

Insurance and Inequality with Persistent Private Information*

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

7TH SEPTEMBER 2018

Abstract

We study optimal insurance contracts for an agent with Markovian private information. Our main results characterize the implications of constrained efficiency for long-run welfare and inequality. Under minimal technical conditions, there is *Absolute Immiseration*: in the long run, the agent's consumption and utility converge to their lower bounds. When types are persistent and utility is unbounded below, there is *Relative Immiseration*: low-type agents are immiserated at a faster rate than high-type agents, and “pathwise welfare inequality” grows without bound. These results extend and substantially generalize the hallmark findings from the classic literature with iid types, suggesting that the underlying forces are robust to a broad class of private information processes. The proofs rely on novel recursive techniques and martingale arguments. When the agent has CARA utility, we also analytically and numerically characterize the short-run properties of the optimal contract. Persistence gives rise to qualitatively novel short-run dynamics and allocative distortions (or “wedges”) and, quantitatively, induces less efficient risk-sharing. We compare properties of the wedges to their counterparts in the dynamic taxation literature.

Keywords: Absolute immiseration; relative immiseration; dynamic contracting; recursive contracts; principal-agent problem; persistent private information.

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

(*) This paper was previously circulated as “Misery, Persistence, and Growth” by Bloedel and Krishna. We thank Gabriel Carroll, V.V. Chari, Sebastian Di Tella, Ed Green, Pablo Kurlat, Paul Milgrom, Shunya Noda, Ilya Segal, Andy Skrzypacz, and seminar participants at Stanford for useful comments and conversations. We also thank the co-editor, Giuseppe Moscarini, and three anonymous referees for detailed suggestions that greatly improved the paper.

(†) Stanford University. Email: abloedel@stanford.edu

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

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

1. Introduction

Problems of insurance and distribution are inextricably linked. Financial markets, tax systems, and social insurance programs all serve to facilitate risk-sharing and thereby protect against the many ups and downs of economic life, such as shocks to earnings, spells of unemployment, and unexpected changes in health and productivity. At the same time, these institutions are at the center of ongoing debates over the sources and consequences of growing income and wealth inequality in developed economies. For instance, a recent Pew Research Center survey¹ shows that Europeans and Americans view inequality as among the world’s greatest dangers. More broadly, the tradeoff between insurance and inequality occupies a central role in economic thought: as Lucas (1992) puts it, “...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.”

We revisit this tradeoff by building on the influential line of work that takes a mechanism design approach to study *optimal* insurance arrangements in the presence of private information. The main distributional findings of the classic studies are striking: absent participation constraints, the insured agents become completely impoverished (Thomas and Worrall (1990)) and cross-sectional inequality increases without bound (Atkeson and Lucas (1992)). Thus, *at the optimum and in the long run*, there is effectively *no tradeoff* between insurance and inequality. These *immiseration* results are “often regarded as being the hallmark result[s] of dynamic social contracting in the presence of private information” (Kocherlakota (2010, p. 70)). Due to their extreme and perhaps counter-intuitive nature, they have also generated a substantial literature aimed at understanding the robustness of the underlying mechanisms.

While the classic literature focuses exclusively on the special case of iid types, the data suggest that the relevant risks are not only inherently *dynamic*, but also *highly persistent*.² Recent theoretical work has emphasized the importance of persistent private information for *short-run* distortions in optimal mechanisms (see, eg, Farhi and Werning (2013) and Golosov, Troshkin and Tsyvinski (2016)). But surprisingly little is known about its role in shaping *long-run* outcomes. Given the fundamental nature of the underlying tradeoffs, it is important to understand the long-run distributional implications of constrained efficiency in more realistic and general settings involving persistent private information. Are the classic immiseration results, and the mechanisms that underlie them, robust? A primary goal of this paper is to answer this question.

To that end, we study a model of dynamic insurance in which the agent’s privately observed type evolves according to a fully connected Markov chain on a finite state space

(1) See <http://www.pewglobal.org/2014/10/16/middle-easterners-see-religious-and-ethnic-hat-red-as-top-global-threat/>.

(2) For instance, Storesletten, Telmer and Yaron (2004a) find that labor earnings approximately follow a random walk. Storesletten, Telmer and Yaron (2004b) and Meghir and Pistaferri (2004) emphasize the time-varying risk of labor income, which is inconsistent with the iid model.

(which need not exhibit positive serial correlation). A risk-neutral principal designs an infinite-horizon insurance contract for a risk-averse agent with the goal of minimizing costs. Both parties fully commit to the contract at the initial date, and the agent cannot save or borrow outside of the contract. Following the seminal work of Thomas and Worrall (1990), our baseline model interprets the agent’s private types as shocks to his endowment.³

In this context, we make four contributions. Our first contribution, Theorem 3, shows that, under minimal technical conditions, the optimal contract leads to *Absolute Immiseration*: in the long run, the agent’s consumption and utility converge to their lower bounds, so he becomes impoverished in *absolute* terms. The primary contribution of this result is its *generality*. Since the classic studies — and despite longstanding theoretical interest in the implications of persistence, going back to Fernandes and Phelan (2000) — there has been very little progress on long-run convergence results outside of the iid setting. The generality of Theorem 3 is thus noteworthy from a *technical* perspective. To our knowledge, it is the first result of its kind (even outside the insurance literature) to be established in a generic setting, where the only substantive assumptions on the type process are the Markov property and finiteness of the type space.

Perhaps more importantly, the generality of Theorem 3 is also useful from a *conceptual* perspective. The proof — which builds on martingale convergence ideas pioneered by Thomas and Worrall (1990) — relies on only very basic properties of the contracting environment and type process, helping us identify the driving forces of Absolute Immiseration in a “non-parametric” way. While the original insight of Absolute Immiseration is by now a textbook topic, the recent literature — namely, the work of Zhang (2009) and Williams (2011) — has generated some puzzling results and raised questions about the fragility of the underlying mechanisms. By delineating a wide range of environments in which Absolute Immiseration holds and identifying the robustness of its underlying forces, Theorem 3 takes a substantive step toward explaining when and why it may fail. In particular, we suggest that the asymptotic behavior of “impulse response functions” (Pavan, Segal and Toikka (2014)) is the key property of the type process, and that many other details of the contracting environment are irrelevant (see Section 6.2).

Our second contribution, Theorem 4, shows that, when the type process exhibits positive serial correlation (ie, is *persistent*) and utility is unbounded below, the optimal contract also induces *Relative Immiseration*: low-type agents are immiserated at a faster rate than high-type agents, so that low types become impoverished in *relative* terms and “pathwise welfare inequality” grows without bound. Imagine two agents, A and B , who have observed the same sequence of realized endowments up through period $t - 1$. In period t , agent A receives a higher endowment than agent B . How much better off is A than B going forward, as measured by

(3) This is not essential. We discuss in Section 6 how our results extend to other insurance settings with different sources of private information, such as those considered in the dynamic taxation literature where the agent has private information concerning his productivity.

the difference of their continuation utilities in the optimal contract? Theorem 4 states that, as $t \rightarrow \infty$, the impact of this last endowment shock on continuation utility grows without bound. The familiar intuition behind Absolute Immiseration is that, due to risk aversion, it is less expensive for the principal to make utility *vary* in response to the agent’s report (ie, provide incentives) when the average *level* of consumption is lower; it is thus optimal for consumption to drift down over time. The intuition behind Relative Immiseration pushes this logic further: simply put, the principal does not waste the chance to provide high-powered incentives as they become affordable. Even if the consumption process converges, it retains enough noise to make the variance of utility explode in the limit. This notion of Relative Immiseration is novel to our analysis and, as we discuss in Section 5.3, differs in important ways from extant results on unbounded long-run inequality (as in Atkeson and Lucas (1992)).

Our third contribution is methodological. Along the way to proving Theorems 3 and 4, we develop new recursive techniques and martingale arguments that can be applied to a range of contracting problems. First, we begin our analysis with a recursive formulation of the contracting problem. The recursive approach with Markovian types dates back to Fernandes and Phelan (2000), but the added complexity created by persistence has long been identified as a major obstacle for obtaining substantive economic results. We make new analytical headway by using a slightly different state variable than Fernandes and Phelan (2000), which consists of a vector of type-contingent (or *interim*) promised utilities. This change yields both conceptual and technical dividends.⁴ Conceptually, the interim approach allows us to keep track only of *on-path* quantities, and is key to formalizing our notion of Relative Immiseration (see Section 5.3). Technically, we show how to characterize structural properties of implementable and optimal contracts at a much higher level of generality than existing work.⁵ For instance, our Theorem 1 characterizes the recursive domain⁶ — an essential piece of the problem solution — for any finite number of types and a broad class of utility functions and type processes; in an important set of special cases, we even provide a closed form solution for the domain. Aside from being a central component in the proofs of Theorems 3 and 4, this result is also essential for understanding the underlying economics of the incentive problem, and closed form solutions for the domain can substantially simplify the task of obtaining numerical solutions for optimal contracts.

Second, the martingale-based proof of Theorem 3 builds on this recursive formulation. While we have emphasized the conceptual simplicity afforded by the martingale approach, the proof itself requires new, and occasionally subtle, arguments. The main new difficulties

-
- (4) To be clear, we do *not* suggest that it is impossible to derive our results using the Fernandes and Phelan (2000) formulation. Indeed, it must be possible, as there is a one-to-one mapping between the two formalisms. Rather, we emphasize that our approach makes important aspects of the problem more transparent, and certain aspects of the analysis more straightforward. For specific examples, see footnotes 35 and 58.
 - (5) The literature following Fernandes and Phelan (2000) either (i) does not develop theoretical results beyond an abstract recursive representation, or (ii) restricts attention to the special case of binary types. We discuss this line of work, and the complementary “first-order approach,” below in Section 1.1.
 - (6) That is, set of implementable promised utility vectors.

arise from the multi-dimensionality of the state space in our recursive formulation — a necessary economic consequence of incentive compatibility with persistent types, not merely a mathematical artifact. We show how to overcome these issues by focusing on convergence of dual variables in the principal’s problem, which satisfy a kind of “renewal” property implied by ergodicity of the underlying Markov process. These ideas can be applied to many other contracting settings where long-run convergence is of interest.

Our fourth, and final, contribution is a detailed study of the short-run dynamics and allocative distortions (or *wedges*) induced by the optimal contract in Section 7. Focusing on the case in which the agent has CARA utility and persistent, binary types, we provide a detailed description of the optimal contract through a combination of analytical and numerical characterizations. Persistence, by endowing the agent with an additional source of information rents, induces qualitatively novel dynamics and distortions in the promised utility process. An important *order-independence* property from the iid case is overturned: with persistence, the agent cares *when* low endowment shocks occur, and is worse off when they occur earlier in the contracting relationship. Persistence also induces qualitatively novel dynamic behavior of the wedges, and certain monotonicity properties of the iid solution are overturned. Quantitatively, the wedges tend to be larger when types are more persistent, suggesting that persistence leads to less efficient risk sharing. We discuss how our findings in the pure insurance setting compare to recent results in the dynamic taxation literature. In short, both settings exhibit similar promised utility dynamics, but the wedges differ in important and subtle ways.

The rest of the paper is organized as follows. After discussing the related literature, we lay out the baseline model in Section 2 and formulate the recursive contracting problem in Section 3. Section 4 presents our main structural results on implementable and optimal contracts. Section 5 presents our main results on Absolute and Relative Immiseration. In Section 6, we discuss how these results can be further generalized and their importance for the literature. Section 7 presents our analysis of short-run dynamics and allocative distortions in the special case of CARA utility. Finally, Section 8 concludes. All proofs are contained in the appendices.

1.1. Related Literature

Insurance and immiseration: We build on the classic literature that studies dynamic insurance contracts when the agent has iid private information. Green (1987) and Thomas and Worrall (1990) develop the recursive approach to dynamic screening in this context using the agent’s *ex ante* promised utility as a state variable, and show that the optimal contract leads to (absolute) immiseration.⁷ Atkeson and Lucas (1992) study a “general equilibrium” version of this contracting problem with date-by-date resource constraints and show that the cross-sectional distribution of promised utility fans out over time, so that inequality increases without bound. While Green (1987) and Atkeson and Lucas (1992) study special cases in

(7) Spear and Srivastava (1987) introduce recursive methods in a closely related moral hazard setting.

which (nearly) closed form solutions are available, we follow the approach of Thomas and Worrall (1990) by developing martingale arguments that apply much more generally.

Given the extreme nature of the immiseration results, a number of papers have studied their robustness under different modeling assumptions.⁸ When the agent faces binding participation constraints, Atkeson and Lucas (1995) and Phelan (1995) show that there exists a non-degenerate stationary distribution for promised utility, so that inequality remains bounded in the long run.⁹ Farhi and Werning (2007) and Phelan (2006) obtain similar results using models without participation constraints, but in which the principal is more patient than the agent. In a dynastic economy with endogenous fertility, Hosseini, Jones and Shourideh (2013) show that there is not immiseration in consumption but that, under certain conditions, there is immiseration in family size — ie, family size may converge to zero in the long run. In a production economy with endogenous growth, Khan and Ravikumar (2001) show that inequality increases without bound, though consumption increases over time as the economy grows.

Each of these papers (which all assume iid types) identifies a novel economic force that renders immiseration suboptimal. From a more technical perspective, Phelan (1998) argues that immiseration hinges on details of the agent’s utility function, namely, that it needn’t occur with probability one if the agent’s (positive) marginal utility of consumption is bounded away from zero (which is ruled out under standard Inada conditions that we assume).¹⁰ The aforementioned recent studies of Zhang (2009) and Williams (2011) suggest that immiseration may also be fragile in the face of persistent private information. We discuss these two papers, along with the related work of Strulovici (2011), in detail in Section 6.2.

Recursive contracts with Markovian types: There are three key precedents to the recursive approach taken here. As already mentioned, Fernandes and Phelan (2000) are the first to use recursive methods when types are persistent. Their state variable consists of promised utility and a vector of *threat-point* utilities, all of which are *ex ante* quantities, in that they do not condition on the agent’s current report. Threat-point utility is the continuation utility promised to the agent if he had lied about his previous type, and as such never arises on-path. In contrast, we use a vector of *interim* promised utilities that condition on the agent’s current report, all of which arise on-path. Doepke and Townsend (2006) extend the recursive approach to more general settings with both hidden information and hidden actions, where the interaction between agency frictions can lead to a severe curse of dimensionality. Their main innovation is

-
- (8) The broader literature on recursive contracts generated by those early papers is too vast to do justice to here. See Chapters 20-21 of Ljungqvist and Sargent (2012) for a textbook treatment of the classic models and their importance in macroeconomics, and Golosov, Tsyvinski and Werquin (2016) for a more recent survey that discusses additional applications.
 - (9) Thomas and Worrall (1988) and Kocherlakota (1996), among others, study risk-sharing with participation constraints (interpreted as limited commitment) under *symmetric* information. The tradeoffs and resulting consumption dynamics in such models are very different.
 - (10) Instead, promised utility and consumption almost surely converge *either* to their lower bounds *or* to their upper bounds — ie, *polarization* occurs with probability one.

showing how to maintain computational efficiency by constructing bounds on off-path utility gains. Finally, Zhang (2009) extends the ideas of Fernandes and Phelan (2000) to a continuous time setting where the agent’s private information follows a finite-state Markov jump process, which is the natural continuous time limit of our framework.¹¹

A few very recent papers, either concurrent with or subsequent to ours, use recursive methods to characterize optimal contracts in principal-agent settings with *risk neutral* agents and *limited transfers*.¹² The contemporaneous work of Fu and Krishna (2017) uses the same techniques as the present paper to study a cash-flow diversion model of firm financing. Subsequently, Krasikov and Lamba (2018) use the same techniques to study a very closely related screening model of repeated procurement.¹³ Both papers allow monetary transfers but impose a limited liability constraint, which implies that the optimal contract converges to the first-best (in finite time), as at some point it is optimal for the principal to “sell the firm” to the agent. Guo and Hörner (2018), in independent and contemporaneous work, study a dynamic allocation problem without transfers that is, in some sense, the appropriate risk-neutral analogue of our baseline insurance model. In both papers, the principal aims to maximize efficiency and controls *only* the agent’s consumption in each period, so that (i) all incentives must be provided dynamically through allocative distortions, and (ii) low-value types want to imitate high-value types, who receive larger flow allocations.¹⁴ While they use the same recursive representation, the agent’s risk-neutrality and the indivisibility of the underlying good results in very different results and leads Guo and Hörner (2018) to rely on different mathematical arguments. Consistent with Phelan (1998), their optimal contract leads to polarization in the long run, and they prove this through a detailed construction of the optimal contract instead of using martingale techniques.

-
- (11) Casting things directly in continuous time adds tractability, but also requires the Markov process to be persistent and forces Zhang to treat components of his promised utility state variable asymmetrically (as “persistent” vs. “transitional” utilities). Our discrete time model nests the case of iid and negatively correlated states, and allows us to treat components of our state variable symmetrically, which we find conceptually clearer.
 - (12) By contrast, the dynamic mechanism design literature almost universally assumes perfectly transferable utility. Monetary transfers can be used to (i) transfer promised utility across different dates without distorting the allocation, and (ii) extract value before information becomes private. Thus, without exogenous limitations on transfers or risk aversion, the agent typically only earns rents for his *initial* private information (Esö and Szentes (2017)), and distortions often vanish in the long run (Battaglini (2005)). From a technical perspective, transferable utility in conjunction with the first-order approach allows one to substitute out transfers and move the agent’s promised utility directly into the principal’s objective function, which obviates the need for recursive methods. This substitution step typically cannot be carried out with restricted transfers or risk aversion — see, eg, Bergemann and Välimäki (2018) for a discussion.
 - (13) Fu and Krishna (2017) is essentially a Markovian version of Clementi and Hopenhayn (2006), while Krasikov and Lamba (2018) is essentially a Markovian version of Krishna, Lopomo and Taylor (2013).
 - (14) This is *opposite* the pattern of binding constraints in settings with transfers (perhaps subject to limited liability), where high types want to imitate low types to prevent some of their information rents from being extracted. (In our model, “high value” agents are those with low endowments which, due to risk aversion, means that their marginal utility is high.)

Importantly, all of the above papers either (i) do not develop theoretical results beyond an abstract recursive representation, or (ii) restrict attention to the special case of *persistent* and *binary* types.¹⁵ To the best of our knowledge, our paper is the first in this line of work to derive substantive results — either structural (Theorems 1–2) or economic (Theorems 3–4) — in a general Markovian setting, and this added generality is our primary methodological contribution.

A complementary line of work develops the first-order approach (FOA) in dynamic settings, facilitating the study of models with a continuum of types. Pavan, Segal and Toikka (2014) develop the FOA in a general mechanism design setting, potentially with multiple agents and non-stationary private information processes, while Kapička (2013) (in discrete time) and Williams (2011) (in continuous time) focus on single-agent insurance settings with Markovian types. While all three papers provide sufficient conditions under which the FOA is valid, it is known to be fragile. For example, Battaglini and Lamba (2018) show that contracts obtained via the FOA typically violate global incentive constraints even in a simple quasi-linear setting of monopolistic screening. This is likely to be true in our setting (and other non-quasi-linear ones like it) as well, where incentive constraints are typically even more difficult to deal with and the FOA more difficult to verify. Our analysis partially overcomes this fragility by directly incorporating all downward incentive constraints (see Assumption NHB and the discussion in Section 6.1).

Optimal dynamic taxation: Among the many applications of the FOA, the most relevant to our work are Farhi and Werning (2013) and Golosov, Troshkin and Tsyvinski (2016), who provide detailed analyses of the short-run allocative distortions — in particular, the “labor wedge” — that arise under optimal dynamic tax schemes. While they emphasize different aspects of the problem — Farhi and Werning (2013) focus on time-series properties, while Golosov, Troshkin and Tsyvinski (2016) focus on cross-sectional properties, namely, the shape of distortions as a function of type within a period — both papers find that the distortions depend critically on the autocorrelation structure of the type process.¹⁶ As we discuss in Section 7, despite the different settings, some of our results about the short-run dynamics of the insurance wedge mirror findings about mean-reversion of the labour wedge in Farhi and Werning (2013). But the nature of the contributions are fundamentally different, as their focus on continuously-distributed types and use of the FOA makes their approach more amenable

(15) Fernandes and Phelan (2000) and Doepke and Townsend (2006) follow route (i) and numerically compute solutions in particular examples. Zhang (2009), Fu and Krishna (2017), Krasikov and Lamba (2018), and Guo and Hörner (2018) all follow route (ii). Two less related papers, Broer, Kapička and Klein (2017) and Halac and Yared (2014), also restrict attention to binary types. Broer, Kapička and Klein (2017), in concurrent work, use a slightly different recursive approach to study insurance with limited enforcement constraints, but focus on numerical experiments. Halac and Yared (2014) study a dynamic delegation problem using the recursive formulation of Fernandes and Phelan (2000).

(16) See also Albanesi and Sleet (2006) for a detailed analysis in the iid case, and Kocherlakota (2010) for an excellent overview of the optimal dynamic taxation literature.

to clean closed-form solutions, while our primary focus is on general and “non-parametric” results.

Moreover, the nature of intertemporal distortions is quite different across most models of dynamic taxation and our pure insurance setting. Golosov, Kocherlakota and Tsyvinski (2003) emphasize the generality of the inverse Euler equation and corresponding positive “intertemporal wedge” in the workhorse model with “separable” utility, in which the agent’s marginal utility of consumption is independent of his private information. In that case, consumption utilities serve as “type-independent numeraire” that the principal can use to transfer value across periods. By contrast, the agent’s marginal utility is private information in our pure insurance setting. As we discuss in Section 6.2, the appropriate martingale in our setting generalizes the inverse Euler equation, and these separabilities are not important for long-run convergence results. These differences do, however, matter for the properties short-run intertemporal distortions, as we discuss in Sections 7.2 and 7.3.

Robust long-run predictions: At a thematic level, our work connects to at least two other papers that emphasize the conceptual importance of studying the *robustness* of the *long-run* behavior of optimal contracts, albeit in rather different settings. First, in concurrent work, Garrett, Pavan and Toikka (2018) identify long-run properties of allocative distortions in a monopolistic screening setting that are “robust,” in the sense that they are valid even when the FOA fails, and discuss how these properties depend on features of the type process. Our Theorems 3 and 4 are established at a similar level of generality to most of their results, and also do not rely on the FOA (again, see Section 6.1). From a technical perspective, both papers establish results concerning convergence *in probability*, but this similarity is superficial as the underlying arguments are completely different (see footnote 55 in Section 5.2). Pavan (2016) points to these kinds of results as an important open direction for dynamic mechanism design more broadly.

Second, Albanesi and Armenter (2012) study the determinants of long-run intertemporal distortions in a broad class of second-best economies, including models of constrained-optimal risk-sharing with Markovian private information. They emphasize the importance of (i) “permanent” intertemporal distortions, whereby the agent’s Euler equation is distorted in the same direction at each history (as is the case when the Inverse Euler Equation holds), and (ii) a unified sufficient condition, the “front-loading principle,” that rules out permanent intertemporal distortions in the second-best. Absolute Immiseration in Theorem 3, and the martingale convergence arguments that underlie it, are closely related to a generalized version of their front-loading principle (see our Sections 5.2 and 6.2, and their Section 5.2.2). But the contributions are fundamentally different: while we *prove* long-run convergence, they *assume* it and study properties of the limit. Moreover, we emphasize in Section 7 that permanent intertemporal distortions in the sense of Albanesi and Armenter (2012) only arise in special classes of private information models. Indeed, they fail to arise even in the simplest special

cases of our pure insurance setting.

2. Baseline Model

2.1. Environment

A single risk-averse agent (he) faces an uncertain endowment stream, and a risk-neutral principal (she) is prepared to provide insurance. Time is discrete and infinite, indexed by $t \in \mathbb{N}$.¹⁷ We begin by describing the primitives of the environment. Assumptions [DARA](#), [NHB](#), and [Markov](#), stated below, *hold for the remainder of paper* with the exception of Section [6.1](#), where we discuss in more detail their importance for our results and how they can be relaxed.

Preferences: Both the principal and agent discount the future at common rate $\alpha \in (0, 1)$. The agent has utility function $U : \mathcal{C} \rightarrow \mathbb{R}$ over consumption, where the domain of consumption $\mathcal{C} \subseteq \mathbb{R}$, $\bar{c} := \sup \mathcal{C} = +\infty$, and $\underline{c} := \inf \mathcal{C}$ may be finite or infinite. Let $\mathcal{U} := U(\mathcal{C}) \subseteq \mathbb{R}$ denote the range of feasible utilities. We require 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 on the interior of \mathcal{C} , and satisfies the Inada conditions $\lim_{c \rightarrow \underline{c}} U'(c) = +\infty$ and $\lim_{c \rightarrow \bar{c}} U'(c) = 0$;
- (b) It is bounded above and unbounded below. In particular, $\lim_{c \rightarrow \bar{c}} U(c) = 0$ and $\lim_{c \rightarrow \underline{c}} U(c) = -\infty$. Thus, $\mathcal{U} = \mathbb{R}_{--}$;
- (c) It has *decreasing absolute risk aversion* (DARA). In particular, the mapping $c \mapsto -\log(U'(c))$ is (weakly) concave.¹⁸

Assumption [DARA](#), which is fairly weak, appears also in Thomas and Worrall (1990) and serves to simplify the analysis. In particular, 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. Assumption [DARA](#) is satisfied, for example, when $U(\cdot)$ is of the CARA or CRRA class.

Information: The agent receives an endowment of $\omega_i \in \mathbb{R}$ in each period, 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 good is perishable, so the agent cannot save his endowment for consumption in later periods. The principal does not observe these endowment shocks and must rely on the agent's reports.

(17) We adopt the convention that $0 \in \mathbb{N}$.

(18) When $U(\cdot)$ is twice differentiable, this definition is equivalent to the coefficient of absolute risk aversion.

Assumption 2 (NHB). There is *No Hidden Borrowing*. Thus, the agent may *not* over-state his endowment in any period.

