The Dual Approach to Recursive Optimization: Theory and Examples*

Matthias Messner[†] Nicola Pavoni[‡] Christopher Sleet[§]

September 15, 2013

Abstract

We bring together the theories of duality and dynamic programming. We show that the dual of a separable dynamic optimization problem can be recursively decomposed. We provide a dual version of the principle of optimality and give conditions under which the dual Bellman operator is a contraction with the optimal dual value function its unique fixed point. We relate primal and dual problems, address computational issues and give examples.

JEL codes: C61, C73, D82, D86, E61.

Keywords: Dynamic Contracts, Duality, Dynamic Programming.

1 Introduction

Many dynamic economic optimization problems have a recursive structure that makes them amenable to solution via dynamic programming. This structure allows the original problem to be decomposed into a family of simpler sub-problems linked by state variables. The set of state variables consistent with a non-empty constraint correspondence is called

*We thank Musab Kurnaz for expert assistance with the calculations. We are grateful for the comments of seminar participants at Concordia University, Mannheim University, the Bank of Portugal, the 2013 SAET meetings. This paper supersedes Messner, Pavoni, Sleet (2011), which developed the dual recursive method in a simpler two period setting, and Messner, Pavoni, Sleet (2012a), which constitutes our first attempt to derive contractive properties of the dual Bellman operator we introduce here. Pavoni gratefully acknowledges financial support from the European Research Council, Starting Grant #210908.

[†]Bocconi University and IGIER, 20136 Milano, Italy; matthias.messner@unibocconi.it.

[‡]Department of Economics, Bocconi University and IGIER, Via Roentgen 1, I-20136, Milan, Italy; IFS and CEPR, London; pavoni.nicola@gmail.com.

[§]Tepper School of Business, Carnegie Mellon University, Pittsburgh PA 15217; csleet@andrew.cmu.edu.

the "effective" state space and is a key component of a problem's recursive formulation. In many settings the effective state space is not given explicitly: it must be recovered as part of the solution to the problem. This complicates the application of dynamic programming methods and, following Marcet and Marimon (2011), has prompted economists to adopt recursive formulations that replace or supplement standard "primal" state variables with "dual" ones. Examples include, inter alia, Kehoe and Perri (2002), Marimon and Quadrini (2006), Acemoğlu, Golosov, and Tsyvinski (2010), Chien, Cole, and Lustig (2011) and Aiyagari, Marcet, Sargent, and Seppälä (2002). Despite their widespread use, thorough analysis of these methods is limited and their application has often been ad hoc. This paper develops a new recursive dual approach to dynamic optimization that blends elements of the theories of duality and dynamic programming. It shows that (i) a large class of dynamic optimization problems in economics have recursive duals, (ii) such recursive duals relocate the analysis to a more convenient dual state space that is often easy to characterize and (iii) the associated dual Bellman operator is contractive on an appropriate function space. Sufficient conditions for the dual and, hence, the recursive dual to characterize the solution of the original (primal) problem are given. For situations in which these sufficient conditions are not satisfied, a numerical check of optimality is proposed. Numerical implementation of the recursive dual method is discussed and various economic examples and applications provided.

The paper begins with a family of recursive (primal) optimizations that encompasses many economic applications. These optimizations feature objective and constraint functions that can be expressed in terms of recursively-evolving "summaries" of past and future actions. In the context of particular applications, such summaries have interpretations as capital, utility promises or inflation; they may be backward-looking (i.e. functions of past actions and shocks and an initial condition) or forward-looking (i.e. functions of future actions). In recursive formulations of primal problems, they serve as state variables. Primal optimization problems may be re-stated using a Lagrangian. In the re-stated problem a sup-inf operation over choices and Lagrange multipliers replaces a sup operation over choices alone. By interchanging the sup and inf operations a dual inf-sup problem is obtained. We show that if the correct Lagrangian is chosen, the recursive structure of the primal is inherited by the dual with, in the latter case, co-states (i.e. multipliers on laws of motion for primal states) serving as dual state variables.

We use this structure to recover a dual Bellman operator. The dual Bellman updates candidate value functions via "inf-sup" operations over Lagrange multipliers and actions. Specifically, at each dual state and current multiplier combination, an "inner" supremum operation is performed over current actions. Then, at each current dual state, an outer infimum operation over multipliers gives the updated value function. We show that *without further assumptions* the dual Bellman gives necessary conditions for optimal dual values and policies and, under mild additional restrictions, sufficient conditions for such values and policies. In short, we recover a dual principle of optimality. The key step in the derivation is an interchange of an infimum operation over future multipliers with a supremum operation over current actions. To ensure this interchange does not modify values or solutions (in the absence of further assumptions), it is essential to associate a Lagrangian with the problem that is rich enough to allow all non-linearities in constraints and in the objective to be contained in the Lagrangian's "current" terms. In general this requires explicitly incorporating laws of motion for primal state variables into the Lagrangian.

An attractive aspect of the recursive dual is that in some important cases, in particular, when the primal state space is bounded, the "effective" dual state space is readily identified as all of $\mathbb{R}^N$, where N is the number of dual state variables. Thus, for the dual problem, the difficulty of determining the state space is resolved. In addition, dual value functions are positively homogenous of degree one. Consequently, in calculations, the dual state space may be identified with the unit circle (in $\mathbb{R}^N$).

The recursive dual features an unbounded value function and an unbounded constraint correspondence. This combination creates a challenge for the standard approach to establishing contractivity of the Bellman operator. For problems with unbounded value functions, a common procedure following Wessels (1977), is to show that there is a set of functions closed and bounded with respect to a more permissive weighted sup norm[1] that contains the optimal value function and on which the Bellman is a contractive self-map. However, this approach requires that the continuation state variables and, hence, the continuation value function cannot vary "too much" on the graph of the constraint correspondence. Since the dual Bellman operator permits the choice of multipliers from an unbounded set, this condition is only guaranteed in the dual setting if additional non-binding constraints on multipliers are found. Instead, we show that the Bellman is contractive with respect to an alternative metric on a space of functions sandwiched between two (unbounded) functions.[2] We show through examples that such bounding functions are often available. A further difficulty is that the unboundedness and, hence, non-compactness of the set of feasible multipliers disrupts the application of the Theorem of the Maximum. However, it is easy to show that the optimal value function is convex. When it is every-

[1] A weighted sup norm on a set of functions $\mathscr{F}$ with common domain X is a function $\|\cdot\|_w : \mathscr{F} \to \mathbb{R}$ of the form $\|f\|_w = \sup_X \left|\frac{f(x)}{w(x)}\right|$ for some $w : X \to \mathbb{R}_{++}$.

[2] The argument combines the concavity of the dual Bellman, properties of the metric and of the sandwich. It adapts ideas of Rinćon-Zapatero and Rodríguez-Palmero (2003). The novelty lies in the application of this argument to the dual setting to which it seems very suited.

where real-valued as well, appeals can be made to the continuity properties of convex functions to establish continuity of the optimal value function.

The recursive dual formulation permits solution of the dual rather than the original primal problem. It remains to relate them. Weak duality results imply that the dual and, hence, the recursive dual supplies an upper bound for payoffs from the primal problem. Consequently, with no further assumptions the recursive dual gives welfare bounds for optimal policies or policy improvements. For concave problems, possibly after relaxation of the equality constraints describing laws of motion for state variables, we may appeal directly to known duality results to relate the dual (and, hence, again the recursive dual) more tightly to the original primal. These results give sufficient conditions on primitives for dual and primal values and, sometimes, solutions to coincide. When theoretical sufficient conditions for equality of dual and primal values and solutions are not available, because, for example of non-concavities, we propose a numerical procedure for checking whether a dual solution solves the original primal problem.

The paper proceeds as follows. After a brief literature review, Section 2 introduces a general class of stochastic, infinite-horizon problems. Economic examples are given in Section 3. Section 4 presents a primal recursive formulation for a sub-class of these problems and points out difficulties in applying it. In Section 5, the primal problem is paired with a dual problem and a recursive formulation of the latter obtained. A Bellman-type principle of optimality for the dual problem is established; Section 6 gives a contraction result for recursive dual problems. The important class of problems with laws of motion and constraints that are quasi-linear in (primal) states is considered in Section 7. Section 8 relates primal and dual problems. Numerical implementation is discussed and a numerical example given in Section 9.

Literature Our method is related to, but distinct from, that of Marcet and Marimon (1999) (revised: Marcet and Marimon (2011)). These authors propose solving dynamic optimizations by recursively decomposing a saddle point operation. They restrict attention to concave problems with constraints (including laws of motion) that are linear in forward-looking state variables. They substitute forward-looking states out of the problem using their laws of motion and absorb a subset of constraints into a Lagrangian. Laws of motion for backward-looking primal states (e.g. capital) are left as explicit restrictions. They then recursively decompose a saddle point of this Lagrangian (on the constraint set defined by the backward-looking laws of motion).

In contrast, our approach cleanly separates dualization of the primal from recursive decomposition of the dual and shows that the latter is available under rather weak separability conditions, much weaker than those imposed by Marcet and Marimon (2011). Our *theoretical* sufficient conditions for equality of optimal dual and primal values and solutions are stronger than those guaranteeing recursive decomposition. However, even here we can dispense with several of Marcet-Marimon's restrictions. The requirements that constraints are linear in forward-looking state variables and that *every* continuation problem has a saddle can be dropped. Moreover, when these theoretical conditions are not satisfied, we propose a numerical procedure for checking primal optimality of a dual solution.

For some problems, Marcet and Marimon (2011)'s recursive saddle Bellman operator is available and resembles our dual Bellman.[3] In others, it is not available or is available, but is quite different from ours. In particular, all of the examples considered in this paper either cannot be handled by Marcet and Marimon (2011)'s formulation or would be handled differently. The difference in the handling of backwards-looking state variables across our approach and that of Marcet and Marimon (2011) is not a detail. Our treatment of these variables is essential for the contractivity of the dual Bellman. This result relies on the concavity of the Bellman operator, which is always true for our formulation, but not theirs.[4]

Messner, Pavoni, and Sleet (2012b) consider the relationship between primal and dual Bellman operators. They restrict attention to concave problems without backward-looking state variables and with laws of motion that are linear in forward-looking ones. Thus, their setting is much less general than the present one; it excludes many economically relevant problems such as default with capital accumulation, risk sharing with non-expected utility and optimal monetary policy, all of which are considered here. In addition, they do not provide contraction results or a numerical implementation. In a similar setting to Messner, Pavoni, and Sleet (2012b), Cole and Kubler (2012) show how recursive methods using dual variables may be extended to give sufficient conditions for an optimal primal solution under weak concavity conditions. In addition, they derive a contraction result using a weighted sup-norm. They do so by obtaining additional non-binding constraints on multipliers and, hence, continuation states. However, the restrictions on primitives for these additional constraints to be non-binding appear strong.

[3]However, even in these cases, our Bellman operator implements a fairly straightforward inf-sup operation, whereas theirs involves a more difficult saddle point operation.

[4]Underpinning this is the fact that our dual formulation relies entirely on dual state variables; Marcet and Marimon (1999) dualize a subset of constraints and rely on a mixture of dual and primal state variables.

2 Decision Maker's Problem

This section describes an abstract recursive choice problem that can be specialized to give many problems considered in the literature. In particular, it encompasses many dynamic contracting and optimal policy problems. Concrete examples are given in Section 3.

Shocks and Action Plans Let $\mathcal{S} = \{1, \ldots, n_s\}$, with element s, denote a finite set of shocks.[5] Shock histories of length $t = 1, \ldots, \infty$ are denoted $s^t \in \mathcal{S}^t$. Let $\mathcal{A} \subset \mathbb{R}^{n_a}$, with element a, denote a set of actions available to a decision-maker. The decision-maker's action choices at each history are collected into an *action plan*: $\alpha = \{a_t\}_{t=0}^{\infty}$, with $a_0 \in \mathcal{A}$ and, $\forall t \in \mathbb{N}$, $a_t : \mathcal{S}^t \to \mathcal{A}$. The s^t-continuation of an action plan α is denoted $\alpha|s^t = \{a_{t+\tau}(s^t, \cdot)\}_{\tau=0}^{\infty}$. Plans are restricted to a set $\mathscr{A} \subset 2^{\mathcal{A}^\infty}$ such that if $\alpha \in \mathscr{A}$, then for all t and s^t, $\alpha|s^t \in \mathscr{A}$. Let $\mathcal{R}(\mathcal{S})$ denote the set of probability distributions on $\mathcal{S}$ and $Q : \mathcal{S} \times \mathcal{A} \to \mathcal{R}(\mathcal{S})$ a transition that maps current shock-action pairs to probability distributions over the subsequent period's shocks. Together Q, a seed shock s_0 and an action plan α induce a probability distribution over shocks and actions in all periods.

Constraints The set $\mathscr{A}$ is supplemented with additional constraints involving explicit functions of actions. These functions depend on recursively evolving "summaries" of past and future actions. Such summaries serve as states in primal recursive formulations, where they often have concrete economic interpretations as, inter alia, capital stocks, utility promises or inflation rates. In the dual setting multipliers on the laws of motion for these summaries will serve as states.

We distinguish between summaries of past and future actions. Let $\mathcal{K} \subset \mathbb{R}^{n_k}$ be a bounded set. Given a plan α, summaries of *past* actions and shocks $K_{t+1}(\alpha, s^t)$ are constructed recursively from a function $W^K : \mathcal{K} \times \mathcal{S} \times \mathcal{A} \to \mathcal{K}$ according to:

$$K_{t+1}(\alpha, s^t) = W^K[K_t(\alpha, s^{t-1}), s_t, a_t(s^t)], \tag{1}$$

with $K_0(\alpha, s^{-1}) = \bar{k}$ an initial seed state. In the sequel, we call summaries of past actions and shocks *backward-looking* state variables. In many economic models physical or human capital are naturally formalized as a backward-looking state variables.

Summaries of *future* actions $V(s_t, \alpha|s^t) \in \mathbb{R}^{n_v+1}$ are constructed recursively from a pair of functions $W^V : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{n_v+1} \to \mathbb{R}^{n_v+1}$ and $M^V : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{n_s(n_v+1)} \to \mathbb{R}^{n_v+1}$. The first is a time aggregator that gives the current summary as a function of current actions

[5] The restriction to a finite set of shocks streamlines our presentation by avoiding measure-theoretic complications, but is not essential for our main results.

and a certainty equivalent of future summaries; the second is a stochastic aggregator that generates the certainty equivalent. Future summaries are given by a function $V : \mathcal{S} \times \mathscr{A} \rightarrow \mathbb{R}^{n_v+1}$ satisfying the fixed point condition:

$$V(s_t, \alpha|s^t) = W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), V'(\alpha|s^t)]]. \tag{2}$$

where $V'(\alpha|s^t) = \{V(s, \alpha|(s^t, s))\}_{s=1}^{n_s} \in \mathbb{R}^{(n_v+1)n_s}$ is a vector of continuation summaries. In many examples, V gives the continuation payoffs of a group of agents facing incentive constraints. If these agents have time additive-expected utility preferences, then $W^V[s, a, m] = f(s, a) + \delta m$ and $M^V[s, a, v'] = \sum_{s' \in \mathcal{S}} v'(s')Q(s|s')$. However, our formulation allows us to accommodate problems in which agents have non-time additive or non-expected utility preferences or, indeed, problems in which the forward-looking variables are not payoffs at all (see Section 3).

To ensure the future summaries $V(s_t, \alpha|s^t)$ are well defined and that (2) admits a fixed point in a suitable space of functions, the following restrictions are imposed on W^V and M^V.

Assumption 1. *W^V is increasing and continuous in its third argument. $W^V[\cdot, \cdot, 0]$ is bounded and there is a $\bar{\delta} \in [0, 1)$ such that for all m and $m' \in \mathbb{R}^{n_v+1}$:*

$$\sup_{\mathcal{S} \times \mathcal{A}} \|W^V[s, a, m] - W^V[s, a, m']\| < \bar{\delta}\|m - m'\|,$$

with $\|\cdot\|$ the Euclidean metric (on $\mathbb{R}^{n_v+1}$).

If $v' \in \mathbb{R}^{n_s(n_v+1)}$ and $\kappa \in \mathbb{R}^{n_v+1}$, then we will write $v' + \kappa$ for $v' + (\kappa\ \kappa \cdots \kappa) \in \mathbb{R}^{n_s(n_v+1)}$.

Assumption 2. *For each $(s, a) \in \mathcal{S} \times \mathcal{A}$ and $\kappa \in \mathbb{R}^{n_v+1}$, (i) $M^V[s, a, \cdot]$ is increasing, (ii) $M^V[s, a, \kappa] = \kappa$ and (iii) $M^V[s, a, \cdot]$ is constant sub-additive: for all $v' \in \mathbb{R}^{n_s(n_v+1)}$, $M^V[s, a, v' + \kappa] \leq M^V[s, a, v'] + \kappa$.*

Existence and uniqueness of a bounded function V satisfying (2) follows from Assumptions 1 and 2 and is shown in Appendix A. In the remainder of the paper, summaries of future actions $V(s_t, \alpha|s^t)$ are called *forward-looking* state variables. Let $\mathcal{V} := V(\mathcal{S} \times \mathscr{A})$, i.e. $\mathcal{V}$ is the (bounded) codomain of V.

Constraints are constructed from state variables according to for all $t \in \{0\} \cup \mathbb{N}$, $s_t \in \mathcal{S}$ and $s^t \in \mathcal{S}^t$,

$$H[K_t(\alpha, s^{t-1}), s_t, a_t(s^t), V'(\alpha|s^t)] \geq 0, \tag{3}$$

where $H : \mathcal{K} \times \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s} \to \mathbb{R}^{n_h}$ is bounded. In applications these inequalities capture incentive and resource constraints. We assume throughout that the decision-maker's constraint set is non-empty for some combination of initial state variables.

Objective and Problem The decision-maker's objective, $U : \mathcal{S} \times \mathscr{A} \to \mathbb{R}$, is given by an aggregator over the forward-looking state variables:

$$U(s_0, \alpha) = F[s_0, V(s_0, \alpha)],$$

where $F[s_0, \cdot]$ is non-decreasing. For example, $V(s_0, \alpha) = \{V^i(s_t, \alpha|s^t)\}_{i=0}^{n_v} \in \mathbb{R}^{n_v+1}$ may give the payoffs of agents $i = 0, \ldots, n_v$ and F may attach (possibly state contingent) Pareto weights to these agents. U is then interpreted as a planner's payoff.

The decision-maker's primal problem is:

$$P^* = \sup F[s_0, V(s_0, \alpha)] \tag{P}$$

subject to $\forall t$, s^t, (3). We follow the usual convention $\sup \emptyset = -\infty$.

3 Examples and Variations

Our framework accommodates many examples from the literature. Below we give three that highlight the scope of our method. The first is a limited commitment problem similar to that studied by Kocherlakota (1996) except that we assume agents have non-expected utility Epstein-Zin preferences. Consequently, the law of motion for forward-looking states, in this case the agents' utilities, is non-linear in V'. The next example is a limited commitment problem with physical capital (and standard preferences); it features a backward-looking state variable. The third is an optimal monetary policy problem similar to those considered in Woodford (2003). This problem also features a non-linear law of motion for the forward-looking state variables. All of these examples are outside of the formulation of Messner, Pavoni, and Sleet (2012b) (which features no backward-looking states and linear laws of motion for forward-looking ones) and Marcet and Marimon (2011) (which allows for backward-looking states, but also assumes linear laws of motion for forward-looking ones and restricts attention to concave problems). In addition, Marcet and Marimon (2011) treat backward-looking states quite differently to us.

Example 1 (*Risk sharing with limited commitment and Epstein-Zin preferences*)**.** Two agents share risk. They face shocks to their endowments and to their utility options from separation. Let $\gamma : \mathcal{S} \to \mathbb{R}_+$ give the joint endowment of the agent pair in each shock state and $w : \mathcal{S} \to \mathbb{R}^2$, $w(s) = \{w^i(s)\}_{i=1,2}$, their outside utility options. Let $\mathcal{A} = \mathbb{R}^2_+$ denote a set of possible consumptions for the agents. There are no backward-looking state variables. The continuation payoffs of the two agents, $V^i(s, \alpha)$, $i = 0, 1$, constitute forward-looking state variables. They evolve according to Equation (2) with aggregators:

$$W^V[s, a, m] = \Big(\frac{1-\delta}{1-\mu}(a^i)^{1-\mu} + \delta m^i\Big)_{i=0,1} \qquad M^V[s, v'] = \Big(\Big\{\sum_{s' \in \mathcal{S}} v^{i\prime}(s')^{\sigma} Q(s|s')\Big\}^{\frac{1}{\sigma}}\Big)_{i=0,1}.$$

Boundedness and concavity of these aggregators is assured if $\mu, \sigma \in (0, 1)$. The resource and incentive constraints are collected into a single function:

$$H[s, a, v'] = \begin{pmatrix} W^V[s, a, M^V[s, v']] - w(s) \\ \gamma(s) - \sum_{i=1}^{2} a^i \end{pmatrix} \geq 0.$$

Finally, the decision-maker is a planner who attaches Pareto weight λ^i to the i-th agent. Her objective is $F[s_0, V(s_0, \alpha)] = \sum_{i=0}^{1} \lambda^i V^i(s_0, \alpha)$. □

Example 2 (*Default with capital accumulation*)**.** A lender (agent 0) extends credit to a borrower (agent 1) who can accumulate capital and can default. Let $a = (a^0, a^1) \in A \subset \mathbb{R} \times \mathbb{R}_+$ denote a pair of consumptions for the lender and the borrower. Their (bounded below) utility functions are denoted $f^i(a^i)$, $i = 0, 1$. The borrower operates a risky technology $\gamma : \mathbb{R}_+ \times \mathcal{S} \to \mathbb{R}_+$ that maps the capital stock and current shock to output. γ is assumed bounded. The borrower is free to default and take an outside utility $w : \mathbb{R}_+ \times \mathcal{S} \to \mathbb{R}$ that depends upon the amount of capital she has sunk into the technology and the current shock. The lender and borrower's utilities and capital constitute forwards and backwards-looking state variables with aggregators:

$$W^V[s, a, m] = \Big(f^i(a^i) + \delta m^i\Big)_{i=0,1} \qquad M^V[s, v'] = \Big(\sum_{s' \in \mathcal{S}} v^{i\prime}(s') Q(s|s')\Big)_{i=0,1};$$

and

$$W^K[k, s, a] = \gamma(k, s) - \sum_{i=0,1} a^i.$$

The incentive ("no default") constraint is given by:

$$f^1(a^1) + \delta \sum_{s' \in \mathcal{S}} v^{1\prime}(s') Q(s|s') - w(k, s) \geq 0,$$

and the resource constraint by:

$$W^K[k,s,a] \geq 0.$$

As in the previous example, these may be collected into a single function H. The objective is given by the Pareto sum: $F[s_0, V(s_0,\alpha)] = \sum_{i=0,1} \lambda^i V^i(s_0,\alpha)$. □

Example 3 (*Optimal monetary policy*)**.** The government's social objective over sequences of output $\alpha = \{a_t\}_{t=0}^{\infty}$ and inflation $\{\Delta p_t\}_{t=0}^{\infty}$ is given by $\sum_{t=0}^{\infty} \delta^t L(a_t, \Delta p_t)$ with $L : \mathbb{R}^2 \to \mathbb{R}$ continuous.[6] Output sequences are restricted to $\mathscr{A} := \mathcal{A}^{\infty}$, with $\mathcal{A} = [\underline{a}, \overline{a}]$ a bounded interval. Inflation evolves according to a simple New Keynesian Philips Curve,

$$\Delta p_t = \kappa a_t + \delta \Delta p_{t+1},$$

with the terminal condition $\lim_{t\to\infty} \sum_{t=0}^{\infty} \delta^t \Delta p_t = 0$. Consequently, given an output plan α, inflation at t is $\Delta p_t = V^1(\alpha|t) := \kappa \sum_{\tau=0}^{\infty} \delta^{\tau} a_{t+\tau}$ and the government's continuation payoff: $V^0(\alpha|t) := \sum_{\tau=0}^{\infty} \delta^t L(a_{t+\tau}, V^1(\alpha|t+\tau))$. The government's payoff and inflation serve as forward-looking state variables; there are no backward-looking state variables in this problem. Adopting our previous notation and letting $v = (v^0, v^1)$ and $v' = (v^{0\prime}, v^{1\prime})$ denote, respectively, current and future pairs of payoff and inflation, the (non-linear) aggregator W^V is given by:

$$v = W^V[a, v'] = \begin{pmatrix} L(a, \kappa a + \delta v^{1\prime}) \\ \kappa a \end{pmatrix} + \delta v'.$$

There is no H function in this case and the social objective is simply $F[V(\alpha)] = V^0(\alpha)$. □

3.1 Variations

Our framework also accommodates dynamic (hidden action) moral hazard problems with general recursive preferences and the timing assumed in Hopenhayn and Nicolini (1997).[7] Small modifications of our basic framework admit other economic problems considered in the literature. We briefly describe two of these.

