma' \in (\mathcal{U} \times K)^d$. Each of the $[\mathbf{PK}_j]$ constraints is linear, so clearly γ'' satisfies all of them at $\mathbf{v}''$. Consider constraint $[\mathbf{IC}_{ij}]$, with $i > j$. By definition of γ and γ' , we have

$$v_i'' \geq a\psi(u_j, i, j) + (1 - a)\psi(u'_j, i, j) + \alpha \mathbf{E}^{\mathbf{f}_i} [a\mathbf{w}_j + (1 - a)\mathbf{w}'_j]$$

and, since $\psi(\cdot, i, j)$ is convex by part (a) of Lemma G.1, $a\psi(u_j, i, j) + (1 - a)\psi(u'_j, i, j) \geq \psi(u''_j, i, j)$. Thus, γ'' satisfies $[\mathbf{IC}_{ij}]$ at $\mathbf{v}''$. Hence, $\gamma'' \in \Gamma(\mathbf{v}'', K)$, ie, $\Gamma(\cdot, K)$ has convex graph when K is convex.

For part (b), suppose that K has decreasing returns and let $\gamma \in \Gamma(\mathbf{v}, K)$. Let $a \in (0, 1)$ and define $\mathbf{v}' := a\mathbf{v} \in \mathcal{U}^d$ and $\gamma' := a\gamma \in (\mathcal{U} \times K)^d$. Each of the $[\mathbf{PK}_j]$ constraints is linear, so clearly γ' satisfies all of them at $\mathbf{v}'$. Consider constraint $[\mathbf{IC}_{ij}]$, with $i > j$. By definition of γ , we have

$$v'_i \geq a\psi(u_j, i, j) + \alpha \mathbf{E}^{\mathbf{f}_i} [a\mathbf{w}_j]$$

and, by part (c) of Lemma G.1, $a\psi(u_j, i, j) \geq \psi(au_j, i, j)$. Thus, γ' satisfies $[\mathbf{IC}_{ij}]$ at $\mathbf{v}'$. Hence, $\gamma' \in \Gamma(\mathbf{v}', K)$, ie, $\Gamma(\cdot, K)$ has graph with decreasing returns when K has decreasing returns.

For part (c), mimic the proof of part (b) above, invoking part (d) of Lemma G.1 instead of part (c) of that lemma. $\square$

Proposition G.3. The following hold:

- (a) There exists a largest recursive domain D . It is nonempty, convex, has decreasing returns, and is unbounded below.
- (b) $D^* \subseteq D$ has all of the above properties.
- (c) The constraint correspondence $\Gamma : D \rightarrow (\mathcal{U} \times D)^d$ is nonempty-valued, has convex graph, and is continuous.

Proof. Begin with part (a). It is easy to see that the operator $\mathcal{B}_\circ$ is monotone increasing with respect to set inclusion. Because $(2^{\mathcal{U}^d}, \supseteq)$ defines a complete lattice, $\mathcal{B}_\circ$ has a largest fixed point D by Tarski's Fixed Point Theorem.

To establish that D is nonempty and unbounded below, define

$$B := \{c \in \mathbb{R} : c + \omega_i \in (\underline{c}, \infty) \ \forall i \in S\}$$

Clearly $\inf B = \underline{c} - \omega_1$ and $\sup B = +\infty$. For any $c \in B$, let $\mathbf{v}^c \in \mathcal{U}^d$ denote the vector of contingent promised utilities induced by the stochastic consumption process $(\omega^{(t)} + c)_{t=0}^\infty$. (Thus, $\mathbf{v}^0$ is the promised utility vector under autarky.) Define $D_c := \{\mathbf{v}^c\}$ for each $c \in B$. It is easy to see that $\mathcal{B}_\circ(D_c) = D_c$, and thus $\cup_{c \in B} D_c \subseteq D$, for each $c \in B$. Since $\mathcal{U} \subseteq \mathbb{R}_{--}$, it follows that

$$v_i^c \leq U(\omega_1 + c) \cdot T_1$$

for all $i \in S$, where

$$T_1 := \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t \mathbf{1}(s^{(t+1)} = 1) \mid s^{(0)} = i \right]$$

and $T_1 > 0$ by Assumption [Markov](#). Moreover, $\lim_{c \rightarrow \inf B} U(\omega_1 + c) = -\infty$ under Assumption [DARA](#). Thus, for any $r \in \mathbb{R}_{--}$, there exists $c \in B$ such that $\max_{i \in S} v_i^c < r$. Thus, $\cup_{c \in B} D_c$ is unbounded below. Now, it follows from the definition of D that $D \supseteq \cup_{c \in B} D_c$. It follows that D is nonempty and unbounded below.

We now proceed to establish convexity and decreasing returns. Define $\mathcal{K}_d^{cdr} \subseteq 2^{\mathcal{U}^d}$ to be the collection of all subsets of $\mathcal{U}^d$ that are convex and have decreasing returns. It is easy to see that $\mathcal{K}_d^{cdr}$ is a complete lattice, with the meet operation defined by set intersection and the join operation defined by the natural notion of “convex, decreasing-returns hull,” namely

$$\bigvee_{\lambda \in \Lambda} K_\lambda := \bigcap \left\{ K \in \mathcal{K}_d^{cdr} : K \supseteq \bigcup_{\lambda \in \Lambda} K_\lambda \right\}$$

It follows from Lemma [G.2](#) that $\mathcal{B}_\circ : \mathcal{K}_d^{cdr} \rightarrow \mathcal{K}_d^{cdr}$ is well-defined and monotone increasing. Thus, by Tarski’s Fixed Point Theorem, $\mathcal{B}_\circ$ has a largest fixed point $D' \in \mathcal{K}_d^{cdr}$. But since $\mathcal{U}^d \in \mathcal{K}_d^{cdr}$, it is easy to see that $D \in \mathcal{K}_d^{cdr}$ and hence $D = D'$. (See, eg, the transfinite recursion in Lemma 1 of [Echenique \(2005\)](#).) This concludes the proof of part (a).

Now consider part (b). Define scalar multiplication and convex combination of recursive contracts pointwise. It follows from the definition that if ξ satisfies agent transversality at $\mathbf{v} \in D$, then for any $a > 0$, $a\xi$ satisfies agent transversality at $a\mathbf{v} \in D$. Similarly, if ξ and ξ' satisfy agent transversality at, respectively, $\mathbf{v} \in D$ and $\mathbf{v}' \in D$, then for any $a \in (0, 1)$, $a\xi + (1 - a)\xi'$ satisfies agent transversality at $a\mathbf{v} + (1 - a)\mathbf{v}'$. It follows from these observations and the fact that D is convex and has decreasing returns, which was established in part (a), that D^* inherits both of these properties. Moreover, the vectors $\mathbf{v}^c$

defined in the proof of part (a) are clearly elements of D^* , which establishes that D^* is nonempty and unbounded below.

Finally, consider part (c). This follows immediately from part (a) of the present proposition and part (a) of Lemma G.2, namely the facts that D is convex and $\Gamma(\cdot) = \Gamma(\cdot, D)$ as defined in [G.3]. Continuity is easy to see from the definitions. $\square$

G.2. Fixed Points of Operators

While we are ultimately interested in (largest) fixed points of $\mathcal{B}_\circ$, it is convenient to work with some auxiliary operators. Consider first the operator $\mathcal{B}$, defined as follows for any $K \subset \mathbb{R}_-^d$:

[G.4]

$$\mathcal{B}(K) := \{ \mathbf{v} \in \mathbb{R}_-^d : \exists (\mathbf{u}, \mathbf{w}) \text{ that implements } \mathbf{v} \text{ with } (u_i, \mathbf{w}_i) \in \mathbb{R}_- \times K \text{ for all } i \in S \}$$

It is easy to see that $\mathcal{B}_\circ(K) \subset \mathcal{B}(K)$ for all K , so that the largest fixed point of $\mathcal{B}_\circ$ is a subset of the largest fixed point of $\mathcal{B}$. To compute the largest fixed point of $\mathcal{B}$, consider in turn the family of operators defined as follows for any $K \subset \mathcal{U}^d$:

$$[\text{G.5}] \quad \mathcal{B}_{ij}(K) := \{ \mathbf{v} \in \mathbb{R}_-^d : \exists (u_j, \mathbf{w}_j) \in \mathbb{R}_- \times K \text{ s.t. } [\text{PK}_j] \text{ and } [\text{IC}_{ij}] \text{ hold at } \mathbf{v} \}$$

Notice that for any set K , we have

$$\mathcal{B}(K) = \bigcap_{i,j \in S: i > j} \mathcal{B}_{ij}(K)$$

We may also use the $\mathcal{B}_{ij}$ to define a further auxiliary operator $\mathcal{B}^\dagger : 2^{\mathbb{R}_-^d} \rightarrow 2^{\mathbb{R}_-^d}$ as

$$\mathcal{B}^\dagger(K) = \bigcap_{j=1}^{d-1} \mathcal{B}_{j+1,j}(K)$$

where $\mathcal{B}_{j+1,j}$ is defined in [G.5]. It is easy to see that $\mathcal{B}(K) \subset \mathcal{B}^\dagger(K)$ for all K , so that the largest fixed point of $\mathcal{B}$ is a subset of the largest fixed point of $\mathcal{B}^\dagger$.

We note the following property for future reference. Let $\mathcal{K} \subseteq 2^{\mathbb{R}_-^d}$ denote the collection of subsets $K \subseteq \mathbb{R}_-^d$ that are closed convex cones.

Lemma G.4. Suppose utility is CARA. The maps $\mathcal{B} : \mathcal{K} \rightarrow \mathcal{K}$ and $\mathcal{B}^\dagger : \mathcal{K} \rightarrow \mathcal{K}$ are well-defined.

Proof. Immediate from part (c) of Lemma G.2 and the definition of the constraint correspondence $\Gamma(\cdot)$ (see [3.1]). $\square$

G.2.1. Some Fixed Points

Let $V_d := \{\mathbf{v} \in \mathbb{R}_-^d : v_1 < v_2 < \dots < v_d\}$ and let $\text{cl } V_d$ denote its closure; thus, $\text{cl } V_d = \{\mathbf{v} \in \mathbb{R}_-^d : v_1 \leq v_2 \leq \dots \leq v_d\}$. We claim that, if the type process satisfies FOSD, $\mathcal{B}(\text{cl } V_d) = \text{cl } V_d$, $\mathcal{B}_o(\text{cl } V_d) = V_d$, and that $\mathcal{B}_o(V_d) = V_d$. We establish these facts below.

Lemma G.5. Consider $[\mathbf{PK}_j]$ and $[\mathbf{IC}_{ij}]$ where $i > j$ and let $\mathbf{v} \in \text{cl } V_d$.

- (a) Suppose the type process satisfies FOSD. Suppose, in addition, that $v_i = v_j$. Any solution to $[\mathbf{PK}_j]$ and $[\mathbf{IC}_{ij}]$ with $\mathbf{w}_j \in \text{cl } V_d$ must satisfy $u_j = 0$ and $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] = \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]$.
- (b) In general, if $v_{j+1} > v_j$, there exist $u_j < 0$ and $\mathbf{w}_j \in \text{cl } V_d$ with $\mathbf{w}_j \ll \mathbf{0}$ such that $[\mathbf{PK}_j]$ holds and $[\mathbf{IC}_{ij}]$ holds as a strict inequality for all $i > j$. In particular, there exists such a solution in which $\mathbf{w}_j = w^j \mathbf{1}$ for some $w^j < 0$.

Proof. To see (a), let $(u_j, \mathbf{w}_j)$ satisfy both $[\mathbf{PK}_j]$ and $[\mathbf{IC}_{ij}]$. Subtract $[\mathbf{PK}_j]$ from each side of $[\mathbf{IC}_{ij}]$ to obtain

$$0 \leq \psi(u_j, i, j) - u_j + \alpha(\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] - \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]) \leq v_i - v_j = 0$$

Recall from Lemma G.1 that $\psi(u, i, j) > u$ for all $u \in \mathcal{U}$ and $\psi(0, i, j) = 0$. Moreover, $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] \geq \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]$ because $\mathbf{w}_j \in \text{cl } V_d$ and $\mathbf{f}_i$ first-order stochastically dominates $\mathbf{f}_j$. Thus, the above display implies that $u_j = 0$ and $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] = \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]$.

To see (b), notice that setting $u_j = 0$ and $\mathbf{w}_j = w^j \mathbf{1}$ with $w^j = v_j/\alpha$ means that $[\mathbf{PK}_j]$ holds as an equality, while $[\mathbf{IC}_{ij}]$ holds as a *strict* inequality for all $i > j$ (because $v_{j+1} > v_j$ and $\mathbf{v} \in \text{cl } V_d$). Now, increase w_j by a small amount while reducing u_j simultaneously to ensure that $[\mathbf{PK}_j]$ still holds. Because all the constraints $[\mathbf{IC}_{ij}]$ held as strict inequalities, this perturbation also satisfies these constraints as strict inequalities. $\square$

Lemma G.6. Fix $w^j < 0$. Then, for all $\varepsilon > 0$, there exists $\mathbf{w}_j \in V_d$ such that (i) $\mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j] = w^j$, and (ii) $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] - \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j] < \varepsilon$ for all $i > j$. If the type process satisfies FOSD, there exists $\mathbf{w}_j \in V_d$ that satisfies (i), (ii), and (iii) $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] \geq \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]$.

Proof. Fix $\tilde{\varepsilon} > 0$ as a parameter to be optimised later. Let $w_{jk} := w^j - \tilde{\varepsilon}/2^k$ for $k = 1, \dots, d-1$, and $w_{jd} := w^j + \delta$, where $\delta := \frac{\tilde{\varepsilon}}{f_{jd}} \sum_{k=1}^{d-1} f_{jk}/2^k$. It is easy to see that $\mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j] = w^j$, establishing (i).

To see (ii), note that $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] - \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j] \leq w_{jd} - w_{j1} = \delta + \tilde{\varepsilon} = \frac{\tilde{\varepsilon}}{f_{jd}} \left[\sum_{k=1}^{d-1} f_{jk}/2^k + f_{jd} \right] < \tilde{\varepsilon}/f_{jd}$. So by choosing $\tilde{\varepsilon} = \varepsilon \cdot f_{jd}$, we establish (ii).

If the type process satisfies FOSD, it follows immediately that $\mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_j] \geq \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j]$ because $w_{j1} < \dots < w_{jd}$. $\square$

Lemma G.7. The set V_d satisfies:

- (a) $\mathcal{B}(\text{cl } V_d) \supseteq \text{cl } V_d$.
- (b) $\mathcal{B}^\dagger(\text{cl } V_d) \supseteq \text{cl } V_d$.
- (c) $\mathcal{B}_\circ(V_d) \supseteq V_d$.

Proof. First, consider part (a). Let $\mathbf{v} \in \text{cl } V_d$. By part (b) of Lemma G.5, for each $j = 1, \dots, n-1$, there exist $(u_j, \mathbf{w}_j) \in \mathbb{R}_- \times \text{cl } V_d$ such that $[\mathbf{PK}_j]$ holds, and $[\mathbf{IC}_{ij}]$ holds for all $i > j$. Thus, $\text{cl } V_d \subseteq \mathcal{B}(\text{cl } V_d)$, which proves part (a). Part (b) follows immediately from part (a), as $\mathcal{B}^\dagger(K) \supseteq \mathcal{B}(K)$ for all $K \subseteq \mathbb{R}_-^d$.

Finally, consider part (c). Let $\mathbf{v} \in V_d$. Then, $v_j < v_{j+1}$ for all $j = 1, \dots, d-1$. By part (b) of Lemma G.5, there exist $u_j < 0$ and $\mathbf{w}_j \in \text{cl } V_d$ with $\mathbf{w}_j \ll \mathbf{0}$ such that $[\mathbf{PK}_j]$ holds and $[\mathbf{IC}_{ij}]$ holds as a strict inequality for all $i > j$. By Lemma G.6, we can replace $\mathbf{w}_j$ by some $\mathbf{w}'_j \in V_d$ so that $[\mathbf{PK}_j]$ still holds, and $[\mathbf{IC}_{ij}]$ holds as a strict inequality for all $i > j$. Because $\mathbf{w}'_j \in V_d$, it follows that $\mathcal{B}_\circ(V_d) \supseteq V_d$, as claimed. $\square$

Proposition G.8. The following hold:

- (a) There exists a largest fixed point K_* of $\mathcal{B}$. It satisfies $K_* \supseteq \text{cl } V_d$.
- (b) There exists a largest fixed point $K_*^\dagger$ of $\mathcal{B}^\dagger$. It satisfies $K_*^\dagger \supseteq K_* \supseteq \text{cl } V_d$.
- (c) The largest fixed point D of $\mathcal{B}_\circ$ satisfies $D \supseteq V_d$.

Proof. All three operators $\mathcal{B}$, $\mathcal{B}^\dagger$, and $\mathcal{B}_\circ$ are monotone increasing. Thus, by Tarski's Fixed Point Theorem, $\mathcal{B}$ and $\mathcal{B}^\dagger$ each has a largest fixed point in $\mathbb{R}_-^d$, K_* and $K_*^\dagger$, respectively. Existence of D , the largest fixed point of $\mathcal{B}_\circ$ in $\mathcal{U}^d$, was established in part (a) of Proposition G.3 by an identical argument.

Define $\Phi := \{K \subseteq \mathbb{R}_-^d : K \supseteq \text{cl } V_d\}$ and $\Phi_\circ := \{K \subseteq \mathcal{U}^d : K \supseteq V_d\}$. Both Φ and $\Phi_\circ$ define complete lattices when equipped with the set inclusion partial order. It follows from monotonicity and Lemma G.7 that the mappings $\mathcal{B} : \Phi \rightarrow \Phi$, $\mathcal{B}^\dagger : \Phi \rightarrow \Phi$, and $\mathcal{B}_\circ : \Phi_\circ \rightarrow \Phi_\circ$ are well-defined. Thus, by Tarski's Fixed Point Theorem, $\mathcal{B}$ and $\mathcal{B}^\dagger$ each has a largest fixed point in Φ and $\mathcal{B}_\circ$ has a largest fixed point in $\Phi_\circ$. It follows immediately that $K_*, K_*^\dagger \in \Phi$ and $D \in \Phi_\circ$. $\square$

Proposition G.9. Suppose that the type process satisfies FOSD. Then:

- (a) $\mathcal{B}(\text{cl } V_d) = \text{cl } V_d$.
- (b) $\mathcal{B}^\dagger(\text{cl } V_d) = \text{cl } V_d$.
- (c) $\mathcal{B}_\circ(\text{cl } V_d) \subseteq V_d \subsetneq \text{cl } V_d$.

Proof. First, consider part (a). By part (a) of Lemma G.7, it suffices to show that $\mathcal{B}(\text{cl } V_d) \subset \text{cl } V_d$ under the additional FOSD assumption. To that end, let $\mathbf{v} \in \mathbb{R}_-^d$ and consider $i > j$. Then, subtracting $[\mathbf{PK}_j]$ from $[\mathbf{IC}_{ij}]$, we find that

$$0 \leq \psi(u_j, i, j) + \alpha(\mathbf{E}^{f_i}[\mathbf{w}_j] - \mathbf{E}^{f_j}[\mathbf{w}_j]) \leq v_i - v_j$$

which implies $v_i \geq v_j$, as (i) $\psi(u_j, i, j) - u_j \geq 0$, and (ii) $E^{f_i}[\mathbf{w}_j] \geq E^{f_j}[\mathbf{w}_j]$ because f_i first order stochastically dominates f_j and $\mathbf{w}_j \in \text{cl } V_d$. Thus, if $\mathbf{v} \in \mathcal{B}(\text{cl } V_d)$, we must have $v_i \geq v_j$ whenever $i > j$, ie, $\mathcal{B}(\text{cl } V_d) \subset \text{cl } V_d$.

Next, consider part (b). By part (b) of Lemma G.7, it suffices to show that $\mathcal{B}^\dagger(\text{cl } V_d) \subset \text{cl } V_d$ under the additional FOSD assumption. The proof of this fact is identical to the proof for part (a).

Finally, consider part (c). It is immediate that $\mathcal{B}_\circ(\text{cl } V_d) \subseteq \mathcal{B}(\text{cl } V_d) = \text{cl } V_d$ by monotonicity. Take any $\mathbf{v} \in \text{cl } V_d \setminus V_d$. It is easy to see that there exists some $j \in S$ such that either (i) $v_j = 0$ or (ii) $v_j = v_{j+1}$ or both. In case (i), [PK_j] implies that

$$0 = u_j + \alpha E^{f_j}[\mathbf{w}_j] \leq u_j$$

where the inequality follows from $\mathbf{w}_j \in \text{cl } V_d$. It follows that, in case (i), it must be that $u_j = 0$ and hence $\mathbf{v} \notin \mathcal{B}_\circ(\text{cl } V_d)$. In case (ii), by part (a) of Lemma G.5, if $(u_j, \mathbf{w}_j) \in \mathbb{R}_- \times \text{cl } V_d$ satisfy [IC_{ij}] and [PK_j], it must necessarily be the case that $u_j = 0$. Again, it follows that $\mathbf{v} \notin \mathcal{B}_\circ(\text{cl } V_d)$, which completes the proof of the claim. $\square$

Proposition G.10. If the type process satisfies FOSD, then V_d is a fixed point of $\mathcal{B}_\circ$.