Assumption **NHB** is motivated by the ideas that (i) endowments are partially verifiable and (ii) the agent does not have access to a market outside of his relationship with the principal. For example, before receiving any transfers from the principal, the agent might be required to deposit some fraction of his endowment in an account that the principal can monitor. If the agent is not able to borrow units of the consumption good without the principal knowing — ie, if there is No Hidden Borrowing — then he can deposit at most his true endowment.

Type Process: We make one substantive assumption on $(\omega^{(t)})_{t \in \mathbb{N}}$, the (stochastic) *endowment process*.¹⁹

Assumption 3 (Markov). The agent’s types follow a first-order Markov process with transition probabilities

$$\mathbf{P}(\omega^{(t+1)} = j \mid \omega^{(t)} = i) = f_{ij}$$

and the Markov process is *fully connected*. That is, the transition probabilities satisfy $f_{ij} > 0$ for all $i, j \in S$.

The transition probabilities may be represented as a $d \times d$ transition matrix with rows $\{\mathbf{f}_i\}_{i=1}^d$, where \mathbf{f}_i denotes the distribution over tomorrow’s states if today’s state is $i \in S$. Note that Assumption **Markov** does not place any substantive restrictions on the serial correlation properties of the type process. We will often assume that types are *persistent* in one of the following senses.

Definition 2.1. The Markov process satisfies:

- (a) *FOSD* if \mathbf{f}_i first-order stochastically dominates \mathbf{f}_j whenever $i > j$.
- (b) *MLRP* if the transition probabilities are non-decreasing in the monotone likelihood ratio order, ie, if the ratio f_{ki}/f_{kj} is non-decreasing in k whenever $i > j$.
- (c) The *pseudo-renewal* property if there exists a probability distribution $\pi \in \Delta(S)$ such that $f_{ij} = \pi_j$ whenever $i \neq j$.²⁰
- (d) *UPR* (uniform pseudo-renewal) if it satisfies the pseudo-renewal property and $(f_{ii} - \pi_i)(f_{jj} - \pi_j) \geq 0$ for all $i, j \in S$.
- (e) *PPR* (persistent pseudo-renewal) if it satisfies FOSD and the pseudo-renewal property.

It is easy to see that both MLRP and PPR imply FOSD, and that all three nest the case of iid types. It is also easy to see that PPR implies UPR. When $d = 2$, every Markov chain satisfies UPR. It is also easy to see that MLRP, PPR, and FOSD are all equivalent when $d = 2$.

(19) Throughout, we use the notation $(x^{(t)})_{t \in \mathbb{N}}$ or, more simply $(x^{(t)})$, to denote stochastic processes.

(20) This definition is from Hörner, Mu and Vielle (2017). To our knowledge, it first appeared as “Condition A” in Renault, Solan and Vielle (2013).

Thus, each of these three conditions is a strict generalization of the persistence conditions assumed in models with binary types.²¹ One interpretation of PPR is that the discrete time process is generated by sampling a continuous-time process; shocks occur at random dates in continuous time and, when they do, the type is re-drawn according to a fixed distribution. MLRP is a well-understood notion of positive serial correlation and, though it is stronger than FOSD, it is satisfied by many parametric classes of distributions used in applications.²²

2.2. The Contracting Problem

At the initial date, $t = 0$, the principal offers a long-term insurance contract to the agent. 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. We study constrained-efficient risk-sharing schemes. The principal’s objective is to minimize costs,²³ given some (possibly degenerate) prior belief over the agent’s initial type, and subject to delivering a pre-specified schedule of lifetime utilities to the agent and providing appropriate incentives. In particular, at time $t = 0$ the principal promises an agent with initial type $i \in S$ exactly $v_i \in \mathcal{U}$ lifetime utilities, summarized by the vector of *contingent* promised utilities $\mathbf{v} := (v_1, \dots, v_d)$.

Intuitively, a *sequential contract* specifies consumption allocations (transfers of the consumption good from the principal to the agent, or vice versa) conditional on histories of previous allocations and messages from the agent. By the Revelation Principle, it is without loss to consider direct revelation sequential contracts. Any sequential contract offered by the principal must (i) give exactly v_i lifetime utilities to an agent with initial type $i \in S$ (the *promise keeping* constraint) and (ii) make truthful reporting of endowment shocks an optimal strategy for the agent in the induced decision problem (the *incentive compatibility* constraint).

We refer to the principal’s problem of choosing a sequential contract to minimize costs subject to promise keeping and incentive compatibility the *sequential problem*, [SP]. The details of its formulation are standard and thus relegated to Appendix S.1. For purposes of

-
- (21) MLRP and PPR are “nearly orthogonal” generalizations in the following sense. Any process that satisfies both conditions must have $f_{ii} = \pi_i$ for all “interior” states $i \notin \{1, d\}$. If d is large and the $\{\pi_i\}_{i=2}^{d-1}$ are not too small, the transition matrix is nearly iid. If d is large and $\pi_1 + \pi_d \approx 1$, then the transition probabilities are close to those of a two-state process on the state space $\{1, d\}$.
- (22) For example, random walks and mean-reverting processes with normal, log-normal, or Pareto transitions satisfy the MLRP. These processes are commonly used in dynamic insurance and taxation models with continuous types — see, eg, Williams (2011), Golosov, Troshkin and Tsyvinski (2016), Farhi and Werning (2013), and Kapička (2013). Our setup encompasses these processes when their state spaces have been truncated and discretized.
- (23) As is standard, we work in partial equilibrium so that there is *no* explicit resource constraint. We think of the principal as having access to a linear storage technology with rate of return $R = 1/\alpha$. Of course, under appropriate boundedness conditions, the principal’s problem (more precisely, the efficiency problem [Eff_{*i*}] defined in Section 5.1) is dual to the problem of a utilitarian planner with an intertemporal resource constraint — see, eg, Golosov, Tsyvinski and Werquin (2016).

analysis, it is much more convenient and tractable to view contract and report choices as Markov Decision Processes (MDPs) with simple state spaces. Thus, we develop an alternative *recursive* formulation of the principal’s contracting problem. We proceed directly to this formulation in Section 3.1, deferring a discussion of its relation to the sequential problem [SP] to Section 3.2 and, in more detail, Appendix S.1.3.

3. Recursive Contracts

3.1. The Recursive Problem

When the type process is iid, Green (1987) and Thomas and Worrall (1990) show that the appropriate state variable for the principal’s MDP is the agent’s promised utility — ie, the lifetime expected utility starting from the given period that he would obtain if he were truthful in all future periods. In the Markovian setting, promised utility is not a sufficient state variable because the agent’s *true* type determines both his current marginal utility of consumption *and his preferences over continuation contracts*, so that both aspects of preferences are private information. The principal therefore needs additional instruments to screen the agent through continuation contracts. Our formulation of the principal’s problem uses the pair (\mathbf{v}, s) of *contingent* promised utilities and yesterday’s (reported) type as state variables. Here, s is the previous period’s *reported* type, and $\mathbf{v} := (v_1, \dots, v_d)$ where v_i denotes the utility promised to the agent conditional on reporting type i today.²⁴ Importantly, v_i is the lifetime utility promised to a type- i agent *assuming he reports truthfully going forward*.²⁵

A transfer of c_i from the principal to an agent with endowment ω_i delivers to the agent $u_i := U(c_i + \omega_i)$ flow utiles. Thus, any such transfer is equivalent to a *flow utility allocation* of u_i to an type- i agent at a cost $C(u_i, i) := C(u_i) - \omega_i$, where $C(u) := U^{-1}(u)$. By Assumption DARA, $C(\cdot)$ is strictly increasing, strictly convex, and continuously differentiable, and satisfies the Inada conditions $\lim_{u \rightarrow -\infty} C'(u) = 0$ and $\lim_{u \rightarrow 0} C'(u) = \infty$. Because consumption is unbounded above, it also satisfies $\lim_{u \rightarrow 0} C(u) = +\infty$. Define the function $\psi : \mathcal{U} \times S \times S \rightarrow \mathcal{U}$ by

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

which is the flow utility for an agent of type i who claims to be of type $j \geq i$. If an agent of type i reports truthfully, he receives flow utility $\psi(u, i, i) = u$.

(24) Thus, at every step in the principal’s MDP, the state variable essentially is the same as the initial conditions in [SP] — namely, a vector of contingent promised utilities and a prior over today’s type. In the recursive formulation, $s \in S$ induces the “prior” \mathbf{f}_s .

(25) Note that 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 $s^{(t+1)} = i$. The stochastic process $(s^{(t)})_{t=0}^{\infty}$ is the *type process*. With a slight abuse of terminology, we will often refer to the endowment and type processes interchangeably, despite the minor timing discrepancy.

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 denotes the contingent utility vector promised to the agent if he reports his current type to be $i \in S$. Moreover, a reported type of $i \in S$ means that the principal's state variable in the next period is (\mathbf{w}_i, i) .

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 instant, assuming he reports truthfully in the future. In the above notation, these *recursive constraints* are, for all $i, j \in S$ with $i > j$,²⁶ the *promise keeping* and *incentive compatibility* conditions, rendered as

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

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 .) Importantly, notice in **[IC_{ij}]** that, even if the agent lies today, his expectation over tomorrow's type is still governed by his true current type. This set of constraints is independent of the previous report s . 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. Thus, our formulation solves the issue of the agent's private preferences over continuation contracts in our setting.²⁷

The next, and essential, step in the recursive formulation is to specify which \mathbf{w}_i 's are feasible for the principal to offer to the agent, ie, which promised utility vectors are *implementable*. In general, there exist $\mathbf{v} \in \mathcal{U}^d$ and $(u_i, \mathbf{w}_i)_{i \in S}$ that satisfy all of the recursive constraints at \mathbf{v} , but for which there does *not* exist any $(u'_i, \mathbf{w}'_i)_{i \in S}$ that satisfy the recursive constraints at one (or more) of the \mathbf{w}_j . Clearly, then, that \mathbf{w}_j should not have been considered feasible in the first place.

Definition 3.1. A set $D' \subseteq \mathcal{U}^d$ is a *recursive domain* if, for every $\mathbf{v} \in D'$ there is a tuple $(u_i, \mathbf{w}_i)_{i \in S}$ satisfying the recursive constraints such that $u_i \in \mathcal{U}$ and $\mathbf{w}_i \in D'$ 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, if it exists, is denoted D .²⁸

Thus, the largest recursive domain D characterizes the implementable promised utility vectors, and a promised utility vector \mathbf{v} should be considered feasible if, and only if, $\mathbf{v} \in D$. In Section 4.2, we show that a largest recursive domain exists and characterize its properties.

(26) Assumption **NHB**, No Hidden Borrowing, implies that we needn't consider **[IC_{ij}]** with $j > i$.

(27) Because the endowment process is Markovian, the agent's incentives depend only on his current type, and not on the history of his past true or reported types. Thus, the Markovian structure is essential for the present recursive formulation. See, eg, Pavan, Segal and Toikka (2014) for a discussion of these and related points.

(28) Clearly, if D exists it is unique. We are implicitly using the fact that $\mathcal{U} = \mathbb{R}_{--}$ in defining recursive domains to be subsets of \mathcal{U}^d ; in general, one can simply normalize flow payoffs by $1 - \alpha$, as is standard.

A *recursive contract* is a map $\xi : D \times S \times S \rightarrow \mathcal{U} \times D$, written as $\xi(\mathbf{v}, s, i) = (\xi^f(\mathbf{v}, s, i), \xi^c(\mathbf{v}, s, i))$, where $\xi^f(\mathbf{v}, s, j) = u_j(\mathbf{v}, s) \in \mathcal{U}$ provides *flow* consumption utilities to today's report of ω_j , 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 $(\xi(\mathbf{v}, s, i))_{i \in S} \in \Gamma(\mathbf{v})$, where the correspondence $\Gamma : D \rightrightarrows (\mathcal{U} \times D)^d$ defined by

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

is the principal's *constraint correspondence*. 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*.²⁹

The principal's *recursive problem*³⁰ is to choose the recursive contract that minimizes the lifetime expected cost of the induced allocation, subject to the recursive constraints at each step:³¹

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

Note that the expectation is taken with respect to the true probability measure over paths of endowment types, which is the measure over reported paths induced by the agent selecting the truthful reporting strategy. 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 *recursively optimal* if it attains the infimum in $[\mathbf{RP}]$.³²

(29) In particular, $u_\xi^{(t)} := \xi^f(\mathbf{v}^{(t)}, s^{(t)}, s^{(t+1)}) = u_{s^{(t+1)}}^{(t)}(\mathbf{v}^{(t)}, s^{(t)})$ and $\mathbf{v}_\xi^{(t)} := \xi^c(\mathbf{v}^{(t-1)}, s^{(t-1)}, s^{(t)}) = \mathbf{w}_{s^{(t)}}(\mathbf{v}^{(t-1)}, s^{(t-1)})$. The transition probabilities of these processes are determined by the agent's reporting strategy and the underlying Markovian type process.

(30) Note well that the principal's problem here is recursive because we are considering *recursive* contracts, instead of the *sequential* contracts considered in Appendix S.1. A more apt terminology, familiar from Stokey, Lucas and Prescott (1989), would be to call $P(\mathbf{v}, s)$ the principal's sequential value function over recursive contracts. However, for the sake of brevity, and because $P(\mathbf{v}, s)$ satisfies a Bellman equation $[\mathbf{FE}]$, we shall refer to it as the value function for the recursive problem.

(31) With a slight abuse of notation, the expectation operator \mathbf{E} serves two purposes. When given a superscript \mathbf{f}_i and a vector argument such as \mathbf{w}_j , it denotes a dot product as described below the statements of $[\mathbf{PK}_i]$ and $[\mathbf{IC}_{ij}]$. Without a superscript, and with a scalar argument, it denotes the expectation operator corresponding to the probability measure $\mathbf{P} \in \Delta(S^\infty)$ induced by the transition probabilities defined in Section 2.

(32) Recursive contracts, as we have defined them, are (i) *deterministic*, in that they do not involve extraneous randomization, and (ii) *stationary* in that they do not depend explicitly on the public history or on time. It is a standard result that both restrictions are without loss of optimality in our setting. There is no gain to stochastic mechanisms because the cost function $C(\cdot)$ is convex and the constraint correspondence $\Gamma(\cdot)$ has convex graph (see part (b) of Theorem 1), and it is easy to see from the Bellman equation in Theorem 2 that *some* recursively optimal contract is stationary. When the value function is strictly convex, which is guaranteed

3.2. Optimality for the Agent

Before proceeding, two points concerning [RP] must be addressed. First, given a recursive contract $\xi \in \Xi(\mathbf{v})$, it is possible, even under truth-telling and despite the promise keeping constraints [PK_{*i*}] constraints holding at each step, that $v_j \neq \mathbf{E} \left[\sum_{t=0}^{\infty} \alpha^t \tilde{u}_{\xi}^{(t)} \mid s^{(0)} = s_j \right]$. In this case, we say that ξ does not *deliver promises*. This can happen if the promised utility process grows too quickly — essentially, if the principal violates the analogue of a no-Ponzi-scheme condition. Second, every recursive contract ξ induces a decision problem for the agent, and it is possible that truth-telling is not an optimal strategy. In particular, although the incentive constraints [IC_{*i*}] deter one-shot deviations — making truth-telling an *unimprovable* strategy in the agent’s decision problem — more complicated, infinite-length deviations may yet be profitable. This can happen if the utility process induced by the contract does not satisfy a “continuity at infinity” condition.

In either case, the recursive contract essential fails to do what it purports: to deliver a specified amount of promised utility and to induce truth-telling. By extension, in either case, a recursive contract may induce an allocation that does not satisfy the “full” set of constraints embodied in the sequential problem [SP] (roughly, that truth-telling is a globally optimal strategy). Here, we state a standard sufficient condition for a recursive contract to both deliver promises and be continuous at infinity, and thus also to induce an allocation that is feasible in the sequential problem [SP]. Further details and discussion can be found in Appendix S.1.3.

Let $\mathcal{H} := S^{\infty}$ denote the space of all infinite sequences, or *paths*, of endowment 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

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

where $(\mathbf{v}^{(t)}(h))_{t=0}^{\infty}$ denotes the (deterministic) sequence of contingent promises along the path $h \in \mathcal{H}$.³³ Any feasible recursive contract ξ that satisfies [TVC] is said to be [TVC]-*implementable*.

Lemma 3.2. If a recursive contract ξ is [TVC]-implementable at $\mathbf{v} \in D$, then:

(a) It delivers promises at \mathbf{v} , ie,

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

for all $i \in S$.

by [TVC]-regularity, the unique recursively optimal contract strictly dominates all non-deterministic and non-stationary mechanisms. See Sections 4.3 and 4.4, and especially part (e) of Theorem 2), as well as footnote 34.

(33) Agent transversality is a slightly weaker sufficient condition than the one given in Theorem 9.2 of Stokey, Lucas and Prescott (1989, pp. 246-247) and the notion of a “lower convergent” utility process in Kreps (1977).

- (b) Truthtelling after *every* history is an optimal strategy for the agent.
- (c) The induced allocation \tilde{u}_ξ is feasible in [SP].

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

The proof of Lemma 3.2 is in Supplementary Appendix S.1.2. Let $\Xi^*(\mathbf{v}) \subseteq \Xi(\mathbf{v})$ denote the set of feasible recursive contracts that are initialized and [TVC]-implementable at $\mathbf{v} \in D$. Let

$$D^* := \{\mathbf{v} \in D : \Xi^*(\mathbf{v}) \neq \emptyset\}$$

denote the set of contingent promises that can be generated by some [TVC]-implementable contract.³⁴

4. Implementability and Optimality

4.1. Full Information Benchmark

Before proceeding to analyze the recursive problem [RP], it is useful to briefly describe the first-best optimal contract that arises when there is full information, ie, when the principal is able to observe the agent's endowment types in each period. This is an important benchmark, and properties of its solution are essential for characterizing the optimal contract under hidden information.

There are no incentive constraints under full information, so every promised utility vector $\mathbf{v} \in \mathcal{U}^d$ can be implemented and it is optimal to provide full insurance. The optimal 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) $v_1^{(t)} = \dots = v_d^{(t)}$.

Since it will be referenced below, we note that the principal's value function in the full information problem is denoted $Q^* : \mathcal{U}^d \times S \rightarrow \mathbb{R}$. In Supplementary Appendix S.2, we characterize its properties and formalize the above discussion concerning the first-best optimal contract.

(34) We could similarly define $\Xi^\dagger(\mathbf{v})$ to be the set of feasible recursive contracts that are initialized and deliver promises at $\mathbf{v} \in D$. Then, $D^\dagger := \{\mathbf{v} \in D : \Xi^\dagger(\mathbf{v}) \neq \emptyset\}$. All subsequent results — except for point (iv) in part (e) of Theorem 2 — if we replace [TVC] with [DP], D^* with D^\dagger , and $\Xi^*(\cdot)$ with $\Xi^\dagger(\cdot)$ everywhere. (For example, Condition R.4 could be weakened to the obvious analogue, [DP]-regularity.) These hypotheses are weaker than the ones stated in the main text, as [DP] is a *necessary* condition for \tilde{u}_ξ to be feasible in [SP], while [TVC]-implementability is a *sufficient* condition by Lemma 3.2. We state everything in terms of the stronger [TVC]-implementability condition because it simplifies some statements, economizes on notation, and is no more difficult to verify than [DP].

4.2. Implementability

The first step in the analysis of the recursive problem [RP] is to characterize implementable promised utilities (ie, the sets D and D^*) and implementable recursive contracts (ie, the constraint correspondence $\Gamma(\cdot)$). Theorem 1, stated below, characterizes these objects. Parts (a)–(c) establishes existence and other basic properties that are valid for all primitives satisfying our basic Assumptions DARA, NHB, and Markov. Parts (d)–(f) provide sharper characterizations under additional restrictions on the primitives, and are stated in order of decreasing generality.

We emphasize that this step is essential and that the sets D and D^* should be viewed as important components of the solution to [RP].³⁵ Conceptually, a tight characterization of D sheds light on how the primitives — ie, characteristics of preferences and the type process — shape the incentive constraints, and is thus critical for understanding the forces underlying the optimal contract. Technically, all of our subsequent results rely heavily on properties of D and D^* established here.

Theorem 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 Markov process, D is independent of the discount factor $\alpha \in (0, 1)$ and the utility function $U(\cdot)$.³⁶
- (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 (ie, if $\mathbf{v} \in D$, then $a\mathbf{v} \in D$ for all $a \in (0, 1]$), and is unbounded below (ie, for all $k < 0$ there exists some $\mathbf{v} \in D$ such that $\mathbf{v} \leq k\mathbf{1}$).
- (d) If the Markov process satisfies FOSD, then $D^* \subseteq V_d$.
- (e) If the Markov process satisfies either MLRP or PPR, then $D = V_d$.
- (f) If the Markov process satisfies either MLRP or PPR and, in addition, the agent has CARA utility, then $D = V_d = D^*$.

(35) Different recursive formulations yield different domains. For example, while the domain in Fernandes and Phelan (2000) is (necessarily) isomorphic to ours, it is defined contingent on the previous period's report. That is, their domain is actually a function $W : S \rightarrow 2^{\mathcal{U}^d}$, and the collection of sets $\{W(s)\}_{s \in S}$ must be solved for jointly. Consider the case where $d = 2$ and the Markov process satisfies MLRP. By part (e) of Theorem 1, our domain is the set V_2 , which is independent of the transition probabilities (within the MLRP class). For a given $\mathbf{v} \in V_2$ and $s \in S$, let $v^p(s) := \mathbf{E}^{\mathbf{f}^s}[\mathbf{v}]$ denote the *ex ante* promised utility, and $v^\dagger(s) := \mathbf{E}^{\mathbf{f}^{3-s}}[\mathbf{v}]$ denote the *threat point* utility, the utility the agent gets if he lied in the last period about his type. For each $s \in S$, the set $W(s)$ consists of tuples of the form $(v^p(s), v^\dagger(s))$. It is easy to see that $W(1) = \begin{bmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{bmatrix} V$ and $W(2) = \begin{bmatrix} f_{21} & f_{22} \\ f_{11} & f_{12} \end{bmatrix} V$, which exhibits the isomorphism. Thus, $W(1)$ and $W(2)$ are cones and are symmetric about the diagonal in \mathbb{R}_+^2 . Importantly, the sets $W(s)$ depend on the transition probabilities (even within the MLRP class). For example, if types are iid, then $W(1) = W(2) = \{(t, t) : t < 0\}$, while with positive serial correlation, they are disjoint open cones.

(36) That is, it is independent of the utility function within the class allowed for by Assumption DARA.

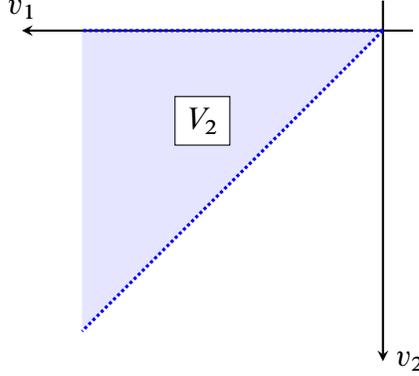


Figure 1: The largest recursive domain $D = V_2$ in the $d = 2$ and FOSD case.

The proof of Theorem 1 is in Appendix A. While the proof is quite involved, it is worth sketching both as a source of intuition for the theorem and to highlight some novel technical contributions.

The most basic pieces of the theorem are parts (b) and (c) and the statements regarding existence, non-emptiness, and convexity in part (a). We establish existence of D by characterizing it as the largest fixed point of a monotone set-valued operator \mathcal{B}_\circ . The convexity and decreasing returns properties in parts (a) and (c) follow from part (c) of Assumption DARA, which, in particular, implies that the $\psi(\cdot, i, j)$ functions in the incentive constraints $[\text{IC}_{ij}]$ are convex. Non-emptiness of D and D^* , as well as the unboundedness property of D^* , are established by direct construction.

Before explaining the remaining pieces of part (a) — namely, the statements that D is open, conic, contains V_d , and depends only on the Markov process — and the sharper characterizations in parts (e) – (f), it is useful to first understand the special role of the set V_d . For simplicity, consider the case where $d = 2$ and the Markov process satisfies FOSD (which is illustrated in Figure 1). Subtracting both sides of the promise keeping constraint $[\text{PK}_i]$ ($i = 1$) from the incentive constraint $[\text{IC}_{ij}]$ (setting $j = 1$ and $i = 2$) yields

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

To see that the iid information rent is positive, notice that $\psi(u_1, 2, 1) = U(\omega_2 + C(u_1) - \omega_1) > u_1$. If the contingent continuation utility $\mathbf{w} \in V_2$, then the Markov information rent is also positive whenever \mathbf{f}_2 first order stochastically dominates \mathbf{f}_1 . This implies that if $\mathbf{w}_1 \in V_2$, then we must have $v_2 > v_1$. But iterating on this idea — essentially by noting that \mathbf{w}_1 must also have some implementation — we can conclude that as long as the recursive contract actually delivers the promised utility \mathbf{v} to the agent (and, in particular, if $[\text{TVC}]$ holds), we must always have $v_2 > v_1$. This last argument is in fact the substance of the proof of part (d) of Theorem 1.

While this calculation suggests that D and V_d coincide whenever the Markov process satisfies FOSD, it merely establishes an inner bound for D and is thus far from a full proof.

Even with CARA utility, while it is easy to see that D must be a cone, it is not clear *which* cone it equals.³⁷ Indeed, even in the CARA case, it is perhaps unexpected that the domain would equal a set, such as V_d , that is independent of both the transition probabilities \mathbf{f}_i and the endowment sizes ω_i . With general preferences satisfying Assumption **DARA**, it is not even clear that D should be a cone.³⁸

The remainder of the proof fills in these gaps by studying more detailed properties of the operator \mathcal{B}_\circ . This requires several steps. In the first step, we show, by direct calculation, that (i) $V_d \subseteq \mathcal{B}_\circ(V_d)$ for general Markov processes and (ii) V_d is a fixed point of \mathcal{B}_\circ when the Markov process satisfies FOSD. Point (i) implies that $V_d \subseteq D$ always; when FOSD fails, the inclusion is typically strict because the iid and Markov information rents act in “opposite directions.”³⁹ When the Markov process satisfies FOSD, these sources of information rents act in the same direction, but point (ii) still leaves open the question whether V_d is the *largest* fixed point of \mathcal{B}_\circ .

