

Computing Dynamic Optimal Mechanisms When Hidden Types Are Markov*

Kenichi Fukushima

Yuichiro Waki

University of Wisconsin, Madison

University of Queensland

September 25, 2011

Abstract

Consider a dynamic mechanism design problem in which the agent's hidden type follows an N -state Markov chain as in Fernandes and Phelan (2000). If the Markov chain's transition probabilities are mixtures of $K(\leq N)$ densities whose mixture proportions encapsulate the dependence on the previous state, there exists a well-behaved recursive formulation of the problem with K , as opposed to N , continuous state variables. This result makes it possible to formulate computationally tractable models in which the hidden type process contains both persistent and transitory components, each of which can take many distinct values.

1 Introduction

In recent years, dynamic mechanism design theory has found applications in a variety of fields ranging from finance to public economics. As a result of this success, there is now ongoing interest in making these applications quantitative. One technical challenge in doing so however has been computational: it has proved challenging to compute dynamic optimal mechanisms in many realistically parametrized models using existing methods.

To describe the source of this challenge, it is useful to start with some background. The first models to highlight the role of multiperiod contracts when informational asymmetries

*An earlier version of the paper was entitled "Computing Dynamic Optimal Mechanisms When Private Shocks Are Persistent." We thank workshop and conference participants, the reviewers, Johannes Hörner, Larry Jones, Narayana Kocherlakota, Ellen McGrattan, Kevin Wiseman, and especially Chris Phelan for useful communications. Valuable computational resources were provided by the Minnesota Supercomputing Institute during an early stage of the project.

impede risk sharing were presented by Radner (1981) and Townsend (1982). These papers showed how superior insurance can be achieved through contracts that span multiple periods. The crux of their findings was that such contracts can give people stronger incentives to behave honestly by making resource transfers history dependent. This insight was an important one, but the implied history dependence also made it challenging to obtain sharper characterizations of optimal contracts.

A breakthrough in conquering this challenge was achieved by Green (1987), Spear and Srivastava (1987), and Thomas and Worrall (1990), who showed how optimal multiperiod contracts can be characterized as solutions to well-behaved¹ dynamic programming problems with small state spaces. Their essential step was to take continuation utilities as an endogenous state variable, an idea related to one explored by Abreu, Pearce, and Stacchetti (1990) in the context of repeated games. Somewhat remarkably, this made it possible to compute and describe such contracts by keeping track of a single variable instead of entire histories. The resulting recursive formulation led to a substantial clarification of the mechanics and served as a useful foundation for a number of theoretical studies.

In its original form, however, this formulation was still limited by the fact that it is valid only if privately observed shocks are taken to be serially independent. This restriction is problematic for many quantitative purposes because serial dependence is often a stark feature of reality. Indeed, one naturally thinks that such variables as income or productivity contain sizable persistent components, and it seems fair to say that this view now enjoys strong empirical support.

Motivated by this limitation, Fernandes and Phelan (2000) subsequently developed a generalized formulation which is applicable to settings in which the agent's hidden type is Markov.² The main complication involved in this generalization is that when types are serially dependent, the agent's continuation utility depends on his true type which is not observable to the planner. As a result, providing correct incentives in a recursive manner requires the planner to promise and deliver continuation utilities to *all* possible types that she may be facing. This observation led Fernandes and Phelan to develop a recursive formulation which tracks continuation utility profiles—functions $U_t(\cdot)$ that return the agent's valuation of the continuation contract as a function of what his true type might have been in the previous period—as the endogenous state variable.

While Fernandes and Phelan's formulation was a conceptually important one, using it for quantitative purposes proved to be challenging for computational reasons. The difficulty

¹Well-behavedness is important. In general, it is always possible to use a bijection between \mathbb{R}^n and \mathbb{R} to hide any history dependence. Doing so is not very useful however because the resulting formulation will be ill-behaved, with discontinuous value functions and policy functions.

²See Doepke and Townsend (2006) and Zhang (2009) for further developments along this line.

here comes from the need to track $U_t(\cdot)$ as a state variable, which implies that in order to solve a model in which the agent’s hidden type follows an N state Markov chain, one needs to work with a Bellman equation with N continuous state variables and a set operator that maps subsets of \mathbb{R}^N into subsets of \mathbb{R}^N (to compute the state space which is unknown in advance). Thus, although the computations remain manageable for $N = 2$, they quickly become infeasible as N is increased. This is problematic in many instances, as setting $N = 2$ severely restricts the cross sectional distribution of types to one concentrated at two values and rules out multivariate processes such as those containing both persistent and transitory components. As a result, many researchers interested in quantitative applications have avoided this issue by using models with simple (i.i.d. or fixed) type processes or short (two to three period) time horizons.

In this paper, we suggest a new way of ameliorating this computational difficulty. Our main theoretical result is that if the agent’s probability of transiting from state θ_- to state θ can be written as $\sum_{k=1}^K p_k(\theta)w_k(\theta_-)$, where p_k is a probability density over θ and w_k is a weight on density p_k that can depend on θ_- , the mechanism design problem admits a well-behaved recursive formulation with K continuous state variables regardless of N . This formulation therefore achieves a dimensionality reduction relative to Fernandes and Phelan’s when $K < N$. We then show that this result makes it possible to formulate computationally tractable models in which the agent’s hidden type follows a process with both persistent and transitory components, each of which can take many distinct values.

Our approach, however, does seem to be limited in terms of its ability to handle highly persistent processes. In particular, numerical examples suggest that obtaining close approximations of an $N = 15$ state Markov chain which discretizes an AR(1) process with autoregressive coefficient $\rho = 0.95$ requires $K \geq 4$. And while 4 dimensions is much better than 15 dimensions, it is not good enough in many computational environments. It remains to be seen whether if this will remain so in the future. For the time being, we simply point out that, as long as persistence is not too high, it is possible to mitigate this problem by making a modest change in the empirical definition of a model period (e.g., by changing it from one year to five years).

An important and increasingly popular way of confronting the computational challenge we address here is to use a dynamic first order approach (FOA), developed by Courty and Li (2000), Abraham and Pavoni (2008), Kapicka (2010), Williams (2011), and Pavan, Segal, and Toikka (2009), among others. Under this approach, one first solves a relaxed version of the problem in which non-local incentive constraints are ignored, and then goes on to check if the solution to the relaxed problem is indeed incentive compatible. To the extent that the ex-post validation succeeds, this approach can be more efficient than ours. This is especially

true in settings with many highly persistent types; Farhi and Werning (2010) report success in such a setting. One limitation of the FOA however is that general sufficient conditions on model primitives that guarantee its success are currently unavailable. This means that the ex-post validation must be carried out on faith and there is no clear guidance on how to proceed when it fails. Our approach does not have this problem, and can therefore be used as a fallback option if the FOA should fail. A second limitation of the FOA is that it is not clear how one would apply it in settings with multi-dimensional types (e.g., the hidden type contains persistent and transitory components). Our approach can handle at least a subset of such environments, and in this respect has broader applicability. And it can be extended to accommodate hidden actions as well (Fukushima and Waki, 2011a).

2 Problem Statement

Consider the following dynamic mechanism design problem, which subsumes versions of several well-known setups as special cases. Examples are given at the end of the section.

There is a planner and an agent, and time flows $t = 0, 1, 2, \dots$. In each period, the agent draws a shock $\theta_t \in \Theta$ and sends a report $r_t \in \Theta$ to the planner. The planner then chooses an outcome $x_t \in X$ given the agent's history of reports. We assume Θ is finite with cardinality N and $X \subset \mathbb{R}^L$ is compact.

The shock process $\{\theta_t\}_{t=0}^{\infty}$ is first order Markov, and the probability of transiting from θ_- to θ is $\pi(\theta|\theta_-)$. Its initial distribution is $\pi(\cdot|\theta_{-1})$, where θ_{-1} is a publicly known value. As well, $\pi(\theta|\theta_-) > 0$ for all θ and θ_- . For $t \geq s \geq 0$ we write $\theta_s^t = (\theta_s, \dots, \theta_t)$, $\theta^t = \theta_0^t$, and $\Pr(\theta_s^t|\theta_{s-1}) = \pi(\theta_t|\theta_{t-1}) \cdots \pi(\theta_s|\theta_{s-1})$.

We define an *allocation* as a sequence $\mathbf{x} = \{x_t\}_{t=0}^{\infty}$, $x_t : \Theta^{t+1} \rightarrow X$ for each t , and let \mathbf{X} denote the set of all allocations.

If allocation \mathbf{x} takes place, that is, if the outcome $x_t(\theta^t)$ occurs after each shock history θ^t , the agent obtains lifetime utility

$$U(\mathbf{x}; \theta_{-1}) = \sum_{t=0}^{\infty} \sum_{\theta^t} \beta^t u(x_t(\theta^t); \theta_t) \Pr(\theta^t|\theta_{-1})$$

where $\beta \in (0, 1)$ and $u : X \times \Theta \rightarrow \mathbb{R}$ has the property that each $u(\cdot; \theta)$ is continuous. We let $V = [\min u(X; \Theta), \max u(X; \Theta)] / (1 - \beta)$ denote the closed interval to which U always belongs. The cost for the planner is

$$C(\mathbf{x}; \theta_{-1}) = \sum_{t=0}^{\infty} \sum_{\theta^t} q^t c(x_t(\theta^t)) \Pr(\theta^t|\theta_{-1})$$

where $q \in (0, 1)$ and $c : X \rightarrow \mathbb{R}$ is continuous.

For the most part we will assume that the environment is *convex*, meaning that X is convex, c is convex, and each $u(\cdot; \theta)$ is linear. It is often possible to obtain this property by a suitable change of variables (see the examples below). More generally, it can be guaranteed by introducing lotteries (following Prescott and Townsend, 1984).

The planner is placed in this environment with the ability to commit and with the obligation to provide lifetime utility U_0 to the agent. Her goal is to fulfill this obligation in a cost-minimizing way, respecting the agent's incentives.