Proof. Lemma G.7 states that $\mathcal{B}_\circ(V_d) \supseteq V_d$. Part (c) of Proposition G.9 states that $\mathcal{B}_\circ(\text{cl } V_d) \subseteq V_d$ and, since the operator $\mathcal{B}_\circ$ is monotone, it follows that $\mathcal{B}_\circ(V_d) \subseteq V_d$. This completes the proof. $\square$

G.2.2. Iterates of $\mathcal{B}^\dagger$ in the CARA Case

Recall (from, eg, part (c) of Lemma G.1) that under CARA utility, $\psi(u, i, j) = (\theta_i/\theta_j) \cdot u$ where $\theta_i := e^{-\omega_i}$. Thus, the incentive constraints [IC_{ij}] are linear in the u_j .

For any $\mathbf{t} = (t_2, \dots, t_d) \in \mathbb{R}_+^{d-1}$, let

$$[\text{G.6}] \quad K_{\mathbf{t}} := \{\mathbf{v} \in \mathbb{R}_-^d : v_{j-1}t_j \leq v_j, \text{ for } j = 2, \dots, d\}$$

It is easy to see that $K_{\mathbf{t}}$ is a closed, convex cone and is a subset of $\mathbb{R}_-^d$.

Proposition G.11. Suppose utility is CARA. There exists $\mathbf{t} = (t_2, \dots, t_d) \geq \mathbf{1}$ with $t_2 \cdot t_3 \cdots t_d < \infty$ such that $\mathcal{B}^\dagger(\mathbb{R}_-^d) = K_{\mathbf{t}}$.

Proof. Let us first consider $\mathcal{B}_{21}(\mathbb{R}_-^d)$. We want to characterize the set of $\mathbf{v} \in \mathbb{R}_-^d$ that can be implemented while satisfying

$$[\text{PK}_1] \quad u_1 + \alpha E^{f_1}[\mathbf{w}_1] = v_1$$

and

$$[\mathbf{IC}_{21}] \quad \frac{\theta_2}{\theta_1} \cdot u_1 + \alpha \mathbf{E}^{\mathbf{f}_2}[\mathbf{w}_1] \leq v_2$$

subject to

$$[\mathbf{Feas}_{21}] \quad u_1 \leq 0 \quad \text{and} \quad \mathbf{w}_1 \in \mathbb{R}_-^d$$

It is easy to see that any $\mathbf{v}$ with $v_2 \geq v_1$ can be implemented. (Indeed, this follows from part (b) of Lemma G.7.) For general $\mathbf{v}$, consider the following problem: What is the maximum value of $u_1 + \alpha \mathbf{E}^{\mathbf{f}_1}[\mathbf{w}_1]$ that can be achieved while satisfying $[\mathbf{IC}_{21}]$ and the feasibility constraints $[\mathbf{Feas}_{21}]$?

Our question can be formulated as the following linear programming problem

$$\begin{aligned} & \max_{u_1, \mathbf{w}_1} u_1 + \alpha \sum_{j=1}^d f_{1j} w_{1j} \\ & \text{s.t.} \quad [\mathbf{IC}_{21}] \text{ and } [\mathbf{Feas}_{21}] \end{aligned}$$

It is easy to see that the solution to this problem consists of setting $u_1 = 0$, letting k^* such that $f_{1k^*}/f_{2k^*} \leq f_{1j}/f_{2j}$ for all $j = 1, \dots, d$, and setting $w_{1j} = 0$ for $j \neq k^*$ while choosing w_{1k^*} so that $[\mathbf{IC}_{21}]$ holds with equality. Intuitively, $[\mathbf{IC}_{21}]$ is a budget constraint and our solution consists of choosing the smallest amount of the ‘quantity’ (u_1 and $\mathbf{w}_1$) where the marginal utility-to-price ratio is smallest, while choosing the largest quantity of everything else.

This implies that the largest value $u_1 + \alpha \sum_{j=1}^d f_{1j} w_{1j}$, and hence v_1 , can take is $v_2 \cdot f_{1k^*}/f_{2k^*}$. That is, $\mathcal{B}_{21}(\mathbb{R}_-^d) := \{\mathbf{v} \in \mathbb{R}_-^d : v_2 \geq t_2 v_1, t_2 := \max[f_{2j}/f_{1j}]\}$. Similarly, we can show that $\mathcal{B}_{j+1,j}(\mathbb{R}_-^d) := \{\mathbf{v} \in \mathbb{R}_-^d : v_{j+1} \geq t_{j+1} v_j, t_{j+1} := \max[f_{j+1,k}/f_{jk}]\}$. Clearly, $1 \leq t_j < \infty$ for every $j = 2, \dots, d$, which proves the claim. $\square$

We will use the following version of Linear Programming Duality and Complementary Slackness in the sequel.⁵³

Theorem G.1 (LP duality and Complementary Slackness). *Let $y^* = (y_1^*, \dots, y_n^*)$ be a feasible solution of the linear program*

$$[\mathbf{P}] \quad \max_y \quad c \cdot y \quad \text{subject to } Ay \leq b \text{ and } y \leq 0$$

(53) Section 4.3 of Vohra (2005) shows that $[\mathbf{P}]$ and $[\mathbf{D}]$ as written here are dual to each other. Alternatively, see Section 6.2 of Matoušek and Gärtner (2007). The version of complementary slackness described in Theorem G.1 is on p. 204 of the Glossary of Matoušek and Gärtner (2007).

and let $q^* = (q_1^*, \dots, q_m^*)$ be a feasible solution of the dual linear program

$$[\mathbf{D}] \quad \min_q \quad b \cdot q \quad \text{subject to } qA \leq c \text{ and } q \geq 0$$

Then, the following are equivalent:

- (a) y^* is optimal for $[\mathbf{P}]$ and q^* is optimal for $[\mathbf{D}]$.
- (b) For all $i = 1, \dots, m$, y^* satisfies the i th constraint of $[\mathbf{P}]$ with equality or $q_i^* = 0$; similarly, for all $j = 1, \dots, n$, q^* satisfies the j th constraint of $[\mathbf{D}]$ with equality or $y_j^* = 0$.

Consider the operator $\mathcal{B}_{j+1,j}(K_t)$ where $t \geq 1$ is finite in every component. We seek to characterize the set of all $\mathbf{v}$ that can be implemented by $(u_j, \mathbf{w}_j)_{j \in S} \in \mathbb{R}_+^d \times K_t$ while satisfying $[\mathbf{PK}_j]$ and $[\mathbf{IC}_{ij}]$ where $i = j + 1$. As in the proof of Proposition G.11, we can state this problem as follows:

$$\begin{aligned}
 [\mathbf{P}_{j+1,j}] \quad & \max_{u_j, \mathbf{w}_j} \quad u_j + \alpha[f_{j1}w_{j1} + f_{j2}w_{j2} + \dots + f_{jd}w_{jd}] \\
 \text{s.t.} \quad & \frac{\theta_{j+1}}{\theta_j} \cdot u_j + \alpha[f_{j+1,1}w_{j,1} + \dots + f_{j+1,d}w_{j,d}] \leq v_{j+1} \\
 & \alpha t_2 w_{j1} - \alpha w_{j2} \leq 0 \\
 & \alpha t_3 w_{j2} - \alpha w_{j3} \leq 0 \\
 & \vdots \\
 & \alpha t_d w_{j,d-1} - \alpha w_{jd} \leq 0 \\
 & x_1, w_{jd} \leq 0
 \end{aligned}$$

It is easy to see that $[\mathbf{P}_{j+1,j}]$ has a solution. Let $\mathbf{q}^{j+1}$ be a feasible solution of the dual linear program

$$\begin{aligned}
 [\mathbf{D}_{j+1,j}] \quad & \min_{\mathbf{q}^{j+1}} \quad q_1^{j+1} \cdot v_{j+1} \\
 \text{s.t.} \quad & \theta_{j+1} q_1^{j+1} \leq \theta_j \\
 & \alpha f_{j+1,1} q_1^{j+1} + \alpha t_2 q_2^{j+1} \leq \alpha f_{j1} \\
 & \alpha f_{j+1,2} q_1^{j+1} - \alpha q_2^{j+1} + \alpha t_3 q_3^{j+1} \leq \alpha f_{j2} \\
 & \vdots \\
 & \alpha f_{j+1,d-1} q_1^{j+1} - \alpha q_{d-1}^{j+1} + \alpha t_d q_d^{j+1} \leq \alpha f_{j,d-1} \\
 & \alpha f_{j+1,d} q_1^{j+1} - \alpha q_d^{j+1} \leq \alpha f_{jd} \\
 & \mathbf{q}^{j+1} \geq \mathbf{0}
 \end{aligned}$$

The dual constraints can be rewritten as follows:

$$\begin{aligned}
[\mathbf{DC}_0] \quad & q_1^{j+1} \leq \frac{\theta_j}{\theta_{j+1}} \\
[\mathbf{DC}_1] \quad & q_1^{j+1} \leq \frac{f_{j1}}{f_{j+1,1}} - \frac{t_2}{f_{j+1,1}} q_2^{j+1} \\
& \vdots \\
[\mathbf{DC}_k] \quad & q_1^{j+1} \leq \frac{f_{jk}}{f_{j+1,k}} + \frac{q_k^{j+1}}{f_{j+1,k}} - \frac{t_{k+1}}{f_{j+1,k}} q_{k+1}^{j+1} \\
& \vdots \\
[\mathbf{DC}_{d-1}] \quad & q_1^{j+1} \leq \frac{f_{j,d-1}}{f_{j+1,d-1}} + \frac{q_{d-1}^{j+1}}{f_{j+1,d-1}} - \frac{t_d}{f_{j+1,d-1}} q_d^{j+1} \\
[\mathbf{DC}_d] \quad & q_1^{j+1} \leq \frac{f_{jd}}{f_{j+1,d}} + \frac{q_d^{j+1}}{f_{j+1,d}}
\end{aligned}$$

where $q_1^{j+1} \geq 0$ and $[\mathbf{DC}_k]$ is the k th dual constraint.

Define inductively the sequence $\mathbf{t}^{(n)} \in \mathbb{R}_+^{d-1}$ as follows: $K_{\mathbf{t}^{(1)}} := \mathcal{B}^\dagger(\mathbb{R}_-^d)$, and for all $n > 1$, $K_{\mathbf{t}^{(n)}} := \mathcal{B}^\dagger(K_{\mathbf{t}^{(n-1)}})$.

Lemma G.12. The sequence $(\mathbf{t}^{(n)}) \subseteq \mathbb{R}_+^{d-1}$ is well-defined and satisfies:

- (a) $\mathbf{t}^{(n)} \geq \mathbf{1}$ for all $n \geq 1$.
- (b) $\mathbf{t}^{(n)} \leq \mathbf{t}^{(n-1)}$ for all $n \geq 1$.
- (c) The largest fixed point of $\mathcal{B}^\dagger$, $K_*^\dagger = \lim_{n \rightarrow \infty} K_{\mathbf{t}^{(n)}}$.
- (d) For each $j = 1, \dots, d-1$, $t_{j+1}^{(n)} = 1/q_1^{j+1}(\mathbf{t}^{(n-1)})$ where $q_1^{j+1}(\mathbf{t}^{(n-1)})$ is the solution to $[\mathbf{D}_{j+1,j}]$ given $\mathbf{t}^{(n-1)}$.

Proof. Proposition G.11 shows that $\mathbf{t}^{(1)}$ is well-defined. The above linear programming formulation shows that $\mathbf{t}^{(n)}$ is well-defined for all $n > 1$ and is characterized as stated in part (d) of the present lemma. It remains to show parts (a)–(c). Part (b) follows from monotonicity of $\mathcal{B}^\dagger$. Part (c) follows from the simple observation that $\mathcal{B}^\dagger$ is order-continuous on the complete lattice $(2^{\mathbb{R}^d}, \supseteq)$ and arguments from Section 4 of Echenique (2005). Finally, suppose part (a) fails, ie, there exists some $n \geq 1$ such that $\mathbf{t}^{(n)} \not\geq \mathbf{1}$. Then, by parts (b) and (c), it follows that $K_*^\dagger \not\supseteq V_d$, which contradicts part (b) of Proposition G.8. $\square$

Corollary G.13. Given any $j = 1, \dots, d-1$ and $\mathbf{t}^{(n)}$ for any $n \geq 1$, constraint $[\mathbf{DC}_0]$ holds as a strict inequality at any optimal solution of $[\mathbf{D}_{j+1,j}]$. Thus, $u_j = 0$ at any optimal solution of $[\mathbf{P}_{j+1,j}]$.

Proof. Towards a contradiction, suppose $[\mathbf{DC}_0]$ holds as an equality at some optimal solution given $\mathbf{t}^{(n)}$. This implies $q_1^{j+1} = \theta_j/\theta_{j+1} > 1$. Then, by part (d) of Lemma G.12,

$t_{j+1}^{(n+1)} < 1$. But this contradicts part (a) of Lemma G.12. That $u_j = 0$ at any optimum then follows from Complementary Slackness (part (b) of Theorem G.1). $\square$

The following lemma describes which constraints bind in the dual program.

Lemma G.14. Let $t^{(n)}$ for some $n \geq 1$ be given. Suppose the dual constraint $[\mathbf{DC}_k]$ binds and $w_{j,k} < 0$. Then, all the prior dual constraints also bind, except perhaps for $[\mathbf{DC}_0]$.

Proof. If the $k - 1$ -th dual constraint does not bind, then by Complementary Slackness (Theorem G.1), it must be that $w_{j,k-1} = 0$, which would violate the requirement that $w_{j,k-1}t_j \leq w_{j,k} < 0$. (Note that $t_j \in \mathbb{R}_+$ by Proposition G.11.) Hence, by induction, all prior constraints, other than $[\mathbf{DC}_0]$, must hold, which proves the claim. $\square$

Lemma G.15. Suppose the type process satisfies MLRP. Then, $q_k^{j+1} > 0$ and $[\mathbf{DC}_k]$ binds for all $k = 1, \dots, d$.

Proof. Let $m \leq d$ be such that $q_m^{j+1} = 0$. Then, $w_{jm} = 0$ is a solution of the primal (as is easily seen from the Lagrangian). But this implies $w_{j,m+1} = 0$ because $w_{jm}t_{m+1} \leq w_{j,m+1}$. Hence, $q_{m+1}^{j+1} = 0$ also. Inductively, we conclude that $q_n^{j+1} = 0$ for all $n > m$.

Let $k < d$ be the largest number such that $q_k^{j+1} > 0$ (and hence, as seen above, $q_i^{j+1} > 0$ for all $i > k$). Then, $w_{jk} < 0$, which implies $[\mathbf{DC}_k]$ holds as an equality. In this case, we have

$$q_1^{j+1} = \frac{f_{jk}}{f_{j+1,k}} + \frac{q_k^{j+1}}{f_{j+1,k}} > \frac{f_{jd}}{f_{j+1,d}}$$

where the inequality is because of MLRP. But this violates $[\mathbf{DC}_d]$ (recall that $q_d^{j+1} = 0$), so we conclude that either $q_k^{j+1} = 0$ for all $k = 1, \dots, d$ or $q_k^{j+1} > 0$ for all $k = 1, \dots, d$. The former is impossible because that would imply $w_{jk} = 0$ is optimal for all $k = 1, \dots, d$, which is impossible by Proposition G.11, as it violates $[\mathbf{IC}_{ij}]$ where $i = j + 1$. Hence, it must be that $q_d^{j+1} > 0$ (which implies $w_{jd} < 0$, as is easily seen from Complementary Slackness in Theorem G.1) and so by Lemma G.14, $[\mathbf{DC}_k]$ binds for all $k = 1, \dots, d$. $\square$

Lemma G.16. Suppose the type process satisfies PPR and $\mathbf{f}_j \neq \mathbf{f}_{j+1}$. In the dual problem $[\mathbf{D}_{j+1,j}]$, $q_k^{j+1} > 0$ for all $k \leq j + 1$ and $q_k^{j+1} = 0$ for all $k > j + 1$. Moreover, $[\mathbf{DC}_k]$ binds for all $k = 1, \dots, j + 1$, while $[\mathbf{DC}_k]$ is slack for all $k = j + 2, \dots, d$.

Proof. It is clear that at the optimum, we must have $q_1^{j+1} \leq 1$. If $[\mathbf{DC}_d]$ binds, we must have $w_{jd} < 0$ (from the Lagrangian or from Complementary Slackness of Theorem G.1), which implies $q_d^{j+1} > 0$. But this means $q_1^{j+1} = \frac{f_{jd}}{f_{j+1,d}} + \frac{q_d^{j+1}}{f_{j+1,d}} = 1 + \frac{q_d^{j+1}}{f_{j+1,d}} > 1$, which is impossible, where $f_{jd} = f_{j+1,d}$ because the type process satisfies PPR. Therefore, $[\mathbf{DC}_d]$

must be slack and we must have $w_{jd} = 0 = q_d^{j+1}$. Inductively, it follows that $[\mathbf{DC}_k]$ is slack for all $k > j + 1$. If $j + 1 = d$, then the claim holds trivially.

Because $\mathbf{f}_j \neq \mathbf{f}_{j+1}$, PPR implies that the type process satisfies, $f_{j,j+1} < f_{j+1,j+1}$. Then, the argument above establishes that if $[\mathbf{DC}_k]$ is slack for $k = j + 1$, it must be slack for all $k < j + 1$, because $f_{jk}/f_{j+1,k} > f_{j,j+1}/f_{j+1,j+1}$ for all $k \neq j + 1$ (by PPR). But this is clearly impossible, as it violates $[\mathbf{IC}_i]$ where $i = j + 1$, as demonstrated in Proposition G.11. This completes the proof. $\square$

Notice that if we have $f_{j,j+1} = f_{j+1,j+1}$, PPR implies that $f_{jk} = f_{j+1,k}$ for all $k = 1, \dots, d$. In this case, the optimal solution to the dual program is $q_1^{j+1} = 1$ and $q_k^{j+1} = 0$ for all $k > 1$. Clearly, all the constraints hold with equality.

Lemma G.17. Fix $\mathbf{t} \in \mathbb{R}_+^{d-1}$ such that $\mathbf{t} \geq \mathbf{1}$ and $t_2 \cdots t_d > 1$, and let $K_{\mathbf{t}}$ be defined as in [G.6] above. If the type process satisfies MLRP, $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$.

Proof. By Lemma G.15, the dual constraints $[\mathbf{DC}_k]$ bind for all $k = 1, \dots, d$ because MLRP holds. A little algebra (eliminating q_k^{j+1} for $k = 2, \dots, d$) tells us that

$$q_1^{j+1} = \frac{f_{j1} + t_2 f_{j2} + \cdots + t_2 \cdots t_d f_{jd}}{f_{j+1,1} + t_2 f_{j+1,2} + \cdots + t_2 \cdots t_d f_{j+1,d}}$$

for all $j = 1, \dots, d - 1$. Recall that the maximal value for $[\mathbf{P}_{j+1,j}]$ is v_j . By the duality theorem for Linear Programming, it must be that $v_j = q_1^{j+1} v_{j+1}$, where the latter is the value of $[\mathbf{D}_{j+1,j}]$. In other words, any $\mathbf{v}$ with $v_{j+1} \geq t'_{j+1} v_j$ is implementable, where $t'_{j+1} = 1/q_1^{j+1}$. This implies $\mathcal{B}^\dagger(K_{\mathbf{t}}) = K_{\mathbf{t}'}$, where $\mathbf{t}' = (t'_2, \dots, t'_d)$. It is easy to see that

$$t'_2(\mathbf{t}) \cdot t'_3(\mathbf{t}) \cdots t'_d(\mathbf{t}) = \frac{f_{d1} + t_2 f_{d2} + \cdots + t_2 \cdots t_d f_{dd}}{f_{11} + t_2 f_{12} + \cdots + t_2 \cdots t_d f_{1d}}$$

To show that $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$, it suffices to show that $t'_2(\mathbf{t}) \cdot t'_3(\mathbf{t}) \cdots t'_d(\mathbf{t}) \neq t_2 \cdots t_d$. But this follows immediately from the observation that

$$\frac{1}{t_2 \cdots t_d} f_{d1} + \frac{1}{t_3 \cdots t_d} f_{d2} + \cdots + f_{dd} < 1 < f_{11} + t_2 f_{12} + \cdots + t_2 \cdots t_d f_{1d}$$

which is true because $\mathbf{t} \geq \mathbf{1}$ and $t_2 \cdots t_d > 1$. $\square$

Lemma G.18. Fix $\mathbf{t} \in \mathbb{R}_+^{d-1}$ such that $\mathbf{t} \geq \mathbf{1}$ and $t_2 \cdots t_d > 1$, and let $K_{\mathbf{t}}$ be defined as in [G.6] above. If the type process satisfies PPR, $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$.