In the second step — which is the most novel and perhaps most involved part of the proof — we focus on the case of CARA utility and show that V_d is indeed the largest fixed point of \mathcal{B}_\circ in two special cases with positive serial correlation. While the operator approach itself is not new,⁴⁰ we show how to construct \mathcal{B}_\circ and its iterates *in closed form* when the Markov process satisfies either MLRP or PPR. To do this, we formulate and solve an auxiliary linear programming problem that characterizes exactly the set of implementable \mathbf{v} when the \mathbf{w}_i are required to satisfy an appropriate set of linear constraints. Under MLRP or PPR, we can determine which constraints in the linear program bind at the optimum, and thus determine its solutions. By iterating this procedure, we obtain better and better outer approximations of the domain D and can verify that no convex cone strictly larger than V_d can constitute a largest fixed point of \mathcal{B}_\circ . Thus, $D = V_d$ under CARA utility and MLRP or PPR. We show that $D^* = V_d$ in these cases as well by explicitly constructing a feasible recursive contract that

(37) The recursive constraints are always linear in contingent continuation utilities \mathbf{w}_i . The absence of wealth effects in the CARA case implies that the ψ functions, and hence the constraints, are also linear in the flow utilities u_i .

(38) The DARA property, part (c) of Assumption **DARA**, implies rather directly that D must have decreasing returns in the sense of part (c) of Theorem 1. But it is not clear that D should have “constant returns,” as required if it is conic.

(39) For example, when $d = 2$ the Markov information rent in $[\mathbf{IC}_{21}^*]$ may be written as $\alpha(f_{22} - f_{12}) \cdot (w_{12} - w_{11})$. If FOSD is not satisfied (ie, if $f_{22} - f_{12} \equiv f_{11} - f_{21} < 0$), any $\mathbf{w}_1 \in V_2$ will confer *negative* Markov information rents. Intuitively, negative serial correlation relaxes the incentive constraints because it implies that private information is short-lived (relative to the iid benchmark). Following the same linear programming procedure used in Appendix A.2.2, the reader may easily verify in this case that $D = \{\mathbf{v} \in \mathbb{R}_{--}^2 : v_2 > (f_{21}/f_{11})v_1\} \supsetneq V_2$.

(40) It dates back to at least Abreu, Pearce and Stacchetti (1990) in the context of repeated games, and was used in Fernandes and Phelan (2000) to establish an existence result analogous to the existence statement in part (a) of Theorem 1. One non-standard aspect of our environment, relative to those papers, is that the range of flow utilities \mathcal{U} is both unbounded and open. This makes establishing even basic properties of D somewhat subtle, and effectively requires transfinite iterations of the \mathcal{B}_\circ operator.

“stops” after finitely-many steps, and is therefore guaranteed to satisfy [TVC].⁴¹

The final step extends these ideas to general utility functions and type processes. The essential idea is to characterize iterates of \mathcal{B}_0 via solutions of auxiliary *concave* (instead of linear) programming problems. The solutions to these concave programs are, perhaps surprisingly, independent of all model primitives aside from the transition probabilities. Thus, in particular, all properties of D established in the much simpler CARA case — namely, the cone property and the fact that $D = V_d$ under PPR/MLRP — immediately extend to general utilities. Moreover, solutions to these programs — which correspond to points in the boundary of D — require that the flow utility terms $u_i = 0$. Thus, the boundary points cannot be implemented, and D must be open. To see an example of this, consider again the $d = 2$ and FOSD case. Let \mathbf{v} lie in the lower boundary of $D = V_2$, so that $v_2 = v_1$. Because the Markov information rent is positive under FOSD, [IC₂₁^{*}] implies that any menu implementing \mathbf{v} must satisfy $u_1 = 0$. This is impossible — it requires the low endowment type to receive infinite consumption — and so \mathbf{v} in this lower boundary cannot be implemented.

4.3. Regularity Conditions

Having characterized *implementable* contracts, the rest of the analysis focuses on characterizing *optimal* contracts. To ensure that the optimization problem in the recursive problem [RP] is sufficiently well-behaved, we require that the environment satisfy a few mild regularity conditions. Our approach here is to state a small set of conditions directly in terms of derived objects, and which can be readily verified on a case-by-case basis for particular parameterizations.

Definition 4.1. 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 real-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

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

R.3 (Constraint Qualification) Let $\Gamma_0(\mathbf{v}) \subseteq \Gamma(\mathbf{v})$ 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 $\mathbf{v} \in D$, $\Gamma_0(\mathbf{v}) \neq \emptyset$.

(41) We conjecture that $D = D^* = V_d$ for any utility function satisfying Assumption DARA when the Markov process satisfies the weaker FOSD condition. Extending part (e) of the theorem (ie, the equality $D = V_d$) is complicated because we are unable to determine, in general, which constraints in the auxiliary LP bind at the optimum. Extending the explicit construction behind part (f) of the theorem (ie, the equality $D^* = V_d$) poses no conceptual challenge but is extremely tedious.

R.4 ([TVC] *Existence*) There exists a recursively optimal contract $\xi^* \in \Xi^*(\mathbf{v}^{(0)})$.⁴²

Conditions R.1–R.3 are mild technical conditions that allow us to establish basic properties of the principal’s recursive problem [RP]. For example, Condition R.2 and the well-posedness criterion in Condition R.1 automatically hold when the consumption domain \mathcal{C} is bounded below, as is the case when the agent has CRRA utility. In general, both Conditions R.1 and R.2 can be verified by constructing real-valued functions that serve as upper and lower bounds for P and checking that the flow cost function $C(\cdot)$ satisfies appropriate growth conditions. We show how to carry out these verification steps for CARA utility in the appendix; the most involved step is constructing a real-valued function to serve as an upper bound for P .⁴³ Condition R.3 is a standard sufficient condition for the existence of Lagrange multipliers, thereby allowing us to use Lagrangian methods, and is guaranteed to hold when the transition probabilities $\{\mathbf{f}_i\}_{i \in \mathcal{S}}$ are FOSD-ordered or affinely independent. Notably, affine independence is, in a particular sense, without loss of generality.⁴⁴ The following lemma records this discussion.

Lemma 4.2. The regularity conditions can be verified in the following cases:

- (a) Condition R.3 holds if either (i) the Markov process satisfies FOSD, or (ii) the transition probabilities $\{\mathbf{f}_i\}_{i \in \mathcal{S}}$ are affinely independent.
- (b) If the agent has CARA utility and the Markov process satisfies MLRP or PPR, the environment is regular.

The proof of Lemma 4.2 is in Supplementary Appendix S.3.1.

Condition R.4, on the other hand, is a somewhat more substantive requirement and is important for the proofs of our main economic results, Theorems 3 and 4. In settings such as ours where \mathcal{U} is not bounded, Condition R.4 typically cannot be verified without solving for the optimal contract in (nearly) closed form. Even with iid types this can only be done in a few special cases, and we do not know of any cases with Markovian types where it is possible. These issues are discussed in more detail in Appendix S.1.3. On the other hand, we do not

(42) $\mathbf{v}^{(0)}$ is the initial condition for the promised utility process. It may be given exogenously, or optimally initialized as in the efficiency problem [Eff₁] described below in Section 5.1.

(43) We use the value function induced by the (suboptimal) [TVC]-implementable contract constructed in the proof of part (f) of Theorem 1. The main difficulty of that construction is precisely ensuring that the contract has finite value.

(44) For any given dimension d , transition matrices that fail affine independence (or, equivalently, linear independence) are non-generic under any standard genericity notion. Of course, important examples, such as the iid benchmark, are non-generic in this sense. Even in these cases, affine independence is without loss in the following sense. But if the affine hull of $\{\mathbf{f}_i\}_{i \in \mathcal{S}}$ has dimension $d' < d$, it is always possible to (i) reduce the dimensionality of the promised utility state variable to d' by “pooling” the promise keeping constraints together with appropriate weights, and (ii) then analyze recursive contracts on this lower-dimensional domain. In the extreme case of iid types ($d' = 1$), this reduces to the scalar state variable of, eg, Thomas and Worrall (1990). The projection onto this lower-dimensional space is linear, and thus preserves convexity, topological, and smoothness properties of D and P , so the analysis can be carried out essentially verbatim after this reduction. Details are available upon request.

know of any examples in which Condition R.4 fails, either, and in light of Lemma 3.2 it should be viewed as a minimal consistency requirement on any solution to [RP]. (As noted in footnote 34, Condition R.4 can actually be weakened slightly.)

4.4. Bellman Equation

The main result of this section shows that the principal's value function P satisfies a standard Bellman equation, characterize its properties, and establish basic properties of the optimal contract that it generates.

Theorem 2. *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 a recursively 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) .⁴⁵
- (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 recursively 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}_{ξ^*} solves the sequence problem [SP].

The proof of Theorem 2 is in Supplementary Appendix S.3. The proof makes clear precisely which of the Conditions R.1–R.4 are used to establish each property stated in the theorem.

While many properties described in Theorem 2 are standard, two points warrant further explanation. First, part (b) highlights important non-monotonicity and boundary properties of the value function that derive from fundamental features of the incentive constraints. The unboundedness of P near the boundaries of D follows from the same logic used to establish that

(45) 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 is open in part (a) of Theorem 1; indeed, these two properties are essentially equivalent.⁴⁶ To get intuition for the non-monotonicity properties, consider once again the case in which $d = 2$ and the Markov process satisfies FOSD (recall Figure 1). An increase in v_1 tightens both the promise keeping constraint $[\mathbf{PK}_i]$ ($i = 1$) and the incentive constraint $[\mathbf{IC}_{21}^*]$, leading to unambiguously higher costs for the principal. An increase in v_2 , on the other hand, tightens $[\mathbf{PK}_i]$ ($i = 2$) and adds slack to $[\mathbf{IC}_{21}^*]$. 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 $[\mathbf{FE}]$ may be decomposed into d separate problems, each of which conditions only on the agent's *current* report.⁴⁷ 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 \mathbf{w}_i only appear in those incentive constraints that deter deviations by *higher* types.⁴⁸

5. The Optimal Contract — Long-Run Properties

We now characterize the long-run properties of the optimal contract. In Section 5.1, we introduce a martingale that plays a central role in our characterization. Then, in Sections 5.2 and 5.3 we present our main results on, respectively, absolute and relative immiseration.

(46) Because P is unbounded even on bounded sets, we are unable to use standard contraction mapping arguments to ensure that P is the *unique* solution to the functional equation $[\mathbf{FE}]$. However, part (a) of Theorem 2 states that P is the *smallest* solution in an appropriate class. This fact is essential both for establishing certain analytical properties of P (eg, that the directional derivative in Proposition 5.1 is non-negative) and for numerical computation.

(47) See Appendix B.3.3 for more details. We emphasize that this separation is not possible in the Fernandes and Phelan (2000) approach for reasons described in footnote 35.

(48) Without Assumption **NHB**, the policy functions $u_i(\cdot)$ and $\mathbf{w}_i(\cdot)$ would depend on (\mathbf{v}, s) through the entire vector of promises \mathbf{v} , as these variables must be chosen to deter deviations by both higher and lower types.

5.1. The Differential Martingale

In the iid setting, Thomas and Worrall (1990) show that the derivative of their value function with respect to (scalar) ex ante promised utility is a martingale. The appropriate martingale in our Markovian setting is a somewhat subtle, but natural, generalization. We 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 . Denote the directional derivative of $P(\cdot, s)$ in direction $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^d$ by $D_{\mathbf{1}}P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} [P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)] / \varepsilon$.

Proposition 5.1. Suppose that the environment is regular. Then the stochastic process induced by the optimal contract, $(D_{\mathbf{1}}P(\mathbf{v}^{(t)}, s^{(t)}))_{t=0}^{\infty}$, is a non-negative martingale. If the environment is [TVC]-regular, then this process is strictly positive.

The proof of Proposition 5.1 is in Appendix B. We will refer to this stochastic process as the *differential martingale*. The statement of the proposition includes three important pieces: the martingale property, the direction $\mathbf{1}$ in which the directional derivative is taken, and the fact that the directional derivative is strictly positive.

The intuition for the martingale property is simply that the principal *optimally smooths costs* over time and across states. To be a bit more precise, 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, this cost of this increase in promises is close to the marginal cost $D_{\mathbf{1}}P(\mathbf{v}, s)$. One way to deliver the additional utility in an incentive-compatible way is to increase each of the continuation promises \mathbf{w}_i to $\mathbf{w}'_i := \mathbf{w}_i + (\varepsilon/\alpha)\mathbf{1}$. To a first-order approximation, the cost of this perturbation is $\sum_{i=1}^d f_{si} D_{\mathbf{1}}P(\mathbf{w}_i, i)$. An envelope theorem argument implies that this perturbation is locally optimal, implying that the marginal costs are equal, giving us precisely the martingale property

$$D_{\mathbf{1}}P(\mathbf{v}, s) = \sum_{i=1}^d f_{si} D_{\mathbf{1}}P(\mathbf{w}_i, i)$$

in Proposition 5.1.

To understand why the directional derivative must be taken in the direction $\mathbf{1}$, notice that this is the *unique* direction of change for \mathbf{v} that increases the agent's *ex ante* continuation utility *while leaving his information rent unchanged*. As in Section 4.2, substituting the promise keeping constraint [PK_{*i*}] ($i = j$) into the incentive constraint [IC_{*ij*}] ($i > j$) gives an alternative way to write the latter:

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

Written this way, 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 $\mathbf{1}$. Another way to see this is by considering perturbations of the \mathbf{w}_i . Suppose the principal wants to change \mathbf{w}_i to $\hat{\mathbf{w}}_i$ in such a way that the expected value to type i is $\varepsilon > 0$, ie, $\mathbf{E}^{f_i} [\hat{\mathbf{w}}_i] - \mathbf{E}^{f_i} [\mathbf{w}_i] = \varepsilon$. In general,

each type values such a perturbation differently, as each (current) type has distinct beliefs concerning future types. 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, ie, that $\mathbf{E}^{f^j} [\hat{\mathbf{w}}_i - \mathbf{w}_i] = \varepsilon$ for all $j \in S$. (When the transition probabilities $\{\mathbf{f}_i\}_{i=1}^d$ are affinely independent, $\mathbf{1}$ is the *unique* direction with this property.)

Finally, the intuition for strict positivity of the directional derivative $D_{\mathbf{1}}P(\mathbf{v}, s)$ is similar. As we have seen in part (b) of Theorem 2, there are typically regions of the domain on which the value function is decreasing in certain directions, so it is *not* true that arbitrary directional derivatives are non-negative. The value function is decreasing or non-monotone in certain directions because increasing a single component of \mathbf{v} has two effects. First, it increases the promise to some type, which mechanically increases costs. Second, it either adds slack to or tightens certain incentive constraints $[\mathbf{IC}_{ij}^*]$. Whenever it adds “enough” slack to the incentive constraints, the second effect dominates and leads to an overall decrease in costs. The above arguments show that, by taking the derivative in direction $\mathbf{1}$, the second effect completely washes out, leaving only the mechanical increase in costs due to higher promised utilities.

Efficiency and the TW Martingale: At this point, it is useful to make two related observations that are useful going forward. First, Markovian private information induces a novel intertemporal distortion that is not present with iid types. Second, when the type process is iid, the differential martingale reduces to the martingale from Thomas and Worrall (1990) (henceforth, the TW martingale).

To see how this is so, consider the question, What is the cost-minimizing way for the principal to give a type i agent at least *ex ante* promised utility $v \in \mathcal{U}$, without consideration of today’s incentives? Formally, this is equivalent to solving the problem⁴⁹

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

We say that solutions to the problems $[\mathbf{Eff}_i]$ are *efficient*. Very informally, efficiency corresponds to a kind of “renegotiation-proofness”: even *after* the agent reveals himself to be of type i and so incentive compatibility in the current period is no longer a concern, there is no way for the principal to reduce her costs while leaving the agent at least as well off, in expectation, going forward.

In general, the optimal contract is typically *not* efficient, as the principal must keep in mind that her choice of \mathbf{w}_j affects the Markov information rent in each $[\mathbf{IC}_{ij}^*]$ with $i > j$. However, the optimal contract *is* efficient after reports of ω_d , ie, for all $(\mathbf{v}, s) \in D$ the policy function $\mathbf{w}_d(\mathbf{v}, s)$ solves $[\mathbf{Eff}_i]$ ($i = d$) for some optimally-chosen $w \in \mathcal{U}$. This is because, under Assumption **NHB**, the variables u_d and \mathbf{w}_d do not enter into any of the $[\mathbf{IC}_{ij}]$, only the promise keeping constraint $[\mathbf{PK}_i]$ ($i = d$).

(49) 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.”

When types are iid, the optimal contract is efficient at *all* histories. To see this, note that in the iid case, the Markov information rent in $[\mathbf{IC}_{ij}^*]$ vanishes. Therefore, the continuation utility vectors \mathbf{w}_i enter the constraints of $[\mathbf{FE}]$ (with each $[\mathbf{IC}_{ij}^*]$ taking the place of $[\mathbf{IC}_{ij}]$) only through the promise keeping constraints $[\mathbf{PK}_i]$; there, they appear only through their expectations $\mathbf{E}^\pi [\mathbf{w}_i]$, where $\pi \in \Delta(S)$ is the common vector of transition probabilities. Thus, optimality clearly implies that, for all $(\mathbf{v}, s) \in D \times S$ and $i \in S$, the policy function $\mathbf{w}_i(\mathbf{v}, s)$ must solve $[\mathbf{Eff}_i]$ (with $\mathbf{f}_i = \pi$ for each $i \in S$).

Therefore, along the path induced by the optimal contract, the principal's value process is given by $K(v^{(t)}, s^{(t)}) = P(\mathbf{v}^{(t)}, s^{(t)})$, where $v := \mathbf{E}^\pi [\mathbf{v}]$. Thus, ignoring the (in this case) irrelevant $s^{(t)}$ argument, K is the value function of Thomas and Worrall (1990) and the process $(K'(v^{(t)}, s^{(t)}))_{t=0}^\infty$ is the TW martingale. It is easy to see from the envelope and first-order conditions for $[\mathbf{Eff}_i]$ (any $i \in S$) that $K'(w, i) = P_j(\mathbf{w}_i, i)/\pi_j$ and, thus, summing over $j \in S$, that $K'(v, s) = D_1 P(\mathbf{v}, s)$. Hence, our differential martingale coincides with the TW martingale when types are iid. This is *not* to say that $K(v, s) = P(\mathbf{v}, s)$ or $K'(v, s) = D_1 P(\mathbf{v}, s)$ for all $\mathbf{v} \in D$, as there are many points $\mathbf{v} \in D$ that could not possibly be efficient.⁵⁰ These objects need only coincide at the optimum.

However, the differential and TW martingales are *not* equivalent. As we show by example in Section 7.3 (where we consider CARA utility and type processes that satisfy MLRP), in the iid setting the TW martingale can be derived via perturbations that are *not* incentive compatible when types are persistent. Thus, while there are other directional derivative processes that define martingales when types are iid, *only* the directional derivative $D_1 P(\mathbf{v}, s)$ is guaranteed to define a martingale for more general type processes.

5.2. Absolute Immiseration

With the differential martingale in hand, we may state our first main economic result. We refer to it as *Absolute Immiseration* because it states that the *level* of the agent's (promised) utility, and thus consumption, tends to its lower bound — ie, the agent becomes impoverished in absolute terms.

Theorem 3 (Absolute Immiseration). *Suppose the environment is $[\mathbf{TVC}]$ -regular. Under the optimal contract:*

- (a) $D_1 P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ almost surely.
- (b) $v_i^{(t)} \rightarrow -\infty$ almost surely for all $i \in S$.
- (c) $u_i^{(t)} \rightarrow -\infty$ and $c_i^{(t)} + \omega_i \rightarrow \underline{c}$ in probability for all $i \in S$.

(50) The first-order conditions in $[\mathbf{Eff}_i]$ imply that all of the partial derivatives $P_j(\mathbf{w}_i, i)$ are positive; this is true even when \mathbf{f}_i in $[\mathbf{Eff}_i]$ is replaced with some arbitrary “prior” $\mathbf{q} \in \Delta(S)$. But as we have seen in part (b) of Theorem 2, $P(\cdot, s)$ is non-monotone in all components of \mathbf{v} aside from v_1 . Thus, there are regions of D on which $P_j(\cdot, s) < 0$ for all $j \neq d$. This is true even in the iid case.

(d) If the type process satisfies UPR, then the convergence in part (c) can be strengthened from “in probability” to “almost surely.”

Thus, neither flow utility nor promised utility possesses a stationary distribution under the optimal contract. Net consumption possesses a stationary distribution if and only if $\underline{c} > 0$, in which case it puts full mass at the lower bound \underline{c} .

The proof of Theorem 3 is in Appendix B.4. We note that even the restriction to UPR type processes in part (d) nests the cases of (i) iid types for any d and (ii) any Markovian type process when $d = 2$. These two cases alone include essentially all extant immiseration-type results in the literature.⁵¹ While the general case covered by part (c) is stated in terms of a weaker mode of convergence (for technical reasons described below), we emphasize that this is still sufficient to rule out any stationary distribution for net consumption other than the Dirac measure at the lower bound \underline{c} . Thus, we conclude that Absolute Immiseration is very robust.

The intuition, familiar from Thomas and Worrall (1990),⁵² is that, due to risk aversion, providing incentives is cheaper for the principal when promised utility is lower. In particular, the principal provides truth-telling incentives by making (flow and continuation) utility vary depending on the reported type. Suppose, for instance, that $d = 2$ and the principal wants to induce a spread of $\varepsilon > 0$ between the flow utility for types 2 and 1, ie, $u_2 - u_1 = \varepsilon$. When the previous report was $s \in S$ and $\varepsilon > 0$ is small, this costs the principal approximately

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

Thus, because $C(\cdot, 2)$ is convex, the cost of making utility variable is an increasing function of the utility level. In the long run, it is optimal for the principal to drive this cost of incentive provision to zero by forcing the level of the agent’s utility as low as possible.

To better understand each part of Theorem 3, it is useful to walk through a sketch of the proof. While the basic idea of using the Martingale Convergence Theorem is familiar from Thomas and Worrall (1990), the details of the argument are significantly more subtle when types are not iid, and is important to understand why.

The proof of part (a) is based three simple ideas. First, the principal optimally smooths costs over time and across states, as captured by the differential martingale from Proposition 5.1. Second, the first-best contract is not implementable when the agent has private information. While this is easy to see, the key is to translate this fact into a statement about *splitting* of the differential martingale. In particular, we show that whenever the state is of the form (\mathbf{v}, d) where $\mathbf{v} = \xi^c(\mathbf{v}', s, d)$ for some $\mathbf{v}' \in D$ — ie, given some initial condition, the optimal contract maps to \mathbf{v} after a d report — it must be that $D_1 P(\mathbf{v}, d) \neq D_1 P(\mathbf{w}_i, i)$ for some $i \in S$. That is, whenever

(51) The one exception we are aware of is the example in Williams (2011) in which the hidden endowment follows a Brownian motion with drift. We provide a detailed comparison in Section 6.2.

(52) Golosov, Kocherlakota and Tsyvinski (2003) provide an especially clear discussion in the related context of “separable” taxation models (but see Section 7 for some important differences between their setup and ours).

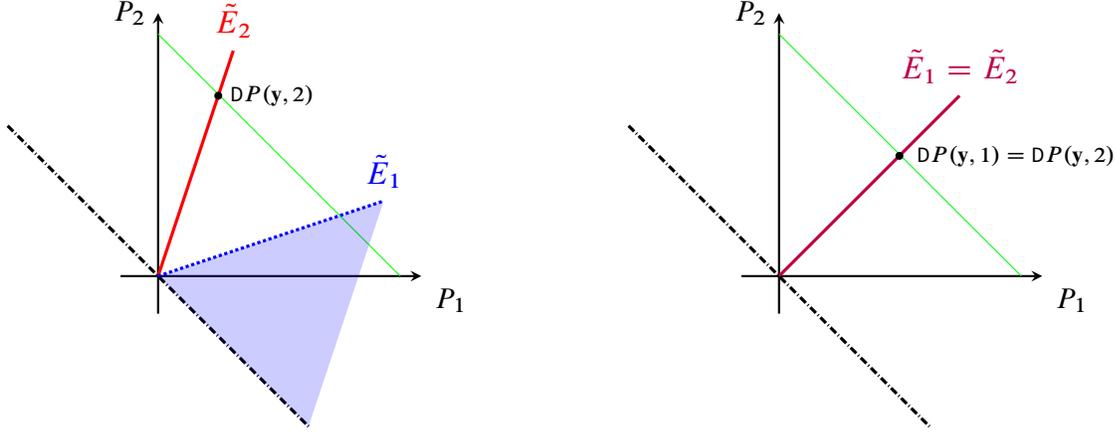


Figure 2: Martingale convergence argument when $d = 2$ and the type process satisfies FOSD (left) and is iid (right).

yesterday's reported endowment level was ω_d , the martingale cannot be constant between today and tomorrow with probability one — ie, it must split — at the optimum. Intuitively, the differential martingale splits when incentive constraints bind. Thus, the differential martingale is constant with probability one if and only if *all* of the incentive constraints are slack, in which case cost-minimization requires selecting the first-best contract. Third, as noted above in Section 5.1, the optimal contract is efficient after reports of the highest endowment type ω_d .

