

NBER WORKING PAPER SERIES

RECURSIVE CONTRACTS AND ENDOGENOUSLY INCOMPLETE MARKETS

Mikhail Golosov
Aleh Tsyvinski
Nicolas Werquin

Working Paper 22012
<http://www.nber.org/papers/w22012>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

Golosov and Tsyvinski thank the NSF for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Mikhail Golosov, Aleh Tsyvinski, and Nicolas Werquin. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Recursive Contracts and Endogenously Incomplete Markets
Mikhail Golosov, Aleh Tsyvinski, and Nicolas Werquin
NBER Working Paper No. 22012
February 2016
JEL No. C61,E2,E61

ABSTRACT

In this chapter we study dynamic incentive models in which risk sharing is endogenously limited by the presence of informational or enforcement frictions. We comprehensively overview one of the most important tools for the analysis such problems — the theory of recursive contracts. Recursive formulations allow to reduce often complex models to a sequence of essentially static problems that are easier to analyze both analytically and computationally. We first provide a self-contained treatment of the basic theory: the Revelation Principle, formulating and simplifying the incentive constraints, using promised utilities as state variables, and analyzing models with persistent shocks using the first-order approach. We then discuss more advanced topics: duality theory and Lagrange multiplier techniques, models with lack of commitment, and martingale methods in continuous time. Finally, we show how a variety of applications in public economics, corporate finance, development and international economics featuring incomplete risk-sharing can be analyzed using the tools of the theory of recursive contracts.

Mikhail Golosov
Department of Economics
Princeton University
111 Fisher Hall
Princeton, NJ 08544
and NBER
golosov@princeton.edu

Nicolas Werquin
Toulouse School of Economics
21 allée de Brienne
31000 Toulouse
nwerquin@gmail.com

Aleh Tsyvinski
Department of Economics
Yale University
Box 208268
New Haven, CT 06520-8268
and NBER
a.tsyvinski@yale.edu

1 Introduction

Dynamic incentive problems are ubiquitous in macroeconomics. The design of social insurance programs by governments, long-run relationships between banks and entrepreneurs, informal insurance contracts against idiosyncratic shocks provided in village economies, sovereign borrowing and lending between countries can all be understood using the theory of dynamic incentives. These models have been widely used in macroeconomics, public economics, international macro, finance, development, or political economy, both for explaining existing patterns in the data and for normative policy analysis. The unifying feature of these models is that, at their essence, they study endogenously incomplete markets, i.e., environments in which risk sharing is constrained by (informational or enforcement) frictions and insurance arrangements arise endogenously.

One of the most important tools used for studying dynamic incentive problems is the theory of recursive contracts. Recursive formulations allow one to reduce often complex models to a sequence of essentially static problems that are easier to analyze both analytically and computationally. This substantially simplifies the analysis and the characterization of the optimal insurance arrangements in rich and realistic environments. The goal of this chapter is to provide an overview of the theory of recursive contracts and give a number of examples of application. The analysis in the theoretical part is self-contained; whenever a textbook approach is not directly applicable (e.g., when the assumptions needed to apply the recursive techniques in Stokey et al. (1989) are not met), we provide the necessary mathematical background. We also discuss the strengths and weaknesses of several alternative approaches to solving dynamic incentive problems that emerged in the literature. In the last part of the chapter we show how the methods of recursive contracts can be used in a variety of applications.

Our paper is organized as follows. Section 2 considers a prototypical dynamic incentive problem – insurance against privately observable idiosyncratic taste shocks under perfect commitment by the principal. The goal of this section is to provide an example of a self-contained, rigorous and relatively general treatment of a dynamic incentive problem. We also use this economy in subsequent sections to illustrate other approaches to the analysis of dynamic incentive problems. In Section 2 we highlight the three main steps in the analysis: first, applying the Revelation Principle to set up a mechanism design problem with incentive constraints; second, simplifying this problem by focusing on one-shot incentive constraints; and third, writing this problem recursively using “promised utilities” as state variables. We then show how this recursive formulation can be used to characterize the properties of the optimal insurance arrangements in our economy. We derive general features of the optimal insurance contract and characterize the long-run behavior of the economy in Section 2.4. We show how to overcome the technical difficulties that arise when the idiosyncratic shocks are persistent in Section 2.5, and conclude by showing how the same techniques can be applied to other dynamic incentive problems, such as moral hazard.

Section 3 considers more advanced topics. We focus on three topics: using Lagrange mul-

multiplier tools in recursive formulations, studying dynamic insurance problems in economies in which the principal has imperfect commitment, and applying martingale techniques to study recursive contracts in continuous time. Section 3.1 discusses the Lagrangian techniques. Using Lagrangians together with the recursive methods of Section 2 greatly expands the class of problems that can be characterized. We first provide an overview of the theory of constrained optimization using Lagrange multipliers, with a particular focus on showing how to use them in the infinite dimensional settings that frequently arise in macroeconomic applications. We then show how to apply these theoretical techniques to incentive problems to obtain several alternative recursive formulations having some advantages relative to those discussed in Section 2. A number of results in this section are new to the dynamic contracts literature. In Section 3.2 we show how to analyze dynamic insurance problems in settings where the principal cannot commit to the contracts. The arguments used to prove simple versions of the Revelation Principle under commitment fail in such an environment; we discuss several ways to generalize it and write a recursive formulation of the mechanism design problem. Our characterization of such problems relies heavily on our analysis of Section 3.1. Finally, in Section 3.3, we show how to analyze a dynamic contracting problem in continuous time using martingale methods and the dynamic programming principle. To keep the analysis self-contained, we start by stating the results of stochastic calculus that we use. Continuous-time methods often simplify the characterization of optimal contracts, allowing for analytical comparative statics and easier numerical analysis of the solution.

Section 4 gives a number of applications of the recursive techniques discussed in Sections 2 and 3 to various environments. We show that these diverse applications share three key features: (i) insurance is endogenously limited by the presence of a friction; (ii) the problem is dynamic; (iii) the recursive contract techniques that we develop in the theoretical sections allow us to derive deep characterizations of these problems. We explain how theoretical constructs such as the incentive constraints and promised utilities can be mapped into concrete economic concepts, and how the predictions of dynamic incentive models can be tested empirically and used for policy analysis. In Section 4.1, we apply the techniques and results of Section 2 to public finance where the endogenous market incompleteness and the limited social insurance arise due to the unobservability of the shocks that agents receive. We derive several central results characterizing the optimal social insurance mechanisms and show how to implement the optimal allocations with a tax and transfer system that arises endogenously, without restricting the system exogenously to a specific functional form. In Section 4.2 we show how recursive techniques can be applied to study the effect of informational frictions on firm dynamics and optimal capital structure. Section 4.3 presents applications of these techniques to study insurance in village economies in developing countries where contracts are limited by enforcement and informational frictions. Section 4.4 discusses applications to international borrowing and lending.

2 A simple model of dynamic insurance

In this section we study a prototypical model of dynamic insurance against privately-observed idiosyncratic shocks. Our goal is to explain the key steps in the analysis and the main economic insights in the simplest setting. The mathematical techniques that we use as well as the economic insights that we obtain extend to many richer and more realistic environments. We discuss examples of such environments in following sections.

2.1 Environment

We consider a discrete-time economy that lasts T periods, where T may be finite or infinite. The economy is populated by a continuum of ex-ante identical agents whose preferences over period- t consumption $c_t \geq 0$ are given by $\theta_t U(c_t)$, where $\theta_t \in \Theta \subset \mathbb{R}_+$ is an idiosyncratic “taste shock” that the individual receives in period t , and U is a utility function.

Assumption 1. *The utility function $U : \mathbb{R}_+ \rightarrow \mathbb{R}$ is an increasing, strictly concave, differentiable function that satisfies the Inada conditions $\lim_{c \rightarrow 0} U'(c) = \infty$ and $\lim_{c \rightarrow \infty} U'(c) = 0$.*

All agents have the same discount factor $\beta \in (0, 1)$. In each period the economy receives e units of endowment which can be freely transferred between periods at rate β .

The idiosyncratic taste shocks are stochastic. We use the notations $\theta^t = (\theta_1, \dots, \theta_t) \in \Theta^t$ to denote a history of realizations of shocks up to period t and $\pi_t(\theta^t)$ to denote the probability of realization of history θ^t . We assume that the law of large numbers holds so that $\pi_t(\theta^t)$ is also the measure of individuals who experienced history θ^t .¹ An individual privately learns his taste shock θ_t at the beginning of period t . Thus, at the beginning of period t an agent knows his history θ^t of current and past shocks, but not his future shocks. This implies that his choices in period t , and more generally all the period- t random variables x_t that we encounter, can only be a function of this history.

Some parts of our analysis use results from probability theory and require us to be more formal about the probability spaces that we use. A standard way to formalize these stochastic processes is as follows.² Let Θ^T be the space of all histories θ^T and let π_T be a probability measure over the Borel subsets $\mathcal{B}(\Theta^T)$ of Θ^T . Thus, $(\Theta^T, \mathcal{B}(\Theta^T), \pi_T)$ forms a probability space. Any period- t random variable is required to be measurable with respect to $\mathcal{B}(\Theta^t)$, that is, for any Borel subset M of \mathbb{R} , $x_t^{-1}(M) = B \times \Theta^{T-t}$, where B is a Borel subset of Θ^t . This formalizes the intuition that the realization of shocks in future periods is not known as of period t .

Until Section 2.5.1 we make the following assumptions about the idiosyncratic taste shocks:

Assumption 2. *The set $\Theta \subset \mathbb{R}_+$ of taste shocks is discrete and finite with cardinality $|\Theta|$. Agents' shocks evolve according to a first-order Markov process, that is, the probability of drawing*

¹The assumption that the law of large numbers holds can be justified formally (see Uhlig (1996), Sun (2006)).

²See Stokey et al. (1989, Chapter 7) for a review of the measure-theoretic apparatus.

type θ_t in period t depends only on the period- $(t-1)$ type:

$$\pi_t(\theta_t | \theta^{t-1}) = \pi(\theta_t | \theta_{t-1}), \quad \forall \theta^{t-1} \in \Theta^{t-1}, \theta_t \in \Theta,$$

where θ_{t-1} is the last component of θ^{t-1} .

We use the notation $\pi_t(\theta^t | \theta^s)$ for $t > s$ to denote the probability of realization of history θ^t up to period t conditional on the realization of history θ^s up to period s , with a convention that $\pi_t(\theta^t | \theta^s) = 0$ if the first s elements of θ^t are not θ^s (history θ^t in period t cannot occur if θ^s was not realized up to period s). We use θ_s^t to denote $(\theta_s, \dots, \theta_t)$. Finally we index the elements of Θ by the subscript (j) for $j \in \{1, \dots, |\Theta|\}$, and assume that $\theta_{(1)} < \theta_{(2)} < \dots < \theta_{(|\Theta|)}$.

We consider the problem of a social planner who chooses consumption allocations³ $c_t : \Theta^t \rightarrow \mathbb{R}_+$ to maximize agents' ex-ante expected utility and has the ability to commit to such allocations in period 0. At this stage we are agnostic about who this planner is. One can think of it as a government that provides insurance to agents, or as some decentralized market arrangement. We study the optimal insurance contract that such a planner can provide given the feasibility constraint and informational constraints. We use the shortcut \mathbf{c} to denote the consumption plan $\{c_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$.

The ex-ante, “period-0” expected utility of all agents is denoted by $U_0(\mathbf{c})$ and is given by

$$U_0(\mathbf{c}) \equiv \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t U(c_t) \right] = \sum_{t=1}^T \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t) \theta_t U(c_t(\theta^t)). \quad (1)$$

Here \mathbb{E}_0 represents the (unconditional) expectation at time 0, before the first-period type θ_1 is known. Under our assumption that resources can be freely transferred between periods, the resource constraint is

$$\mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} c_t \right] \leq \frac{1 - \beta^T}{1 - \beta} e. \quad (2)$$

Note that to write the left hand side of this feasibility constraint we again implicitly invoke the law of large numbers.

When the realizations of the taste shocks are observable by the planner, this problem can easily be solved explicitly. Let $\zeta > 0$ be the Lagrange multiplier on the feasibility constraint. The optimal allocation \mathbf{c}^{fb} in the case where shocks are observable (the “first best” allocation) is a solution to

$$\theta_t U'(c_t^{\text{fb}}(\theta^t)) = \zeta, \quad \forall t \geq 1, \forall \theta^t \in \Theta^t. \quad (3)$$

It is immediate to see that this equation implies that $c_t^{\text{fb}}(\theta^t)$ is independent of period t or the past history of shocks θ^{t-1} , and only depends on the current realization of the shock θ_t . That is, the informationally-unconstrained optimal insurance in this economy gives agents with a higher

³Formally, c_t is a random variable over the probability space $(\Theta^T, \mathcal{B}(\Theta^T), \pi_T)$ that is measurable with respect to $\mathcal{B}(\Theta^t)$.

realization of the shock θ in any period (hence whose current marginal utility of consumption is higher) more consumption than agents with a lower realization of a taste shock.

2.2 The Revelation Principle and incentive compatibility

We are interested in understanding the properties of the best insurance arrangements that a planner can provide in the economy with private information. This insurance can be provided by many different mechanisms: the agents may be required to live in autarky and consume their endowment, or may be allowed to trade assets, or may be provided with more sophisticated arrangements by the planner. A priori it is not obvious how to set up the problem of finding the best mechanism to provide the highest utility to agents. This problem simplifies once we apply the results of the mechanism design literature, in particular the Revelation Principle. Textbook treatments of the Revelation Principle are widely available (see, e.g. Chapter 23 in Mas-Colell et al. (1995)). Here we outline the main arguments behind the Revelation Principle in our context. This overview is useful both to keep the analysis self-contained and to emphasize subtleties that emerge in using the Revelation Principle once additional frictions, such as lack of commitment by the planner, are introduced.