[6] We consider here the deterministic version of the problem as in Woodford (2003). In most of Woodford (2003)'s examples L is a (concave) quadratic approximation to an underlying objective over primitives. For now we place no such restrictions on L.

[7] Under this timing the public signal (a job or unemployment) of a hidden action (job search) is realized in the period after the action is taken. Alternative timing assumptions are possible after modifications of our framework. The modifications are similar to those described in the discussion of hidden information problems below.

Participation Constraints In some problems constraints are supplemented with additional restrictions on the initial values of forward-looking variables. For example, contracting problems often place initial participation constraints on agents. Such constraints are easily incorporated into our basic framework by appending the additional restriction:

$$W^V[s_0, a_0, M^V[s_0, a_0, V'(\alpha)]] - \bar{v} \geq 0, \tag{4}$$

where $\bar{v}$ gives the player's initial outside payoff option. Our basic formulation omits (4), but we point out the small modifications needed to incorporate it.

Hidden Information problems In hidden information problems some or all agents privately observe a shock process. Without loss of generality attention may be restricted to plans that induce agents to truthfully reveal their current shock. This requires incentive constraints that "run across" contemporaneous shock states and, hence, the replacement of $H : \mathcal{K} \times \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s} \to \mathbb{R}^{n_h}$ with $\tilde{H} : \mathcal{K} \times \mathcal{A}^{n_s} \times \mathcal{V}^{n_s \times n_s} \to \mathbb{R}^{n_h}$. For example, consider the simplest case in which a firm induces a worker to reveal whether she is well ($s = 1$) or sick ($s = 2$). The incentive constraints require that well workers reveal their health and are of the form:

$$\begin{aligned} \tilde{H}[\{a_t(s^{t-1}, s)\}, \{V'(\alpha|s^{t-1}, s)\}] = u(1, a_t(s^{t-1}, 1)) + \delta \sum_{s' \in \mathcal{S}} V(s', \alpha|s^{t-1}, 1, s')Q(1, s') \\ - u(1, a_t(s^{t-1}, 2)) - \delta \sum_{s' \in \mathcal{S}} V(s', \alpha|s^{t-1}, 2, s')Q(1, s') \geq 0, \end{aligned}$$

where $a_t(s^t) \in \mathcal{A}$ is now the bundle of consumption and effort prescribed after health history s^t. In this case, $W^V[s, a, m] = u(s, a) + \delta m$ and $M^V[s, v'] = \sum_{s' \in \mathcal{S}} v'(s')Q(s, s')$. If the health shocks are i.i.d., then $M^V[v'] = \sum_{s' \in \mathcal{S}} v'(s')Q(s')$ and it is more convenient to redefine the forward-looking state as the certainty equivalent of V, i.e. as $\tilde{V}(\alpha|s^{t-1}) = M^V[V'(\alpha|s^{t-1})]$. Forward-looking states then evolve according to

$$\tilde{V}(\alpha|s^{t-1}) = M^V[W^V[s_t, a_t(s^t), \tilde{V}(\alpha|s^t)]]$$

and the constraints become:

$$\begin{aligned} \tilde{H}[\{a_t(s^{t-1}, s)\}, \{\tilde{V}(\alpha|s^{t-1}, s)\}] = u(1, a_t(s^{t-1}, 1)) + \delta \tilde{V}(\alpha|s^{t-1}, 1) \\ - u(1, a_t(s^{t-1}, 2)) - \delta \tilde{V}(\alpha|s^{t-1}, 2) \geq 0. \end{aligned}$$

Slightly modified versions of all the results given below hold with $\tilde{H}$ replacing H.

4 Augmented Primal and a Recursive Primal Problem

We define an augmented primal problem in which state variables are introduced as explicit choices rather than as functions of past actions. Our motive for introducing this problem is that it, rather than the original one, is amenable to direct recursive decomposition. We give a recursive formulation that decomposes the augmented problem into sub-problems linked by state variables. The difficulties with this formulation motivate our subsequent dual approach.

4.1 Augmented primal problem

Define a *primal process* π to be a plan α augmented with a process for backwards and forward-looking states $\{k_t, v_t\}_{t=0}^{\infty}$. The set of primal processes is given by:

$$\mathscr{P} = \left\{ \pi = (\alpha, \{k_t, v_t\}_{t=0}^{\infty}) \;\middle|\; \begin{array}{rl} \alpha \in \mathscr{A},\ k_0 \in \mathcal{K}, & \forall t \in \mathbb{N}, k_t : \mathcal{S}^{t-1} \to \mathcal{K}, \\ v_0 \in \mathcal{V}, & \forall t \in \mathbb{N}, v_t : \mathcal{S}^{t} \to \mathcal{V} \end{array} \right\}.$$

The *augmented primal problem* is:

$$\sup F[s_0, v_0] \tag{AP}$$

subject to $\pi \in \mathscr{P}$, $k_0 = \bar{k}$ and $\forall t$, s^t,

$$k_{t+1}(s^t) = W^K[k_t(s^{t-1}), s_t, a_t(s^t)], \tag{5}$$

$$v_t(s^t) = W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), v_{t+1}(s^t)]], \tag{6}$$

and

$$H[k_t(s^{t-1}), s_t, a_t(s^t), v_{t+1}(s^t)] \geq 0. \tag{7}$$

Thus, the augmented primal problem (AP) re-expresses constraints in terms of state processes.[8] We record the following (obvious) fact.

Proposition 1. *If P^* is the optimal value for* (P)*, then it is also the optimal value for* (AP)*. α^* solves* (P) *if and only if there is a state process $\{k_t^*, v_t^*\}_{t=0}^{\infty}$ such that $(\alpha^*, \{k_t^*, v_t^*\}_{t=0}^{\infty})$ solves* (AP)*.*

[8] Participation constraints are incorporated by adding $v_0 \geq \bar{v}$ to the constraint set.

4.2 Recursive Primal Problem

In this section we give a recursive primal formulation of a principal-agent problem.[9] In such a problem, a committed principal possessing no private information designs a contract to motivate a group of agents. A forward-looking variable V^0 defining the principal's payoff function is the objective and does not enter the constraints. It is convenient to exploit this structure by separating the principal's payoff V^0 from the other forward-looking variables (typically utility promises to agents) and redefining $V := \{V^i\}_{i=1}^{n_v}$ to exclude V^0. The problem becomes:

$$\sup V^0(s_0, \alpha)$$

subject to $\forall t$, s^t, (3) with V (and H, W^V and M^V) redefined. The augmented version of this problem is:

$$\sup V^0(s_0, \alpha) \tag{PA}$$

subject to $\pi \in \mathscr{P}$, $k_0 = \bar{k} \in \mathcal{K}$ and (5) to (7) with v_t redefined to exclude v_t^0, i.e. $v_t = \{v_t^i\}_{i=1}^{n_v}$.

The aggregators W^K, W^V and M^V may be used to decompose (PA) into a family of sub-problems linked by elements in $\mathcal{S}$, $\mathbb{R}^{n_k}$ and $\mathbb{R}^{n_v}$. It is useful to identify "state spaces" on which these sub-problems are well-posed (i.e. have non-empty constraint sets). To that end define the "endogenous state space" $\mathcal{X}$ to be the largest subset of $\mathcal{K} \times \mathcal{S} \times \mathcal{V}$ satisfying the recursion:

$$\mathcal{X} = \left\{ (k,s,v) \;\middle|\; \begin{array}{ll} \exists (a,k',v') \in \mathcal{A} \times \mathcal{K} \times \mathcal{V}, & k' = W^K[k,s,a], \\ v = W^V[s,a,M^V[s,a,v']], & H[k,s,a,v'] \geq 0, \\ \text{and } \forall s' \in \mathcal{S}, & (k',s',v'(s')) \in \mathcal{X} \end{array} \right\}. \tag{8}$$

Crucially, while $\mathcal{K}$ and $\mathcal{V}$ are given exogenously or are easy to find, $\mathcal{X}$ is often neither. In addition, let:

$$\Gamma(k,s,v) = \left\{ (a,k',v') \in \mathcal{A} \times \mathcal{K} \times \mathcal{V}^{n_s} \;\middle|\; \begin{array}{l} k' = W^K[k,s,a], v = W^V[a,s,M^V[a,s,v']], \\ H[k,s,a,v'] \geq 0 \text{ and } (k',s',v'(s')) \in \mathcal{X} \end{array} \right\}.$$

Define:

$$\mathcal{V}(k,s) = \{v : (k,s,v) \in \mathcal{X}\}$$

[9] The principal-agent problem is a special case of (P). The more general problem (P) also has primal recursive formulations, see Kocherlakota (1996), Rustichini (1998) and, especially, Messner, Pavoni, and Sleet (2012b), Section 7. However, since our goal here is to briefly review the recursive primal approach and point out its limitations, we restrict ourselves to a recursive primal treatment of the simpler principal-agent problem.

and let $W^{V,0}$ and $M^{V,0}$ denote the time and stochastic aggregators for V^0.

Proposition 2. *Let $P_0^* \in \mathbb{R} \cup \{-\infty\}$ be the optimal value for problem (PA). Then:*

$$P_0^* = \sup_{\mathcal{V}(\bar{k},s_0)} P^*(\bar{k}, s_0, v_0), \tag{9}$$

where P^* satisfies the recursion, for each $(k,s,v) \in \mathcal{X}$,

$$P^*(k,s,v) = \sup_{\Gamma(k,s,v)} W^{V,0}[s,a,M^{V,0}[s,a,P^*(k',v')]], \tag{10}$$

with $P^*(k',v') = \{P^*(k',s',v'(s'))\}_{s'=1}^{n_s}$. In addition, $(\alpha^*, \{k_t^*, v_t^*\}_{t=0}^{\infty})$ solves (PA) if and only if (i) $k_0^* = \bar{k}$ and $v_0^* \in G_0^*$ and (ii) for all $t \in \mathbb{N}, s^t \in \mathcal{S}^t$, $(a_t^*(s^t), k_{t+1}^*(s^t), v_{t+1}^*(s^t)) \in G^*(k_t^*(s^{t-1}), s_t, v_t^*(s^{t-1}))$, where:

$$G_0^* := \operatorname*{argmax}_{\mathcal{V}(\bar{k},s_0)} P^*(\bar{k}, s_0, v) \text{ and}$$

$$G^*(k,s,v) := \operatorname*{argmax}_{\Gamma(k,s,v)} W^{V,0}[s,a,M^{V,0}[s,a,P^*(k',v')]].$$

Proof. See Appendix B. □

Note that the role of the 'first stage problem' (9) is to provide an optimal initial condition v_0^* for the forward-looking state variables; (10) then gives the primal Bellman equation. As Proposition 2 indicates $\mathcal{X}$ is generally part of the solution to the problem along with P_0^* and P^*. Stokey, Lucas, and Prescott (1989) document problems in which $\mathcal{X}$ *is* determined exogenously as a primitive of the problem. However, for many other problems, in particular those with forward-looking state variables, $\mathcal{X}$ is given implicitly and recovering it (i.e. solving the fixed point problem defined by (8)) is a major complication. This motivates the dual approach.

5 Recursive Dual

We begin this section by defining a Lagrangian for (AP). The Lagrangian involves the product of constraint values with multipliers. We collect the former into an object called a constraint process and the latter into an object called a dual process. Definitions of these follow.

5.1 Lagrangians and Dual Problems

As a preliminary, we make a small adjustment to the definition of a primal process. Recall that backwards-looking state variables k_t were previously defined to be s^{t-1}-measurable. To fully exploit the recursive structure in the Lagrangian it is convenient to allow these variables to be s^t-measurable and to enforce s^{t-1}-measurability via their law of motion. Thus, from now on, unless further restricted, each t-dated variable (including k_t) in a primal process $\pi = \{k_t, a_t, v_t\}_{t=0}^{\infty}$ is s^t-measurable.

A constraint process evaluates constraint functions inclusive of laws of motion at a given primal process. For each primal process π, let $z_0^K(\pi) = \bar{k} - k_0$ and, for all $t \in \mathbb{N}$ and $s^t \in \mathcal{S}^t$, let:

$$z_t^K(\pi)(s^t) = W^K[k_{t-1}(s^{t-1}), s_{t-1}, a_{t-1}(s^{t-1})] - k_t(s^t).$$

Then $\{z_t^K(\pi)\}_{t=0}^{\infty}$ gives the values of the law of motion for backward-looking constraints (inclusive of the initial condition) at π. Similarly, define for all $t \in \{0\} \cup \mathbb{N}, s^t \in \mathcal{S}^t$,

$$z_t^V(\pi)(s^t) = W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), v_{t+1}(s^t)]] - v_t(s^t)$$

and $z_t^H(\pi)(s^t) = H[k_t(s^{t-1}), s_t, a_t(s^t), v_{t+1}(s^t)]$. Then $\{z_t^V(\pi)\}_{t=0}^{\infty}$ and $\{z_t^H(\pi)\}_{t=0}^{\infty}$ give the values of the forward-looking law of motion and H constraints at π. These terms are collected into the *constraint process* $\zeta(\pi) = \{z_t^j(\pi)\}_{t=0, j\in\mathcal{J}}^{\infty}$, $\mathcal{J} := \{K, V, H\}$. The boundedness assumptions placed on primitives and the countable number of constraints ensure that for all $\pi \in \mathscr{P}$, $\zeta(\pi) \in \ell_\infty$.[10]

A dual process contains summable ("countably additive") multipliers for the various constraints facing the decision-maker. Let $\theta^K = \{q_t^K\}_{t=0}^{\infty}$, with $q_t^K : \mathcal{S}^t \to \mathbb{R}^{n_k}$, denote multipliers (co-states) for the backward-looking law of motion and $\theta^V = \{q_t^V\}_{t=0}^{\infty}$, with $q_t^V : \mathcal{S}^t \to \mathbb{R}^{n_v+1}$, multipliers (co-states) for the forward-looking law of motion. Let $\theta^H = \{q_t^H\}_{t=0}^{\infty}$, with $q_t^H : \mathcal{S}^t \to \mathbb{R}^{n_h}$, denote multipliers for the H-constraints. Collect these various multipliers into a *dual process* $\theta = \{\theta^j\}_{j\in\mathcal{J}}$ and define the set of (bounded) dual processes:

$$\mathcal{Q} = \left\{ \theta \,\middle|\, \sum_{\mathcal{J}} \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} \bar{\delta}^{t-1} \|q_t^j(s^t)\| < \infty \right\},$$

[10] We use ℓ_∞ to denote the set of sup-norm bounded, vector valued sequences: $\{\{x_n\}_{n=1}^{\infty} | x_n \in \mathbb{R}^m, \sup_{n\in\mathbb{N}} \|x_n\| < \infty\}$, where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^m$. In our setting, $m = \sum_{j\in\mathcal{J}} n_j + 1$ and $x_n = \{z_{t(n)}^j(s^t(n))\}_{j\in\mathcal{J}}$ for some enumeration of histories $s^t(n)$.

with $\bar{\delta} \in (0,1)$ the discount from the aggregator W^V. Define the Lagrangian:

$$\mathscr{L}(\pi,\theta) = F[s_0, v_0] + \langle \theta, \zeta(\pi) \rangle,$$

where $\langle \theta, \zeta(\pi) \rangle = \sum_{\mathcal{J}} \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} \bar{\delta}^t \{q_t^j(s^t) \cdot z_t^j(\pi)(s^t)\}$ and $\cdot$ is the usual vector dot product.

The decision-maker's augmented primal problem (AP) may be re-expressed as a sup-inf problem:

$$P_0^* := \sup_{\mathscr{P}} \inf_{\mathscr{Q}} \mathscr{L}(\pi,\theta). \tag{SI}$$

Its dual interchanges the infimum and supremum operations:

$$D_0^* := \inf_{\mathscr{Q}} \sup_{\mathscr{P}} \mathscr{L}(\pi,\theta). \tag{IS}$$

Discussion of the relation between these problems is deferred until Section 8. Instead in the remainder of this section we pursue a recursive formulation of (IS).[11]

5.2 Recursive Dual

The recursive dual formulation decomposes (IS) into sub-problems linked by co-state variables. We introduce some preliminary notation and concepts. We call $p = (a, k, v')$ a *current primal choice* where $a \in \mathcal{A}$ is a current action, $k \in \mathcal{K}$ is a *current* backwards-looking state and $v' \in \mathcal{V}^{n_s}$ is a tuple of *continuation* forward-looking states, one for each future shock s'. Current primal choices belong to $\mathcal{P} = \mathcal{A} \times \mathcal{K} \times \mathcal{V}^{n_s}$.[12] We call $q = (q^H, y')$ a *current dual choice* where $q^H \in \mathbb{R}_+^{n_h}$ is a current H-constraint multiplier and $y' = (q^{K\prime}, q^{V\prime}) \in \mathbb{R}^{n_s(n_k+n_v+1)}$ is a tuple of co-states for the next period's backward and forward-looking laws of motion. Current dual choices belong to $\mathcal{Q} = \mathbb{R}_+^{n_h} \times \mathbb{R}^{n_s(n_k+n_v+1)}$. Let $y = (q^K, q^V) \in \mathcal{Y} := \mathbb{R}^{n_k+n_v+1}$ denote a pair of co-state variables on current laws of motion.

The Lagrangian in (IS) may be expanded as:

$$\begin{aligned} D_0^* = \inf_{\mathscr{Q}} \sup_{\mathscr{P}} F[s_0, v_0] - q_0^V \cdot \{v_0 - W^V[s_0, a_0, M^V[s_0, a_0, v_1]]\} + q_0^K \cdot (\bar{k} - k_0) \\ + \bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta(\pi) | s_1 \rangle, \end{aligned} \tag{11}$$

with $\langle \theta, \zeta(\pi) | s_1 \rangle = \sum_{\mathcal{J}} \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} \bar{\delta}^t q_{t+1}^j(s_1, s^t) \cdot z_{t+1}^j(\pi)(s_1, s^t)$ the continuation of $\langle \theta, \zeta(\pi) \rangle$

[11] An initial participation constraint may be incorporated by appending $z_{-1}^V = v_0 - \bar{v}$ and multiplier q_{-1}^V to the constraint and dual process respectively.

[12] Our notation convention is to use calligraphic letters $\mathcal{P}$ for sets of current actions and script letters $\mathscr{P}$ for sets of stochastic processes.

after the realization of the first period shock s_1. Removing $F[s_0, v_0] - q_0^V \cdot v_0 + q_0^K \cdot \bar{k}$ from (11) and fixing the initial co-states $y_0 = (q_0^K, q_0^V)$ gives the following *continuation dual problem*:

$$D^*(s_0, y_0) = \inf_{\mathscr{Q}(y_0)} \sup_{\mathscr{P}(v_0)} -q_0^K \cdot k_0 + q_0^V \cdot W^V[s_0, a_0, M^V[s_0, a_0, v_1]] \qquad (12)$$
$$+ q_0^H \cdot H[k_0, s_0, a_0, v_1] + \bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta | s_1 \rangle,$$

where $\mathscr{Q}(y_0)$ omits $y_0 = (q_0^K, q_0^V)$ from $\mathscr{Q}$, $\mathscr{P}(v_0)$ omits v_0 from $\mathscr{P}$. Collecting terms in (12) involving the initial current primal choice $p_0 = (a_0, k_0, v_1)$ gives the current "dual" payoff J:

$$J(s_0, y_0; q_0, p_0) = -q_0^K \cdot k_0 + q_0^V \cdot W^V[s_0, a_0, M^V[s_0, a_0, v_1]] + q_0^H \cdot H[k_0, s_0, a_0, v_1]$$
$$- \bar{\delta} \sum_{s_1 \in \mathcal{S}} q_1^V(s_1) \cdot v_1(s_1) + \bar{\delta} \sum_{s_1 \in \mathcal{S}} q_1^K(s_1) \cdot W^K[k_0, s_0, a_0]. \qquad (13)$$

Note that the terms in the second line of (13) are extracted from $\bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta | s_1 \rangle$ in (12). Below we give explicit economic interpretations of the terms in J in the context of examples. Proposition 3 relates D_0^*, D^* and J and gives the key dynamic programming result for dual value functions.

Proposition 3 (Value functions)**.** *The value D_0^* satisfies:*

$$D_0^* = \inf_{\mathcal{Y}} \sup_{\mathcal{V}} F[s_0, v] - q^V \cdot v + q^K \cdot \bar{k} + D^*(s_0, q^K, q^V), \qquad (14)$$

with for all $(s, y) \in \mathcal{S} \times \mathcal{Y}$,

$$D^*(s, y) = \inf_{\mathcal{Q}} \sup_{\mathcal{P}} J(s, y; q, p) + \bar{\delta} \sum_{s' \in \mathcal{S}} D^*\left(s', y'(s')\right), \qquad (15)$$

where $y'(s') = (q^{K\prime}, q^{V\prime})(s')$.

Proof. See Appendix C. □

The first stage problem (14) generates the initial co-states; (15) then gives the dual Bellman equation. Moving from the dual problem (IS) to the recursive dual problems (14) and (15) involves interchanging an infimum operation over future dual variables with a supremum operation over current primal ones. In general interchanging such operations alters optimal values. But here the additive separability of the Lagrangian in these two sets of variables ensures that it does not. See the proof of Proposition 3 for details. Note if the laws of motion for backward or forward-looking states are non-linear in these states, then it is necessary to work with the Lagrangian of the augmented primal to ensure this separability.