Proof. Consider the dual problem $[\mathbf{D}_{j+1,j}]$ for some $j = 1, \dots, d - 1$. By Lemma G.16, it follows that the dual constraints $[\mathbf{DC}_k]$ bind for all $k = 1, \dots, j + 1$. A little algebra tells us that

$$q_1^{j+1} = \frac{f_{j1} + t_2 f_{j2} + \cdots + t_2 \cdots t_{j+1} f_{j,j+1}}{f_{j+1,1} + t_2 f_{j+1,2} + \cdots + t_2 \cdots t_{j+1} f_{j+1,j+1}}$$

Recall that the maximal value for $[\mathbf{P}_{j+1,j}]$ is v_j . By the duality theorem for Linear Programming, it must be that $v_j = q_1^{j+1} v_{j+1}$, where the latter is the value of $[\mathbf{D}_{j+1,j}]$. In other words, any $\mathbf{v}$ with $v_{j+1} \geq t'_{j+1} v_j$ is implementable, where $t'_{j+1} = 1/q_1^{j+1}$. This implies $\mathcal{B}^\dagger(K_{\mathbf{t}}) = K_{\mathbf{t}'}$, where $\mathbf{t}' = (t'_2, \dots, t'_d)$. To show that $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$, it suffices to show that if $t_{j+1} > 1$, then $t'_{j+1} < t_{j+1}$.

To see that this is the case, observe that $t'_{j+1} < t_{j+1}$ if, and only if, $\varphi(t_{j+1}) > 0$, where

$$\begin{aligned} \varphi(t_{j+1}) := & t_{j+1}^2 [t_2 \cdots t_j f_{j,j+1}] + t_{j+1} [f_{j1} + t_2 f_{j2} + \cdots + t_2 \cdots t_j f_{jj} - t_2 \cdots t_j f_{j+1,j+1}] \\ & - [f_{j+1,1} + t_2 f_{j+1,2} + \cdots + t_2 \cdots t_j f_{j+1,j}] \end{aligned}$$

The function φ is a quadratic in t_{j+1} . It is clear that $\varphi(0) < 0$, and a simple calculation establishes that $\varphi(1) = 1$. Moreover, φ is convex, which implies $\varphi(t_{j+1}) > 0$ whenever $t_{j+1} > 1$, which establishes the claim. $\square$

G.2.3. Iterates of $\mathcal{B}$ in the CARA Case

Define inductively the sequence of sets $K^{(n)}$ as follows: $K^{(0)} := \mathbb{R}_-^d$, and for all $n \geq 1$, $K^{(n)} := \mathcal{B}(K^{(n-1)})$.

Lemma G.19. The sequence $(K^{(n)})$ satisfies:

- (a) For each $n \in \mathbb{N}$, $K^{(n)}$ is a closed, convex cone in $\mathbb{R}_-^d$.
- (b) $K^{(n)} \subseteq K^{(n-1)}$ for all $n \geq 1$.
- (c) The largest fixed point of $\mathcal{B}$, $K_* = \lim_{n \rightarrow \infty} K^{(n)}$ and is a closed, convex cone.

Proof. Part (a) follows from induction and Lemma G.4. Part (b) follows from monotonicity of $\mathcal{B}$. The limit statement in part (c) follows from the simple observation that $\mathcal{B}$ is order-continuous on the complete lattice $(2^{\mathbb{R}_-^d}, \supseteq)$ and arguments in Section 4 of Echenique (2005). That K_* is a closed, convex cone follows from the limit statement and monotonicity of the sequence as described in part (b). $\square$

We would like to establish analogues of Lemma G.12 and Corollary G.13 for this more general setting. To do so, we show how to compute iterates of $\mathcal{B}$ through solutions of conic programs. (As will become clear, due to the multiple incentive constraints, in contrast to the LP analysis in the previous section, it is no longer clear that the iterates of $\mathcal{B}$ are *polyhedral* cones; even if they are, the generating polyhedra are more difficult to explicitly describe. Since we will not need to compute the iterates in closed form, we adopt the more general approach of conic, instead of linear, programming.)

Given a convex cone $C \subseteq \mathbb{R}^n$, the *dual cone* is $C^* := \{y \in \mathbb{R}^n : y \cdot x \geq 0 \ \forall x \in C\}$. We will use the following version of Conic Programming Strong Duality in the sequel.⁵⁴

Theorem G.2 (Conic Programming Strong Duality). *Let $J \subseteq \mathbb{R}^n$ and $L \subseteq \mathbb{R}^m$ be closed convex cones, and let J^* and L^* denote the respective dual cones. Let $b \in \mathbb{R}^m$ and $A \in \mathbb{R}^{m \times n}$. Consider the primal conic program*

$$[\mathbf{P}^{\text{cone}}] \quad \max_y \quad c \cdot y \quad \text{subject to } b - Ay \in L \text{ and } y \in J$$

and the dual conic program

$$[\mathbf{D}^{\text{cone}}] \quad \min_q \quad b \cdot q \quad \text{subject to } A^\top q - c \in J^* \text{ and } q \in L^*$$

Suppose that $[\mathbf{P}^{\text{cone}}]$ satisfies the Slater condition, ie, there exists $z \in J$ such that $b - Az \in \text{int}L$. Then, if $[\mathbf{P}^{\text{cone}}]$ is feasible and has finite value:

- (a) The dual problem $[\mathbf{D}^{\text{cone}}]$ is feasible.
- (b) The values of $[\mathbf{P}^{\text{cone}}]$ and $[\mathbf{D}^{\text{cone}}]$ are equal (and finite).

We will also use the following piece of notation in the sequel: for any $v \in \mathbb{R}_-^d$ and $j = 1, \dots, d-1$, define $v_{>j} := (v_{j+1}, \dots, v_d) \in \mathbb{R}_-^{d-j}$.

Given some closed, convex cone $K \subseteq \mathbb{R}_i^d$, some $j = 1, \dots, d-1$ and $v_{>j} \in \mathbb{R}_i^{d-j}$, consider the following conic programming problem:

$$\begin{aligned}
[\mathbf{P}_j^{\text{full}}] \quad & \max_{u_j, \mathbf{w}_j} \quad u_j + \alpha[f_{j1}w_{j1} + f_{j2}w_{j2} + \dots + f_{jd}w_{jd}] \\
\text{s.t.} \quad & \frac{\theta_{j+1}}{\theta_j} \cdot u_j + \alpha[f_{j+1,1}w_{j,1} + \dots + f_{j+1,d}w_{j,d}] \leq v_{j+1} \\
& \vdots \\
& \frac{\theta_d}{\theta_j} \cdot u_j + \alpha[f_{d,1}w_{j,1} + \dots + f_{d,d}w_{j,d}] \leq v_d \\
& u_j \leq 0 \\
& \mathbf{w}_j \in K
\end{aligned}$$

It is easy to see that this is a primal conic programming problem of the form described in $[\mathbf{P}^{\text{cone}}]$, with $y = (u_1, \mathbf{w}_1)$, $b = v_{>j}$, $L = \mathbb{R}_+^d$, $J = \mathbb{R}_- \times K$, and the vector c and matrix A easy to read off from the problem. Thus, the dual cones are $L^* = L$ and $J = \mathbb{R}_- \times K^*$. It

---

(54) This statement of conic duality is from Theorem 4.7.1 of Gärtner and Matousek (2012, p. 62).

follows that the dual conic program may be written as

$$\begin{aligned}
[\mathbf{D}_j^{\text{full}}] \quad & \min_{\mathbf{q}^j \in \mathbb{R}^{d-j}} \sum_{k=j+1}^d q_{k-j}^j v_k \\
[\mathbf{DC}_0^{\text{full}}] \quad & \text{s.t.} \quad \sum_{k=j+1}^d q_{k-j}^j \theta_k \leq \theta_j \\
[\mathbf{DC}_{K^*}^{\text{full}}] \quad & \alpha \left[\mathbf{f}_{j+1}, \mathbf{f}_{j+2}, \dots, \mathbf{f}_d \right] \mathbf{q}^j - \alpha \mathbf{f}_j \in K^* \\
& \mathbf{q}^j \geq \mathbf{0}
\end{aligned}$$

where $\left[\mathbf{f}_{j+1}, \mathbf{f}_{j+2}, \dots, \mathbf{f}_d \right]$ denotes the $d \times (d-j)$ matrix whose k th column is given by the (column) vector of transition probabilities from state $j+k$, ie, $\mathbf{f}_{j+k} = (f_{j+k,1}, \dots, f_{j+k,d})^\top$.

Let $MP_j : \mathbb{R}_i^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_-^{d+1}$ denote the argmax correspondence from $[\mathbf{P}_j^{\text{full}}]$, as a function of $(\mathbf{v}_{>j}, K)$. (Recall that $\mathcal{K}$ denotes the space of closed convex cones in $\mathbb{R}_-^d$.) Let $VP_j : \mathbb{R}_i^{d-j} \times \mathcal{K} \rightarrow \mathbb{R}_-$ denote the value function for $[\mathbf{P}_j^{\text{full}}]$. Similarly, let $MD_j : \mathbb{R}_i^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_+^{d-j}$ and $VD_j : \mathbb{R}_i^{d-j} \times \mathcal{K} \rightarrow \mathbb{R}_-$ denote the argmin correspondence and value function from $[\mathbf{D}_j^{\text{full}}]$.

Lemma G.20. Let $j \leq d-1$ be given. For any $\{0\} \neq K \in \mathcal{K}$, the following hold:

- (a) For any $\mathbf{v}_{>j} \in \mathbb{R}_i^{d-j}$, the Slater condition and strong conic programming duality (conclusions (a) and (b) of Theorem G.2) hold.
- (b) $VP_j(\cdot, K)$ and $VD_j(\cdot, K)$ are well-defined (ie, finite-valued) continuous, concave, and HD1.
- (c) $MP_j(\cdot, K)$ and $MD_j(\cdot, K)$ are nonempty-valued and upper hemi-continuous. $MD_j(\cdot, K)$ is HD0.
- (d) $VP_j(\mathbf{v}_{>j}, K) = VD_j(\mathbf{v}_{>j}, K) < 0$ if and only if $\mathbf{v}_{>j} \neq \mathbf{0}$.

Proof. First, observe that $[\mathbf{P}_j^{\text{full}}]$ is clearly feasible and satisfies the Slater condition, as it is always possible to set $u_j = 0$ and scale up $\mathbf{w}_j$ along some ray contained in K so that all of the incentive constraints hold as strict inequalities. Note that the following choice is always feasible in $[\mathbf{P}_j^{\text{full}}]$: $\mathbf{w}_j = \mathbf{0}$ and $u_j = \theta_j \cdot \min_{k=j+1, \dots, d} [v_k / \theta_k]$. This policy has finite value $\theta_j \cdot \min_{k=j+1, \dots, d} [v_k / \theta_k]$. (The value is always bounded above by 0.) Thus, it is without loss of optimality to restrict attention to $(u_j, \mathbf{w}_j) \in R(\mathbf{v}_{>j})$ where

$$R(\mathbf{v}_{>j}) := \left\{ (u_j, \mathbf{w}_j) \in \mathbb{R}_-^{d+1} : u_j + \alpha [f_{j1}w_{j1} + f_{j2}w_{j2} + \dots + f_{jd}w_{jd}] \geq \theta_j \cdot \min_{k=j+1, \dots, d} [v_k / \theta_k] \right\}$$

is nonempty and compact. Define the correspondence $M : \mathbb{R}_i^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_-^{d+1}$ by

$$M(\mathbf{v}_{>j}, K) := \left\{ (u_j, \mathbf{w}_j) \in R(\mathbf{v}_{>j}) : (u_j, \mathbf{w}_j) \text{ is feasible in } [\mathbf{P}_j^{\text{full}}] \text{ at } (\mathbf{v}_{>j}, K) \right\}$$

It is easy to see that, for fixed K , $M(\cdot, K)$ is nonempty- and compact-valued and continuous.

Part (a) thus follows from Theorem G.2. The constraint set in $[\mathbf{D}_j^{\text{full}}]$ independent of $\mathbf{v}_{>j}$ and is clearly compact (due to $[\mathbf{DC}_0^{\text{full}}]$ and $\mathbf{q}^j \geq 0$). Part (c) thus follows from these observations and Berge's Theorem of the Maximum, as does continuity of the value functions stated in part (b). Concavity follows from the fact that $VD_j(\cdot, K)$ is the lower envelope of linear functions and strong duality from part (a). HD1 follows from the simple observation that $MD_j(\cdot, K)$ must be HD0, ie, $\mathbf{q}^j \in MD_j(\mathbf{v}_{>j}, K)$ if and only if $\mathbf{q}^j \in MD_j(a\mathbf{v}_{>j}, K)$ for all $a \geq 0$.

Finally, part (d) is obvious upon inspection of the constraints in $[\mathbf{P}_j^{\text{full}}]$. $\square$

Now we may state the analogue of Lemma G.12.

Lemma G.21. The sequence $(K^{(n)})$ satisfies:

$$[\text{G.7}] \quad K^{(n)} = \{ \mathbf{v} \in \mathbb{R}_-^d : v_j \leq VP_j(\mathbf{v}_{>j}, K^{(n-1)}) \text{ for all } j = 1, \dots, d-1 \}$$

Thus, the largest fixed point of $\mathcal{B}$, K_* , satisfies

$$[\text{G.8}] \quad K_* = \{ \mathbf{v} \in \mathbb{R}_-^d : v_j \leq VP_j(\mathbf{v}_{>j}, K_*) \text{ for all } j = 1, \dots, d-1 \}$$

Proof. The first part is obvious from the definition of the conic problem. The second part follows from continuity of the value function $VP_j(\cdot, \cdot)$. $\square$

We remark that parts (b) and (c) of Lemma G.20 imply that $K^{(n)}$ and K_* indeed satisfy all the properties laid out in Lemma G.19.

Define the following sets:

$$\begin{aligned} A_j &:= \{ \mathbf{v} \in \mathbb{R}_-^d : v_j = 0 \} \\ B_j &:= \{ \mathbf{v} \in \mathbb{R}_-^d : v_j = VP_j(\mathbf{v}_{>j}, K_*) \} \\ A &:= \bigcup_{j=1}^d A_j \\ B &:= \bigcup_{j=1}^{d-1} B_j \end{aligned}$$

Lemma G.22. The boundary of K_* satisfies $\text{bd } K_* \subseteq A \cup B$. The interior of K_* satisfies

$$[\text{G.9}] \quad \text{int} K_* = \{ \mathbf{v} \in \mathbb{R}_-^d : v_j < VP_j(\mathbf{v}_{>j}, K_*) \text{ for all } j = 1, \dots, d-1 \}$$

Proof. This is an immediate consequence of Lemma G.21. $\square$

Finally, the desired analogue to Corollary G.13.

Lemma G.23. Suppose that $K = K^{(n)}$ for some $n \in \mathbb{N}$ and $\mathbf{v}_{>j} \neq \mathbf{0}$. Then for any $(u_j, \mathbf{w}_j) \in MP_j(\mathbf{v}_{>j}, K)$ and $\mathbf{q}^j \in MD_j(\mathbf{v}_{>j}, K)$:

- (a) $0 \preceq \mathbf{q}^j \leq 1$.
- (b) $[\mathbf{DC}_0^{\text{full}}]$ holds as a strict inequality.
- (c) $u_j = 0$.

Proof. By part (d) of Lemma G.20, $\mathbf{v}_{>j} \neq \mathbf{0}$ implies that $VD_j(\mathbf{v}_{>j}, K) < 0$. Thus it cannot be that $\mathbf{q}^j = \mathbf{0}$. Let $I_0 := \{k > j : v_k = 0\}$. There exists some $k > j$ such that $k \notin I_0$. Thus, each of the incentive constraints in $[\mathbf{p}_j^{\text{full}}]$ with $k \in I_0$ must hold as strict inequalities. Thus, by complementary slackness, $q_{k-j}^j = 0$ for all $k \in I_0$. Suppose that there exists $k^* \notin I_0$ such that $q_{k^*-j}^j > 1$. From the definition of $[\mathbf{DC}_0^{\text{full}}]$, for any $\mathbf{v}' \in \mathcal{B}(K)$ with $\mathbf{v}'_{>j} = a\mathbf{v}_{>j}$ for some $a > 0$, must be that $v'_j \leq \sum_{k=j+1}^d q_{k-j}^j v'_k \leq q_{k^*-j}^j v'_{k^*}$ (here we are also using HD0 of the dual argmin correspondence from part (c) of Lemma G.20). But this is equivalent to $v'_{k^*} \geq (1/q_{k^*-j}^j)v'_j$, which implies either (i) that $v'_{k^*} = 0$, or (ii) $v'_{k^*} > v'_j + \varepsilon$ for some $\varepsilon > 0$. Point (i) contradicts the hypothesis that $v_{k^*} < 0$. By Lemma G.19, point (ii) implies that $K_* \not\subseteq \text{cl } V_d$, contradicting part (a) of Proposition G.8.

This proves part (a). Part (b) is an immediate consequence of part (a) and the fact that $\theta_k < \theta_j$ for all $k > j$. Part (c) is then a consequence of complementary slackness. $\square$

G.2.4. Iterates of $\mathcal{B}$ in the General Case

We now turn to the case of general utility functions $U(\cdot)$ satisfying Assumption DARA. Unlike the case of CARA utility, it is no longer immediate that $\mathcal{B}$ maps convex cones to convex cones. Below, we show that this is nevertheless true.

Given some closed, convex cone $K \subseteq \mathbb{R}_-^d$, some $j = 1, \dots, d-1$, and $\mathbf{v}_{>j} \in \mathbb{R}_-^{d-j}$, consider the following optimization problem:

$$\begin{aligned}
[\mathbf{p}_j^{\text{DARA}}] \quad & \max_{u_j, \mathbf{w}_j} \quad u_j + \alpha[f_{j1}w_{j1} + f_{j2}w_{j2} + \dots + f_{jd}w_{jd}] \\
\text{s.t.} \quad & \psi(u_j, j+1, j) + \alpha[f_{j+1,1}w_{j,1} + \dots + f_{j+1,d}w_{j,d}] \leq v_{j+1} \\
& \vdots \\
& \psi(u_j, d, j) + \alpha[f_{d,1}w_{j,1} + \dots + f_{d,d}w_{j,d}] \leq v_d \\
& u_j \leq 0 \\
& \mathbf{w}_j \in K
\end{aligned}$$

Note that problem $[\mathbf{p}_j^{\text{DARA}}]$ is a strict generalization of $[\mathbf{p}_j^{\text{full}}]$. It is also a *concave* optimization problem because, by Lemma G.1, the constraint set is closed and convex.

Proposition G.24. Let $K \in \{K^{(n)}\}_{n \in \mathbb{N}} \cup \{K_*\}$. Then, the set $\mathcal{B}(K^{(n)})$ is independent of the discount factor $\alpha \in (0, 1)$ and the utility function $U(\cdot)$ (which satisfies Assumption 1).

Proof. We will prove the claim for $n = 0$, and the proposition then follows from induction on n . (By Lemma G.4 or, alternatively, Lemmas G.20 and G.21, each $K^{(n)}$ is indeed a closed, convex cone.) Let $n = 0$ so that $K = \mathbb{R}_+^d$. Clearly, K is independent of α and U .

Let $\mathbf{v}_{>j} \neq \mathbf{0}$ be given. We will show that any optimal solution to $[\mathbf{P}_j^{\text{DARA}}]$ entails $u_j = 0$. The Lagrangian for this problem is

$$[\text{G.10}] \quad \mathcal{L}(u_j, \mathbf{w}_j, \mathbf{q}^j, \mathbf{v}_{>j}) := u_j + \alpha \mathbf{E}^{\mathbf{f}_j}[\mathbf{w}_j] + \sum_{k=j+1}^d q_{k-j}^j (v_k - \psi(u_j, k, j) - \alpha \mathbf{E}^{\mathbf{f}_k}[\mathbf{w}_j])$$

As in the proof of Lemma G.20, it is easy to see that $[\mathbf{P}_j^{\text{DARA}}]$ satisfies the Slater condition. Thus, standard Lagrangian theorems (eg, Corollary 1 on p. 219 and Theorem 2 on p. 221 of Luenberger (1969)) apply, and state that $(u_j^*, \mathbf{w}_j^*)$ is optimal in $[\mathbf{P}_j^{\text{DARA}}]$ if and only if there exists some $\mathbf{q}^{*j}$ such that this primal-dual pair constitutes a saddle point of $\mathcal{L}$, ie,

$$[\text{G.11}] \quad \mathcal{L}(u_j, \mathbf{w}_j, \mathbf{q}^{*j}, \mathbf{v}_{>j}) \leq \mathcal{L}(u_j^*, \mathbf{w}_j^*, \mathbf{q}^{*j}, \mathbf{v}_{>j}) \leq \mathcal{L}(u_j^*, \mathbf{w}_j^*, \mathbf{q}^j, \mathbf{v}_{>j})$$

for all $(u_j, \mathbf{w}_j) \in \mathbb{R}_+ \times K, \mathbf{q}^j \in \mathbb{R}_+^{d-j}$

Moreover, it is well known that the saddle set is a product set whenever it is nonempty.