To formulate the planner's problem, we invoke the revelation principle and say that an allocation \mathbf{x} is *feasible* if it satisfies two conditions. The first condition says that it is optimal for the agent to report truthfully. Formally, define a reporting strategy as a sequence $\mathbf{r} = \{r_t\}_{t=0}^\infty$, $r_t : \Theta^{t+1} \rightarrow \Theta$ for each t , and let \mathbf{R} be the set of all reporting strategies. We say that \mathbf{x} is *incentive compatible* if

$$U(\mathbf{x}; \theta_{-1}) \geq U(\mathbf{x} \circ \mathbf{r}; \theta_{-1}), \quad \forall \mathbf{r} \in \mathbf{R} \quad (1)$$

where $\mathbf{x} \circ \mathbf{r} = \{x_t \circ r^t\}_{t=0}^\infty$, $r^t = (r_0, \dots, r_t)$. It is straightforward to check that the set of incentive compatible allocations does not depend on θ_{-1} .

The second condition says that the agent indeed gets lifetime utility U_0 :

$$U(\mathbf{x}; \theta_{-1}) \geq U_0. \quad (2)$$

We say that \mathbf{x} satisfies *promise keeping* if this holds.

The *planning problem* given initial condition (θ_{-1}, U_0) is then to choose an allocation \mathbf{x} so as to minimize $C(\mathbf{x}; \theta_{-1})$ subject to feasibility. We assume U_0 is such that this problem has a non-empty constraint set.

Example (Hidden Income). When $L = 1$, $\Theta \subset \mathbb{R}$, $u(x; \theta) = v(x + \theta)$, and $c(x) = x$, the model specializes to the hidden endowment model of Green (1987) and Thomas and Worrall (1990). In this model, θ is the agent's hidden income and x is an additional income transfer; thus the agent's consumption is $x + \theta$. A standard way of obtaining convexity is to assume CARA utility $v(c) = -\exp(-\gamma c)$ ($\gamma > 0$) and use the change of variable $\tilde{x} = -\exp(-\gamma x)$.

Example (Hidden Tastes). When $L = 1$, $\Theta \subset \mathbb{R}_{++}$, $u(x; \theta) = \theta v(x)$, and $c(x) = x$, the model specializes to a version of the Atkeson and Lucas (1992) hidden taste shock model. Here, x is the agent's consumption and θ is a taste shock representing his urgency to consume, say due to illness. With the change of variable $\tilde{x} = v(x)$, the environment becomes convex.

Example (Hidden Skills). When $L = 2$, $\Theta \subset \mathbb{R}_{++}$, $u(x; \theta) = v_1(x_1) - v_2(x_2/\theta)$, and $c(x) = x_2 - x_1$, the model specializes to a dynamic extension of Mirrlees (1971), versions of which are used in a literature overviewed by Kocherlakota (2010). Here, x_1 is consumption, x_2 is labor output, and θ is a hidden skill level. The idea is that if the agent exerts effort e , he produces $x_2 = \theta e$ and incurs disutility $v_2(e) = v_2(x_2/\theta)$. Convexity obtains when $v_2(e) = e^\gamma$ ($\gamma \geq 1$) under the change of variables $(\tilde{x}_1, \tilde{x}_2) = (v_1(x_1), v_2(x_2))$.

3 Recursive Formulation

This section presents our recursive formulation of the planning problem. Our starting point is the following version of the one-shot deviation principle. (All proofs are in the Appendix.)

Lemma 1. *An allocation \mathbf{x} is incentive compatible if and only if*

$$\begin{aligned} u(x_t(\theta^{t-1}, \theta_t); \theta_t) + \beta \sum_{s=t+1}^{\infty} \sum_{\theta_{t+1}^s} \beta^{s-(t+1)} u(x_s(\theta^{t-1}, \theta_t, \theta_{t+1}^s); \theta_s) \Pr(\theta_{t+1}^s | \theta_t) \\ \geq u(x_t(\theta^{t-1}, \theta'_t); \theta_t) + \beta \sum_{s=t+1}^{\infty} \sum_{\theta_{t+1}^s} \beta^{s-(t+1)} u(x_s(\theta^{t-1}, \theta'_t, \theta_{t+1}^s); \theta_s) \Pr(\theta_{t+1}^s | \theta_t) \end{aligned} \quad (3)$$

for all t , θ^{t-1} , θ_t , and θ'_t .

In the special case with i.i.d. shocks, the conditional probabilities in the second terms of both sides of (3) are independent of the agent's true type θ_t . Green (1987), Spear and Srivastava (1987), and Thomas and Worrall (1990) exploited this property and rewrote (3) as

$$u(x_t(\theta^{t-1}, \theta_t); \theta_t) + \beta U_{t+1}(\theta^{t-1}, \theta_t) \geq u(x_t(\theta^{t-1}, \theta'_t); \theta_t) + \beta U_{t+1}(\theta^{t-1}, \theta'_t)$$

where

$$U_t(\theta^{t-1}) = \sum_{s=t}^{\infty} \sum_{\theta_t^s} \beta^{s-t} u(x_s(\theta^{t-1}, \theta_t^s); \theta_s) \Pr(\theta_t^s)$$

is the agent's continuation utility after history θ^{t-1} . They then used these two conditions to rewrite the planning problem as a standard dynamic programming problem, taking $U_t(\theta^{t-1})$ as the state variable.

In the more general case where the shocks are not i.i.d., (3) needs to be written as:

$$u(x_t(\theta^{t-1}, \theta_t); \theta_t) + \beta U_{t+1}(\theta^{t-1}, \theta_t; \theta_t) \geq u(x_t(\theta^{t-1}, \theta'_t); \theta_t) + \beta U_{t+1}(\theta^{t-1}, \theta'_t; \theta_t)$$

where

$$U_t(\theta^{t-1}; \theta_-) = \sum_{s=t}^{\infty} \sum_{\theta_t^s} \beta^{s-t} u(x_s(\theta^{t-1}, \theta_t^s); \theta_s) \Pr(\theta_t^s | \theta_-)$$

describes the agent's continuation utility profile—a function $U_t(\theta^{t-1}; \cdot) : \Theta \rightarrow \mathbb{R}$ that returns the agent's continuation utility as a function of what his true type might have been in the previous period. This means that if we are to follow the above approach, we must track the N dimensional variable $U_t(\theta^{t-1}; \cdot)$. The reason for this is simple: When deciding whether to misreport his type today, the agent compares the immediate gains from doing so and the long term consequences. But unless the shocks are i.i.d., the agent's valuation of the continuation allocation depends on his current type, which is not observable to the planner. It follows that in order to correctly provide incentives, the planner must promise and deliver continuation utilities to *all* possible types that she may be facing. Building on this observation, Fernandes and Phelan (2000) developed a recursive formulation which takes $U_t(\theta^{t-1}; \cdot)$ as the endogenous state variable.

While Fernandes and Phelan's approach was a conceptually straightforward response to the situation just described, it is not a convenient one for computations as it leads to a curse of dimensionality which makes the formulation essentially unusable when N is large. In what follows we describe how this limitation may be overcome when the shock process is taken to have a special structure. What we show is that by exploiting that structure, it is possible to track the continuation utility profile $U_t(\theta^{t-1}; \cdot)$ more efficiently and thereby obtain a recursive formulation of the problem with a smaller state space.

Definition. The transition kernel π has an *order K mixture representation* if we can write:

$$\pi(\theta | \theta_-) = \sum_{k=1}^K p_k(\theta) w_k(\theta_-), \quad (4)$$

where $K \in \{1, \dots, N\}$ and $(p, w) : \Theta \rightarrow \mathbb{R}_+^K \times \mathbb{R}_+^K$ satisfies $\sum_{\theta} p_k(\theta) = 1$ for each k and $\sum_k w_k(\theta_-) = 1$ for each θ_- . Let Π_K denote the set of all transition kernels which have an order K mixture representation.

In words, (4) says that the conditional densities over θ , $\pi(\cdot | \theta_-)$, can be represented as mixtures of K densities $\{p_k\}_{k=1}^K$ where the mixture proportions $\{w_k\}_{k=1}^K$ encapsulate their dependence on θ_- . Equivalently, it says that the transition matrix Π is the product of an $N \times K$ matrix (with elements w) and a $K \times N$ matrix (with elements p), so that Π is of rank K . One can always write π in this way with $K = N$,³ but it is also true that there is

³Let $p_k(\theta) = \pi(\theta | k)$ and let $w_k(\theta_-)$ be the indicator of $k = \theta_-$. Our recursive formulation reduces to Fernandes and Phelan's under this trivial representation.

a non-trivial class of π 's for which one can do the same for $K < N$, and our interest in this representation stems from the latter fact.

Let us now imagine, given (4), that in each period the agent transits between states via a fictitious ‘‘interim state’’ k : he goes from θ_- to k with probability $w_k(\theta_-)$ and then from k to θ with probability $p_k(\theta)$. Then if we look at the vector of continuation utilities starting from the interim states:

$$a_t(\theta^{t-1}) = \sum_{\theta_t} \left\{ u(x_t(\theta^t); \theta_t) + \beta \sum_{s=t+1}^{\infty} \sum_{\theta_{t+1}^s} \beta^{s-(t+1)} u(x_s(\theta^t, \theta_{t+1}^s); \theta_s) \Pr(\theta_{t+1}^s | \theta_t) \right\} p(\theta_t) \quad (5)$$

we can see that by the law of iterated expectations:

$$U_t(\theta^{t-1}; \cdot) = \sum_{k=1}^K a_{kt}(\theta^{t-1}) w_k(\cdot). \quad (6)$$

Hence the K -dimensional variable $a_t(\theta^{t-1})$ carries all relevant information contained in the N -dimensional $U_t(\theta^{t-1}; \cdot)$. Our idea, naturally suggested by this and $K \leq N$, is to use $a_t(\theta^{t-1})$ instead of $U_t(\theta^{t-1}; \cdot)$ as our record-keeping device.