Remark 1. The function J may be interpreted as an augmented Hamiltonian. Suppose that $W^V[s,a,m] = u(s,a) + \bar{\delta}m$, $M^V[s,a,v'] = \sum_{\mathcal{S}} v'(s')Q(s|s')$ and $H = 0$, then J reduces to:

$$J(s,y;q,p) = \left\{\bar{\delta}\sum_{\mathcal{S}} q^{K\prime}(s') - q^K\right\} \cdot k + \bar{\delta}\sum_{\mathcal{S}}\left\{q^V Q(s|s') - q^{V\prime}(s')\right\} \cdot v'(s') + q^V \cdot u(s,a) + \bar{\delta}\sum_{\mathcal{S}} q^{K\prime}(s') \cdot \{W^K[k,s,a] - k\}.$$

The terms in the second line isolate the current action and correspond to a classical Hamiltonian. J augments this with additional terms involving adjustments to the shadow value of backward and forward-looking states. Assuming differentiability of W^K and differentiating with respect to k gives a discrete time analogue of the co-state equation from optimal control. In our more general setting, current resource and incentive conditions are explicitly incorporated into J via the H function and linearity of J in the forward-looking states is not assumed.

Remark 2. Our recursive dual formulation relies entirely dual co-state variables y_t to summarize the past. This contrasts with Marcet and Marimon (2011) who dualize a subset of constraints and make use of a mixture of primal and dual variables to summarize histories.

Remark 3. The primal "state" variables k and v continue to appear in the recursive dual problem. This allows us to accommodate non (quasi-)linear laws of motion for such variables in our framework. However, they are no longer passed between sub-problems in the recursive dual setting and in this sense no longer function as state variables.[13]

Definition 1. Let $\mathscr{F}$ denote the set of proper functions $D : \mathcal{S} \times \mathcal{Y} \to \mathbb{R} \cup \{\infty\}$ that are not everywhere infinite valued. Define the *dual Bellman operator* $\mathbf{B} : \mathscr{F} \to \mathscr{F}$, $\forall (s,y) \in \mathcal{S} \times \mathcal{Y}$,

$$\mathbf{B}(D)(s,y) = \inf_{\mathcal{Q}} \sup_{\mathcal{P}} J(s,y;q,p) + \bar{\delta} \sum_{s' \in \mathcal{S}} D\left(s', y'(s')\right).$$

The following theorem recasts D^* as a fixed point of $\mathbf{B}$. It is an immediate corollary of Proposition 3.

[13] Notice also that they are restricted to the exogenous $\mathcal{K} \times \mathcal{V}^{n_s}$ and not the endogenous $\mathcal{X}$. Choices of primal states inconsistent with $\mathcal{X}$ are (finitely) penalized via the Lagrangian.

Theorem 1. $D^* = \mathbf{B}(D^*)$.

To make the preceding discussion concrete we revisit the examples.

Example 1 (*Risk sharing with limited commitment and Epstein-Zin preferences*)**.** This example lacks backward-looking state variables. The initial period problem is:

$$D_0^* = \inf_{\mathbb{R}^2} \sup_{\mathcal{V}} \; \lambda \cdot v_0 - q^V \cdot v_0 + D^*(s_0, q^V),$$

where recall λ is a pair of exogenous Pareto weights, q^V is a pair of initial co-states and v_0 is a pair of utility promises drawn from the exogenous feasibility set $\mathcal{V}$. The recursive dual problem is as in (15), but without backward-looking states or co-states and with the current dual function:

$$J(s, q^V; q, p) = \sum_{i=0,1} (q^{V,i} + q^{H,i}) \left\{ \frac{1-\delta}{1-\mu}(a^i)^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} v^{i\prime}(s')^{\sigma} Q(s|s') \right)^{\frac{1}{\sigma}} \right\} \tag{16}$$
$$- \sum_{i=0,1} q^{H,i} w^i(s) - q^{H,2} \left(\sum_{i=0,1} a^i - \gamma(s) \right) - \delta \sum_{s' \in \mathcal{S}} q^{V\prime}(s') \cdot v'(s').$$

The function, J incorporates the "shadow value" of delivering utility to the agents (inclusive of relaxation of the incentive constraints) less the shadow costs of resources and continuation utility promises. □

Example 2 (*Default with capital accumulation*)**.** In this case,

$$D_0^* = \inf_{\mathbb{R}^3} \sup_{\mathcal{V}} \; \lambda \cdot v_0 - q^V \cdot v_0 + q^K \cdot \bar{k} + D^*(s_0, q^K, q^V).$$

The recursive dual problem is as in (15), but now with current dual function:

$$J(s, q^K, q^V; q, p) = - q^K \cdot k + q^{V,0} \Big\{ f^0(a^0) + \delta \sum_{s' \in \mathcal{S}} v^{0\prime}(s') Q(s|s') \Big\} \tag{17}$$
$$+ (q^{V,1} + q^{H,1}) \Big\{ f^1(a^1) + \delta \sum_{s' \in \mathcal{S}} v^{1\prime}(s') Q(s|s') \Big\} - q^{H,1} w(k, s)$$
$$+ (\delta \sum_{s'} q^{K\prime}(s') + q^{H,2}) \Big(\gamma(k, s) - \sum_{i=0,1} a^i \Big) - \delta \sum_{s' \in \mathcal{S}} q^{V\prime}(s') \cdot v'(s').$$

Many of the terms in (17) have similar interpretations to those in (16). In addition, $-q^K \cdot k$ is the shadow cost of using k of the backward-looking state in the present and $\delta \sum_{s' \in \mathcal{S}} q^{K\prime}(s')(\gamma(k, s) - \sum_{i=0,1} a^i)$ is the shadow benefit of delivering $\gamma(k, s) - \sum_{i=0,1} a^i$ of this state variable into the future. □

Example 3 (*Optimal monetary policy*)**.** In this case, the period 0 dual period value is given by:

$$D_0^* = \inf_{\mathbb{R}^2} \sup_{\mathcal{V}} \quad v_0^0 - q^V \cdot v_0 + D^*(q^V),$$

where $v_0 = (v_0^0, v_0^1)$ is the period 0 government payoff and inflation, while the current dual function specializes to:

$$J(q^V; q, p) = q^{V,0}\left\{L(a, \kappa a + \delta v^{1\prime}) + \delta v^{0\prime}\right\} + q^{V,1}\left\{\kappa a + \delta v^{1\prime}\right\} - \delta q^{V\prime} \cdot v'. \tag{18}$$

Here J incorporates the shadow value of delivering payoff to the government and inflation less the cost of shadow cost of future payoff and inflation promises. □

State Spaces Specialized to the principal-agent case Proposition 3 supplies a dual analogue of the value function component of Proposition 2 (the policy function component follows below). It relocates the dynamic programming to a state space of dual co-state variables. As previously emphasized, determining the endogenous set of feasible states in the recursive primal setting is problematic and adds another layer of calculation. The next result shows that in the dual setting (with bounded primal variables), the dual value function D^* is finite-valued on all of $\mathcal{S} \times \mathcal{Y}$ $(= \mathcal{S} \times \mathbb{R}^{n_k+n_v+1})$. Thus, the effective dual state space, on which choice sets are non-empty and value functions finite, is immediately determined.

Proposition 4. *$D^* : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$.*

Proof. See Appendix C. □

The immediate determination of the state space is an important advantage of the dual approach. In addition, it is easily verified that each $D^*(s, \cdot)$ is positively homogenous of degree one (see Lemma 1 below). This has the advantage that once the dual value functions $D^*(s, \cdot)$ are determined on the unit circle $\mathcal{C} = \{y \in \mathcal{Y} | \|y\| = 1\}$, then they are determined everywhere via positive scaling. From a practical point of view, the state space may be identified with $\mathcal{S} \times \mathcal{C}$. To make this concrete, consider Example 2. In this example, there are two co-states (associated with capital and borrower payoffs) and the effective dual state space is simply S copies of the unit circle in $\mathbb{R}^2$. In contrast, in the primal formulation of the problem the state space is an implicit subset of $\mathbb{R}_+ \times \mathbb{R}$ describing the set of incentive-feasible capitals and borrower payoffs. This would have to be calculated separately adding an extra layer of calculation. We take up the issue of how to approximate value functions on $\mathcal{C}$ in Section 9. Less positively the homogeneity of candidate value functions combined with the unboundedness of the current dual set $\mathcal{Q}$ (i.e. the set of current dual choices in (15)) disrupts the conventional approach to proving that $\mathbf{B}$ is a contraction. We address this issue in Section 6.

Policies We now turn to policies. For arbitrary sets C and E and function $g : C \times E \to \overline{\mathbb{R}}$, define the argminmax operation

$$\underset{C|E}{\operatorname{argminmax}}\, g = \left\{ (c^*, e^*) \;\middle|\; c^* \in \underset{C}{\operatorname{argmin}} \sup_E g(c,e) \text{ and } e^* \in \underset{E}{\operatorname{argmax}}\, g(c^*, e) \right\}.$$

The solution to the sequential dual (IS) is given by:

$$\Lambda^{IS} := \underset{\mathscr{Q}|\mathscr{P}}{\operatorname{argminmax}}\, \mathscr{L}(\pi, \theta).$$

On the other hand, the solution to the recursive dual is described by a set:

$$G_0^{IS} = \underset{\mathcal{Y}|\mathcal{V}}{\operatorname{argminmax}}\, F[s_0, v] - q^V \cdot v + q^K \cdot \bar{k} + D^*(s_0, q^K, q^V)$$

and a correspondence

$$G^{IS}(s, y) = \underset{\mathcal{Q}|\mathcal{P}}{\operatorname{argminmax}}\;\; J(s, y; q^H, y', p) + \bar{\delta} \sum_{s' \in \mathcal{S}} D^*\left(s', y'(s')\right).$$

Any element (θ^*, π^*) in $\Lambda^{IS} \subset \mathscr{Q} \times \mathscr{P}$ implies an initial $(v_0^*, y_0^*) = (v_0^*, q_0^{K*}, q_0^{V*})$ and a sequence of multipliers and choices $\{q_t^*, p_t^*\}_{t=0}^{\infty}$, with $q_t^* = (q_t^{H*}, y_{t+1}^*) = (q_t^{H*}, q_{t+1}^{K*}, q_{t+1}^{V*})$. On the other hand, such a sequence can be recovered from G_0^{IS} and G^{IS}: $(y_0, v_0) \in G_0^{IS}$ and $(q_t(s^t), p_t(s^t)) \in G^{IS}(s_t, y_t(s^t))$ for each t, s^t. The next proposition relates policies from the dual and the recursive dual.

Proposition 5 (Policy functions). *$(\theta^*, \pi^*) \in \Lambda^{IS}$ only if $(q_0^{K*}, p_0^{V*}) \in G_0^{IS}$ and for each $t \in \mathbb{N}$, $s^t \in \mathcal{S}^t$, $(q_t^{H*}(s^t), y_{t+1}^*(s^t), p_t^*(s^t)) \in G^{IS}(s_t, y_t^*(s^t))$. Conversely, $(\theta^*, \pi^*) \in \Lambda^{IS}$ if $(q_0^{K*}, p_0^{V*}) \in G_0^{IS}$, for each $t \in \mathbb{N}$, $s^t \in \mathcal{S}^t$, $(q_t^{H*}(s^t), y_{t+1}^*(s^t), p_t^*(s^t)) \in G^{IS}(s_t, y_t^*(s^t))$ and:*

$$\lim_{T \to \infty} \bar{\delta}^{T+1} \sum_{\mathcal{S}^{T+1}} D^*(s_{T+1}, y_{T+1}^*(s^{T+1})) \geq 0. \tag{T}$$

Proof. Appendix C. □

Example 1 (*Risk sharing with limited commitment and Epstein-Zin preferences; Policies*)**.** From (15) and (16) it follows that the consumptions $a = (a^0, a^1)$ are chosen to solve the "Pareto problems":

$$\max_{\mathbb{R}_+} \quad (q^{V,i} + q^{H,i})\frac{1-\delta}{1-\mu}(a^i)^{1-\mu} - q^{H,2}a^i, \qquad i = 0, 1.$$

It is easily shown that:

$$a^i = \frac{r_i}{1+r_i}\gamma(s),$$

where $r_i = \left(\frac{q^{V,i}+q^{H,i}}{q^{V,j}+q^{H,j}}\right)^{\frac{1}{\mu}}$, $j = 0, 1, j \neq i$. The continuation forward-looking states are chosen to solve:

$$\max_{\mathcal{V}} \sum_{i=0,1} (q^{V,i} + q^{H,i}) \left(\sum_{s' \in \mathcal{S}} v^{i\prime}(s')^{\sigma} Q(s|s') \right)^{\frac{1}{\sigma}} - \sum_{i=0,1} \sum_{s' \in \mathcal{S}} q^{V,i\prime}(s') v^{i\prime}(s'). \tag{19}$$

If the boundaries implied by $\mathcal{V}$ are non-binding, then (19) implies that the co-states (endogenous Pareto weights) evolve as:

$$\left[\sum_{s' \in \mathcal{S}} \left(\frac{q^{V,i\prime}(s')}{Q(s'|s)} \right)^{\frac{\sigma}{\sigma-1}} Q(s'|s) \right]^{\frac{\sigma-1}{\sigma}} = q^{V,i} + q^{H,i},$$

i.e. the stochastic aggregator of an agent's (normalized) continuation Pareto weights is adjusted upwards if the multiplier on her incentive constraint $q^{H,i}$ is positive.[14] If $\sigma \in (0, 1)$, then the stochastic aggregator is concave and increments to low valued continuation utilities are more valuable than increments to high valued ones. Consequently, in contrast to the standard expected utility case, following a binding incentive constraint (a positive $q^{H,i}$ value), the agent's continuation Pareto weight is increased more in low continuation utility states than in high, i.e. the incremental reward to keep the agent inside the risk sharing arrangement is skewed towards these states. To see this note that the sub-problem (19) implies (absent binding boundaries):

$$\frac{q^{V,i\prime}(s')}{Q(s|s')} = (q^{V,i} + q^{H,i}) \left(\frac{v^{i\prime}(s')}{\{\sum_{s'' \in \mathcal{S}} v^{i\prime}(s'')^{\sigma} Q(s|s'')\}^{\frac{1}{\sigma}}} \right)^{\sigma-1}. \qquad \square$$

An explicit solution to Example 1 is computed in Section 9. Solutions to the other

[14] Note that the utility certainty equivalent uses the power σ, but the Pareto weight aggregator uses the dual or conjugate exponent to σ, $\frac{\sigma}{\sigma-1}$.

examples are discussed in Section 7 where we exploit or impose additional structure.

6 Contraction

This section establishes sufficient conditions for $\mathbf{B}$ to be contractive on an appropriate space of functions. The combination of an unbounded dual value function and an unbounded dual constraint correspondence[15] is an obstacle to conventional approaches to proving contractivity.[16] Following Thompson (1963), Marinacci and Montrucchio (2010) and especially Rinćon-Zapatero and Rodríguez-Palmero (2003), we pursue a different approach. The basic idea is to restrict attention to spaces of functions having a certain scaling property. Specifically, for any distinct pair g_1, g_2, scaleability requires a positive number $b \in \mathbb{R}_+$ satisfying $bg_1 \geq g_2$. Distances between function pairs (g_1, g_2) are then identified with the log of the smallest scaling factor b such that both $bg_1 \geq g_2$ and $bg_2 \geq g_1$. Scaleability of a set of candidate value functions is ensured via a renormalization involving bounding value functions that are themselves scaleable (after renormalization). Since the optimal dual value function is convex and positively homogenous (see below) in co-states, we restrict attention to candidate value functions with these properties. Consequently, it is sufficient for us to have scaleability on the unit circle in the co-state space (i.e. on a compact set) and to define our distance measures accordingly. The interval of convex, positively homogenous functions between the bounding value functions is a complete metric space. If $\mathbf{B}$ is a self-map on this interval, then contractivity follows from monotonicity and concavity of $\mathbf{B}$, the properties of the bounding value functions and the homogeneity of candidate value functions.[17]

The following definition is useful.

Definition 2. A function $D : \mathcal{Y} \to \mathbb{R}$ is *sub-linear* if (i) $D(\cdot)$ is convex and (ii) $D(\cdot)$ is positively homogeneous of degree 1. A function $D : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$ is sub-linear if each $D(s, \cdot)$ is sub-linear.

Lemma 1 indicates the importance of the previous definition for our setting.

Lemma 1. *(i) D^* is sub-linear. (ii) If $D : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$ is sub-linear, then $\mathbf{B}(D)$ is sub-linear.*

Proof. See Appendix D. □

[15]The current dual choice set is $\mathcal{Q} = \mathbb{R}^{n_h}_+ \times \mathbb{R}^{n_s(n_k+n_v+1)}$.

[16]When the optimal value function is unbounded and the constraint correspondence compact-valued it is often possible to prove contractivity on a space of weight norm bounded functions. In the dual setting, this approach is disrupted by the unboundedness of the constraint correspondence (for multipliers and co-states).

[17]Thus, Blackwell's Theorem is avoided.

Once again, let $\mathcal{C} = \{y \in \mathcal{Y} | \|y\| = 1\}$ denote the unit circle in $\mathbb{R}^{n_k+n_v+1}$. The key assumption ensuring contractivity is the following.

Assumption 3 (Bounds). *There is a triple of functions $\underline{\underline{D}} : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$, $\underline{D} : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$ and $\overline{D} : \mathcal{S} \times \mathcal{Y} \to \mathbb{R}$ and a pair of numbers ε_0, $\varepsilon_1 > 0$ such that for each s, $\underline{\underline{D}}(s, \cdot)$ is continuous and positively homogeneous of degree 1, both $\underline{D}(s, \cdot)$ and $\overline{D}(s, \cdot)$ are continuous and for all $(s, y) \in \mathcal{S} \times \mathcal{C}$, (i) $\underline{\underline{D}}(s, y) + \varepsilon_0 \leq \underline{D}(s, y) \leq \overline{D}(s, y)$, (ii) $\underline{D}(s, y) \leq \mathbf{B}(\underline{D})(s, y)$ and $\mathbf{B}(\overline{D})(s, y) \leq \overline{D}(s, y)$ and (iii) $\underline{\underline{D}}(s, y) + \varepsilon_1 < \mathbf{B}(\underline{\underline{D}})(s, y)$.*

We discuss the selection of bounding functions in the context of specific examples below. Note, however, if $\underline{\underline{D}}$ satisfies Assumption 3 (iii) and $\underline{\underline{D}} \leq D^* \leq \overline{D}$, then, from the monotonicity of $\mathbf{B}$ and Theorem 1, for all $(s, y) \in \mathcal{S} \times \mathcal{C}$,

$$\underline{\underline{D}}(s, y) + \varepsilon < \mathbf{B}(\underline{\underline{D}})(s, y) \leq \mathbf{B}(D^*)(s, y) = D^*(s, y) \leq \overline{D}(s, y).$$

Thus, if each $\mathbf{B}(\underline{\underline{D}})(s, \cdot)$ is continuous, then $\underline{D}$ may be set equal to $\mathbf{B}(\underline{\underline{D}})$. Given a triple of functions $\underline{\underline{D}}$, $\underline{D}$ and $\overline{D}$ satisfying Assumption 3, let:

$$\mathscr{G} = \{D : \mathcal{S} \times \mathcal{Y} \to \mathbb{R} | D \text{ is sub-linear and } \underline{D} \leq D \leq \overline{D}\}.$$

Define the "Thompson-like" metric $d : \mathscr{G} \times \mathscr{G} \to \mathbb{R}_+$ according to:

$$\begin{aligned} d(D_1, D_2) &= \sup_{\mathcal{S} \times \mathcal{C}} \left| \ln \left(\frac{D_1(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right) - \ln \left(\frac{D_2(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right) \right| \\ &\leq \sup_{\mathcal{S} \times \mathcal{C}} \ln \left(\frac{\overline{D}(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right) < \infty, \end{aligned}$$

where the finiteness stems from Assumption 3.[18] That $(\mathscr{G}, d)$ is complete metric space is shown next.

Lemma 2. *$(\mathscr{G}, d)$ is a complete metric space.*

Proof. See Appendix D. ☐

Proposition 6 verifies that $\mathbf{B}$ is contraction on $\mathscr{G}$. It relies on the concavity (and monotonicity) of $\mathbf{B}$ rather than any discounting-type conditions. This makes it well suited to the present setting where concavity of $\mathbf{B}$ is easy to show, but discounting (with respect to a suitable bounding norm) is not.

[18] In particular, it follows from $\overline{D}(s, y) \geq \underline{D}(s, y) \geq \underline{\underline{D}}(s, y) + \varepsilon_0$, the compactness of $\mathcal{C}$ and the continuity of the functions $\underline{\underline{D}}$, $\underline{D}$ and $\overline{D}$.

Proposition 6. *Let Assumption 3 hold. There is a $\rho \in [0,1)$ such that for all $D_1, D_2 \in \mathcal{G}$, $d(\mathbf{B}(D_1), \mathbf{B}(D_2)) \leq \rho d(D_1, D_2)$, i.e. $\mathbf{B}$ is a contraction on $(\mathcal{G}, d)$ with modulus of contraction ρ.*

Proof. See Appendix D. □

Application of the contraction mapping theorem yields that $\mathbf{B}$ admits a unique fixed point in $\mathcal{G}$.

Theorem 2. *Let Assumption 3 hold and assume that $\underline{D} \leq D^* \leq \overline{D}$. D^* is the unique fixed point of $\mathbf{B}$ in $\mathcal{G}$. Also, there is a $\rho \in [0,1)$ such that for any $D_0 \in \mathcal{G}$, $\mathbf{B}^n(D_0) \overset{d}{\to} D$ with $d(\mathbf{B}^n(D_0), D^*) \leq \rho^n d(D_0, D^*) \leq \rho^n d(\overline{D}, \underline{D})$.*

Proof. D^* is sub-linear by Lemma 1 and bounded below by $\underline{D}$ and above by $\overline{D}$ by assumption. Thus, $D^* \in \mathcal{G}$. Also by Lemma 1, if $D \in \mathcal{G}$, then $\mathbf{B}(D)$ is sub-linear and by the montonicity of $\mathbf{B}$ and Assumption 3 it is bounded below by $\underline{D}$ and above by $\overline{D}$. Thus, $\mathbf{B} : \mathcal{G} \to \mathcal{G}$. By Proposition 6, it is contractive on $\mathcal{G}$. The results in the theorem then stem from the contraction mapping theorem. □

Application of Theorem 2 requires bounding functions satisfying Assumption 3. Such functions are often easy to derive in the context of particular applications using actions and "large" values from the bounding set of state variables that strictly satisfy current constraints. The following examples illustrate.[19]

Example 1 (*Risk sharing with limited commitment and Epstein-Zin preferences*)**.** Let $\mathcal{V} = [\underline{v}, \overline{v}]^2$ and assume there is a resource-feasible consumption profile that gives each agent strictly more than autarky if combined with autarkic continuation payoffs and strictly less than the best possible payoff if combined with best possible continuation payoffs, i.e. an $\tilde{a} \in \mathcal{A}^{n_s}$ and a $\xi > 0$ such that for each $s \in \mathcal{S}$, $\gamma(s) > \sum_{i=0}^{1} \tilde{a}^i(s)$, and for each $s \in \mathcal{S}$ and $i \in \{0, 1\}$,

$$\overline{v} - \xi \geq \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta\overline{v} > \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta\Big\{\sum_{s' \in S} w^{i\prime}(s')^{\sigma} Q(s|s')\Big\}^{\frac{1}{\sigma}} > w^i(s) + \xi. \tag{20}$$

[19]For application of Theorem 2, it is sufficient to know (i) that bounding functions satisfying Assumption 3 exist and (ii) that a given function D_0 lies between them and can thus serve as an initial condition in a value iteration. Explicit calculation of the bounding functions is unnecessary. This contrasts with results relying on monotone (not contractive) operators, which require an upper or lower bound to the true value function as an initial condition. In addition, as always, the contraction result allows us to calculate error bounds and rates of convergence and is, thus, an improvement on results relying only on monotone iterations and pointwise convergence of iterates.