To see how these three ideas coalesce into a proof of part (a) of Theorem 3, consider Figure 2, which illustrates the special case of binary types. The rays \tilde{E}_i , which have slopes f_{i2}/f_{i1} , characterize the solutions of $[\mathbf{Eff}_i]$ in terms of dual variables — namely, the derivative $DP(\mathbf{w}_i, i)$, which is pinned down by the first-order conditions in $[\mathbf{Eff}_i]$. Thus, we call them *efficiency rays*.⁵³ When types are iid (the right panel in Figure 2), the rays \tilde{E}_1 and \tilde{E}_2 coincide and, as a consequence of the efficiency property discussed in Section 5.1, the derivative process satisfies $DP(\mathbf{v}^{(t)}, s^{(t)}) \in \tilde{E}_i$ for all $t \in \mathbb{N}$. Thus, on the optimal path, the derivative DP and the promised utility vector \mathbf{v} are *uniquely* pinned down by the value of the differential martingale, as illustrated in the figure. The argument familiar from Thomas and Worrall (1990) then implies that the differential martingale must converge to zero.

When types are *not* iid (the left panel in Figure 2), (i) the efficiency rays do not coincide and (ii) since the optimal contract is not efficient after $i = 1$ reports, the derivative $DP(\mathbf{w}_1, 1)$ does not lie in \tilde{E}_1 . Instead, it lies in the shaded region below \tilde{E}_1 , and its distance from \tilde{E}_1 is determined by how tightly the incentive constraint $[\mathbf{IC}_{ij}]$ ($i = 2, j = 1$) binds.⁵⁴ Thus, the value of the derivative DP and the promised utility vector \mathbf{v} is *not* uniquely pinned down by the value of the differential martingale. In general, the *iso-differential sets*

$$I_M := \{\mathbf{v} \in D : D_1 P(\mathbf{v}, s) = M\}$$

(53) See Appendix B.3.1 for details. Similar objects appear also in Zhang (2009), among others.

(54) See [C.11] and Lemma C.4 in Appendix C.

for $M > 0$ are *surfaces* in D and, even on the optimal path, it is possible that $(\mathbf{v}^{(t)})$ travels along such a surface. Thus, convergence of the differential martingale does *not* directly imply that the derivative or promised utility processes converge. We emphasize that this is not merely a mathematical complication. Rather, it is a fundamental consequence of incentive compatibility in the presence of persistent private information.

We overcome these difficulties by observing that the derivative process satisfies a kind of “renewal” property: though it takes excursions away from the efficiency rays during strings of consecutive low-type reports, it returns to the ray \tilde{E}_2 after the first high-type report, and then the cycle starts again. We thus look along the subsequence of dates at which $s^{(t)} = 2$. (Part of the following argument is illustrated in the left-hand panel of Figure 2.) Fix some path and denote this (deterministic) subsequence of dates by $(\tau^t)_{t=0}^\infty$. By the Martingale Convergence Theorem, unless the chosen path is non-generic, the differential martingale must converge to some finite, non-negative number along this path. If the limit is some number $C > 0$, then since $DP(\mathbf{v}^{(\tau^t)}, s^{(\tau^t)}) \in \tilde{E}_2$ for all $t \in \mathbb{N}$, it must be that $\mathbf{v}^{(\tau^t)} \rightarrow \mathbf{y} \in D$ with $DP(\mathbf{y}, 2) \in \tilde{E}_2$ and $D_1P(\mathbf{y}, 2) = C$. Because every state in the Markov process for endowments is recurrent, unless the chosen path is non-generic it features $s^{(\tau^t+1)} = 1$ for infinitely many t and $s^{(\tau^t+1)} = 2$ for infinitely many t . Because the differential martingale converges along the entire sequence, and by continuity of the policy functions and directional derivative, it follows that $D_1P(\mathbf{y}, 2) = D_1P(\xi^*(\mathbf{y}, 2, 1), 1) = D_1P(\xi^*(\mathbf{y}, 2, 2), 2) = C$. But we have argued above that this is not possible, as it implies that the first-best is implementable. Thus, the limit of the differential martingale must be zero, completing the proof sketch of part (a).

The remainder of the proof, which establishes parts (b)–(d) of Theorem 3, works to connect convergence properties of the differential martingale to convergence properties of flow and promised utilities. Again, this exercise is substantially complicated by the fact that the iso-differential sets are surfaces in D . The proof is fairly involved, but the main idea is to show that convergence of the differential martingale to zero implies convergence of the vector of Lagrange multipliers to the zero vector. When the type process satisfies UPR (which, recall, automatically holds in the $d = 2$ example above), we are able to show that, roughly speaking, the rate of convergence of the multipliers can be controlled uniformly by the rate of convergence of the differential martingale. This turns out to be enough to reach the desired conclusion. The general case, covered by part (c) of Theorem 3, is much more subtle, and we are unable to establish strong enough uniformity properties to guarantee almost sure convergence of the multipliers. Instead, we show how to use probabilistic arguments, and the “renewal” reasoning described above, to establish convergence in probability. Roughly speaking, while we cannot rule out that the vector of multipliers take infinitely-many excursions away from the zero vector on any given path, under Assumption [Markov](#) the type process mixes quickly enough to guarantee that such excursions are very rare.⁵⁵

(55) We are only aware of one other paper in the dynamic contracting literature — namely, Garrett, Pavan and Toikka (2018) — that establishes results concerning convergence in probability, as opposed to almost sure convergence, of the optimal contract. By design, their “variational approach” cannot *in principle* establish

5.3. Relative Immiseration

Imagine two agents, A and B , who have received the same sequence of realized endowments up through period $t - 1$. In period t , agent A receives a higher endowment than agent B . How much better off is A than B going forward, as measured by the difference between their continuation utilities? Our second main economic result shows that, as t increases, the impact on continuation utilities of this last report grows without bound. We refer to it as *Relative Immiseration* because it states that the *difference* in promised utilities across *different types* of agents grows — ie, low endowment types become impoverished relative to high endowment types.

Theorem 4 (Relative Immiseration). *Suppose the environment is [TVC]-regular and the Markov process satisfies FOSD. Under the optimal contract:*

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

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

in probability.

(b) *If, in addition, Absolute 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*

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

almost surely. In particular, these conclusions hold if the Markov process satisfies PPR.

The proof of Theorem 4 is in Appendix B.5. Mathematically, Relative Immiseration is a consequence of Absolute Immiseration. In particular, the incentive constraint [IC $_{ij}^*$] with $j = i - 1$ is

$$v_i - v_{i-1} \geq \psi(u_{i-1}, i, i - 1) - u_{i-1} + \alpha (\mathbf{E}^{f_i} [\mathbf{w}_{i-1}] - \mathbf{E}^{f_{i-1}} [\mathbf{w}_{i-1}])$$

When the type process satisfies FOSD, Proposition 1 implies that the Markov information rent term is non-negative, and parts (c) and (d) of Theorem 3 imply that the iid information rent term, which is always non-negative, grows without bound (either in probability or almost surely). It follows from the above incentive constraint that the difference $v_i - v_{i-1}$ must also grow without bound. The proof formalizes this argument, and establishes the statements regarding conditional variances as easy corollaries.

pathwise properties of the optimal contract. Thus, their results are fundamentally different from part (c) of Theorem 3, which is stated in terms of convergence in probability only due to our technical limitations. The proof of part (c) actually involves establishing quite strong pathwise properties of the optimal contract, and we conjecture that the result can be strengthened from convergence “in probability” to “almost surely.”

While Absolute Immiseration in Theorem 3 is driven by the fact lower *levels* of promised utility make it cheaper for the principal to provide incentives through *variability* of (flow and promised) utility, this does not tell us *how* variable utility is in the limit. Relative Immiseration in Theorem 4 states that, as the cost of incentive provision decreases to zero, the variability of promised utility increases without bound. Intuitively, the principal does not waste the ability to provide high-powered incentives as they become affordable. This observation, while simple, has important implications for long-run inequality under the optimal contract.

In particular, Relative Immiseration formalizes the idea that “pathwise inequality” increases over time, and without bound, under the optimal contract. Any given sequence of types up through date $t - 1$ is summarized by a state variable $(\mathbf{v}^{(t)}, s^{(t)})$. 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

$$[5.1] \quad \mathbf{E} \left[v_{s^{(t+1)}}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right] = \sum_{i=1}^d f_{s^{(t)}, i} v_i^{(t)}$$

and conditional variance

$$[5.2] \quad \text{Var} \left(v_{s^{(t+1)}}^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right) = \sum_{i=1}^d f_{s^{(t)}, i} \left(v_i^{(t)} - \mathbf{E} \left[v_i^{(t)} \mid \mathbf{v}^{(t)}, s^{(t)} \right] \right)^2$$

The conditional mean is just the “ex ante” promised utility at date t (as in Thomas and Worrall (1990)), and Theorem 3 shows that it decreases without bound on the optimal path. Under the first-best (recall Section 4.1), the conditional variance is always zero because consumption is perfectly stabilized. The first half of part (b) of Theorem 4 states that, in the second-best and under FOSD, the function $i \mapsto v_i^{(t)}$ is (i) strictly increasing and (ii) becomes arbitrarily steep in the long run. Thus, idiosyncratic and transient shocks — namely, the realized endowment in period t — translate to large, permanent differences in welfare. As a consequence, the second half of part (b) of Theorem 4 says that the conditional variance in [5.2] becomes arbitrarily large in the long run. Thus, the agent’s uncertainty about his future prospects, before observing his realized endowment in a given period, increases without bound.

Thus, Relative Immiseration represents one sense in which private information leads to a severe breakdown of insurance in the long run. But variability of contingent promised utility is not the only measure of insurance quality, and Theorem 4 does *not* imply that either (i) the differences in net consumption $(c_{i+1}^{(t)} + \omega_{i+1}) - (c_i^{(t)} + \omega_i)$ or (ii) any of the “wedges” used to quantify allocative distortions (see Section 7) diverge. Indeed, neither is true. When the consumption domain is bounded below, parts (c) and (d) of Theorem 3 imply that net consumption converges to the lower bound \underline{c} . Thus, *at the limit*, net consumption is completely stable, always at its lower bound, and the difference between consumption transfers is bounded — in particular, $|c_{i+1}^* - c_i^*| = |\omega_{i+1} - \omega_i|$. In Section 7, we study the dynamics of the insurance and intertemporal wedges, which are standard measures of departures from the first-best

allocation. When the agent has CARA utility, the wedges are “scale invariant” and exhibit a strong form of mean reversion. Indeed, as we have argued above, Absolute and Relative Immiseration are driven by the fact that, to minimize costs, the principal aims to provide maximal incentives with minimal variability of consumption itself.⁵⁶ Crucially, the facts that the utility function is (i) concave and (ii) unbounded below imply that small variations in net consumption lead to very large variations in flow utility when the level of consumption is low. Thus, Theorem 4 shows that, even as net consumption converges to the constant \underline{c} , it retains enough noise to make the variance of utility explode.

We note that, while the statements about conditional variances in Theorem 4 look superficially similar to existing results in the literature, they differ in important ways. A central finding of Atkeson and Lucas (1992) is that the *unconditional* (ie, from the $t = 0$ perspective) variance of *ex ante* promised utility (as in [5.1]) grows without bound.⁵⁷ In an exact analogue of our model with CARA utility and endowment shocks that are driven by a Brownian motion with drift, Williams (2011) shows that the distribution of *ex ante* promised utility spreads out over time, as the *unconditional* variance of the Brownian motion naturally tends to infinity. Both of these results essentially say that a society’s *cross-sectional* inequality is *expected* to grow over time. Importantly, in both papers, promised utility fans out over time due to the *cumulative* effect of all previous shocks; in Williams (2011), this follows fairly mechanically from the fact that the private information process itself has unbounded unconditional variance. Moreover, both results concern *ex ante* promised utility — not promised utility contingent on type. In contrast, Theorem 4 — namely, the stronger statement in part (b) — says that, *along (almost) every path*, the impact on inequality of an *incremental* piece of private information increases without bound. To capture the notion of an incremental piece of private information, it is essential that Theorem 4 is stated in terms of promised utilities contingent on type.⁵⁸

6. Discussion: Robustness of Immiseration

6.1. Revisiting Basic Assumptions

While the baseline model laid out in Section 2 is fairly general, our analysis does rely on three types of assumptions (enumerated below). Given the fundamental nature of the questions we consider, it is important to understand the role that each plays in obtaining our results.

(56) Kocherlakota (2010, p. 70) argues that this is the general principle underlying all immiseration-type results. The general “front-loading principle” emphasized by Albanesi and Armenter (2012) is also related.

(57) To be precise, this is true when utility is logarithmic. In the other cases they consider, the variance of *appropriate functions of* promised utility explodes.

(58) This is one *conceptual* advantage of using interim variables, as opposed to *ex ante* variables as in Fernandes and Phelan (2000) or even Thomas and Worrall (1990).

Preferences and Source of Private Information: We have focused on the case in which the agent’s private information concerns his endowment, but this is not essential. Here, we mention a few other cases of particular interest.

Multiplicative taste shocks: The agent’s utility function over consumption c is $U(c, \theta) = \theta u(c)$, where θ denotes a subjective taste parameter. The principal minimizes the cost of providing consumption to the agent. This specification is adopted in many papers starting with Atkeson and Lucas (1992). When the range of $u(\cdot)$ is open, bounded above, and unbounded below, this class of models is essentially identical to our hidden endowment model with CARA utility.

Productivity shocks with separable preferences: The agent’s utility function over consumption c and labor ℓ is $U(c, \ell, \theta) = u(c) - v(\ell, \theta)$, where θ is a productivity parameter. The principal provides consumption to the agent but collects his output $\theta\ell$, and thus aims to minimize the present discounted value of consumption minus output, $c - \theta\ell$. This specification is the workhorse of the optimal dynamic taxation literature — see, eg, Golosov, Kocherlakota and Tsyvinski (2003), Albanesi and Sleet (2006), Zhang (2009), and Farhi and Werning (2013) for recent contributions, and Kocherlakota (2010) for a survey.

Productivity shocks with non-separable preferences: The agent’s utility function over consumption c and labor ℓ is $U(c, \ell, \theta)$, where θ is a productivity parameter. The principal’s objective is the same as above. This clearly generalizes the separable specification described above, and is studied in, eg, Farhi and Werning (2013) and Golosov, Troshkin and Tsyvinski (2016).

In each case, it is straightforward to formulate the contracting problem recursively as in Sections 3 and 4. Under “appropriate technical conditions,” it is also straightforward to extend parts (a), (c), and (d) of Theorem 3 by following the same proof steps almost verbatim.⁵⁹ The proof of part (b), as written, does use the fact that the utility function $U(\cdot)$ is bounded above, but it can easily be extended if the flow utility process satisfies weaker uniform integrability conditions. This is a primary virtue of the martingale approach, which isolates the driving forces and does not hinge on irrelevant functional form assumptions or other fine details of the environment.

Indeed, the “appropriate technical conditions” needed to extend Theorem 3 essentially consist only of monotonicity, concavity, smoothness, and Inada conditions embedded in part (a) of Assumption DARA. It is well known that these are nearly necessary conditions for Theorem 3, and in each of the above specifications, appropriate analogues of these conditions must be adopted for Absolute Immiseration to hold more generally.⁶⁰ On the other hand, parts (b)

(59) The first-best contracts in these models need not fully stabilize consumption, as is the case with endowment shocks (see Section 4.1), but they are still not implementable and satisfy appropriate self-generation properties (see Appendix B.3).

(60) The monotonicity and concavity properties guarantee that the first-best is not implementable, the smoothness conditions imply smoothness of the principal’s value function, and the Inada conditions are important for reasons discussed in Phelan (1998).

and (c) of Assumption [DARA](#) are purely technical (as are the regularity conditions described in Section 4.3, which we do not mention further). As has been previously noted, part (c) of Assumption [DARA](#) is used to guarantee that the space of feasible contracts is convex, and thus that (i) there is no gain to using randomized contracts and (ii) first-order optimality conditions are sufficient. These properties are automatically satisfied in the multiplicative taste shock specification, and can be similarly assumed in the productivity shock specifications.

The primary role of part (b) of Assumption [DARA](#) — in particular, the requirement that the utility function $U(\cdot)$ is unbounded below — is to ensure that the range of feasible flow utilities \mathcal{U} is open, and thereby rule out corner solutions. When corner solutions are possible, the directional derivative process may not define a martingale globally, and additional care is required. In an example with iid and multiplicative taste shocks, Golosov, Tsyvinski and Werquin (2016) show (in the proof of their Proposition 6) how to deal with these issues when $\mathcal{U} = \mathbb{R}_+$, and thus has a single boundary point. In the Markovian setting, the boundary of the largest recursive domain D is a surface in \mathbb{R}^d (typically with no closed-form description), and it thus seems unavoidable that one must rely on idiosyncratic arguments adapted to the case at hand. Of course, if one can show that the optimal contract is always interior (in an appropriate sense), then our martingale arguments apply directly. We would find it very surprising if such boundary issues, by themselves, could lead to a failure of Absolute Immiseration.

Importantly, part (b) of Assumption [DARA](#) also requires that the utility function is bounded above. This serves as a technical simplification in a few proofs, but its main role is in determining the shape of the largest recursive domain D in Theorem 1. If, instead, the utility function were bounded below and unbounded above, then D would be a cone in the positive orthant \mathbb{R}_+^d ; if it were unbounded both above and below (eg, logarithmic utility), then D would span the entire space \mathbb{R}^d . While these differences are not essential for Absolute Immiseration, they do have important implications for the robustness of Relative Immiseration. In particular, Theorem 4 — or, at least, the proof technique we employ — uses the fact that $D^* \subseteq V_d$ when the Markov process satisfies FOSD (this is part (d) of Theorem 1). For example, if instead we considered a model where the utility function was bounded *below*, then Relative Immiseration would not occur. Even with our Assumption [DARA](#) in place, it is possible that Relative Immiseration would fail when the Markov process exhibits negative serial correlation, and thus Markov information rents could be negative. Thus, we conclude that, while Relative Immiseration holds in the broad class of environments covered by Theorem 4 (and appropriate analogues in the taste- and productivity-shock specifications), it *is* sensitive to details of the type process and utility function in ways that Absolute Immiseration is not.

Markov Process: Assumption [Markov](#) plays two substantive roles. The recursive representation itself (Sections 3 and 4) plainly relies on finiteness of the type space and the Markov property. Beyond this, the joint assumption of finiteness and full connectedness is mainly for technical convenience. The existence of a highest endowment type does play a role in the

proof of Theorem 3, but finiteness of the type space itself seems less important. We conjecture that Theorem 3 holds for a larger class of ergodic Markov processes on compact type spaces, though it is not clear that such a generalization could add new insights.⁶¹ To be sure, Theorems 3 and 4 will not hold if the Markov process has absorbing states, so ergodicity is critical.⁶²

Feasible Reporting Strategies: No Hidden Borrowing (Assumption NHB) is the most substantive assumption, as it limits the agent’s set of feasible reporting strategies. While we do not view it as economically very restrictive and believe that Theorems 3 and 4 would hold without it, the proofs would need to be augmented.⁶³ We justify No Hidden Borrowing on three main grounds. First, constraints on under-reporting seem to be the relevant concern empirically — see, eg, Feldman and Slemrod (2007) for evidence from US tax data. Second, as noted in Section 2, it is a natural technological assumption on the principal’s verification technology, and is commonly assumed in a variety of settings in the literature.⁶⁴ Like the standard assumption of no hidden savings, it endows the principal with full power to save and borrow on the agent’s behalf; extensions in which No Hidden Borrowing are relaxed can be viewed as studying the interaction between the optimal mechanism and unmonitored market activity outside of the mechanism. Finally, an alternative view is that No Hidden Borrowing is, like the FOA, simply a relaxation of a “true” problem that also incorporates upward incentive constraints. In this light, No Hidden Borrowing corresponds to a *less* relaxed problem than

-
- (61) With finite types, irreducibility and aperiodicity of the chain are sufficient for ergodicity. Any Markov process can be approximated (in a variety of metrics) by finite-state, full-connected processes. Establishing convergence of policy functions — and preservation of the limiting behavior of optimal policies — is, however, much more involved.
- (62) If there did exist such an absorbing state $a \in S$ then, conditional on reaching a , the contracting problem becomes static and the optimal contract will self-generate at a single point. Thus, if a is reached with positive probability given the prior distribution over initial types, immiseration cannot possibly occur with probability one. See Golosov and Tsyvinski (2006) for a model of optimal insurance with absorbing types.
- (63) The recursive representation itself is easy to extend. The shape of the largest recursive domain D would change if we included upward incentive constraints, but could be characterized using exactly the same techniques. The martingale convergence argument behind part (a) of Theorem 3 is somewhat more delicate: in principle, it is possible that *both* upward and downward incentive constraints could *simultaneously* bind for multiple consecutive periods in such a way that martingale convergence to some interior value does not imply implementation of the first-best. Given Absolute Immiseration, Theorem 4 would continue to hold.
- (64) See Phelan (1998) for an insurance model with endowment shocks and Golosov and Tsyvinski (2007) for a taxation model with productivity shocks. Fernandes and Phelan (2000) argue that No Hidden Borrowing is essentially without loss of generality when the consumption domain is bounded below, and rely on this assumption in their numerical examples. No Hidden Borrowing is essentially built into cash-flow diversion models of firm financing — see, eg, Clementi and Hopenhayn (2006), DeMarzo and Sannikov (2006), and Fu and Krishna (2017). Williams (2011), in insurance models of endowment shocks and taste shocks, relies on a stronger notion of No Hidden Borrowing (even when verifying the validity of the FOA). Namely, the agent cannot over-report the *difference* between today’s and yesterday’s types — even if he had been under-reporting in the past, so that an over-report of the difference would still correspond to an under-report of the type itself.

the FOA, as it incorporate *more* constraints, and will thus tend to be “valid” in a larger set of environments. Indeed, all results in the text are correct, essentially verbatim, if we replace No Hidden Borrowing with the FOA.⁶⁵

6.2. Comparison to Zhang (2009) and Williams (2011)

Zhang (2009) and Williams (2011) — substantively, the two closest papers to ours — reach strikingly different conclusions about long-run properties of the optimal contract. This raises doubts as to whether, or to what extent, the main intuitions from insurance models with iid types — especially the forces underlying the immiseration result — extend to settings with persistent private information. Theorem 3 helps to clarify these puzzling findings.

Zhang (2009) studies a dynamic taxation model that, aside from being cast in continuous time, is a special case of “separable productivity shock” extension of our model discussed in Section 6.1. He focuses on the special case in which the agent’s private information follows a two-state Markov jump process with symmetric transition probabilities,⁶⁶ and shows that (absolute) immiseration occurs. Williams (2011) studies two examples — one of endowment shocks as in Thomas and Worrall (1990), and one of taste shocks as in Atkeson and Lucas (1992) — in which the type process follows a Brownian motion with drift.⁶⁷ In both examples, (absolute) immiseration does *not* occur, and consumption actually *drifts up* over time — the polar opposite conclusion. Both papers focus on the cases of CRRA or CARA utility, and rely heavily on the implied homogeneity properties to derive their long-run results by characterizing the optimal contract in (nearly) closed form.

Both papers attribute these differences to whether or not the *inverse Euler equation* (IEE) holds; it does in Zhang (2009) but does not in Williams (2011). But the IEE by itself is neither necessary nor sufficient for immiseration-type results. Golosov, Kocherlakota and Tsyvinski (2003) show that the IEE holds for essentially arbitrary private information processes when the agent’s marginal utility for consumption is independent of his private information. But they also point out that this separability is essentially a necessary condition for the IEE, and therefore it does not hold in the class of models considered here (in the baseline case) or in the examples of Williams (2011). Thus the IEE cannot be necessary for immiseration-type results. As for sufficiency, note that the IEE says *nothing* about the “non-monetary” part of

(65) We do not spell out the details in the appendix, but the reader can easily verify this fact. In our setting, the FOA corresponds to including only the “local downward” incentive constraints, ie, $[IC_{ij}]$ for $j = i - 1$. In the iid setting of Thomas and Worrall (1990), only local downward constraints bind but, as suggested by Battaglini and Lamba (2018), that needn’t be true when types are highly persistent.

(66) This is the continuous-time limit of our model when $d = 2$ and $f_{11} = f_{22}$. Such symmetry of transitions is non-generic for any number of types.

(67) Williams (2011) also studies an example where endowments follow an OU process, but Strulovici (2011) points out that Williams’ purported optimal contract in this case is actually strictly dominated by the optimal renegotiation-proof contract.

the allocation (labor, in taxation models). Thus, although the martingale convergence theorem applies to the inverse marginal-utility-of-consumption process under the IEE, this has no direct implications for the rest of the allocation or the promised utility process.⁶⁸

Given the restriction to Markovian type processes, the differential martingale from Proposition 5.1 is the appropriate generalization of the IEE: it embodies the same principle of optimal cost-smoothing for the principal, reduces to the IEE when appropriate separability conditions hold, and is a consequence of the envelope theorem that *always* holds (ie, for essentially any preference specification, subject to regularity conditions).⁶⁹ The arguments underlying Theorem 3 demonstrate that basic cost-smoothing considerations, embodied by the differential martingale, are what drives (absolute) immiseration. Moreover, the discussion above in Section 6.1 illustrates that these mechanisms are robust to many details of the contracting environment.

What, then, leads to such divergent results? The hidden endowment example in Williams (2011) has CARA utility, No Hidden Borrowing, and MLRP-ordered transitions — just like a special case of our baseline model — but differs in that the type process (i) is modeled a continuous-time diffusion, (ii) is unbounded and continuously distributed, and (iii) is not ergodic (though it is recurrent) and has non-vanishing “impulse response functions” (in the language of Pavan, Segal and Toikka (2014)) that are constant and equal to one. Which of these differences is critical? It is straightforward to verify that the assumption of a finite type space in the classic iid setting of Thomas and Worrall (1990) is not essential, and that their (absolute) immiseration result extends to at least some unbounded type spaces under mild technical conditions.⁷⁰ Thus, unboundedness alone does not appear to be critical. Moreover, Strulovici (2011), considers the same hidden endowment model as Williams (2011) but assumes that types follow an OU process (with non-trivial mean reversion), which satisfies points (i) and (ii) but not point (iii). (OU processes are ergodic and have exponentially-decaying impulse responses.)