This leads us to seek a recursive formulation of the planning problem that takes a_t as the endogenous state variable. Toward this end, let us abuse notation slightly and write $a_t(\theta^{t-1}; \mathbf{x})$ to describe the mapping from \mathbf{x} to $a_t(\theta^{t-1})$ defined by (5). Let us also write $a_0(\mathbf{x}) = a_0(\theta^{-1}; \mathbf{x})$, as this is independent of θ_{-1} . Then consider the following minimization problem indexed by the initial condition $(\theta_{-1}, a_0) \in \Theta \times V^K$:

$$J^*(\theta_{-1}, a_0) = \inf_{\mathbf{x}} C(\mathbf{x}; \theta_{-1})$$

subject to

$$U(\mathbf{x}; \theta_{-1}) \geq U(\mathbf{x} \circ \mathbf{r}; \theta_{-1}), \quad \forall \mathbf{r} \in \mathbf{R} \quad (7)$$

$$a_0(\mathbf{x}) = a_0. \quad (8)$$

We call this the *auxiliary planning problem* starting from (θ_{-1}, a_0) , and let $A^* \subset V^K$ denote the set of a_0 's for which its constraint set is non-empty (note that this set is the same for all θ_{-1}). It is straightforward to check that if

$$a_0^* \in \arg \min_{a_0 \in A^*} J^*(\theta_{-1}, a_0) \quad \text{s.t.} \quad a_0 \cdot w(\theta_{-1}) \geq U_0, \quad (9)$$

then a solution to the auxiliary planning problem starting from (θ_{-1}, a_0^*) is a solution to the

planning problem starting from (θ_{-1}, U_0) .

The reason for introducing the auxiliary planning problem is that it has a stationary recursive structure which the original planning problem does not.

Lemma 2. *An allocation \mathbf{x} satisfies the constraints of the auxiliary planning problem if and only if there exists $\mathbf{a} = \{a_t\}_{t=0}^\infty$, $a_t : \Theta^t \rightarrow A^*$, such that (\mathbf{x}, \mathbf{a}) satisfies*

$$u(x_t(\theta^{t-1}, \theta_t); \theta_t) + \beta a_{t+1}(\theta^{t-1}, \theta_t) \cdot w(\theta_t) \geq u(x_t(\theta^{t-1}, \theta'_t); \theta_t) + \beta a_{t+1}(\theta^{t-1}, \theta'_t) \cdot w(\theta_t) \quad (10)$$

$$a_t(\theta^{t-1}) = \sum_{\theta_t} \{u(x_t(\theta^t); \theta_t) + \beta a_{t+1}(\theta^t) \cdot w(\theta_t)\} p(\theta_t) \quad (11)$$

for all t , θ^t , θ'_t and $a_0(\theta^{-1}) = a_0$.

The upshot of this lemma is that the auxiliary planning problem is equivalent to a problem in which one minimizes $C(\mathbf{x}; \theta_{-1})$ by choice of (\mathbf{x}, \mathbf{a}) subject to the constraints (10), (11), and $a_0(\theta^{-1}) = a_0$. It is easy to see that this rewritten problem is a standard dynamic programming problem with state space $\Theta \times A^*$.

To solve this problem using recursive methods, we first need to know what the set A^* is. For this we define an operator B , which maps $A \subset V^K$ into $B(A) \subset V^K$ defined as:

$$B(A) = \{a \in V^K \mid \exists (x, a^+) \in F(a; A)\} \quad (12)$$

where $F(a; A)$ is the set of function pairs $(x, a^+) : \Theta \rightarrow X \times A$ satisfying:

$$u(x(\theta); \theta) + \beta a^+(\theta) \cdot w(\theta) \geq u(x(\theta'); \theta) + \beta a^+(\theta') \cdot w(\theta), \quad \forall \theta, \theta' \in \Theta$$

$$a = \sum_{\theta} \{u(x(\theta); \theta) + \beta a^+(\theta) \cdot w(\theta)\} p(\theta).$$

Proposition 3. *A^* is a non-empty and compact set, and is the largest fixed point of B . If $A_0 \subset V^K$ is a compact set satisfying $A_0 \supset B(A_0) \supset A^*$ (one example being $A_0 = V^K$) then $B^n(A_0)$ is decreasing in n and $\bigcap_{n=0}^\infty B^n(A_0) = A^*$. If $A_0 \subset V^K$ satisfies $A^* \supset B(A_0) \supset A_0$ (one example being $A_0 = \{a_0(\bar{\mathbf{x}})\}$ where $\bar{\mathbf{x}}$ is a constant allocation), then $B^n(A_0)$ is increasing in n and $\text{cl}(\bigcup_{n=0}^\infty B^n(A_0)) = A^*$. If the environment is convex, B maps convex sets into convex sets and A^* is convex.*

The proof of this result is mostly an application of arguments due to Abreu, Pearce, and Stacchetti (1990). However the third part—which ensures convergence of $B^n(A_0)$ to A^* from below—is not, and as far as we know there is no general counterpart of this in the context of repeated games. We use this part later in section 4 in developing our numerical implementation.

Give this, we can now formulate and solve a Bellman equation for the problem. Define an operator T which maps $J : \Theta \times A^* \rightarrow \mathbb{R}$ into $TJ : \Theta \times A^* \rightarrow \mathbb{R}$, defined as:

$$TJ(\theta_-, a) = \inf_{(x, a^+) \in F(a; A^*)} \sum_{\theta} \{c(x(\theta)) + qJ(\theta, a^+(\theta))\} \pi(\theta|\theta_-). \quad (13)$$

Let $\|\cdot\|$ denote the supremum norm on the space of bounded real valued functions on $\Theta \times A^*$. The following standard properties hold:

Proposition 4. *J^* is a bounded lower semicontinuous function, and $\|T^n J - J^*\| \rightarrow 0$ as $n \rightarrow \infty$ for any bounded $J : \Theta \times A^* \rightarrow \mathbb{R}$. There exists a function $g^* : \Theta \times A^* \rightarrow (X \times A^*)^\Theta$ which attains the infimum on the right hand side of (13) when $J = J^*$, and for any such g^* the allocation \mathbf{x}^* defined recursively by $(x_t^*(\theta^t), a_{t+1}^*(\theta^t)) = g^*(\theta_{t-1}, a_t^*(\theta^{t-1}))(\theta_t)$ solves the auxiliary planning problem starting from $(\theta_{-1}, a_0^*(\theta^{-1}))$. If the environment is convex, each $J^*(\theta_-, \cdot)$ is convex.*

In summary, we obtain the following algorithm for solving the planner's problem: First, use Proposition 3 to iteratively compute A^* . Then solve the Bellman equation in Proposition 4 using that A^* . Finally, solve (9) to get a_0^* and roll out \mathbf{x}^* from there using the policy function g^* . This procedure is similar to one suggested by Fernandes and Phelan (2000), but better suited for numerical computations thanks to the smaller state space (and, to a lesser extent, the smaller number of control variables).

This algorithm is readily adapted to finite, \bar{t} -period settings ($t = 0, \dots, \bar{t}-1$) as follows: (i) compute a sequence of sets $\{A_t^*\}_{t=0}^{\bar{t}}$ ($A_t^* \subset \mathbb{R}^K$), as $A_{\bar{t}}^* = \{0\}$ and $A_t^* = B(A_{t+1}^*)$; (ii) compute a sequence of value functions $\{J_t^*\}_{t=0}^{\bar{t}}$ ($J_t^* : \Theta \times A_t^* \rightarrow \mathbb{R}$), as $J_{\bar{t}}^* \equiv 0$ and $J_t^* = TJ_{t+1}^*$; (iii) generate \mathbf{x}^* using the policy functions $\{g_t^*\}_{t=0}^{\bar{t}-1}$ ($g_t^* : \Theta \times A_t^* \rightarrow (X \times A_{t+1}^*)^\Theta$) which solve the minimization problems in the definition of TJ_{t+1}^* . Of course we can allow u , c , and π to be time-dependent in this case and proceed as above using the time-dependent analogs of B and T .

4 Numerical Implementation

We next describe a procedure for numerically implementing our scheme on a computer. The procedure works for general convex environments, and is designed with an emphasis on robustness. We limit our discussion to the infinite horizon case; adapting it to settings with finite horizons is straightforward.

The first step is to compute a polytope $\hat{A}^* \subset \mathbb{R}^K$ that approximates $A^* \subset \mathbb{R}^K$. A procedure for this is the following, which adapts Judd, Yeltekin, and Conklin's (2003) inner

Algorithm 1 : Solve the planning problem.

I. Compute \hat{A}^* :

1. Compute \hat{A}_I (inner approximation of A^*).

(a) Set $\hat{A}_I^{(0)} = \{a_0(\bar{\mathbf{x}})\}$ where $\bar{\mathbf{x}}$ is a constant allocation.

(b) For $n \geq 0$, let $\hat{A}_I^{(n+1)} = \hat{B}_I(\hat{A}_I^{(n)})$.

(c) If $\hat{A}_I^{(n+1)}$ and $\hat{A}_I^{(n)}$ are sufficiently close, set $\hat{A}_I = \hat{A}_I^{(n+1)}$ and go to step 2. Otherwise, set $n \rightsquigarrow n + 1$ and go back to part (b).

2. Compute \hat{A}_O (outer approximation of A^*).

(a) Set $\hat{A}_O^{(0)} = V^K$.

(b) For $n \geq 0$, let $\hat{A}_O^{(n+1)} = \hat{B}_O(\hat{A}_O^{(n)})$.

(c) If $\hat{A}_O^{(n+1)}$ and $\hat{A}_O^{(n)}$ are sufficiently close, set $\hat{A}_O = \hat{A}_O^{(n+1)}$ and go to step 3. Otherwise, set $n \rightsquigarrow n + 1$ and go back to part (b).

3. If \hat{A}_I is close enough to \hat{A}_O , set $\hat{A}^* = \hat{A}_I$ and stop. Otherwise, enlarge H and retry.

II. Compute (\hat{J}^*, \hat{g}^*) via value iteration and compute \hat{a}_0^* using (9) with \hat{J}^* replacing J^* .

ray and outer hyperplane approximation methods to our setting.

First, let H be a finite collection of vectors $h \in \mathbb{R}^K$ satisfying $0 \in \text{co}(H)$. Then define two operators \hat{B}_I and \hat{B}_O which map polytopes $\hat{A} \subset \mathbb{R}^K$ into polytopes

$$\begin{aligned}\hat{B}_I(\hat{A}) &= \text{co}(\{a_0(\bar{\mathbf{x}}) + hl(h, \hat{A})\}_{h \in H}) \\ \hat{B}_O(\hat{A}) &= \{a \in \mathbb{R}^K \mid h \cdot a \geq z(h, \hat{A}), \forall h \in H\}\end{aligned}$$

where $l(h, \hat{A})$ and $z(h, \hat{A})$ are defined in terms of linear programs:

$$\begin{aligned}l(h, \hat{A}) &= \max\{l \in \mathbb{R}_+ : a_0(\bar{\mathbf{x}}) + hl \in B(\hat{A})\} \\ z(h, \hat{A}) &= \min\{h \cdot a : a \in B(\hat{A})\}\end{aligned}$$

These operators are monotone ($\hat{A} \subset \hat{A}' \implies \hat{B}_I(\hat{A}) \subset \hat{B}_I(\hat{A}')$ and $\hat{B}_O(\hat{A}) \subset \hat{B}_O(\hat{A}')$) and approximate B from inside and outside in the sense that $\hat{B}_I(\hat{A}) \subset B(\hat{A}) \subset \hat{B}_O(\hat{A})$ for any \hat{A} . As well, $a_0(\bar{\mathbf{x}}) \in \hat{B}_I(\{a_0(\bar{\mathbf{x}})\})$.