Hurwicz (1960, 1972) provided a general framework to study various arrangements of allocation provision in environments with private information. He showed that such arrangements can be represented as abstract communication mechanisms. Consider an arbitrary message space M that consists of a collection of messages m . Each agent observes his shock θ_t and sends a (possibly random) message $m_t \in M$ to the principal. The agent's reporting strategy in period t is a map $\tilde{\sigma}_t : \Theta^t \rightarrow \Delta(M)$. The planner in turn chooses a (possibly stochastic) allocation rule $\tilde{c}_t : M^t \rightarrow \Delta(\mathbb{R}_+)$, where $\Delta(\mathbb{R}_+)$ denotes a probability measure on \mathbb{R}_+ . The strategies $\tilde{\sigma} = \{\tilde{\sigma}_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$ and $\tilde{c} = \{\tilde{c}_t(m^t)\}_{t \geq 1, m^t \in M^t}$ induce a measure over the consumption paths $\{c_t\}_{t \geq 1} \in \mathbb{R}_+^T$, which we denote by $\tilde{c} \circ \tilde{\sigma}$. The expected utility of each agent is then equal to $\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}} \left[\sum_{t=1}^T \beta^{t-1} U(c_t) \right]$, where the superscript in $\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}}$ means that the expectation is computed using the probability distribution $\tilde{c} \circ \tilde{\sigma}$ over the paths $\{c_t\}_{t \geq 1}$. The strategy $\tilde{\sigma}$ is *incentive compatible* for the agent if^{4,5}

$$\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}} \left[\sum_{t=1}^T \beta^{t-1} \theta_t U(c_t) \right] - \mathbb{E}^{\tilde{c} \circ \tilde{\sigma}'} \left[\sum_{t=1}^T \beta^{t-1} \theta_t U(c_t) \right] \geq 0, \quad \forall \tilde{\sigma}'. \quad (4)$$

⁴When T is allowed to be infinite, these sums may not be well-defined for all σ , and we require (4) to hold as $\limsup_{T \rightarrow \infty}$.

⁵Note that the constraints (4) also include all the constraints that ensure that $\tilde{\sigma}$ is optimal after any history t, θ^t , i.e., $\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}} \left[\sum_{s=t}^T \beta^{s-t} \theta_s U(c_s) \mid \theta^t \right] - \mathbb{E}^{\tilde{c} \circ \tilde{\sigma}'} \left[\sum_{s=t}^T \beta^{s-t} \theta_s U(c_s) \mid \theta^t \right] \geq 0, \quad \forall \tilde{\sigma}'$.

A mechanism $\tilde{\Gamma} = (M, \tilde{\mathbf{c}} \circ \tilde{\sigma})$ is incentive compatible if it satisfies (4), and *feasible* if it satisfies

$$\mathbb{E}^{\tilde{\mathbf{c}} \circ \tilde{\sigma}} \left[\sum_{t=1}^T \beta^{t-1} c_t \right] \leq \frac{1 - \beta^T}{1 - \beta} e. \quad (5)$$

The key insight behind the Revelation Principle is that any outcome $\tilde{\mathbf{c}} \circ \tilde{\sigma}$ of an incentive compatible and feasible mechanism can be achieved as the outcome of a *direct truthtelling* mechanism, in which agents report their types directly to the principal. Define a direct mechanism as a reporting strategy $\sigma_t : \Theta^t \rightarrow \Theta$. Define a truthtelling strategy σ^{truth} as $\sigma_t^{truth}(\theta^{t-1}, \theta) = \theta$ for all θ^{t-1}, θ . The key observation is that there exists $\mathbf{c} = \{c_t\}_{t \geq 1}$, with $c_t : \Theta^t \rightarrow \Delta(\mathbb{R}_+)$ for each t , such that the (induced) measure $\mathbf{c} \circ \sigma^{truth}$ replicates the measure $\tilde{\mathbf{c}} \circ \tilde{\sigma}$.⁶

Theorem 1. (*Revelation Principle.*) *The outcome of any incentive compatible and feasible mechanism $\tilde{\Gamma} = (M, \tilde{\mathbf{c}} \circ \tilde{\sigma})$ is also the outcome of an incentive compatible and feasible direct truthful mechanism $\Gamma = (\Theta, \mathbf{c} \circ \sigma^{truth})$.*

Proof. By construction, we have

$$\mathbb{E}^{\mathbf{c} \circ \sigma^{truth}} \left[\sum_{t=1}^T \beta^{t-1} c_t \right] = \mathbb{E}^{\tilde{\mathbf{c}} \circ \tilde{\sigma}} \left[\sum_{t=1}^T \beta^{t-1} c_t \right],$$

so that the truthtelling strategy satisfies (5). Any alternative strategy σ' induces a measure $\mathbf{c} \circ \sigma'$ which replicates the measure $\tilde{\mathbf{c}} \circ \tilde{\sigma}'$ for some strategy $\tilde{\sigma}'$ in the original mechanism. Therefore

$$\mathbb{E}^{\mathbf{c} \circ \sigma^{truth}} \left[\sum_{t=1}^T \beta^{t-1} \theta_t U(c_t) \right] - \mathbb{E}^{\mathbf{c} \circ \sigma'} \left[\sum_{t=1}^T \beta^{t-1} \theta_t U(c_t) \right] \geq 0, \quad \forall \sigma'. \quad (6)$$

This concludes the proof. \square

We can simplify our analysis further by showing that there is no loss of generality in focusing on *deterministic* direct mechanisms, where each history of reports yields a deterministic consumption allocation (rather than a measure) $\mathbf{c}_t^{\det} : \Theta^t \rightarrow \mathbb{R}_+$. We show:

Proposition 1. *For any incentive compatible and feasible direct mechanism $\Gamma = (\Theta, \mathbf{c} \circ \sigma^{truth})$ there exists an incentive compatible, feasible, deterministic direct mechanism $(\Theta, \mathbf{c}^{\det} \circ \sigma^{truth})$ that achieves the same ex-ante utility.*

Proof. Consider any incentive compatible and feasible, but possibly stochastic, direct mechanism

⁶The proof of this observation is straightforward. For simplicity, assume that $\tilde{\sigma}$ and $\tilde{\mathbf{c}}$ involve randomization over a finite number of elements after each history and let $\tilde{\sigma}^t(m^t|\theta^t)$ be the probability that agent θ^t sends a history of messages m^t , and $\tilde{c}_t(x|m^t)$ be the probability that the principal delivers consumption x to an agent with a reported history m^t . Then c_t is simply defined by $c_t(x|\theta^t) \equiv \sum_{m^t} \tilde{c}_t(x|m^t) \tilde{\sigma}^t(m^t|\theta^t)$. Given this definition of \mathbf{c} the payoff of any strategy $\tilde{\sigma}$ in the original mechanism $(M, \tilde{\mathbf{c}} \circ \tilde{\sigma})$ can be replicated by a strategy σ .

$\Gamma = (\Theta, \mathbf{c} \circ \sigma^{truth})$. Define a deterministic consumption allocation $\mathbf{c}_t^{\det} : \Theta^t \rightarrow \mathbb{R}_+$ implicitly by

$$U(\mathbf{c}_t^{\det}(\theta^t)) = \mathbb{E}^{\mathbf{c} \circ \sigma^{truth}} [U(c_t) | \theta^t], \quad \forall t \geq 1, \theta^t \in \Theta^t, \quad (7)$$

where the right hand side is the expected consumption given at time t under the mechanism Γ to the agent who reports the history θ^t . Since U is concave by Assumption 1, Jensen's inequality implies that

$$\mathbb{E}^{\mathbf{c} \circ \sigma^{truth}} [c_t | \theta^t] \geq \mathbf{c}_t^{\det}(\theta^t), \quad \forall \theta^t,$$

hence the mechanism $(\Theta, \mathbf{c}^{\det} \circ \sigma^{truth})$ is feasible. By construction, we have that for all t, θ^t , $\mathbb{E}^{\mathbf{c} \circ \sigma^{truth}} [U(c_t) | \theta^t] = \mathbb{E}^{\mathbf{c}^{\det} \circ \sigma^{truth}} [U(c_t) | \theta^t]$, since the conditional expectation in (7) implies that for any report the agent receives the same utility under \mathbf{c} and under \mathbf{c}^{\det} , hence the mechanism is incentive compatible. This concludes the proof. \square

With a slight abuse of notation we will use $\mathbf{c} = \{c_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$ instead of \mathbf{c}_t^{\det} . The incentive constraint in the deterministic direct mechanism can be written simply as

$$\sum_{t=1}^T \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t) \theta_t [U(c_t(\theta^t)) - U(c_t(\sigma^t(\theta^t)))] \geq 0 \text{ for all } \forall \sigma'. \quad (8)$$

The proof of the Revelation Principle requires very few assumptions except the ability of the social planner to commit to the long-term contract in period 0. Theorem 1 and Proposition 1 are very powerful results that provide a simple way to find informationally-constrained optimal allocations. In particular, such allocations are a solution to the problem

$$V(e) \equiv \sup_{\mathbf{c}} \sum_{t=1}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t) \theta_t U(c_t(\theta^t)) \quad (9)$$

subject to (2), (8).

If the supremum of this problem is attained by some vector \mathbf{c}^* , any insurance arrangement in which agents consume \mathbf{c}^* in equilibrium is efficient.

In Sections 2.3, 2.4 and 2.5 we focus on describing general methods to solve the maximization problem defined in (9). We give examples of specific insurance arrangements when discussing various applications in Section 4.

2.3 Recursive formulation with i.i.d. shocks

The analysis of the solution to problem (9) is significantly simplified if shocks are independently and identically distributed. In more general Markov settings, many of the same arguments continue to hold but they are more cumbersome, and analytical results are more difficult to obtain. For this reason we first focus on i.i.d. shocks and discuss general Markov shocks in Section 2.5.

Assumption 3. Types θ are independent and identically distributed, that is, $\pi_t(\theta_t | \theta^{t-1}) = \pi(\theta_t)$. Without loss of generality we assume that $\mathbb{E}[\theta] = \sum_{\theta \in \Theta} \pi(\theta) \theta = 1$.

2.3.1 Main ideas in a finite-period economy

In an economy with a finite number of periods, the maximization problem (9) is defined over a closed and bounded set, because the feasibility constraint imposes that for all t, θ^t , we have $0 \leq c_t(\theta^t) \leq \beta \frac{1-\beta^T}{1-\beta} (\beta \min_{\theta \in \Theta} \pi(\theta))^{-T} e$. In finite dimensions closed and bounded sets are compact and therefore by Weierstrass' theorem the maximum of problem (9) is achieved, so that we can replace the “sup” with a “max”. Moreover, it is easy to see that at the optimum the feasibility constraint must hold with equality.

We want to simplify the set of the incentive constraints in problem (9). Equation (8) should hold for all possible reporting strategy σ' . The set of such strategies is large; it consists of all strategies in which an agent misreports his type in some (possibly all) states only in period 1, all strategies in which he misreports his types in some states in periods 1 and 2, and so on. Most of these constraints are redundant. We say that σ'' is a one-shot deviation strategy if $\sigma''(\theta^{t-1}, \theta_t) \neq \theta_t$ for only one θ^t . It turns out that if (8) is satisfied for one-shot deviations, it is satisfied for all deviations in a finite period economy. Formally, we can write a *one-shot* incentive constraint (see Green (1987)) as: for all $\theta^{t-1}, \theta, \hat{\theta}$,

$$\begin{aligned} & \theta U(c_t(\theta^{t-1}, \theta)) + \beta \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi_{t+s}(\theta^{t+s}) \theta_{t+s} U(c_{t+s}(\theta^{t-1}, \theta, \theta_{t+1}^{t+s})) \\ & \geq \theta U(c_t(\theta^{t-1}, \hat{\theta})) + \beta \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi_{t+s}(\theta^{t+s}) \theta_{t+s} U(c_{t+s}(\theta^{t-1}, \hat{\theta}, \theta_{t+1}^{t+s})). \end{aligned} \quad (10)$$

Proposition 2. Suppose that T is finite and shocks are i.i.d. An allocation \mathbf{c} satisfies (8) if and only if it satisfies (10).

Proof. That (8) implies (10) is clear, since (10) considers a strict subset of the possible deviations. To show the converse, consider any reporting strategy σ' . Suppose that the last period in which the agent misreports his type is period t . By (10), for any θ_t the agent gets higher utility from reporting his type truthfully in that period than from deviating. Therefore, the strategy σ'' which coincides with σ' in the first $t-1$ periods and reveals types truthfully from period t onwards gives higher utility to the agent than σ' . Backward induction then implies that truthtelling gives higher utility than σ' , establishing the result. \square

Proposition 2 simplifies the maximization problem (9) by replacing the constraint set (8) with a smaller number of constraints (10). This simplified problem is still too complicated to be solved directly. We next show how to re-write this problem recursively to reduce it to a sequence of essentially static problems which can be easily analyzed analytically and computationally.

We take several intermediate steps to re-write constraints (2) and (10). First, observe that the constraint set defined by equations (2) and (10) is not convex. Although much of the analysis can be done for a non-convex maximization problem, we can obtain convexity by a simple change of variables: instead of choosing consumption $c_t(\theta^t)$ we can choose utils $u_t(\theta^t) \equiv U(c_t(\theta^t))$. The resource cost of providing u units of utils is $C(u) = U^{-1}(u)$, where the cost function C , defined on the range of U , is increasing, differentiable and strictly convex by Assumption 1. Let \underline{u} and \bar{u} be the (possibly infinite) greatest lower bound and smallest upper bound of U . Observe that $\lim_{u \rightarrow \underline{u}} C(u) = 0$ and $\lim_{u \rightarrow \bar{u}} C(u) = \infty$. We use \mathbb{U} to denote the domain of C , which is (\underline{u}, \bar{u}) if the utility function is unbounded below and $[\underline{u}, \bar{u})$ if it is bounded below. Given this change of variables, the incentive constraint (10) becomes linear in $\mathbf{u} = \{u_t(\theta^t)\}_{t, \theta^t}$, while the resource constraint becomes

$$\mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} C(u_t) \right] \leq \frac{1 - \beta^T}{1 - \beta} e,$$

which defines a convex set of feasible \mathbf{u} .