Let $\mathbf{q}^{*j}$ denote a solution to the dual conic problem $[\mathbf{D}_j^{\text{full}}]$. (Existence is guaranteed by part (c) of Lemma G.23.) A necessary first-order condition for $(u_j, \mathbf{w}_j), \mathbf{q}^{*j}$ to constitute a saddle point as in [G.11] is

$$1 - \sum_{k=j+1}^d q_{k-j}^{*j} \psi'(u_j, k, j) \geq 0$$

with equality if $u_j < 0$. But if $u_j < 0$, then $\psi'(u_j, k, j) \in (0, 1)$ and thus part (a) of Lemma G.23 implies that the inequality in the above display must be strict. This is a contradiction, so it must be that $u_j = 0$ if $(u_j, \mathbf{w}_j), \mathbf{q}^{*j}$ constitutes a saddle point of the Lagrangian. Now let $(u_j^*, \mathbf{w}_j^*)$ denote a solution to the primal conic problem $[\mathbf{P}_j^{\text{full}}]$. (By part (c) of Lemma G.23, $u_j^* = 0$.) It is then easy to see that $(u_j^*, \mathbf{w}_j^*), \mathbf{q}^{*j}$ is a saddle point of the Lagrangian; if not, these would not be primal-dual solutions of $[\mathbf{P}_j^{\text{full}}] - [\mathbf{D}_j^{\text{full}}]$. It follows from the product structure of the saddle set that *any* solution $(u_j, \mathbf{w}_j)$ to $[\mathbf{P}_j^{\text{DARA}}]$ satisfies $u_j = 0$, proving the claim.

In light of the above argument and part (c) of Lemma G.23, it is easy to see that the saddle set of the Lagrangian is simply the product of the primal-dual solutions of the conic program $[\mathbf{P}_j^{\text{full}}] - [\mathbf{D}_j^{\text{full}}]$. Thus, by Lemma G.21, the iterate $\mathcal{B}(K)$ given U is equal to the iterate given CARA utility. Thus, $\mathcal{B}(K)$ is independent of the utility function U .

Now we show independence with respect to the discount factor $\alpha \in (0, 1)$. Recall the constraint $[\mathbf{DC}_{K^*}^{\text{full}}]$ from the dual conic program $[\mathbf{D}_j^{\text{full}}]$. Note that, because K^* is a cone (dual to K), $[\mathbf{DC}_{K^*}^{\text{full}}]$ is equivalent to

$$\left[\mathbf{f}_{j+1}, \mathbf{f}_{j+2}, \dots, \mathbf{f}_d \right] \mathbf{q}^j - \mathbf{f}_j \in K^*$$

Thus, the constraint set in $[\mathbf{D}_j^{\text{full}}]$ is independent of α ; hence, the solutions and values are also independent. The claim then follows from the above observation that the Lagrangian saddle set corresponds exactly to the primal-dual solutions of the conic program. $\square$

Proposition G.25. The following hold:

- (a) K_* , the largest fixed point of $\mathcal{B}$, is independent of α and U .
- (b) $D = \mathcal{B}_\circ(K_*) = \text{int}K_*$. Thus, D is an open, convex cone and is independent of α and U .

Proof. Part (a) follows immediately from Proposition G.24, Lemma G.21, and part (c) of Lemma G.19.

For part (b), consider problem $[\mathbf{P}_j^{\text{DARA}}]$ given K_* . By Proposition G.24, K_* is characterized by Lemmas G.21 and G.22. Recall the sets A and B defined before Lemma G.22. Clearly, $\mathbf{v} \in A$ implies that $\mathbf{v} \notin \mathcal{B}_\circ(K_*)$. Consider $\mathbf{v} \in B$. It follows from the proof of Proposition G.24 that any menu $(u_j, \mathbf{w}_j)_{j \in S} \in (\mathbb{R}_- \times K_*)^d$ that implements $\mathbf{v}$ must satisfy $u_j = 0$ for some $j = 1, \dots, d-1$. Thus, $\mathbf{v} \notin \mathcal{B}_\circ(K_*)$.

It therefore follows from Lemma G.22 that $\mathcal{B}_\circ(K_*) \subseteq \text{int}K_*$. We claim that $\mathcal{B}_\circ(\text{int}K_*) \supseteq \text{int}K_*$. Let $\mathbf{v} \in K_*$, so that by Lemma G.22 $v_d < 0$ and $v_j < VP_j(\mathbf{v}_{>j}, K_*)$ for all $j = 1, \dots, d-1$. Consider the concave program $[\mathbf{P}_j^{\text{DARA}}]$ (given this $\mathbf{v}$ and K_*). The constraint set is a convex set with nonempty interior (this follows from the same simple arguments used to establish the Slater condition in the proofs of part (a) of Lemma G.20 and Proposition G.24), and thus for any feasible $(u_j, \mathbf{w}_j)$ there exists a sequence $(u_j^n, \mathbf{w}_j^n) \rightarrow (u_j, \mathbf{w}_j)$ such that each $(u_j^n, \mathbf{w}_j^n)$ lies in the interior of the feasible set. The objective function is continuous, so there exists some $N \in \mathbb{N}$ such that for any $n \geq N$ we have $v_j < u_j^n + \alpha \mathbf{E}^j[\mathbf{w}_j^n]$ for all $j = 1, \dots, d-1$. Thus, $\mathbf{v} \in \mathcal{B}_\circ(\text{int}K_*)$. Thus, we have proved the claim.

Now, monotonicity of $\mathcal{B}_\circ$ implies that $\mathcal{B}_\circ(K_*) = \mathcal{B}_\circ(\text{int}K_*) = \text{int}K_*$ and thus also that $D = \text{int}K_*$, which completes the proof. $\square$

G.3. A Useful Suboptimal Contract

In this SA, we consider the case in which the agent has utility is of the CARA class and his type process satisfies either MLRP or PPR. The goal is construct a (suboptimal) recursive

contract $\zeta = (\zeta^f, \zeta^c)$ that is [TVC]-implementable starting from any $\mathbf{v} \in D$. This serves as an intermediate step towards proving both part (f) of Theorem B.1 (see SA-G.4 below). We also show that ζ generates well-defined and finite value for the principal starting from any $\mathbf{v} \in D$, which is needed to prove part (b) of Lemma B.2 in Appendix B.2; this is nontrivial because the principal's flow cost function $C(\cdot)$ is unbounded both above and below.

For simplicity, for the remainder of (only) the present section, we restrict attention to the $d = 2$ case; extending to the general case for $d \geq 3$ is straightforward but notationally very cumbersome. Let $\beta := f_{22} - f_{12} = f_{11} - f_{21}$. (The equality $f_{22} - f_{12} = f_{11} - f_{21}$ is an identity that follows from $f_{11} + f_{12} = f_{21} + f_{22} = 1$.) Because we are considering the MLRP/PPR case types and the process is fully connected by Assumption Markov, we have $\beta \in [0, 1)$. Because we are considering CARA utility of the form $U(c) := -e^{-\rho c}$ for some $\rho > 0$, it is convenient to work directly with the variables $x_i := u_i/\theta_i$ where $\theta_i := e^{-\rho \omega_i}$. Thus, with a slight abuse of notation, we will let recursive contracts specify flow allocations in terms of the x_i instead of the u_i .

G.3.1. Some Derived Parameters and Relations Between Them

Let $t_0 := \theta_2/\theta_1$. It is easy to see that $t_0 \in (0, 1)$. Define $\eta_1 := f_{21} - t_0 f_{11}$ and $\eta_2 := f_{22} - t_0 f_{12}$. It is easy to see that $\eta_2 > 0$. However, we do not have sufficient information to sign η_1 . Let $\lambda_* := \alpha f_{21}/(1 - \alpha f_{22})$ and $\lambda^* := (\alpha \eta_1 + t_0)/(1 - \alpha \eta_2)$.

We shall now record some facts about these newly defined parameters.

Lemma G.26. We have $\eta_1 + t_0 \eta_2 > 0$. In particular, the quadratic $\psi(t) = -f_{12}t^2 + (f_{22} - f_{11})t + f_{21} > 0$ for all $t \in (0, 1)$, and $\psi(t_0) = \eta_1 + t_0 \eta_2$.

Proof. Of course, if $\eta_1 \geq 0$, then the inequality always holds. We shall show that it holds even if $\eta_1 < 0$.

Notice that

$$\begin{aligned} \eta_1 + t_0 \eta_2 &= f_{21} - t_0 f_{11} + t_0 (f_{22} - t_0 f_{12}) \\ &= -f_{12}t_0^2 + (f_{22} - f_{11})t_0 + f_{21} \\ &=: \psi(t_0) \end{aligned}$$

where $\psi(t)$ is the quadratic polynomial defined above. Notice that $\psi(1) = 0$, while $\psi(0) = f_{21} > 0$. Also observe that $\psi'(t) = -2f_{12}t + (f_{22} - f_{11})$, which implies that $\psi'(1) = -2f_{12} + f_{22} - f_{11} = \beta - 1 < 0$.

Because ψ is a quadratic, it can have at most two real zeros. But 1 is a zero of ψ , which implies that the other zero is real and (strictly) negative. Therefore, for all $t \in [0, 1)$, $\psi(t) > 0$, which completes the proof. $\square$

Lemma G.27. The following inequalities hold:

- (a) $0 < \lambda_*, \lambda^* < 1$.
- (b) $t_0 < \lambda^*$.
- (c) $\lambda_* < \lambda^*$.

Proof. Part (a) follows immediately from basic assumptions in the model, so we shall only establish parts (b) and (c).

To see (b), notice that $t_0 \leq \lambda^*$ if, and only if, $t_0 - t_0\alpha\eta_2 \leq \alpha\eta_1 + t_0$, which holds if, and only if, $\alpha(\eta_1 + t_0\eta_2) > 0$, which is true by Lemma G.26.

To see (c), notice that $\alpha f_{21}/(1 - \alpha f_{22}) < (\alpha\eta_1 + t_0)/(1 - \alpha\eta_2)$ if, and only if, $\alpha f_{21} - \alpha^2 f_{21}\eta_2 < \alpha\eta_1 + t_0 - \alpha^2 f_{22}\eta_1 - \alpha t_0 f_{22}$. This is equivalent to requiring that

$$\begin{aligned}\xi(\alpha) &:= \alpha^2(f_{21}\eta_2 - f_{22}\eta_1) + \alpha(\eta_1 - f_{21} + f_{22}t_0) + t_0 \\ &= \alpha^2 t_0(f_{22}f_{11} - f_{12}f_{21}) + \alpha t_0(f_{22} - f_{11}) + t_0 > 0\end{aligned}$$

It is easy to see that $f_{22}f_{11} \geq f_{12}f_{21}$. We also have $\alpha t_0(f_{22} - f_{11}) + t_0 = t_0[1 + \alpha(f_{22} - f_{11})] > 0$ because $1 + \alpha(f_{22} - f_{11}) > 0$. This implies $\xi(\alpha) > 0$ for all $\alpha \in (0, 1)$, as claimed. $\square$

Let $\mu^* := (f_{21} + \lambda^* f_{22})/(f_{11} + \lambda^* f_{12})$. It is easy to see that $\mu^* > 0$. We claim

Lemma G.28. With μ^* and λ^* defined as above, $\mu^* > \lambda^*$.

Proof. $\mu^* = (f_{21} + \lambda^* f_{22})/(f_{11} + \lambda^* f_{12}) > \lambda^*$ if, and only if, $f_{21} + \lambda^* f_{22} > f_{11}\lambda^* + \lambda^{*2} f_{12}$. This holds if, and only if, $\psi(\lambda^*) > 0$, where ψ is the quadratic defined in Lemma G.26. But this is true because $\lambda^* \in (0, 1)$ and $\psi > 0$ on $(0, 1)$. $\square$

Consider, again, the equations

$$[\mathbf{PK}_1] \quad \theta_1 x_1 + \alpha f_{11} w_{11} + \alpha f_{12} w_{12} = v_1$$

$$[\mathbf{PK}_2] \quad \theta_2 x_2 + \alpha f_{21} w_{21} + \alpha f_{22} w_{22} = v_2$$

$$[\mathbf{IC}_{21}] \quad \theta_2 x_1 + \alpha f_{21} w_{11} + \alpha f_{22} w_{12} \leq v_2$$

Multiply $[\mathbf{PK}_1]$ by $t_0 = \theta_2/\theta_1 < 1$ and subtract it from $[\mathbf{IC}_{21}]$. This results in

$$\alpha w_{11}(f_{21} - t_0 f_{11}) + \alpha w_{12}(f_{22} - t_0 f_{12}) \leq v_2 - t_0 v_1$$

which we rewrite as

$$[\mathbf{IC}_{21}^*] \quad \alpha \eta_1 w_{11} + \alpha \eta_2 w_{12} \leq v_2 - t_0 v_1$$

Notice that $[\mathbf{PK}_1]$ and $[\mathbf{IC}_{21}^*]$ together imply $[\mathbf{IC}_{21}]$. To see this, multiply $[\mathbf{PK}_1]$ by $t_0 = \theta_2/\theta_1 < 1$ and add it to $[\mathbf{IC}_{21}^*]$ to obtain $[\mathbf{IC}_{21}]$.

G.3.2. Suboptimal Strategy – Part 1

Let $V^{(1)} := \{\mathbf{v} \in V_2 : \mathbf{v} = (v_1, \lambda v_1), \lambda_* < \lambda \leq \lambda^*\}$ be a cone in the domain. We shall now show that every point in $V^{(1)}$ is self-generating in the sense that there exists a feasible strategy such that given a starting point $\mathbf{v} \in V^{(1)}$, we can satisfy promise keeping and incentive compatibility while staying at $\mathbf{v}$ in all subsequent periods.

For $\mathbf{v} \in V^{(1)}$ and $s \in S$, define the menu $\zeta(\mathbf{v}, s) = (\zeta^f(\mathbf{v}, s), \zeta^c(\mathbf{v}, s))$ by $\zeta^f(\mathbf{v}, s, i) := x_i(\mathbf{v})$, where $x_i(\mathbf{v})$ is independent of s and to be solved for later, and $\zeta^c(\mathbf{v}, s, i) := \mathbf{v}$, which is independent of both s and i . When $\mathbf{v} \in V^{(1)}$ is given, we will often simply denote the flow policy by x_i .

It suffices to show the following:

Lemma G.29. Let $\mathbf{v} \in V^{(1)}$. Then, there exist $x_1, x_2 < 0$ such that $\zeta(\mathbf{v}, s)$, defined as above with these x_i , solves [PK₂], [PK₁], and [IC₂₁].

Proof. There exists $x_2 < 0$ such that $(x_2, \mathbf{v})$ satisfies [PK₂] if, and only if,

$$\alpha(f_{21} + \lambda f_{22})v_1 > \lambda v_1$$

which, in turn, holds if, and only if, $\lambda > \alpha f_{21}/(1 - \alpha f_{22}) = \lambda_*$, which is true by assumption.

Similarly, there exists $x_2 < 0$ such that $(x_2, \mathbf{v})$ satisfies [PK₁] if, and only if,

$$\alpha(f_{11} + \lambda f_{12})v_1 > v_1$$

which, in turn, holds if, and only if, $\lambda < (1 - \alpha f_{11})/\alpha f_{22}$, which is always true because $\lambda < 1 < (1 - \alpha f_{11})/\alpha f_{22}$.

Finally, note that [IC₂₁[★]] does not depend on the choice of (x_1, x_2) . Thus, it suffices to show that [IC₂₁[★]] is satisfied with $\mathbf{w}_1 = \mathbf{v}$. This is true if and only if

$$\alpha(\eta_1 + \lambda \eta_2)v_1 \leq (\lambda - t_0)v_1$$

Because $v_1 < 0$, this holds if, and only if, $\alpha(\eta_1 + \lambda \eta_2) \geq (\lambda - t_0)$ which is equivalent to $\lambda \leq (\alpha \eta_1 + t_0)/(1 - \alpha \eta_2) = \lambda^*$. This holds because $\mathbf{v} = (v_1, \lambda v_1) \in V^{(1)}$. This completes the proof. $\square$

G.3.3. Suboptimal Strategy — Part 2

Recall μ^* defined above in SA-G.3.1, and define

$$V^{(2)} := \{\mathbf{v} \in V_2 : \mathbf{v} = (v_1, \mu v_1), 0 < \mu < \mu^*\}$$

We will now show that it is possible to implement $\mathbf{v} \in V^{(2)}$ with a menu such that $\mathbf{w}_i \in V^{(1)}$ for $i = 1, 2$. More precisely:

Lemma G.30. Let $\mathbf{v} \in V^{(2)}$. Then, there exists $(x_i, \mathbf{w}_i)_{i=1,2} \ll \mathbf{0}$ that implements $\mathbf{v}$ such that $\mathbf{w}_i \in V^{(1)}$ for $i = 1, 2$.

Proof. From Lemma G.29, we can implement all $\mathbf{v} \in V^{(1)}$ in this manner. Therefore, it suffices to consider $\mathbf{v} \in V^{(2)} \setminus V^{(1)}$. (Recall that $V_1 \subset V_2$ because $\lambda^* < \mu^*$, as proved in Lemma G.28.) We can readily choose $\mathbf{w}_2 \in V^{(1)}$ and $x_2 < 0$ such that $\theta_2 x_2 + \mathbf{E}^{f_2}[\mathbf{w}_2] = v_2$, ie, such that [PK₂] holds. Notice that this can be done independently of our choice of $(x_1, \mathbf{w}_1)$, as $(x_2, \mathbf{w}_2)$ does not enter in to either [PK₁] or [IC₂₁].

Let $\mathbf{v} = (v_1, v_2 = \mu v_1)$. In order to prove the lemma, we need to find $w_{11} < 0$ and $\lambda \in (\lambda_*, \lambda^*]$ such that

$$\begin{aligned}\alpha(f_{11} + \lambda f_{12})w_{11} &> v_1 \\ \alpha(\eta_1 + \lambda \eta_2)w_{11} &\leq v_2 - t_0 v_1 = (\mu - t_0)v_1\end{aligned}$$

where the first inequality is merely [PK₁] (with $x_1 < 0$ treated as a slack variable) while the second reflects [IC₂₁^{*}]. This system has a solution if, and only if,

$$\frac{1}{\alpha(f_{11} + \lambda f_{12})}v_1 < \frac{\mu - t_0}{\alpha(\eta_1 + \lambda \eta_2)}v_1$$

which has a solution if, and only if,

$$\frac{1}{\alpha(f_{11} + \lambda f_{12})} > \frac{\mu - t_0}{\alpha(\eta_1 + \lambda \eta_2)}$$

This has a solution if, and only if,

$$\begin{aligned}\mu &< \frac{\eta_1 + \lambda \eta_2 + t_0(f_{11} + \lambda f_{12})}{(f_{11} + \lambda f_{12})} \\ &= \frac{f_{21} - t_0 f_{11} + \lambda(f_{22} - t_0 f_{12}) + t_0 f_0 - t_0 \lambda f_{12}}{(f_{11} + \lambda f_{12})} \\ &= \frac{f_{21} + \lambda f_{22}}{(f_{11} + \lambda f_{12})} \\ &\leq \mu^* \quad \text{whenever } \lambda \leq \lambda^*\end{aligned}$$

which is true by the assumption that $\mathbf{v} \in V^{(2)}$. Therefore, we can choose any $\lambda \in (\lambda_*, \lambda^*]$ for which there exists $w_{11} < 0$ so that [PK₁] and [IC₂₁^{*}] (or equivalently [IC₂₁]) hold.

Notice that given $\mathbf{w}_1 \in V^{(1)}$ with the above properties, we can always find $x_1 < 0$ such that [PK₁] holds with equality. This completes the proof. $\square$

Let $(x_i(\mathbf{v}), \mathbf{w}_i(\mathbf{v}))_{i=1,2}$ be the solution described in Lemma G.30, for some given $\mathbf{v} \in V^{(2)}$. In this manner, for each $\mathbf{v} \in V^{(2)}$ and $s \in S$ we may define the menu $\zeta(\mathbf{v}, s)$ by $\zeta^f(\mathbf{v}, s, i) := x_i(\mathbf{v})$ and $\zeta^c(\mathbf{v}, s, i) := \mathbf{w}_i(\mathbf{v})$, both of which are independent of s .

G.3.4. Suboptimal Strategy — Part 3

Let $V^{(3)} := \{\mathbf{v} \in V_2 : \mathbf{v} = (v_1, \mu v_1), \mu^* \leq \mu < 1\}$, so that $V_2 = V^{(2)} \cup V^{(3)}$. We shall now show that it is possible to construct a contract (really, a sequence of menus) so that for any $\mathbf{v} \in V^{(3)}$, the continuation utilities reach $V^{(1)}$ in finitely many steps, regardless of the sequence of shocks. In particular, this time can be taken to be independent of the sequence of shocks under consideration.

Let $W \subset V_2$ be a fixed subset and fix an initial $\mathbf{v} \in V_2$ and $s^{(0)} \in S$. A contract is *W-amenable* at $(\mathbf{v}, s^{(0)})$ if there is a finite natural number N (that depends on W and $(\mathbf{v}, s^{(0)})$) such that for any sequence of shocks $s^{(1)}, \dots, s^{(N)} \in S^N$, the continuation utility reaches W in at most N steps.