Part I of Algorithm 1 describes how to compute \hat{A}^* using these operators. It follows from Proposition 3 and the above properties of \hat{B}_I and \hat{B}_O that: (i) the iterations in steps 1 and 2 converge monotonically; (ii) $\hat{A}^* = \hat{A}_I \subset A^* \subset \hat{A}_O$; and (iii) $\hat{A}^* \subset B(\hat{A}^*)$. It follows from

(ii) that step 3 of the algorithm provides an accuracy check with error bounds. Property (iii) ensures that the dynamic programming part below does not involve optimizations over empty constraint sets.

The next step, stated as Part II of Algorithm 1, is to solve the Bellman equation and obtain a numerical approximation of (J^*, g^*, a_0^*) , denoted $(\hat{J}^*, \hat{g}^*, \hat{a}_0^*)$, using \hat{A}^* as the state space. The only obvious approach for this part is to use value function iteration, interpolating the candidate value function in each step.

While this step is more or less standard, there are two important details. The first is how to construct a grid on the computed state space \hat{A}^* . An approach here is to first compute the half spaces whose intersection equals \hat{A}^* (for example using Barber, Dobkin, and Huhdanpaa’s (1996) Qhull package) and then use Smith’s (1984) “hit-and-run” procedure to generate pseudo random grid points on \hat{A}^* that are asymptotically uniformly distributed. The second detail concerns which interpolation scheme to use. This is a non-trivial issue because the non-rectangularity of the domain \hat{A}^* and the potential non-smoothness of the value function J^* can cause many standard methods to behave poorly, with undesirable consequences on numerical stability and solution quality (cf. Judd, 1998, p. 438). One option here is to use an approach described in Fukushima and Waki (2011b) which is designed to handle problems of this sort in a robust manner.

5 An Illustration

Let us finally use a simple example to illustrate the potential of our approach. We focus here on highlighting some key ideas and limitations; for full-blown quantitative applications, see Fukushima (2010) and Waki (2011).

We consider an optimal lending problem with hidden income and CARA utility $v(c) = -\exp(-\gamma c)$ as in the first example from section 2. Log income y_t follows a finite state version of an ARMA(1,1) process as in Storesletten, Telmer, and Yaron (2004):

$$y_t = \kappa_t + \tau_t \tag{14}$$

$$\kappa_t = \rho \kappa_{t-1} + \epsilon_t \tag{15}$$

where $\{\tau_t\}_{t=0}^\infty$ and $\{\epsilon_t\}_{t=0}^\infty$ are independent i.i.d. processes. Thus y_t is a function of a two-dimensional type vector $\theta_t = (\kappa_t, \tau_t)$ whose persistent component κ_t follows an N^κ state Markov process with transition probabilities $\mu(\kappa|\kappa_-)$ and whose transitory component τ_t is an N^τ state i.i.d. process with density $\phi(\tau)$. The hidden type θ_t therefore follows an $N = N^\kappa \times N^\tau (\geq 4)$ state Markov chain, which makes the problem challenging to solve using

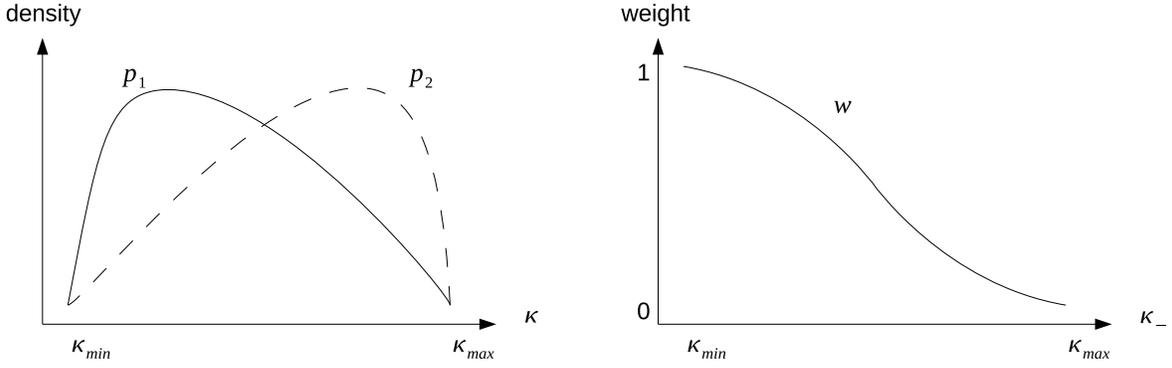


Figure 1: Example of a persistent shock process with large N and $K = 2$.

Fernandes and Phelan's (2000) N dimensional recursive formulation. In the following we show how our results can help mitigate this problem.

We begin by pointing out that we can always achieve a dimensionality reduction from N to N^κ using our approach. To see this, observe that the transition probabilities for the hidden type $\theta = (\kappa, \tau)$ can be written:

$$\pi(\kappa, \tau | \kappa_-, \tau_-) = \mu(\kappa | \kappa_-) \phi(\tau)$$

and that this fits the format (4) with $K = N^\kappa$, $p_k(\kappa, \tau) = \mu(\kappa | k) \phi(\tau)$, and $w_k(\kappa_-, \tau_-)$ equal to the indicator function of $k = \kappa_-$. The type process therefore has an order N^κ mixture representation. This alone eliminates two or more continuous state variables.

Reducing the dimensionality below N^κ is not always possible but it is in some special cases. Figure 1 illustrates the general idea using an example where $K = 2 \ll N^\kappa$. Here, μ has the structure

$$\mu(\kappa | \kappa_-) = p_1(\kappa) w(\kappa_-) + p_2(\kappa) (1 - w(\kappa_-)),$$

and the figure depicts the densities p_k ($k = 1, 2$) (left panel) and the weight w (right panel). The κ and κ_- values on the horizontal axes are allowed to take a large number of values N^κ . To see how this works, first suppose the agent has the lowest possible value of κ_- today, in which case he draws his κ tomorrow from p_1 with certainty. Then as his κ_- is increased, his next-period draw of κ comes from p_2 with higher and higher probability, until the highest possible value is reached and his draw comes exclusively from p_2 . Using this information it is not too difficult to visualize how the conditional density over κ , $\mu(\cdot | \kappa_-)$, varies with κ_- , and see from there that the process exhibits positive autocorrelation and mean reversion like a stationary AR(1).

To quantify how far this idea can take us, we next undertake a numerical exercise. We start with a “target” process for $\{\kappa_t\}$, specified as a 15 state Tauchen (1986) discretization of (15) assuming $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$, and let μ^* denote its transition kernel and $\bar{\mu}^*$ its invariant distribution. We set $\rho = 0.95$ and $\sigma_\epsilon^2 = 0.13^2$ at an annual frequency following Storesletten, Telmer, and Yaron (2004). We then ask how well we can approximate this target process using $\mu_K \in \Pi_K$ for small K . Specifically, we choose our approximating process μ_K so that:

$$\mu_K \in \arg \min_{\mu \in \Pi_K} \sum_{\kappa_-} \left\{ \sum_{\kappa} \log \left(\frac{\mu^*(\kappa|\kappa_-)}{\mu(\kappa|\kappa_-)} \right) \mu^*(\kappa|\kappa_-) \right\} \bar{\mu}^*(\kappa_-), \quad (16)$$

where the objective function is the Kullback-Leibler divergence adapted to the present Markov setting.

Figure 2 compares μ^* and $\{\mu_K\}_{K=1}^4$ along several dimensions for a quinquennial model ($\rho = 0.95^5 \approx 0.77$, $\sigma_\epsilon^2 = 0.13^2 \times \sum_{i=1}^5 0.95^{2(i-1)} \approx 0.26^2$). The top four panels report population moments of κ_t . Panel (a) depicts the density functions of the stationary distributions, which all turn out to be nearly identical. Panel (b) depicts the conditional means $E[\kappa|\kappa_-]$ as a function of κ_- . Here we can see how the lack of persistence with μ_1 (i.i.d. shocks) is remedied as we increase K ; as one might expect from (16), each μ_K attains a better match at κ_- values with high probability under the stationary distribution (cf. panel (a)). These properties are reflected in the autocorrelations, depicted in panel (c). Panel (d) depicts the conditional variances $\text{Var}[\kappa|\kappa_-]$ as functions of κ_- . The discrepancy again gets smaller as we increase K , although the improvement is less uniform. The latter effect arises because the conditional variances tend to increase at those κ_- values that lie between the peaks of the densities $\{p_k\}_{k=1}^K$.

The bottom four panels of figure 2 compare μ^* and the μ_K 's in terms of the sample paths generated by identical forcing variables.⁴ The discrepancy between the paths from μ^* and μ_1 here is quite evident. The path from μ_2 tracks that from μ^* much closer, although it shows a tendency to overshoot, reflecting its excessive conditional variance (cf. panel (d)). The discrepancy is further attenuated with μ_3 and μ_4 , where the approximation quality looks, at least to our eyes, quite good.

Figure 3 reproduces figure 2 for the annual model ($\rho = 0.95$, $\sigma_\epsilon^2 = 0.13^2$). Comparing figures 2 and 3, we can see that although the qualitative characteristics remain similar, the higher persistence here makes it harder to obtain close approximations with small K . To understand this result, let us refer back to figure 1, which is essentially a schematic

⁴Specifically, we let $\{Z_t\}$ be an i.i.d. draw from $\mathcal{U}[0, 1]$, set κ_1 to the value of the inverse c.d.f. of $\bar{\mu}$ at Z_1 , and construct subsequent κ_t 's recursively by setting κ_t equal to the value of the inverse c.d.f. of $\mu(\cdot|\kappa_{t-1})$ at Z_t .