The second simplification is to define a *continuation (or promised) utility* variable

$$v_t(\theta^t) \equiv \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta^{t+s} | \theta^t) \theta_{t+s} u_{t+s}(\theta^t, \theta^{t+s}). \quad (11)$$

Using repeated substitution we get

$$v_t(\theta^t) = \sum_{\theta \in \Theta} \pi(\theta) [\theta u_{t+1}(\theta^t, \theta) + \beta v_{t+1}(\theta^t, \theta)], \quad \forall \theta^t, \quad (12)$$

where we use the convention $v_T = 0$. Given this definition we can re-write the incentive constraints (10) as

$$\theta u_t(\theta^{t-1}, \theta) + \beta v_t(\theta^{t-1}, \theta) \geq \theta u_t(\theta^{t-1}, \hat{\theta}) + \beta v_t(\theta^{t-1}, \hat{\theta}), \quad \forall \theta^{t-1}, \theta, \hat{\theta}. \quad (13)$$

We are now ready to simplify our analysis by observing that while the original maximization problem does not have an obvious recursive structure, its *dual* does. Our arguments imply that the maximization problem (9) can be re-written as the maximization of the planner's objective over $(\mathbf{u}, \mathbf{v}) = (\{u_t(\theta^t)\}_{t, \theta^t}, \{v_t(\theta^t)\}_{t, \theta^t})$ subject to the constraints (2), (12) and (13). Let $(\mathbf{u}^*, \mathbf{v}^*)$ be the solution to that problem and v_0 be the value of the maximum. Then, by standard duality arguments, $(\mathbf{u}^*, \mathbf{v}^*)$ also minimizes the cost of providing (\mathbf{u}, \mathbf{v}) subject to the incentive compatibility constraints and the "promise-keeping constraint"

$$\mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t u_t(\theta^t) \right] = v_0. \quad (14)$$

Using the definition of v_1 (θ^1), this constraint can be rewritten as

$$v_0 = \sum_{\theta \in \Theta} \pi(\theta) [\theta u_1(\theta) + \beta v_1(\theta)]. \quad (15)$$

Define the set $\Gamma(v_0)$ as

$$\Gamma(v_0) = \{(\mathbf{u}, \mathbf{v}) : (12), (13), (15) \text{ hold}\}. \quad (16)$$

We thus obtain that $(\mathbf{u}^*, \mathbf{v}^*)$ is the solution to

$$K_0(v_0) \equiv \max_{(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)} \mathbb{E}_0 \left[- \sum_{t=1}^T \beta^{t-1} C(u_t) \right]. \quad (17)$$

The key simplification allowed by this formulation is that it can be easily solved using recursive techniques. Let $K_T(\cdot) \equiv -C(\cdot)$, which has domain $\mathbb{V}_T = \mathbb{U}$. Define the functions K_t for $t > 0$ and their domains \mathbb{V}_t recursively by

$$K_t(v) \equiv \max_{\{(u(\theta), w(\theta))\}_{\theta \in \Theta}} \sum_{\theta \in \Theta} \pi(\theta) [-C(u(\theta)) + \beta K_{t+1}(w(\theta))] \quad (18)$$

subject to the promise keeping constraint:

$$v = \sum_{\theta \in \Theta} \pi(\theta) [\theta u(\theta) + \beta w(\theta)], \quad (19)$$

and the incentive compatibility constraint:

$$\theta u(\theta) + \beta w(\theta) \geq \theta u(\hat{\theta}) + \beta w(\hat{\theta}), \quad \forall \theta, \hat{\theta}, \quad (20)$$

and

$$u(\theta) \in \mathbb{U}, \quad w(\theta) \in \mathbb{V}_{t+1}.$$

Equation (18) defines the domain of K_t , denoted by \mathbb{V}_t . It is easy to verify that it is either $\left[\frac{1-\beta^{T-t+1}}{1-\beta} \underline{u}, \frac{1-\beta^{T-t+1}}{1-\beta} \bar{u} \right)$ or $\left(\frac{1-\beta^{T-t+1}}{1-\beta} \underline{u}, \frac{1-\beta^{T-t+1}}{1-\beta} \bar{u} \right]$, depending on whether the utility function is bounded below or not. It is easy to see that the function K_0 defined in (17) satisfies (18) for $t = 0$. Standard arguments establish that K_t is a continuous, strictly decreasing, strictly concave, and differentiable function. For any value $v \in \mathbb{V}_t$, let $\vec{u}_{v,t} = \{u_{v,t}(\theta)\}_{\theta \in \Theta}$ and $\vec{w}_{v,t} = \{w_{v,t}(\theta)\}_{\theta \in \Theta}$ denote the solution (i.e., the argmax) of the Bellman equation (18). We call $(\vec{u}_{v,t}, \vec{w}_{v,t})$ the *policy functions* of the Bellman equation. Given our assumption that C is strictly convex, these policy functions are unique for each v, t .

We can now describe how to find the solution to (17). The main simplification comes from the fact that if we know the optimal value $v_t^*(\theta^t)$ after any history θ^t , we can find the optimal allocations in the nodes following θ^t without having to know the optimal allocations in any other node. We start with $t = 0$. Since $K_0(v)$ is (minus) the amount of resources required to achieve

the expected utility v , the initial value v_0 must satisfy $K_0(v_0) = -\frac{1-\beta^T}{1-\beta}e$. The constrained-optimal utility allocation in period 1 for an agent with shock θ_1 , $u_1^*(\theta_1)$, is then given by the policy function $u_{v_0,1}(\theta_1)$, and his expected utility starting from period 2, $v_1^*(\theta_1)$, is $w_{v_0,1}(\theta_1)$. The optimal utility allocation in period two for a history of shocks (θ_1, θ_2) , $u_2^*(\theta_1, \theta_2)$, is then given by $u_{w_{v_0,1}(\theta_1),2}(\theta_2)$, and similarly for $v_2^*(\theta_1, \theta_2)$. This way we can use forward induction to find the solution to (9), $(\mathbf{u}^*, \mathbf{v}^*)$. We say that the solution $(\mathbf{u}^*, \mathbf{v}^*)$ is *generated by the policy functions* of the Bellman equation (18) given v_0 .

2.3.2 Extension to an infinite period economy

In the previous section we showed a simple way to characterize the solution to a dynamic contracting problem recursively when the number of periods is finite. For many applications it is more convenient to work with infinite periods for at least two reasons. One is that many problems do not have a natural terminal period and the assumption of infinite periods is more convenient. The second reason is that the assumption of infinite periods allows us to obtain sharp insights about the economic forces behind the optimal provision of incentives that are more difficult to see in finite-period economies.

The key step in the analysis of Section 2.3.1 consisted of setting up the dual problem (17) and its recursive representation (18). The infinite period analogue of the sequential dual problem is

$$K(v_0) \equiv \sup_{\mathbf{u}} \mathbb{E}_0 \left[- \sum_{t=1}^{\infty} \beta^{t-1} C(u_t) \right] \quad (21)$$

subject to (8), (14).

We now show that the value function K defined in (21) can be written recursively, and that the solution to this recursive formulation can, under some conditions, recover the maximum to our primal problem (9). Let $\bar{v} = \frac{1}{1-\beta}\bar{u}$, $\underline{v} = \frac{1}{1-\beta}\underline{u}$, and $\mathbb{V} = [\underline{v}, \bar{v}]$ if the utility is bounded below and $\mathbb{V} = (\underline{v}, \bar{v})$ otherwise.⁷ We denote by $B(v)$ the set of pairs $(\vec{u}, \vec{w}) = (\{u(\theta)\}_{\theta \in \Theta}, \{w(\theta)\}_{\theta \in \Theta})$ that satisfy the constraints of the recursive problem, i.e.,

$$B(v) \equiv \left\{ (\vec{u}, \vec{w}) \in \mathbb{U}^{|\Theta|} \times \mathbb{V}^{|\Theta|} : (19), (20) \text{ hold} \right\}. \quad (22)$$

We first prove an infinite period analogue of the Bellman equation (18). Several of the arguments are based on those in Farhi and Werning (2007).

⁷In our benchmark taste shock model it is easy to find the domain of K that we denote by \mathbb{V} . Any constant consumption sequence is incentive compatible, so it satisfies (8). Since the consumption set is bounded below by 0, the greatest lower bound for the set \mathbb{V} must be $\underline{v} = U(0) \sum_{\theta \in \Theta} \pi(\theta) \theta = \frac{1}{1-\beta}\underline{u}$, where we used the normalization $\mathbb{E}\theta = 1$. If $U(0)$ is finite, so is \underline{v} . Similarly, since the consumption set is unbounded above, $\bar{v} = \frac{1}{1-\beta}\bar{u}$ is the least upper bound of \mathbb{V} . Since (8) and (14) define a convex set, any $v_0 \in (\underline{v}, \bar{v})$ can be attained by incentive compatible allocations, which establishes that $\mathbb{V} = [\underline{v}, \bar{v}]$ if the utility is bounded below and $\mathbb{V} = (\underline{v}, \bar{v})$ otherwise. It is not always possible to characterize the domain of the value function in such a simple way. The general way to characterize the set \mathbb{V} is described in Proposition 8.

Proposition 3. *Suppose that the utility function satisfies Assumption 1, shocks are i.i.d. and $T = \infty$. Then K satisfies the Bellman equation*

$$K(v) = \max_{(\vec{u}, \vec{w}) \in B(v)} \sum_{\theta \in \Theta} \pi(\theta) [-C(u(\theta)) + \beta K(w(\theta))]. \quad (23)$$

Proof. We first show that the maximum in problem (23) is well-defined. That is, for any $v \in \mathbb{V}$, there exist (\vec{u}_v, \vec{w}_v) that maximize the right hand side of (23) within the set $B(v)$ defined in (22). To do so we restrict the optimization over (\vec{u}, \vec{w}) to a compact set. Since the right hand side of (23) is a function of two variables (\vec{u}, \vec{w}) that is continuous, this implies that it reaches its maximum.

The allocation (\vec{u}', \vec{w}') defined by $u'(\theta) = (1 - \beta)v$ and $w'(\theta) = v$ for all $\theta \in \Theta$ satisfies the constraints (19) and (20) and yields value $-C((1 - \beta)v) + \beta K(v) \equiv \underline{K}_v$. Therefore the r.h.s. of the Bellman equation is larger than \underline{K}_v . Now suppose that for some θ , $w(\theta)$ is such that $\beta\pi(\theta)K(w(\theta)) < \underline{K}_v$. Then we have

$$\sum_{\theta \in \Theta} \pi(\theta) [-C(u(\theta)) + \beta K(w(\theta))] < \underline{K}_v,$$

a contradiction. Thus we can restrict the search to $\{w(\theta) \text{ s.t. } \beta\pi(\theta)K(w(\theta)) \geq \underline{K}_v\}$ and, similarly, to $\{u(\theta) \text{ s.t. } -\pi(\theta)C(u(\theta)) \geq \underline{K}_v\}$. Moreover, we have $\lim_{u \rightarrow \bar{u}} -C(u) = -\infty$ and $\lim_{v \rightarrow \bar{v}} K(v) = -\infty$. To show the latter, consider the function $\bar{K}(v)$ which maximizes the objective function (21) subject to delivering lifetime utility $v_0 = v$, without the incentive constraints. Obviously $\bar{K}(v) \geq K(v)$. We easily obtain that the solution to this relaxed problem is $\bar{K}(v) = -\frac{1}{1-\beta} \mathbb{E}[C(C'^{-1}(\gamma_v \theta))]$ where $\gamma_v > 0$ is the multiplier on the promise-keeping constraint. We have $\mathbb{E}[\theta C'^{-1}(\gamma_v \theta)] = (1 - \beta)v$, so $\lim_{v \rightarrow \bar{v}} \gamma_v = \infty$ and hence $\lim_{v \rightarrow \bar{v}} \bar{K}(v) = -\infty$. This implies that $\lim_{v \rightarrow \bar{v}} K(v) = -\infty$, and therefore the previous arguments lead to upper bounds $\bar{u}_\theta, \bar{w}_\theta$ for $u(\theta)$ and $w(\theta)$, respectively. Moreover, $\mathbb{E}[\theta u(\theta) + \beta w(\theta)]$ goes to $-\infty$ if $u(\theta)$ or $w(\theta)$ go to $-\infty$, because of the upper bounds $\bar{u}_\theta, \bar{w}_\theta$. This contradicts the promise keeping constraint, and thus gives us lower bounds $\underline{u}_\theta, \underline{w}_\theta$ for all θ . Therefore, we can restrict the search for $\{u(\theta), w(\theta)\}_{\theta \in \Theta}$ to the compact set $\prod_{\theta \in \Theta} [\underline{u}_\theta, \bar{u}_\theta] \times [\underline{w}_\theta, \bar{w}_\theta]$. This concludes the proof that the maximum in the right hand side of (23) is attained.