Lemma G.31. Let $\mathbf{v} \in V^{(3)}$ and fix $s^{(0)}$. Then, there exists a contract that is $V^{(1)}$ -amenable at $(\mathbf{v}, s^{(0)})$.

Proof. Let $\mathbf{v} = (v_1, v_2 = \mu_0 v_1) \in V^{(3)}$ where $\mu_0 \in [\mu^*, 1)$. We will show that linear inequalities derived from the constraints possess suitable solutions; any collection of such solutions will serve as the desired $V^{(1)}$ -amenable contract.

First, consider the case where $s^{(1)} = 2$. Let $\mathbf{w}_2 = (w_{21}, \lambda w_{21})$, so that we can re-write [PK₂] as

$$\alpha(f_{21} + \lambda f_{22})w_{21} > \mu v_1$$

(where x_2 is treated as a slack variable). It is immediate that we can choose $\lambda \in (\lambda_*, \lambda^*]$ and $w_{21} < 0$ that satisfies this inequality, and hence [PK₂] with some appropriately chosen $x_2 < 0$. Thus, if the first shock is $s^{(1)} = 2$, then the continuation utility $\mathbf{w}_2 \in V^{(1)}$.

Next, consider the case where $s^{(1)} = 1$. Let $\mathbf{w}_1 = (w_{11}, \mu_1 w_{11})$, so that [PK₁] and [IC₂₁[★]] can (respectively) be re-written as

$$\begin{aligned} \alpha(f_{11} + \mu_1 f_{12})w_{11} &> v_1 \\ \alpha(\eta_1 + \mu_1 \eta_2)w_{11} &\leq (\mu_0 - t_0)v_1 \end{aligned}$$

(where $x_1 < 0$ is treated as a slack variable). We can find $w_{11} < 0$ that solves the above inequalities if, and only if,

$$\frac{1}{f_{11} + \mu_1 f_{12}} > \frac{\mu_0 - t_0}{\eta_1 + \mu_1 \eta_2}$$

which, in turn, holds if, and only if,

$$\mu_1(\eta_2 - f_{12}(\mu_0 - t_0)) > f_{11}(\mu_0 - t_0) - \eta_1$$

It is easy to see that $\eta_2 - f_{12}(\mu_0 - t_0) = f_{22} - f_{12}\mu_0 > 0$. Moreover, $f_{11}(\mu_0 - t_0) - \eta_1 = f_{11}\mu_0 - f_{21}$.

Notice that $f_{11}\mu_0 - f_{21}$ can be either positive or negative. If it is negative (or zero), then we can always find $\mu_1 \in (\lambda_*, \lambda^*]$ such that $\mu_1(\eta_2 - f_{12}(\mu_0 - t_0)) > 0$, which implies that continuation utility $\mathbf{w}_1 \in V^{(1)}$ in the following period.

However, if $f_{11}\mu_0 - f_{21} > 0$, we can set $\mu_1 := (f_{11}\mu_0 - f_{21})/(f_{22} - f_{12}\mu_0)$. Then,

$$\begin{aligned}\mu_0 - \mu_1 &= \mu_0 - (f_{11}\mu_0 - f_{21})/(f_{22} - f_{12}\mu_0) \\ &= \frac{1}{f_{22} - f_{12}\mu_0} [-f_{12}\mu_0^2 + \mu_0(f_{22} - f_{11}) + f_{21}] \\ &= \frac{1}{f_{22} - f_{12}\mu_0} \psi(\mu_0) \\ &> 0\end{aligned}$$

because $\psi(t) > 0$ for $t \in (0, 1)$. Thus, $\mu_0 > \mu_1$. Proceeding iteratively, we can find a decreasing sequence where $\mu_n := (f_{11}\mu_{n-1} - f_{21})/(f_{22} - f_{12}\mu_{n-1})$. Because $\psi(\cdot)$ is a quadratic, the sequence $(\psi(\mu_n))$ lies in $[\psi(0), \psi(\mu_0)]$. Therefore,

$$\begin{aligned}\mu_{n-1} - \mu_n &= \frac{1}{f_{22} - f_{12}\mu_{n-1}} \psi(\mu_{n-1}) \\ &\geq \frac{1}{f_{22}} \min[\psi(0), \psi(\mu_0)] > 0\end{aligned}$$

so the decreasing sequence has a difference uniformly bounded away from zero. Therefore, there exists a finite N such that $\mu_N < \mu^*$. Then, if continuation utility is in $V^{(2)}$, we can use the menu defined in Lemma G.30 that then implements a transition to $V^{(1)}$ in one more step.

Any contract constructed from the solutions to these linear inequalities is one we desired, completing the proof. $\square$

G.3.5. Suboptimal Strategy — Summary

We shall now show that there exists a contract that has finite cost for the principal and is [TVC]-implementable starting from any $\mathbf{v} \in V_2$.

Proposition G.32. There exists a recursive contract ζ that (i) is [TVC]-implementable starting from any $\mathbf{v} \in V_2$ and (ii) induces a well-defined and finite-valued cost function $Q^\zeta : V_2 \times S \rightarrow \mathbb{R}$.

Proof. The proof relies on Lemmas G.29, G.30, and G.31. Starting at $(\mathbf{v}, s)$, the contract defined in Lemmas G.30 and G.31 guides continuation promised utility to $V^{(1)}$ in finitely many steps, where the finite bound is independent of the sequence of shocks. Then, upon reaching $V^{(1)}$, continuation utility stays at the same point regardless of the shock – this is

Lemma G.29. Since there are only finitely many different utility levels for any starting point (v, s) , [TVC] clearly holds and the cost to the principal of this contract is well-defined and finite, which proves the claim. $\square$

G.4. Final Steps: Proof of Theorem B.1

We prove each part of the theorem in turn.

Proof of Part (a). Existence, non-emptiness, and convexity of D follow from part (a) of Proposition G.3. Part (c) of Proposition G.7 establishes that $D \supseteq V_d$ (and, hence, also non-emptiness of D). Openness and the independence properties follow from part (b) of Proposition G.25. $\square$

Proof of Part (b). This is exactly part (c) of Proposition G.3. $\square$

Proof of Part (c). This is exactly part (b) of Proposition G.3. $\square$

Proof of Part (d). This was proved in OA-D.1 of the main text. $\square$

Proof of Part (e). By part (b) of Proposition G.25, it suffices to establish the claim for the case of CARA utility. Thus, consider iterates of the operator $\mathcal{B}^\dagger$ under CARA utility. By Lemma G.12, the iterates of the operator $\mathcal{B}^\dagger$ starting with the set $\mathbb{R}_+^d$ are characterized by the sequence $(K_{t^{(n)}})$ and $K_*^\dagger$, the largest fixed point of $\mathcal{B}^\dagger$, is the limit of this sequence. That is, $K_*^\dagger = K_{t^*}$ where $t^* := \lim_{n \rightarrow \infty} t^{(n)}$.

But we have shown in Lemmas G.17 and G.18 that under either MLRP or PPR, it cannot be that $t^* > 1$. Thus, by part (a) of Lemma G.12, $t^* = 1$. Thus, $K_*^\dagger = \text{cl}V_d$. It follows from part (b) of Proposition G.8 and part (a) of Proposition G.9 that $K_* = \text{cl}V_d$. Thus, part (c) of Proposition G.9 and Proposition G.10 imply that $D = V_d$ (alternatively, this last step can be deduced from part (b) of Proposition G.25). $\square$

Proof of Part (f). Given part (e) of Theorem B.1, proved above, it suffices to show that $D^* = V_d$ under the additional assumption of CARA utility. This was established in Proposition G.32 of SA-G.3. $\square$

H. Proofs for Sections 3.2 and 3.3

This SA collects proofs related to the regularity conditions from Subsection 3.2 and the Bellman equation from Subsection 3.3. We first prove Lemma B.2 in SA-H.1. In SA-H.2, we state and exposit Theorem H.1, which extends the treatment of the Bellman equation from Subsection 3.3. We prove most of Theorem H.1 in SA-H.3 (part (c) of that theorem is proved separately in OA-C.3.3 of BKL).

H.1. Proof of Lemma B.2

We begin with the proof of part (a) of Lemma B.2. (Note that, if $\mathbf{v} \in V_d$, then part (b) of Lemma G.5 and Lemma G.6 together imply that $\Gamma_\circ(\mathbf{v}) \neq \emptyset$. The following proof extends the same reasoning to the potential case in which $D \not\supseteq V_d$.)

Proof of Part (a), Case (i). Let the type process satisfy FOSD. Let $\mathbf{v} \in D$ and consider any menu $(u_i, \mathbf{w}_i)_{i \in S} \in \Gamma(\mathbf{v})$. We will construct a perturbation of this menu that lies in $\Gamma_\circ(\mathbf{v})$, thus establishing nonemptiness.

Let $\ell \in S$ be given. By hypothesis, $(u_\ell, \mathbf{w}_\ell) \in \mathcal{U} \times D$ satisfies [PK_i] ($i = \ell$) and each of the [IC_{ij}] ($i > j = \ell$). Let $\varepsilon > 0$, and let $w'_{\ell k} := w_{\ell k} - \varepsilon/2^k$ for $k = 1, \dots, d$. Now, let $u'_\ell := u_\ell + \delta$ such that $\delta = \alpha \varepsilon \mathbf{E}^{\mathbf{f}_\ell}[(\frac{1}{2}, \dots, \frac{1}{2^d})]$. By construction, $(u'_\ell, \mathbf{w}'_\ell)$ satisfies [PK_i] ($i = \ell$). Note that we may take $\varepsilon > 0$ sufficiently small so that $\mathbf{w}'_\ell \in D$ and $u_\ell \in \mathcal{U}$, as both D and $\mathcal{U}$ are open sets (by part (a) of Theorem B.1).

Recall from Lemma G.1 that for all $i > \ell$ we have $\psi'(u, i, \ell) < 1$ since $u < 0$, which means that $\psi(u_\ell + \delta, i, \ell) - \psi(u_\ell, i, \ell) = \int_0^\delta \psi'(u + y, i, \ell) dy < \delta$. Therefore, for all $i > \ell$, we have

$$[\text{H.1}] \quad [\psi(u_\ell + \delta, i, \ell) - \psi(u_\ell, i, \ell)] - \alpha \varepsilon \mathbf{E}^{\mathbf{f}_i}[(\frac{1}{2}, \dots, \frac{1}{2^d})] < 0$$

by construction of $\delta > 0$ and because $\mathbf{E}^{\mathbf{f}_i}[(\frac{1}{2}, \dots, \frac{1}{2^d})] > \mathbf{E}^{\mathbf{f}_\ell}[(\frac{1}{2}, \dots, \frac{1}{2^d})]$ due to the assumption that $\mathbf{f}_i$ first order stochastically dominates $\mathbf{f}_\ell$ for all $i > \ell$. Then, we have

$$\begin{aligned} \psi(u'_\ell, i, \ell) + \alpha \mathbf{E}^{\mathbf{f}_i}[\mathbf{w}'_\ell] &= \psi(u_\ell, i, \ell) + [\psi(u'_\ell, i, \ell) - \psi(u_\ell, i, \ell)] + \alpha \mathbf{E}^{\mathbf{f}_i}[\mathbf{w}_\ell - \varepsilon(\frac{1}{2}, \dots, \frac{1}{2^d})] \\ &\leq v_i + [[\psi(u'_\ell, i, \ell) - \psi(u_\ell, i, \ell)] - \alpha \varepsilon \mathbf{E}^{\mathbf{f}_i}[(\frac{1}{2}, \dots, \frac{1}{2^d})]] \\ &< v_i \end{aligned}$$

where the weak inequality is because $(u_\ell, \mathbf{w}_\ell)$ satisfies [IC_{ij}] ($i > j = \ell$), and the strict inequality is because of [H.1]. Thus, it is clear that $(u'_\ell, \mathbf{w}'_\ell) \in \Gamma_\circ(\mathbf{v})$. $\square$

Proof of Part (a), Case (ii). Let the transition probabilities $\{\mathbf{f}_i\}_{i=1}^d$ be affinely independent. This implies that each of the $\mathbf{f}_j$ is an extreme point of the convex hull $\text{co}(\{\mathbf{f}_i\}_{i=1}^d)$. Thus, by the Separating Hyperplane Theorem, for each $\ell \in S$ there exists a vector $\mathbf{a}_\ell \in \mathbb{R}^d$ such that $\mathbf{E}^{\mathbf{f}_i}[\mathbf{a}_\ell] > \mathbf{E}^{\mathbf{f}_\ell}[\mathbf{a}_\ell] > 0$ for all $i > \ell$. We may then adopt exactly the same proof used in Case (i) above by simply replacing the vector $(\frac{1}{2}, \dots, \frac{1}{2^d})$ with the vectors $\mathbf{a}_\ell$ when perturbing $(u_\ell, \mathbf{w}_\ell)$. $\square$

Towards the proof of part (b) of Lemma B.2, we next present a series of lemmas. Lemma H.1, stated below, does not rely on the CARA or PPR/MLRP assumptions.

Lemma H.1. There exists a real-valued function $Q^L : D \times S \rightarrow \mathbb{R}$ such that $Q^L \leq P$ on $D \times S$.

Proof. Lemma E.1 states that the first-best value function Q^* is real-valued on $\mathcal{U}^d \times S$. Let Q^L denote the restriction of Q^* to $D \times S$. By definition of Q^* and P , we have that $Q^L \leq P$ on $D \times S$. $\square$

Lemma H.2. Suppose that the agent has CARA utility and that the type process satisfies either MLRP or PPR. Then there exists a real-valued function $Q^U : D \times S \rightarrow \mathbb{R}$ such that $P \leq Q^U$ on $D \times S$.

Proof. By part (f) of Theorem B.1, we know that under these conditions $D^* = D = V_d$. The finite upper bound is a consequence of the (suboptimal) [TVC]-implementable contract ζ constructed in SA-G.3. By Proposition G.32, the value function induced by the contract ζ is real-valued, $Q^\zeta : D \times S \rightarrow \mathbb{R}$. Moreover, by the definition of Q^ζ and P , we have that $P \leq Q^\zeta$ on $D \times S$. Thus, let $Q^U := Q^\zeta$. $\square$

Lemma H.3, stated below, does not rely on the CARA or PPR/MLRP assumptions. It will also be used later in the proof of Theorem H.1.

Lemma H.3. There exists some $\eta \in \mathbb{R}_{++}$ such that, for each $\mathbf{v} \in D$ and $(u_i, \mathbf{w}_i)_{i \in S} \in \Gamma(\mathbf{v})$,

$$\begin{aligned} u_i &\in (v_i, 0) \\ w_{ij} &\in (\eta v_i, 0) \end{aligned}$$

for all $i, j \in S$.

Proof. The upper bounds follow directly from Assumptions DARA. The lower bounds are simple consequences of the promise keeping constraints [PK_i]. First, we have $v_i < u_i$ because $\mathbf{w}_i \in D \subseteq \mathbb{R}_{--}^d$ by Assumption DARA. Second, we have $v_i < \alpha f_{ij} w_{ij}$ because $u_i \in \mathbb{R}_{--}^d$ by Assumption DARA. The lemma follows by defining $\eta := 1/(\alpha \min_{i,j \in S} f_{ij})$, which is well-defined and finite by Assumption Markov. $\square$

Lemma H.4. Suppose that the agent has CARA utility. Then regularity Condition R.2 holds.

Proof. We first bound the growth rate of the principal's flow cost function. Under CARA utility, $U(c) = -e^{-c}$ (we set the risk aversion coefficient $\rho := 1$ for notational simplicity), the cost function is given by $C(u) := U^{-1}(u) = -\log(-u)$. Let $\theta_i = e^{-\omega_i}$. Then it is easy to see that $C(u, i) = C(u/\theta_i)$. Fix an initial promise $\mathbf{v} \in \mathcal{U}^d$, a seed type $s \in S$, and a recursive contract $\xi \in \Xi^{\text{FB}}(\mathbf{v})$. Let $(\mathbf{v}^{(t)})$ and, for each $i \in S$, $(u_i^{(t)})$ and $(\mathbf{w}_i^{(t)})$ denote the

stochastic processes for promises, flow utility policies, and continuation utility policies induced by ξ initialized at $\mathbf{v}$. Fix a path $h \in \mathcal{H}$. Recall the constant $\eta > 0$ defined in Lemma H.3. Along the path h , we have

$$\begin{aligned}
C(u_i^{(t)}(h)) &= -\log \left(\left| \frac{v_i^{(t)}(h) - \alpha \mathbb{E}^{\mathbf{f}_i} [\mathbf{w}_i^{(t)}(h)]}{\theta_i} \right| \right) \\
\text{[H.2]} \quad &\geq -\log \left(\frac{1}{\theta_i} \cdot \left[|v_i^{(t)}(h)| + \sum_{j=1}^d |w_{ij}^{(t)}(h)| \right] \right) \\
\text{[H.3]} \quad &\geq \log(\min_{i \in S} \theta_i) - \log \left(\eta^t \cdot \max_{i \in S} |v_i^{(0)}| \cdot [1 + d \cdot \eta] \right)
\end{aligned}$$

where the first inequality follows from the triangle inequality and monotonicity of the $\log(\cdot)$ function, and the second inequality is a t -fold iteration of the bounds from Lemma H.3. This simplifies to

$$C(u_i^{(t)}(h), i) \geq \log(\min_{i \in S} \theta_i) - t \log(\eta) - \log(\max_{i \in S} |v_i|) - \log(1 + d \cdot \eta)$$

Note that the RHS is independent of the path $h \in \mathcal{H}$ and the current type $i \in S$. Thus, taking the α -discounted sum, it is easy to see that there exist $\kappa_1, \kappa_2 \in \mathbb{R}$ such that

$$\text{[H.4]} \quad \inf_{\xi \in \Xi^{\text{FB}}(\mathbf{v}), h \in \mathcal{H}} \sum_{t=0}^{\infty} \alpha^t C(u_{s^{(t+1)}(h)}^{(t)}(h), s^{(t+1)}(h)) \geq \kappa_1 + \kappa_2 \log(\max_{i \in S} |v_i|)$$

The discounted sum on the LHS is the realized cost to the principal of ξ along path h . The inequality in [H.4] gives a lower bound on *pathwise* costs that is uniform across paths and feasible contracts. Note also that the constants κ_1, κ_2 are independent of both the path h and the initial condition $(\mathbf{v}, s)$. An immediate consequence of [H.4] is that

$$Q^*(\mathbf{v}, s) \geq \kappa_1 + \kappa_2 \log(\max_{i \in S} |v_i|)$$

on $D \times S$. Let a (potentially different) path $h \in \mathcal{H}$ be given. The above inequality holds for all $(\mathbf{v}^{(t)}(h), s^{(t)}(h))$, ie,

$$Q^*(\mathbf{v}^{(t)}(h), s^{(t)}(h)) \geq \kappa_1 + \kappa_2 \log(\max_i |v_i^{(t)}(h)|)$$

Iterating on the same growth-rate bounds from Lemma H.3 used to derive [H.3], we may bound the RHS of the above inequality from below as follows:

$$\text{[H.5]} \quad \kappa_1 + \kappa_2 \log(\max_i |v_i^{(t)}(h)|) \geq \kappa_1 + \kappa_2 \cdot t \cdot \log(\eta) + \kappa_2 \cdot \log(\max_{i \in S} |v_i^{(0)}|)$$

The RHS of [H.5] is independent of the path $h \in \mathcal{H}$. Then, using the fact that $\alpha^t \cdot t \rightarrow 0$, we have

$$\liminf_{t \rightarrow \infty} \alpha^t \left[\inf_{h \in \mathcal{H}} Q^*(\mathbf{v}^{(t)}(h), s^{(t)}(h)) \right] \geq \lim_{t \rightarrow \infty} \alpha^t \left[\kappa_1 + \kappa_2 \cdot t \cdot \log(\eta) + \kappa_2 \cdot \log(\max_{i \in S} |v_i^{(0)}|) \right] = 0$$

which is the desired condition. $\square$

We now collect these pieces into a proof of part (b) of Lemma B.2.

Proof of Part (b). Lemmas H.1 and H.2 together imply that P must be real-valued whenever it is well-defined. Note that the inequality [H.4] in the proof of Lemma H.4 implies that P is everywhere well-defined. In particular, [H.4] implies that the lifetime discounted costs generated by any feasible recursive contract may be assumed non-negative along *each* path $h \in \mathcal{H}$ (up to the addition of a constant that depends only on the parameters κ_1, κ_2 and the initial promise $\mathbf{v} \in D$). Thus, the *expected* lifetime cost of any feasible recursive contract — and, thus, also P — is well-defined. This establishes that Condition R.1 is satisfied. Condition R.2 is satisfied by Lemma H.4. Condition R.3 is satisfied by case (i) of part (a) of Lemma B.2, proved above. $\square$

H.2. Bellman Equation

The following result generalizes and extends Proposition 3.3 from Subsection 3.3 of the main text, which then follows as an immediate consequence. Parts of the following theorem are used repeatedly in intermediate steps of the proof of Theorem 4.1 contained in OA-C of BKL.