Set:

$$\overline{D}(s,q^V) = \sum_{i=0,1} q^{V,i}\phi^i(q^{V,i},s), \qquad \phi^i(q^{V,i},s) := \begin{cases} \overline{v} & q^{V,i} \geq 0 \\ \underline{v} & q^{V,i} < 0 \end{cases}$$

and

$$\underline{\underline{D}}(s,q^V) = \sum_{i=0,1} \{q^{V,i}\psi^i(q^{V,i},s) + |q^{V,i}|\xi\}, \qquad \psi^i(q^{V,i},s) := \begin{cases} w(s) & q^{V,i} \geq 0 \\ \overline{v} & q^{V,i} < 0. \end{cases}$$

It is easy to see that $\overline{D}$ is sub-linear, $\underline{\underline{D}}$ is continuous and positively homogenous and for all $(s,q^V) \in \mathcal{S} \times \mathcal{C}$, $\underline{\underline{D}}(s,q^V) < \overline{D}(s,q^V)$. In addition, for $\overline{v}$ large enough, these definitions and (20) also ensure $\underline{\underline{D}} < D^* \leq \overline{D}$. In Appendix D, we show that given (20), there exists an $\varepsilon > 0$ such that for all $(s,q^V) \in \mathcal{S} \times \mathcal{C}$, $\underline{\underline{D}}(s,q^V) + \varepsilon < \mathbf{B}(\underline{\underline{D}})(s,q^V)$. $\underline{D}$ may be set equal to $\mathbf{B}(\underline{\underline{D}})$ and the conditions of Assumption 3 are satisfied. □

Example 2 (*Default with capital accumulation*)**.** We give mild conditions that ensure the existence of bounding functions satisfying Assumption 3 for the default problem. To economize on space we do so only for the problem without shocks: $\gamma(k,s) \equiv \gamma(k)$ and $w(k,s) \equiv w(k)$. Assume a $\bar{k}$ such that $\gamma(\bar{k}) = \bar{k} > 0$. Let $\mathcal{V} = [\underline{v}^0, \overline{v}^0] \times [\underline{v}^1, \overline{v}^1]$, $[0, \bar{k}] \subset \mathcal{K}$ and $\mathcal{A} = \mathcal{A}^0 \times \mathcal{A}^1$, with $\mathcal{A}^i$ the action set of agent i. Suppose there is a small $\xi^V > 0$ and an $\tilde{a}^1 \in \mathcal{A}^1$ such that the following inequalities are satisfied:

$$\overline{v}^1 - \xi^V \geq f^1(\tilde{a}^1) + \delta\overline{v}^1 > f^1(\tilde{a}^1) + \delta w(\bar{k}) > f^1(\tilde{a}^1) + \delta w(0) \geq w(\bar{k}) > w(0) + \xi^V. \quad (21)$$

In addition, for some small $\xi^K > 0$, suppose that $\underline{a}^0 = -\tilde{a}^1 - \xi^K$ and $\overline{a}^0 = \gamma(\bar{k}) - \tilde{a}^1 - \xi^K$, are in $\mathcal{A}^0$ and, hence, feasible for agent 0. Note that negative values for a^0 are natural if agent 0 is a risk neutral lender. Assume also that for $\tilde{a}^0 \in \{\underline{a}^0, \overline{a}^0\}$, $\overline{v}^0 - \xi^V \geq f^0(\tilde{a}^0) + \delta\overline{v}^0 > f^0(\tilde{a}^0) + \delta\underline{v}^0 > \underline{v}^0 + \xi^V$. Let:

$$\psi^K(q^K) := \begin{cases} -\bar{k} & q^K \geq 0 \\ 0 & q^K < 0, \end{cases} \quad \psi^0(q^{V,0}) := \begin{cases} \underline{v}^0 & q^{V,0} \geq 0 \\ \overline{v}^0 & q^{V,0} < 0, \end{cases} \quad \psi^1(q^{V,1}) := \begin{cases} w(0) & q^{V,1} \geq 0 \\ \overline{v}^1 & q^{V,1} < 0. \end{cases}$$

In Appendix D we show that the following are valid bounding functions:

$$\overline{D}(q^K,q^V) = \sum_{i=0,1} q^{V,i}\phi^{V,i}(q^{V,i}) + q^K\psi^K(-q^K), \qquad \phi^{V,i}(q^{V,i}) := \begin{cases} \overline{v}^i & q^{V,i} \geq 0 \\ \underline{v}^i & q^{V,i} < 0, \end{cases}$$

$$\underline{\underline{D}}(q^K, q^V) = \sum_{i=0,1} \{q^{V,i}\psi^{V,i}(q^{V,i}) + |q^{V,i}|\xi^V\} + q^K\psi^K(q^K) - |q^K|\xi^K.$$

□

Example 3 (*Optimal monetary policy*). Let $\mathcal{V} = \prod_{i=0,1}[\underline{v}^i, \overline{v}^i]$ denote a set of possible government payoffs and inflation rates. Assume an $\tilde{a} \in \mathcal{A} = [\underline{a}, \overline{a}]$ and $\xi > 0$ such that:

$$\begin{pmatrix} \overline{v}^0 - \xi \\ \overline{v}^1 - \xi \end{pmatrix} \geq \begin{pmatrix} L(\tilde{a}, \kappa\tilde{a} + \delta\overline{v}^1) \\ \kappa\tilde{a} \end{pmatrix} + \delta \begin{pmatrix} \overline{v}^0 \\ \overline{v}^1 \end{pmatrix}$$
$$\geq \begin{pmatrix} L(\tilde{a}, \kappa\tilde{a} + \delta\underline{v}^1) \\ \kappa\tilde{a} \end{pmatrix} + \delta \begin{pmatrix} \underline{v}^0 \\ \underline{v}^1 \end{pmatrix} > \begin{pmatrix} \underline{v}^0 + \xi \\ \underline{v}^1 + \xi \end{pmatrix}. \tag{22}$$

It may be verified that:

$$\overline{D}(q^V) = \sum_{i=0}^{1} q^{V,i}\phi^i(q^{V,i}), \qquad \phi^i(q^{V,i}) = \begin{cases} \overline{v}^i & q^{V,i} \geq 0 \\ \underline{v}^i & q^{V,i} < 0, \end{cases}$$

$$\underline{\underline{D}}(q^V) = \sum_{i=0}^{1}\{q^{V,i}\psi^i(q^{V,i}) + |q^{V,i}|\xi\}, \qquad \psi^i(q^{V,i}) = \begin{cases} \underline{v}^i & q^{V,i} \geq 0 \\ \overline{v}^i & q^{V,i} < 0. \end{cases}$$

and $\underline{D} = \mathbf{B}(\underline{\underline{D}})$ satisfy all desired conditions.[20] □

7 Quasi-linearity in backward and forward state variables

Many problems have aggregators and constraint functions that are quasi-linear in k or v or both.[21] Exploiting this structure can lead to considerable simplification. Specifically, it is possible to work with the dual of the original rather than the augmented problem, i.e. (P) rather than (AP). Backward and forward primal states k_t and v_t are then removed from the analysis along with the equality constraints describing their evolution. In addition, the co-state variables q_t^V are no longer explicit choices (they are determined as functions of q^H multipliers). All of this simplifies optimizations. Below we describe the modified recursive dual problems that emerge, first for problems in which all laws of motion are quasi-linear in primal states and then, via an example, for those in which some are.

[20] The verification is similar to that given for the limited commitment case in Appendix D.

[21] For example, Messner, Pavoni, and Sleet (2012b) only considers simplified problems of this sort.

7.1 Fully quasi-linear problems

Assume that W^K is quasi-linear in k:

$$W^K[k, s, a] = A(s)k + B(s, a),$$

for some functions $A : \mathcal{S} \to \mathbb{R}^{n_k}_+$ and $B : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^{n_k}$. The functions K_{t+1} are defined to be consistent with this aggregator:

$$K_{t+1}(\bar{k}, \alpha | s^t) = \sum_{\tau=0}^{t} \prod_{j=\tau+1}^{t} A(s_j) B(s_\tau, a_\tau(s^\tau)) + \prod_{j=0}^{t} A(s_j)\bar{k}, \tag{23}$$

where in what follows it is useful to make the dependence of backward-looking states on the initial value $\bar{k} \in \mathcal{K} \subset \mathbb{R}^{n_k}$ explicit. The requirement that the K_t functions and $\mathcal{K}$ are bounded is no longer imposed.[22]

Assume that the shock transition is independent of any action: $Q : \mathcal{S} \to \mathcal{R}(S)$ and let $Q^t(s_0, s^t)$ denote the induced probability over history s^t given seed shock s_0. Turning to forward-looking states, assume that W^V is quasi-linear in m:

$$W^V[s, a, m] = f(s, a) + \delta m, \tag{24}$$

with $f : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^{n_v+1}$ a bounded function and δ a non-negative $n_v + 1$ diagonal matrix with elements bounded by $\bar{\delta}$. Assume that M^V is linear in v':

$$M^V[s, v'] = \sum_{s' \in \mathcal{S}} v'(s')Q(s, s'). \tag{25}$$

V is now defined to be consistent with these aggregators:

$$V(s_0, \alpha) = \sum_{t=0}^{\infty} \delta^t \sum_{s^t \in \mathcal{S}^t} f(s_t, a_t(s^t))Q^t(s_0, s^t).$$

The composition of W^V and M^V is quasi-linear in v'. The constraint function H is obtained from an $n_h \times n_k$ matrix N^K, a function $h : \mathcal{K} \times \mathcal{S} \times \mathcal{A} \to \mathbb{R}^{n_h}$ and a family of $n_h \times (n_v + 1)$ matrices $N^V(s, s')$:

$$H[k, s, a, v'] = N^K(s)k + h(s, a) + \sum_{s' \in \mathcal{S}} N^V(s, s')v'(s')Q(s, s').$$

[22] In many applications K_t is capital and $\mathcal{K} = \mathbb{R}^{n_k}_+$.

Finally, F is assumed to be linear in forward states v: $F[s,v] = q_0^V \cdot v$. Combining these assumptions, (P) becomes:

$$\sup q_0^V \cdot V(s_0, \alpha) \tag{QL-P}$$

subject to $\alpha \in \mathscr{A}$ and for all t, s^t,

$$N^K(s_t)K_t(\bar{k}, \alpha | s^{t-1}) + h(s_t, a_t(s^t)) + \sum_{s' \in \mathcal{S}} N^V(s,s')V(s', \alpha | s^t, s')Q(s_t, s') \geq 0.$$

Various problems satisfy these types of assumptions.

Example 2 (*Default with capital accumulation, AK version*)**.** We specialize Example 2 to place it within the quasilinear framework. Since $W^K[k,s,a] = \gamma(k,s) - \sum_{i=0,1} a^i$, quasi-linearity of W^K in k follows if the production function is specialized to $\gamma(k,s) = \gamma(s)k$. Quasi-linearity of the incentive ("no default") constraint in k and v follows if the default value is given by $w(k,s) = w(s)k$.[23] Then our general notation specializes to: $A(s) = \gamma(s)$, $B(a) = -\sum_{i=0,1} a^i$,

$$N^K(s) = \begin{pmatrix} -w(s) \\ \gamma(s) \end{pmatrix}; \quad h(a) = \begin{pmatrix} f(a^1) \\ -\sum_{i=0,1} a^i \end{pmatrix}; \quad \text{and} \quad N^V = \begin{pmatrix} \delta \\ 0 \end{pmatrix}.$$

□

Note that a small modification of the assumptions in Example 2 that removes capital and incorporates an incentive constraint for agent 0 gives the limited commitment model of Kocherlakota (1996). On the other hand, the removal of the incentive constraints gives a standard AK growth model.

Since (QL-P) incorporates only the H-function constraints, the constraint process is given simply by $\zeta^H(k,\alpha) = \{z_t^H(k,\alpha)\}_{t=0}^{\infty}$, with

$$z_t^H(k,\alpha)(s^t) = N^K(s_t)K_t(k, \alpha | s^t) + h(s_t, a_t(s^t)) + \sum_{s' \in \mathcal{S}} N^V(s,s')V(s', \alpha | s^t, s')Q(s_t, s').$$

Because neither the K_t functions nor the h function need be bounded, constraint processes need not belong to the normed space ℓ_∞. However, we assume that B, A, h and $\mathscr{A}$ are such that for all $k \in \mathcal{K}$ and $\alpha \in \mathscr{A}$, $\zeta^H(k,\alpha)$ satisfies $\sup_{t,s^t} \|z_t^H(k,\alpha)(s^t)\| / M_t(s^t) < \infty$ for some positive-valued process $\{M_t\}$.[24] Consequently, the constraint process normalized by

[23] This assumption on default values is made in Cooley, Marimon, and Quadrini (2004).

[24] In bounded problems with $\bar{A} := \max_{\mathcal{S}} A(s) \leq 1$ and h bounded, the natural candidate for $\{M_t\}$ is the constant process $\forall t, s^t$, $M_t(s^t) = 1$. In growth models with positive valued A process and $\bar{A} := \max_{\mathcal{S}} A(s) > 1$, the natural candidate is $\forall t, s^t$, $M_t(s^t) = A^t(s^t) := \prod_{j=0}^t A(s_j)$.

$\{M_t\}$ is in ℓ_∞.

The Lagrangian

$$\mathcal{L}(\alpha, \theta^H) = q_0^V \cdot V(s_0, \alpha) + \langle \theta^H, \zeta^H(\bar{k}, \alpha) \rangle$$

allows the following dual problem to be associated with (QL-P):

$$D_0^* = \inf_{\mathscr{Q}} \sup_{\mathscr{A}} \mathcal{L}(\alpha, \theta^H), \tag{26}$$

with $\mathscr{Q} = \{\theta^H | q_t^H(s^t) \geq 0, \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} q_t^H(s^t) M_t(s^t) < \infty\}$. Crucially, the quasi-linearity of the aggregators ensures that this Lagrangian has the necessary separability for a recursive dual approach.

As a first step, we recover the continuation dual problem from (26). The definition of ζ^H, (23) and straightforward algebra implies that the dual problem (26) can be rewritten as:

$$\begin{aligned} D_0^* &= \inf_{\mathscr{Q}} \sup_{\mathscr{A}} \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} q_t^H(s^t) \cdot N^K(s_t) A^t(s^t) \cdot \bar{k} + q_0^V \cdot V(s_0, \alpha) + \langle \theta^H, \zeta^H(0, \alpha) \rangle \\ &= \inf_{\mathbb{R}^{n_k}} \left\{ q_0^K \cdot \bar{k} + \inf_{\mathscr{Q}|q_0^K} \sup_{\mathscr{A}} q_0^V \cdot V(s_0, \alpha) + \langle \theta^H, \zeta^H(0, \alpha) \rangle \right\}, \end{aligned} \tag{27}$$

where $\mathscr{Q}|q_0^K = \{\theta^H \in \mathscr{Q} | q_0^K := \mathcal{T}(\theta^H)\}$ and $\mathcal{T}$ is the linear map $\mathcal{T}(\theta^H) = \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} \{q_t^H(s^t) \cdot N^K(s_t) A^t(s^t)\}$. In the first line of (27) terms involving $\bar{k}$ are factored out of the Lagrangian, while in the second line the infimum over dual processes is broken into two steps: an infimum over a co-state for the backward-looking state followed by a (more) constrained infimum over dual processes. Equation (27) motivates the following choice of continuation dual problem, for each $y \in \mathcal{Y} = \{(q^K, q^V) \in \mathbb{R}^{n_k} \times \mathbb{R}^{n_v+1} : \mathscr{Q}|q^K \neq \emptyset\}$,

$$D^*(s, y) = \inf_{\mathscr{Q}|q^K} \sup_{\mathscr{A}} q^V \cdot V(s, \alpha) + \langle \theta^H, \zeta^H(0, \alpha) \rangle. \tag{28}$$

Now the dual co-state space $\mathcal{Y} = \{(q^K, q^V) \in \mathbb{R}^{n_k} \times \mathbb{R}^{n_v+1} : \mathscr{Q}|q^K \neq \emptyset\} = \{q^K \in \mathbb{R}^{n_k} : \mathscr{Q}|q^K \neq \emptyset\} \times \mathbb{R}^{n_v+1}$ may be a proper subset of $\mathbb{R}^{n_k} \times \mathbb{R}^{n_v+1}$. In particular, to guarantee a continuation dual problem with a non-empty constraint set q^K must be in the range of the linear map $\mathcal{T}$ on $\mathscr{Q}$, i.e. there must be a non-negative valued process $\theta^H \in \mathscr{Q}$ such that $q^K = \mathcal{T}(\theta^H)$. Since $\mathscr{Q}$ is a cone and $\mathcal{T}$ is linear, the range of $\mathcal{R}$ is also a cone and in several relevant applications is easy to find. [25]

We now turn to the recursive form of (28). This is modified in several ways from previous sections. First, the terms $\{k_t, v_t\}$ are substituted out of the problem; second, the co-state variable q^V is no longer chosen directly, rather it evolves as a function of initial values and accumulated multipliers q^H. On the other hand, in general, $q^{K\prime}$ must still be picked: it is a forward-looking variable and is not (generally) determined by past q^H multipliers. Define the current dual correspondence $\mathcal{Q} : \mathbb{R}^{n_k} \to 2^{\mathbb{R}^{n_h+n_s n_k}}$,

$$\mathcal{Q}(q^K) = \left\{ (q^H, q^{K\prime}) \in \mathbb{R}^{n_h}_+ \times \mathbb{R}^{n_k \times n_s} \middle| q^K = q^H \cdot N^K(s) + A(s) \sum_{\mathcal{S}} q^{K\prime}(s') \right\},$$

and the current dual objective $J : \mathcal{S} \times \mathbb{R}^{n_v+1} \times \mathbb{R}^{n_h}_+ \times \mathbb{R}^{n_k \times n_s} \times \mathcal{A} \to \mathbb{R}$,

$$J(s, q^V; q^H, q^{K\prime}, a) = q^V \cdot f(s, a) + \sum_{s' \in \mathcal{S}} q^{K\prime}(s') \cdot B(s, a) + q^H \cdot h(s, a).$$

Finally, define the law of motion for co-states q^V:

$$\phi(s, q^V; q^H)(s') = \frac{1}{\bar{\delta}} \left\{ \delta \cdot q^V + q^H N^V(s, s') \right\}.$$

The recursive dual problem for this case is described in the following proposition.

Proposition 7 (Value functions)**.** *The value function D^*_0 satisfies:*

$$D^*_0 = \inf_{\mathbb{R}^{n_k}} D^*(s_0, q^K, q^V_0) + q^K \cdot \bar{k}. \tag{29}$$

with for all $(s, q^K, q^V) \in \mathcal{S} \times \mathbb{R}^{n_k}_+ \times \mathbb{R}^{n_v+1}$,

$$D^*(s, q^K, q^V) = \inf_{\mathcal{Q}(q^K)} \sup_{\mathcal{A}} J(s, q^V; q^H, q^{K\prime}, a) + \bar{\delta} \sum_{s' \in \mathcal{S}} D^*(s', q^{K\prime}(s'), \phi(s, q^V; q^H)(s')) Q(s, s'). \tag{30}$$

The proof is essentially the same as Proposition 3 and is omitted. Notice that in (29) the initial condition for the costate q^K is picked, whilst that for q^V is pinned down by the parameter q^V_0;[26] (30) gives the dual Bellman. Comparison of Propositions 3 and 7 and

[25]For example, in AK growth models, for all s, $N^K(s) = 1$, and $\mathcal{T}(\mathscr{Q}) = \mathbb{R}^{n_k}_+$. In limited commitment models without capital, for all s, $N^K(s) = 0$ and $\mathcal{T}(\mathscr{Q}) = \{0\}$ (and backward-looking state variables and their co-states may be omitted). In Example 2, $N^K(s) = (-w(s) \quad \gamma(s))$, $\mathcal{T}(\mathscr{Q}) = \mathbb{R}^{n_k}$ and $\mathcal{Y} = \mathbb{R}^{n_k+n_v+1}$ once more. The function D^* remains sub-linear. Hence, for practical purposes the effective state space can be identified with $\mathcal{C} \cap \mathcal{T}(\mathscr{Q})$.

[26]Thus, the co-state q^K for the backward-looking state k is forward-looking and the the co-state q^V for the

the terms defining the Bellman in each (for example, comparison of the corresponding J functions) reveals how exploitation of quasilinearity simplifies matters.

Example 2 (*Default with capital accumulation, AK version*)**.** To make the preceding discussion concrete, consider again the default model with linear production. Applying (30), the dual Bellman is:

$$\begin{aligned} D^*(s, q^K, q^V) \quad = \quad & \inf_{\mathcal{Q}(q^K)} \sup_{\mathcal{A}} \; q^{V,0} f^0(a^0) + (q^{V,1} + q^{H,1}) f^1(a^1) - \Big(\sum_{s' \in \mathcal{S}} q^{K\prime}(s') + q^{H,2} \Big) \sum_{i=0,1} a^i \\ & + \delta \sum_{s' \in \mathcal{S}} D^*(s', q^{K\prime}(s'), \phi(s, q^V; q^H)(s')) Q(s, s'). \end{aligned}$$

with $\mathcal{Q}(q^K) = \{(q^H, q^{K\prime}) | q^K = -q^{H,1} w(s) + \left(q^{H,2} + \sum_{\mathcal{S}} q^{K\prime}(s')\right) \gamma(s)\}$ and

$$\phi(s, q^V; q^H)(s') = \begin{pmatrix} q^{V,0} \\ q^{V,1} + q^{H,1} \end{pmatrix}.$$

Thus, the weights on the borrower's current utility $f(a^1)$ and, via the updating function ϕ, future utility are augmented by the multiplier on her incentive constraint $q^{H,1}$. The constraint set $\mathcal{Q}(q^K)$ reveals the evolution of the co-state q^K, the shadow value of capital. This value is depressed to the extent that capital tightens the incentive constraint $-q^{H,1} w(s)$, but enhanced to the extent that capital relaxes the current resource constraint or augments the future capital stock $(q^{H,2} + \sum_{\mathcal{S}} q^{K\prime}(s'))\gamma(s)$. $\square$

Remark 4. In Example 2, the weight on the lender's payoffs $q^{V,0}$ remains constant at its initial value. This is typical of (principal agent) problems in which one forward-looking state variable v^0 does not enter the H-constraint. For these problems, the corresponding co-state $q^{V,0}$ may be removed as an explicit state variable and the state space reduced in dimension from $\mathbb{R}^{n_k+n_v+1}$ to $\mathbb{R}^{n_k+n_v}$. The cost is that positive homogeneity of the value function is lost.