Next, we show that K , the solution to (21), satisfies the Bellman equation (23). We start by showing that the left hand side is weakly smaller than the right hand side. Suppose that for some v , we have

$$K(v) > \max_{(\vec{u}, \vec{w}) \in B(v)} \sum_{\theta \in \Theta} \pi(\theta) [-C(u(\theta)) + \beta K(w(\theta))].$$

Thus there exists $\varepsilon > 0$ such that

$$K(v) \geq \mathbb{E}[-C(u(\theta)) + \beta K(w(\theta))] + \varepsilon, \quad \forall (\vec{u}, \vec{w}) \in B(v).$$

Now consider an allocation $\mathbf{u} = \{u_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$ that satisfies incentive compatibility (8) and delivers lifetime utility v . We can write $\mathbf{u} = (\{u_1(\theta_1)\}_{\theta_1 \in \Theta}, \{\mathbf{u}_2(\theta_1)\}_{\theta_1 \in \Theta})$, where for all θ_1 , $\mathbf{u}_2(\theta_1) = \{u_t(\theta_1, \theta_2^t)\}_{t \geq 2, \theta_2^t \in \Theta^{t-1}}$. Let $w_2(\theta_1)$ denote the lifetime utility achieved by $\mathbf{u}_2(\theta_1)$.⁸ The pair $(\vec{u}_1, \vec{w}_2) = (\{u_1(\theta_1)\}_{\theta_1 \in \Theta}, \{w_2(\theta_1)\}_{\theta_1 \in \Theta})$ satisfies (19) and (20), i.e., $(\vec{u}_1, \vec{w}_2) \in B(v)$. Thus, the previous inequality implies that

$$\begin{aligned} K(v) &\geq \mathbb{E}[-C(u_1(\theta_1)) + \beta K(w_2(\theta_1))] + \varepsilon \\ &\geq \mathbb{E}\left[-C(u_1(\theta_1)) + \beta \mathbb{E}_1\left[-\sum_{t=2}^{\infty} \beta^{t-2} C(u_t(\theta_1, \theta_2^t))\right]\right] + \varepsilon \\ &= \mathbb{E}_0\left[-\sum_{t=1}^{\infty} \beta^{t-1} C(u_t(\theta^t))\right] + \varepsilon, \end{aligned}$$

where the second inequality follows from the definition of $K(w_2(\theta_1))$, since the allocation $\mathbf{u}_2(\theta_1)$ satisfies (8) and yields $w_2(\theta_1)$. Since this reasoning holds for any allocation \mathbf{u} that satisfies (8) and delivers v , we get a contradiction.

Next we show the reverse inequality. Note that by definition of the supremum in (21), for all v and $\varepsilon > 0$ there exists an allocation $\tilde{\mathbf{u}}^{v, \varepsilon} = \{\tilde{u}_t^{v, \varepsilon}(\theta^t)\}$ that satisfies (8) and delivers v with cost

$$\mathbb{E}_0\left[-\sum_{t=1}^{\infty} \beta^{t-1} C(\tilde{u}_t^{v, \varepsilon}(\theta^t))\right] > K(v) - \varepsilon.$$

Let

$$(\vec{u}_v, \vec{w}_v) \in \arg \max_{(\vec{u}, \vec{w}) \in B(v)} \mathbb{E}[-C(u(\theta)) + \beta K(w(\theta))].$$

Consider the incentive compatible allocation \mathbf{u} defined by $u_1(\theta^1) = u_v(\theta_1)$ for all $\theta_1 \in \Theta$ and

⁸Note that the continuation utilities, and in particular $w_2(\theta)$ for all θ , are well defined. Indeed, if not, then for some $s \geq 0$, $U_s^+ \equiv \lim_{T \rightarrow \infty} \mathbb{E}_s \left[\sum_{t=1}^T \beta^{t-1} \{\theta_t u_t(\theta^t) \vee 0\} \right] = \infty$. Since the cost function is convex, we have $C(u) \geq -B + A \{\theta u \vee 0\}$ for some $A, B > 0$, and hence $\lim_{T \rightarrow \infty} \mathbb{E}_s \left[\sum_{t=1}^T \beta^{t-1} C(u_t(\theta^t)) \right] \geq -B + A U_s^+ = \infty$. This implies

$$\lim_{T \rightarrow \infty} \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} C(u_t(\theta^t)) \right] = \sum_{\theta^s \in \Theta^s} \pi_0(\theta^s) \lim_{T \rightarrow \infty} \mathbb{E}_s \left[\sum_{t=1}^T \beta^{t-1} C(u_t(\theta^t)) \right] = \infty,$$

which contradicts the feasibility constraint (2).

$u_t(\theta^t) = \tilde{u}_t^{w_v(\theta_1), \varepsilon}(\theta^t)$ for all $t \geq 2$, $\theta^t \in \Theta^t$. We have

$$\begin{aligned}
K(v) &\geq \mathbb{E}_0 \left[-\sum_{t=1}^{\infty} \beta^{t-1} C(u_t(\theta^t)) \right] \\
&= \mathbb{E}_0 \left[-C(u_v(\theta_1)) + \beta \mathbb{E}_1 \left[-\sum_{t=2}^{\infty} \beta^{t-2} C(\tilde{u}_t^{w_v(\theta_1), \varepsilon}(\theta^t)) \right] \right] \\
&\geq \mathbb{E}_0 [-C(u_v(\theta_1)) + \beta K(w_v(\theta_1))] - \beta \varepsilon \\
&= \max_{(\vec{u}, \vec{w}) \in B(v)} \mathbb{E} [-C(u(\theta)) + \beta K(w(\theta))] - \beta \varepsilon.
\end{aligned}$$

Since $\varepsilon > 0$ was arbitrary, we can let $\varepsilon \rightarrow 0$ in this inequality. We have thus shown that the value function (21) of the dual planner's problem satisfies the Bellman equation (23). \square

The function K inherits the same properties as the function K_t in the finite period version of this economy.

Lemma 1. *Suppose that the utility function satisfies Assumption 1, shocks are i.i.d. and $T = \infty$. Then K is continuous on \mathbb{V} , strictly concave, strictly decreasing, differentiable, with $\lim_{v \rightarrow \underline{v}} K(v) = \lim_{v \rightarrow \underline{v}} K'(v) = 0$ and $\lim_{v \rightarrow \bar{v}} K(v) = \lim_{v \rightarrow \bar{v}} K'(v) = -\infty$.*

Proof. The objective function in (21) is concave and the constraint set is convex, therefore, K is weakly concave. To show the strict concavity of K , pick any $v^a, v^b \in \mathbb{V}$ such that $v^a \neq v^b$, and let $(\vec{u}_{v^a}, \vec{w}_{v^a})$ and $(\vec{u}_{v^b}, \vec{w}_{v^b})$ be the corresponding policy functions that maximize the right hand side of (23). The incentive constraint (20) implies that $\vec{u}_{v^a} \neq \vec{u}_{v^b}$. Let $v^\alpha \equiv \alpha v^a + (1 - \alpha) v^b$, and $(\vec{u}_{v^\alpha}, \vec{w}_{v^\alpha})$ be the corresponding policy function. Since (19) and (20) are linear in $u(\theta)$ and $w(\theta)$, we obtain that

$$(\alpha \vec{u}_{v^a} + (1 - \alpha) \vec{u}_{v^b}, \alpha \vec{w}_{v^a} + (1 - \alpha) \vec{w}_{v^b}) \in B(v^\alpha).$$

Thus K satisfies

$$\begin{aligned}
K(v^\alpha) &= \sum_{\theta \in \Theta} \pi(\theta) [-C(u_{v^\alpha}(\theta)) + \beta K(w_{v^\alpha}(\theta))] \\
&\geq \sum_{\theta \in \Theta} \pi(\theta) [-C(\alpha u_{v^a}(\theta) + (1 - \alpha) u_{v^b}(\theta)) + \beta K(\alpha w_{v^a}(\theta) + (1 - \alpha) w_{v^b}(\theta))],
\end{aligned}$$

so that by the strict concavity of $-C$ and the weak concavity of K we get

$$\begin{aligned}
K(v^\alpha) &> \alpha \sum_{\theta \in \Theta} \pi(\theta) [-C(u_{v^a}(\theta)) + \beta K(w_{v^a}(\theta))] \\
&\quad + (1 - \alpha) \sum_{\theta \in \Theta} \pi(\theta) [-C(u_{v^b}(\theta)) + \beta K(w_{v^b}(\theta))] = \alpha K(v^a) + (1 - \alpha) K(v^b),
\end{aligned}$$

Therefore K is strictly concave.

The weak concavity implies that K is continuous on the interior of \mathbb{V} (Exercise 4.23 in Rudin (1976)). To show the continuity of K on \mathbb{V} it remains to show that $\lim_{v \rightarrow \underline{v}} K(v) = K(\underline{v})$ when the utility is bounded below. Since the only feasible solution that delivers \underline{v} has $u_t(\theta^t) = \underline{u}$ for all t, θ^t , we have in this case $K(\underline{v}) = -\frac{1}{1-\beta} C(\underline{u}) = 0$. Therefore showing the continuity at \underline{v} is equivalent to showing that $\lim_{v \rightarrow \underline{v}} K(v) = 0$. Let $\underline{K}(v) = -\frac{1}{1-\beta} C((1-\beta)v)$ be the cost of delivering $u_t(\theta^t) = (1-\beta)v$ independently of θ^t . Since this allocation is incentive compatible, we have $0 \geq K(v) \geq \underline{K}(v)$ for all v . \underline{K} is continuous on \mathbb{V} with $\lim_{v \rightarrow \underline{v}} \underline{K}(v) = 0$, therefore $\lim_{v \rightarrow \underline{v}} K(v) = 0$.

We already showed that $\lim_{v \rightarrow \bar{v}} K(v) = -\infty$ in the proof of Proposition 3.

To show the strict monotonicity, for any $v_0^a < v_0^b$ pick $v \in (v_0^a, v_0^b)$ and $\alpha_v \in (0, 1)$ such that $v_0^a = \alpha_v v + (1 - \alpha_v) v_0^b$. Since K is strictly concave, we have $K(v_0^a) > \alpha_v K(v) + (1 - \alpha_v) K(v_0^b)$. Letting $v \rightarrow \underline{v}$ in this inequality, we obtain $K(v_0^a) \geq (1 - \alpha_{\underline{v}}) K(v_0^b) \geq K(v_0^b)$, and hence K is weakly decreasing. But then using $K(v) \geq K(v_0^b)$ in the previous inequality leads to $K(v_0^a) > \alpha_v K(v_0^b) + (1 - \alpha_v) K(v_0^b) = K(v_0^b)$, so that K is strictly decreasing.

Next we show the differentiability of the cost function K in the case where the utility is unbounded. A slightly different perturbational argument can be used to establish the differentiability when the utility function is bounded, taking care of the situations when the optimum is at the corners (see, e.g., Farhi and Werning (2007)). Fix an interior v and define, for all $x \in (-\varepsilon, \varepsilon)$ for some small $\varepsilon > 0$,

$$L_v(x) = \sum_{\theta \in \Theta} \pi(\theta) [-C(u_v(\theta) + x) + \beta K(w_v(\theta))].$$

The allocation (\vec{u}_x, \vec{w}_x) with $u_x(\theta) = u_v(\theta) + x$ and $w_x(\theta) = w_v(\theta)$ for all θ is incentive compatible and delivers lifetime utility $v + x$. Therefore, for all x we have $L_v(x) \leq K(v + x)$, with equality if $x = 0$. Since $L_v(\cdot)$ is concave and differentiable on $(-\varepsilon, \varepsilon)$ (because $-C(\cdot)$ is), the Benveniste-Scheinkman theorem (Benveniste and Scheinkman (1979), or Theorem 4.10 in Stokey et al. (1989)) implies that K is differentiable at v and we have $K'(v) = L'_v(0)$. Direct calculation of $L'_v(0)$ shows that

$$K'(v) = \sum_{\theta \in \Theta} \pi(\theta) [-C'(u_v(\theta))] \leq 0. \quad (24)$$

The bounds $\underline{K}(v) \leq K(v) \leq \bar{K}(v)$ (see the proof of Proposition 3) and the limits $\lim_{v \rightarrow \underline{v}} \underline{K}'(v) = 0$ and $\lim_{v \rightarrow \bar{v}} \bar{K}'(v) = -\infty$ imply that $\lim_{v \rightarrow \underline{v}} K'(v) = 0$ and $\lim_{v \rightarrow \bar{v}} K'(v) = -\infty$. \square

Finally, we are ultimately interested in recovering a solution to the problem (9). Analogous to the finite period case, we call the solution to (23) a *policy function* and denote it by (\vec{u}_v, \vec{w}_v) . For any initial v_0 these functions *generate* (\mathbf{u}, \mathbf{v}) as in Section 2.3.1.

Proposition 4. *Suppose that the utility function satisfies Assumption 1, shocks are i.i.d. and $T = \infty$. Let v_0 be defined by $K(v_0) = -\frac{e}{1-\beta}$. If the sequence (\mathbf{u}, \mathbf{v}) generated by the policy*

functions to the Bellman equation (23) given v_0 satisfies

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 [\beta^t v_t (\theta^t)] = 0 \quad (25)$$

and

$$\limsup_{t \rightarrow \infty} \mathbb{E}_0 [\beta^t v_t (\sigma^t (\theta^t))] \geq 0, \quad \forall \sigma \quad (26)$$

then (\mathbf{u}, \mathbf{v}) achieves the supremum of the primal maximization problem (9).

Proof. Let (\mathbf{u}, \mathbf{v}) denote the allocations generated by the policy functions (\vec{u}, \vec{w}) starting at v_0 . First, we show that (\mathbf{u}, \mathbf{v}) achieves the supremum of the *dual* problem (21), i.e., that (\mathbf{u}, \mathbf{v}) satisfies the constraints (14) and (8) and attains $K(v_0)$. To see that the constraint (14) is satisfied, note that by repeated substitution, (\mathbf{u}, \mathbf{v}) satisfies

$$v_0 = \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t u_t (\theta^t) \right] + \beta^T \mathbb{E}_0 [v_T (\theta^T)].$$

If (\mathbf{u}, \mathbf{v}) satisfies (25), then taking limits as $T \rightarrow \infty$ (see footnote 8 for the existence of the limit on the right hand side) leads to

$$v_0 = \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t (\theta^t) \right].$$

To see that (\mathbf{u}, \mathbf{v}) satisfies the incentive compatibility constraint (8), consider any reporting strategy σ . Since the policy functions (\vec{u}, \vec{w}) that generate (\mathbf{u}, \mathbf{v}) satisfy (20), repeated substitution implies that (\mathbf{u}, \mathbf{v}) satisfies

$$v_0 \geq \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t u (\sigma^t (\theta^t)) \right] + \beta^T \mathbb{E}_0 [v_T (\sigma^T (\theta^T))].$$

If the condition (26) is satisfied, taking limits implies that

$$\limsup_{T \rightarrow \infty} \left\{ v_0 - \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t U (c_t (\sigma^t (\theta^t))) \right] \right\} \geq 0 \quad \forall \sigma,$$

establishing that (\mathbf{u}, \mathbf{v}) satisfies (8).