(68) This is the primary reason why we focus on Markovian types and rely on recursive methods. The variational arguments used in Golosov, Kocherlakota and Tsyvinski (2003) to obtain the IEE are not, to our knowledge, sufficient to prove our theorems.

(69) Williams (2011) uses a recursive approach with a two-dimensional state variable consisting of *ex ante* promised utility and *marginal* promised utility based on the FOA. In that setting, the appropriate analogue of the differential martingale is *the partial derivative of the principal’s value function with respect to ex ante promised utility*, which represents the same kind of perturbation whereby information rents are held fixed. It is easy to verify that this partial derivative process defines a martingale under the optimal contract. This is also true in other models based on the FOA (both in discrete and continuous time), such as Farhi and Werning (2013) and Kapička (2013), among others.

(70) Consider, for example, a variant of their model where the agent has CARA utility and the set of endowment levels is countably infinite and unbounded. Suppose that the distribution over types in each period has full support, so that the type process is ergodic and each type is (positive) recurrent. Suppose also that (a) an optimal contract exists and is continuous in ex ante promised utility, and (b) the value function is well-defined, finite-valued, strictly concave, and continuously differentiable. Then essentially the same proof (eg, a slight variant on the proof of Proposition 6 in Golosov, Tsyvinski and Werquin (2016)) can be used to establish (absolute) immiseration.

He derives the optimal renegotiation-proof contract and shows that absolute immiseration occurs. When there is no mean-reversion — so the endowment process is a Brownian motion with drift — the optimal renegotiation-proof contract coincides with the full-commitment optimal contract derived in Williams (2011).⁷¹ Thus, it also seems unlikely that continuous time or continuously-distributed types are critical.⁷²

The key difference therefore appears to be point (iii). We have already noted in Section 6.1 that ergodicity of the type process is important for Absolute Immiseration. Impulse responses that are constant and equal to one imply that a change in today’s type has a permanent effect on all future types. When the process is ergodic, as is the case in our setting, those effects decay to zero in the long run — ie, impulse responses vanish asymptotically. In quasi-linear settings (such as the monopolistic screening problem), it is known that the difference between vanishing and non-vanishing impulse responses is often a key determinant for the long-run properties of allocative distortions.⁷³ Thus, it is not entirely surprising to see such differences emerge in the insurance setting as well. We believe that further and more definitive exploration of these points is an important direction for future research.

7. Short-Run Dynamics with CARA Utility

In this section, we turn to the short-run dynamics and insurance properties of the optimal contract. For simplicity, we focus on the case in which the agent has CARA utility. We present both analytical properties and numerical results.

To obtain an explicit description of the optimal contract, we specialize to the case of binary and persistent types, ie, $d = 2$ and the Markov process satisfies MLRP. (All proofs are in Appendix C. All of the analytical properties, and corresponding intuitions, extend to the general $d \geq 2$ case under MLRP.) Recall from Lemma 4.2 and Theorem 1 that, in this case, the environment is regular and $D = D^* = V_2$. We also assume [TVC]-regularity throughout.

The agent’s utility function takes the exponential form $U(c) = -e^{-\rho c}$, where $\rho > 0$ is the coefficient of absolute risk aversion. It is convenient to define $\theta_i := e^{-\rho \omega_i}$ for each $i \in S$, so that $U(\omega_i + c_i) = \theta_i U(c_i)$.⁷⁴ In the numerical computations, we fix the preference parameters

(71) Strulovici (2011) does not state any results in the case without mean-reversion, but it is straightforward to verify this claim by following the calculations in the paper.

(72) The notion of renegotiation-proofness in Strulovici (2011) essentially requires the contract to be efficient (ie, solve the analogue of our [Eff_i]) after every history. Thus, many of the subtleties discussed in the proof sketch of Theorem 3 in Section 5.2 do not arise. It is therefore difficult to *directly* compare his results to Theorem 3, which considers the full-commitment optimal contract.

(73) Bergemann and Strack (2015) derive closed-form solutions in several continuous-time examples that illustrate this point, and also reference the key discrete-time precedents. Garrett, Pavan and Toikka (2018) emphasize the importance of ergodicity as a sufficient condition for (a) distortions to vanish in the long run and (b) impulse response functions (more precisely, what they call “marginal handicaps”) to vanish in the long run, in expectation.

(74) It is a standard observation that, given CARA utility, the model with hidden endowments is equivalent to

$\alpha, \rho, \theta_1, \theta_2$ and vary the transition probabilities $\mathbf{f}_1, \mathbf{f}_2$. We focus on the case of symmetric transitions (ie, $f_{11} = f_{22}$) as in Zhang (2009), and consider three different persistence levels: iid, medium persistence, and high persistence.⁷⁵

7.1. Promised Utility Dynamics

We begin by analyzing the dynamics of promised utility, turning to properties of the flow allocation in the next section. The main simplification implied by CARA utility is that the optimal contract is “scale invariant:”

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

Thus, the range of solutions to the efficiency problem $[\mathbf{Eff}_i]$ defines a ray $E_i \subset V_2$ ($i = 1, 2$).⁷⁶ Analytically, we can partially describe (i) the dynamics of promised utility relative to these rays, and (ii) the monotonicity properties of promised utility after different types of reports. These properties are summarized in Facts 2 and 3 below.

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 $\mathbf{v} \in E$, $\mathbf{w}_1(\mathbf{v}, s) \ll \mathbf{v} \ll \mathbf{w}_2(\mathbf{v}, s)$.*

If the initial $\mathbf{v} \notin E$, then it jumps on to E after the first period, regardless of whether the initial report is high or low, and then stays on E forever after.⁷⁷ Thus, as discussed in Section 5.1, the promised utility vector is *always* efficient (on the optimal path) when types are iid. Moreover, once on E , promised utility increases (in both components) after high reports and decreases after low reports. This is the standard *co-insurance* property, whereby the agent is effectively punished in the future for claiming to have a low endowment and rewarded for admitting he has a high endowment.

These properties are illustrated in Figure 3. The figure also makes clear (i) that the policy function $\mathbf{w}_2(\mathbf{v}, s)$ depends only on component v_2 , and (ii) that when \mathbf{v} lies *above* E , the policy function $\mathbf{w}_1(\mathbf{v}, s)$ depends only on the component v_1 . Point (i) is a necessary consequence of

one with multiplicative taste shocks.

- (75) 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.
- (76) E_i is the pre-image, under the derivative mapping $DP(\cdot, i)$, of the ray \tilde{E}_i described in Section 5.2 and depicted in Figure 2.
- (77) Clearly, the ray E can be parameterized by expected promised utility as in Thomas and Worrall (1990).

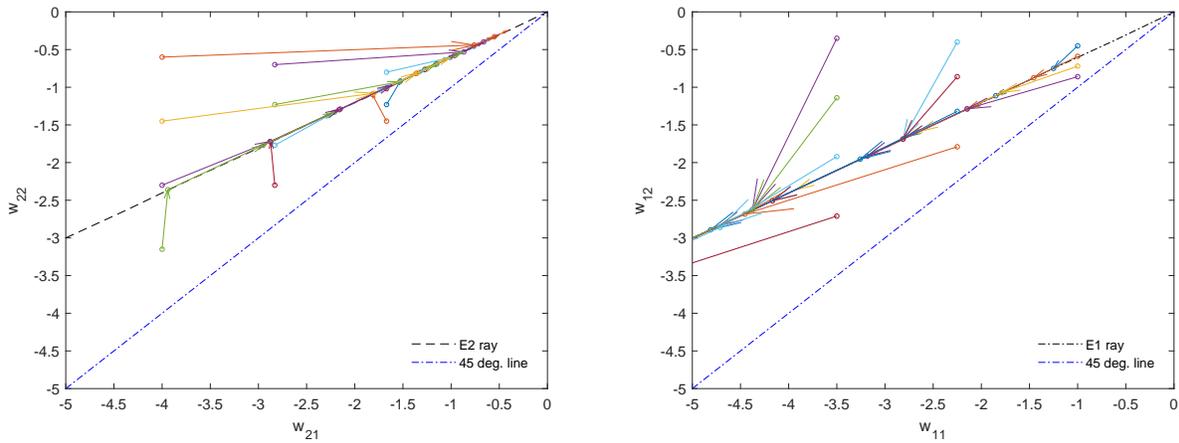


Figure 3: IID. Promised utility dynamics after consecutive high reports (left) and consecutive low reports (right).

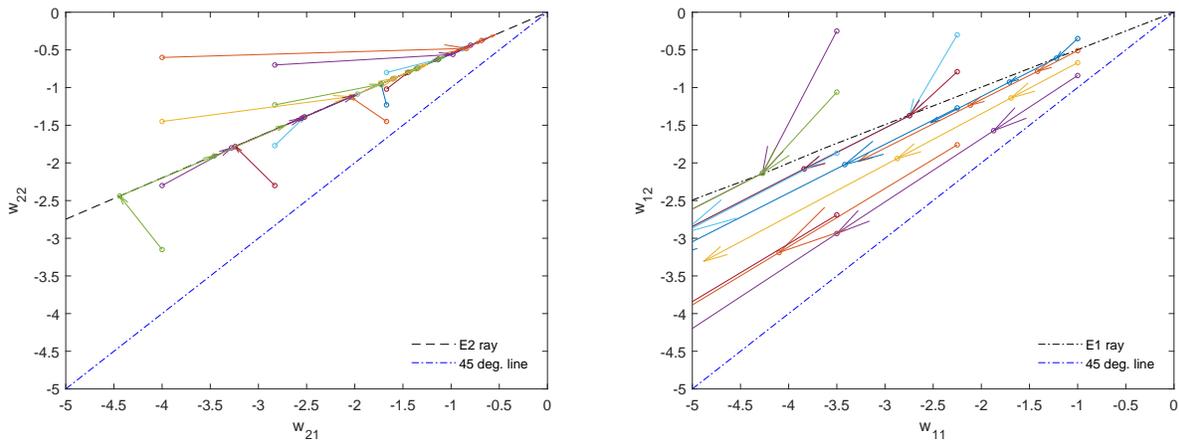


Figure 4: Medium persistence. Promised utility dynamics after consecutive high reports (left) and consecutive low reports (right).

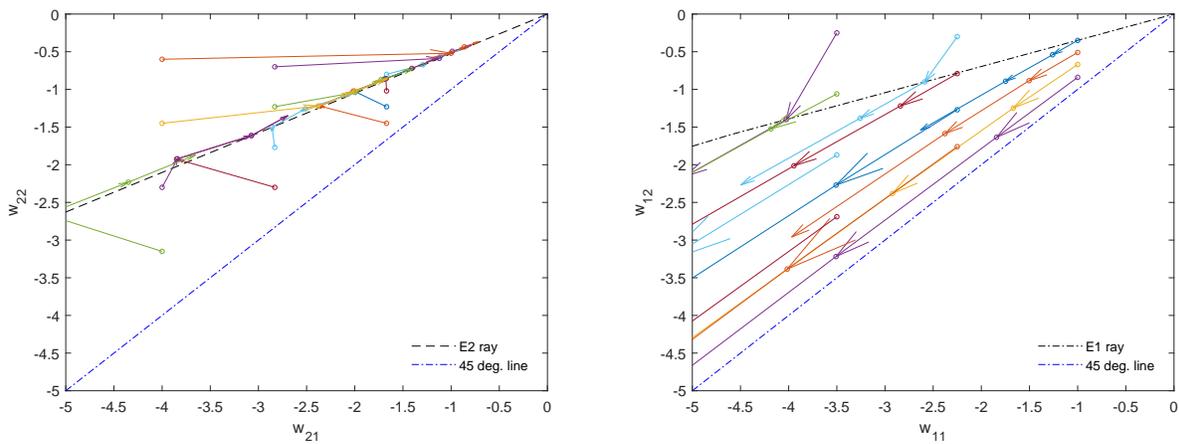


Figure 5: High persistence. Promised utility dynamics after consecutive high reports (left) and consecutive low reports (right).

parts (c) and (e) of Theorem 2. Point (ii) reveals that, above E , $v_2 - v_1$ is large enough that the incentive constraint $[\mathbf{IC}_{21}^*]$ does not bind. Thus, in this region, \mathbf{w}_1 effectively enters into only the promise keeping constraint $[\mathbf{PK}_i]$ ($i = 1$).

Fact 3. *With positive persistence:*⁷⁸

- (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 .*

With persistence, the efficient rays do not coincide. E_1 lies above E_2 because a high (low) report is less (more) likely under the measure \mathbf{f}_1 than under \mathbf{f}_2 , and 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. As discussed in Section 5.1, \mathbf{w}_2 does not contribute to any Markov information rents, and thus the underlying tradeoffs are very similar to those in the iid case. Part (d), however, only arises under positive persistence. The intuition comes from the Markov information rent term in $[\mathbf{IC}_{21}^*]$, which may be written as $\alpha(f_{22} - f_{12}) \cdot (w_{12} - w_{11})$, where $f_{22} > f_{12}$. In order to relax the incentive constraint, 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 .⁷⁹ Economically, the principal uses persistence to screen types by committing to punish high-type deviators who mis-report as low types for a single period.

These features are illustrated in Figures 4 and 5, which also elucidate two additional features of the optimal contract (that we cannot establish analytically).⁸⁰ First, note 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 after consecutive low reports, and then jumps back to E_2 after the next high report. Thus, on the optimal path, we always have $\mathbf{w}_1(\mathbf{v}, 1) \ll \mathbf{v}$, which extends the monotonicity properties from the iid case. These dynamics closely mirror those in Zhang (2009), who also restricts attention to symmetric transitions.

Second, we see that the promised utility vector *moves southwest faster* after low reports when *persistence is higher*. 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

(78) That is, when $f_{11} - f_{21} \equiv f_{22} - f_{12} > 0$.

(79) This is consistent with intuitions in Kapička (2013) and Williams (2011), who describe how the principal (inefficiently) lowers *marginal* promised utility to provide incentives.

(80) As in the iid case, we see that the incentive constraint $[\mathbf{IC}_{21}^*]$ is slack in the region above E_1 . In general, characterizing the dynamics after consecutive low reports is very difficult.

achieve the same schedule of information rents *in the current period*. This is a force that tends to move \mathbf{w}_1 toward the diagonal. But this also leads to a tighter incentive constraint (ie, the left-hand side of $[\mathbf{IC}_{21}^*]$) *in the next period*. To compensate for this additional incentive cost, the principal pushes the level of tomorrow's promised utility down further.⁸¹ We can see this more precisely 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

$$[7.1] \quad \underbrace{\frac{P_1(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},1}} - \frac{P_2(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},2}}}_{\text{MC of compression}} = \mathbf{1}(s^{(t)} = 1) \cdot \underbrace{\frac{(f_{11} - f_{21})}{f_{s^{(t)},1} \cdot f_{s^{(t)},2}} \cdot \frac{\mu_{21}(\mathbf{v}^{(t-1)}, s^{(t-1)})}{f_{s^{(t-1)},s^{(t)}}}}_{\text{MB of compression}}$$

The MC of compression is the *Marginal Cost* of moving the promised utility vector toward the diagonal while holding ex ante promised utility $\mathbf{E}^{\mathbf{f}_{s^{(t)}}}[\mathbf{v}^{(t)}]$ 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 (ie, convexity of P), it is optimal for the principal to have promised utility drift away from E_1 gradually (ie, over multiple periods) 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,⁸² 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 appears to be borne out in the numerical examples.

The need to manage Markov information rents in this manner also introduces a new form of history-dependence relative to the iid case. An important result in Thomas and Worrall (1990) is that, when types are iid and utility is CARA, the *order* of reported types does not matter:⁸³

Fact 4. *Suppose that types are iid and the initial $\mathbf{v}^{(0)} \in E$. Then, $\mathbf{v}^{(t)}(h)$ depends on the realized length- t history of reports $h^t = (s^1, \dots, s^t)$ only through the **number** of high and low reports, but **not the order** in which they were made.⁸⁴*

This is easy to see from Facts 1 and 2. Because the policy functions are HD1 and promised utility travels along the single ray E , there must exist numbers $b_1, b_2 > 0$ such

(81) This does not necessarily imply that *ex ante* promised utility decreases more quickly when persistence is greater, as the slopes of the level sets of the form $\{\mathbf{v} : \mathbf{E}^{\mathbf{f}_s}[\mathbf{v}] = M\}$ vary with the transition probabilities. But, at least in these parametrized examples, it is the case that *ex ante* promised utility decreases more quickly with persistence. This can be seen in the right-hand panels of Figures 6–8, where the “Ev” series illustrate the dynamics of $\mathbf{E}^{\mathbf{f}_{s^{(t)}}}[\mathbf{v}^{(t)}]$.

(82) We show in [C.12] that this weighted multiplier does not depend on the value of $s^{(t-1)}$.

(83) This fact is also implicit in Green (1987).

(84) Recall that s^n denotes the type of report made in period $n - 1$, which then becomes part of the period n state variable.

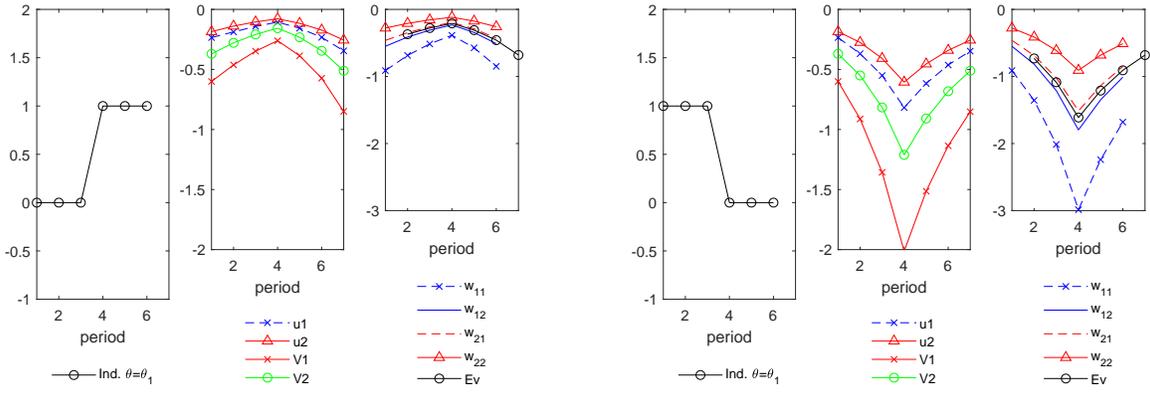


Figure 6: IID. Short-run dynamics along two report sequences.

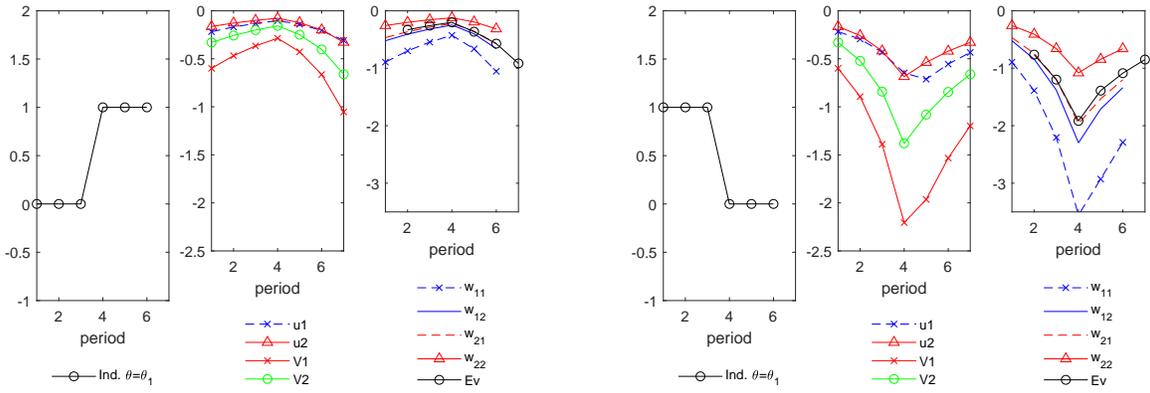


Figure 7: Medium persistence. Short-run dynamics along two report sequences.

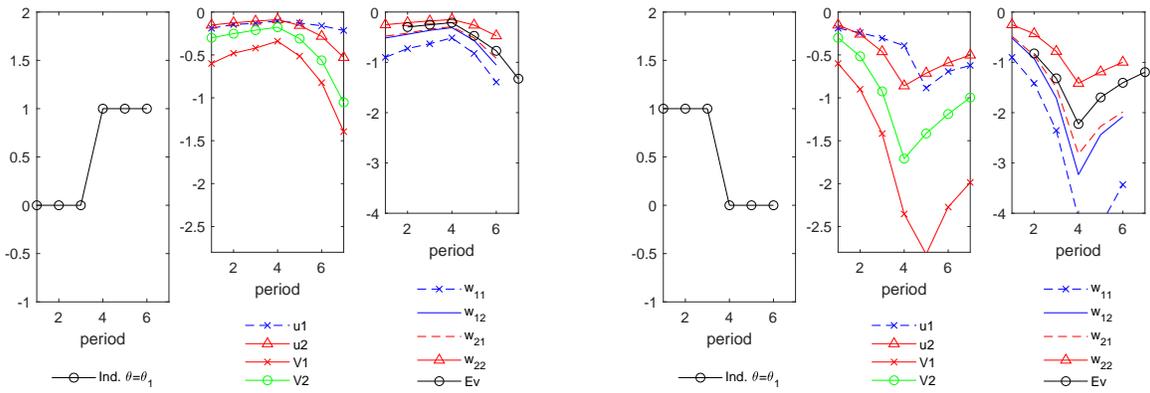


Figure 8: High persistence. Short-run dynamics along two report sequences.

that $w_i(\mathbf{v}, s) = b_i \cdot \mathbf{v}$ whenever $\mathbf{v} \in E$. Thus, under the hypotheses of Fact 4, it must be that $\mathbf{v}^{(t)} = b_1^{N^{(t)}} \cdot b_2^{t-N^{(t)}} \cdot \mathbf{v}^0$ where $N^{(t)}$ is the (random) number of low reports in periods $0, 1, \dots, t-1$. Thus, the order of reports does not matter. This 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.⁸⁵ 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 $u_i(\mathbf{v}^{(7)}, s^{(7)})$.

The intuition for this order-independence result in Thomas and Worrall (1990) is based on the absence of wealth effects under CARA utility. But, as illustrated in Figures 7 and 8, this result also depends critically on the assumption of iid types. 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.⁸⁶ This appears to be 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 $w_{22}(\mathbf{v}, 1) - w_{21}(\mathbf{v}, 1)$. We have argued above that, with persistence, the incentive constraint effectively tightens after consecutive low reports. Here, we see that it is evidently 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.

7.2. Wedge Dynamics

Next, we 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. While not emphasized in the early work on risk sharing, these are the main objects of study in the literature optimal dynamic taxation.

Definition 7.1. The *wedges* are defined as follows:⁸⁷

-
- (85) The initial condition for \mathbf{v} is drawn from the ray E_2 in each of Figures 6–8.
 - (86) 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.
 - (87) These definitions are the same as those in Zhang (2009). Some papers use slightly different, but equivalent, definitions of the wedges. For example, Farhi and Werning (2013) refer to $W_{int}/(W_{int} + 1)$ as the “inter-temporal wedge,” and Golosov, Troshkin and Tsyvinski (2016) refer to the same object as the “savings distortion.”

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

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

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

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

where $\mathbf{w}_i := \mathbf{w}_i(\mathbf{v}, s)$.

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, $C(u_1, 1) + \omega_1 < C(u_2, 2) + \omega_2$. The intertemporal wedge is positive if and only if the optimal contract acts as an implicit tax on savings. Under CARA utility, the insurance wedge at (\mathbf{v}, s) simplifies to

$$[7.2] \quad W_{\text{ins}}(\mathbf{v}, s) = \frac{u_1(\mathbf{v}, s)}{u_2(\mathbf{v}, s)} - 1$$

and the intertemporal wedge at (\mathbf{v}, s) and given report i simplifies to

$$[7.3] \quad W_{\text{inter}}(\mathbf{v}, s, i) = \frac{\mathbf{E}^{f_i} [u_j(\mathbf{w}_i, i)]}{u_i(\mathbf{v}, s)} - 1$$

where, again, $\mathbf{w}_i := \mathbf{w}_i(\mathbf{v}, s)$. Thus, $W_{\text{ins}}(\mathbf{v}, s) > 0$ if and only if $u_2(\mathbf{v}, s) > u_1(\mathbf{v}, s)$ and $W_{\text{inter}}(\mathbf{v}, s, i) > 0$ if and only if $u_i(\mathbf{v}, s) > \mathbf{E}^{f_i} [u_j(\mathbf{w}_i, i)]$, ie, flow utility has a “negative drift” conditional on (\mathbf{v}, s, i) .

The following two basic facts follow from homogeneity properties of the optimal contract and the first-order optimality conditions.

Fact 5. *When types are iid, the process $(W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)}))_{t=0}^{\infty}$ is constant and strictly positive. The process $(W_{\text{inter}}(\mathbf{v}^{(t)}, s^{(t)}, s^{(t+1)}))_{t=0}^{\infty}$ takes on at most two values, depending on whether $s^{(t+1)} = 1$ or $s^{(t+1)} = 2$.*

Fact 6. *When types are persistent, $W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)}) = W_{\text{ins}}(\mathbf{v}^{(t+1)}, s^{(t+1)}) > 0$ whenever $s^{(t)} = s^{(t+1)} = 2$. Similarly, $W_{\text{inter}}(\mathbf{v}^{(t)}, s^{(t)}, s^{(t+1)}) = W_{\text{inter}}(\mathbf{v}^{(t+1)}, s^{(t+1)}, s^{(t+2)})$ (not necessarily > 0) whenever $s^{(t+1)} = s^{(t+2)} = 2$.*

In short, the insurance wedge is constant *and positive* whenever promised utility moves along a single efficiency ray, as is the case (i) when types are iid, or (ii) when types are persistent and there are consecutive high reports. The intertemporal wedges is constant after consecutive high reports and, in the iid case, it depends *only* on the current report.