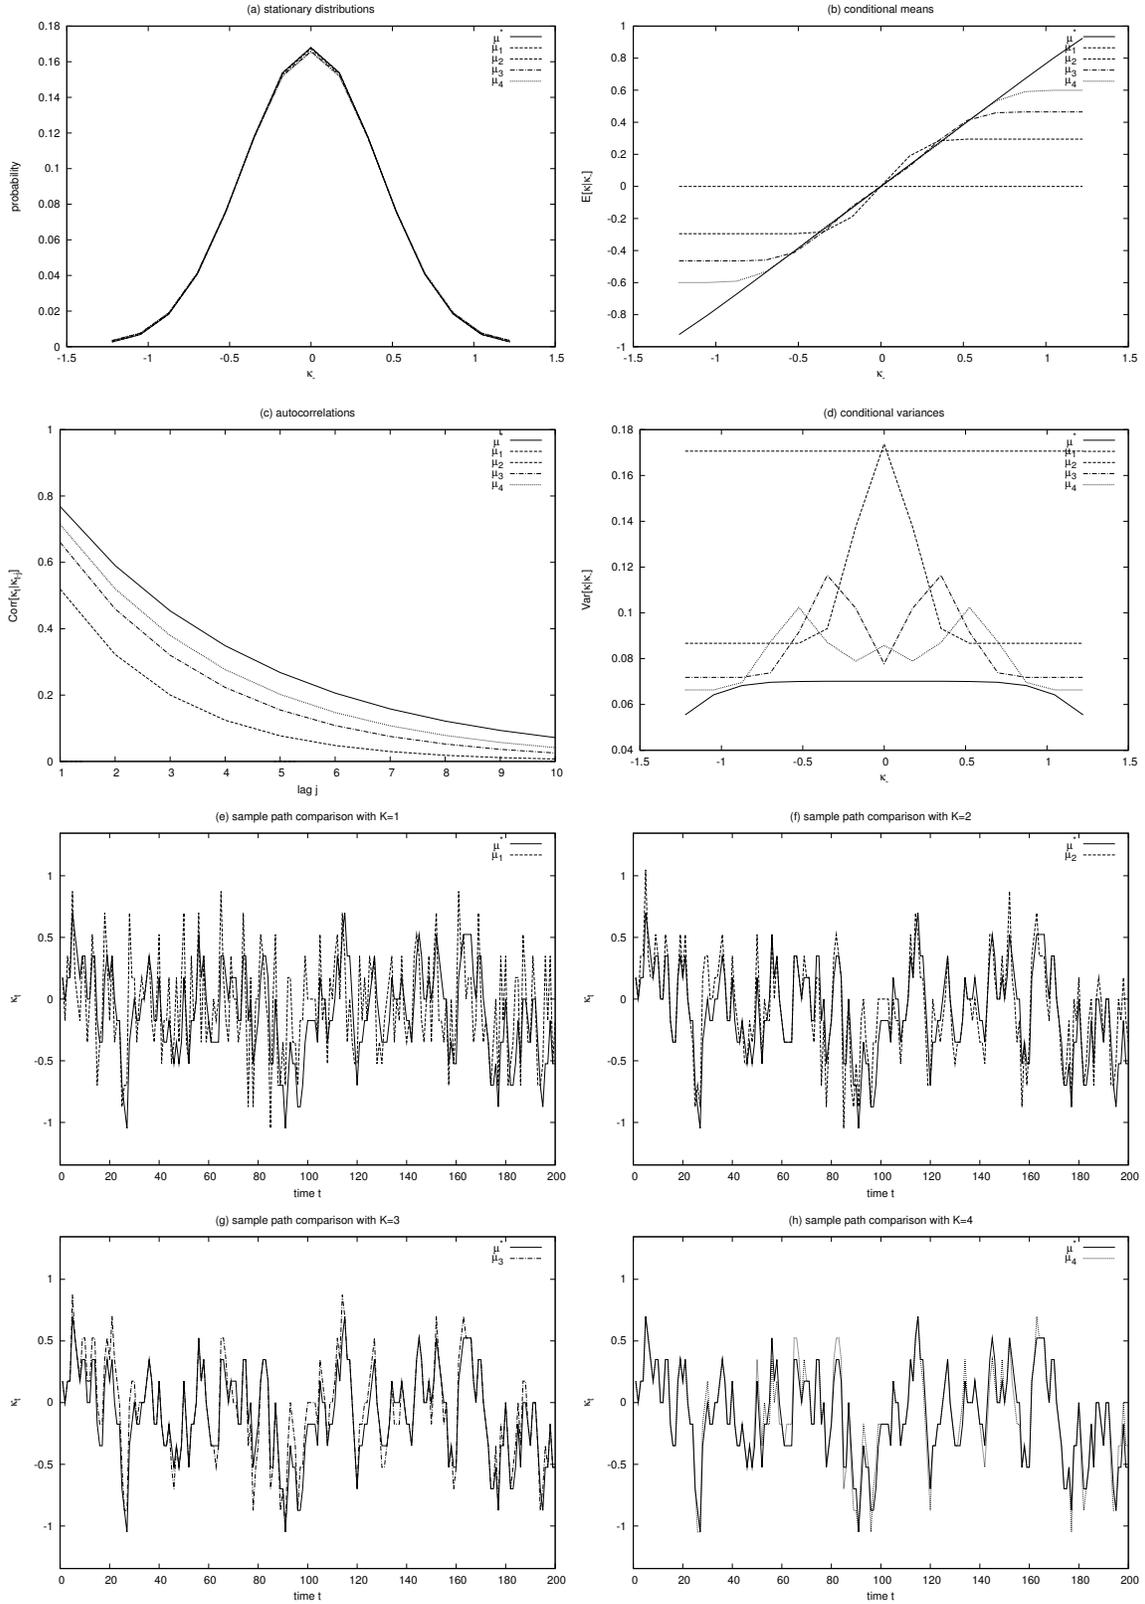


Figure 2: Statistical comparison of μ^* and μ_K for quinquennial model.