7.2 Partially quasi-linear problems

The preceding analysis extends to problems in which aggregators are quasi-linear in a subset of state variables. Rather than developing this in full, we describe the application to Example 3 (the optimal monetary policy problem). Recall that in this example, there are

forward-looking state v is backward-looking.

only forward-looking states and the aggregator W^V is given by:

$$\begin{pmatrix} v^0 \\ v^1 \end{pmatrix} = W^V\left[a, \begin{pmatrix} v^{0'} \\ v^{1'} \end{pmatrix}\right] = \begin{pmatrix} L(a, \kappa a + \delta v^{1\prime}) + \delta v^{0\prime} \\ \kappa a + \delta v^{1\prime} \end{pmatrix},$$

with v^0 the government's payoff, v^1 inflation and a output. The forward-looking state describing the government's future payoff $v^{0\prime}$ enters W^V in a quasi-linear way and can be substituted out. In contrast, the forward-looking state describing inflation $v^{1\prime}$ enters non-linearly and cannot be so removed. After substitution of v^0, the problem becomes:

$$\sup \sum_{t=0}^{\infty} \delta^t L(a_t, \kappa a_t + \delta v^1_{t+1})$$

subject to, for all t, $v^1_t = \kappa a_t + \delta v^1_{t+1}$. This leads to the dual problem:

$$D^*_0 = \inf_{\mathscr{Q}} \sup_{\mathscr{P}} \sum_{t=0}^{\infty} \delta^t L(a_t, \kappa a_t + \delta v^1_{t+1}) + \sum_{t=0}^{\infty} \delta^t q^{V,1}_t \{\kappa a_t + \delta v^1_{t+1} - v^1_t\}, \tag{31}$$

where $\mathscr{Q}$ is the set of inflation co-state sequences $\{q^{V,1}_t\}$ and $\mathscr{P}$ the set of inflation-output sequences $\{v^1_t, a_t\}_{t=0}^{\infty}$. Notice that in (31) the co-state on the government's payoff $q^{V,0}$ is initialized to and remains at 1. This is a principal-agent type problem. Using arguments similar to before the initial problem specializes to:

$$D^*_0 = \inf_{\mathbb{R}} \sup_{\mathcal{V}^1} \quad -q^{V,1}_0 v^1_0 + D^*(1, q^{V,1}_0),$$

where $\mathcal{V}^1 = \frac{\kappa}{1-\delta}[\underline{a}, \overline{a}]$ is the set of possible inflations, while the dual Bellman equation becomes:

$$D^*(1, q^{V,1}) = \inf_{\mathbb{R}} \sup_{\mathcal{A}\times\mathcal{V}^1} \quad L(a, \kappa a + \delta v^{1\prime}) + q^{V,1}(\kappa a + \delta v^{1\prime}) - q^{V,1\prime} v^{1\prime} + \delta D^*(1, q^{V,1\prime}).$$

In the latter the inner supremum is over current output-inflation pairs (a, v^1), while the infimum operation is over the future inflation co-state $q^{V,1\prime}$.

Quadratic Case If (the negative of) the loss function L is specialized to be quadratic, an explicit closed form solution of the dual problem is available. Let $\underline{a} = 0$, and

$$L(x, z) = -\frac{1}{2}\left\{x^2 + \lambda z^2\right\},$$

with $\lambda > 0$. Then, a standard 'guess and verify' exercise confirms that the value function for this problem satisfies:

$$D^*(1, q^{V,1}) = \frac{1}{2\chi}(\max\{0, q^{V,1}\})^2,$$

with $\chi > 0$ the positive root of a quadratic equation.[27] Optimal dual policy functions are then easily obtained. For $q^{V,1} \geq 0$, they are linear in the dual co-state and are given by:

$$\begin{pmatrix} q^{V\prime}(q^{V,1}) \\ a(q^{V,1}) \\ \pi'(q^{V,1}) \end{pmatrix} = \begin{pmatrix} 1 \\ \frac{\kappa}{\lambda} \\ \frac{1}{\chi} \end{pmatrix} \xi q^{V,1},$$

with $\xi = \frac{1}{1+\frac{\kappa^2}{\lambda}+\frac{\delta}{\chi}} \in (0,1)$. For $q^{V,1} < 0$, $q^{V,1\prime}(q^{V,1}) = q^{V,1}$, and $a(q^{V,1}) = v^1(q^{V,1}) = 0$. The problem is strictly concave and policies are single valued. Hence, from Proposition 9 below and the subsequent discussion, the solution to the dual problem delivers necessary and sufficient conditions for the solution to the original problem.

8 Relating Primal and Dual

The preceding sections established that optimal dual values and solutions may be recovered from the recursive dual. They also showed that the dual Bellman operator was contractive. Consequently, if the dual supplies an optimal value and optimal policies for the original primal problem, then the recursive dual does as well and the primal may be solved via dual value iteration. In this section, we briefly discuss conditions for the sequential dual and primal problems to have common values and policies.

8.1 Saddles and recursive dual policies

Without further restriction, classical weak duality implies that the optimal dual value bounds the optimal primal value: $D_0^* \geq P_0^*$. Thus, with no further assumptions the recursive dual gives welfare bounds for optimal policies or policy improvements.

A well known sufficient condition for equality of optimal values, albeit not on prim-

[27]It solves: $\chi = \frac{1-\delta+\frac{\kappa^2}{\lambda}+\sqrt{\left(1-\delta+\frac{\kappa^2}{\lambda}\right)^2+4\frac{\kappa^2}{\lambda}\delta}}{2\frac{\kappa^2}{\lambda}}$.

itives, is that the Lagrangian admits a saddle point.[28] Saddle existence also ensures that the dual policy set includes all primal solutions. However, the converse is not true: additional restrictions are required to ensure that the "finite penalization" implicit in the dual problem is "sharp enough" to pin down only primal solutions. The following propositions summarize the situation for the general problem considered in Section 5.[29]

Proposition 8 (Policy functions; Necessity)**.** *Assume that the Lagrangian $\mathscr{L}$ admits a saddle point. Then: (i) (Equality of values) $D_0^* = P_0^*$. (ii) (Necessity conditions for policies) If π^* solves* (AP)*, then there is a corresponding optimal dual sequence θ^* such that $(q_0^{K*}, q_0^{V*}, v_0^*) \in G_0^{IS}$ and for all t and s^t, $(q_t^{H*}(s^t), y_{t+1}^*(s^t), p_t^*(s^t)) \in G_t^{IS}(s_t, y_t^*(s^t))$.*

Proof. See Appendix E. ☐

Proposition 8 only requires that $\mathscr{L}$ admits a saddle point. It does *not* require that the Lagrangian associated with every $(s^t, y_t(s^t))$-continuation problem has a saddle, as is the case in Marcet and Marimon (2011). Proving, or numerically checking, the existence of a saddle for $\mathscr{L}$, while non-trivial, is less demanding than doing so for all possible histories.

Sufficiency of (recursive) dual policies for primal attainment requires additional assumptions. We say that a set of primal processes $\mathscr{P}'$ shares a plan α if for each $\pi \in \mathscr{P}'$ there is a process $\{v_t, k_t\}_{t=0}^{\infty}$ such that $\pi = (\alpha, \{v_t, k_t\}_{t=0}^{\infty})$.

Proposition 9 (Policy functions; Sufficiency)**.** *Assume that $\mathscr{Q}^* \times \mathscr{P}^*$, the set of saddle points of $\mathscr{L}$, is non-empty and that for each $\theta^* \in \mathscr{Q}^*$, the set of primal processes $\operatorname{argmax}_{\mathscr{P}} \mathscr{L}(\cdot, \theta^*)$ shares a plan α^*. Then: (i) α^* is the unique solution of* (P) *and (ii) if a pair (π, θ) with $\pi = (\alpha, \{k_t, v_t\}_{t=0}^{\infty})$ satisfies $(q_0^K, q_0^V, v_0) \in G_0^{IS}$, for all t and s^t, $(q_t^H(s^t), y_{t+1}(s^t), p_t(s^t)) \in G_t^{IS}(s_t, y_t(s^t))$ and* (T)*, then α solves* (P) *(and equals α^*).*

Proof. See Appendix E. ☐

We apply Proposition 9 in Section 9 to show that recursive dual policies are sufficient for primal optimality in a parameterized version of Example 1.

[28]For a real-valued function defined on a product set $g : C \times E \to \mathbb{R}$, the set of saddle points is:

$$\underset{C|E}{\operatorname{saddle}}\, g = \left\{ (c^*, e^*) \,\middle|\, c^* \in \underset{C}{\operatorname{argmin}}\, g(c, e^*) \text{ and } e^* \in \underset{E}{\operatorname{argmax}}\, g(c^*, e) \right\}.$$

[29]Similar results hold for the quasi-linear case considered in the preceding section (see Messner, Pavoni, and Sleet (2011), Section 3 for details).

8.2 Concave Problems

The literature gives various sufficient conditions on primitives ensuring equality of optimal values and saddle existence.[30] Consider for a moment the original (non-augmented) optimization (P). This problem omits laws of motion for primal states and is written entirely in terms of plans rather than primal processes. The constraints may be collected together as:

$$\zeta^H(\alpha) = \{H(K_t(\alpha, s^t), a_t(s^t), V(\alpha, s^t))\}_{t,s^t} \geq 0.$$

Since there are a countable number of constraints and H is bounded, $\zeta^H : \mathscr{A} \to \ell_\infty$. We may associate a Lagrangian $\tilde{\mathcal{L}}$ with this problem:

$$\tilde{\mathcal{L}}(\alpha, \theta^H) = F[s_0, V(s_0, \alpha)] + \langle q^H, \zeta^H(\alpha)\rangle, \tag{32}$$

where q^H belongs to $\ell_\infty^\star$, the dual space of ℓ^∞ and $\langle q^H, \zeta^H(\alpha)\rangle$ is the evaluation of q^H at $\zeta^H(\alpha)$.[31] The Lagrangian $\tilde{\mathcal{L}}$ can be used to define primal and dual problems for (P) directly. A well known sufficient condition for these problems to have equality of optimal values and a minimizing dual multiplier (so called "strong duality") is that (i) the objective and constraints are concave and (ii) the evaluation of the constraints at some primal choice lies in the interior of the constraint space's closed non-negative cone (a Slater condition).[32] If, in addition, a solution to (P) exists then it and the minimizing multiplier constitute a saddle point. Since the objective and constraints are constructed from compositions of functions, a standard assumption guaranteeing concavity is that F, H, W^K, W^V and M^V are jointly concave in their arguments and either quasi-linear or non-decreasing in the primal states k and v'. Stronger strict concavity restrictions ensure uniqueness of the primal solution and sufficiency of the dual solution for primal optimality.

A difficulty is that the preceding result guarantees the existence of a minimizing multiplier q^H in $\ell_\infty^\star$. It is much more convenient to establish such existence in $\ell_1 \subset \ell_\infty^\star$, the space of summable sequences $\{\{q_t^H\} : \sum_{t=0}^\infty \sum_{\mathcal{S}^t} \|q_t^H(s^t)\| < \infty\}$, and, hence, to obtain existence of a saddle point of the Lagrangian:

$$\mathcal{L}(\alpha, \theta^H) = F[s_0, V(s_0, \alpha)] + \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} q_t^H(s^t) \cdot H(K_t(\alpha, s^t), a_t(s^t), V(\alpha, s^t)). \tag{33}$$

[30] Luenberger (1969) and Rockafellar (1974), especially Section 7, are good references for the theory in infinite dimensional settings.

[31] We discuss $\ell_\infty^\star$ briefly below. It is the space of bounded continuous functionals on ℓ_∞ and, as is well known, equals the space of all signed charges of bounded variation on the power set $2^{\mathbb{N}}$.

[32] And the constraint space, i.e. the codomain of ζ^H, has a non-negative cone with non-empty interior because it is the set ℓ_∞.

In fact, following an argument of Ponstein (1981), the structure of the constraints enables us to do this and, hence, obtain saddle existence for $\mathcal{L}$ rather than $\tilde{\mathcal{L}}$ under the conditions given above.[33]

Such saddle existence results are directly applicable to the quasi-linear case discussed in Section 7.[34] However, as noted, for more general problems with laws of motion that are non-linear in states $\mathcal{L}$ is not suitable for dual recursive decomposition. Instead the Lagrangian $\mathscr{L}$ from the augmented problem is needed. An apparent difficulty is that $\mathscr{L}$ incorporates (possibly non-linear) equality constraints for the laws of motion for states. Thus, standard conditions for saddle existence are not applicable.[35] If, however, H is non-decreasing in k and v', then the equality constraints can be relaxed to inequalities. The relaxation does not modify optimal values or solutions. Strong duality for the relaxed problem is then established under the standard concavity and monotonicity assumptions on F, H, W^K, W^V and M^V described previously. For more details on relaxation including weaker conditions for its validity, see Appendix F.[36]

8.3 Ex Post Check

The following elementary proposition gives a sufficient condition for primal optimality in terms of the optimal dual value. Importantly, the condition does not rely on any concavity assumption on the problem. We call a process $(\hat{\pi}, \hat{\theta}) = (\hat{\alpha}, \{\hat{k}_t, \hat{v}_t\}_{t=0}^{\infty}, \hat{\theta})$ *a candidate plan* if it is obtained from the policy correspondence: $(q_0^K, q_0^V, v_0) \in G_0^{IS}$, and $\forall$ t and s^t, $(q_t^H(s^t), y_{t+1}(s^t), p_t(s^t)) \in G_t^{IS}(s_t, y_t(s^t))$.

Proposition 10. *Suppose a candidate plan $(\hat{\pi}, \hat{\theta})$ satisfies: (i) $F[s_0, V(s_0, \hat{\alpha})] \geq D_0^*$ and (ii) $\hat{\pi}$ is feasible for* (AP)*. Then $\hat{\pi}$ is optimal for* (AP) *and $D_0^* = P_0^*$. If in addition to (i)-(ii), $(\hat{\pi}, \hat{\theta})$ satisfies condition* (T)*, then $(\hat{\pi}, \hat{\theta})$ is a saddle for the Lagrangian associated with problem* (AP)*.*

Proof. See Appendix E. □

Despite its simplicity, Proposition 10 is the basis of a useful ex post check of primal optimality. Suppose the recursive dual problem has been solved and a fixed point $\hat{D}$ of the operator $\mathbf{B}$ obtained. If the conditions of Assumption 3 hold and $\hat{D}$ lies between the

[33] We defer this technical argument to Messner, Pavoni, and Sleet (2013).

[34] With the slight modification that the constraint space is set to $\{\{x_t\} : \sup \|x_t(s^t)\| / M_t(s^t) < \infty\}$ and the multiplier space to $\{\{q_t^H\} : \sum_{t=0}^{\infty} \sum_{\mathcal{S}^t} \|q_t^H(s^t) M_t(s^t)\| < \infty\}$ to accommodate unbounded H functions.

[35] Even if the constraints stemming from the laws of motion are re-expressed as pairs of inequalities, the Slater condition and, unless these laws of motion are linear, concavity is lost.

[36] The dual and recursive dual of the relaxed problem are slightly modified to restrict co-states to be non-negative.

bounding functions $\underline{D}$ and $\overline{D}$, then $\hat{D} = D^*$. Consequently, the value D_0^* and a candidate plan $\hat{\pi}$ may be recovered. Proposition 10 then provides sufficient conditions for $\hat{\pi}$ to be a solution to (AP) and for the existence of a saddle point of the associated Lagrangian.

In practice, D_0^*, G_0^{IS} and G^{IS} must be approximated via, say, a numerical implementation of the value iteration described in Theorem 2, and the conditions (i) and (ii) in Proposition 10 checked numerically to within some acceptable level of tolerance. We describe a numerical implementation of the value iteration next.

9 Numerical Implementation

We use Example 1 (limited commitment risk sharing) to illustrate the numerical implementation of the recursive dual problem.

Applicability of the Recursive Dual If the agent's initial Pareto weights are non-negative, $\lambda = (\lambda^1, \lambda^2) \in \mathbb{R}^2_+$, then without loss of generality the law of motion for utility promises may be relaxed to an inequality in either the augmented primal problem or its dual:

$$\frac{1-\delta}{1-\mu}(a_t^i(s^t))^{1-\mu} + \delta\Big(\sum_{s' \in \mathcal{S}} (v_{t+1}^i(s^t, s'))^{\sigma} Q(s_t|s')\Big)^{\frac{1}{\sigma}} - v_t^i(s_t) \geq 0.$$

Given $\mu, \sigma \in (0,1)$, the functions describing these constraints (and the limited commitment constraints) are strictly concave. If a primal solution exists and there is a primal process strictly satisfying all constraints, then standard results and an argument of Ponstein (1981), establishes the existence of a saddle point for the Lagrangian $\mathscr{L}$ with co-states restricted to be non-negative. Consequently, by our previous results, the recursive dual gives the optimal primal value and necessary and sufficient conditions for optimal primal policies.

Value iteration and function approximation We limit attention to value functions defined on a domain of shocks and non-negative co-states ("Pareto weights"), $\mathcal{S} \times \mathbb{R}^2_+$ and modify definitions accordingly.[37] The bounding value functions $\underline{\underline{D}}$, $\underline{D}$ and $\overline{D}$ are as in Section 6, but restricted to this domain. The definitions of $\mathscr{G}$ and $\mathbf{B}$ become:

$$\mathscr{G} = \{D : \mathcal{S} \times \mathbb{R}^2_+ \to \mathbb{R} | D \text{ is sublinear, each } D(s, \cdot) \text{ is continuous and } \underline{D} \leq D \leq \overline{D}\}.$$

[37] This restricts us to concave and economically interesting continuation problems in which no agent gets a negative weight.

and

$$\mathbf{B}(D)(s,y) = \inf_{\mathcal{Q}_+} \sup_{\mathcal{P}} J(s,y;q,p) + \delta \sum_{s' \in \mathcal{S}} D(s', y'(s')),$$

with $\mathcal{Q}_+ = \mathbb{R}^4_+$ replacing $\mathbb{R}^2_+ \times \mathbb{R}^2$ and J as in (16). Theorem 2, very slightly modified to incorporate the domain restriction, ensures that D^* may be calculated via an iteration of $\mathbf{B}$ from any $D_0 \in \mathscr{G}$. Implementation of this iteration requires approximation of the value functions. Our approximation procedure exploits the sub-linearity of dual value functions.[38]

If $g : \mathbb{R}^2_+ \to \mathbb{R}$ is sub-linear, then for all $y \in \mathbb{R}^2_+$,

$$g(y) = \max\{m \cdot y | \forall y' \in \mathcal{C}_+, m \cdot y' \leq g(y')\}.$$

Let $\hat{\mathcal{C}}^I_+ := \{y_i\}^I_{i=1} \subset \mathcal{C}_+$ denote a set of $I > 1$ distinct points in $\mathcal{C}_+$. Then g is bounded above by $\hat{g}^I$, where $\hat{g}^I$ is defined by the less restricted problem:

$$g(y) \leq \hat{g}^I(y) := \max\{m \cdot y | \forall y_i \in \hat{\mathcal{C}}^I_+, m \cdot y_i \leq g(y_i)\}. \tag{34}$$

The function $\hat{g}^I$ is continuous and sub-linear and $\hat{g}^I(y^i) = g(y^i)$ at each $y^i \in \mathcal{C}^I_+$. In addition, $\hat{g}^I(y)$ is easily found by solving the simple linear programming problem in (34). A sequence of sets $\hat{\mathcal{C}}^I_+$, $I = 2, 3, \ldots$, may be constructed with $\hat{\mathcal{C}}^I_+ \subset \hat{\mathcal{C}}^{I+1}_+$ and $\hat{\mathcal{C}}^\infty_+ = \cup_I \hat{\mathcal{C}}^I_+$ dense in $\mathcal{C}_+$.[39] If g is also continuous, then it is readily verified that the corresponding sequence of approximating functions $\hat{g}^I$ converges pointwise to g from above.[40] Moreover, by Dini's theorem it converges uniformly on $\mathcal{C}_+$ and, hence, in the Thompson-like metric d to g.

This procedure may be used to approximate sub-linear functions $D \in \mathscr{G}$ from above. It is easy to implement, may be integrated into the value iteration and involves approximation on a simple state space.

Remark 5. As we have previously remarked Example 1 is outside the scope of Marcet

[38] For facts about sub-linear functions used below consult Florenzano and Van (2001)

[39] For example, the set of points in $\mathcal{C}_+$ with rational coordinates is dense in $\mathcal{C}$, see Schmutz (2008) for an explicit construction.

[40] It clearly converges at all points in $\hat{\mathcal{C}}^\infty$ and if $(1,0)$ and $(0,1)$ are in $\hat{\mathcal{C}}^\infty_+$ at these two points. Choose a point $y \in \mathcal{C}_+ \cap$ int $\mathbb{R}^2_+$. Let y^1_n and y^2_n be two sequences in $\cup_I \hat{\mathcal{C}}^I_+$ converging to y and such that $y = \lambda_n a_n y^1_n + (1-\lambda_n) b_n y^2_n$, with $\lambda_n \in (0,1)$, $a_n, b_n \in \mathbb{R}_+$ and $a_n, b_n \downarrow 1$, i.e. $a_n y^1_n$ and $b_n y^2_n$ lie either side of y on the tangent to $\mathcal{C}_+$ passing through y. There is a sequence $\{I_n\}$ such that $\hat{g}^{I_n}(y^1_n) = g(y^1_n)$ and $\hat{g}^{I_n}(y^2_n) = g(y^2_n)$. By the sub-linearity of g and each $\hat{g}^{I_n}$, we have $g(y) \leq g^{I_n}(y) \leq \lambda_n \hat{g}^{I_n}(a_n y^1_n) + (1-\lambda_n)\hat{g}^{I_n}(b_n y^2_n) = \lambda_n a_n \hat{g}^{I_n}(y^1_n) + (1-\lambda_n) b_n \hat{g}^{I_n}(y^2_n) = \lambda_n a_n g(y^1_n) + (1-\lambda_n) b_n g(y^2_n)$. Since g is continuous, $y^i_n \to y$ and $a_n, b_n \downarrow to 1$, it follows that the last term in the string of inequalities converges to $g(y)$. Thus, the sequence of functions converges pointwise on $\mathcal{C}_+$ and by the positive homogeneity of the functions on $\mathbb{R}^2_+$.

and Marimon (2011). Recursive primal formulations of limited commitment problems are available. Such problems (without Epstein-Zin preferences) have often been handled numerically by maximizing the payoff to one player subject to incentive and utility promise-keeping constraints. This leads to a Bellman-type operator in which the continuation value function of the maximized player enters the constraint set. Although this Bellman operator has monotonicity properties it is not a contraction. Such a problem also involves an endogenous state space of utility promises. With only two players, this state space is an interval and, therefore, easy to approximate. But as more players are introduced representation and approximation of this set becomes more difficult.

Our function approximation procedure is similar to Judd, Yeltekin, and Conklin (2003)'s outer approximation method. They use piecewise linear approximations to support functions of payoff sets in their analysis of repeated games.