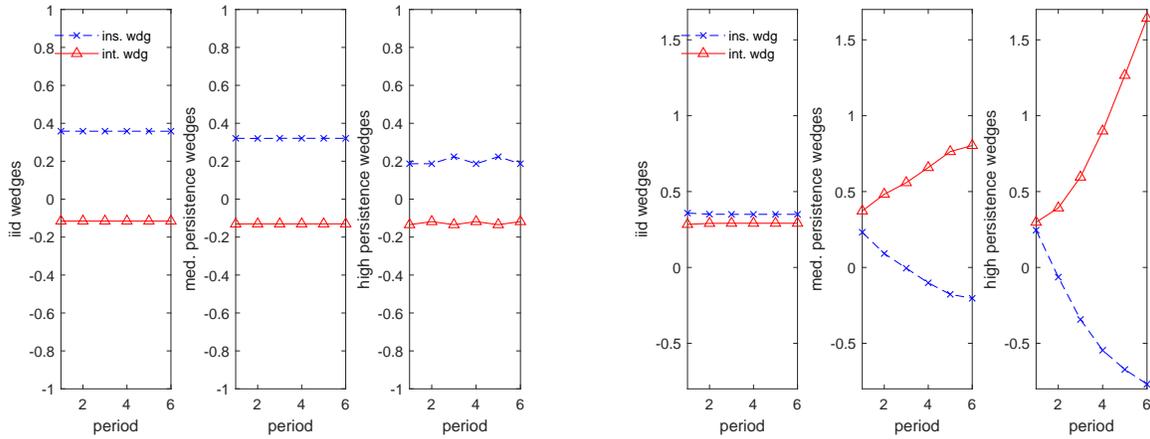


Figure 9: Wedge dynamics after consecutive high reports (left) and consecutive low reports (right).

Figure 9 illustrates the time series behavior of the wedges at each of three persistence levels. Facts 5 and 6 are clearly borne out in the figure.⁸⁸ Several other salient features warrant discussion. First, the wedges are quantitatively similar across the persistence levels when promised utility is efficient (ie, after consecutive high reports). The intertemporal wedge is *negative* in this case. Second, relative to the iid benchmark of Thomas and Worrall (1990), persistence induces qualitatively new wedge dynamics after low endowment reports. The intertemporal wedge *increases* while the insurance wedge *decreases* after low reports, and the insurance wedge quickly becomes *negative*.⁸⁹ Moreover, the wedges increase/decrease more quickly at higher persistence levels. In the high persistence case, the intertemporal wedge increases by an order of magnitude, and the absolute value of the insurance wedge more than doubles, after just 6 consecutive low reports. This last feature is (very) broadly consistent with Zhang (2009), who shows that average wedges are much larger with persistence types. But the insurance wedge is always positive in Zhang (2009), which is very different from our findings.⁹⁰

We now explain the intuition behind these features, considering each in turn.

Insurance wedge: The salient features of the insurance wedge are driven by the same forces that determine the dynamics of promised utility. It is useful to observe that, on the optimal path, the flow utilities and Lagrange multipliers (recall [7.1]) satisfy

-
- (88) In the left-hand panel, there are 6 consecutive high reports, while in the right-hand panel there are 6 consecutive low reports. The initial condition for \mathbf{v} in each case is drawn from the appropriate ray E_2 . The small bit of variation in the high persistence, high report figure is due to numerical error.
- (89) Given the simplified expression for the insurance wedge in [7.2], we can also see from Figures 6–8 that the insurance wedge may be negative in the persistent cases.
- (90) A more detailed comparison is difficult because he does not study time-series behavior of the wedges and, as discussed below, the wedges behave in fundamentally different ways in separable settings.

$$\begin{aligned}
& \underbrace{C' \left(u_2(\mathbf{v}^{(t)}, s^{(t)}), 2 \right) - C' \left(u_1(\mathbf{v}^{(t)}, s^{(t)}), 1 \right)}_{\text{direct MC of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})} = \\
[7.4] \quad & \underbrace{\left(\frac{f_{s^{(t)},1}}{f_{s^{(t)},2}} + \frac{\theta_2}{\theta_1} \right) \cdot \frac{\mu(\mathbf{v}^{(t)}, s^{(t)})}{f_{s^{(t)},1}}}_{\text{MB of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})} - \underbrace{\mathbf{1}(s^{(t)} = 1) \cdot \frac{(f_{11} - f_{21})}{f_{s^{(t)},1} \cdot f_{s^{(t)},2}} \cdot \frac{\mu_{21}(\mathbf{v}^{(t-1)}, s^{(t-1)})}{f_{s^{(t-1)},s^{(t)}}}}_{\text{indirect MC of } \uparrow W_{\text{ins}}(\mathbf{v}^{(t)}, s^{(t)})}
\end{aligned}$$

Increasing the insurance wedge always leads to a direct MC (on the left-hand side of [7.4]), 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 [7.4]. Notice that this term is identical to the MB of compression (of contingent promised utilities) in [7.1]. 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 $u_2(\mathbf{w}_1(\mathbf{v}^{(t-1)}, s^{(t-1)}), 1) - u_1(\mathbf{w}_1(\mathbf{v}^{(t-1)}, s^{(t-1)}), 1)$. 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 Section 7.1, the indirect MC term in 7.4 suggests that these effects should be larger when persistence is higher. Indeed, this is precisely what we see in the computed numerical examples.

It is useful to compare these insurance wedge dynamics to the findings of Farhi and Werning (2013), who study the dynamics of the labour wedge in (both separable and non-separable) productivity shock models of optimal dynamic taxation. Their central finding is that, roughly speaking, the labour wedge inherits the same mean-reversion properties as the productivity shock process. While direct comparisons are difficult given the differences in settings,⁹¹ our results seem to be broadly consistent with theirs. In particular, our insurance wedge *does* exhibit a kind of mean reversion: it becomes negative after consecutive low reports, but becomes positive again (and the process “renews”) after a high report. Moreover, the indirect MC term in [7.4] — and thus the size of the insurance wedge — depends directly on the degree of persistence $f_{11} - f_{21}$, which is akin to their result that the labour wedge mean-reverts at the same rate as the underlying type process.

(91) Aside from considering a different class of preferences, source of private information, and type of wedge, they also assume that types are continuously distributed and pay special attention to the case where the logarithm of the productivity parameter follows an OU process.

Intertemporal wedge: 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 should expect his consumption to have “positive drift” in the short run; conversely, receiving a low endowment should imply consumption to have short run “negative drift.” Figures 6–8 show that this is indeed the case in our examples. Thus, 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 negative at low endowment reports because the agent wants to *save* in the face of a downward-sloping consumption profile.⁹²

The fact that the intertemporal wedge increases after consecutive low reports is a consequence of the same forces underlying the dynamics of promised utility discussed above. To optimally manage Markov information rents, consumption has a *more negative* drift after low reports when persistence is higher. This is compounded by the fact that persistence implies that the agent puts more weight on the event of having another low endowment shock in the next period.

This contrasts sharply with the separable productivity shock model (as in Zhang (2009)), which always has a positive intertemporal wedge due to the IEE.⁹³ Golosov, Kocherlakota and Tsyvinski (2003) also emphasize that negative intertemporal distortions can arise in pure insurance models (though their example is based on an extreme form of persistence), while Golosov, Troshkin and Tsyvinski (2016) make a similar observation for non-separable productivity shock models.⁹⁴ The difference is that, in the separable model, consumption utilities serve as a “type-independent numeraire,” much like money in quasi-linear models. The IEE represents the optimal allocation of this numeraire over time, given the principal’s convex cost function (the inverse of consumption utility). Outside of the separable model, the timing of consumption is an important screening device, so cost smoothing and incentive provision are inextricably linked.

(92) The reader can verify that, in the iid case, this feature is implied by the analytical solution from Proposition 6 in Thomas and Worrall (1990), at least when endowments are variable enough (ie, when $\omega_2 - \omega_1 > 0$ is sufficiently large). Details are available upon request.

(93) It seems that many intuitions in the broader literature are based on this positive intertemporal wedge result, even though it is special to the separable model. Albanesi and Armenter (2012), for example, emphasize the importance of “permanent intertemporal distortions,” meaning an intertemporal wedge that does not change sign along the optimal path. Clearly, our pure insurance model does not have this feature.

(94) In particular, it involves increasing impulse response functions (in the language of Pavan, Segal and Toikka (2014)). Similarly, the intertemporal wedge is *always* negative in the hidden endowment example in Williams (2011), where impulse responses are constant. (He finds that consumption has positive drift, and marginal utility is a convex function of consumption.) Golosov, Troshkin and Tsyvinski (2016) show that “lifetime” savings distortions must be positive in a broad class of non-separable productivity shock models, but this result hinges on the assumption of a finite time horizon.

7.3. Intertemporal Cost-Smoothing

It is useful to take a more detailed look at how (i) persistence of private information and (ii) the absence of a “type-independent numeraire” shape the principal’s intertemporal cost smoothing at the optimum. With CARA utility, $\psi'(u, 2, 1) = \theta_2/\theta_1$. Fix some (\mathbf{v}, s) and consider the cost of increasing ex ante promised utility $\mathbf{E}^{\mathbf{f}_s}[\mathbf{v}]$ by some amount $\varepsilon > 0$. Recall from Section 5.1 that the martingale property of $D_1 P(\mathbf{v}, s)$ in Proposition 5.1 is based on perturbing each of the \mathbf{w}_i in direction $\varepsilon \cdot \mathbf{1}$ while leaving the flow utilities u_i unchanged. We consider two different perturbations here.

The first perturbation, which also changes \mathbf{v} to $\mathbf{v}' := \mathbf{v} + \varepsilon \mathbf{1}$, is to increase each of the flow utilities u_i by ε but leave the \mathbf{w}_i unchanged. This must be locally optimal by the envelope theorem, giving us the formula

$$[7.5] \quad D_1 P(\mathbf{v}, s) = \underbrace{\mathbf{E}^{\mathbf{f}_s} [C'(u_i, i)]}_{\text{direct cost}} - \underbrace{\left(1 - \frac{\theta_2}{\theta_1}\right) \cdot \mu(\mathbf{v}, s)}_{\text{incentive cost}}$$

where $\mu(\mathbf{v}, s) \geq 0$ is a Lagrange multiplier on the incentive constraint $[\mathbf{IC}_{ij}]$ ($i = 2, j = 1$). Notice that, because $C'(u_i, i) = 1/U'(c_i + \omega_i)$, the first “direct cost” term on the right is actually the agent’s *expected inverse marginal utility*. If expected inverse marginal utility were a martingale, it would be reminiscent of (but distinct from) the IEE. However, this object does *not* satisfy the martingale property because the underlying perturbation gives different deviation incentives to each type. To understand the second “incentive cost” term on the right, note that this perturbation increases the left-hand side of the incentive constraint $[\mathbf{IC}_{ij}]$ ($i > j$) by ε while increasing the right-hand side by approximately $\psi'(u_j, i, j) \cdot \varepsilon = (\theta_i/\theta_j) \cdot \varepsilon < \varepsilon$. Thus, it adds total slack $1 - \theta_i/\theta_j$ to $[\mathbf{IC}_{ij}]$, and the second term in [7.5] collects the marginal (shadow) cost reduction achieved by relaxing the incentive constraints in this way.

Second, there is a different way to perturb (only) flow utilities that leaves the slack in *the current period’s* incentive constraints invariant. Consider perturbing each u_i to $u'_i := u_i + \varepsilon \theta_i / \mathbf{E}^{\mathbf{f}_s}[\theta_j]$. It is easy to see that this increases each side of the $[\mathbf{IC}_{ij}]$ equally and perturbs \mathbf{v} to $\mathbf{v}' := \mathbf{v} + \varepsilon \cdot \hat{\theta}_s$, where for each $s \in S$, the vector $\hat{\theta}_s \in \mathbb{R}_{++}^d$ has i th component $\theta_i / \mathbf{E}^{\mathbf{f}_s}[\theta_j]$. Clearly, $\mathbf{E}^{\mathbf{f}_s}[\mathbf{v}' - \mathbf{v}] = \varepsilon$, as desired. The marginal cost of such a perturbation is, after a short calculation,

$$[7.6] \quad D_{\hat{\theta}_s} P(\mathbf{v}, s) = \underbrace{\mathbf{E}^{\mathbf{f}_s} [C'(u_i, i)]}_{\text{(expected) inverse marginal utility}} + \underbrace{\frac{\text{Cov}^{\mathbf{f}_s}[\theta_i, C'(u_i, i)]}{\mathbf{E}^{\mathbf{f}_s}[\theta_i]}}_{\text{incentive provision}}$$

The second term on the right vanishes under the first-best, where the flow utility process is constant, and thus measures the cost of incentive provision.

How are these perturbations connected? It turns out that they are both optimal when types are iid, but the second is not optimal when types are persistent. Recall the cost function

$K(v, s)$ from the efficiency problem [Eff_i] from Section 5.1, which coincides with the value function from Thomas and Worrall (1990) when types are iid.

Lemma 7.2. Suppose that the agent has CARA utility and the Markov process satisfies MLRP. Let $(\mathbf{v}^{(t)}, s^{(t)})_{t=0}^{\infty}$ be the process induced by the optimal contract. Let $v^{(t)} := \mathbf{E}^{\mathbf{f}_{s^{(t)}}} [\mathbf{v}^{(t)}]$.

(a) If types are iid, then

$$K'(v^{(t)}, s^{(t)}) = D_{\mathbf{1}} P(\mathbf{v}^{(t)}, s^{(t)}) = D_{\hat{\theta}_{s^{(t)}}} P(\mathbf{v}^{(t)}, s^{(t)})$$

and thus the process $\left(D_{\hat{\theta}_{s^{(t)}}} P(\mathbf{v}^{(t)}, s^{(t)}) \right)_{t=0}^{\infty}$ is a martingale.

(b) If types are *not* iid, then

$$D_{\hat{\theta}_{s^{(t)}}} P(\mathbf{v}^{(t)}, s^{(t)}) \geq D_{\mathbf{1}} P(\mathbf{v}^{(t)}, s^{(t)})$$

and the process $\left(D_{\hat{\theta}_{s^{(t)}}} P(\mathbf{v}^{(t)}, s^{(t)}) \right)_{t=0}^{\infty}$ is *not* a martingale.

The proof of Lemma 7.2 is in Appendix C.4. This gives one sense in which the TW martingale $K'(v^{(t)}, s^{(t)})$ and the differential martingale are *not* equivalent. In the iid setting, optimal intertemporal cost-smoothing can be achieved by perturbations different from those underlying the differential martingale. But these perturbations are not guaranteed to be optimal (ie, induce a process satisfying the martingale property) outside the iid case. Intuitively, the perturbation underlying the process $D_{\hat{\theta}_s} P(\mathbf{v}, s)$ does not take into account the principal's need to manage Markov information rents.

8. Concluding Remarks

We close with just a few remarks regarding directions for future research. First, we argued in Section 6.2 that ergodicity and asymptotically-vanishing impulse response functions appear to be the critical properties of the type process that determine the long-run properties of the optimal contract — namely, whether or not Absolute Immiseration occurs. While Theorem 3 helped us go some way toward making this case, much remains to be done. A definitive answer to our conjecture seems to be important for a more complete understanding of optimal insurance in private information economies. But, given the generality of the present analysis, it seems that new techniques will be needed to tackle these questions.

Second, it seems important to better understand how the behavior of the wedges differ between the kind of pure insurance setting studied here and the redistribution problem considered in optimal dynamic taxation. Our findings in Section 7 make some progress in this direction, but perhaps more analytical progress could be made using the FOA in a model with continuous types and continuous time.

Finally, many of our techniques — namely, the arguments behind the domain characterization in Theorem 1 and the convergence proof in Theorem 3 — apply directly to other models

of dynamic contracting with Markovian types. For example, they can be used to generalize the corresponding results in the firm financing model of Fu and Krishna (2017) and the repeated procurement model of Krasikov and Lamba (2018), thus paving the way for a deeper understanding of dynamic incentives in these settings.

Appendices

The appendices are organized as follows. Appendix A contains the proof of Theorem 1. Appendix B contains proofs from Section 5, including the proofs for Proposition 5.1, Theorem 3, and Theorem 4. Appendix C contains proofs and derivations for the material in Section 7.

A Supplementary Appendix contains additional proofs and results. Material pertaining to the sequential problem, [SP], and its relation to the recursive problem, [RP], is relegated to Supplementary Appendix S.1.1. Supplementary Appendix S.2 contains details pertaining to the first-best problem from Section 4.1. Supplementary Appendix S.3 contains all other proofs from Section 4, including that of Theorem 2. Finally, Supplementary Appendices S.4 and S.5 contain auxiliary results used in the proofs of Theorems 1 and 3, respectively.

A. Proofs Concerning the Recursive Domain

Appendix A, which is divided into several parts, is concerned with proving Theorem 1. Appendix A.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 Appendix A.2, which is focused mainly on two auxiliary operators, \mathcal{B} and \mathcal{B}^\dagger . First, Appendix A.2.1 establishes several results relating fixed points of the three operators to each other and to the set V_d . Second, Appendices A.2.2 and A.2.3 focus on the special case of CARA utility. Appendix A.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 Markov process satisfies MLRP or PPR. Appendix A.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, Appendix A.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 Appendix A.2.3. Finally, Appendix A.3 collects these various pieces into a proof of Theorem 1.

A.1. Preliminary Observations

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

Lemma A.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

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

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 \mathcal{C} 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 **[A.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 $K \subset \mathcal{U}^d$:

$$[\mathbf{A.2}] \quad \mathcal{B}_\circ(K) := \{\mathbf{v} \in \mathcal{U}^d : \exists (\mathbf{u}, \mathbf{w}) \text{ that implements } \mathbf{v} \text{ with } (u_i, \mathbf{w}_i) \in \mathcal{U} \times K \text{ for all } i \in S\}$$

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

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

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 [A.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 A.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 $\boldsymbol{\gamma} \in \Gamma(\mathbf{v}, K)$ and $\boldsymbol{\gamma}' \in \Gamma(\mathbf{v}', K)$. Let $a \in (0, 1)$ and define $\mathbf{v}'' := a\mathbf{v} + (1-a)\mathbf{v}' \in \mathcal{U}^d$ and $\boldsymbol{\gamma}'' := a\boldsymbol{\gamma} + (1-a)\boldsymbol{\gamma}' \in (\mathcal{U} \times K)^d$. Each of the [PK_j] constraints is linear, so clearly $\boldsymbol{\gamma}''$ satisfies all of them at \mathbf{v}'' . Consider constraint [IC_{ij}], with $i > j$. By definition of $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$, we have

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

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

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

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

and, by part (c) of Lemma A.1, $a\psi(u_j, i, j) \geq \psi(au_j, i, j)$. Thus, $\boldsymbol{\gamma}'$ satisfies [IC_{ij}] at \mathbf{v}' . Hence, $\boldsymbol{\gamma}' \in \Gamma(\mathbf{v}', K)$, 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 A.1 instead of part (c) of that lemma. \square

Proposition A.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

$$B := \{c \in \mathbb{R} : c + \omega_i \in \mathcal{C} \ \forall i \in S\}$$

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

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

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 **A.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 **A.2**, namely the facts that D is convex and $\Gamma(\cdot) = \Gamma(\cdot, D)$ as defined in **[A.3]**. Continuity is easy to see from the definitions. \square

A.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 $K \subset \mathbb{R}_-^d$:

$$\mathbf{[A.4]} \quad \mathcal{B}(K) := \{ \mathbf{v} \in \mathbb{R}_-^d : \exists (\mathbf{u}, \mathbf{w}) \text{ that implements } \mathbf{v} \text{ with } (u_i, \mathbf{w}_i) \in \mathbb{R}_- \times K \text{ for all } i \in S \}$$

It is easy to see that $\mathcal{B}_o(K) \subset \mathcal{B}(K)$ for all K , so that the largest fixed point of \mathcal{B}_o 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 $K \subset \mathcal{U}^d$:

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

Notice that for any set K , we have

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

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 [A.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 A.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 A.2 and the definition of the constraint correspondence $\Gamma(\cdot)$ (see [3.1]). \square

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

- (a) Suppose the Markov 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 A.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 A.6. Fix $w^j < 0$. Then, for all $\varepsilon > 0$, there exists $\mathbf{w}_j \in V_d$ such that (i) $\mathbf{E}^{f_j}[\mathbf{w}_j] = w^j$, and (ii) $\mathbf{E}^{f_i}[\mathbf{w}_j] - \mathbf{E}^{f_j}[\mathbf{w}_j] < \varepsilon$ for all $i > j$. If the Markov process satisfies FOSD, there exists $\mathbf{w}_j \in V_d$ that satisfies (i), (ii), and (iii) $\mathbf{E}^{f_i}[\mathbf{w}_j] \geq \mathbf{E}^{f_j}[\mathbf{w}_j]$.

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 $\mathbf{E}^{f_j}[\mathbf{w}_j] = w^j$, establishing (i).

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

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

Lemma A.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 A.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 A.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 A.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 A.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 A.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 Φ_\circ define complete lattices when equipped with the set inclusion partial order. It follows from monotonicity and Lemma A.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 Φ_\circ . It follows immediately that $K_*, K_*^\dagger \in \Phi$ and $D \in \Phi_\circ$. \square

Proposition A.9. Suppose that the Markov 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 A.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

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

which implies $v_i \geq v_j$, as (i) $\psi(u_j, i, j) - u_j \geq 0$, and (ii) $\mathbf{E}^i[\mathbf{w}_j] \geq \mathbf{E}^j[\mathbf{w}_j]$ because \mathbf{f}_i first order stochastically dominates \mathbf{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 A.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), $[\mathbf{PK}_j]$ implies that

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

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 A.5, if $(u_j, \mathbf{w}_j) \in \mathbb{R}_- \times \text{cl } V_d$ satisfy $[\mathbf{IC}_{ij}]$ and $[\mathbf{PK}_j]$, 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 A.10. If the Markov process satisfies FOSD, then V_d is a fixed point of \mathcal{B}_\circ .

Proof. Lemma A.7 states that $\mathcal{B}_\circ(V_d) \supseteq V_d$. Part (c) of Proposition A.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

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

Recall (from, eg, part (c) of Lemma A.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 $[\mathbf{IC}_{ij}]$ are linear in the u_j .

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

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

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

Proposition A.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

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

and

$$[\mathbf{IC}_{21}] \quad \frac{\theta_2}{\theta_1} \cdot u_1 + \alpha \mathbf{E}^{f_2}[\mathbf{w}_1] \leq v_2$$

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 A.7.) For general \mathbf{v} , consider the following problem: What is the maximum value of $u_1 + \alpha \mathbf{E}^{f_1}[\mathbf{w}_1]$ 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

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

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.⁹⁵

Theorem 5 (LP duality and Complementary Slackness). *Let $y^* = (y_1^*, \dots, y_n^*)$ be a feasible solution of the linear program*

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

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:

(95) 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 5 is on p. 204 of the Glossary of Matoušek and Gärtner (2007).

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

Consider the operator $\mathcal{B}_{j+1,j}(K_{\mathbf{t}})$ where $\mathbf{t} \geq 1$ is finite in every component. We seek to characterize the set of all \mathbf{v} that can be implemented by $(u_j, \mathbf{w}_j)_{j \in \mathcal{S}} \in \mathbb{R}_-^d \times K_{\mathbf{t}}$ while satisfying **[PK_j]** and **[IC_{ij}]** where $i = j + 1$. As in the proof of Proposition A.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 **[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 $\mathbf{q}^{j+1} \geq \mathbf{0}$ 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 A.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 A.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 A.8. \square

Corollary A.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 A.12, $t_{j+1}^{(n+1)} < 1$. But this contradicts part (a) of Lemma A.12. That $u_j = 0$ at any optimum then follows from Complementary Slackness (part (b) of Theorem 5). \square

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

Lemma A.14. Let $\mathbf{t}^{(n)}$ 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 5), 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 A.11.) Hence, by induction, all prior constraints, other than $[\mathbf{DC}_0]$, must hold, which proves the claim. \square

Lemma A.15. Suppose the Markov 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 A.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 5) and so by Lemma A.14, $[\mathbf{DC}_k]$ binds for all $k = 1, \dots, d$. \square

Lemma A.16. Suppose the Markov 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 5), 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 Markov 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 Markov 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 $[\mathbf{IC}_{ij}]$ where $i = j + 1$, as demonstrated in Proposition A.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 A.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 [A.6] above. If the Markov process satisfies MLRP, $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$.

Proof. By Lemma A.15, the dual constraints [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 [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 [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 $t_2 \cdots t_d > 1$. □

Lemma A.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 [A.6] above. If the Markov process satisfies PPR, $\mathcal{B}^\dagger(K_{\mathbf{t}}) \neq K_{\mathbf{t}}$.

Proof. Consider the dual problem [D_{j+1,j}] for some $j = 1, \dots, d-1$. By Lemma A.16, it follows that the dual constraints [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 [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 [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. □

A.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 A.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 A.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 A.12 and Corollary A.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 $C^* := \{\mathbf{y} \in \mathbb{R}^n : \mathbf{y} \cdot \mathbf{x} \geq 0 \forall \mathbf{x} \in C\}$. We will use the following version of Conic Programming Strong Duality in the sequel.⁹⁶

Theorem 6 (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*

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

and the dual conic program

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

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 $\mathbf{v} \in \mathbb{R}_-^d$ and $j = 1, \dots, d-1$, define $\mathbf{v}_{>j} := (v_{j+1}, \dots, v_d) \in \mathbb{R}_-^{d-j}$.