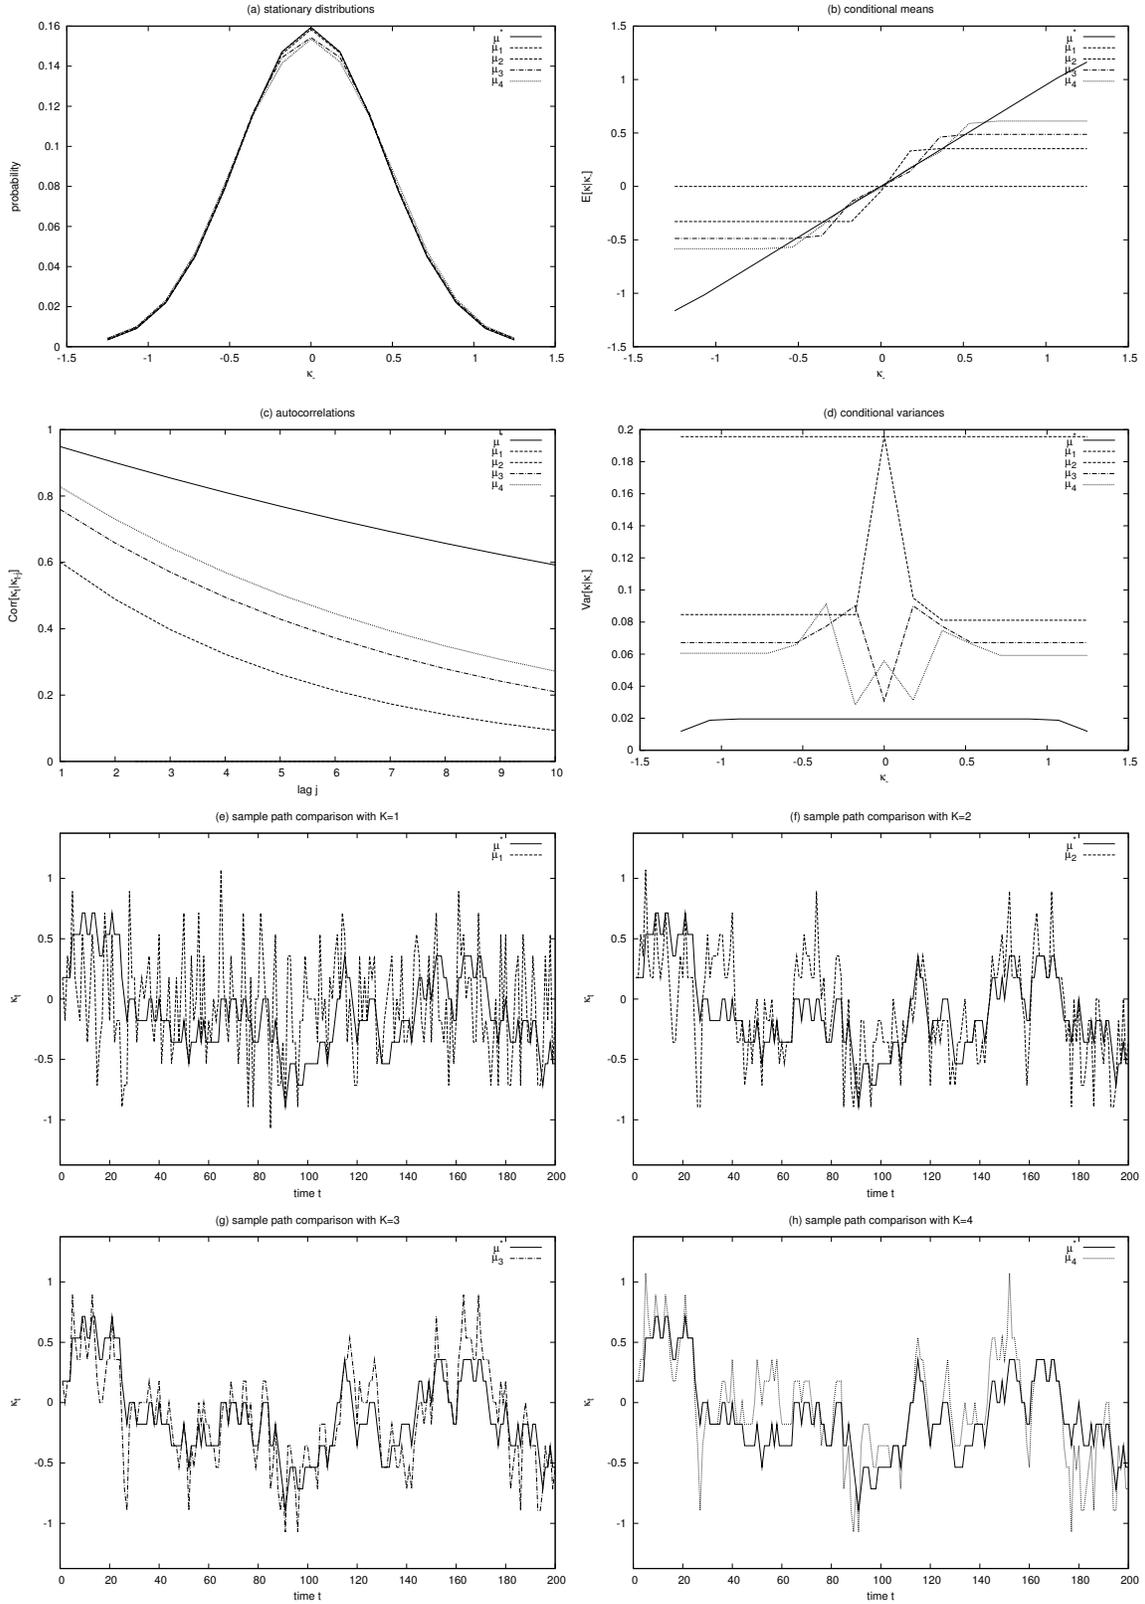


Figure 3: Statistical comparison of μ^* and μ_K for annual model.

representation of μ_2 . The figure reveals that in order to obtain high persistence, we need p_1 and p_2 to be sufficiently distinct and the graph of w to be sufficiently steep. But if we carry these properties to extremes, the stationary distribution will no longer have the unimodal form that $\bar{\mu}^*$ does. The insufficient persistence of μ_2 follows from this tension and the fact that the approximation method (16) tries to closely match the stationary distribution (cf. panel (a) of figures 2 and 3).

So far we have focused on the statistical properties of $\{\kappa_t\}$ under each μ_K and how they compare with those under μ^* . These comparisons provide information on how flexible each Π_K is in a statistical sense, that is, what kinds of persistence properties can be captured using processes with order K ($< N$) mixture representations.

A different but equally interesting question however is this: Suppose we knew μ^* is the “true” transition kernel for $\{\kappa_t\}$ but we computed the optimal mechanism under μ_K . Would the computed mechanism be “close” to the optimal mechanism under μ^* ? We unfortunately cannot provide a complete answer to this question as it requires solving the mechanism design problem under μ^* , which is a recursive problem with a 15 dimensional state space. We can, however, provide a partial answer by examining how close the solutions under μ^* and each μ_K are in a simplified version of the model with a short time horizon (which can be solved sequentially). We pursue this next.

We thus consider a two period version of the model which we interpret as standing for a person’s 40 year working career. Each period then stands for 20 years, so we set $\beta = q = 0.95^{20} \approx 0.36$, $\rho = 0.95^{20} \approx 0.36$, and $\sigma_\epsilon^2 = 0.13^2 \times \sum_{i=1}^{20} 0.95^{2(i-1)} \approx 0.39^2$. We normalize $U_0 = -1$ and set the agent’s absolute risk aversion coefficient to $\gamma = 0.5$; the implied relative risk aversion is about 1 on average for the range of consumption values that we observe. We abstract from the transitory shocks τ_t . We then solve the mechanism design problem for each $\mu \in \{\mu^*, \mu_1, \dots, \mu_4\}$ and compare the optimal transfers from the planner to the agent which we denote $z_t^*(\theta^t; \mu)$.

The first four columns of table 1 summarize our main findings. The first two rows report the maximum errors in each z_t^* :

$$\max_{\theta^t} |z_t^*(\theta^t; \mu_K) - z_t^*(\theta^t; \mu^*)|,$$

expressed as fractions of average income. It turns out that the biggest errors here occur with low probability, and as a result the average errors,

$$\sum_{\theta^t} |z_t^*(\theta^t; \mu_K) - z_t^*(\theta^t; \mu^*)| \Pr(\theta^t),$$

	$\rho = 0.95$ annually				$\rho = 0.99$ annually			
	$K = 1$	$K = 2$	$K = 3$	$K = 4$	$K = 1$	$K = 2$	$K = 3$	$K = 4$
max z_1 error	0.1632	0.0951	0.0358	0.0078	1.3299	1.1756	1.0355	0.8791
max z_2 error	0.3369	0.1894	0.0692	0.0151	1.4453	1.2396	0.9974	0.7730
avg z_1 error	0.0320	0.0041	0.0007	0.0003	0.1171	0.0444	0.0219	0.0150
avg z_2 error	0.0809	0.0104	0.0016	0.0008	0.2972	0.0874	0.0385	0.0267
cost error	0.0054	0.0009	0.0001	0.0000	0.0539	0.0124	0.0069	0.0039

Table 1: Comparison of optimal mechanisms under μ^* and μ_K .

reported in the next two rows as fractions of average income, are up to several orders of magnitude smaller. As we can see, these errors drop rapidly as we increase K and reach levels that appear small enough for many purposes at $K = 3$ or 4 . Similar properties hold for errors in the optimal cost:

$$\left| \sum_{t=1}^2 \sum_{\theta^t} q^t z_t^*(\theta^t; \mu_K) \Pr(\theta^t) - \sum_{t=1}^2 \sum_{\theta^t} q^t z_t^*(\theta^t; \mu^*) \Pr(\theta^t) \right|,$$

reported in the final row as fractions of average present value income (with q discounting). Overall, we find these results encouraging.

We observed previously that approximating μ^* by μ_K (with small K) appears challenging when κ is highly persistent. In the last four columns of table 1 we revisit this point by showing how the results change if we increase ρ from 0.95 to 0.99 in annual terms. As we can see, there is indeed a significant increase in the errors compared to the baseline case.

Finally, we tried solving the same problem using the first order approach. We found the approach to be valid for the $\rho = 0.95$ case. For the $\rho = 0.99$ case it was invalid, but the solution to the relaxed problem turned out to be close to the true solution (closer than with $K = 4$). These results support our tentative view that the first order approach may be more effective than ours when the hidden type is one dimensional and highly persistent.

6 Conclusion

At an abstract level, the essence of our finding is the observation that one can efficiently track conditional expectations over time by carefully choosing the timing convention if the exogenous forcing variables follow a Markov process with a low-order mixture representation. We have elaborated on how to exploit this fact for computational purposes in the context of dynamic mechanism design. This interpretation of our analysis suggests that a similar approach may prove useful in other contexts as a dimensionality reduction technique when

there is a need to track a conditional expectation as a state variable.

A Proofs

This section collects the proofs. Many of the arguments are standard but we include them for completeness.

A.1 Proof of Lemma 1

First suppose (3) did not hold for some t , θ^{t-1} , θ_t , θ'_t . Then if we define a reporting strategy $\mathbf{r} \in \mathbf{R}$ by $r_t(\theta^t) = \theta'_t$ and $r_s(\theta^s) = \theta_s$ for all $(s, \theta^s) \neq (t, \theta^t)$ we have

$$U(\mathbf{x}; \theta_{-1}) - U(\mathbf{x} \circ \mathbf{r}; \theta_{-1}) = \beta^t [(\text{L.H.S. of (3)}) - (\text{R.H.S. of (3)})] \Pr(\theta^t | \theta_{-1}) < 0,$$

which violates incentive compatibility.

Next suppose \mathbf{x} satisfies (3) and let $\mathbf{r} \in \mathbf{R}$ be an arbitrary reporting strategy. To show that $U(\mathbf{x}; \theta_{-1}) \geq U(\mathbf{x} \circ \mathbf{r}; \theta_{-1})$, let $W_t(\mathbf{x} \circ \mathbf{r}; \theta_{-1})$ denote the utility the agent gets from following \mathbf{r} for the first t periods and then reverting back to truth telling:

$$\begin{aligned} W_t(\mathbf{x} \circ \mathbf{r}; \theta_{-1}) &= \sum_{s=0}^t \sum_{\theta^s} \beta^s u(x_s(r^s(\theta^s)); \theta_s) \Pr(\theta^s | \theta_{-1}) \\ &\quad + \sum_{s=t+1}^{\infty} \sum_{\theta^s} \beta^s u(x_s(r^t(\theta^t), \theta_{t+1}^s); \theta_s) \Pr(\theta^s | \theta_{-1}). \end{aligned}$$

We claim that $U(\mathbf{x}; \theta_{-1}) \geq W_0(\mathbf{x} \circ \mathbf{r}; \theta_{-1}) \geq \dots \geq W_t(\mathbf{x} \circ \mathbf{r}; \theta_{-1})$ for any t and that $W_t(\mathbf{x} \circ \mathbf{r}; \theta_{-1}) \rightarrow U(\mathbf{x} \circ \mathbf{r}; \theta_{-1})$ as $t \rightarrow \infty$. The first statement follows from (3) and mathematical induction. The second statement follows from:

$$|U(\mathbf{x} \circ \mathbf{r}; \theta_{-1}) - W_t(\mathbf{x} \circ \mathbf{r}; \theta_{-1})| \leq \beta^{t+1} \times \text{length}(V) \rightarrow 0, \quad \text{as } t \rightarrow \infty.$$

The result follows.

A.2 Proof of Lemma 2

It is enough to show that an allocation \mathbf{x} satisfies the constraints in the auxiliary planning problem starting from (θ_{-1}, a_0) if and only if there is a sequence $\mathbf{a} = \{a_t\}_{t=0}^{\infty}$, $a_t : \Theta^t \rightarrow A^*$, such that (\mathbf{x}, \mathbf{a}) satisfies the constraints in the statement of the lemma.

First suppose \mathbf{x} satisfies the constraints in the auxiliary planning problem starting from (θ_{-1}, a_0) . Define \mathbf{a} by (5). Condition (10) then follows from the incentive compatibility of \mathbf{x} , Lemma 1, and (6). Condition (8) follows from (5) and (6). Finally, $a_{t+1}(\theta^t) \in A^*$ follows from the incentive compatibility of the continuation allocation $\mathbf{x}|_{\theta^t} := \{x_{t+j+1}(\theta^t, \cdot)\}_{j \geq 0}$ and $a_{t+1}(\theta^t) = a_0(\mathbf{x}|_{\theta^t})$.