Theorem H.1. *Suppose the environment is regular. Then the principal's value function $P : D \times S \rightarrow \mathbb{R}$ satisfies the functional equation*

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

and, for each $s \in S$, $P(\cdot, s)$ is convex and continuously differentiable. Moreover:

- (a) P is the pointwise smallest solution to [FE] that lies pointwise above Q^* , the first-best value function.
- (b) P is strictly increasing in v_1 and non-monotone in v_i for all $i > 1$. For any sequence $(\mathbf{v}^n) \subset D$ such that $\mathbf{v}^n \rightarrow \text{bd } D$, the boundary of D , we have $P(\mathbf{v}^n, s) \rightarrow +\infty$.
- (c) There exists an optimal contract ξ^* such that, for each $i \in S$, the functions $\xi^{*f}(\cdot, \cdot, i)$ and $\xi^{*c}(\cdot, \cdot, i)$ depend on $(\mathbf{v}, s)$ only through $(v_i, \dots, v_d)$.⁵⁵

---

(55) That is $\xi^{*f}(\mathbf{v}, s, i) = \xi^{*f}(\mathbf{v}', s', i)$ for all $(\mathbf{v}, s), (\mathbf{v}', s') \in D \times S$ such that $v_j = v'_j$ for all $j \geq i$, and similarly for ξ^{*c} .

- (d) *The policy correspondence derived from [FE] is nonempty-, compact-, and convex-valued and is upper hemicontinuous.*
- (e) *If, in addition, the environment is [TVC]-regular, then (i) each $P(\cdot, s)$ is strictly convex, (ii) there exists a unique optimal contract ξ^* , which satisfies the independence properties stated in part (c), (iii) for each $s \in S$, $\xi^*(\cdot, s)$ is a continuous function, and (iv) the induced allocation $\tilde{u}_{\xi^*}$ solves the sequence problem [SP].*

The proof of Theorem H.1 is in SA-H.3 below. The proof makes clear precisely which of the Conditions R.1–R.4 are used to establish each part of the theorem.

While many properties described in Theorem H.1 are standard, two points in particular warrant mention. First, part (b) establishes that the value function is non-monotone in each v_i with $i > 1$. To understand this, the case in which $d = 2$ and the type process satisfies FOSD, in which case the domain $D = V_2$ as discussed in Subsection 3.1 (see Figure 1). An increase in v_1 tightens both the promise keeping constraint [PK_{*i*}] ($i = 1$) and the incentive constraint [IC₂₁^{*}], leading to unambiguously higher costs for the principal. An increase in v_2 , on the other hand, tightens [PK_{*i*}] ($i = 2$) and adds slack to [IC₂₁^{*}]. The first effect increases costs, while the second effect lowers costs; depending on the current state $(\mathbf{v}, s)$, either of these effects can dominate. When $v_1 \approx 0$ (ie, $\mathbf{v}$ is close to the upper boundary of V_2), the first effect dominates and P is increasing in v_2 . When $\mathbf{v}$ is close enough to the diagonal (ie, the lower boundary of V_2), the second effect dominates and P is decreasing in v_2 .

Second, part (c) states that the optimal contract does not depend on all components of the state variable $(\mathbf{v}, s)$. The independence of the policy functions on the previous report $s \in S$ follows directly from the interim nature of the promised utility state variable. Because the constraint set $\Gamma(\mathbf{v})$ does not depend on s , it is easy to see that the optimization problem in [FE] may be decomposed into d separate problems, each of which conditions only on the agent's *current* report. (This property is used in OA-C.3.3 of BKL as part of the proof of Theorem 4.1.) Thus, $\mathbf{v} \in D$ is a “sufficient statistic” for $(\mathbf{v}, s)$ for the purpose of computing optimal flow and continuation utilities. Of course, the agent's previous report $s \in S$ is essential for determining the principal's value at each history, and is thus a necessary component of the principal's state variable. The fact that the policy functions for type $i \in S$ depend on the current promised utility $\mathbf{v}$ *only* through its components $(v_i, v_{i+1}, \dots, v_d)$ is a consequence of Assumption NHB, which implies that u_i and w_i only appear in those incentive constraints that deter deviations by *higher* types.

H.3. Proof of Theorem H.1

We will establish the relevant properties through a series of lemmas, and collect them into a proof of the theorem at the end.

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

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

Lemma H.5. Suppose that Condition R.1 holds.⁵⁶ Then the principal's value function, P , satisfies the Bellman equation

$$[BE] \quad P = TP$$

Proof. This follows from standard Principle of Optimality arguments, eg, a straightforward adaptation of Theorem 4.2 of Stokey, Lucas and Prescott (1989, p. 71). $\square$

Lemma H.6. Suppose that Conditions R.1 and R.2 hold. Then P is the smallest fixed point of the Bellman operator [T] in the order interval $[Q^*, \infty]$ (endowed with the pointwise order).

Proof. By Lemma H.5, it suffices to show that any fixed point $Q = TQ$ with $Q \geq Q^*$ satisfies $Q \geq P$. Because $P \geq Q^*$ by definition of these functions, Condition R.2 implies that the appropriate analogue of condition 2.17 in Kamihigashi (2014) holds (with his upper bound $\bar{v}$ replaced by the lower bound Q^* and the sign of the limit term reversed, as is appropriate given our focus on minimization and his focus on maximization). Then, an elementary adaptation of Proposition 2.1 in Kamihigashi (2014) establishes the claim. $\square$

Lemma H.7. Suppose that Condition R.1 holds. Then $P(\cdot, s)$ is convex and continuous on D for each $s \in S$.

Proof. Part (b) of Theorem B.1 states that the constraint correspondence $\Gamma(\cdot)$ has a convex graph. Thus, the feasibility correspondence $\Xi(\cdot)$ in [RP] has convex graph (when convex combinations of recursive contracts are defined pointwise). In addition, the return functions $C(\cdot, s)$ are convex. It follows from the definition of [RP] that $P(\cdot, s)$ is convex. Continuity of D follows immediately, as D is open by part (a) of Theorem B.1. $\square$

---

(56) The requirement of finiteness in Condition R.1 is not, strictly speaking, necessary. But it is sufficient and, in particular, rules out the possibility that P takes on both values $\pm\infty$, in which case the operator [T] need not be well-defined.

We will use the following notation in the coming proofs. For each menu $(\mathbf{u}, \mathbf{w}) := (u_i, \mathbf{w}_i)_{i \in S} \in (\mathcal{U} \times D)^d$ and $s \in S$, define

$$H(P, (\mathbf{u}, \mathbf{w}), s) := \sum_{i=1}^d f_{si} [C(u_i, i) + \alpha P(\mathbf{w}_i, i)]$$

Thus, in particular, the Bellman equation [BE] may be written as $P(\mathbf{v}, s) = \inf_{(\mathbf{u}, \mathbf{w}) \in \Gamma(\mathbf{v})} H(P, (\mathbf{u}, \mathbf{w}), s)$. Define the *policy correspondence* by $\Gamma^* : D \times S \rightrightarrows (\mathcal{U} \times D)^d$ by

$$[\text{H.6}] \quad \Gamma^*(\mathbf{v}, s) := \arg \min_{(\mathbf{u}, \mathbf{w}) \in \Gamma(\mathbf{v})} H(P, (\mathbf{u}, \mathbf{w}), s)$$

Lemma H.8. Suppose that Condition R.1 holds. For any convergent sequence $(\mathbf{v}^n)_{n=0}^\infty \subset D$ with limit point $\mathbf{v}^* \in \text{bd } D$, the boundary of D , we have $P(\mathbf{v}^n, s) \rightarrow +\infty$ for each $s \in S$.

Proof. Let $s \in S$ be fixed throughout. Let $(\mathbf{u}^n, \mathbf{w}^n) := (u_i^n, \mathbf{w}_i^n)_{i \in S}$ denote a sequence of menus with the following properties. First, $(\mathbf{u}^n, \mathbf{w}^n) \in \Gamma(\mathbf{v}^n)$ for all $n \in \mathbb{N}$. Now, because the sequence $(\mathbf{v}^n)$ converges and is thus bounded, it follows from Lemma H.3 that there exists a bounded set $R \subset (\mathcal{U} \times D)^d$ such that $(\mathbf{u}^n, \mathbf{w}^n) \in R$ for all $n \in \mathbb{N}$. Second, suppose that there exists some $\varepsilon > 0$ such that each $(\mathbf{u}^n, \mathbf{w}^n)$ is ε -optimal, ie,

$$[\text{H.7}] \quad H^n := H(P, (\mathbf{u}^n, \mathbf{w}^n), s) \leq P(\mathbf{v}, s) + \varepsilon$$

for all $n \in \mathbb{N}$. Such an $\varepsilon > 0$ and a corresponding sequence of menus exists by virtue of Lemma H.5.

We claim that $H^n \rightarrow +\infty$, which implies that $P(\mathbf{v}^n, s) \rightarrow +\infty$. This claim follows from two facts. First, we assert that there exists some $K \in \mathbb{R}$ such that

$$\inf_{i \in S, (\mathbf{u}, \mathbf{w}) \in R} P(\mathbf{w}_i, i) \geq K$$

To see this, observe that $\inf_{i \in S, (\mathbf{u}, \mathbf{w}) \in \text{cl } R} Q^*(\mathbf{w}_i, i) > -\infty$ by Lemma E.1, where Q^* is the first-best value function. The first fact then follows from the fact that $P \geq Q^*$ on $D \times S$. Second, we assert that $\mathbf{u}^n \rightarrow \mathbf{0}$. This follows from the structure of the domain D — see, in particular, the statements and proofs of Lemma G.23 and Proposition G.24. It is then easy to see that the claim that $H^n \rightarrow +\infty$ follows, as we have

$$H^n \geq \sum_{i=1}^d f_{si} C(u_i, i) + \alpha K$$

by the first fact, and $C(u_i^n, i) \rightarrow +\infty$ for all $i \in S$ by the second fact. This completes the proof. $\square$

Following Kreps (1977), we say that a recursive contract ξ is *conserving* if it attains the infimum in [FE] for each $(\mathbf{v}, s) \in D \times S$. (Note that, by definition, if ξ is conserving then it must be feasible.) It is easy to see that if ξ is optimal, then it must be conserving (though, the converse requires additional argument).

A *Kakutani correspondence* is a correspondence that is nonempty-, compact-, and convex-valued and is continuous.

Lemma H.9. Suppose that Condition R.1 holds. For each $(\mathbf{v}, s) \in D \times S$, the infimum in [BE] is attained. That is, there exists a conserving recursive contract ξ^* . Moreover, the policy correspondence Γ^* is nonempty-, convex-, and compact-valued and is upper hemi-continuous.

Proof. We will construct a Kakutani correspondence and show it is without loss of optimality to restrict attention to menus that lie in this correspondence. Let $(\mathbf{v}, s) \in D \times S$ be given. Define

$$L(\mathbf{v}, s) := \{(\mathbf{u}, \mathbf{w}) \in \Gamma(\mathbf{v}) : H(P, (\mathbf{u}, \mathbf{w})) \leq P(\mathbf{v}, s) + 1\}$$

Define $K(\mathbf{v}, s) := \text{cl } L(\mathbf{v}, s)$, the closure of $L(\mathbf{v}, s)$. It follows from Lemma H.3 that $\Gamma(\mathbf{v})$ is a bounded set. Thus, it follows that $L(\mathbf{v}, s)$ is bounded and thus that $K(\mathbf{v}, s)$ is compact. It is also true that $K(\mathbf{v}, s) \neq \emptyset$ because P satisfies the Bellman equation (Lemma H.5), and thus there exist ε -optimal menus for each $\varepsilon > 0$. Moreover, it is easy to see that the function $H(P, \cdot) : \Gamma(\mathbf{v}) \rightarrow \mathbb{R}$ is convex, and thus that $K(\mathbf{v}, s)$ is a convex set (since $\Gamma(\mathbf{v})$ is convex by part (b) of Theorem B.1).

We assert that the implied correspondence

$$\begin{aligned} K : D \times S &\rightrightarrows (\mathcal{U} \times D)^d \\ (\mathbf{v}, s) &\mapsto K(\mathbf{v}, s) \end{aligned}$$

is a well-defined Kakutani correspondence. To show that it is a well-defined correspondence, we must show that $K(\mathbf{v}, s) \in (\mathcal{U} \times D)^d$ for all $(\mathbf{v}, s) \in D \times S$. Clearly we must have $K(\mathbf{v}, s) \subseteq \text{cl } \Gamma(\mathbf{v})$. If $K(\mathbf{v}, s) \cap (\text{cl } \Gamma(\mathbf{v}) \setminus \Gamma(\mathbf{v})) \neq \emptyset$, then there exists some sequence $((\mathbf{u}^n, \mathbf{w}^n))_{n=0}^\infty \subseteq L(\mathbf{v}, s)$ such that $(\mathbf{u}^n, \mathbf{w}^n) \rightarrow \text{bd } (\mathcal{U} \times D)^d$. (It is easy to see that $(\text{cl } \Gamma(\mathbf{v}) \setminus \Gamma(\mathbf{v})) \subseteq \text{bd } (\mathcal{U} \times D)^d$.) By the same arguments as in the proof of Lemma H.8, $H(P, (\mathbf{u}^n, \mathbf{w}^n)) \rightarrow +\infty$, which contradicts the definition of $L(\mathbf{v}, s)$ and finiteness of $P(\mathbf{v}, s)$ (Condition R.1). Thus, K is a well-defined correspondence. To show that it is Kakutani, it suffices to show that it is continuous. But this is easy to see, as $\Gamma(\cdot)$ is a continuous correspondence, P is a continuous function (Lemma H.7), and each $H(P, \cdot, s)$ is a continuous function on $(\mathcal{U} \times D)^d$ by virtue of the continuity of the $C(\cdot, i)$ and P .

It follows that the argmin correspondence M^* defined by

$$M^*(\mathbf{v}, s) := \arg \min_{(\mathbf{u}, \mathbf{w}) \in K(\mathbf{v}, s)} H(P, (\mathbf{u}, \mathbf{w}), s)$$

is continuous and nonempty- and compact-valued by Berge's Theorem, and convex-valued because K is Kakutani and $H(P, \cdot, s)$ is convex. We showed above that $K(\mathbf{v}, s) \subseteq \Gamma(\mathbf{v})$. It is also easy to see that the policy correspondence Γ^* (defined in [H.6]) satisfies $\Gamma^*(\mathbf{v}, s) \subseteq K(\mathbf{v}, s)$. Thus, it follows that $M^*(\mathbf{v}, s) = \Gamma^*(\mathbf{v}, s)$ for all $(\mathbf{v}, s) \in D \times S$. This completes the proof. $\square$

Lemma H.10. Suppose that Conditions R.1 and R.2 hold. Any conserving contract ξ^* is optimal, ie, is optimal in the principal's recursive problem [RP]. Hence, there exists a optimal contract.

Proof. Lemma H.9 shows that a conserving contract exists under Condition R.1, so it suffices to show that any such contract is optimal for the principal. But optimality is a direct consequence of Condition R.2 and an adaptation of Theorem 4.5 in Stokey, Lucas and Prescott (1989). In particular, because P satisfies the Bellman equation [BE] (Lemma H.5) and the infimum is attained at each step (Lemma H.9), induction gives

$$P(\mathbf{v}, s) = \mathbb{E} \left[\sum_{t=0}^T \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right] + \mathbb{E} \left[\alpha^{T+1} P(\mathbf{v}^{(T+1)}, s^{(T+1)}) \middle| s^{(0)} = s \right]$$

where $(u^{(t)})_{t=0}^\infty$ is the flow utility process induced by conserving recursive contract (ie, some measurable selection from the policy correspondence). Now, Condition R.2 implies that

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

because $P \geq Q^*$ on $D \times S$. Thus, for some sufficiently large T , there exists a constant $K \in \mathbb{R}$ such that for all $t \geq T$,

$$\alpha^t \left[\inf_{h \in \mathcal{H}} P(\mathbf{v}^{(t)}(h), s^{(t)}(h)) \right] + K \geq 0$$

Thus, we may apply Fatou's Lemma to get

$$\begin{aligned} 0 &\leq \mathbb{E} \left[\liminf_{T \rightarrow \infty} \alpha^{T+1} P(\mathbf{v}^{(T+1)}, s^{(T+1)}) \middle| s^{(0)} = s \right] \\ &\leq \liminf_{T \rightarrow \infty} \mathbb{E} \left[\alpha^{T+1} P(\mathbf{v}^{(T+1)}, s^{(T+1)}) \middle| s^{(0)} = s \right] \end{aligned}$$

Combining the above displays, we get

$$[\text{H.8}] \quad P(\mathbf{v}, s) \geq \liminf_{T \rightarrow \infty} \mathbb{E} \left[\sum_{t=0}^T \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right]$$

By Condition [R.1](#) (which requires that the double integral be well-defined and finite), we have

$$\begin{aligned} \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t \max \{0, C(u^{(t)}, s^{(t+1)})\} \middle| s^{(0)} = s \right] &< +\infty \\ \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t \max \{0, -C(u^{(t)}, s^{(t+1)})\} \middle| s^{(0)} = s \right] &< +\infty \end{aligned}$$

and thus,

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t |C(u^{(t)}, s^{(t+1)})| \middle| s^{(0)} = s \right] < +\infty$$

Then, by Fubini's Theorem,

$$\lim_{T \rightarrow \infty} \mathbb{E} \left[\sum_{t=0}^T \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right] = \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right]$$

Then it follows from [\[H.8\]](#) that $P(\mathbf{v}, s) \geq \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right]$. The opposite inequality holds by the definition of P in [\[RP\]](#). Thus, $P(\mathbf{v}, s) = \mathbb{E} \left[\sum_{t=0}^{\infty} \alpha^t C(u^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right]$ and hence the conserving contract is optimal. $\square$

Lemma H.11. Suppose that Conditions [R.1](#) and [R.2](#) hold. Then $P(\cdot, s)$ is continuously differentiable on D for each $s \in S$.

Proof. From Lemma [H.9](#), the infimum in the Bellman equation is attained. Moreover, the domain D is open by part (a) of Theorem [B.1](#). The claim then follows from a straightforward adaptation of the Benveniste and Scheinkman theorem (see, eg, Theorems 4.10 and 9.10 in Stokey, Lucas and Prescott (1989)). $\square$

Lemma H.12. Suppose that Conditions [R.1](#) and [R.4](#) hold. Then $P(\cdot, s)$ is *strictly* convex for each $s \in S$.

Proof. Let ξ^* be a [\[TVC\]](#)-implementable optimal contract, the existence of which is guaranteed by Condition [R.4](#). Let $\mathbf{v}, \mathbf{v}' \in D$ be given such that $\mathbf{v} \neq \mathbf{v}'$. For fixed $\eta \in (0, 1)$, define $\mathbf{v}'' := \eta \mathbf{v} + (1 - \eta) \mathbf{v}'$. Let $(\mathbf{u}^{(t)}, \mathbf{w}^{(t)})_{t=0}^{\infty}$ and $(\mathbf{u}'^{(t)}, \mathbf{w}'^{(t)})_{t=0}^{\infty}$ denote the stochastic processes of menus generated by ξ^* starting from $\mathbf{v}$ and $\mathbf{v}'$, respectively. For each $t \in \mathbb{N}$, let

$$(\mathbf{u}''^{(t)}, \mathbf{w}''^{(t)}) := \eta(\mathbf{u}^{(t)}, \mathbf{w}^{(t)}) + (1 - \eta)(\mathbf{u}'^{(t)}, \mathbf{w}'^{(t)})$$

Thus, the process $(\mathbf{u}''^{(t)}, \mathbf{w}''^{(t)})_{t=0}^{\infty}$ denotes the *history-wise* convex combination of the processes $(\mathbf{u}^{(t)}, \mathbf{w}^{(t)})_{t=0}^{\infty}$ and $(\mathbf{u}'^{(t)}, \mathbf{w}'^{(t)})_{t=0}^{\infty}$. Because ξ^* is feasible and delivers promises (ie, satisfies [DP], by virtue of Lemma A.1) and because the constraint correspondence Γ has convex graph (by part (b) of Theorem B.1), it follows that the process $(\mathbf{u}''^{(t)}, \mathbf{w}''^{(t)})_{t=0}^{\infty}$ is feasible and delivers promises starting from $\mathbf{v}'' \in D$.