Numerical Example Figure 1 illustrates value and policy functions from a numerical example. In the example, the discount factor δ is set to 0.8. The preference parameters μ and σ are set to 0.5 and 0.8 respectively. Two shock states are assumed with Markov transition $Q(1,1) = Q(2,2) = 0.8$. Two agents are assumed. In the first state, agent 1's outside option is 0 and agent 2's is set to 1.25. These are reversed in state 2. Output is constant at 1 across the states. These parameters determine $\underline{D}$ and $\overline{D}$. The bounding function $\underline{\underline{D}}$ is set equal to $\underline{D} - \varepsilon$. The dual Bellman is a contraction and ε chosen to ensure that it has a modulus of contraction of $\rho = 0.9$ with respect to the implied Thompson-like metric.

Figure 1a gives the value function on $\mathbb{R}^2_+$. In each iteration this function is evaluated at a finite number of points and then approximated on its entire domain as described above. The remainder of the figure shows optimal policies for agent 1 as a function of the shock s and co-state/Pareto variable $q^{V,1}$, with $q^{V,2}$ set so that $1 = \sqrt{(q^{V,1})^2 + (q^{V,2})^2}$ and $(q^{V,1}, q^{V,2})$ is in $\mathcal{C}_+$. As Figure 1c shows, the agent's incentive multiplier is positive in state 2 for low co-state/Pareto weight values $q^{V,1}$, otherwise, it is zero. Only for this combination of a high outside option (state 2) and a low Pareto weight, does the agent's incentive constraint bind. As Figure 1b shows, the agent's consumption is 0.4 (i.e. 40% of the endowment) for this combination. In contrast, in state $s = 1$, the agent receives a share of the endowment that becomes arbitrarily small as the agent's Pareto weight becomes small. On other hand in this state, the agent's share of the endowment reaches a maximum value of 0.6 for larger values of her Pareto weight (and correspondingly smaller values of the other agent's Pareto weight). For these values, agent 2's incentive constraint binds. Implications for agent 1's next period Pareto weight are illustrated in Figure 1d.

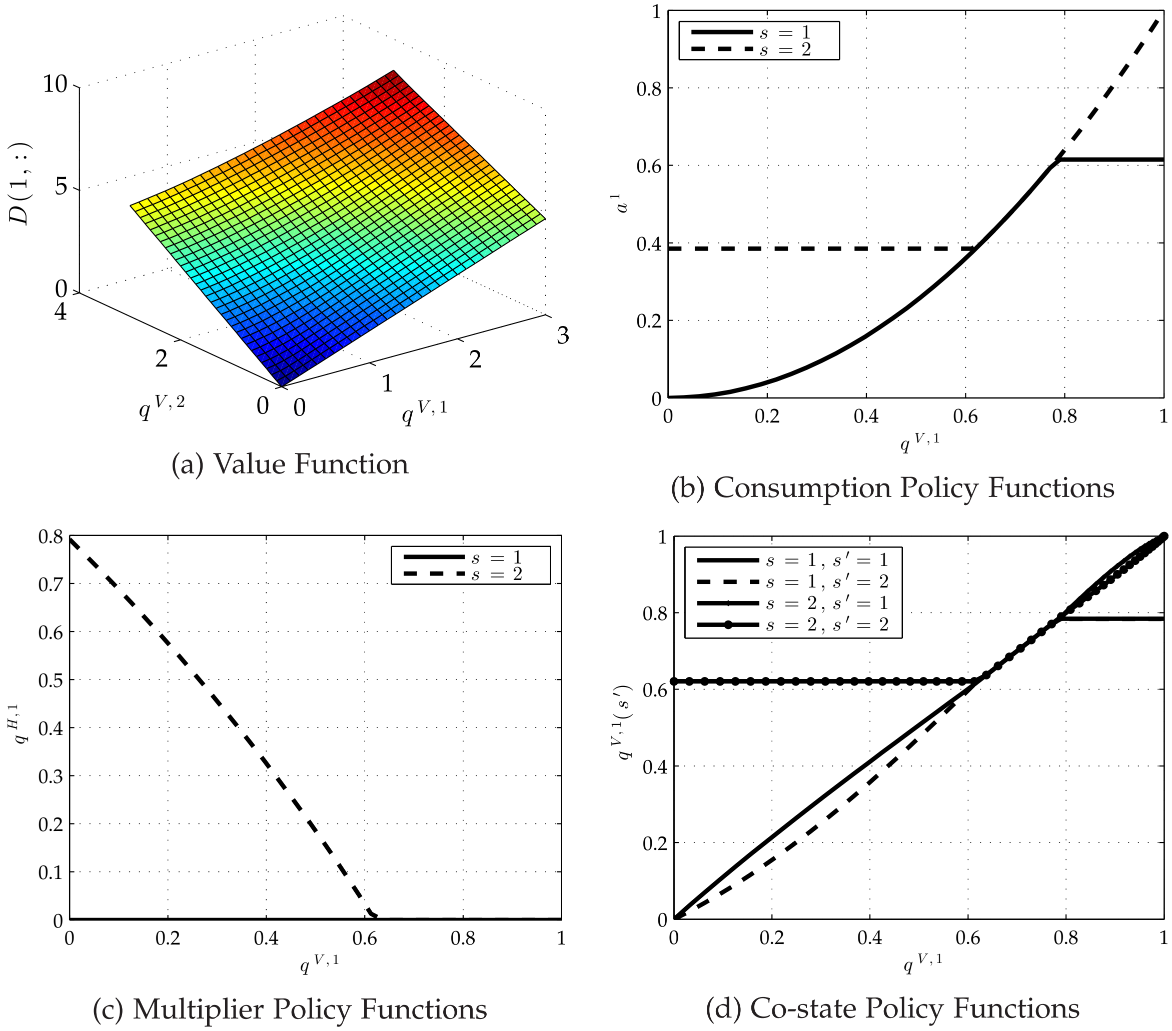

(a) Value Function

(b) Consumption Policy Functions

(c) Multiplier Policy Functions

(d) Co-state Policy Functions

Figure 1: Value and Policy Functions for the Limited Commitment Problem

10 Conclusion

In many settings the (primal) state space of a dynamic economic problem is defined implicitly and must be recovered as part of the solution to the problem. This complicates the application of recursive methods. Associated dual problems have recursive formulations in which co-states are used to keep track of histories of past or feasible future actions. If the primal state space is bounded, then the dual (co-)state space is immediately determined as $\mathbb{R}^N$ (or, perhaps, $\mathbb{R}^N_+$). Despite the unboundedness of the dual value functions and the lack of a bounded constraint correspondence, contractivity of the dual Bellman operator (with respect to the modified Thompson metric) may be established if suitable bounding functions are available. In many problems they are.

References

Acemoğlu, D., M. Golosov, and A. Tsyvinski (2010). Dynamic mirrlees taxation under political economy constraints. *Review of Economic Studies 77*, 841–881.

Aiyagari, R., A. Marcet, T. Sargent, and J. Seppälä (2002). Optimal taxation without state-contingent debt. *Journal of Political Economy 110*, 1220–1254.

Chien, Y., H. Cole, and H. Lustig (2011). A multiplier approach to understanding the macro implications of household finance. *Review of Economic Studies 78*, 199–234.

Cole, H. and F. Kubler (2012). Recursive contracts, lotteries and weakly concave pareto sets. *Review of Economic Dynamics 15*(4), 479–500.

Cooley, T., R. Marimon, and V. Quadrini (2004). Aggregate consequences of limited contract enforceability. *Journal of Political Economy 112*(4), 817–847.

Florenzano, M. and C. L. Van (2001). *Finite Dimensional Convexity and Optimization*. Springer Studies in Economic Theory 13.

Hopenhayn, H. and J. Nicolini (1997). Optimal unemployment insurance. *Journal of Political Economy 105*, 412–438.

Judd, K., S. Yeltekin, and J. Conklin (2003). Computing supergame equilibria. *Econometrica 71*, 1239–1254.

Kehoe, P. and F. Perri (2002). International business cycles with endogenous incomplete markets. *Econometrica 70*, 907 – 928.

Kocherlakota, N. (1996). Implications of efficient risk sharing without commitment. *Review of Economic Studies 63*, 595–609.

Luenberger, D. (1969). *Optimization by Vector Space Methods*. New York, John Wiley & Sons.

Marcet, A. and R. Marimon (1999). Recursive contracts. Working paper.

Marcet, A. and R. Marimon (2011). Recursive contracts. Working Paper.

Marimon, R. and V. Quadrini (2006). Competition, innovation and growth with limited commitment. NBER Working Paper 12474.

Marinacci, M. and L. Montrucchio (2010). Unique solutions for stochastic recursive utilities. *Journal of Economic Theory 145*, 1776–1804.

Messner, M., N. Pavoni, and C. Sleet (2011). On the dual approach to recursive optimization. IGIER Working Paper 423.

Messner, M., N. Pavoni, and C. Sleet (2012a). Contractive dual methods for incentive problems. IGIER Working Paper 466.

Messner, M., N. Pavoni, and C. Sleet (2012b). Recursive methods for incentive problems. *Review of Economic Dynamics 15*(4), 501–525.

Messner, M., N. Pavoni, and C. Sleet (2013). Countably additive multipliers in incentive problems. Working Paper.

Ponstein, J. (1981). On the use of purely finitely additive multipliers in mathematical programming. *Journal of Optimization Theory and Applications 33*, 37–55.

Rinćon-Zapatero, J. and C. Rodríguez-Palmero (2003). Existence and uniqueness of solutions to bellman equation in the unbounded case. *Econometrica 71*, 1519–1555.

Rockafellar, R. (1974). *Conjugate Duality and Optimization*. SIAM Society for Applied and Industrial Mathematics.

Rockafellar, T. and R. Wets (1998). *Variational Analysis*. Springer-Verlag.

Rustichini, A. (1998). Dynamic programming solution of incentive-constrained problems. *Journal of Economic Theory 78*, 329–354.

Schmutz, E. (2008). Rational points on the unit sphere. *Central European Journal of Mathematics 6*(3), 482–487.

Stokey, N., R. Lucas, and E. Prescott (1989). *Recursive Methods in Economic Dynamics*. Harvard University Press.

Thompson, A. (1963). On certain contraction mappings in a partially ordered vector space. *Proceedings of the American Mathematical Society 14*, 438–443.

Wessels, J. (1977). Markov programming by successive approximation with respect to weighted supremum norms. *Journal of Mathematical Analysis and Applications 58*, 326–355.

Woodford, M. (2003). *Interest and Prices*. Princeton University Press.

Appendix

A Construction of Payoffs

A function $g : \mathcal{S} \times \mathscr{A} \to \mathbb{R}^{n_v+1}$ is bounded if $\|g\|_\infty := \sup_{\mathcal{S}\times\mathscr{A}} \|g(s,\alpha)\|$. Let $\mathcal{G}$ denote the set of bounded functions $g : \mathcal{S} \times \mathscr{A} \to \mathbb{R}^{n_v+1}$. For $g \in \mathcal{G}$, define $T^V(g)(s,\alpha)$ according to:

$$T^V(g)(s,\alpha) = W^V[s, a_0, M^V[s, a_0, g'(\alpha)]],$$

where $g'(\alpha) = \{g(s', \alpha|s')\}_{s'=1}^{n_s}$.

Lemma 3. *$T^V : \mathcal{G} \to \mathcal{G}$ and is contractive.*

Proof. Let $g \in \mathcal{G}$. The monotonicity of M^V and the fact that $M^V[s, a, \cdot]$ maps constant-valued random variables to their constant values implies:

$$\sup_{\mathcal{S}\times\mathcal{A}} |M^V[s, a, g'(\alpha)]| \leq \|g\|_\infty I,$$

where I is the $n_v + 1$-unit vector. The boundedness and discounting properties of W^V imply:

$$\sup_{\mathcal{S}\times\mathscr{A}} \|W^V[s, a, M^V[s, a, g'(\alpha)]]\| \leq \sup_{\mathcal{S}\times\mathscr{A}} \|W^V[s, a, 0]\| + \bar{\delta}\|g\|_\infty < \infty.$$

We deduce that $T(g) \in \mathcal{G}$. Monotonicity of T follows from monotonicity of each $W^V[s, a, \cdot]$ and $M^V[s, a, \cdot]$. Let $g, \tilde{g} \in \mathcal{G}$, then from the monotonicity and sub-additivity of M^V:

$$M^V[s, a, g'(\alpha)] - M^V[s, a, \tilde{g}'(\alpha)] \leq M^V[s, a, \tilde{g}'(\alpha) + \|g - \tilde{g}\|_\infty] - M^V[s, a, \tilde{g}'(\alpha)] \leq \|g - \tilde{g}\|_\infty.$$

By the monotonicity and discounting properties of W^V, for each (s, a, α),

$$\begin{aligned} W^V[s, a, &M^V[s, a, g'(\alpha)]] - W^V[s, a, M^V[s, a, \tilde{g}'(\alpha)]] \\ &\leq W^V[s, a, M^V[s, a, \tilde{g}'(\alpha)] + \|\tilde{g} - g\|_\infty] - W^V[s, a, M^V[s, a, \tilde{g}'(\alpha)]] \leq \bar{\delta}\|\tilde{g} - g\|_\infty. \end{aligned}$$

Hence, T^V satisfies a discounting property and, by Blackwell's theorem, is a contraction on $\mathcal{G}$. ☐

It follows from Lemma 3, the completeness of $\mathcal{G}$ and the contraction mapping theorem that T^V has a unique fixed point on $\mathcal{G}$. V is identified with this function. By placing additional continuity restrictions on W^V and M^V, the previous result may be strengthened to show that V, the unique fixed point on $\mathcal{G}$, is continuous.

B Proofs for Section 4

Proof of Proposition 2. For $(k, s, v) \in \mathcal{X}$, let:

$$\Omega(k, s, v) = \left\{ \pi \in \mathscr{P} \;\middle|\; \begin{array}{l} k_0 = k, s_0 = s, v_0 = v, \\ k_{t+1}(s^t) = W^K[k_t(s^{t-1}), s_t, a_t(s^t)], \\ v_t(s^t) = W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), v_{t+1}(s^t)]], \\ H[k_t(s^{t-1}), a_t(s^t), v_{t+1}(s^t)] \geq 0 \end{array} \right\}.$$

It follows from definitions that $\Omega(k, s, v) = \{\pi | k_0 = k, v_0 = v, (a_0, k_1, v_1) \in \Gamma(k, s, v), \pi|s' \in \Omega(k_1, s', v_1(s'))\}$, where $\pi|s'$ is the continuation of π after $s_1 = s'$. For $(k, s, v) \in \mathcal{X}$, let $P^*(k, s, v) = \sup_{\Omega(k,s,v)} V^0(s, \alpha)$. Define $\mathbf{B}$ on the domain $\mathscr{F} = \{P : \mathcal{X} \to \mathbb{R}\}$ as, for $P \in \mathscr{F}$ and $(k, s, v) \in \mathcal{X}$, $\mathbf{B}(P)(k, s, v) = \sup_{\Gamma(k,s,v)} W^{V,0}[s, a, M^{V,0}[s, a, P'(k', v')]]$, with $P'(k'v') =$

$\{P(k', s', v'(s'))\}_{s' \in \mathcal{S}}$. We verify that for $(k, s, v) \in \mathcal{X}$, $P^*(k, s, v) = \mathbf{B}(P^*)(k, s, v)$. Suppose $P^*(k, s, v) > \sup_{\Gamma(k,s,v)} W^{V,0}[s, a, M^{V,0}[s, a, P^{*\prime}(k', v')]]$. Since $(k, s, v) \in \mathcal{X}$, $\Omega(k, s, v) \neq \emptyset$ and there is a $\pi \in \Omega(k, s, v)$ with $V^0(s, \alpha) > \sup_{\Gamma(k,s,v)} W^{V,0}[s, a, M^{V,0}[s, a, P^{*\prime}(k', v')]]$. But $\pi \in \Omega(k, s, v)$, thus $(a_0, k_1, v_1) \in \Gamma(k, s, v)$ and $\pi|s' \in \Omega(k_1, s', v_1(s'))$. From the monotonicity of $W^{V,0}$ in its third argument and the definition of P^*, $V^0(s, \alpha) = W^{V,0}[s, a_0, M^{V,0}[s, a_0, V^{0\prime}(\alpha)]] \leq W^{V,0}[s, a_0, M^{V,0}[s, a_0, P^{*\prime}(k_1, v_1)]] \leq \sup_{\Gamma(k,s,v)} W^{V,0}[s, a, M^{V}[s, a, P^{*\prime}(k', v')]]$. This is a contradiction and so $P^*(k, s, v) \leq \mathbf{B}(P^*)(k, s, v)$. Next suppose that $P^*(k, s, v) < \sup_{\Gamma(k,s,v)} W^{V,0}[s, a, M^{V,0}[s, a, P^{*\prime}(k', v')]]$. Then, since $(k, s, v) \in \mathcal{X}$, $\Gamma(k, s, v)$ is non-empty and there is a triple $(a, k', v') \in \Gamma(k, s, v)$ with $P^*(k, s, v) < W^{V,0}[s, a, M^{V,0}[s, a, P^{*\prime}(k', v')]]$. Since $W^{V,0}$ and $M^{V,0}$ are continuous in their third arguments, there is a family $\pi|s'$ with for each $\pi|s' \in \Omega(k', s', v')$ and with associated plans $\alpha|s'$ satisfying $P^*(k, s, v) < W^{V,0}[s, a, M^{V,0}[s, a, V^{0\prime}(\alpha)]] = V^0(s, \alpha)$. But the definition of $\Omega(k, s, v)$ implies that $\pi = (k, v, a, \{\pi|s'\}) \in \Omega(k, s, v)$. Hence, $V^0(s, \alpha) \leq \sup_{\Omega(k,s,v)} V(s, \alpha') = P^*(k, s, v)$, another contradiction. Thus, $P^*(k, s, v) \geq \mathbf{B}(P^*)(k, s, v)$. Combining inequalities and noting that (k, s, v) was arbitrary in $\mathcal{X}$, it follows that $P^* = \mathbf{B}(P^*)$ as required. By a very similar argument, $P_0^* = \sup_{\mathcal{V}(\bar{k}, s_0)} P^*(\bar{k}, s_0)$.

Similar reasoning to that above establishes that any solution to (PA) satisfies (i) and (ii) from the proposition. Conversely, let $\pi^* = (\alpha^*, \{k_t^*, v_t^*\})$ be a primal process satisfying (i) and (ii) in the proposition. Feasibility of π^* for (PA) is immediate. Also,

$$
\begin{aligned}
|P^*(k_0^*, s_0, v_0^*) - V^0(s_0, \alpha^*)| &= |W^{V,0}[s_0, a_0^*, M^{V,0}[s_0, a_0^*, \{P^*(k_1^*(s_0), s_1, v_1^*(s_0))\}_{s_1 \in \mathcal{S}}]] \\
&\qquad - W^{V,0}[s_0, a_0^*, M^{V,0}[s_0, a_0^*, \{V^0(s_1, \alpha^*|s_1)\}_{s_1 \in \mathcal{S}}]]| \\
&\leq \delta \max_{s_1 \in \mathcal{S}} |P^*(k_1^*(s_0), s_1, v_1^*(s_0)) - V^0(s_1, \alpha^*|s_1)| \\
&\leq \ldots \leq \delta^t \max_{s^t \in \mathcal{S}} |P^*(k_t^*(s^{t-1}), s_t, v_t^*(s^{t-1})) - V^0(s_t, \alpha^*|s^t)|, \qquad (35)
\end{aligned}
$$

where the first equality uses property (ii) in the proposition and (10), the first inequality uses the sub-additivity of $M^{V,0}$ and the discounting property of $W^{V,0}$. The final inequality follows from an iteration of these arguments. The boundedness of P^* and V^0 and $\delta \in (0, 1)$ then implies that the final term in (35) converges to 0 as t converges to ∞. Hence, $P^*(k_0^*, s_0, v_0^*) = V^0(s_0, \alpha^*)$. Then using property (i) in the proposition and (9), we have that $P_0^* = P^*(k_0^*, s_0, v_0^*) = V^0(s_0, \alpha^*)$ and π^* is a solution to (PA). □

C Proofs for Section 5

Proof of Proposition 3. We have:

$$
\begin{aligned}
D_0^* &= \inf_{\mathcal{Q}} \sup_{\mathcal{P}} \mathscr{L}(\pi,\theta) \\
&= \inf_{\mathcal{Q}} \sup_{\mathcal{P}} F[s_0,v_0] + q_0^K \cdot (\bar{k} - k_0) + q_0^V \cdot (W^V[s_0,a_0,M^V[s_0,a_0,v_1]] - v_0) \\
&\qquad\qquad + q_0^H \cdot H[k,s_0,a_0,v_1] + \bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta(\pi) | s_1 \rangle \\
&= \inf_{\mathcal{Y}} \sup_{\mathcal{V}} q_0^K \cdot \bar{k} + F[s_0,v_0] - q_0^V \cdot v_0 \\
&\qquad\qquad + \inf_{\mathcal{Q}(q_0^K,q_0^V)} \sup_{\mathscr{P}(v_0)} -q_0^K \cdot k_0 + q_0^V \cdot W^V[s_0,a_0,M^V[s_0,a_0,v_1]] \\
&\qquad\qquad\qquad + q_0^H \cdot H[k,s_0,a_0,v_1] + \bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta(\pi) | s_1 \rangle,
\end{aligned}
$$

which combined with the definition of D^* gives the first equality in the proposition. For each $(s,y) = (s,q^K,q^V) \in \mathcal{S} \times \mathcal{Y}$,

$$
\begin{aligned}
D^*(s,y) &= \inf_{\mathcal{Q}(y)} \sup_{\mathscr{P}(v_0)} -q^K \cdot k_0 + q^V \cdot W^V[s,a_0,M^V[s,a_0,v_1]] \\
&\qquad\qquad + q_0^H \cdot H[k_0,s,a_0,v_1] + \bar{\delta} \sum_{s' \in \mathcal{S}} \langle \theta, \zeta(\pi) | s' \rangle \\
&= \inf_{\mathcal{Q}(y)} \sup_{\mathscr{P}(v_0)} -q^K \cdot k_0 + q^V \cdot W^V[s,a_0,M^V[s,a_0,v_1]] + q_0^H \cdot H[k_0,s,a_0,v_1] \\
&\qquad\qquad - \bar{\delta} \sum_{s' \in \mathcal{S}} q_1^V(s') \cdot v_1(s') + \bar{\delta} \sum_{s' \in \mathcal{S}} q_1^K(s') \cdot \{W^K[k_0,s,a_0] - k_1(s')\} \\
&\qquad\qquad + \bar{\delta} \sum_{s' \in \mathcal{S}} \Big\{ q_1^V(s') \cdot W^V[s',a_1(s'),M^V[s',a_1(s'),v_2(s')]] \\
&\qquad\qquad\qquad + q_1^H(s') \cdot H[k_1,s',a_1(s'),v_2(s')] + \bar{\delta} \sum_{s'' \in \mathcal{S}} \langle \theta, \zeta(\pi) | s', s'' \rangle \Big\}.
\end{aligned}
$$