Next suppose (\mathbf{x}, \mathbf{a}) satisfies the given conditions. From (8) it follows that

$$a_t(\theta^{t-1}) = \sum_{\theta_t} \{u(x_t(\theta^t); \theta_t) + \beta a_{t+1}(\theta^t) \cdot w(\theta_t)\} p(\theta_t).$$

Iterating forward on this and using the fact that $\{a_t\}_{t=0}^\infty$ is a bounded sequence, we can see that (\mathbf{x}, \mathbf{a}) satisfies (5). It follows from this and $a_0(\theta^{-1}) = a_0$ that \mathbf{x} satisfies (8). As well, (6), (10), and Lemma 1 together imply (7).

A.3 Proof of Proposition 3

The following lemmas are analogous to those in Abreu, Pearce, and Stacchetti (1990).

Lemma 5. *If $A \subset V^K$ satisfies $A \subset B(A)$, then $B(A) \subset A^*$.*

Proof. Suppose A satisfies the hypotheses, and let $a \in B(A)$. Using this, we can roll out an allocation \mathbf{x} as follows. First, for period 0, use $a \in B(A)$ to construct $(x_0(\cdot), a_1(\cdot)) \in F(a; A)$. Note that $a_1(\theta_0) \in A \subset B(A)$ for all θ_0 . Then, for periods $t \geq 1$ and given histories θ^{t-1} , proceed inductively by using $a_t(\theta^{t-1}) \in B(A)$ to construct $(x_t(\theta^{t-1}, \cdot), a_{t+1}(\theta^{t-1}, \cdot)) \in F(a_t(\theta^{t-1}); A)$.

To finish the proof, observe that (\mathbf{x}, \mathbf{a}) thus constructed satisfies conditions in Lemma 2, with A^* replaced by A . The second half of the proof goes through, which verifies that \mathbf{x} is incentive compatible and satisfies $a = a_0(\mathbf{x})$. It follows that $a \in A^*$. \square

Lemma 6. $B(A^*) = A^*$.

Proof. Given Lemma 5, it is enough to show that $A^* \subset B(A^*)$. So let $a \in A^*$ and let \mathbf{x} be an incentive compatible allocation satisfying $a = a_0(\mathbf{x})$. Define $a_1 : \Theta \rightarrow A$ by $a_1(\theta_0) = a_0(\mathbf{x}|_{\theta_0})$, all θ_0 . It is then easy to see that $(x_0(\cdot), a_1(\cdot)) \in F(a; A^*)$, implying $a \in B(A^*)$. \square

Lemma 7. *If $A \subset A' \subset V^K$, then $B(A) \subset B(A')$.*

Proof. Immediate from the definition of B . \square

Lemma 8. *If A is compact, so is $B(A)$.*

Proof. Suppose A is compact. Clearly $B(A) \subset V^K$ is bounded. To see that it is closed, pick a convergent sequence $\{a_n\}_{n=1}^\infty \subset B(A)$ and let a denote its limit. Then for each n , there exists $(x_n, a_n^+) \in F(a_n; A)$. Because $\{x_n, a_n^+\}_{n=1}^\infty$ can be viewed as sequence of finite-dimensional vectors in the compact set $X^N \times A^N$, it has a convergent subsequence whose limit we denote by (x, a^+) . Using continuity and the closedness of A , we can see that $(x, a^+) \in F(a; A)$, implying $a \in B(A)$. \square

Part 1. To see that A^* is non-empty, consider a constant allocation $\bar{\mathbf{x}}$ which repeats $\bar{x} \in X$. Its constancy implies incentive compatibility, and $a_0(\bar{\mathbf{x}}) \in V^K$. Hence $a_0(\bar{\mathbf{x}}) \in A^*$.

Next, we verify the compactness of A^* . Boundedness follows from $A^* \subset V^K$. To prove closedness, note that because $\text{cl}(A^*)$ is compact, so is $B(\text{cl}(A^*))$ by Lemma 8. As well, $A^* = B(A^*) \subset B(\text{cl}(A^*))$ by Lemmas 6 and 7. Combining, we have $\text{cl}(A^*) \subset B(\text{cl}(A^*))$. Lemma 5 then implies $\text{cl}(A^*) \subset A^*$, which proves the claim.

Part 2. Suppose $V^K \supset A_0 \supset B(A_0) \supset A^*$ and let $A_n = B^n(A_0)$ for each $n = 1, 2, \dots$. Using Lemmas 7 and 8, we can see that $A_n \downarrow \cap_{n=0}^\infty A_n =: A_\infty$ and that each A_n as well as A_∞ is compact. By Lemmas 6 and 7, $A^* \subset A_\infty$. We show $A_\infty \subset A^*$ by verifying $A_\infty \subset B(A_\infty)$ (cf. Lemma 5). So let $a \in A_\infty$. Then because $A_\infty \subset B(A_n) \subset V^K$ for each n , we can construct a sequence of function pairs $\{x_n, a_n^+\}_{n=0}^\infty$ such that for each n there holds $(x_n, a_n^+) \in F(a; A_n)$. As in the proof of Lemma 8, we can pick a convergent subsequence and let (x, a^+) denote the limit. Because each $a_n^+(\theta) \in A_n$ and $\{A_n\}_{n=0}^\infty$ is a sequence of compact sets converging down to the compact set A_∞ , we have that $a^+(\theta) \in A_\infty$ for each θ . From this and continuity it follows that $(x, a^+) \in F(a; A_\infty)$. Hence $a \in B(A_\infty)$.

Now consider setting $A_0 = V^K$. Compactness of A_0 is evident. The set inclusion $A^* \subset B(A_0)$ follows from $A^* \subset A_0$ and Lemmas 6 and 7. To see that $B(A_0) \subset A_0$, let $a \in B(A_0)$ and let $(x, a^+) \in F(a; A_0)$. We then have

$$a = \sum_{\theta} \{u(x(\theta); \theta) + \beta a^+(\theta) \cdot w(\theta)\} p(\theta) \in V^K = A_0.$$

Part 3. Suppose $A_0 \subset B(A_0) \subset A^*$. By Lemmas 6 and 7 and compactness of A^* , $B^n(A_0)$ is increasing in n and satisfies $\text{cl}(\cup_{n=0}^\infty B^n(A_0)) \subset A^*$.

To prove $A^* \subset \text{cl}(\cup_{n=0}^\infty B^n(A_0))$, pick any $a \in A^*$. We will construct a sequence in $\cup_{n=0}^\infty B^n(A_0)$ that converges to a . For this, let $a' \in A_0 (\subset A^*)$. By the definition of A^* there exist incentive compatible allocations \mathbf{x} and \mathbf{x}' such that $a = a_0(\mathbf{x})$ and $a' = a_0(\mathbf{x}')$. Then for each $n \geq 1$, do the following. First let $\mathbf{x}^n = \{x_t^n\}_{t=0}^\infty$ be an allocation constructed by truncating \mathbf{x} after n periods and then appending \mathbf{x}' . Thus for $t > n$:

$$(x_0^n(\theta^0), \dots, x_t^n(\theta^t)) = (x_0(\theta^0), \dots, x_n(\theta^n), x_0'(\theta_{n+1}^{n+1}), \dots, x_{t-n-1}'(\theta_{n+1}^t)). \quad (17)$$

Next let $\mathbf{r}^n = \{r_t^n\}_{t=0}^\infty$ be an optimal reporting strategy for the agent given \mathbf{x}^n (i.e., one that maximizes $U(\mathbf{x}^n \circ \mathbf{r}^n; \theta_{-1})$). Note that we can take r_t^n for $t > n$ to be truth-telling, thanks to the incentive compatibility of \mathbf{x}^n and (17). Then let $\hat{\mathbf{x}}^n = \mathbf{x}^n \circ \mathbf{r}^n$. By construction, $\hat{\mathbf{x}}^n$ is incentive compatible, $a_0(\hat{\mathbf{x}}^n) \geq a_0(\mathbf{x}^n)$, and $a_{n+1}(\theta^n; \hat{\mathbf{x}}^n) = a'$ for all θ^n .

We next verify $a_0(\hat{\mathbf{x}}^n) \in \cup_{n=0}^\infty B^n(A_0)$ for all n . For this, note that for each t and θ^{t-1} :

$$a_t(\theta^{t-1}; \hat{\mathbf{x}}^n) = \sum_{\theta_t} \{u(\hat{x}_t^n(\theta^t); \theta_t) + \beta a_{t+1}(\theta^t; \hat{\mathbf{x}}^n) \cdot w(\theta_t)\} p(\theta_t), \quad (18)$$

and by the incentive compatibility of $\hat{\mathbf{x}}^n$:

$$\begin{aligned} & u(\hat{x}_t^n(\theta^t); \theta_t) + \beta a_{t+1}(\theta^t; \hat{\mathbf{x}}^n) \cdot w(\theta_t) \\ & \geq u(\hat{x}_t^n(\theta^{t-1}, \theta'_t); \theta_t) + \beta a_{t+1}(\theta^{t-1}, \theta'_t; \hat{\mathbf{x}}^n) \cdot w(\theta_t), \quad \forall (\theta_t, \theta'_t) \in \Theta \times \Theta. \end{aligned} \quad (19)$$

Using (18) and (19) at $t = n$ and $a_{n+1}(\theta^n; \hat{\mathbf{x}}^n) = a'$, we obtain $a_n(\theta^{n-1}; \hat{\mathbf{x}}^n) \in B(\{a'\})$ for all θ^{n-1} . From here we can use induction on (18) and (19) for $t = n - 1, \dots, 0$ to get $a_0(\hat{\mathbf{x}}^n) \in B^{n+1}(\{a'\})$. Lemma 7 and the fact that $B^n(A_0)$ is increasing in n imply $B^{n+1}(\{a'\}) \subset B^{n+1}(A_0) \subset \cup_{n=0}^\infty B^n(A_0)$.