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

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 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 [\[Pcone\]](#), with $y = (u_1, \mathbf{w}_1)$, $b = \mathbf{v}_{>j}$, $L = \mathbb{R}_+^d$, $J = \mathbb{R}_- \times K$, 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 follows that the dual conic program may be written as

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

where $[\mathbf{f}_{j+1}, \mathbf{f}_{j+2}, \dots, \mathbf{f}_d]$ 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 $MP_j : \mathbb{R}_+^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_+^{d+1}$ denote the argmax correspondence from [\[P_j^{full}\]](#), as a function of $(\mathbf{v}_{>j}, K)$. (Recall that \mathcal{K} denotes the space of closed convex cones in \mathbb{R}_+^d .) Let $VP_j : \mathbb{R}_+^{d-j} \times \mathcal{K} \rightarrow \mathbb{R}_-$ denote the value function for [\[P_j^{full}\]](#). Similarly, let $MD_j : \mathbb{R}_+^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_+^{d-j}$ and $VD_j : \mathbb{R}_+^{d-j} \times \mathcal{K} \rightarrow \mathbb{R}_-$ denote the argmin correspondence and value function from [\[D_j^{full}\]](#).

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

- For any $\mathbf{v}_{>j} \in \mathbb{R}_+^{d-j}$, the Slater condition and strong conic programming duality (conclusions (a) and (b) of [Theorem 6](#)) hold.
- $VP_j(\cdot, K)$ and $VD_j(\cdot, K)$ are well-defined (ie, finite-valued) continuous, concave, and HD1.
- $MP_j(\cdot, K)$ and $MD_j(\cdot, K)$ are nonempty-valued and upper hemi-continuous. $MD_j(\cdot, K)$ is HD0.
- $VP_j(\mathbf{v}_{>j}, K) = VD_j(\mathbf{v}_{>j}, K) < 0$ if and only if $\mathbf{v}_{>j} \neq \mathbf{0}$.

Proof. First, observe that [\[P_j^{full}\]](#) 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 [\[P_j^{full}\]](#): $\mathbf{w}_j = \mathbf{0}$

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 $M : \mathbb{R}_-^{d-j} \times \mathcal{K} \rightrightarrows \mathbb{R}_-^{d+1}$ 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 6. 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 A.12.

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

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

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

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

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 A.20 imply that $K^{(n)}$ and K_* indeed satisfy all the properties laid out in Lemma A.19.

Define the following sets:

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

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

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

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

Finally, the desired analogue to Corollary A.13.

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

Proof. By part (d) of Lemma A.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 $q_{k^*-j} > 1$. From the definition of $[\text{DC}_0^{\text{full}}]$, 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 A.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 A.19, point (ii) implies that $K_* \not\subseteq \text{cl } V_d$, contradicting part (a) of Proposition A.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

A.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:

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

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 A.1, the constraint set is closed and convex.

Proposition A.24. Let $K \in \left\{K^{(n)}\right\}_{n \in \mathbb{N}} \cup \{K^*\}$. 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 A.4 or, alternatively, Lemmas A.20 and A.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

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

As in the proof of Lemma A.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,

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

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 A.23.) A necessary first-order condition for $(u_j, \mathbf{w}_j), \mathbf{q}^{*j}$ to constitute a saddle point as in [A.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 A.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 A.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 A.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 A.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 A.25. The following hold:

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

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

For part (b), consider problem $[\mathbf{P}_j^{\text{DARA}}]$ given K_* . By Proposition A.24, K_* is characterized by Lemmas A.21 and A.22. Recall the sets A and B defined before Lemma A.22. Clearly, $\mathbf{v} \in A$ implies that $\mathbf{v} \notin \mathcal{B}_o(K_*)$. Consider $\mathbf{v} \in B$. It follows from the proof of Proposition A.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, $\mathbf{v} \notin \mathcal{B}_o(K_*)$.

It therefore follows from Lemma A.22 that $\mathcal{B}_o(K_*) \subseteq \text{int}K_*$. We claim that $\mathcal{B}_o(\text{int}K_*) \supseteq \text{int}K_*$. Let $\mathbf{v} \in K_*$, so that by Lemma A.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 A.20 and Proposition A.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, $\mathbf{v} \in \mathcal{B}_o(\text{int}K_*)$. Thus, we have proved the claim.

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

A.3. Proof of Theorem 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 A.3. Part (c) of Proposition A.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 A.25. \square

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

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

Proof of Part (d). Suppose that the Markov process satisfies FOSD. Consider the IC constraint $[\mathbf{IC}_{ij}^*]$. 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 $[\mathbf{IC}_{ij}^*]$ (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 displays 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 $i > j$, ie, $D^* \subset V$. □

Proof of Part (e). By part (b) of Proposition A.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 A.12, the iterates of the operator \mathcal{B}^\dagger starting with the set \mathbb{R}^d are characterized by the sequence $(K_{\mathbf{t}^{(n)}})$ and K_*^\dagger , the largest fixed point of \mathcal{B}^\dagger , is the limit of this sequence. That is, $K_*^\dagger = K_{\mathbf{t}^*}$ where $\mathbf{t}^* := \lim_{n \rightarrow \infty} \mathbf{t}^{(n)}$.

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

Proof of Part (f). Given part (e) of Theorem 1, proved above, it suffices to show that $D^* = V_d$ under the additional assumption of CARA utility. The proof for the $d = 2$ case is straightforward but fairly tedious, and is thus relegated to Appendix S.4.1, which culminates in the statement of Proposition S.4.7. The general $d \geq 2$ case is completely analogous, but is notationally very cumbersome and thus omitted. □

B. Proofs from Section 5

Appendix B is divided into several parts, which are organized as follows. First, in Appendix B.1 we lay out the Lagrangian and optimality conditions for the principal's problem. in [FE]. These are referenced heavily in the subsequent parts and Appendix C. Appendix B.2 proves Proposition 5.1.

Appendices B.3 and B.4 prove Theorem 3. First, Appendix B.3 establishes a set of intermediate results — culminating in Lemmas B.5 and B.15 — that are needed to establish the martingale convergence result in part (a) of Theorem 3. Appendix B.4 then contains the convergence proof itself.

Finally, Appendix B.4 contains the proof of Theorem 4.

B.1. Optimality Conditions

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

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

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

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

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.)

Theorem 2 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, eg, Exercise 7 on p. 236 and Theorem 2 on p. 221 of Luenberger (1969)). The Lagrangian for this problem is

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

where $\lambda_i \in \mathbb{R}$ is the multiplier on the promise keeping constraint $[\mathbf{PK}_i]$ and $\mu_{ij} \geq 0$ is the multiplier on the incentive constraint $[\mathbf{IC}_{ij}]$. 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

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

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

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

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

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

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

B.2. Proof of Proposition 5.1

Suppose the environment is regular. Because P is continuously differentiable by Theorem 2, 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 $D_{\mathbf{1}}P(\mathbf{v}^{(t)}, s^{(t)})$ then follow from non-negativity and finiteness of the directional derivative and finiteness of S .

As for the martingale property, summing the [Env_{*i*}] over $i \in S$ delivers

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

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

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

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

$$\sum_{i \in S} f_{si} D_{\mathbf{1}}P(\mathbf{w}_i, i) = \sum_{i=1}^d \lambda_i$$

which, combined with [B.1], gives the martingale property

$$D_{\mathbf{1}}P(\mathbf{v}, s) = \sum_{i=1}^d f_{si} D_{\mathbf{1}}P(\mathbf{w}_i, i)$$

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

Lemma B.1. Suppose that the environment is regular. Then the directional derivative $D_{\mathbf{1}}P(\cdot, s)$ is non-negative for each $s \in S$. If the environment is [TV_C]-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 would follow directly from the definition

$$D_{\mathbf{1}}P(\mathbf{v}, s) := \lim_{\varepsilon \downarrow 0} \frac{P(\mathbf{v} + \varepsilon \mathbf{1}, s) - P(\mathbf{v}, s)}{\varepsilon}$$

(Note that Lemma S.3.11 implies that the directional derivative exists.) Hence, it suffices to show that $P(\cdot, s)$ is non-decreasing in direction $\mathbf{1}$. Lemma S.2.2 shows that the first-best value function Q^* is non-decreasing (indeed, strictly increasing) in this direction. We will show that P inherits this property from Q^* . The proof is order-theoretic. (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 not able to show that such convergence takes place in general.)

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 7. 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 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 8. The lattice Φ is complete.

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

$$\overline{f}(\mathbf{v}, s) - \delta \geq \overline{f}(\mathbf{v}', s)$$

By definition of the supremum, there exists some $f \in F$ such that $f(\mathbf{v}, s) > \overline{f}(\mathbf{v}, s) - \delta$. Combined with the above display and the definition of Φ , this implies that

$$f(\mathbf{v}, s) > \overline{f}(\mathbf{v}', s) \geq f(\mathbf{v}', s)$$

which contradicts the fact that f is non-decreasing in the direction $\mathbf{1}$ by virtue of $f \in F \subset \Phi$. \square

Claim 9. $T : \Phi \rightarrow \Phi$ 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,

$$\begin{aligned} \delta + TQ(\mathbf{v}, s) &\geq \sum_{i \in S} f_{si} [C(u_i, i) + \alpha Q(\mathbf{w}_i, i)] \\ &\geq \sum_{i \in S} f_{si} [C(u_i, i) + \alpha Q(\mathbf{w}_i - \frac{\varepsilon}{\alpha} \mathbf{1}, i)] \\ &\geq TQ(\mathbf{v} - \varepsilon \mathbf{1}, s) \end{aligned}$$

where the second 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 Lemma S.3.6 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 $D_{\mathbf{1}}P(\mathbf{v}, s) = 0$. 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 Lemma S.3.12, and is non-increasing by Lemma B.1. Hence, the hypothesis that $D_{\mathbf{1}}P(\mathbf{v}, s) = 0$ implies that $f(y) \equiv f(0)$ for all $y \in Y$. But this contradicts strict convexity of P and is therefore impossible. It follows that $D_{\mathbf{1}}P$ is strictly positive on $D \times S$. \square

B.3. Intermediate steps toward the proof of Theorem 3

B.3.1. The Efficiency Problem

Recall the *efficiency problem* [Eff_{*i*}] from Section 5.1, 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 B.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. The proof, which is standard but somewhat tedious, follows from the same arguments as the proof of Theorem 2 in Supplementary Appendix S.3. Thus, we simply point to the appropriate lemmas there, and let the reader fill in the easy details. For existence, see Lemmas S.3.9 and S.3.10. For properties (including upper hemi-continuity) of the policy correspondence, see again Lemma S.3.10. For strict convexity of the value function, see Lemmas S.3.7 and S.3.12. Uniqueness of the minimizer follows from strict convexity, and continuity of the policy function follows from uniqueness and upper hemi-continuity of the policy correspondence. Continuous differentiability follows as in Lemma S.3.11. Together, this establishes parts (a) and (b) and most of part (c). To see that $K(\cdot, i)$ is strictly increasing, 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 5.1 (or Lemma B.1), $P(\mathbf{w}', i) < K(w, i)$. Thus, by revealed preference, $K(w', i) < K(w, i)$, which proves strict monotonicity.

The Inada conditions in part (c) follow from properties of the value function for the analogue of [Eff_{*i*}] in the full-information, which is defined as [Eff_{*i*}^{FB}] in Supplementary Appendix S.2. That

value function is called $K^*(w, i)$, and clearly satisfies $K^* \leq K$ on $\mathcal{U} \times S$. Lemma S.2.3 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 S.2.3 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

$$E_i := \left\{ \mathbf{v} \in D : \mathbf{v} = \mathbf{w}^\dagger(w, i) \text{ for some } w \in \mathcal{U} \right\}$$

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 [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 B.2, the unique solution to [Eff_{*i*}] is characterized by the first-order conditions

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

and the envelope condition

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

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

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

is the image of E_i under the derivative mapping $DP(\cdot, i)$. Moreover, by summing the first-order conditions [FOC_{*j*}-Eff_{*i*}] over $j \in S$ and combining with the envelope condition [Env_{*j*}-Eff_{*i*}], we get

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

Remark 10. By comparing the optimality conditions for [RP] (enumerated in Appendix B.1) with the optimality conditions for [Eff_{*i*}] stated above, it is easy to see 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 *efficient* in the iid case. This observation, together with equation [B.3] above, formalize the discussion regarding the TW martingale at the end of Section 5.1.

B.3.2. Facts Concerning the Differential Martingale

Lemma B.3. Suppose the environment is [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 $D P(\cdot, s)$ is injective. To see this, notice first that because P is strictly convex, by Theorem 7.21 of Van Tiel (1984), $D P(\cdot, s)$ is *strictly monotone* in the sense that for all $\mathbf{v}, \mathbf{v}' \in D$ such that $\mathbf{v} \neq \mathbf{v}'$, $\langle D P(\mathbf{v}, s) - D P(\mathbf{v}', s), \mathbf{v} - \mathbf{v}' \rangle > 0$. But now suppose $D P(\cdot, s)$ is not injective, so that there are $\mathbf{v}, \mathbf{v}' \in D$ distinct such that $D P(\mathbf{v}, s) = D P(\mathbf{v}', s)$. But this would imply that $0 = \langle \mathbf{0}, \mathbf{v} - \mathbf{v}' \rangle = \langle D P(\mathbf{v}, s) - D P(\mathbf{v}', s), \mathbf{v} - \mathbf{v}' \rangle > 0$, which is a contradiction. As D is open by Theorem 1 and the derivative $D P(\cdot, s)$ is continuous on D by Theorem 2, it follows from Brouwer's Invariance of Domain Theorem — see, for instance, Theorem 2B.3 on p. 172 Hatcher (2001) — that D is homeomorphic to $D P(D, s) = Y_s$. \square

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

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

Lemma B.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 [B.3] and part (c) of Lemma B.2. \square

Lemma B.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

$$\begin{aligned} f_{dd} \left(\lambda_d(\mathbf{v}, s) + \sum_{k=1}^{d-1} \mu_{dk}(\mathbf{v}, s) \right) &= f_{sd} P_d(\mathbf{w}_d, d) \\ &= f_{sd} \lambda_d(\tilde{\mathbf{w}}_d, d) \end{aligned}$$

where the first line is the FOC for w_{dd} at state (\mathbf{v}, s) and the second line follows from the d th envelope condition at state (\mathbf{w}_d, d) . It follows that

[B.4]

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

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

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

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

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

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

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

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

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

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

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

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 B.4 and [E_{*i*}], this term vanishes. Thus, plugging the first line of [B.9] into [B.7] and rearranging delivers [MS_{*i*}]. \square

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

- (1) $D_1 P(\tilde{\mathbf{w}}_i, i) = D_1 P(\mathbf{w}_d, d)$ 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 B.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 (ie, $\mu_{dk}(\mathbf{v}, s) \geq 0$ 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

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

It follows from [B.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

B.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*:

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

subject to

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

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

for all $j \in S$ with $j > i$. That is, suppose the agent reports to be 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 $[\mathbf{PK}_i]$ for type i and incentive compatibility $[\mathbf{IC}_{ji}^*]$ 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 $[\mathbf{PK}_i]$ ($i = d$).

For each $i \in S$, define

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

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

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

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

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

$$P(\mathbf{v}, s) = \sum_{i=1}^d \inf_{(u_i, \mathbf{w}_i) \in \Gamma_i(\mathbf{v})} [C(u_i, i) + \alpha P(\mathbf{w}_i, i)]$$

from which the lemma immediately follows. □

Lemma B.9. Suppose the environment is regular. Then:

- (a) There exists a recursively 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) .⁹⁷

(97) 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 a recursively optimal contract is established in Lemma S.3.10 in Supplementary Appendix S.3. The existence of a recursively optimal contract with the desired properties then follows immediately from Lemma B.8.

Now consider part (b). By part (e) of Theorem 2 (proved independently in Supplementary Appendix S.3), there exists a unique recursively optimal contract ξ^* . The result then follows from part (a) of the present lemma. \square

Lemma B.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*_j] ($j > i$) hold as strict inequalities.

Proof. Part (a) follows from the definition of [FE- Q^i], convexity and continuous differentiability of $P(\cdot, i)$ (Theorem 2). 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 B.10, the solutions of problem [FE- Q^i] are characterized by saddle points of the Lagrangian (see, eg, Exercise 7 on p. 236 and Theorem 2 on p. 221 of Luenberger (1969))

[L_i]

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

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*_j].

Lemma B.11. Suppose the environment is regular. At the optimum:

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

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

for all $i \in S$.

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

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

for all $i \in S$.

- (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 B.8 and comparison of optimality conditions from the “ex ante” problem in Appendix B.1 and those derived from the “interim” Lagrangians $[\mathbf{L}_i]$, $i \in S$. Begin with part (a). The envelope conditions from $[\mathbf{L}_i]$ read

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

for all $i, j \in S$. It follows from Lemma B.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 [B.15], that

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

Substituting the ex ante envelope condition [Env $_i$] ($i = j$) into [B.16] delivers

$$\lambda_j(\mathbf{v}, s) = f_{sj} \left(\eta_j(\mathbf{v}) - \sum_{k=j+1}^d \sigma_{kj}(\mathbf{v}) \right) - \sum_{k=1}^{j-1} f_{sk} \sigma_{jk}(\mathbf{v})$$

Dividing both sides through by f_{sj} and replacing the dummy index j with i delivers [B.13].

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

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

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

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

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

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

Simplifying the above display yields [B.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 [B.14] with $i = d$ reads

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

where the second line 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 [B.14] with $i = \ell$ reads

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

where the second line follows because (i) holds at \mathbf{v} and the third line 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 B.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

$$\text{[B.19]} \quad \frac{\lambda_i(\mathbf{v}, r)}{f_{ri}} = \eta_i(\mathbf{v})$$

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 B.11 implies that $\sigma_{ij}(\mathbf{v}) = 0$ for all $i, j \in S$ with $i > j$. Plugging this into [B.13] yields [B.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 for all $s', i \in S$

$$\frac{\lambda_i(\mathbf{v}, s')}{f_{s'i}} = \hat{\eta}(\mathbf{v})$$

for all $s', i \in S$. Using the envelope conditions [Env v_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

B.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}_ξ solves the first-best problem [FB] given initial condition (\mathbf{v}, s) .

Lemma B.13. The first-best contract (see Section 4.1 and Supplementary Appendix S.2) is not feasible at any $\mathbf{v} \in D$.

Proof. By Lemma S.2.2 in Supplementary Appendix S.2, 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 B.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 recursively optimal contract self-generates at \mathbf{v} ;
- (ii) The recursively 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 v_i], [FOC u_i], [FOC w_{ij}], respectively) reduce to

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

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

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

It is easy to see that [B.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 B.12 imply that $\mathbf{v} \in E_{s'}$ for all $s' \in S$. Now, plugging [B.19] from Lemma B.12 into [B.20] and [B.22], and invoking part (iii) of Lemma B.12, delivers

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

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

where, as in the proof of Lemma B.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 B.9, that \mathbf{w}_i is independent of s , so [B.23] and [B.24] do not depend on the given s at all. Thus, take an arbitrary $s' \in S$. Summing over $i \in S$ in [B.23] delivers

$$D_1 P(\mathbf{v}, s') = \sum_{i=1}^d f_{s'i} \hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$$

Similarly, set $i = s'$ in [B.24] and sum over $j \in S$ to get

$$D_1 P(\mathbf{w}_{s'}, s') = \sum_{i=j}^d f_{s'j} \hat{\eta}(\mathbf{v}) = \hat{\eta}(\mathbf{v})$$

Now, we have established that $\mathbf{v}, \mathbf{w}_{s'} \in E_{s'}$. By Lemma B.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 [B.20], [B.21], and [B.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 B.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}_{ξ^*} is constant and equal to \hat{u} when initialized at (\mathbf{v}, s) . It follows from Lemma S.2.2 (see the Supplementary Appendix S.2) that ξ^* implements the first-best at (\mathbf{v}, s) . \square

Lemma B.15. Suppose the environment is [TVFC]-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 B.14 and B.7 imply that the optimal contract implements the first best at (\mathbf{w}_d, d) . But this is impossible by B.13, a contradiction. \square

B.4. Proof of Theorem 3

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.

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 S.5.3 in Appendix S.5.

Proof of Part (a). Suppose that the environment is [TVFC]-regular, as hypothesized by the theorem. Proposition 5.1 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 B.4 and B.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 Lemma S.3.11, 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 B.15, B.13, and B.14. Thus, $D_1 P(\mathbf{v}^t, s^t) \rightarrow 0$, completing the proof of part (a). \square

We prove parts (b) – (d) by first characterizing the limit properties of the Lagrange multipliers (the dual variables), and then translating these into statements concerning the allocation (the primal variables). This is done in a series of lemmas, presented next. Let $\mathbf{v}(\mathbf{v}, s) \in \mathbb{R}^{d(d+1)/2}$ 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 $(\mathbf{v}^{(t)})_{t=0}^\infty$, where $\mathbf{v}^{(t)} := \mathbf{v}(\mathbf{v}^{(t)}, s^{(t)})$ 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 B.16. Suppose the environment is [TVC]-regular. Under the optimal contract, $\mathbf{v}^{(\tau^{(t)})} \rightarrow \mathbf{0}$ and $\boldsymbol{\lambda}^{(\tau^{(t)+1})} \rightarrow \mathbf{0}$ almost surely.

Proof of Lemma B.16. Fix some path $h \in \mathcal{F}$ along which the differential martingale converges to zero. By part (a) of Theorem 3, the set of such paths has full measure, and is thus sufficient for establishing almost sure convergence. By [E_{*i*}] and Lemmas B.4 and B.5, convergence of the differential martingale implies that the derivative process converges along the subsequence $\{\tau_t\}_{t=0}^\infty$, and part (a) of Theorem 3 (proved above) requires that $DP(\mathbf{v}^{\tau_t}, s^{\tau_t}) \rightarrow \mathbf{0}$. By the envelope conditions [Env_{*i*}], this translates to

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

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 [FOCw_{*ij*}] with $i = 1$ at state (\mathbf{v}^{τ_t}, d) is

$$f_{d1} P_j(\tilde{\mathbf{w}}_1^{\tau_t}, 1) = f_{1j} (\lambda_1(\mathbf{v}^{\tau_t}, d) + 0) - \sum_{k=2}^d f_{kj} \mu_{k1}(\mathbf{v}^{\tau_t}, d)$$

Since $\boldsymbol{\mu}(\cdot) \geq \mathbf{0}$ on $D \times S$, it follows from [B.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 Markov process is fully connected (Assumption Markov), it

follows that

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

$$\text{[B.27]} \quad \boldsymbol{\mu}_{*,1}(\mathbf{v}^{\tau_t}, d) \rightarrow \mathbf{0}$$

Inductive step: Let $m \in S$. Suppose we have shown that $DP(\tilde{\mathbf{w}}_\ell^{\tau_t}, \ell) \rightarrow \mathbf{0}$ and $\boldsymbol{\mu}_{*,\ell}(\mathbf{v}^{\tau_t}, d) \rightarrow \mathbf{0}$ for all $\ell < m$. The first order condition **[FOCw_{ij}]** with $i = m$ at state (\mathbf{v}^{τ_t}, d) is

$$f_{dm} P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) = f_{mj} \left(\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d) \right) - \sum_{k=m+1}^d f_{kj} \mu_{km}(\mathbf{v}^{\tau_t}, d)$$

By **[B.25]** and the supposition, it follows that

$$\lambda_m(\mathbf{v}^{\tau_t}, d) + \sum_{k=1}^{m-1} \mu_{mk}(\mathbf{v}^{\tau_t}, d) \rightarrow \mathbf{0}$$

Since $\boldsymbol{\mu}(\cdot) \geq \mathbf{0}$ on $D \times S$, it follows from the previous two displays that $f_{dm} P_j(\tilde{\mathbf{w}}_m^{\tau_t}, m) \rightarrow 0$ 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 Markov process is fully connected (Assumption **Markov**), it follows that

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

$$\text{[B.29]} \quad \boldsymbol{\mu}_{*,m}(\mathbf{v}^{\tau_t}, d) \rightarrow \mathbf{0}$$

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

Lemma B.17. Suppose the environment is **[TV C]**-regular. Under the optimal contract and for all $k \in \mathbb{N}$, $\mathbf{v}^{(\tau^{(t)}+k)} \rightarrow \mathbf{0}$ and $\boldsymbol{\lambda}^{(\tau^{(t)}+k+1)} \rightarrow \mathbf{0}$ almost surely.

Proof of Lemma B.17. Suppose we have shown the desired convergence for all $k = 0, \dots, m-1$. Replicating the proof of Lemma **B.16** 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 **B.16** serving as the base step. \square

Lemma B.18. Suppose the environment is **[TV C]**-regular. Under the optimal contract, $\mathbf{v}^{(t)} \rightarrow \mathbf{0}$ in probability.