We will establish strict convexity by expanding the right-hand side of [FE]. It follows from the facts (i) that $\mathbf{v} \neq \mathbf{v}'$, (ii) that ξ^* delivers promises, and (iii) that the type process is fully connected (by Assumption Markov) that there exists some $T \in \mathbb{N}$ and two vectors $\mathbf{a}, \mathbf{b} \in \mathcal{U}^d$ with $\mathbf{a} \neq \mathbf{b}$ such that $\mathbf{P}(\mathbf{u}^{(T)} = \mathbf{a}) > 0$ and $\mathbf{P}(\mathbf{u}^{(T)} = \mathbf{b}) > 0$. Now, expanding the right-hand side of [FE] for T periods, we have

$$\begin{aligned} & \eta P(\mathbf{v}, s) + (1 - \eta) P(\mathbf{v}', s) \\ &= \sum_{t=0}^T \alpha^t \mathbf{E} \left[\eta C(u_{s^{(t+1)}}^{(t)}, s^{(t+1)}) + (1 - \eta) C(u_{s^{(t+1)}}'^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right] \\ &+ \alpha^T \mathbf{E} \left[\eta P(\mathbf{w}_{s^{(T)}}'^{(T+1)}, s^{(T+1)}) + (1 - \eta) P(\mathbf{w}_{s^{(T+1)}}'^{(T)}, s^{(T+1)}) \middle| s^{(0)} = s \right] \\ &> \sum_{t=0}^T \alpha^t \mathbf{E} \left[C(u_{s^{(t+1)}}''^{(t)}, s^{(t+1)}) \middle| s^{(0)} = s \right] + \alpha^T \mathbf{E} \left[P(\mathbf{w}_{s^{(T+1)}}''^{(T)}, s^{(T+1)}) \middle| s^{(0)} = s \right] \\ &\geq P(\mathbf{v}'', s) \end{aligned}$$

The strict inequality follows from convexity of $P(\cdot, s)$ (established in Lemma H.7), strict convexity of the $C(\cdot, i)$, and the fact that $\mathbf{P}(\mathbf{u}^{(T)} = \mathbf{a}) > 0$ and $\mathbf{P}(\mathbf{u}^{(T)} = \mathbf{b}) > 0$ where $\mathbf{a} \neq \mathbf{b}$. The weak inequality follows from the definition of $P(\cdot, s)$. This establishes strict convexity. $\square$

Lemma H.13. Suppose that Condition R.1 holds. For each $s \in S$, the value function $P(\cdot, s)$ is (a) strictly increasing in the component v_1 and (b) non-monotone in the component v_j for all $j = 2, \dots, d$.

Proof. Consider part (a). Let $\mathbf{v} \in D$ and $s \in S$ be given, and define $\mathbf{v}' := (v_1 - \varepsilon, v_2, \dots, v_d)$ for some sufficiently small $\varepsilon > 0$ so that $\mathbf{v}' \in D$. Let $(u_i^*, \mathbf{w}_i^*)_{i \in S}$ be an optimal menu at $(\mathbf{v}, s)$ (such a menu exists by Lemma H.9). It is easy to see that the menu $(u_i', \mathbf{w}_i')_{i \in S}$ defined by $u_1' := u_1^* - \varepsilon$, $u_j' := u_j^*$ for all $j \geq 2$, and $\mathbf{w}_i' := \mathbf{w}_i^*$ for all $i \in S$ satisfies $(u_i', \mathbf{w}_i')_{i \in S} \in \Gamma(\mathbf{v}')$. It satisfies the promise keeping constraints [PK_i] by construction, and adds (potentially zero) slack to all of the incentive constraints [IC_{ij}] ($i > j$). (Note that v_1 does not appear in any of the incentive constraints.) This menu is strictly less costly than the optimal menu at $\mathbf{v}$, so by revealed preference we have $P(\mathbf{v}', s) < P(\mathbf{v}, s)$. This proves part (a).

Part (b) follows immediately from part (a) of Theorem B.1, Proposition G.11, and Lemma H.8 above. In particular, those first two results imply that $D \subseteq K_{t(1)}$ (the cone $K_{t(1)}$ is constructed in Proposition G.11). It is easy to see from the structure of that set that, for any $j \geq 2$ and fixed $\mathbf{v}_{-j} := (v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_d)$, the set $\{y \in \mathcal{U} : (y, \mathbf{v}_{-j}) \in D\}$ is bounded. Non-monotonicity in the j th component then follows from Lemma H.8. $\square$

We may now collect the pieces above to prove the theorem.

Proof of Theorem H.1. Lemmas H.5 and H.9 together imply that P satisfies [FE]. Convexity and continuous differentiability follow from Lemmas H.7 and H.11, respectively. Part (a) of the theorem follows from Lemma H.6. Part (b) of the theorem follows from Lemmas H.8 and H.13. For part (c), existence of an optimal contract follows from Lemma H.10. The independence properties follow from part (a) of Lemma C.9 in OA-C.3.3. Part (d) follows from Lemma H.9.

Finally, consider part (e). Point (i), strict convexity, follows from Lemma H.12. Given strict convexity and Lemma H.9, the policy correspondence Γ^* is singleton-valued, so there exists a unique conserving contract and hence at most one optimal contract. By Lemma H.10, it is indeed the unique optimal contract, establishing point (ii). Point (iii), continuity, follows from uniqueness and the upper hemi-continuity in Lemma H.9. Point (iv) follows from Lemma A.1. $\square$

I. A Solved Example

This SA studies the short-run dynamics and insurance properties of the optimal contract in a solved example. In short, we find that (i) the optimal contract is not renegotiation-proof and exhibits rich forms of history-dependence when types are not iid, (ii) distortions relative to first-best insurance are larger when persistence is greater, and (iii) these properties are driven by the principal's need to manage the agent's Markov information rents in $[\mathbf{IC}_{ij}^*]$.

To most simply isolate the role of persistence, we focus on the case in which $d = 2$, the type process satisfies FOSD and has symmetric transitions (ie, $f_{22} = f_{11} > f_{12} = f_{21}$), and the agent has CARA utility, so that $U(c) = -e^{-\rho c}$ where $\rho > 0$ is the coefficient of absolute risk aversion.⁵⁷ It is convenient to define the marginal utility parameter $\theta_i := e^{-\rho \omega_i}$ for each $i \in S$, so that $U(\omega_i + c_i) = \theta_i U(c_i)$. Under these assumptions, the environment is

(57) The CARA assumption allows for direct comparison to the iid benchmark, namely, the solved example in Thomas and Worrall (1990). The assumptions on the type process also allow for a direct comparison to Zhang's (2009) optimal taxation model.

regular and the recursive domain is $D = V_2$, as illustrated in Figure 1.⁵⁸ We additionally assume that Condition R.4 holds, so that the environment is [TVC]-regular. All proofs and derivations for this section are in SA-J. In the numerical computations, we fix the preference parameters $\alpha, \rho, \theta_1, \theta_2$ and consider three different persistence levels for the type process: iid, medium persistence, and high persistence.⁵⁹

We begin by analyzing the dynamics of promised utility in SA-I.1 and then turn to properties of the flow allocation in SA-I.2.

I.1. Promised Utility Dynamics

The main simplification implied by CARA utility is that the optimal contract is “scale invariant” in the following sense:

Fact 1. *For any persistence level, the optimal contract is homogenous of degree one (HD1) in $\mathbf{v}$, and the solutions to the efficiency problems [Eff_{*i*}] are HD1 in $\mathbf{v}$.*

Thus, the range of solutions to the efficiency problem [Eff_{*i*}] defines a ray E_i in the domain V_2 . We will focus on the dynamics of promised utility relative to these rays. The following result, illustrated in Figure 3, characterizes the iid benchmark.

Fact 2. *When types are iid:*

- (a) *The rays E_1 and E_2 coincide. (Call the single ray E .)*
- (b) *The policy functions satisfy $\mathbf{w}_1(\mathbf{v}, s) \in E$ and $\mathbf{w}_2(\mathbf{v}, s) \in E$ for all $(\mathbf{v}, s) \in V_2 \times S$.*
- (c) *Given $\mathbf{v} \in E$, $\mathbf{w}_1(\mathbf{v}, s) \ll \mathbf{v} \ll \mathbf{w}_2(\mathbf{v}, s)$.*

The key property of iid types is that there no Markov information rent to manage. This leads to part (b) of Lemma 2, which says that optimal contract is *always* efficient (along the optimal path). Part (c) is the standard *co-insurance* property: to deter the agent from misreporting as a low endowment type, promised utility decreases after low reports.

The following result shows that, when types are persistent, the optimal contract is efficient *only* after high endowment reports (see Figures 4 and 5 for illustration).

Fact 3. *With positive persistence (ie, when $f_{11} - f_{21} \equiv f_{22} - f_{12} > 0$):*

-
- (58) For formal statements of these points, see, respectively, Lemma B.2 in Appendix B.2 and Theorem B.1 in Appendix B.1.
  - (59) The fixed preferences parameters are $\alpha = 0.5, \rho = 1, \theta_1 = 5$, and $\theta_2 = 0.1$. The transition probabilities are $f_{11} = f_{22} = 0.5$ for iid, $f_{11} = f_{22} = 0.65$ for medium persistence, and $f_{11} = f_{22} = 0.8$ for high persistence. These three transition matrices are Blackwell ordered, and can thus be interpreted as sampling from a single (continuous-time) Markov process at different intervals, where the more persistent chains correspond to more frequent sampling.

Figure 3: IID. Promised utility dynamics after 3 consecutive high reports (left) and 3 consecutive low reports (right). Different colors represent paths starting from different initial conditions.

Figure 4: Medium persistence. Promised utility dynamics after 3 consecutive high reports (left) and 3 consecutive low reports (right). Different colors represent paths starting from different initial conditions.

Figure 5: High persistence. Promised utility dynamics after 3 consecutive high reports (left) and 3 consecutive low reports (right). Different colors represent paths starting from different initial conditions.

- (a) *The ray E_1 lies strictly above the ray E_2 .*
- (b) *The policy function $\mathbf{w}_2$ satisfies $\mathbf{w}_2(\mathbf{v}, s) \in E_2$ for all $(\mathbf{v}, s) \in V_2 \times S$.*
- (c) *Given $\mathbf{v} \in E_2$, $\mathbf{w}_2(\mathbf{v}, s) \gg \mathbf{v}$.*
- (d) *On the optimal path, the policy function $\mathbf{w}_1(\mathbf{v}^{(t)}, s^{(t)})$ lies strictly below the ray E_1 .*

Part (a) follows from the simple observations that (i) a high (low) report is less (more) likely under the distribution f_1 than under f_2 , and (ii) it is cost-efficient to implement a given level of *expected* promised utility through promises to the relatively less likely type. Parts (b) and (c) of Fact 3 are reminiscent of the same parts of Fact 2. This is to be expected for, as argued above, the optimal contract must be efficient after $d = 2$ reports under Assumption NHB.

Part (d) is the key difference from the iid benchmark: to reduce the Markov information rent $\alpha(f_{22} - f_{12}) \cdot (w_{12} - w_{11})$ in $[\mathbf{IC}_{21}^*]$, the principal *compresses* the difference $w_{12} - w_{11}$ by moving $\mathbf{w}_1$ closer to the diagonal, and thus *below* the ray E_1 . Because types are persistent, by providing *inefficiently high insurance* to low types in the next period (full insurance corresponds to promised utilities that lie on the diagonal), the principal is able to punish high-type deviators who mis-report as low types for a single period. This decreases information rents *in the current period*, but also leads to a tighter incentive constraint (ie, the left-hand side of $[\mathbf{IC}_{21}^*]$) *in the next period*.⁶⁰ To compensate for this additional incentive cost, the principal pushes the level of tomorrow's promised utility down further still after consecutive low reports. Indeed, Figures 4 and 5 indicate that the region in V_2 below E_2 is absorbing, and the optimal path of promised utility follows a “triangular” path in this region. Starting from E_2 , if there is a high report, promised utility moves up along E_2 toward the origin. If instead there is a low report, promised utility moves to the *southwest* and *below* E_2 . Promised utility continues to move to the southwest — away from E_1 and the origin, and toward the diagonal — after consecutive low reports, and then jumps back to E_2 after the next high report.

Quantitatively, the key comparative static is that the promised utility vector *moves southwest faster* after low reports when *persistence is higher* (compare Figures 4 and 5). More precisely, the *slope* of the promised utility path after low reports is *steeper* at higher persistence levels. Intuitively, as persistence increases (ie, the difference $f_{22} - f_{12} > 0$ increases) the difference $w_{12} - w_{11}$ needs to be *more compressed* to achieve the same schedule of information rents.

Finally, an important qualitative implication of persistence is a new form of history-dependence that is absent in the iid case. Thomas and Worrall (1990) show that, when types are iid and utility is CARA, the *order* of reported types does not matter:

(60) See J.13 in SA-J.4 for a formal expression of this point.

Fact 4. Fix $\mathbf{v}^{(0)} \in E$, and let h be a history of realized types and $h^t = (s^1, \dots, s^t)$ be its length- t restriction. If types are iid, then $\mathbf{v}^{(t)}(h)$ depends on h^t only through the **number** of high and low reports.

This is easy to see from Facts 1 and 2, which together imply that promised utility must increase (decrease) by a fixed proportion after high (low) reports. It is illustrated in Figure 6, which plots the time-series behavior of promised utility and the policy functions along two report sequences, each of which have three low reports and three high reports but differ in the order. In the terminal period, $t = 7$, we see that $\mathbf{v}^{(7)}$ is the same across the two sequences, as are the flow utility policy functions $u_i(\mathbf{v}^{(7)}, s^{(7)})$.

Thomas and Worrall (1990) base the intuition for this order-independence on the absence of wealth effects under CARA utility. But, as illustrated in Figures 7 and 8, it also depends critically on the efficiency of the promised utility process under iid types.⁶¹ In both figures, we see that the terminal promised utility vector and policy functions for flow utility end up *lower* when the low reports occur *earlier* in the sequence. Moreover, this effect is quantitatively larger in the high persistence case.⁶² This is driven by the following feature: if, after multiple consecutive low reports the next report is high, then the promised utility vector jumps back to E_2 by moving to the *northwest*. (See the left-hand panel of Figure 5 and the $\mathbf{v}^{(t)}$ time series in the right-hand panel of Figure 8.) That is, when $\mathbf{v}$ is far enough below E_2 , we have $w_{22}(\mathbf{v}, 1) > v_2$ while $w_{21}(\mathbf{v}, 1) < v_1$. Thus, promised utility for low endowment types continues to decrease even after the first high endowment report, therefore increasing the difference $w_{22}(\mathbf{v}, 1) - w_{21}(\mathbf{v}, 1)$. Intuitively, following a sequence of consecutive low reports (which leads to the incentive constraint being tightened), it is optimal to *re-introduce slack into* $[\mathbf{IC}_{21}^*]$ after the first high report. Thus, it matters whether a given high report is preceded by a high or low report.

I.2. Wedge Dynamics

We now turn to the dynamics of the flow utilities and the insurance properties of the optimal contract. A standard way to characterize allocative distortions in settings with non-transferable utility is via *wedges* that represent departures from the first-best allocation.

-
- (61) As shown in Strulovici (2020), this order-independence also holds under renegotiation-proof contracts when types are persistent in a continuous-type variant of the present model (see also Bloedel, Krishna and Strulovici 2020). We therefore deduce that it is the fact that the optimal contract is not renegotiation-proof under persistence, rather than the persistence of types per se, that leads to the order-dependence exhibited here.
- (62) In these examples, only the terminal values of $v_i^{(7)}$ and $u_i(\mathbf{v}^{(7)}, s^{(7)})$ for $i = 1$ are noticeably different across the two sequences. We have generated other (unreported) examples in which *all* of the terminal variables depend on the report order.

Figure 6: IID. Short-run dynamics along two report sequences with fixed initial condition $\mathbf{v}^{(0)} \in E_2$.

Figure 7: Medium persistence. Short-run dynamics along two report sequences with fixed initial condition $\mathbf{v}^{(0)} \in E_2$.

Figure 8: High persistence. Short-run dynamics along two report sequences with fixed initial condition $\mathbf{v}^{(0)} \in E_2$.

The *insurance wedge* at state $(\mathbf{v}, s)$ is

$$[\text{Ins-W}] \quad W_{\text{ins}}(\mathbf{v}, s) := \frac{U'(C(u_1(\mathbf{v}, s), 1) + \omega_1)}{U'(C(u_2(\mathbf{v}, s), 2) + \omega_2)} - 1 = \frac{u_1(\mathbf{v}, s)}{u_2(\mathbf{v}, s)} - 1$$

and the *intertemporal wedge* at state $(\mathbf{v}, s)$ and given report $i \in S$ is

$$[\text{Inter-W}] \quad W_{\text{inter}}(\mathbf{v}, s, i) := \frac{\mathbf{E}^{\mathbf{f}_i} [U'(C(u_j(\mathbf{w}_i, i), j) + \omega_j)]}{U'(C(u_i(\mathbf{v}, s), i) + \omega_i)} - 1 = \frac{\mathbf{E}^{\mathbf{f}_i} [u_j(\mathbf{w}_i, i)]}{u_i(\mathbf{v}, s)} - 1$$

where $\mathbf{w}_i := \mathbf{w}_i(\mathbf{v}, s)$ and the simplified expressions follow from properties of CARA utility. The insurance wedge measures departures from perfect *intratemporal* consumption smoothing, while the intertemporal wedge measures departures from the agent's Euler equation (given rate of return $R = 1/\alpha$). Both wedges are zero in the first-best allocation. The insurance wedge is positive if and only if there is *under-insurance*, namely, $u_2(\mathbf{v}, s) > u_1(\mathbf{v}, s)$. The intertemporal wedge is positive if and only if the optimal contract acts as an implicit tax on savings, so that $u_i(\mathbf{v}, s) > \mathbf{E}^{\mathbf{f}_i} [u_j(\mathbf{w}_i, i)]$, ie, flow utility has a “negative drift” conditional on $(\mathbf{v}, s, i)$.

The following result, illustrated in Figure 9, characterizes the iid benchmark.

Fact 5. *When types are iid, the insurance wedge is constant and strictly positive. The intertemporal wedge takes on (at most) two values, depending on whether the current report is high or low; it is always positive at low reports but may be negative at high reports.*

With iid types, there is *always* under-insurance because the agent only earns the iid information rent in $[\mathbf{IC}_{ij}^*]$. To understand the sign of the intertemporal wedge, recall that promised utility drifts up (to the northeast in Figures 3–5) after consecutive high reports, and down (to the southwest) after consecutive low reports. This suggests that, conditional on receiving a high endowment, the agent expects his consumption to have “positive drift” in the short run; conversely, receiving a low endowment leads to a “negative drift.” Figure 6 shows that this is indeed the case in our examples. Intuitively, the intertemporal wedge is negative at high endowment reports because the agent would like to *borrow* in the face of an upward-sloping consumption profile, and it is positive at low endowment reports because the agent wants to *save* in the face of a downward-sloping consumption profile.

With persistent types, these properties continue to hold qualitatively when the contract is efficient, i.e., after consecutive high reports.

Fact 6. *When types are persistent, the insurance wedge is constant and strictly positive after consecutive high reports. The intertemporal wedge is also constant (but not necessarily positive) after consecutive high reports.*

Figure 9: Wedge dynamics after 6 consecutive high reports (left) and 6 consecutive low reports (right). Initial $\mathbf{v}^{(0)} \in E_2$.

Fact 6 is borne out in the left-hand panel of Figure 9, which also illustrates that the wedges are quantitatively similar across all three persistence levels.

However, Figure 9 also highlights that the wedges behave very differently in the iid and persistent cases after consecutive low reports. When types are persistent, the intertemporal wedge *increases* along strings of consecutive low reports (instead of remaining constant as in the iid case), and increases *faster* when persistence is higher. The reasoning is the same as in the discussion of part (d) of Fact 3: optimal management of Markov information rents leads promised utility and consumption to have a *more negative* drift after low reports when persistence is higher. On the other hand, the insurance wedge *decreases* and quickly becomes *negative* after consecutive low reports. Thus, the contract *over-insures* the agent, and the wedges negatively co-vary, when types are persistent. Intuitively, compression of the Markov information rent at date t requires the principal to compress the difference of flow utilities $u_2 - u_1$ at date $t + 1$. When this motive is strong enough, it can be optimal to over-insure the agent.⁶³

J. Proofs for SA-I

This SA presents proofs of Facts 1–6 and derivations of various formulas stated in SA-I. It is a standing assumption throughout the present section that the agent has CARA utility: $U(c) = -e^{-\rho c}$ for some $\rho > 0$. Recall that $\theta_i := e^{-\rho \omega_i}$ so that $\psi(u, i, j) = (\theta_i / \theta_j)u$. Throughout, let ξ^* denote the optimal contract (which exists and is unique under [TVC]-regularity), and let $u_i(\mathbf{v}, s) := \xi^{*f}(\mathbf{v}, s, i)$ and $\mathbf{w}_i(\mathbf{v}, s) := \xi^{*c}(\mathbf{v}, s, i)$ denote the policy

(63) See [J.14] in SA-J.4 for a formal expression of this point.

functions.

J.1. Homogeneity Properties

We begin with homogeneity properties of the value function and optimal contract. These facts do not require any assumptions on $d \geq 2$ or the type process beyond Assumption [Markov](#).