Note that once $q_0 = (q_0^H, q_1^K, q_1^V)$ is chosen, $p_0 = (k_0, a_0, v_1)$ is independent of the remaining dual variables. Consequently, conditional on $q_0 = (q_0^H, q_1^K, q_1^V)$, the infimum over these dual variables and the supremum over p_0 may be interchanged to give:

$$\begin{aligned}
D^*(s,y) = &\inf_{\mathcal{Q}} \sup_{\mathcal{P}} \; -q_0^K \cdot k_0 + q_0^V \cdot W^V[s, a_0, M^V[s, a_0, v_1]] \\
&+ q_0^H \cdot H[k_0, s, a_0, v_1] - \bar{\delta} \sum_{s' \in \mathcal{S}} q_1^V(s') v_1(s') + \sum_{s' \in \mathcal{S}} q_1^K(s') \cdot W^K[k_0, s, a_0] \\
&+ \bar{\delta} \sum_{s' \in \mathcal{S}} \inf_{\mathscr{Q}(y_1(s'))} \sup_{\mathscr{P}(v_1(s'))} \Big\{ - q_1^K(s') \cdot k_1(s') + q_1^V(s') \cdot W^V[s', a_1(s'), M^V[s', a_1(s'), v_2(s')]] \\
&\qquad + q_1^H(s') \cdot H[k_1, s', a_1(s'), v_2(s')] + \bar{\delta} \sum_{s'' \in \mathcal{S}} \langle \theta, \zeta(\pi) | s', s'' \rangle \Big\}. \\
= &\inf_{\mathcal{Q}} \sup_{\mathcal{P}} \; -q_0^K \cdot k_0 + q_0^V \cdot W^V[s, a_0, M^V[s, a_0, v_1]] + q_0^H \cdot H[k_0, s, a_0, v_1] \\
&- \bar{\delta} \sum_{s' \in \mathcal{S}} q_1^V(s') \cdot v_1(s') + \bar{\delta} \sum_{s' \in \mathcal{S}} q_1^K(s') \cdot W^K[k_0, s, a_0] + \bar{\delta} \sum_{s' \in \mathcal{S}} D^* \left(s', y_1(s')\right).
\end{aligned}$$

Combining the last equality with the definition of J gives the second equality in the proposition. ☐

Proof of Proposition 4. Choose an arbitrary $(s,y) = (s, q^K, q^V) \in \mathcal{S} \times \mathcal{Y} = \mathcal{S} \times \mathbb{R}^{n_k + n_v + 1}$. Since $0 \in \mathscr{Q}$ is a feasible multiplier choice for the infimum in the continuation problem (12):

$$\begin{aligned}
D^*(s,y) = &\inf_{\mathscr{Q}(y)} \sup_{\mathscr{P}(v_0)} \; -q^K \cdot k_0 + q^V \cdot W^V[s, a_0, M^V[s, a_0, v_1]] \\
&\qquad + q^H \cdot H[k_0, s, a_0, v_1] + \sum_{s' \in \mathcal{S}} \langle \theta, \zeta(\pi) | s' \rangle \\
\leq &\sup_{\mathcal{P}} -q^K \cdot k_0 + q^V \cdot W[s, a_0, M[s, a_0, v_1]] < \infty,
\end{aligned}$$

where the last inequality uses the boundedness of $\mathcal{K} \times \mathcal{V}$. On the other hand, for a fixed feasible primal process $\pi' \in \mathscr{P}$ and an arbitrary dual process in $\mathscr{Q}$,

$$\begin{aligned}
&\sup_{\mathscr{P}} \; - q^K \cdot k_0 + q^V \cdot W^V[s, a_0, M^V[s, a_0, v_1]] + q^H \cdot H[k_0, s, a_0, v_1] + \sum_{s' \in \mathcal{S}} \langle \theta, \zeta(\pi) | s' \rangle \\
&\qquad \geq -q^K \cdot k_0' + q^V \cdot W^V[s, a_0', M^V[s, a_0', v_1']] > -\infty.
\end{aligned}$$

And so, $D^*(s,y) \geq -q^K \cdot k_0' + q^V \cdot W^V[s, a_0', M^V[s, a_1', v_1']] > -\infty$. Hence, $D^*(s,y)$ is real and, since (s,y) was arbitrary, D^* is real-valued on $\mathcal{S} \times \mathcal{Y}$. ☐

Proof of Proposition 5. (Only if) Let $J_0(y_0, v_0) = F[s_0, v_0] - q_0^V \cdot v_0 + q_0^K \cdot \bar{k}$. Using this definition and that of J implies:

$$\mathscr{L}(\pi, \theta) = J_0(y_0, v_0) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J(s_t, y_t(s^t); q_t(s^t), p_t(s^t)).$$

Thus, if $(\theta^*, \pi^*) \in \Lambda^{IS}$, then:

$$\pi^* \in \underset{\mathscr{P}}{\operatorname{argmax}}\, J_0(y_0^*, v_0) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J(s_t, y_t^*(s^t); q_t^*(s^t), p_t(s^t)).$$

The above maximization can be decomposed into a collection of static maximizations with $v_0^* \in \operatorname{argmax}_{\mathcal{V}} J_0(y_0^*, v_0)$ and $p_t^*(s^t) \in \operatorname{argmax}_{\mathcal{P}} J(s_t, y_t^*(s^t); q_t^*(s^t), p_t(s^t))$. Let $J_0^*(y) = \sup_{\mathcal{V}} J_0(y, v_0)$ and $J^*(s, y; q) = \sup_{\mathcal{P}} J(s, y; q, p)$. Then:

$$\begin{aligned} D_0^* &= J_0^*(y_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t)) \\ &\leq J_0^*(y_0) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t(s^t); q_t(s^t)), \qquad \theta \in \mathscr{Q}. \end{aligned}$$

In particular, the inequality holds for all θ with initial element y_0^* and so, since $D^*(s, y) = \inf_{\mathscr{Q}(y)} \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t(s^t); q_t(s^t))$,

$$D_0^* = J_0^*(y_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t)) \leq J_0^*(y_0^*) + D^*(s_0, y_0^*).$$

Conversely, since the continuation of θ^* lies in $\mathscr{Q}(y)$, the reverse inequality holds: $D_0^* = J_0^*(y_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t)) \geq J_0^*(y_0^*) + D^*(s_0, y_0^*)$. Hence $D_0^* = J^*(y_0^*) + D^*(s_0, y_0^*)$ and y_0^* attains the minimum in (14). Consequently, $(y_0^*, v_0^*) \in G_0^{IS}$. Pursuing the same argument at successive histories gives $(q_t^*(s^t), y_{t+1}^*(s^t))$ attains the minimum in (15) at $(s_t, y_t^*(s^t))$ and so $(q_t^*(s^t), p_t^*(s^t)) \in G^{IS}(s_t, y_t^*(s^t))$.

(If) Suppose (π^*, θ^*) is such that $(q_0^{K*}, p_0^{V*}, v_0^*) \in G_0^{IS}$ and for each $t \in \mathbb{N}$, $s^t \in \mathcal{S}^t$, $(q_t^{H*}(s^t), y_{t+1}^*(s^t), p_t^*(s^t)) \in G^{IS}(s_t, y_t^*(s^t))$. The definitions of G_0^{IS} and G^{IS} imply that $J_0(y_0^*, v_0^*) = \sup_{\mathcal{V}} J_0(y_0^*, v_0)$ and $J(s_t, y_t^*(s^t); q_t^*(s^t), p_t^*(s^t)) = \sup_{\mathcal{P}} J(s_t, y_t^*(s^t); q_t^*(s^t), p_t(s^t))$. Hence, for arbitrary $\pi \in \mathscr{P}$,

$$\begin{aligned} \mathscr{L}(\pi^*, \theta^*) &= J_0(y_0^*, v_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^{t-1} \sum_{\mathcal{S}^t} J(s_t, y_t^*(s^t); q_t^*(s^t), p_t^*(s^t)) \\ &\geq J_0(y_0^*, v_0) + \sum_{t=0}^{\infty} \bar{\delta}^{t-1} \sum_{\mathcal{S}^t} J(s_t, y_t^*(s^t); q_t^*(s^t), p_t(s^t)) \\ &= \mathscr{L}(\pi, \theta^*). \end{aligned}$$

And so $\pi^* \in \operatorname{argmax}_{\mathscr{P}} \mathscr{L}(\pi, \theta^*)$. Let $J_0^*(y) = \sup_{\mathcal{V}} J_0(y, v_0)$ and $J^*(s, y; q) = \sup_{\mathcal{P}} J(s, y; q, p)$. Then:

$$D_0^* = \inf_{\mathscr{Q}} \sup_{\mathscr{P}} \mathscr{L}(\pi, \theta) = \inf_{\mathscr{Q}} J_0^*(y_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^{t-1} \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t))$$

The definitions of G_0^{IS} and G^{IS} imply:

$$D_0^* = J_0^*(y_0^*) + D^*(s_0, y_0^*)$$

and

$$D^*(s_t, y_t^*(s^t)) = J^*(s_t, y_t^*(s^t); q_t^*(s^t)) + \bar{\delta} \sum_{s' \in \mathcal{S}} D^*(s', y_{t+1}^*(s^t, s')).$$

Consequently, we have:

$$D_0^* = J_0^*(y_0^*) + \sum_{t=0}^{T} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t)) + \bar{\delta}^{T+1} \sum_{\mathcal{S}^{T+1}} D^*(s_{T+1}, y_{T+1}^*(s^{T+1})).$$

Taking the limit as T goes to infinity and using the condition in the proposition implies that

$$D_0^* \geq J_0^*(y_0^*) + \sum_{t=0}^{\infty} \bar{\delta}^t \sum_{\mathcal{S}^t} J^*(s_t, y_t^*(s^t); q_t^*(s^t)).$$

But since $\theta^* \in \mathscr{Q}$ and so is feasible for the minimization defining Λ^{IS}, the reverse inequality holds and θ^* attains the minimum as required. □

D Proofs for Section 6

Proof of Lemma 1. We begin with simple general result. Let Ψ, Φ and Ω denote vector spaces and $L : \Psi \times \Phi \times \Omega \to \mathbb{R}$ a real-valued function. Assume that for each $\omega \in \Omega$, $L(\cdot, \cdot, \omega)$ is sub-linear. For $\psi \in \Psi$, let: $\Lambda(\psi) = \inf_{\Phi} \sup_{\Omega} L(\psi, \phi, \omega)$. We prove that Λ is sub-linear. To begin with we first show that Λ is convex. Let ψ_1 and ψ_2 be elements of Ψ and $\lambda \in [0, 1]$. Let $\psi_\lambda = \lambda\psi_1 + (1-\lambda)\psi_2$. Assume that the infimum defining Λ is attained at ψ_i by some ϕ_i^*, $i = 1, 2$. This assumption simplifies the exposition and can easily be dropped. Let $\phi_\lambda^* = \lambda\phi_1^* + (1-\lambda)\phi_2^*$. Then:

$$\begin{aligned}
\lambda\Lambda(\psi_1) + (1-\lambda)\Lambda(\psi_1) &= \lambda \inf_{\Phi} \sup_{\Omega} L(\psi_1; \phi, \omega) + (1-\lambda) \inf_{\Phi} \sup_{\Omega} L(\psi_2; \phi, \omega) \\
&= \lambda \sup_{\Omega} L(\psi_1; \phi_1^*, \omega) + (1-\lambda) \sup_{\Omega} L(\psi_2; \phi_2^*, \omega) \\
&\geq \sup_{\Omega} \{\lambda L(\psi_1; \phi_1^*, \omega) + (1-\lambda) L(\psi_2; \phi_2^*, \omega)\} \\
&\geq \sup_{\Omega} L(\psi_\lambda; \phi_\lambda^*, \omega) \geq \inf_{\Phi} \sup_{\Omega} L(\psi_\lambda; \phi, \omega) = \Lambda(\psi_\lambda),
\end{aligned}$$

where the second inequality uses the convexity of $L(\cdot, \cdot, \omega)$. Thus, Λ is convex. Next we show homogeneity. Suppose that $\psi \in \Psi$ and $\lambda > 0$. Then:

$$\Lambda(\lambda\psi) = \inf_{\Phi} \sup_{\Omega} L(\lambda\psi; \phi, \omega) = \lambda \inf_{\Phi} \sup_{\Omega} L(\psi; \phi/\lambda, \omega) = \lambda\Lambda(\psi),$$

where the second equality uses the positive homogeneity of $L(\cdot,\cdot,\omega)$. Thus, Λ is positively homogenous of degree 1 and, combining results, sub-linear.

(i) For fixed $s_0 \in \mathcal{S}$, define the "continuation Lagrangian":

$$M(y_0; q_0^H, \{\theta|s_1\}, p, \{\zeta(\pi)|s_1\}) = -q_0^K \cdot k_0 + q_0^V \cdot W^V[s_0, a_0, M^V[s_0, a_0, v_1]] \\ + q_0^H \cdot H[k_0, s_0, a_0, v_1] + \bar{\delta} \sum_{s_1 \in \mathcal{S}} \langle \theta, \zeta(\pi)|s_1 \rangle,$$

Setting $\psi = y_0$, $\Psi = \mathcal{Y}$, $\phi = (q_0^H, \{\theta|s_1\}_{s_1 \in \mathcal{S}})$, $\Phi = \mathscr{Q}(y_0)$, $\omega = (q, \{\pi|s_1\}_{s_1 \in \mathcal{S}})$, $\Omega = \mathscr{P}(v_0)$ and $L(\psi; \phi, \omega) = M(y_0; q_0^H, \{\theta|s_1\}, p, \{\zeta|s_1\})$, it follows that for each ω, $L(\cdot;\cdot,\omega)$ is linear and, hence, sub-linear. Applying the general result from the first part of the proof, $D^*(s_0,\cdot)$ is sub-linear. Since s_0 was arbitrary in $\mathcal{S}$, D^* is sub-linear.

(ii) It is easy to verify that for each (s,p), $J(s,\cdot;\cdot,p)$ is linear and, hence, sub-linear. Assume that D is sub-linear. Then for each (s,p), $J(s,\cdot;\cdot,p) + \bar{\delta} \sum_{s' \in \mathcal{S}} D(s',\cdot)$ is sub-linear. Consequently, the logic from the first part of the proof establishes that $\mathbf{B}(D)(s,\cdot)$,

$$\mathbf{B}(D)(s,y) = \inf_{\mathcal{Q}} \sup_{\mathcal{P}} J(s,y;q,p) + \bar{\delta} \sum_{s' \in \mathcal{S}} D\left(s', y'(s')\right),$$

is sub-linear. Since s was arbitrary in $\mathcal{S}$, $\mathbf{B}(D)$ is sub-linear. ☐

Proof of Lemma 2. Evidently, $(\mathscr{G}, d)$ is a metric space. Let $\{D_n\}$ be a Cauchy sequence in $\mathscr{G}$. Thus, as $n, m \to \infty$,

$$d(D_n, D_m) = \sup_{\mathcal{S}\times\mathcal{C}} \left| \ln\left(\frac{D_n(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) - \ln\left(\frac{D_m(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) \right| \to 0.$$

For each $n \in \mathbb{N}$, define $g_n : \mathcal{S} \times \mathcal{C} \to \mathbb{R}$ according to: $g_n(s,y) = \ln\left(\frac{D_n(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right)$, $(s,y) \in \mathcal{S} \times \mathcal{C}$. Let $\underline{g} = 0$ and $\overline{g} = \ln\{(\overline{D} - \underline{\underline{D}})/(\underline{D} - \underline{\underline{D}})\}$. It follows that $\{g_n\}$ is Cauchy with respect to the sup-norm and that for each n, $\underline{g} \le g_n \le \overline{g}$. By the completeness of the continuous, bounded functions from $\mathcal{C}$ to $\mathbb{R}$, $\{g_n\}$ converges in the sup-norm to a function g_∞, with each $g_\infty(s,\cdot)$ continuous and bounded and $\underline{g} \le g_\infty \le \overline{g}$. Use g_∞ to define the homogeneous function D_∞ as:

$$D_\infty(s,y) = \|y\| \left\{ \underline{\underline{D}}\left(s, \frac{y}{\|y\|}\right) + \exp\left\{ g_\infty\left(s, \frac{y}{\|y\|}\right)\right\} \left(\underline{D}\left(s, \frac{y}{\|y\|}\right) - \underline{\underline{D}}\left(s, \frac{y}{\|y\|}\right)\right)\right\}.$$

By construction $\underline{D} \le D_\infty \le \overline{D}$ and $D_n \xrightarrow{d} D_\infty$. Since D_∞ is the pointwise limit of a sequence of sub-linear and, hence, convex functions, it too is convex. Hence, it is in $\mathscr{G}$. ☐

Proof of Proposition 6. Let $\mathscr{G}_0$ denote the interval of real-valued functions between $\underline{\underline{D}}$ and $\overline{D}$. Let D_1 and D_2 be any pair of functions in $\mathscr{G}_0$ and let $\lambda \in [0,1]$. Define for each

$(s, y, q) \in \mathcal{S} \times \mathcal{Y} \times \mathcal{Q}$, $J^*(s, y, q) = \sup_{\mathcal{P}} J(s, y; q, p)$. Then, for each $(s, y) \in \mathcal{S} \times \mathcal{Y}$,

$$\begin{aligned}
\mathbf{B}(\lambda D_1 + (1-\lambda) D_2)(s, y) &= \inf_{\mathcal{Q}} J^*(s, y, \eta, y') + \sum_{s' \in \mathcal{S}} \{\lambda D_1 (s', y'(s')) + (1-\lambda) D_2 (s', y'(s'))\} \\
&= \inf_{\mathcal{Q}} \lambda \Big\{ J^*(s, y, \eta, y') + \sum_{s' \in \mathcal{S}} D_1 (s', y'(s')) \Big\} \\
&\qquad + (1-\lambda) \Big\{ J^*(s, y, \eta, y') + \sum_{s' \in \mathcal{S}} D_2 (s', y'(s')) \Big\} \\
&\geq \lambda \inf_{\mathcal{Q}} \Big\{ J^*(s, y, \eta, y') + \sum_{s' \in \mathcal{S}} D_1 (s', y'(s')) \Big\} \\
&\qquad + (1-\lambda) \inf_{\mathcal{Q}} \Big\{ J^*(s, y, \eta, y') + \sum_{s' \in \mathcal{S}} D_2 (s', y'(s')) \Big\} \\
&= \lambda \mathbf{B}(D_1)(s, y) + (1-\lambda) \mathbf{B}(D_2)(s, y).
\end{aligned}$$

Thus, $\mathbf{B}$ is concave on $\mathscr{G}_0$. Let $D_1, D_2 \in \mathscr{G} \subset \mathscr{G}_0$. By definition of d, for each $(s, y) \in \mathcal{S} \times \mathcal{C}$,

$$\ln \left(\frac{D_2(s, y) - \underline{\underline{D}}(s, y)}{\underline{D} - \underline{\underline{D}}} \right) \leq \ln \left(\frac{D_1(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right) + d(D_1, D_2).$$

Taking the exponential of each side and rearranging gives:

$$\exp\{-d(D_1, D_2)\} \left(\frac{D_2(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right) \leq \left(\frac{D_1(s, y) - \underline{\underline{D}}(s, y)}{\underline{D}(s, y) - \underline{\underline{D}}(s, y)} \right).$$

But, by Assumption 3 (i), $\underline{D} - \underline{\underline{D}} > 0$ and so, after rearrangement,

$$D_1(s, y) \geq \exp\{-d(D_1, D_2)\} D_2(s, y) + (1 - \exp\{-d(D_1, D_2)\}) \underline{\underline{D}}(s, y). \tag{36}$$

Since D_1, D_2 and $\underline{\underline{D}}$ are positively homogeneous of degree 1, this inequality holds at all $(s, y) \in \mathcal{S} \times \mathcal{Y}$. Then, by monotonicity and concavity of $\mathbf{B}$ (on $\mathscr{G}_0$),

$$\begin{aligned}
\mathbf{B}(D_1) &\geq \mathbf{B}(\exp\{-d(D_1, D_2)\} D_2 + (1 - \exp\{-d(D_1, D_2)\}) \underline{\underline{D}}) \\
&\geq \exp\{-d(D_1, D_2)\} \mathbf{B}(D_2) + (1 - \exp\{-d(D_1, D_2)\}) \mathbf{B}(\underline{\underline{D}}).
\end{aligned} \tag{37}$$

By assumption there is a $\varepsilon_1 > 0$ such that for each $(s, y) \in \mathcal{S} \times \mathcal{C}$, $\mathbf{B}(\underline{\underline{D}})(s, y) > \underline{\underline{D}}(s, y) + \varepsilon_1$. For $(s, y) \in \mathcal{S} \times \mathcal{C}$, define:

$$\lambda(s, y) := \frac{\varepsilon_1}{\overline{D}(s, y) - \underline{\underline{D}}(s, y)}.$$

Since $\overline{D}(s, y) \geq \mathbf{B}(\overline{D})(s, y) \geq \mathbf{B}(\underline{\underline{D}})(s, y) > \underline{\underline{D}}(s, y) + \varepsilon_1$, $\lambda(s, y) \in (0, 1)$. Now, for each $s \in \mathcal{S}$, $\underline{\underline{D}}(s, \cdot)$ and $\overline{D}(s, \cdot)$ are continuous. Thus, $\lambda(s, \cdot)$ is continuous and since $\mathcal{C}$ is compact, there is a $\lambda^* = \min_{\mathcal{S}\times\mathcal{C}} \lambda(s,y) \in (0,1)$. Then, for all $(s,y) \in \mathcal{S}\times\mathcal{C}$,

$$\begin{aligned}
\mathbf{B}(\underline{\underline{D}})(s,y) &> \underline{\underline{D}}(s,y) + \varepsilon_1 = \lambda(s,y)\overline{D}(s,y) + (1-\lambda(s,y))\underline{\underline{D}}(s,y) \\
&\geq \lambda^*\overline{D}(s,y) + (1-\lambda^*)\underline{\underline{D}}(s,y) \\
&\geq \lambda^*\mathbf{B}(D_2)(s,y) + (1-\lambda^*)\underline{\underline{D}}(s,y),
\end{aligned} \tag{38}$$

where the first inequality is by assumption, the first equality uses the definition of $\lambda(s,y)$, the second inequality uses the definition of λ^* and $\overline{D} \geq \underline{\underline{D}}$ and the final inequality uses $\overline{D} \geq \mathbf{B}(\overline{D}) \geq \mathbf{B}(D_2)$. Combining (37) with (38) gives for all $(s,y) \in \mathcal{C}$,

$$\begin{aligned}
\mathbf{B}(D_1)(s,y) \geq & \exp\{-d(D_1,D_2)\}\mathbf{B}(D_2)(s,y) + (1-\exp\{-d(D_1,D_2)\}) \\
& \times [\lambda^*\mathbf{B}(D_2)(s,y) + (1-\lambda^*)\underline{\underline{D}}(s,y)].
\end{aligned}$$

Letting $r := \exp\{-d(D_1,D_2)\} + (1-\exp\{-d(D_1,D_2)\})\lambda^*$, then gives for $(s,y) \in \mathcal{S}\times\mathcal{C}$:

$$\frac{\mathbf{B}(D_1)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)} \geq r \frac{\mathbf{B}(D_2)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}.$$