Proof of Lemma B.18. Define the stochastic processes $(\delta^{(t)})_{t=0}^\infty$ and $(L^{(t)})_{t=0}^\infty$ by $\delta^{(t)} := \|\mathbf{v}^{(t)}\|$, 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 $\mathbf{v}^{(t)}(h)$ 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:

$$\begin{aligned} A_{\varepsilon,t} &:= \{h \in \mathcal{H} : \delta^{(t)}(h) > \varepsilon\} \\ B_{k,t} &:= \{h \in \mathcal{H} : L^{(t)}(h) > k\} \\ C_{\varepsilon,k,t} &:= \cup_{T \geq t} [A_{\varepsilon,T} \cap B_{k,T}^c] \end{aligned}$$

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

$$\begin{aligned} \mathbf{P}(A_{\varepsilon,t}) &= \mathbf{P}(A_{\varepsilon,t} \cap B_{k,t}^c) + \mathbf{P}(A_{\varepsilon,t} \cap B_{k,t}) \\ &\leq \mathbf{P}(C_{\varepsilon,k,t}) + \mathbf{P}(B_{k,t}) \end{aligned}$$

where the second line 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 B.16 and B.17 imply that $\delta^{(\tau^{(t)}+m)} \rightarrow 0$ almost surely for all $m = 0, \dots, k$. Thus, it follows that

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

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 [B.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, ie, 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 11. 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 $k \in \mathbb{N}$,

$$\lim_{t \rightarrow \infty} \mathbf{P} \left(s^{(t-m)} \neq d \ \forall m = 0, \dots, k \mid s^{(t)} = i \right) = 1 - H_i(k)$$

Proof of Claim 11. It is clear that the time-reversed process is fully connected, and thus each state is recurrent, ie, $\mathbf{Q} \left(T_d^R < \infty \mid r^{(0)} = i \right) = 1$ for all $i \in S$. It follows from the Bounded Convergence Theorem that $\lim_{k \rightarrow \infty} H_i(k) = 1$ 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 $\mathbf{P} \left(s^{(T)} = i \right) > 0$ for all $T \geq t$. (This is possible because $\pi_i > 0$.) Then,

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

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

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

where the first line follows from the Markov property (for the forward chain) and the second line follows from the definition of the time-reversed transition probabilities. Combining the two displays above and noting that $\lim_{t \rightarrow \infty} \mathbf{P} \left(s^{(t-k)} = j_k \right) = \pi_{j_k}$ and $\lim_{t \rightarrow \infty} \mathbf{P} \left(s^{(t)} = i \right) = \pi_i$ 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 11. Notice that we may write

$$\mathbf{P} (B_{k,t}) = \sum_{i=1}^d \mathbf{P} \left(s^{(t)} = i \right) \cdot \mathbf{P} \left(s^{(t-m)} \neq d \ \forall m = 0, \dots, k \mid s^{(t)} = i \right)$$

By part (b) of Claim 11 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 B.19. Suppose the environment is [TVC]-regular. Suppose that the Markov process satisfies UPR. Under the optimal contract, $\mathbf{v}^{(t)} \rightarrow \mathbf{0}$ almost surely.

Proof of Lemma B.19. Fix some path $h = (s^t)_{t=0}^\infty \in \mathcal{F}$ along which the differential martingale converges to zero. By part (a) of Theorem 3, the set of such paths has full measure, and is thus sufficient for establishing almost sure convergence. Recall that the first-order condition [FOC w_{ij}] at state (\mathbf{v}, s) reads

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

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

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

(This is [B.2], from the proof of Proposition 5.1 in Appendix B.2.) Substituting [B.32] into [B.31] and rearranging delivers

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

Now, since the Markov 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 [B.33] reduces to

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

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 Markov process is fully connected (Assumption Markov), it follows that

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

for all $i, j \in S$ (with equality when $j \leq i$). Now, part (a) of Theorem 3 states that $D_1 P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow 0$ almost surely. It follows from this, the martingale property (Proposition 5.1), and non-negativity of the directional derivative $D_1 P(\cdot, \cdot)$ on $D \times S$ (Lemma B.1) that $D_1 P(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow 0$ almost surely. Thus, [B.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 $D_1 P(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow \mathbf{0}$ \mathbf{P} -a.s. for all $i \in S$.

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

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

which, by the same argument, implies that

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

Since $\sum_{j=1}^d P_j(\mathbf{w}_i(\mathbf{v}, s), i) = D_1 P(\mathbf{w}_i(\mathbf{v}, s), i) \rightarrow 0$, this implies that $D_1 P(\mathbf{w}_i(\mathbf{v}^{(t)}, s^{(t)}), i) \rightarrow \mathbf{0}$ \mathbf{P} -a.s. for all $i \in S$.

We may now complete the proof. In either case, note that because the process $(\mathbf{v}^{(t)})_{t=0}^{\infty}$ is itself generated by the policy functions at the optimum — so that along each path $h \in \mathcal{H}$, $\mathbf{v}^{(t)}(h) = \mathbf{w}_i(\mathbf{v}^{(t-1)}(h), s^{(t-1)}(h))$ for some $i \in S$ — it must be that $D_1 P(\mathbf{v}^{(t)}, s^{(t)}) \rightarrow \mathbf{0}$ almost surely. Replicating the proof of Lemma B.16 with “ τ_t ” replaced everywhere by “ t ” establishes that $\mathbf{v}^{(t)} \rightarrow \mathbf{0}$ almost surely, as desired. \square

Lemma B.20. Let $(n^{(t)})_{t=0}^{\infty}$ be a non-decreasing sequence of random times. If $\mathbf{v}^{(n^{(t)})} \rightarrow \mathbf{0}$ almost surely (in probability), then $u_i^{(n^{(t)})} \rightarrow -\infty$ almost surely (in probability) for all $i \in S$.

Proof of Lemma B.20. Let $i \in S$ be given. The first-order condition [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 Markov 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 B.21. 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 B.21. Denote the recursive contract by ξ , and its policy functions by $\xi^f(\mathbf{v}, s, i) := u_i(\mathbf{v}, s)$ and $\xi^c(\mathbf{v}, s, i) = \mathbf{w}_i(\mathbf{v}, s)$. Let $i \in S$ be fixed. Iterating the promise keeping constraints [PK $_i$] one step ahead delivers

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

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

We can now wrap up the proof of Theorem 3.

Proof of parts (b)–(d). Part (b) follows from Lemma B.16, Lemma B.20 with the process $(n^{(t)})_{t=0}^{\infty}$ defined by $n^{(t)} := \tau^{(t)}$, and Lemma B.21. Part (c) follows from Lemmas B.18 and B.20. Part (d) follows from Lemmas B.19 and B.20. \square

B.5. Proof of Theorem 4

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

$$\begin{aligned} v_{j+1} - v_j &\geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) + \alpha \left(\mathbf{E}^{\mathbf{f}_{j+1}} [\mathbf{w}_j] - \mathbf{E}^{\mathbf{f}_j} [\mathbf{w}_j] \right) \\ \text{[B.37]} \quad &\geq U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) \end{aligned}$$

where the second inequality follows from the FOSD assumption and part (d) of Theorem 1. It is thus sufficient to show that the difference of flow utilities in [B.37] grows without bound.

By part (c) of Theorem 3, $u_j^{(t)} \rightarrow -\infty$ 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,

$$\left[U(\omega_{j+1} + C(u_j^{(t)}, j)) - U(\omega_j + C(u_j^{(t)}, j)) \right] \rightarrow \infty$$

in probability because $U(\underline{c} + (\omega_{j+1} - \omega_j)) > -\infty$.

Case 2: Suppose now that $\underline{c} = -\infty$. Observe that

$$\begin{aligned} U(\omega_{j+1} + C(u_j, j)) - U(\omega_j + C(u_j, j)) &= \int_{\omega_j + C(u_j, j)}^{\omega_{j+1} + C(u_j, j)} U'(y) dy \\ &\geq (\omega_{j+1} - \omega_j) U'(\omega_{j+1} + C(u_j, j)) \end{aligned}$$

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 (d) of Theorem 3, a sufficient condition for $u_j^{(t)} \rightarrow -\infty$ (each $j \in S$) almost surely is that the Markov process satisfies UPR. Note that PPR is exactly the intersection of FOSD and UPR. The claim concerning conditional variances then follows as a trivial consequence. \square

C. Proofs for Section 7

In this appendix, we provide proofs of Facts 1–6 and Lemma 7.2, and derivations of various formulas stated in Section 7.

It is a standing assumption throughout this Appendix C that the agent has CARA utility with coefficient of absolute risk aversion $\rho > 0$, so that $U(c) = -e^{-\rho c}$. Recall that $\theta_i := e^{-\rho \omega_i}$ so that $\psi(u, i, j) = (\theta_i / \theta_j) u$. Unless stated, we make no assumptions concerning $d \geq 2$ or the Markov process beyond Assumption Markov.

Throughout, let ξ^* denote the optimal contract (which 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 functions.

C.1. Homogeneity Properties

We begin with homogeneity properties of the value function and optimal contract.

Lemma C.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

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

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

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

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

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

and (ii) that the graph of the constraint correspondence $\Gamma(\cdot)$ is a convex cone by part (c) of Lemma A.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 C.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 $i \in S$,

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

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

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

(c) For each $i \in S$, the efficiency set E_i (defined in Appendix B.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:

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

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 B.3. Part (d) follows immediately from part (c) and examination of the optimality conditions in Appendix B.1. \square

Lemma C.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 C.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

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

Proof. By Remark 10, 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 C.1 and part (d) of Lemma C.2. \square

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

Lemma C.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 B.13 and B.14. The second statement follows from the first and part (d) of Lemma C.2. \square

Lemma C.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 u_i], and since $\mathbf{v} \in E_s$, it follows from [$\tilde{\mathbf{E}}_i$] 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 C.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 C.2. \square

C.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 [TVC]-regular environment.)

Proof of Fact 1. This is an immediate corollary of part (b) of Lemma C.1. □

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 simply to keep things self contained.

Proof of Fact 2. Part (a) follows from part (d) of Lemma C.2. Part (b) follows from Remark 10.

Consider part (c). Because D is a convex cone in \mathcal{U}^d (part (a) of Theorem 1) and E is a ray in D (part (d) of Lemma C.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 B.6 and C.4 (noting that $E_2 = E$ by part (d) of Lemma C.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 C.2, the fact that E is “upward sloping,” and Lemma B.5 that $\mathbf{w}_2(\mathbf{v}, s) \gg \mathbf{v}$, as desired. Moreover, by the martingale property (Proposition 5.1), we have $D_1 P(\mathbf{w}_1(\mathbf{v}, s), 1) < D_1 P(\mathbf{v}, s)$. Remark 10, 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). □

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

Consider part (c). By Lemmas B.6 and C.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 C.2, the fact that E_2 is “upward sloping,” and Lemma B.5 that $\mathbf{w}_2(\mathbf{v}, s) \gg \mathbf{v}$, as desired.

To prove parts (a) and (d), we rely on the following lemma. Recall that multipliers $\eta_i(\mathbf{v})$ and $\sigma_{ij}(\mathbf{v})$ from the “interim” problem described in Appendix B.3.3 (just below [L_{*i*}]).

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

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

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

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

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

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

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

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

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

and η_2 at $i = 1$ reduces to

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

Thus, by the envelope formulas $[\text{Env}_i]$ and the definition of $[\tilde{E}_i]$, $\mathbf{v} \in E_s$ 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 $[\text{FOC}_{w_{ij}}]$ for $i = 1$ delivers

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

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

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

Now, in the $d = 2$ case, $[\text{B.14}]$ 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

Now, with Lemma C.6 in hand, we can complete the proof of parts (a) and (c) of Fact 3. Let there be positive persistence, ie, $f_{22} - f_{12} \equiv f_{11} - f_{21} > 0$, as specified in the fact. Consider part (a). Let $\mathbf{v} \in E_2$. Part (c) of Lemma C.6, Lemma C.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 $[\text{C.9}]$ and $[\text{C.10}]$ (together with the envelope formulas $[\text{Env}_i]$), 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 (Theorem 2) 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 proves part (a). The proof of part (d) is similar, invoking part (b) of Lemma C.6. \square

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

Proof of Fact 5. Constancy on the insurance is an immediate corollary of Lemma C.3, part (b) of Lemma C.1, and the formula in [7.2]. Similarly, that the intertemporal wedge takes on at most two values is an immediate corollary of Lemma C.3, part (b) of Lemma C.1, and the formula in [7.3]. Strict positivity of the insurance wedge follows from the formula in [7.2] and Lemma C.5. \square

Proof of Fact 6. The constancy properties follow from the formulas in [7.2] and [7.3], part (b) of Lemma C.1, Lemma B.4, and part (c) of Lemma C.2. The strict positivity of the insurance wedge then follows from C.5. \square

C.4. Other Derivations

We now derive [7.1] and [7.4], and prove Lemma 7.2. (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 Markov process satisfies MLRP.)

Consider [7.1]. First, take [FOC 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 B.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 [7.1].

Consider [7.4]. 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

$$C'(u_2(\mathbf{v}, s), 2) - C'(u_1(\mathbf{v}, s), 1) = \left(\frac{f_{s1}}{f_{s2}} + \frac{\theta_2}{\theta_1} \right) \cdot \frac{\mu_{21}(\mathbf{v}, s)}{f_{s1}} + \frac{P_2(\mathbf{v}, s)}{f_{s2}} - \frac{P_1(\mathbf{v}, s)}{f_{s1}}$$

[7.4] then follows from substituting [7.1] into the above display and arranging time indices appropriately.

Finally, we prove the lemma.

Proof of Lemma 7.2. Multiply both sides of [FOC u_i] ($i = 1$) through by θ_1 to get

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

Multiply both sides of [FOC u_i] ($i = 2$) through by θ_2 to get

$$f_{s2} \theta_2 C'(u_2(\mathbf{v}, s), 2) = \theta_2 \lambda_2(\mathbf{v}, s) + \theta_2 \mu_{21}(\mathbf{v}, s)$$

Adding these displays gives

$$\begin{aligned} \text{[C.13]} \quad \mathbf{E}^{\mathbf{f}_s} [\theta_i \cdot C'(u_i(\mathbf{v}, s), i)] &= \sum_{i=1}^2 \theta_i \lambda_i(\mathbf{v}, s) \\ &= \mathbf{D}_{\boldsymbol{\theta}} P(\mathbf{v}, s) \end{aligned}$$

where $\theta := (\theta_1, \theta_2)$ and the second line follows from the envelope formulas [Env_{*i*}]. Notice that using the identify

$$\text{Cov}^{\mathbf{f}_s} [\theta_i, C'(u_i(\mathbf{v}, s), i)] := \mathbf{E}^{\mathbf{f}_s} [\theta_i \cdot C'(u_i(\mathbf{v}, s), i)] - \mathbf{E}^{\mathbf{f}_s} [\theta_i] \cdot \mathbf{E}^{\mathbf{f}_s} [C'(u_i(\mathbf{v}, s), i)]$$

and dividing [C.13] through by $m_s := \mathbf{E}^{\mathbf{f}_s} [\theta_i]$ delivers the expression in [7.6].

Let us compute

$$\begin{aligned} D_{\hat{\theta}_s} P(\mathbf{v}, s) - D_1 P(\mathbf{v}, s) &= \sum_{i=1}^2 \left(\frac{\theta_i}{m_s} - 1 \right) \lambda_i(\mathbf{v}, s) \\ &= \sum_{i=1}^2 \left(\frac{f_{si} \theta_i}{m_s} - f_{si} \right) \frac{\lambda_i(\mathbf{v}, s)}{f_{si}} \\ &= \left(\frac{m_s}{m_s} - 1 \right) \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} + \left(\frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} - \frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} \right) \cdot \left(\frac{f_{s2} \theta_2}{m_s} - f_{s2} \right) \\ &= f_{s2} \cdot \left(\frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} \right) \cdot \left(1 - \frac{\theta_2}{m_s} \right) \end{aligned}$$

Now, when the Markov process is iid, we know from Remark 10 (and the envelope formulas [Env_{*i*}]) that

$$\frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} = \frac{P_1(\mathbf{v}, s)}{f_{s1}} - \frac{P_2(\mathbf{v}, s)}{f_{s2}} = 0$$

Thus, in the iid case, we have $D_{\hat{\theta}_s} P(\mathbf{v}, s) = D_1 P(\mathbf{v}, s)$ at all (\mathbf{v}, s) reached on the optimal path. The discussion at the end of Section 5.1 and [B.3] then complete the proof of part (a).

Now consider part (b). On the optimal path, we know from [7.1] and Lemma C.4 that

$$\frac{\lambda_1(\mathbf{v}, s)}{f_{s1}} - \frac{\lambda_2(\mathbf{v}, s)}{f_{s2}} = \frac{P_1(\mathbf{v}, s)}{f_{s1}} - \frac{P_2(\mathbf{v}, s)}{f_{s2}} > 0$$

whenever $s = 1$. Since $\theta_2 < \theta_1$ implies that $\theta_2 < m_s$, it follows that $D_{\hat{\theta}_s} P(\mathbf{v}, s) - D_1 P(\mathbf{v}, s) \geq 0$ on the optimal path (with equality at $s = 2$ and strict inequality at $s = 1$). This also establishes that the $D_{\hat{\theta}_s} P(\mathbf{v}, s)$ process does not satisfy the martingale property, for when $s = 2$, we have

$$\begin{aligned} D_{\hat{\theta}_2} P(\mathbf{v}, 2) &= D_1 P(\mathbf{v}, 2) \\ &= \sum_{i=1}^2 f_{2i} D_1 P(\mathbf{w}_i(\mathbf{v}, 2), i) \\ &< \sum_{i=1}^2 f_{2i} D_{\hat{\theta}_s} P(\mathbf{w}_i(\mathbf{v}, 2), i) \end{aligned}$$

□

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 pp. 20, 53).

- Albanesi, Stefania, and Roc Armenter. 2012. ‘Intertemporal Distortions in the Second Best’. *Review of Economic Studies* 79 (4): 1271–1307. (Cit. on pp. 9, 33, 49).
- Albanesi, Stefania, and Christopher Sleet. 2006. ‘Dynamic Optimal Taxation with Private Information’. *Review of Economic Studies* 73 (1): 1–30. (Cit. on pp. 8, 34).
- 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. 6).
- 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, 5, 33, 34, 37).
- Battaglini, Marco. 2005. ‘Long-Term Contracting with Markovian Consumers’. *American Economic Review* 95 (3): 637–658. (Cit. on p. 7).
- Battaglini, Marco, and Rohit Lamba. 2018. *Optimal Dynamic Contracting: The First-Order Approach and Beyond*. Tech. rep. (Cit. on pp. 8, 37).
- Bergemann, Dirk, and Philipp Strack. 2015. ‘Dynamic Revenue Maximization: A Continuous Time Approach’. *Journal of Economic Theory* 159:819–853. (Cit. on p. 39).
- Bergemann, Dirk, and Juuso Välimäki. 2018. ‘Dynamic Mechanism Design: An Introduction’. *Journal of Economic Literature* forthcoming. (Cit. on p. 7).
- 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. 8).
- Clementi, Gian Luca, and Hugo A. Hopenhayn. 2006. ‘A Theory of Financing Constraints and Firm Dynamics’. *Quarterly Journal of Economics* 121:229–265. (Cit. on pp. 7, 36).
- DeMarzo, Peter, and Yuliy Sannikov. 2006. ‘Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model’. *Journal of Finance* 61 (6): 2681–2724. (Cit. on p. 36).
- 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 pp. 6, 8).
- 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. 55, 61, 64).
- Esö, Péter, and Bal’asz Szentes. 2017. ‘Dynamic contracting: An irrelevance theorem’. *Theoretical Economics* 12:109–139. (Cit. on p. 7).
- Farhi, Emmanuel, and Iván Werning. 2007. ‘Inequality and Social Discounting’. *Journal of Political Economy* 115 (3): 365–402. (Cit. on p. 6).
- . 2013. ‘Insurance and Taxation over the Lifecycle’. *Review of Economic Studies* 80 (2): 596–635. (Cit. on pp. 2, 8, 12, 34, 38, 45, 48).
- 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. 36).
- 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. 3, 4, 6–8, 18, 20, 24, 26, 33, 36).

- Fu, Shiming, and R. Vijay Krishna. 2017. ‘Dynamic Financial Contracting with Persistence’. *RAND Journal of Economics* forthcoming. (Cit. on pp. 7, 8, 36, 52).
- Garrett, Daniel, Alessandro Pavan and Juuso Toikka. 2018. *Robust Predictions in Dynamic Screening*. Tech. rep. Northwestern University. (Cit. on pp. 9, 30, 39).
- Gärtner, Bernd, and Jiri Matousek. 2012. *Approximation Algorithms and Semidefinite Programming*. Springer-Verlag. (Cit. on p. 64).
- 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. 9, 28, 34, 37, 38, 49).
- Golosov, Mikhail, Maxim Troshkin and Aleh Tsyvinski. 2016. ‘Redistribution and Social Insurance’. *American Economic Review* 106 (2): 359–386. (Cit. on pp. 2, 8, 12, 34, 45, 49).
- Golosov, Mikhail, and Aleh Tsyvinski. 2006. ‘Designing optimal disability insurance: A case for asset testing’. *Journal of Political Economy* 114 (2): 257–279. (Cit. on p. 36).
- . 2007. ‘Optimal taxation with endogenous insurance markets’. *The Quarterly Journal of Economics* 122 (2): 487–534. (Cit. on p. 36).
- Golosov, Mikhail, Aleh Tsyvinski and Nicolas Werquin. 2016. ‘Recursive Contracts and Endogenously Incomplete Markets’. In *Handbook of Macroeconomics, Volume 2*. (Cit. on pp. 6, 12, 35, 38).
- 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. 5, 13, 43).
- Guo, Yingni, and Johannes Hörner. 2018. *Dynamic Allocation without Money*. Tech. rep. (Cit. on pp. 7, 8).
- Halac, Marina, and Pierre Yared. 2014. ‘Fiscal rules and discretion under persistent shocks’. *Econometrica* 82 (5): 1557–1614. (Cit. on p. 8).
- Hatcher, Allen. 2001. *Algebraic Topology*. Cambridge University Press. (Cit. on p. 76).
- Hörner, Johannes, Xiaosheng Mu and Nicolas Vieille. 2017. *Keeping Your Story Straight: Truth-telling and Liespotting*. Tech. rep. Yale University. (Cit. on p. 11).
- 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. 6).
- 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 pp. 8, 12, 38, 42).
- Khan, Aubhik, and Bashkar Ravikumar. 2001. ‘Growth and risk-sharing with private information’. *Journal of Monetary Economics* 47 (3): 499–521. (Cit. on p. 6).
- Kocherlakota, Narayana R. 1996. ‘Implications of Efficient Risk Sharing without Commitment’. *Review of Economic Studies* 63 (4): 595–609. (Cit. on p. 6).
- Kocherlakota, Narayana R. 2010. *The New Dynamic Public Finance*. Princeton, NJ: Princeton University Press. (Cit. on pp. 2, 8, 33, 34).

- Krasikov, Iliia, and Rohit Lamba. 2018. *A Theory of Dynamic Contracting with Financial Constraints*. Tech. rep. Penn State. (Cit. on pp. 7, 8, 52).
- 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. 16).
- Krishna, R. Vijay, Giuseppe Lopomo and Curtis R. Taylor. 2013. ‘Stairway to Heaven or Highway to Hell: Liquidity, Sweat Equity, and the Uncertain Path to Ownership’. *RAND Journal of Economics* 44 (1): 104–127. (Cit. on p. 7).
- Ljungqvist, Lars, and Thomas J. Sargent. 2012. *Recursive Macroeconomic Theory*. 3rd ed. Cambridge, Ma: MIT Press. (Cit. on p. 6).
- 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. 68, 71, 79).
- Matoušek, Jiří, and Bernd Gärtner. 2007. *Understanding and Using Linear Programming*. New York: Springer. (Cit. on p. 59).
- Meghir, Costas, and Luigi Pistaferri. 2004. ‘Income variance dynamics and heterogeneity’. *Econometrica* 72 (1): 1–32. (Cit. on p. 2).
- Pavan, Alessandro. 2016. ‘Dynamic mechanism design: Robustness and endogenous types’. In *2015 World Congress of the Econometric Society*. (Cit. on p. 9).
- Pavan, Alessandro, Ilya Segal and Juuso Toikka. 2014. ‘Dynamic Mechanism Design: A Myersonian Approach’. *Econometrica* 82 (2): 601–653. (Cit. on pp. 3, 8, 14, 38, 49).
- Phelan, Christopher. 1998. ‘On the Long Run Implications of Repeated Moral Hazard’. *Journal of Economic Theory* 79:174–191. (Cit. on pp. 6, 7, 34, 36).
- . 2006. ‘Opportunity and social mobility’. *Review of Economic Studies* 73 (2): 487–504. (Cit. on p. 6).
- . 1995. ‘Repeated Moral Hazard and One-Sided Commitment’. *Journal of Economic Theory* 66:488–506. (Cit. on p. 6).
- Renault, Jerome, Eilon Solan and Nicolas Vieille. 2013. ‘Dynamic Sender-Receiver Games’. *Journal of Economic Theory* 148 (2): 502–534. (Cit. on p. 11).
- Shiryayev, Albert N. 1995. *Probability*. Springer. (Cit. on p. 83).
- Spear, Stephen E, and Sanjay Srivastava. 1987. ‘On Repeated Moral Hazard with Discounting’. *Review of Economic Studies* 54 (4): 599–617. (Cit. on p. 5).
- 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. 15, 16).
- Storesletten, Kjetil, Christopher I. Telmer and Amir Yaron. 2004a. ‘Consumption and risk sharing over the life cycle’. *Journal of Monetary Economics* 51 (3): 609–633. (Cit. on p. 2).
- . 2004b. ‘Cyclical dynamics in idiosyncratic labor market risk’. *Journal of Political Economy* 112 (3): 695–717. (Cit. on p. 2).

- Strulovici, Bruno. 2011. *Renegotiation-proof contracts with moral hazard and persistent private information*. Tech. rep. Discussion Paper, Center for Mathematical Studies in Economics and Management Science. (Cit. on pp. 6, 37–39).
- 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, 3, 5, 6, 10, 13, 22, 25–29, 32, 33, 37, 38, 40, 43, 45, 47, 49, 51, 93).
- . 1988. ‘Self-Enforcing Wage Contracts’. *Review of Economic Studies* 55 (4): 541–553. (Cit. on p. 6).
- Van Tiel, Jan. 1984. *Convex Analysis*. Wiley. (Cit. on p. 76).
- Vohra, Rakesh V. 2005. *Advanced Mathematical Economics*. New York, NY: Routledge. (Cit. on p. 59).
- Williams, Noah. 2011. ‘Persistent Private Information’. *Econometrica* 79 (4): 1233–1275. (Cit. on pp. 3, 6, 8, 12, 28, 33, 36–39, 42, 49).
- Zhang, Yuzhe. 2009. ‘Dynamic Contracting with Persistent Shocks’. *Journal of Economic Theory* 144:635–675. (Cit. on pp. 3, 6–8, 29, 34, 37, 40, 42, 45, 47, 49).