We next show that (\mathbf{u}, \mathbf{v}) attains $K(v_0)$. Repeatedly applying the Bellman equation (23) yields

$$K(v_0) = \mathbb{E}_0 \left[- \sum_{t=1}^T \beta^{t-1} C (u_t (\theta^t)) \right] + \beta^T \mathbb{E}_0 [K(v_T (\theta^T))].$$

Since $\limsup_{T \rightarrow \infty} \beta^T \mathbb{E}_0 [K(v_T(\theta^T))] \leq 0$ we obtain

$$K(v_0) \leq \mathbb{E}_0 \left[- \sum_{t=1}^{\infty} \beta^{t-1} C(u_t(\theta^t)) \right].$$

But (\mathbf{u}, \mathbf{v}) satisfies the constraints of problem (21), thus $K(v_0) \geq \mathbb{E}_0 [-\sum_{t=1}^{\infty} \beta^{t-1} C(u_t(\theta^t))]$. Therefore (\mathbf{u}, \mathbf{v}) achieves the supremum of the dual problem (21).

Second, we show that the maximum to the dual problem (21) is also a maximum to the primal problem (9). Since (\mathbf{u}, \mathbf{v}) delivers v_0 which satisfies $-K(v_0) = \frac{e}{1-\beta}$, \mathbf{u} satisfies the feasibility constraint (2) and therefore $V(e) \geq v_0$. Suppose that this inequality is strict, so that there exists $(\mathbf{u}', \mathbf{v}')$ that delivers lifetime utility $v'_0 > v_0$, is incentive compatible and satisfies $\mathbb{E}_0 [\sum_{t=1}^{\infty} \beta^{t-1} C(u'_t)] \leq \frac{e}{1-\beta}$. The continuity and strict monotonicity of K (Lemma 1) imply that $-K(v'_0) > -K(v_0) = \frac{e}{1-\beta}$. Since $\mathbb{E}_0 [\sum_{t=1}^{\infty} \beta^{t-1} C(u'_t)] \geq -K(v'_0)$, this establishes a contradiction. \square

If the utility function is bounded, then the limiting conditions (25) and (26) are automatically satisfied and Proposition 4 implies simultaneously that the supremum to problem (21) is attained and that it can be recovered from the policy functions of the Bellman equation (23). When the utility function is unbounded, an extra step is needed to verify that the policy functions generate a solution that satisfies the ‘‘continuity at infinity’’ conditions (25) and (26). These conditions are standard in the statement of the one-shot deviation principle to ensure that the agent cannot gain by some infinite sequence of deviations (see, e.g., Theorem 4.2 in Fudenberg and Tirole (1991)). We show in an example in Section 2.4 how to verify ex-post these conditions with unbounded utilities.

The analysis above can be simplified if we modify Assumption 1 and assume that the domain of U is compact. In this case $C(\cdot)$ is a bounded function on a compact set $[\underline{u}, \bar{u}]$. The results of Propositions 3 and 4 can be proven immediately using standard contraction mapping arguments (see Chapter 9 in Stokey et al. (1989)). Moreover, the results of Stokey et al. (1989) show that the function K that satisfies the functional equation (23) is the unique fixed point of the Bellman operator defined on the space of continuous and bounded functions by

$$\mathcal{B}(k)(v) = \max_{(\vec{u}, \vec{w}) \in B(v)} \sum_{\theta \in \Theta} \pi(\theta) [-C(u(\theta)) + \beta k(w(\theta))],$$

and that for all bounded and continuous k_0 the sequence $\{k_n\}_{n \geq 0}$ defined by $k_n = \mathcal{B}^n k_0$ for all n converges to K . This characterization is useful to compute the solution to the problem numerically.

2.4 Characterization of the solution with i.i.d. shocks

In this section we characterize the solution to the Bellman equation (23). At the end of this section we provide a simple example showing how to verify the limiting conditions (25) and (26) when the utility function is unbounded.

For simplicity, we assume that θ can take only two values, $\Theta = \{\theta_{(1)}, \theta_{(2)}\}$, with $\theta_{(1)} < \theta_{(2)}$. The incentive constraints (13) with two shocks reduce to

$$\theta_{(1)}u(\theta_{(1)}) + \beta w(\theta_{(1)}) \geq \theta_{(1)}u(\theta_{(2)}) + \beta w(\theta_{(2)}), \quad (27)$$

and

$$\theta_{(2)}u(\theta_{(2)}) + \beta w(\theta_{(2)}) \geq \theta_{(2)}u(\theta_{(1)}) + \beta w(\theta_{(1)}). \quad (28)$$

Proposition 5. *Suppose that the utility function satisfies Assumption 1, shocks are i.i.d., $|\Theta| = 2$ and $T = \infty$. The constraint (27) binds, and the constraint (28) is slack for all interior v . Moreover $u_v(\theta_{(1)}) \leq u_v(\theta_{(2)})$ and $w_v(\theta_{(1)}) \geq v \geq w_v(\theta_{(2)})$, with strict inequalities for all interior v . The policy functions $u_v(\theta)$, $w_v(\theta)$ are continuous in v for all $\theta \in \Theta$. If $w_v(\theta_{(2)})$ is interior, the policy functions satisfy*

$$K'(v) = \mathbb{E}[K'(w_v)] = \mathbb{E}[-C'(u_v)], \quad \forall v. \quad (29)$$

Proof. The proof proceeds by guessing that the constraint (28) is slack and solving a relaxed problem (23) in which this constraint is dropped. We then verify ex post that (28) is satisfied. The strict concavity of the objective function in (23) and the convexity of the constraint set then implies that the solution to the relaxed problem is the unique solution to the original problem.

Let $\xi_v \geq 0$ and $\gamma_v \geq 0$ be the Lagrange multipliers on the incentive compatibility constraint (27) and the promise-keeping constraint (19) in the relaxed problem. The first order conditions with respect to $u(\theta_{(1)})$ and $u(\theta_{(2)})$ are

$$\pi(\theta_{(1)})C'(u_v(\theta_{(1)})) - \xi_v\theta_{(1)} \geq \gamma_v\pi(\theta_{(1)})\theta_{(1)}, \quad (30)$$

$$\pi(\theta_{(2)})C'(u_v(\theta_{(2)})) + \xi_v\theta_{(1)} \geq \gamma_v\pi(\theta_{(2)})\theta_{(2)}, \quad (31)$$

where these constraints hold with equality if $u_v(\theta_{(1)}) > \underline{u}$ and $u_v(\theta_{(2)}) > \underline{u}$, respectively. Similarly, the first order conditions with respect to $w(\theta_{(1)})$ and $w(\theta_{(2)})$ are

$$-\pi(\theta_{(1)})K'(w_v(\theta_{(1)})) - \xi_v \geq \gamma_v\pi(\theta_{(1)}), \quad (32)$$

$$-\pi(\theta_{(2)})K'(w_v(\theta_{(2)})) + \xi_v \geq \gamma_v\pi(\theta_{(2)}), \quad (33)$$

where these constraints hold with equality if $w_v(\theta_{(1)}) > \underline{v}$ and $w_v(\theta_{(2)}) > \underline{v}$, respectively.

We first show that $u_v(\theta_{(1)})$, $w_v(\theta_{(1)})$ are interior for all interior v . (We show below that $u_v(\theta_{(2)})$ is also interior.) Suppose that $u_v(\theta_{(1)}) = \underline{u}$. Since $C'(\underline{u}) = 0$, (30) implies that

$\xi_v = \gamma_v = 0$. If $u_v(\theta_{(2)}) > \underline{u}$, then (31) would hold with equality, implying $C'(u_v(\theta_{(2)})) = 0$, a contradiction. Thus we have $u_v(\theta_{(1)}) = u_v(\theta_{(2)}) = \underline{u}$, and the same reasoning implies that $w_v(\theta_{(1)}) = w_v(\theta_{(2)}) = \underline{v}$, which contradicts the promise keeping constraint (19) when v is interior. Therefore we must have $u_v(\theta_{(1)}) > \underline{u}$ so that (30) holds with equality. An identical reasoning implies that $w_v(\theta_{(1)}) > \underline{v}$, so that (32) holds with equality.

We now show that $\xi_v > 0$ for all interior v . If $\xi_v = 0$, then (32) and (33) imply that $w_v(\theta_{(2)}) \geq w_v(\theta_{(1)})$ by the concavity of K . Moreover (30) and (31) with $\theta_{(2)} > \theta_{(1)}$ imply that $u_v(\theta_{(2)}) > u_v(\theta_{(1)})$. This violates the incentive constraint (27), and hence $\xi_v > 0$ if $v > \underline{v}$. This implies that the constraint (27) holds with equality for all $v > \underline{v}$, and it also trivially holds as an equality for $v = \underline{v}$.

We show next that the solution to the relaxed problem satisfies (28). Suppose not, i.e.,

$$\theta_{(2)}u_v(\theta_{(2)}) + \beta w_v(\theta_{(2)}) < \theta_{(2)}u_v(\theta_{(1)}) + \beta w_v(\theta_{(1)}).$$

Sum this equation with (27) which holds with equality, to obtain $u_v(\theta_{(2)}) < u_v(\theta_{(1)})$, and thus $w_v(\theta_{(2)}) > w_v(\theta_{(1)}) > \underline{v}$. This implies that (33) holds with equality. But (32) and (33) with $\xi_v \geq 0$ then imply that $w_v(\theta_{(2)}) \leq w_v(\theta_{(1)})$, a contradiction. Therefore the incentive constraint (28) is satisfied in the relaxed problem for all v . Moreover, if v is interior, the same reasoning with $\xi_v > 0$ implies that (28) is slack.

Summing the incentive constraints (27) and (28) implies $u_v(\theta_{(2)}) \geq u_v(\theta_{(1)})$ and hence $w_v(\theta_{(1)}) \geq w_v(\theta_{(2)})$. In particular, $u_v(\theta_{(2)})$ is interior for all interior v , and (31) holds with equality. Now suppose $v > \underline{v}$. If $w_v(\theta_{(2)}) = \underline{v}$, then $w_v(\theta_{(1)}) > w_v(\theta_{(2)})$. If $w_v(\theta_{(2)}) > \underline{v}$, then (33) holds with equality, and (32) with $\xi_v > 0$ yields $w_v(\theta_{(1)}) > w_v(\theta_{(2)})$ by the strict concavity of K . We then obtain $u_v(\theta_{(1)}) < u_v(\theta_{(2)})$ from (27). When v is interior, we saw that $u_v(\theta)$ is interior for all θ and therefore Benveniste-Scheinkman arguments (see the arguments leading to equation (24)) establish that $K'(v) = \mathbb{E}[-C'(u_v)] = -\gamma_v$. This equation also holds at the boundary $v = \underline{v}$ since in this case both sides of this expression are equal to zero. Therefore, equations (32) and (33), assuming that $w_v(\theta_{(2)})$ is interior, imply that $-K'(w_v(\theta_{(1)})) > -K'(v)$ and $-K'(w_v(\theta_{(2)})) < -K'(v)$ respectively, so that $w_v(\theta_{(2)}) < v < w_v(\theta_{(1)})$.

Next we show that the policy functions are continuous in v . The objective function in (23) is continuous and strictly concave on $\mathbb{U}^2 \times \mathbb{V}^2$. Following the same steps as in the proof of Proposition 3, we restrict the optimization over (\vec{u}, \vec{w}) to a compact set $\mathbb{X} \subset \mathbb{U}^2 \times \mathbb{V}^2$. The constraint set $B(\cdot) : \mathbb{V} \rightarrow \mathbb{X}$ defined in (22) is then a continuous, compact-valued and convex-valued correspondence. Thus, by the theorem of the Maximum (see, e.g., Theorem 3.6 and Exercise 3.11a in Stokey et al. (1989)), the function (\vec{u}_v, \vec{w}_v) is continuous in v .

We now prove equation (29). We saw above that for all $v \geq \underline{v}$, $K'(v) = \mathbb{E}[-C'(u_v)] = -\gamma_v$. Moreover, when v is interior, summing the conditions (30) and (31) (which both hold with equality), and the conditions (32) and (33) (the former holds with equality, the latter does as

well if $w_v(\theta_{(2)}) > \underline{v}$ yields $\mathbb{E}[K'(w_v)] \leq \mathbb{E}[-C'(u_v)]$, with equality if $w_v(\theta_{(2)})$ is interior. Finally, this equation holds with equality if $v = \underline{v}$.