Hence, taking logs, for $(s,y) \in \mathcal{S}\times\mathcal{C}$,

$$\ln\left(\frac{\mathbf{B}(D_1)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) \geq \ln r + \ln\left(\frac{\mathbf{B}(D_2)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right).$$

But from the definition of r and Jensen's inequality:

$$\ln r \geq (1-\lambda^*)\ln\exp\{-d(D_1,D_2)\} + \lambda^*\ln 1 = -(1-\lambda^*)d(D_1,D_2).$$

Thus, for $(s,y) \in \mathcal{S}\times\mathcal{C}$,

$$(1-\lambda^*)d(D_1,D_2) \geq -\ln r \geq \ln\left(\frac{\mathbf{B}(D_2)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) - \ln\left(\frac{\mathbf{B}(D_1)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right). \tag{39}$$

Repeating the argument with D_1 and D_2 interchanged and combining with (39) implies that for all $(s,y) \in \mathcal{S}\times\mathcal{C}$,

$$(1-\lambda^*)d(D_1,D_2) \geq \left|\ln\left(\frac{\mathbf{B}(D_2)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) - \ln\left(\frac{\mathbf{B}(D_1)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right)\right|.$$

Consequently, letting $\rho := (1-\lambda^*) \in (0,1)$,

$$\begin{aligned}
\rho d(D_1,D_2) &\geq \sup_{\mathcal{S}\times\mathcal{C}}\left|\ln\left(\frac{\mathbf{B}(D_2)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right) - \ln\left(\frac{\mathbf{B}(D_1)(s,y) - \underline{\underline{D}}(s,y)}{\underline{D}(s,y) - \underline{\underline{D}}(s,y)}\right)\right| \\
&= d(\mathbf{B}(D_1),\mathbf{B}(D_2))
\end{aligned}$$

as desired. $\square$

Bounding value functions for Example 1 Assume as in the main text an $\tilde{a} \in \mathcal{A}^{n_s}$ and a $\xi > 0$ such that for each $s \in \mathcal{S}$, $\gamma(s) > \sum_{i=0}^{1} \tilde{a}^i(s)$, and for each $s \in \mathcal{S}$ and $i \in \{0, 1\}$,

$$\overline{v} - \xi \geq \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta\overline{v} > \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta\Big\{\sum_{s' \in S} w^{i'}(s')^{\sigma} Q(s|s')\Big\}^{\frac{1}{\sigma}} > w^i(s) + \xi.$$

Set:

$$\overline{D}(s, q^V) = \sum_{i=0}^{1} q^{V,i}\phi^i(q^{V,i}, s), \qquad \phi^i(q^{V,i}, s) = \begin{cases} \overline{v} & q^{V,i} \geq 0 \\ \underline{v} & q^{V,i} < 0 \end{cases}$$

and

$$\underline{\underline{D}}(s, q^V) = \sum_{i=0}^{1}\{q^{V,i}\psi^i(q^{V,i}, s) + |q^{V,i}|\xi\}, \qquad \psi^i(q^{V,i}, s) = \begin{cases} w^i(s) & q^{V,i} \geq 0 \\ \overline{v} & q^{V,i} < 0. \end{cases}$$

Given $q^{V\prime} = \{q^{V\prime,i}(s')\}$, define $\psi(q^{V\prime}) = \{\psi^i(q^{V\prime,i}(s'), s)\}$ and note that the above definitions imply for each s and $q^{V\prime}$, $H[s, \tilde{a}(s), \psi(q^{V\prime})] \geq 0$.

$\mathbf{B}(D)$ is given by, for all $(s, q^V) \in \mathcal{S} \times \mathcal{Y}$,

$$\mathbf{B}(D)(s, q^V) = \inf_{\mathcal{Q}} \sup_{\mathcal{P}} \sum_{i=0,1} (q^{V,i} + q^{H,i}) \left\{ \frac{1-\delta}{1-\mu}(a^i)^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} v^{i\prime}(s')^{\sigma} Q(s|s') \right)^{\frac{1}{\sigma}} \right\}$$
$$- \sum_{i=0,1} q^{H,i} w^i(s) - q^{H,2} \left(\sum_{i=0,1} a^i - \gamma(s) \right) - \delta \sum_{s' \in \mathcal{S}} q^{V\prime}(s') \cdot v'(s') + \delta \sum_{s' \in \mathcal{S}} D(s', q^{V\prime}(s')).$$

Setting $D = \overline{D}$, using the definition of $\underline{v}$ and $\overline{v}$ and noting that the dual variables $(q^H, q^{V\prime})$ can always be chosen equal to 0 in the infimum, we have $\mathbf{B}(\overline{D})(s, q^V) \leq \overline{D}(s, q^V)$. On the other hand, setting $D = \underline{\underline{D}}$ and noting that for any s and choice of $(q^H, q^{V\prime})$, $(\tilde{a}(s), \psi(q^{V\prime}))$ is a feasible choice for the supremum with $H[s, \tilde{a}(s), \psi(q^{V\prime})] \geq 0$, we have:

$$\mathbf{B}(\underline{\underline{D}})(s, q^V) \geq \inf_{\mathcal{Q}} \sum_{i=0,1} (q^{V,i} + q^{H,i}) \left\{ \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} \psi^i(q^{V\prime i}(s'), s')^{\sigma} Q(s|s') \right)^{\frac{1}{\sigma}} \right\}$$
$$- \sum_{i=0,1} q^{H,i} w^i(s) - q^{H,2} \left(\sum_{i=0,1} \tilde{a}^i(s) - \gamma(s) \right)$$
$$\geq \inf_{\mathcal{Q}} \sum_{i=0,1} q^{V,i} \left\{ \frac{1-\delta}{1-\mu}(\tilde{a}^i(s))^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} \psi^i(q^{V\prime i}(s'), s')^{\sigma} Q(s|s') \right)^{\frac{1}{\sigma}} \right\}$$

If $q^{V,i} \geq 0$, then

$$\inf_{q^{V\prime i}} q^{V,i} \left\{ \frac{1-\delta}{1-\mu} (\tilde{a}^i(s))^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} \psi^i(q^{V\prime i}(s'), s')^\sigma Q(s|s') \right)^{\frac{1}{\sigma}} \right\}$$
$$\geq q^{V,i} \left\{ \frac{1-\delta}{1-\mu} (\tilde{a}^i(s))^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} w^i(s')^\sigma Q(s|s') \right)^{\frac{1}{\sigma}} \right\} \geq q^{V,i}(w^i(s) + \xi),$$

with the inequality strict if $q^{V,i} > 0$. Similarly, if $q^{V,i} < 0$, then

$$\inf_{q^{V\prime i}} q^{V,i} \left\{ \frac{1-\delta}{1-\mu} (\tilde{a}^i(s))^{1-\mu} + \delta \left(\sum_{s' \in \mathcal{S}} \psi^i(q^{V\prime i}(s'), s')^\sigma Q(s|s') \right)^{\frac{1}{\sigma}} \right\}$$
$$\geq q^{V,i} \left\{ \frac{1-\delta}{1-\mu} (\tilde{a}^i(s))^{1-\mu} + \delta \overline{v} \right\} > q^{V,i}(\overline{v} - \xi).$$

Thus, for all $(s, q^V) \in \mathcal{S} \times \mathcal{C}$, $\mathbf{B}(\underline{\underline{D}})(s, q^V) > \underline{\underline{D}}(s, q^V)$. The continuity of each $\mathbf{B}(\underline{\underline{D}})(s, \cdot)$ follows from the assumptions on $\underline{\underline{D}}$ and an argument of Rockafellar and Wets (1998) (Theorem 1.17, p. 16-17). The continuity of each $\underline{\underline{D}}(s, \cdot)$ and $\mathbf{B}(\underline{\underline{D}})(s, \cdot)$ and the compactness of $\mathcal{C}$ then implies that there is a $\varepsilon > 0$ such that for all $(s, q^V) \in \mathcal{S} \times \mathcal{C}$, $\mathbf{B}(\underline{\underline{D}})(s, q^V) > \underline{\underline{D}}(s, q^V) + \varepsilon$ as required. □

Bounding value functions for Example 2 The verification of Assumption 3 is very similar to Example 1. The main differences are in showing that $\mathbf{B}(\underline{\underline{D}}) > \underline{\underline{D}} + \varepsilon$, $\varepsilon > 0$, on $\mathcal{C}$. We detail the steps below. In the deterministic version of the default model, $\mathbf{B}(D)$ takes the form:

$$\mathbf{B}(D)(q^K, q^V) = \inf_{q^{K\prime}, q^{V\prime}, q^H} \sup_{k, a, v'} - q^K \cdot k + q^V \cdot (f(a) + \delta v') + q^{H,1} \left(f^1(a^1) + \delta v^{1\prime} - w(k) \right)$$
$$-\delta q^{V\prime} \cdot v' + \left(\delta q^{K\prime} + q^{H,2} \right) \left(\gamma(k) - \sum_{i=0,1} a^i \right) + \delta D \left(q^{V\prime}, q^{K\prime} \right).$$

Let $D = \underline{\underline{D}}$ and define $\tilde{a}(q^{K\prime}) = (\tilde{a}^0(q^{K\prime}), \tilde{a}^1)$, with $\tilde{a}^0(q^{K\prime}) = \underline{a}^0$ if $q^{K\prime} > 0$ and $\tilde{a}^0(q^{K\prime}) = \overline{a}^0$ otherwise. Note that given (q^K, q^V) and $(q^H, q^{K\prime}, q^{V\prime})$, $(\bar{k}, \tilde{a}(q^{K\prime}), \psi(q^{V\prime}))$ is a feasible choice for the supremum that satisfies the no default constraint. Also, for all possible $q^{K\prime}$, the component $\delta q^{K\prime}(\gamma(\bar{k}) - \sum_{i=0,1} \tilde{a}^i(q^{K\prime}))$ in the objective function exactly offsets $\delta q^{K'} \psi^K(q^{K\prime}) + |q^K| \xi^K$, the K component of $\underline{\underline{D}}$. Consequently,

$$\mathbf{B}(\underline{\underline{D}})(q^K, q^V) \geq \inf_{q^{K\prime}, q^{V\prime}} -q^K \bar{k} + q^V \cdot \left(f(\tilde{a}(q^{K\prime})) + \delta \psi^i(q^{V\prime i}) \right).$$

Fix $(q^K, q^V) \in \mathcal{C}$. The conditions placed on $\underline{a}^0$, $\overline{a}^0$ and $\tilde{a}^1$ in the main text and the same line of argument used in the preceding example establishes that: $\inf_{q^{K\prime}, q^{V\prime}} \ q^V \cdot$

$\big(f(\tilde{a}(q^{K\prime})) + \delta\psi^i(q^{V\prime i})\big) \geq \sum_{i=0,1}\{\psi^{V,i}(q^{V,i}) + |q^{V,i}|\xi^V\}$, with the inequality strict whenever $q^V \neq 0$. Also:

$$-q^K\bar{k} \geq q^K\psi^K(q^K) - |q^K|\xi^K.$$

Where the last inequality is strict whenever $q^K \neq 0$: if $q^K < 0$, then $(-q^K)\bar{k} > 0 > \psi^K(q^K) - |q^K|\xi^K = q^K\xi^K$; if $q^K > 0$, then $-q^K\bar{k} > q^K\psi^K(q^K) - |q^K|\xi^K = -q^K\bar{k} - q^K\xi^K$. Hence, for all $y \in \mathcal{C}$, $\mathbf{B}(\underline{\underline{D}})(y) > \underline{\underline{D}}(y)$. As before, the continuity of each $\mathbf{B}(\underline{\underline{D}})(\cdot)$ follows from the assumptions on $\underline{\underline{D}}$ and an argument of Rockafellar and Wets (1998). The continuity of each $\underline{\underline{D}}(\cdot)$ and $\mathbf{B}(\underline{\underline{D}})(\cdot)$ and the compactness of $\mathcal{C}$ then implies that there is a $\varepsilon > 0$ such that for all $y \in \mathcal{S} \times \mathcal{C}$, $\mathbf{B}(\underline{\underline{D}})(y) > \underline{\underline{D}}(y) + \varepsilon$ as required. □

E Proofs for Section 8

Proof of Proposition 8. Equality of values follows from the proof of Luenberger (1969), Theorem 2, p. 221, following a small extension to accommodate equality constraints. If π^* solves (AP) and $\mathscr{L}$ admits a saddle point, then, again by a small extension to the proof of Luenberger (1969), Theorem 2, p. 221, there is a θ_0^* that attains the infimum in (IS) and is such that π^* maximizes $\mathscr{L}$. The result then follows from Proposition 5. □

Proof of Proposition 9. Part (i). Since for each $\theta^* \in \mathscr{Q}^*$, every element of $\mathscr{P}^*$ is maximal for $\mathscr{L}(\cdot, \theta^*)$ and since all elements of $\mathrm{argmax}_{\mathscr{P}}\,\mathscr{L}(\cdot, \theta^*)$ share a plan α^*, it follows that all elements of $\mathscr{P}^*$ share a plan α^*. That α^* is a solution for (P) then follows from Luenberger (1969), Theorem 2, p. 221 and Proposition 1. That α^* is the unique solution of (P) follows from Proposition 1 and the fact that all solutions to (AP) belong to $\mathscr{P}^*$ and, hence, all share the plan α^*. Part (ii). Suppose $(\pi, \theta) = (\alpha, \{k_t, v_t\}_{t=0}^{\infty})$ satisfies the condition in part (ii) of the proposition. Then, by Proposition 5, (π, θ) solves (IS). In addition, $\theta \in \mathscr{Q}^*$ and since $\pi \in \mathrm{argmax}_{\mathscr{P}}\,\mathscr{L}$ it follows that $\alpha = \alpha^*$ and is optimal for (P). □

Proof of Proposition 10. Condition (i) and the weak duality inequality imply $F[s_0, V(s_0, \hat{\alpha})] \geq D_0^* \geq P_0^*$. On the other hand, Condition (ii) implies $P_0^* \geq F[s_0, V(s_0, \hat{\alpha})]$ and, hence $\hat{\pi}$ solves (AP) and $D_0^* = P_0^*$. In addition, from Proposition 5, if $(\hat{\pi}, \hat{\theta})$ satisfies Condition (T), then it is a solution to the dual problem (IS). Also, $\mathscr{L}(\hat{\pi}, \hat{\theta}) = \sup_{\pi} \mathscr{L}(\pi, \hat{\theta}) = D_0^* = P_0^* = \inf_{\theta} \mathscr{L}(\hat{\pi}, \theta)$, where the first and second inequalities use the fact that $(\hat{\pi}, \hat{\theta})$ solves the dual, the third uses $D_0^* = P_0^*$ and the fourth the fact that $\hat{\pi}$ solves (AP) and, hence, maximizes $\inf_{\theta} \mathscr{L}(\pi, \theta)$ and attains P_0^*. Thus, $\hat{\pi}$ solves $\max \mathscr{L}(\pi, \hat{\theta})$ and $\hat{\theta}$ solves $\min \mathscr{L}(\hat{\pi}, \theta)$ and $(\hat{\pi}, \hat{\theta})$ is a saddle fof $\mathscr{L}$. □

F Relaxation

We first consider the augmented problem without backward-looking state variables. The relaxed version of this problem is:

$$\sup F[s_0, v_0] \tag{R-AP}$$

subject to $\pi \in \mathscr{P}$ and $\forall t$, s^t,

$$W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), v_{t+1}(s^t)]] \geq v_t(s^t), \tag{40}$$

$$H[s_t, a_t(s^t), v_{t+1}(s^t)] \geq 0. \tag{41}$$

If $F[s_0, \cdot]$ is concave in v_0, for each s, $W^V[s, \cdot, \cdot]$ is jointly concave in (a, m) (recall that by assumption it is increasing in m), $M^V[s, \cdot, \cdot]$ and $H[s, \cdot, \cdot]$ are jointly concave in (a, v') and a Slater condition holds, then equality of primal and dual values and existence of a minimizing multiplier in the dual is established by standard arguments. Assumption 4 gives sufficient conditions for relaxation to leave optimal values and primal solutions unaffected. Below for $x, x' \in \mathbb{R}^n$, we write $x > x'$ if $x \geq x'$ and $x \neq x'$. Also, for $v \in \mathcal{V}^{n_s} \subset \mathbb{R}^{n_s n_v}$ and $d \in \mathbb{R}^{n_v}$, let $v +_{s'} d$ denote the addition of d to the $(s'-1)n_v + 1$ to $s' n_v$ elements of v.

Assumption 4. *For all $(s, a, v) \in \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s}$ and $(s', a', v') \in \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s}$ such that (i) $H[s, a, v], H[s', a', v'] \geq 0$ and (ii) $W^V[s', a', M^V[s', a', v']] > v(s')$, there is a pair of directions $(d_1, d_2) \in \mathbb{R}^{n_v}_+ \times \mathbb{R}^{n_a}$, $d_1 > 0$, such that (i) $(v(s') + d_1, a' + d_2) \in \mathcal{V} \times \mathcal{A}$, (ii) $H[s, a, v +_{s'} d_1] \geq H[s, a, v]$ and $H[s', a' + d_2, v'] \geq H[s', a', v']$, (iii) $W^V[s, a, M^V[s, a, v +_{s'} d_1]] > W^V[s, a, M^V[s, a, v]]$, (iv) $W^V[s', a' + d_2, M^V[s', a' + d_2, v']] \geq v(s') + d_1$.*

Proposition 11. *If (AP) features no backward-looking state variables and Assumption 4 holds, then the optimal value from the relaxed problem (R-AP) equals that from (AP) and any solution to (AP) also solves the relaxed problem. If $F[s_0, \cdot]$ is increasing then, in addition, any solution to the relaxed problem also solves (AP).*

Proof. Let π be feasible for the relaxed problem and suppose at some $\hat{s}^t = (\hat{s}^{t-1}, \hat{s}_t)$, $W^V[\hat{s}_t, a_t(\hat{s}^t), M^V[\hat{s}_t, a_t(\hat{s}^t), v_{t+1}(\hat{s}^t)]] > v_t(\hat{s}^t)$. By Assumption 4, there is a feasible perturbation (d_{1t}, d_{2t}), $d_{1t} > 0$, such that:

$$\begin{aligned} W^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), &M^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), v_t(\hat{s}^{t-1}) +_{\hat{s}_t} d_{1t}]] \\ &> W^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), M^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), v_t(\hat{s}^{t-1})]] \geq v_{t-1}(\hat{s}^{t-1}). \end{aligned}$$

Reset $v_t(\hat{s}^t)$ to $v_t(\hat{s}^t) + d_{1t}$ and $a_t(\hat{s}_t)$ to $a_t(\hat{s}_t) + d_{2t}$. After this adjustment

$$W^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), M^V[\hat{s}_{t-1}, a_{t-1}(\hat{s}^{t-1}), v_t(\hat{s}^{t-1})]] > v_{t-1}(\hat{s}^{t-1}).$$

Repeating the argument at successively shorter histories, there is a (d_{11}, d_{21}), $d_{11} > 0$, such that:

$$W^V[s_0, a_0, M^V[s_0, a_0, v_1 +_{\hat{s}_1} d_{11}]] > W^V[s_0, a_0, M^V[s_0, a_0, v_1]] \geq v_0.$$

Reset v_0 to equal $W^V[s_0, a_0, M^V[s_0, a_0, v_1 +_{\hat{s}_1} d_{11}]]$. Applying this argument at all histories such that $W^V[s_t, a_t(s^t), M^V[s_t, a_t(s^t), v_{t+1}(s^t)]] > v_t(s^t)$ holds, a primal process feasible for (AP) is constructed with initial forward-looking state variable v_0 greater than that of the original process. Since $F[s_0, \cdot]$ is assumed non-decreasing the new primal process has a payoff no less than the original process. Consequently, the optimal payoff from (AP) equals that from the relaxed problem and any solution to (AP) also solves the relaxed problem. On the other hand, if $F[s_0, \cdot]$ is increasing, then the constructed process has a payoff strictly above the original process and so any solution to the relaxed problem must be feasible and, hence, optimal for the original problem. □

The simplest situation in which Assumption 4 is satisfied occurs when H is increasing in its third argument. Then d_2 may be set equal to 0 and $d_1 = W^V[s', a', M^V[s', a', v']] - v(s') > 0$. Since $W^V[s', a', M^V[s', a, v']] \in \mathcal{V}(s)$, (i) is satisfied. Since H is increasing in its third argument (ii) is satisfied. The monotonicity properties of W^V and M^V imply that (iii) holds and (iv) holds by construction. These conditions are satisfied in standard limited commitment problems such as the Epstein-Zin example in Section 3.

The analysis may be extended to problems with backward-looking states. In relaxed problems with backward-looking states, the law of motion for such states is replaced with the inequalities $W^K[k_t(s^{t-1}), s_t, a_t(s^t)] \geq k_{t+1}(s^t)$ (and the law of motion for forward-looking states is relaxed as in (40)). Modify Assumption 4 as in Assumption 5 below.

Assumption 5. *For all* $(k, s, a, v) \in \mathcal{K} \times \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s}$ *and* $(k', s', a', v') \in \mathcal{K} \times \mathcal{S} \times \mathcal{A} \times \mathcal{V}^{n_s}$ *such that (i)* $H[k, s, a, v'], H[k', s', a', v'] \geq 0$ *and (ii) either* $W^V[s', a', M^V[s', a', v']] > v(s')$ *or* $W^K[k, s, a] > k'$*, there is a triple* $(d_0, d_1, d_2) \in \mathbb{R}^{n_k} \times \mathbb{R}^{n_v} \times \mathbb{R}^{n_a}$ *such that (i)* $H[k, s, a, v +_{s'} d_1] \geq H[k, s, a, v]$ *and* $H[k' + d_0, s, a' + d_2, v'] \geq H[k', s', a', v']$ *(ii)* $W^V[s, a, M^V[s, a, v +_{s'} d_1]] > W^V[s, a, M^V[s, a, v]]$*, (iii)* $W^V[s', a' + d_2, M^V[s', a + d_2, v']] > v(s') + d_1$ *and (iv)* $W^K[k, s, a] > k' + d_0$ *and* $W^K[k' + d_0, s', a' + d_2] > W^K[k', s', s']$*.*

The proof of the following proposition is similar to Proposition 11.

Proposition 12. *If (AP) satisfies Assumption 5, then the optimal value from the relaxed problem (with both backward and forward-looking state variables) equals that from (AP) and any solution to (AP) also solves the relaxed problem. If* $F[s_0, \cdot]$ *is increasing then, in addition, any solution to the relaxed problem also solves (AP).*

Proposition 12 is directly applicable to the contracting problem in Cooley, Marimon, and Quadrini (2004). This is a limited commitment problem with default in which the production function is strictly concave, but the outside option affine in capital (e.g. the entrepreneur can sell off some capital after default). The law of motion for capital can be relaxed without affecting the optimal solution (since it will never be optimal to throw resources away, they can always be used to raise consumption of the entrepreneur or the lender). Similarly, the law of motion for utility promises can be relaxed since the constraint H is increasing in such promises (or, since the law of motion is quasi-linear in agent utility promises, it may be substituted out).