To verify that $a_0(\hat{\mathbf{x}}^n) \rightarrow a$ as $n \rightarrow \infty$, we pick an arbitrary subsequence $\{a_0(\hat{\mathbf{x}}^{n'})\}_{n'=1}^\infty$ and show that it has a further subsequence $\{a_0(\hat{\mathbf{x}}^{n''})\}_{n''=1}^\infty$ that converges to a . Toward this end, note that because each $r_t^{n'}$ belongs to a finite set (being a mapping from a finite set to a finite set) there is subindex n'' along which $\mathbf{r}^{n''}$ converges to some $\mathbf{r} = \{r_t\}_{t=0}^\infty$ in the sense that, for all t , $r_t^{n''} = r_t$ for large n'' . Also for each t we have $x_t^{n''} = x_t$ for $n'' \geq t$. This together with the boundedness of u implies $a_0(\hat{\mathbf{x}}^{n''}) = a_0(\mathbf{x}^{n''} \circ \mathbf{r}^{n''}) \rightarrow a_0(\mathbf{x} \circ \mathbf{r})$. Combining this with $a_0(\hat{\mathbf{x}}^{n''}) \geq a_0(\mathbf{x}^{n''})$ and $a_0(\mathbf{x}^{n''}) \rightarrow a$, we obtain $a_0(\mathbf{x} \circ \mathbf{r}) \geq a$. But the incentive compatibility of \mathbf{x} implies $a_0(\mathbf{x} \circ \mathbf{r}) \leq a_0(\mathbf{x}) = a$, so $a_0(\mathbf{x} \circ \mathbf{r}) = a$. The conclusion follows.

Now let $\bar{\mathbf{x}}$ be a constant allocation which repeats $\bar{x} \in X$ and let $A_0 = \{a_0(\bar{\mathbf{x}})\}$. We have $A_0 \subset A^*$ from the incentive compatibility of $\bar{\mathbf{x}}$. From this and Lemmas 6 and 7 we get $B(A_0) \subset A^*$. To see that $A_0 \subset B(A_0)$, note that the constant function pair $(x, a^+) \equiv (\bar{x}, a_0(\bar{\mathbf{x}}))$ satisfies $(x, a^+) \in F(a_0(\bar{\mathbf{x}}); A_0)$ thanks to constancy and

$$a_0(\bar{\mathbf{x}}) = \sum_{\theta} \{u(\bar{x}; \theta) + \beta a_0(\bar{\mathbf{x}}) \cdot w(\theta)\} p(\theta).$$

Part 4. If the environment is convex, B maps convex sets into convex sets. Convexity of A^* then follows from the convexity of V^K and $B^n(V^K) \downarrow A^*$, as verified above.

A.4 Proof of Proposition 4

From Lemma 2 we know that the auxiliary planning problem is equivalent to a standard dynamic programming problem. Standard arguments then imply that J^* is a fixed point of T , and that if $g^* : \Theta \times A^* \rightarrow (X \times A^*)^\Theta$ attains the infimum on the right hand side of (13) when $J = J^*$ then an allocation \mathbf{x} defined recursively by $(x_t^*(\theta^t), a_{t+1}^*(\theta^t)) = g^*(\theta_{t-1}, a_t^*(\theta^{t-1}))(\theta_t)$ solves the auxiliary planning problem (cf. Propositions 9.8 and 9.12 of Bertsekas and Shreve (1996)).

Boundedness of J^* follows from $\min c(X)/(1-q) \leq J^* \leq \max c(X)/(1-q)$. It follows from Blackwell's theorem that T is a monotone contraction on the space of bounded functions $J : \Theta \times A^* \rightarrow \mathbb{R}$. Thus $\|T^n J - J^*\| \rightarrow 0$ for any such J .

We go on to prove that J^* is lower semicontinuous and that the function g^* exists. First identify the set of functions $(X \times A^*)^\Theta$ with $X^N \times A^{*N}$, and note that $F(\cdot; A^*) : A^* \rightrightarrows X^N \times A^{*N}$ is nonempty-valued (by $A^* = B(A^*)$), compact-valued (by the continuity of the constraints and the compactness of $X^N \times A^{*N}$), and upper hemicontinuous (by the continuity of the constraints). Hence, if J is lower semicontinuous, so is TJ (cf. Lemma 17.30 of Aliprantis and Border (2006)). Now consider the constant function $J_* \equiv \min c(X)/(1-q)$. Then by the definition of T , $TJ_* \geq J_*$. Because T is monotone and $\|T^n J_* - J^*\| \rightarrow 0$, it follows that J^* is the pointwise supremum of $\{T^n J_*\}_{n=1}^\infty$. Since each $T^n J_*$ is lower semicontinuous, it follows that J^* is lower semicontinuous. The existence of g^* follows from this and the fact that each $F(a; A^*)$ is non-empty and compact.

Finally, suppose the environment is convex. Fix θ_{-1} and pick $a_0^{(i)} \in A^*$ for $i \in \{1, 2\}$ and $\lambda \in (0, 1)$. From the above, we can construct $\mathbf{x}^{*(i)}$ that solves the auxiliary planning problem starting from $(\theta_{-1}, a_0^{(i)})$, $i \in \{1, 2\}$. Since $\lambda \mathbf{x}^{*(1)} + (1-\lambda)\mathbf{x}^{*(2)}$ is feasible in the problem starting from $(\theta_{-1}, \lambda a_0^{(1)} + (1-\lambda)a_0^{(2)})$, it follows that

$$\begin{aligned} J^*(\theta_{-1}, \lambda a_0^{(1)} + (1-\lambda)a_0^{(2)}) &\leq C(\lambda \mathbf{x}^{*(1)} + (1-\lambda)\mathbf{x}^{*(2)}; \theta_{-1}) \\ &\leq \lambda C(\mathbf{x}^{*(1)}; \theta_{-1}) + (1-\lambda)C(\mathbf{x}^{*(2)}; \theta_{-1}) \\ &= \lambda J^*(\theta_{-1}, a_0^{(1)}) + (1-\lambda)J^*(\theta_{-1}, a_0^{(2)}). \end{aligned}$$

Hence $J^*(\theta_{-1}, \cdot)$ is convex.

References

Abraham, A., and N. Pavoni (2008): "Efficient Allocations with Moral Hazard and Hidden Borrowing and Lending: A Recursive Formulation," *Review of Economic Dynamics*, 11(4),

781–803.

- Abreu, D., D. Pearce, and E. Stacchetti (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58(5), 1041–1063.
- Aliprantis, C. D., and K. C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer-Verlag, Berlin.
- Atkeson, A., and R. E. Lucas, Jr. (1992): “On Efficient Distribution with Private Information,” *Review of Economic Studies*, 59(3), 427–453.
- Barber, C. B., D. P. Dobkin, and H. Huhdanpaa (1996): “The Quickhull Algorithm for Convex Hulls,” *ACM Transactions on Mathematical Software*, 22(4), 469–483.
- Bertsekas, D. P., and S. E. Shreve (1996): *Stochastic Optimal Control: The Discrete-Time Case*. Athena Scientific, Belmont, MA.
- Courty, P., and H. Li (2000): “Sequential Screening,” *Review of Economic Studies*, 67(4), 697–717.
- Doepke, M., and R. M. Townsend (2006): “Dynamic Mechanism Design with Hidden Income and Hidden Actions,” *Journal of Economic Theory*, 126(1), 235–285.
- Farhi, E., and I. Werning (2010): “Insurance and Taxation Over the Life Cycle,” Discussion paper, Harvard and MIT.
- Fernandes, A., and C. Phelan (2000): “A Recursive Formulation for Repeated Agency with History Dependence,” *Journal of Economic Theory*, 91(2), 223–247.
- Fukushima, K. (2010): “Quantifying the Welfare Gains From Flexible Dynamic Income Tax Systems,” Working paper.
- Fukushima, K., and Y. Waki (2011a): “Computing Dynamic Optimal Mechanisms When Hidden Types Are Markov and Controlled by Hidden Actions,” Working paper.
- (2011b): “A Polyhedral Approximation Approach to Concave Numerical Dynamic Programming,” Working paper.
- Green, E. J. (1987): “Lending and the Smoothing of Uninsurable Income,” in *Contractual Arrangements in Intertemporal Trade*, ed. by E. C. Prescott, and N. Wallace, pp. 3–25. University of Minnesota Press, Minneapolis, MN.
- Judd, K. L. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

- Judd, K. L., S. Yeltekin, and J. Conklin (2003): “Computing Supergame Equilibria,” *Econometrica*, 71(4), 1239–1254.
- Kapicka, M. (2010): “Efficient Allocations in Dynamic Private Information Economies with Persistent Shocks: A First Order Approach,” Discussion paper, University of California, Santa Barbara.
- Kocherlakota, N. R. (2010): *The New Dynamic Public Finance*. Princeton University Press.
- Mirrlees, J. A. (1971): “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38(2), 175–208.
- Pavan, A., I. Segal, and J. Toikka (2009): “Dynamic Mechanism Design: Incentive Compatibility, Profit Maximization and Information Disclosure,” Working paper, Northwestern University and Stanford University.
- Prescott, E. C., and R. M. Townsend (1984): “Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard,” *Econometrica*, 52(1), 21–46.
- Radner, R. (1981): “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, 49(5), 1127–1148.
- Smith, R. L. (1984): “Efficient Monte Carlo Procedures for Generating Points Uniformly Distributed Over Bounded Regions,” *Operations Research*, 32(6), 1296–1308.
- Spear, S. E., and S. Srivastava (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54(4), 599–617.
- Storesletten, K., C. I. Telmer, and A. Yaron (2004): “Cyclical Dynamics in Idiosyncratic Labor Market Risk,” *Journal of Political Economy*, 112(3), 695–717.
- Tauchen, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20(2), 177–181.
- Thomas, J., and T. Worrall (1990): “Income Fluctuation and Asymmetric Information: An Example of a Repeated Principal-Agent Problem,” *Journal of Economic Theory*, 51(2), 367–390.
- Townsend, R. M. (1982): “Optimal Multiperiod Contracts and the Gain from Enduring Relationships under Private Information,” *Journal of Political Economy*, 90(6), 1166–1186.
- Waki, Y. (2011): “Efficient Risk Sharing under Endogenous Skill Formation: Quantitative Implications,” Working paper.

Williams, N. (2011): “Persistent Private Information,” *Econometrica*, 79(4), 1233–1275.

Zhang, Y. (2009): “Dynamic Contracting with Persistent Shocks,” *Journal of Economic Theory*, 144(2), 635–675.