Note finally that from (30) and (31), we obtain that for all interior v ,

$$-\theta_{(1)}K'(v) < C'(u_v(\theta_{(1)})) < -K'(v) < C'(u_v(\theta_{(2)})) < -\theta_{(2)}K'(v). \quad (34)$$

□

Proposition 5 highlights the main principle underlying the optimal provision of incentives in dynamic economies. Consider the unconstrained first-best allocation, given by

$$\frac{1}{\theta_{(1)}}C'(u_v^{\text{fb}}(\theta_{(1)})) = \frac{1}{\theta_{(2)}}C'(u_v^{\text{fb}}(\theta_{(2)})). \quad (35)$$

In this case the future continuation allocations are independent of the current realization of the shock: the social planner redistributes resources from agents with shock $\theta_{(1)}$ to agents with shock $\theta_{(2)}$. As we discussed above, this allocation is not incentive compatible when shocks are private information. To provide incentives, the social planner spreads out future promised utilities w_v , rewarding agents who report a lower shock and punishing those who report a higher shock. In exchange, a higher reported shock gives the agent a higher utility today. The bounds (34) imply

$$\frac{1}{\theta_{(2)}}C'(u_v(\theta_{(2)})) < \frac{1}{\theta_{(1)}}C'(u_v(\theta_{(1)})).$$

Therefore the spread in the current period utilities (or consumption allocations) is not as large as in the first best allocation. This reflects the fact that private information makes redistribution more costly. The resources are still being redistributed away from the $\theta_{(2)}$ -type, which is implied by the fact that his incentive constraint binds at the optimum.

Equation (29) shows how the planner allocates the costs of providing incentives over time. Fluctuations in promised utilities are costly due to the concavity of the cost function, and it is optimal to smooth these costs over time. The best smoothing can be achieved if the forecast of future marginal costs of providing incentives based on the current information, $\mathbb{E}_t[K'(v_{t+s})]$, is equal to the current period marginal cost, $K'(v_t)$, so that $K'(v_t)$ is a random walk. This result is a manifestation of the same general principle that underlies consumption smoothing in the permanent income hypothesis (see Friedman (1957), Hall (1978)) or tax smoothing in public finance (see Barro (1979)).

Analogous to other environments with cost smoothing, the random walk nature of $K'(v_t)$ has powerful implications about the long-run properties of the solution.⁹ To derive these implications, we first introduce the notion of *martingale* and the Martingale Convergence Theorem (see Billingsley (1995), Section 35):

⁹See Chamberlain and Wilson (2000) for an analogous result in consumption smoothing models and Aiyagari et al. (2002) for tax smoothing.

Definition 1. Let $X_t(\theta^t)$ be a random variable on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The sequence $\{X_t, \mathcal{F}_t\}_{t=1,2,\dots}$ is a *martingale* if:

- (i) $\{\mathcal{F}_t\}_{t \geq 1}$ is an increasing sequence of σ -algebras,
- (ii) X_t is measurable with respect to \mathcal{F}_t ,
- (iii) $\mathbb{E}_0[|X_t|] < \infty$, and
- (iv) $\mathbb{E}_t[X_{t+1}] = X_t$ with probability 1.

A *submartingale* is defined as above except that the condition (iv) is replaced by $\mathbb{E}_t[X_{t+1}] \geq X_t$. Any martingale is a submartingale. We have the following important result (Theorem 35.5 in Billingsley (1995)):

Theorem 2. (Martingale Convergence Theorem.) Let $\{X_t\}_{t=0}^\infty$ be a submartingale. If $M \equiv \sup_t \mathbb{E}_0[|X_t|] < \infty$, then $X_t \rightarrow X$ with probability 1, where X is a random variable on $(\Omega, \mathcal{F}, \mathbb{P})$ satisfying $\mathbb{E}_0[|X|] \leq M$.

To apply this result in our context, observe that the policy functions induce a law of motion for the distribution of promised utilities over time. For any probability distribution Ψ on \mathbb{V} , define an operator \mathcal{T} as

$$(\mathcal{T}\Psi)(A) \equiv \int_{\mathbb{V}} \left[\sum_{\theta \in \Theta} \pi(\theta) \mathbb{I}_{\{w_v(\theta) \in A\}} \right] \Psi(dv), \quad (36)$$

for all Borel sets $A \subset \mathbb{V}$. $\mathcal{T}\Psi$ defines another probability distribution on \mathbb{V} . This operator allows us to study the dynamics of the distribution of lifetime utilities in our economy. In particular, let Ψ_0 be a probability distribution on \mathbb{V} that assigns probability 1 to v_0 and define Ψ_t recursively as $\Psi_t = \mathcal{T}\Psi_{t-1}$. In other words, initially in period 0 everyone is identical with the same lifetime utility v_0 . Over time idiosyncratic shocks lead to inequality in lifetime promises captured by the distribution Ψ_t in period t . A distribution is *invariant* if it satisfies $\Psi = \mathcal{T}\Psi$.

Suppose that the utility function is unbounded below, so that $w_v(\theta_{(2)})$ is always interior. Consider the random variable $K'(v_t(\theta^t))$ defined recursively on the probability space $(\Theta^\infty, \mathcal{B}(\Theta^\infty), \pi_\infty)$ starting at $K'(v_0)$. The sequence $\{K'(v_t(\theta^t)), \mathcal{B}(\Theta^t)\}_{t=1,2,\dots}$ is a martingale. Indeed, $\{\mathcal{B}(\Theta^t)\}_{t \geq 1}$ is an increasing sequence of σ -algebras, $K'(v_t)$ is measurable with respect to $\mathcal{B}(\Theta^t)$, $\mathbb{E}_0[|K'(v_t)|] = -\mathbb{E}_0[K'(v_t)] = -K'(v_0) < \infty$, and $\mathbb{E}_t[K'(v_{t+1})] = K'(v_t)$ follows from Lemma 5. Hence all the conditions of Definition 1 are satisfied, and Theorem 2 implies that $K'(v_t)$ converges almost surely to a random variable X . That is, for almost all histories $\theta^\infty \in \Theta^\infty$, we have $K'(v_t(\theta^\infty)) \rightarrow X(\theta^\infty)$. The following proposition, whose proof follows Thomas and Worrall (1990), further characterizes the limit of the sequence:

Proposition 6. Suppose that the utility function satisfies Assumption 1, shocks are i.i.d., $|\Theta| = 2$ and $T = \infty$.¹⁰ If the utility function is unbounded below, then $v_t(\theta^t) \rightarrow -\infty$ as $t \rightarrow \infty$ with

¹⁰The condition $|\Theta| = 2$ is not important for this proposition. It is easy to show that the martingale property (29) holds for any number of shocks.

probability 1. If the utility function is bounded below by \underline{v} , then the unique invariant distribution of continuation utilities on \mathbb{V} assigns mass 1 to the lower bound \underline{v} .

Proof. Suppose that the utility function is unbounded below. Then $K'_t(v_t(\theta^t))$ is a martingale, and Theorem 2 implies that for almost all θ^∞ , $K'_t(v_t(\theta^t))$ converges to some random variable $X(\theta^\infty)$. We now show that its limit $X(\theta^\infty)$ is equal to 0 almost surely.

Consider a path θ^t such that $K'_t(v_t(\theta^t)) \rightarrow \kappa < 0$. The sequence $v_t(\theta^t)$ thus converges to $\hat{v} > \underline{v}$, solution to $K'(\hat{v}) = \kappa$. With probability one, the state $\theta_{(2)}$ occurs infinitely often on this path. Take the subsequence composed of the dates $\{t_n\}_{n=1,2,\dots}$ where the state $\theta_{(2)}$ occurs. We have $\lim_{n \rightarrow \infty} v_{t_n-1}(\theta^{t_n-1}) = \hat{v}$ and $\lim_{n \rightarrow \infty} v_{t_n}(\theta^{t_n}) = \hat{v}$. Since the policy function $w_v(\theta)$ is continuous in v for all $\theta \in \Theta$, we obtain

$$\lim_{n \rightarrow \infty} w_{v_{t_n-1}(\theta^{t_n-1})}(\theta_{(2)}) = w_{\hat{v}}(\theta_{(2)}).$$

But $w_{v_{t_n-1}(\theta^{t_n-1})}(\theta_{(2)}) = v_{t_n}(\theta^{t_n-1}, \theta_{(2)}) = v_{t_n}(\theta^{t_n})$, hence we also have

$$\lim_{n \rightarrow \infty} w_{v_{t_n-1}(\theta^{t_n-1})}(\theta_{(2)}) = \hat{v}.$$

This implies that $w_{\hat{v}}(\theta_{(2)}) = \hat{v}$, which contradicts the inequality $w_{\hat{v}}(\theta_{(2)}) < \hat{v}$ proved in Proposition 29.

Now suppose that the utility function is bounded below by $\underline{v} > -\infty$. In this case $-K'_t(v_t(\theta^t))$ is a (possibly unbounded) submartingale. Note that the point \underline{v} is absorbing, i.e. $u_{\underline{v}}(\theta) = \underline{v}$ and $w_{\underline{v}}(\theta) = \underline{v}$ for all $\theta \in \Theta$. Consider an invariant distribution Ψ of continuation utilities on \mathbb{V} , and let $\text{Supp}(\Psi) \subset \mathbb{V}$ denote its support. Let \mathcal{M}_v denote the Markov chain characterizing the law of motion of continuation utilities, starting at v . Define the set $S_1 \subset \text{Supp}(\Psi)$ consisting of all the continuation utility values v for which \mathcal{M}_v reaches \underline{v} in a finite number of steps with positive probability, and the set $S_2 = \text{Supp}(\Psi) \setminus S_1$. By construction of S_1 , every state $v \in S_1 \setminus \{\underline{v}\}$ is transient, so that such a v cannot be in the support of the invariant distribution. Now, for $v \in S_2$, the Markov chain \mathcal{M}_v defines a sequence $\{v_t(\theta^t)\}_{t, \theta^t}$. By construction of S_2 , the process $K'_t(v_t(\theta^t))$ is a martingale and the previous arguments show that v cannot be in the support of the invariant distribution. Therefore $\text{Supp}(\Psi) = \{\underline{v}\}$. \square

The result of Proposition 6 is often referred to as the *immiseration* result. It shows that a feature of the optimal contract is that agents' consumption, c_t^* , goes to 0 with probability 1 as $t \rightarrow \infty$. When the utility function is unbounded below, this implies that agents' utility diverges to $-\infty$; otherwise the only invariant distribution is degenerate and assigns probability 1 to $\frac{U(0)}{1-\beta}$. The fact that $c_t^* \rightarrow 0$ w.p. 1 does *not* mean that everyone's consumption converges to zero. As we saw in Proposition 5, an agent with shock realization $\theta_{(1)}$ always gets a strictly higher promised utility (and hence consumption) in the future. Thus there are always some agents (whose measure goes to zero as $t \rightarrow \infty$) with strictly positive consumption. In Section 3.2 we shut down the intertemporal transfer of resources and show that the immiseration result still

holds in those settings. This will lead to an implication of the optimal provision of incentives: in order to provide incentives for agents to reveal their private information, the planner needs to increase inequality without bounds over time. As time goes to infinity this inequality grows until a measure 0 of agents consume the entire endowment of the economy.

The intuition for this result is as follows. To provide incentives to the agents to reveal information in the current period, the principal needs to commit to increasing inequality (in promised utilities and therefore consumption) in the future. When the interest rates are equal to the discount factor, as we assumed in this section, there are no offsetting forces and inequality under the optimal contract grows over time. In the infinite period economy it approaches an extreme level as $t \rightarrow \infty$ in which only a measure zero of agents have positive consumption. We revisit this result in subsequent sections, especially in Sections 2.4.1 and 3.2.

2.4.1 Existence of a non-degenerate invariant distribution

In this section we show a simple example where there exists a non-degenerate invariant distribution of utilities if additional constraints are imposed. We study the simplest case in which the planner is required to promise future utilities in a compact set $[\underline{w}, \bar{w}]$, where $\underline{v} < \underline{w} < \bar{w} < \bar{v}$. In Section 3.2 we show how similar constraints can emerge from more sophisticated political economy arguments, but for now we simply impose¹¹

$$w(\theta) \in [\underline{w}, \bar{w}], \quad \forall \theta \in \Theta \tag{37}$$

in problem (23). It is easy to see that all the properties of the modified Bellman equation continue to satisfy the results of Lemma 1 except for the fact that $K'(v)$ is now strictly negative and finite on the set $[\underline{w}, \bar{w}]$.

Lemma 2. *Suppose that all the assumptions of Proposition 5 are satisfied and in addition the constraint (37) is imposed. Then there are no absorbing points: $w_v(\theta_{(1)}) > w_v(\theta_{(2)})$ for all $v \in [\underline{w}, \bar{w}]$.*

Proof. Suppose $w_v(\theta_{(1)}) = w_v(\theta_{(2)}) = v \in [\underline{w}, \bar{w}]$ for some v , implying (from (27) and (19)) that $u_v(\theta_{(1)}) = u_v(\theta_{(2)}) = (1 - \beta)v$. Thus $K(v) = \underline{K}(v) \equiv -\frac{1}{1-\beta}C((1 - \beta)v)$, where $\underline{K}(v)$ was defined in the proof of Proposition 3. Subtracting equations (30) and (32) (written as an inequality because of the condition $w_v(\theta_{(1)}) \geq \underline{w}$), we obtain

$$K'(v) \leq -\frac{1}{\theta_{(1)}}C'((1 - \beta)v) = \frac{1}{\theta_{(1)}}\underline{K}'(v).$$

But $\theta_{(1)} < 1$ contradicts the fact that $K(v') \geq \underline{K}(v')$ for all $v' > v$. For $v = \bar{w}$, a similar reasoning with equations (31) and (33) (written as an inequality because of the condition

¹¹Our arguments here adapt Atkeson and Lucas (1995), Phelan (1995). See also Farhi and Werning (2007), Hosseini et al. (2013).

$w_v(\theta_{(2)}) \leq \bar{w}$) implies that $K'(w_v(\theta_{(2)})) \geq \frac{1}{\theta_{(2)}} \underline{K}'(w_v(\theta_{(2)})) + \frac{\xi_v}{\pi(\theta_{(2)})} \left(1 - \frac{\theta_{(1)}}{\theta_{(2)}}\right)$ \square

We then show that a non-degenerate long-run distribution of utilities and consumption exists.