Lemma J.1. Suppose the environment is [\[TVC\]](#)-regular. For all $a \in \mathbb{R}$ and $s \in S$:

(a) The principal's value function satisfies the homogeneity property

$$[\text{J.1}] \quad P(e^{-a}\mathbf{v}, s) = P(\mathbf{v}, s) + \frac{a}{1 - \alpha}$$

(b) For fixed $s \in S$, the optimal contract is homogenous of degree one in $\mathbf{v}$

$$[\text{J.2}] \quad \xi(e^{-a}\mathbf{v}, s) = e^{-a}\xi(\mathbf{v}, s)$$

Proof. Part (a) is immediate from the facts that (i) the return function $C(x) := U^{-1}(x)$ satisfies

$$[\text{J.3}] \quad C(e^{-a}x) = -\log(-e^{-a}x) = -\log(-x) - \log(e^{-a}) = C(x) + a$$

and (ii) that the graph of the constraint correspondence $\Gamma(\cdot)$ is a convex cone by part (c) of Lemma [G.2](#). Part (b) follows from part (a) and point (ii) in the previous sentence (recall that the optimal contract is unique under [\[TVC\]](#)-regularity). $\square$

We also have the following useful homogeneity property for the derivatives of the value function:

Lemma J.2. Suppose the environment is [\[TVC\]](#)-regular. For all $a \in \mathbb{R}$ and $s \in S$:

(a) The partial derivatives $P_i(\cdot, i)$ are homogenous of degree -1 , ie, for each $i \in S$,

$$[\text{J.4}] \quad P_i(e^{-a}\mathbf{v}, s) = e^a P_i(\mathbf{v}, s)$$

(b) The directional derivative $D_1 P(\cdot, s)$ is homogenous of degree -1 , ie,

$$[\text{J.5}] \quad D_1 P(e^{-a}\mathbf{v}, s) = e^a D_1 P(\mathbf{v}, s)$$

(c) For each $i \in S$, the efficiency set E_i (defined in [OA-C.3.1](#)) are rays in D .

(d) If types are iid, there exists a ray E in D such that $E = E_1 = \dots = E_d$.

Proof. For part (a), let $\varepsilon > 0$ and $i \in S$. Let $\hat{e}_i \in \mathbb{R}^d$ denote the i th standard basis vector. We have:

$$\begin{aligned}
P_i(e^{-a}\mathbf{v}, s) &= \lim_{\varepsilon \rightarrow 0} \frac{P(e^{-a}\mathbf{v} + \varepsilon \hat{e}_i, s) - P(e^{-a}\mathbf{v}, s)}{\varepsilon} \\
&= \lim_{\varepsilon \rightarrow 0} \frac{P(e^{-a}[\mathbf{v} + e^a \varepsilon \hat{e}_i], s) - P(e^{-a}\mathbf{v}, s)}{\varepsilon} \\
&= \lim_{\varepsilon \rightarrow 0} \frac{P(\mathbf{v} + e^a \varepsilon \hat{e}_i, s) - P(\mathbf{v}, s)}{\varepsilon} && \text{by part (a) of Lemma J.1} \\
&= e^a \cdot \lim_{\varepsilon \rightarrow 0} \left[\frac{P(\mathbf{v} + e^a \varepsilon \hat{e}_i, s) - P(\mathbf{v}, s)}{e^a \varepsilon} \right] = e^a P_1(\mathbf{v}, s)
\end{aligned}$$

which proves part (a). Part (b) is an immediate corollary of part (a). For part (c), note that $[\tilde{E}_i]$ together with part (a) imply that E_i must be a cone. That it must be a ray then follows from, eg, Lemma C.3. Part (d) follows immediately from part (c) and examination of the optimality conditions in OA-C.1. $\square$

Lemma J.3. Suppose the environment is [TVC]-regular, and that the type process is iid. Let $(\mathbf{v}^{(t)})$ denote the induced promise process induced by the optimal contract. Suppose that the initial $\mathbf{v}^{(0)} \in E$, where E is defined in part (d) of Lemma J.2. Then there exist numbers $b_i \in \mathbb{R}$ for each $i \in S$ such that

$$\mathbf{v}^{(t)} = \left(\prod_{i=1}^d b_i^{N_i^{(t)}} \right) \cdot \mathbf{v}^{(0)}$$

where $(N_i^{(t)})_{t=0}^\infty$ is the counting process defined by

$$N_i^{(t)} := \left| \{T \in \{1, \dots, t\} : s^{(T)} = i\} \right|$$

Proof. By Remark C.1, the optimal contract maps E to E , ie, $\mathbf{w}_i(\mathbf{v}, s) \in E$ for all $i \in S$ and $(\mathbf{v}, s) \in E \times S$. Given this, the lemma is an immediate consequence of part (b) of Lemma J.1 and part (d) of Lemma J.2. $\square$

J.2. Elementary Properties in $d = 2$ Case

Lemma J.4. Suppose the environment is [TVC]-regular, and that $d = 2$. If $\mathbf{v} \in E_s$, then $\mu_{21}(\mathbf{v}, s) > 0$ for each $s \in S$. Thus, if the type process is iid and $\mathbf{v} \in E$, then $\mu_{21}(\mathbf{v}, s) > 0$ for each $s \in S$.

Proof. The first statement is a corollary of Lemmas C.13 and C.14. The second statement follows from the first and part (d) of Lemma J.2. $\square$

Lemma J.5. Suppose the environment is [TVC]-regular, and that $d = 2$. If $\mathbf{v} \in E_s$, then $u_2(\mathbf{v}, s) > u_1(\mathbf{v}, s)$ for each $s \in S$. Thus, if the type process is iid and $\mathbf{v} \in E$, then $u_2(\mathbf{v}, s) > u_1(\mathbf{v}, s)$ for each $s \in S$.

Proof. In the $d = 2$ case, the optimality condition [FOC u_i] for $i = 2$ reduces to

$$C'(u_2(\mathbf{v}, s), 2) = \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} + \frac{\mu_{21}(\mathbf{v}, s)}{f_{s2}}$$

and [FOC u_i] for $i = 1$ reduces to

$$C'(u_1(\mathbf{v}, s), 1) = \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}} \cdot \frac{\theta_2}{\theta_1}$$

Since $\lambda_i(\mathbf{v}, s) = P_i(\mathbf{v}, s)$ by the envelope condition [Env $_i$], and since $\mathbf{v} \in E_s$, it follows from [E $_i$] that $\lambda_2(\mathbf{v}, s)/f_{s2} = \lambda_1(\mathbf{v}, s)/f_{s1}$. Thus, combining the above displays gives

$$C'(u_2(\mathbf{v}, s), 2) - C'(u_1(\mathbf{v}, s), 1) = \left(\frac{f_{s1}}{f_{s2}} + \frac{\theta_2}{\theta_1} \right) \cdot \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}}$$

Thus $C'(u_2(\mathbf{v}, s), 2) - C'(u_1(\mathbf{v}, s), 1) > 0$ by Lemma J.4. The first statement of the lemma follows because $C'(\cdot, i) = C'(\cdot) > 0$ on $\mathcal{U}$. The second statement follows from the first and part (d) of Lemma J.2. $\square$

For the next lemma, recall the multipliers $\eta_i(\mathbf{v})$ and $\sigma_{ij}(\mathbf{v})$ from the “interim” problem defined in OA-C.3.3 (just below [L $_i$]).

Lemma J.6. Suppose the environment is [TVC]-regular and $d = 2$. We have the following:

(a) $\mathbf{v} \in E_1$ if and only if

$$\text{[J.6]} \quad \eta_2(v_2) = \eta_1(\mathbf{v}) + \frac{\sigma_{21}(\mathbf{v})}{f_{12}}$$

(b) Fix $\mathbf{v} \in V_2$. At the optimum, we have

$$\text{[J.7]} \quad \eta_2(w_{12}(\mathbf{v}, s)) = \eta_1(\mathbf{w}_1(\mathbf{v}, s)) + \frac{\sigma_{21}(\mathbf{w}_1(\mathbf{v}, s))}{f_{12}} - \frac{(f_{11} - f_{21})}{f_{11}} \frac{\sigma_{21}(\mathbf{v})}{f_{12}}$$

(c) $\mathbf{v} \in E_2$ if and only if

$$\text{[J.8]} \quad \eta_2(v_2) = \eta_1(\mathbf{v}) + \frac{\sigma_{21}(\mathbf{v})}{f_{22}}$$

Proof. In the $d = 2$ case, [C.13] at $i = 2$ reduces to

$$\text{[J.9]} \quad \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} = \eta_2(\mathbf{v}) - \frac{f_{s1}}{f_{s2}} \sigma_{21}(\mathbf{v})$$

and eqreflambda-eta at $i = 1$ reduces to

$$[\text{J.10}] \quad \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} = \eta_1(\mathbf{v}) + \sigma_{21}(\mathbf{v})$$

Thus, by the envelope formulas $[\text{Env}_i]$ and the definition of $[\tilde{\mathbf{E}}_i]$, $\mathbf{v} \in E_s$ if and only if

$$\eta_2(\mathbf{v}) - \frac{f_{s1}}{f_{s2}} \sigma_{21}(\mathbf{v}) = \eta_1(\mathbf{v}) + \sigma_{21}(\mathbf{v})$$

Part (a) follows from rearranging the above display when $s = 1$. Part (c) follows from the same, when $s = 2$.

For part (b), first observe that a simple rearrangement of the optimality conditions $[\text{FOCw}_{ij}]$ for $i = 1$ delivers

$$[\text{J.11}] \quad \frac{P_1(\mathbf{w}_1(\mathbf{v}, s), 1)}{f_{11}} - \frac{P_2(\mathbf{w}_1(\mathbf{v}, s), 1)}{f_{12}} = \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}} \cdot \frac{f_{11} - f_{21}}{f_{11} \cdot f_{12}}$$

Subtracting $[\text{J.9}]$ from $[\text{J.10}]$, invoking the envelope formulas $[\text{Env}_i]$, and combining with the above display yields

$$[\text{J.12}] \quad \eta_1(\mathbf{v}) + \sigma_{21}(\mathbf{v}) - \eta_2(\mathbf{v}) + \frac{f_{s1}}{f_{s2}} \sigma_{21}(\mathbf{v}) = \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}} \cdot \frac{f_{11} - f_{21}}{f_{11} \cdot f_{12}}$$

Now, in the $d = 2$ case, $[\text{C.14}]$ reduces to $\sigma_{21}(\mathbf{v}) = \mu(\mathbf{v}, s)/f_{s1}$. (Thus, in particular, the weighted multiplier $\mu(\mathbf{v}, s)/f_{s1}$ is independent of s .) Plugging this in to the above display and rearranging yields part (b) of the lemma, completing the proof. $\square$

J.3. Proofs of Facts 1–6

We prove each of the facts in turn. (Recall that they pertain to the case $d = 2$, FOSD process, and $[\text{TVC}]$ -regular environment.)

Proof of Fact 1. This is an immediate corollary of part (b) of Lemma J.1. $\square$

We note that the properties in Fact 2 could be proved easily for any $d \geq 2$ by appealing to arguments in Thomas and Worrall (1990). We give a direct proof in the $d = 2$ case.

Proof of Fact 2. Part (a) follows from part (d) of Lemma J.2. Part (b) follows from Remark C.1.

Consider part (c). Because D is a convex cone in $\mathcal{U}^d$ (part (a) of Theorem B.1) and E is a ray in D (part (d) of Lemma J.2, it is easy to see that the ray E must be “upward sloping.” That is, for any $\mathbf{v}, \mathbf{v}' \in E$ such that $\mathbf{v} \neq \mathbf{v}'$, it must be that either $\mathbf{v} \ll \mathbf{v}'$ or

$\mathbf{v}' \ll \mathbf{v}$. By Lemmas C.6 and J.4 (noting that $E_2 = E$ by part (d) of Lemma J.2), we have $D_1 P(\mathbf{w}_2(\mathbf{v}, s), 2) > D_1 P(\mathbf{v}, s)$. It then follows from part (b) of Lemma J.2, the fact that E is “upward sloping,” and Lemma C.5 that $\mathbf{w}_2(\mathbf{v}, s) \gg \mathbf{v}$, as desired. Moreover, by the martingale property (Proposition 4.3), we have $D_1 P(\mathbf{w}_1(\mathbf{v}, s), 1) < D_1 P(\mathbf{v}, s)$. Remark C.1, together with the argument just used, then implies that $\mathbf{w}_1(\mathbf{v}, s) \ll \mathbf{v}$. This completes the proof of part (c). $\square$

Proof of Fact 3. Part (b) is a corollary of Lemma C.4.

Consider part (c). By Lemmas C.6 and J.4, we have $D_1 P(\mathbf{w}_2(\mathbf{v}, s), 2) > D_1 P(\mathbf{v}, s)$. It then follows from part (b) of Lemma J.2, the fact that E_2 is “upward sloping,” and Lemma C.5 that $\mathbf{w}_2(\mathbf{v}, s) \gg \mathbf{v}$, as desired.

Finally, consider parts (a) and (d). For part (a), let $\mathbf{v} \in E_2$. Part (c) of Lemma J.6, Lemma J.4, and $f_{22} > f_{12}$ imply that

$$\eta_2(v_2) < \eta_1(\mathbf{v}) + \frac{\sigma_{21}(\mathbf{v})}{f_{12}}$$

But, by [J.9] and [J.10] (together with the envelope formulas [Env_i]), this is equivalent to $P_1(\mathbf{v}, 1)/f_{11} > P_2(\mathbf{v}, 1)/f_{12}$. This implies that E_1 lies strictly above E_2 because strict convexity and continuous differentiability of P (Proposition 3.3) imply that the lower contour sets of $P(\cdot, 1)$ are strictly convex, and tangent to the vector $\mathbf{f}_1^\perp := (1/f_{11}, -1/f_{12})$ exactly at points $\mathbf{v} \in E_1$. This completes the proof of part (a). The proof of part (d) is similar, invoking part (b) of Lemma J.6. $\square$

Proof of Fact 4. This is an immediate corollary of Lemma J.3. $\square$

Proof of Fact 5. Constancy on the insurance is an immediate corollary of Lemma J.3, part (b) of Lemma J.1, and the formula in [Ins-W]. Similarly, that the intertemporal wedge takes on at most two values is an immediate corollary of Lemma J.3, part (b) of Lemma J.1, and the formula in [Inter-W]. Strict positivity of the insurance wedge follows from the formula in [Ins-W] and Lemma J.5. $\square$

Proof of Fact 6. The constancy properties follow from the formulas in [Ins-W] and [Inter-W], part (b) of Lemma J.1, Lemma C.4, and part (c) of Lemma J.2. The strict positivity of the insurance wedge then follows from J.5. $\square$

J.4. Additional Details for SA-I

This SA provides more precise intuition for the optimal contract’s short-run dynamics — and, in particular, the optimal management of the agent’s Markov information rents — by

considering the time-series behavior of the following *dual variables*: (i) the derivative $DP(\mathbf{v}, s) = (P_1(\mathbf{v}, s), P_2(\mathbf{v}, s))$ and (ii) the Lagrange multiplier $\mu_{21}(\mathbf{v}, s) \geq 0$ on $[\mathbf{IC}_{ij}]$ ($i = 2, j = 1$). On the optimal path, they satisfy

$$[\mathbf{J.13}] \quad \underbrace{\frac{P_1(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},1}} - \frac{P_2(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},2}}}_{\text{MC of compression}} = \underbrace{\mathbf{1}(s^{(t)} = 1) \cdot \frac{(f_{11} - f_{21})}{f_{s^{(t)},1} \cdot f_{s^{(t)},2}} \cdot \frac{\mu_{21}(\mathbf{v}^{(t-1)}, s^{(t-1)})}{f_{s^{(t-1)},s^{(t)}}}}_{\text{MB of compression}}$$

The MC of compression is the *Marginal Cost* of moving the promised utility vector toward the diagonal while holding *ex ante* promised utility $\mathbf{E}^{f_{s^{(t)}}}[\mathbf{v}^{(t)}]$ constant. The MB of compression is precisely the (shadow) *Marginal Benefit* of relaxing the preceding period's incentive constraint. Due to the cost smoothing motive, it is optimal for the principal to have promised utility drift away from E_1 gradually along strings of consecutive low reports. Holding the value of the weighted multiplier $\mu_{21}(\mathbf{v}^{(t-1)}, s^{(t-1)})/f_{s^{(t-1)},s^{(t)}}$ fixed,⁶⁴ the MB of compression indeed increases in persistence, suggesting that promised utility should indeed drift away from E_1 more quickly. While this ignores the indirect effects of increased persistence on the dual variables, this intuition is borne out in the numerical examples.

We also have:

$$[\mathbf{J.14}] \quad \underbrace{C'(u_2(\mathbf{v}^{(t)}, s^{(t)}), 2) - C'(u_1(\mathbf{v}^{(t)}, s^{(t)}), 1)}_{\text{direct MC of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})} = \underbrace{\left(\frac{f_{s^{(t)},1}}{f_{s^{(t)},2}} + \frac{\theta_2}{\theta_1} \right) \cdot \frac{\mu(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},1}}}_{\text{MB of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})} - \underbrace{\mathbf{1}(s^{(t)} = 1) \cdot \frac{(f_{11} - f_{21})}{f_{s^{(t)},1} \cdot f_{s^{(t)},2}} \cdot \frac{\mu_{21}(\mathbf{v}^{(t-1)}, s^{(t-1)})}{f_{s^{(t-1)},s^{(t)}}}}_{\text{indirect MC of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})}$$

Increasing the insurance wedge always leads to a direct MC (on the left-hand side of $[\mathbf{J.14}]$), which corresponds to the departure from first-best risk-sharing. For any degree of persistence, increasing the insurance wedge also adds slack to the incentive constraint $[\mathbf{IC}_{ij}]$ ($i = 2, j = 1$). In the iid case, these are the only two operative forces, and they lead to a positive insurance wedge or, equivalently, *under-insurance*. When types are persistent, the presence of Markov information rents introduces a countervailing force, represented by the indirect MC of increasing the insurance wedge on the right-hand side of $[\mathbf{J.14}]$. Notice that this term is identical to the MB of compression (of contingent promised utilities) in $[\mathbf{J.13}]$. Roughly speaking, one way to compress the difference $w_{12}(\mathbf{v}^{(t-1)}, s^{(t-1)}) - w_{11}(\mathbf{v}^{(t-1)}, s^{(t-1)})$ is to compress the difference

(64) We show in $[\mathbf{J.12}]$ that this weighted multiplier does not depend on the value of $s^{(t-1)}$.

$u_2(\mathbf{w}_1(\mathbf{v}^{(t-1)}, s^{(t-1)}), 1) - u_1(\mathbf{w}_1(\mathbf{v}^{(t-1)}, s^{(t-1)}), 1)$. That is, provide *more* intratemporal insurance — ie, *decrease* the insurance wedge — at date t , contingent on a date $t - 1$ low report (equivalently, $s^{(t)} = 1$). When this force is strong enough, as is the case after consecutive low reports, the optimal contract *over-insures* the agent by specifying flow allocations $u_1 > u_2$ and, equivalently, $W_{\text{ins}} < 0$. These features can be seen in Figures 7–8 and Figure 9, respectively. As we argued heuristically in SA-I.1, the indirect MC term in J.14 suggests that these effects should be larger when persistence is higher. Indeed, this is precisely what we see in the computed numerical examples.

To conclude, we provide the simple derivations of [J.13] and [J.14]. (We focus on the $d = 2$ case, as in the text. But by following the same calculations, the reader may easily verify that each of these equations and the lemma generalize naturally to the $d \geq 2$ case when the type process satisfies MLRP.)

Consider [J.13]. First, take [FOC $\mathbf{w}_{ij}$] ($i = 1$) and divide through by f_{1j} and f_{s1} to get

$$\frac{P_j(\mathbf{w}_1(\mathbf{v}, s), 1)}{f_{1j}} = \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{f_{2j}}{f_{1j} \cdot f_{s1}} \mu_{21}(\mathbf{v}, s)$$

Then, using the above display, we have

$$\frac{P_1(\mathbf{w}_1(\mathbf{v}, s), 1)}{f_{11}} - \frac{P_2(\mathbf{w}_1(\mathbf{v}, s), 1)}{f_{12}} = \frac{1}{f_{s1}} \cdot \mu_2(\mathbf{v}, s) \cdot \left[\frac{f_{22}}{f_{12}} - \frac{f_{21}}{f_{11}} \right]$$

Cross-multiplying the term in brackets, recalling from Lemma C.4 that

$$\frac{P_1(\mathbf{w}_2(\mathbf{v}, s), 2)}{f_{21}} - \frac{P_2(\mathbf{w}_2(\mathbf{v}, s), 2)}{f_{22}} = 0$$

and appropriately arranging time indices delivers the expression in [J.13].

Consider [J.14]. First, take [FOC u_i] ($i = 1$) and divide through by f_{s1} to get

$$C'(u_1(\mathbf{v}, s), 1) = \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{\theta_2}{\theta_1} \cdot \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}}$$

Then, take [FOC u_i] ($i = 2$) and divide through by f_{s2} to get

$$C'(u_2(\mathbf{v}, s), 2) = \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} + \frac{\mu_{21}(\mathbf{v}, s)}{f_{s2}}$$

Combining these displays, and using the envelope formulas [Env $_i$] to set $\lambda_i(\mathbf{v}, s) = P_i(\mathbf{v}, s)$, gives

$$C'(u_2(\mathbf{v}, s), 2) - C'(u_1(\mathbf{v}, s), 1) = \left(\frac{f_{s1}}{f_{s2}} + \frac{\theta_2}{\theta_1} \right) \cdot \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}} + \frac{P_2(\mathbf{v}, s)}{f_{s2}} - \frac{P_1(\mathbf{v}, s)}{f_{s1}}$$

[J.14] then follows from substituting [J.13] into the above display and arranging time indices appropriately.