Proposition 7. *Suppose that all the assumptions of Proposition 5 are satisfied and in addition the constraint (37) is imposed. Then there exists a unique invariant and non-degenerate distribution Ψ of utilities, and for any initial measure Ψ_0 on the state space, $\mathcal{T}^n(\Psi_0)$ converges to Ψ^* as $n \rightarrow \infty$ at a geometric rate that is uniform in Ψ_0 .*

Proof. The result follows from Theorem 11.12 in Stokey et al. (1989), which holds if condition M p. 348 in Stokey et al. (1989) is satisfied. To show this condition, it is sufficient to show that there exists $\underline{\varepsilon} > 0$ and an integer $N < \infty$ such that, for all $v \in [\underline{w}, \bar{w}]$, $P^N(v, \underline{w}) \geq \underline{\varepsilon}$, where P denotes the transition matrix of the Markov chain \mathcal{M} that characterizes the law of motion of continuation utilities; that is, the probability of reaching \underline{w} starting in N steps from any v is at least as large than $\underline{\varepsilon}$.

To show this we proceed in two steps. First, we prove that if the continuation utility v in the current period is close enough to \underline{w} , receiving a high taste shock $\theta_{(2)}$ implies that the promised utility in the next period is \underline{w} . That is, there exists $\varepsilon > 0$ such that, for all $v \leq \underline{w} + \varepsilon$, we have $w_v(\theta_{(2)}) = \underline{w}$. Suppose by contradiction that this is not the case, and consider a sequence $v_n > \underline{w}$ with $\lim_{n \rightarrow \infty} v_n = \underline{w}$, such that $w_{v_n}(\theta_{(2)}) > \underline{w}$ for all n . The martingale property (29) then writes

$$K'(v_n) = \pi(\theta_{(1)}) K'(w_{v_n}(\theta_{(1)})) + \pi(\theta_{(2)}) K'(w_{v_n}(\theta_{(2)})).$$

Letting $n \rightarrow \infty$ in this equation imposes $w_{\underline{w}}(\theta_{(1)}) = w_{\underline{w}}(\theta_{(2)}) = \underline{w}$, which contradicts Lemma 2.

Second, we prove that there exists $\delta > 0$ such that, for any $v > \underline{w} + \varepsilon$, receiving a high taste shock $\theta_{(2)}$ implies that the promised utility in the next period, $w_v(\theta_{(2)})$, is smaller than $v - \delta$. To show this, note that since $w_v(\theta_{(2)})$ is continuous in v , it is either bounded away from the 45° line for $w > \underline{w} + \varepsilon$, or $w_v(\theta_{(2)}) = v$ for some $v \in [\underline{w} + \varepsilon, \bar{w}]$. By the martingale property, the latter implies that $w_v(\theta_{(1)}) = v$, contradicting Lemma 2.

These results imply that there exists $N < \infty$ such that, for any $v \in [\underline{w}, \bar{w}]$, the promised utility after a sequence of N high taste shocks $\theta_{(2)}$, starting from v , is \underline{w} . This implies that $\underline{\varepsilon} < \pi(\theta_{(2)})^N$ is a uniform lower bound on the probability of being at \underline{w} in N steps. Thus condition M is satisfied in Stokey et al. (1989), which concludes the proof. \square

The immiseration result does not hold in the case where expected discounted utilities from each period onward are constrained by (37), because the lower bound $\underline{w} > \underline{v}$ acts as a reflective (rather than absorbing) barrier (Lemma 2), creating mean reversion that leads to a non-degenerate invariant distribution.

2.4.2 A simple example

In this section we address one remaining issue of our analysis. Proposition 4 showed that the allocation generated by the policy functions of our Bellman equation is the solution to the original problem (9) as long as it satisfies the limiting conditions (25) and (26). These conditions are trivially satisfied if the utility function is bounded, but many convenient functional forms assume unbounded utility. In this section we show a simple example with an unbounded utility function that allows us to easily verify conditions (25) and (26). This example also leads to a characterization of the solution to the Bellman equation “almost” in closed form. We assume a logarithmic utility $U(c) = \ln c$; similar arguments extend to CRRA and CARA preferences. Note that $(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)$ if and only if $(\mathbf{u} - (1 - \beta)v_0, \mathbf{v} - v_0) \in \Gamma(0)$, where $\Gamma(\cdot)$ is defined in (16). We can thus re-write the dual planner’s problem (21) as

$$\begin{aligned} K(v_0) &= \max_{(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)} \mathbb{E}_0 \left[- \sum_{t=1}^{\infty} \beta^{t-1} \exp(u_t) \right] \\ &= \max_{(\tilde{\mathbf{u}}, \tilde{\mathbf{v}}) \in \Gamma(0)} \mathbb{E}_0 \left[- \sum_{t=1}^{\infty} \beta^{t-1} \exp(\tilde{u}_t + (1 - \beta)v_0) \right] = \exp((1 - \beta)v_0) K(0). \end{aligned}$$

This implies that if $\{u_0(\theta), w_0(\theta)\}_{\theta \in \Theta}$ is the solution to the Bellman equation (23) for $v = 0$, then $\{u_0(\theta) + (1 - \beta)v, w_0(\theta) + v\}_{\theta \in \Theta}$ is the solution to (23) for any v . This property allows us to establish bounds on (25) and (26). If we start with some initial v_0 and generate (\mathbf{u}, \mathbf{v}) using the policy functions $(u_v(\theta^t), w_v(\theta^t))$ of the Bellman equation (23) as described in Section 2.3, we have

$$\begin{aligned} v_t(\theta^t) &= w_{v_{t-1}(\theta^{t-1})}(\theta_t) = v_{t-1}(\theta^{t-1}) + w_0(\theta_t) \\ &= w_{v_{t-2}(\theta^{t-2})}(\theta_{t-1}) + w_0(\theta_t) = v_{t-2}(\theta^{t-2}) + w_0(\theta_{t-1}) + w_0(\theta_t) \\ &= \dots = v_1(\theta^1) + w_0(\theta_2) + \dots + w_0(\theta_t) = v_0 + \sum_{s=1}^t w_0(\theta_s). \end{aligned}$$

Let $\underline{A} \equiv \min_{\Theta} \{w_0(\theta)\}$ and $\bar{A} \equiv \max_{\Theta} \{w_0(\theta)\}$, so that $\underline{A} \leq w_0(\theta) \leq \bar{A}$ for all $\theta \in \Theta$. Then $\beta^t(v_0 + t\underline{A}) \leq \beta^t v_t(\theta^t) \leq \beta^t(v_0 + t\bar{A})$ for all t, θ^t . Since $\lim_{t \rightarrow \infty} \beta^t t = 0$, this implies that $\lim_{t \rightarrow \infty} \beta^t v_t(\theta^t) = 0$ for all $\theta^\infty \in \Theta^\infty$, which implies both (25) and (26).

Since the value function K is homogeneous, it is easy to find it “almost” in closed form. Our arguments established that $K(v) = a \exp((1 - \beta)v)$ for some $a < 0$. The parameter a can then be found as a fixed point of the equation

$$a = \max_{(\bar{u}, \bar{w}) \in B(0)} \sum_{\theta \in \Theta} \pi(\theta) [-\exp(u(\theta)) + \beta a \exp((1 - \beta)w(\theta))].$$

The arguments used in this example can be extended to utility functions in the CRRA or CARA classes by observing that if $(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)$ then $(\frac{1}{|v_0|}\mathbf{u}, \frac{1}{|v_0|}\mathbf{v}) \in \Gamma(\frac{v_0}{|v_0|})$.

2.5 Autocorrelated shocks

We now address the case where the taste shocks θ follow a first-order Markov process. The goal of this section is to derive a recursive formulation for the planner's dual problem. We assume that the probabilities of the first period types $\theta_1 \in \Theta$ are given by $\pi(\theta_1 | \theta_{(1)})$, i.e., as if the type realization in period 0 was the seed value $\theta_{(1)}$. This assumption carries no loss of generality and simplifies the exposition. Fernandes and Phelan (2000) show how to write a recursive formulation of the planner's problem.

We define the analogue of the temporary incentive compatibility constraint (10) in the case where shocks are first-order Markov as follows. For all $\theta^{t-1}, \theta, \hat{\theta}$,

$$\begin{aligned} & \theta U(c_t(\theta^{t-1}, \theta)) + \beta \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta^{t+s} | \theta) \theta_{t+s} U(c_{t+s}(\theta^{t-1}, \theta, \theta_{t+1}^{t+s})) \\ & \geq \theta U(c_t(\theta^{t-1}, \hat{\theta})) + \beta \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta^{t+s} | \theta) \theta_{t+s} U(c_{t+s}(\theta^{t-1}, \hat{\theta}, \theta_{t+1}^{t+s})). \end{aligned} \quad (38)$$

The one-shot-deviation result of Proposition 2 extends to the problem with persistent shocks:

Lemma 3. *Suppose that either T is finite, or U is bounded. Suppose moreover that the shocks θ follow a first-order Markov process. An allocation \mathbf{c} satisfies (8) if and only if it satisfies (38).*

Proof. Suppose that (8) is violated for some strategy σ' but (38) holds. If σ' involves misreporting in finitely many nodes, the arguments of Proposition 2 apply directly. If T is infinite and σ' recommends lying at infinitely many nodes, we have, by the previous result,

$$\begin{aligned} & \sum_{t=1}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t) \theta_t U(c_t(\theta^t)) \geq \sum_{t=1}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t) \theta_t U(c_t(\sigma'^t(\theta^t))) - \dots \\ & \dots - \beta^T \left[\sum_{s=1}^{\infty} \sum_{\theta^{T+s} \in \Theta^{T+s}} \beta^{s-1} \pi_{T+s}(\theta^{T+s}) \theta_{T+s} (U(c_{T+s}(\sigma'^{T+s}(\theta^{T+s}))) - U(c_{T+s}(\theta^{T+s}))) \right]. \end{aligned}$$

Since the utility is bounded, the second line converges to zero as $T \rightarrow \infty$, which establishes that if \mathbf{c} satisfies (38) then it satisfies (8). \square

We follow Section 2.3 and re-define our maximization with respect to $u_t(\theta^t)$ rather than $c_t(\theta^t)$. We now emphasize the main differences that persistent shocks introduce. As in Section 2.3, we start by assuming that T is finite.

Finite-Period Economy. In this section we consider the case $T < \infty$. For any history $\theta^t \in \Theta^t$ and any $\theta' \in \Theta$, define $v_t(\theta^t | \theta')$ as

$$v_t(\theta^t | \theta') \equiv \sum_{s=1}^{T-t} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta_{t+1}^{t+s} | \theta') \theta_{t+s} u_{t+s}(\theta^t, \theta_{t+1}^{t+s}), \quad (39)$$

where $(\theta^t, \theta_{t+1}^{t+s})$ denote the histories θ^{t+s} whose first t elements are θ^t . This allows us to write (38) as

$$\theta u_t(\theta^{t-1}, \theta) + \beta v_t(\theta^{t-1}, \theta | \theta) \geq \theta u_t(\theta^{t-1}, \hat{\theta}) + \beta v_t(\theta^{t-1}, \hat{\theta} | \theta), \quad \forall \theta^{t-1}, \theta, \hat{\theta}. \quad (40)$$

Unlike the case of i.i.d. shocks, considered in Section 2.3.1, the continuation utility of an agent who reports θ^t depends not only on the history of reports but also on the true period- t shock θ'_t of the agent. The economic intuition for this result is that when shocks are autocorrelated, the realization of the shock θ'_t is informative about the realization of future shocks from period $t+1$ onward. Repeated substitution allows us to re-write v_t as

$$v_t(\theta^t | \theta') = \sum_{\theta \in \Theta} \pi(\theta | \theta') [\theta u_{t+1}(\theta^t, \theta) + \beta v_{t+1}(\theta^t, \theta | \theta)], \quad \forall \theta^t, \theta', \quad (41)$$

with the convention that $v_T(\theta^T | \theta') = 0$ if T is finite. The initial utility v_0 is given by

$$v_0 = \sum_{\theta \in \Theta} \pi(\theta | \theta_{(1)}) [\theta u_1(\theta) + \beta v_1(\theta | \theta)]. \quad (42)$$

Let $\mathbf{v} = \{v_t(\theta^t | \theta'_t)\}_{t \geq 1, \theta^t \in \Theta^t, \theta'_t \in \Theta}$. The set $\Gamma(v_0)$ is now defined as the set of allocations (\mathbf{u}, \mathbf{v}) that satisfy (40), (41), (42). The direct extension of the arguments of Section 2.3 imply that the optimal incentive compatible allocation (i.e., the solution to the primal problem (9)) exists and is a solution to the dual maximization problem

$$\tilde{K}(v_0) \equiv \max_{(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)} \left[- \sum_{t=1}^T \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta^t | \theta_{(1)}) C(u_t(\theta^t)) \right]. \quad (43)$$

This problem can be written recursively following the same ideas as we used to obtain the Bellman equations (18) and (23), with two differences: (i) the state space is larger when the shocks are autocorrelated, and (ii) the space of feasible values for the state variables is now more difficult to characterize. We show both of these differences using backward induction arguments.

The need for the larger state space can be seen already from the incentive constraints. Each history of reports θ^t has an associated $|\Theta|$ -dimensional vector of “promised utilities” $v_t(\theta^t | \cdot) = \{v_t(\theta^t | \theta_{(j)})\}_{j=1}^{|\Theta|}$, where for each j , $v_t(\theta^t | \theta_{(j)})$ is the promised utility allocated to the agent who reports history θ^t and whose realized type in period t is actually $\theta_{(j)}$. Moreover, the expectation over the future realizations of shocks in period $t+1$ depends on the period- t shock θ_t .

Therefore the state space has dimensionality $|\Theta| + 1$. We now describe the recursive construction of the value function $K_t(v(\theta_{(1)}), \dots, v(\theta_{(|\Theta|)}), \theta_-)$ and its domain $\mathcal{V}_t \times \Theta$. Let \mathcal{V}_T be the space of all vectors $v(\cdot) \in \mathbb{R}^{|\Theta|}$ with the property that there exists some $u \in \mathbb{U}$ such that $v(\theta_{(i)}) = u \sum_{\theta \in \Theta} \pi(\theta | \theta_{(i)}) \theta$ for all $i \in \Theta$. Let $K_T(v(\cdot), \theta) = -C(u)$ for all such $v(\cdot) \in \mathcal{V}_T$. This definition simply captures the fact that in the last period the principal cannot provide any insurance against the period- T shocks (by incentive compatibility (40)), and K_T is then (minus) the cost of the feasible promises that the principal can make in period $T - 1$. For $t > 0$ define K_t recursively as

$$K_t(v(\cdot), \theta_-) = \max_{\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}} \sum_{\theta \in \Theta} \pi(\theta | \theta_-) [-C(u(\theta)) + \beta K_{t+1}(w(\theta|\cdot), \theta)] \quad (44)$$

subject to the promise-keeping constraints

$$v(\theta_{(j)}) = \sum_{\theta \in \Theta} \pi(\theta | \theta_{(j)}) [\theta u(\theta) + \beta w(\theta | \theta)], \quad \forall j \in \{1, \dots, |\Theta|\}, \quad (45)$$

the incentive compatibility constraints

$$\theta u(\theta) + \beta w(\theta | \theta) \geq \theta u(\hat{\theta}) + \beta w(\hat{\theta} | \theta), \quad \forall \theta, \hat{\theta} \in \Theta, \quad (46)$$

and

$$u(\theta) \in \mathbb{U}, \quad w(\theta|\cdot) \in \mathcal{V}_{t+1}, \quad \forall \theta \in \Theta. \quad (47)$$

The domain of K_t is $\mathcal{V}_t \times \Theta$, where \mathcal{V}_t is defined as the set of all $v(\cdot) \in \mathbb{R}^{|\Theta|}$ with the property that there exist $\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}$ such that the constraints (45) and (47) are satisfied.

So far we defined K_t from purely mathematical considerations by observing that the solution to the maximization problem (43) after any history θ^t could be found independently of any other history $\hat{\theta}^t$, as long as we keep track of the vector $v(\cdot)$ and the realization of the period- t shock θ_t . It is useful to describe the economic intuition behind these equations. Equation (46) is simply the incentive constraint, familiar from Section 2.3. Equation (45) for $\theta_{(j)} = \theta_-$ summarizes the expected utility that an agent with period- $(t-1)$ shock θ_- receives in period t . This equation is the analogue of the promise-keeping constraint (19) in the i.i.d. case. Equations (45) for $\theta_{(j)} \neq \theta_-$ are auxiliary “threat-keeping” constraints, which allow us to keep track of the incentives provided in the previous period. Since allocations are incentive-compatible, no agent misrepresents his type “along the equilibrium path” and hence no agent actually obtains utility $v(\theta_{(j)})$ for $\theta_{(j)} \neq \theta_-$. One can think of those $v(\theta_{(j)})$ as threats that the principal chooses in period $(t-1)$ to ensure that agents do not misrepresent their type. Equation (45) in period t ensures that the principal’s subsequent choices are consistent with that threat. The principal chooses a common allocation for all the agents that report θ_- . This common allocation simultaneously delivers utility $v(\theta_-)$ to the agents with true type θ_- (i.e., when expected values are computed using the probabilities $\pi(\theta | \theta_-)$), and $v(\theta_{(j)})$ to the agents with

true types $\theta_{(j)}$ (i.e., when expected values are computed using the probabilities $\pi(\theta | \theta_{(j)})$) for each $j \in \{1, \dots, |\Theta|\}$.

The relationship between the function $K_0(v(\cdot), \theta_-)$ defined in (44) and the function $\tilde{K}_0(v_0)$ defined in (43) is as follows. Observe that there are no auxiliary threat-keeping constraints in the set $\Gamma(v_0)$. It is mathematically equivalent to saying that those constraints are slack. Thus, given our assumption that shocks in period 1 are drawn from $\pi(\cdot | \theta_{(1)})$, the relationship between $K_0(v(\cdot), \theta_-)$ and $\tilde{K}_0(v_0)$ is simply

$$\tilde{K}(v_0) = \max_{v(\cdot) \in \mathcal{V}_0, v(\theta_{(1)}) = v_0} K_1(v(\cdot), \theta_{(1)}). \quad (48)$$

This gives a simple way to find the solution $(\mathbf{u}^*, \mathbf{v}^*)$ for the primal problem (9). The value of this problem should be such that the feasibility constraint holds with equality, which can be found as a solution to $\tilde{K}(v_0) = -\frac{1-\beta^T}{1-\beta} e$. Then from the maximization problem (48) we generate the vector $v_0(\cdot)$. Finally, we use the policy functions to the Bellman equation (44) to generate the solution $(\mathbf{u}^*, \mathbf{v}^*)$, analogous to the i.i.d. case.

Infinite-Period Economy. We now turn to the recursive formulation in the infinite-period economy, $T = \infty$. Assume for simplicity that the utility function is bounded, i.e. $\mathbb{U} = [\underline{u}, \bar{u}]$. Let \mathcal{V} be the set of promised utility vectors $v(\cdot)$ for which there exists an allocation \mathbf{u} such that

$$v(\theta_{(j)}) = \sum_{t=1}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi(\theta^t | \theta_{(j)}) \theta_t u_t(\theta^t), \quad \forall j \in \{1, \dots, |\Theta|\}, \quad (49)$$

and for all $t \geq 1$, for all $\theta^{t-1}, \theta, \hat{\theta}$,

$$\begin{aligned} & \theta u_t(\theta^{t-1}, \theta) + \beta \left\{ \sum_{s=1}^{\infty} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta^{t+s} | \theta) \theta_{t+s} u_{t+s}(\theta^{t-1}, \theta, \theta_{t+1}^{t+s}) \right\} \\ & \geq \theta u_t(\theta^{t-1}, \hat{\theta}) + \beta \left\{ \sum_{s=1}^{\infty} \sum_{\theta^{t+s} \in \Theta^{t+s}} \beta^{s-1} \pi(\theta^{t+s} | \theta) \theta_{t+s} u_{t+s}(\theta^{t-1}, \hat{\theta}, \theta_{t+1}^{t+s}) \right\}. \end{aligned} \quad (50)$$

For any $\theta_- \in \Theta$ and $v(\cdot) \in \mathcal{V}$, the Bellman equation writes

$$K(v(\cdot), \theta_-) = \max_{\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}} \sum_{\theta \in \Theta} \pi(\theta | \theta_-) [-C(u(\theta)) + \beta K(w(\theta|\cdot), \theta)] \quad (51)$$

subject to (45), (46), and $u(\theta) \in \mathbb{U}$, $w(\theta|\cdot) \in \mathcal{V}$ for all θ .

This Bellman equation is a direct extension of the Bellman equation (23) in the i.i.d. case. The need to keep track of a larger number of state variables in the case of general Markov shocks follows from our discussion in the finite period economy. One additional consideration that (51) introduces is that it is defined over a set \mathcal{V} , which needs to be found. The work of Abreu et al.

(1990) provides a method of finding this set and characterizing its properties.

Proposition 8. *The set \mathcal{V} is non-empty, compact and convex. It is the largest bounded fixed point of the operator \mathcal{A} defined for an arbitrary compact set $\tilde{\mathcal{V}} \subset \mathbb{R}^{|\Theta|}$ as*

$$\mathcal{A}\tilde{\mathcal{V}} = \left\{ v(\cdot) \text{ s.t. } \exists \{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta} : (45), (46) \text{ hold and } w(\theta|\cdot) \in \tilde{\mathcal{V}}, \forall \theta \right\}.$$

It is the limit of the monotonically decreasing sequence of compact sets $\{\mathcal{V}_n\}_{n=0,1,\dots}$ defined as $\mathcal{V}_0 = \left[\frac{\theta_{(1)}}{1-\beta} \underline{u}, \frac{\theta_{(|\Theta|)}}{1-\beta} \bar{u} \right]^{|\Theta|}$ and $\mathcal{V}_n = \mathcal{A}\mathcal{V}_{n-1}$ for $n \geq 1$, so that $\mathcal{V} = \lim_{n \rightarrow \infty} \mathcal{V}_n = \bigcap_{n=1}^{\infty} \mathcal{V}_n$.

Proof. The set \mathcal{V} is nonempty because any allocation that is independent of the report is incentive compatible. \mathcal{V} is convex since $v_{\mathbf{u}}(\theta_{(j)})$ is affine in \mathbf{u} for all $j \in \{1, \dots, |\Theta|\}$, where $v_{\mathbf{u}}(\theta_{(j)})$ is defined by the right hand side of (49). The construction of \mathcal{V} as the largest compact fixed point of the operator \mathcal{A} follows from the results of Abreu et al. (1990). Here we give a simple proof that \mathcal{V} is compact and is a fixed point of \mathcal{A} .

Let \mathcal{U} denote the space of allocations $\mathbf{u} = \{u_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$, with $u_t(\theta^t) \in [\underline{u}, \bar{u}]$ for all $t \geq 1$, $\theta^t \in \Theta^t$. Since $|\Theta^t| < \infty$ for all $t \geq 1$, \mathcal{U} is the countable product of the compact metric spaces $[\underline{u}, \bar{u}]$. Embedding \mathcal{U} with the product topology, we obtain that \mathcal{U} is a compact metric space (the compactness follows from a diagonalization argument). A sequence $\mathbf{u}^{(n)}$ in \mathcal{U} converges as $n \rightarrow \infty$ if and only if all of its projections $u_t^{(n)}(\theta^t)$ converge in $[\underline{u}, \bar{u}]$ as $n \rightarrow \infty$.

We now show that \mathcal{V} is compact. Since the utility function is bounded, $v_{\mathbf{u}}$ is bounded and hence \mathcal{V} is bounded. To prove that \mathcal{V} is closed, let $\{\vec{v}^{(n)}\}_{n=1}^{\infty}$ be a Cauchy sequence in \mathcal{V} , and let $\vec{v}^{(\infty)} = \{v^{(\infty)}(\theta_{(j)})\}_{j=1, \dots, J}$ its limit. Let $\{\mathbf{u}^{(n)}\}_{n=1}^{\infty}$ be a sequence of allocations such that $v^{(n)}(\theta_{(j)}) = v_{\mathbf{u}^{(n)}}(\theta_{(j)})$ for all $j \in \{1, \dots, |\Theta|\}$ and all $n \geq 1$. Since \mathcal{U} is compact, $\{\mathbf{u}^{(n)}\}_{n=1}^{\infty}$ contains a convergent subsequence $\{\mathbf{u}^{(\varphi(n))}\}_{n=1}^{\infty}$, denote by $\mathbf{u}^{(\infty)}$ its limit. We have $\mathbf{u}^{(\infty)} \in \mathcal{U}$. Since $v_{\mathbf{u}}(\theta_{(j)})$ is continuous in \mathbf{u} we get, for all $j \in \{1, \dots, |\Theta|\}$,

$$v^{(\infty)}(\theta_{(j)}) = \lim_{n \rightarrow \infty} v^{(\varphi(n))}(\theta_{(j)}) = \lim_{n \rightarrow \infty} v_{\mathbf{u}^{(\varphi(n))}}(\theta_{(j)}) = v_{\mathbf{u}^{(\infty)}}(\theta_{(j)}).$$

Finally, since $\mathbf{u}^{(n)}$ satisfies the incentive constraints (50), by continuity we obtain that $\mathbf{u}^{(\infty)}$ satisfies (50) as well. Thus $\vec{v}^{(\infty)} \in \mathcal{V}$, hence \mathcal{V} is closed. Since $\mathcal{V} \subset \mathbb{R}^{|\Theta|}$, we obtain that \mathcal{V} is compact.

Next we show that \mathcal{V} is a fixed point of \mathcal{A} , that is $\mathcal{A}\mathcal{V} = \mathcal{V}$. First let $\vec{v} \in \mathcal{V}$. Then there exists $\mathbf{u} = \{u_t(\theta^t)\}_{t, \theta^t} \in \mathcal{U}$ that satisfies the incentive constraints (50) and delivers $v(\theta_{(j)}) = v_{\mathbf{u}}(\theta_{(j)})$ for all $j \in \{1, \dots, |\Theta|\}$. Define the allocation rule $\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}$ by $u(\theta) = u_1(\theta)$ and $w(\theta|\theta_{(j')}) = v_{\mathbf{u}_2^\infty(\theta)}(\theta_{(j')})$, where $\mathbf{u}_2^\infty(\theta)$ is the continuation of the allocation \mathbf{u} from period 2 onwards given the history $\theta^1 = \theta$. We have $w(\theta|\cdot) \in \mathcal{V}$ for all $\theta \in \Theta$ because the allocation

$\mathbf{u}_2^\infty(\theta)$ satisfies the incentive compatibility condition (50) after all histories. Moreover, we have

$$\begin{aligned} \theta u(\theta) + \beta w(\theta|\theta) &= \theta u_1(\theta) + \beta v_{\mathbf{u}_2^\infty(\theta)}(\theta) \\ &\geq \theta u_1(\hat{\theta}) + \beta v_{\mathbf{u}_2^\infty(\hat{\theta})}(\theta) = \theta u(\hat{\theta}) + \beta w(\hat{\theta}|\theta), \end{aligned}$$

where the inequality follows from (50). Hence $\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}$ satisfies (46). Finally, by construction $\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}$ satisfies (45). Thus, $\vec{v} \in \mathcal{AV}$ and hence $\mathcal{V} \subset \mathcal{AV}$. For the converse, suppose that $\vec{v} \in \mathcal{AV}$. Then there exists some allocation rule $\{u(\theta), w(\theta|\cdot)\}_{\theta \in \Theta}$ such that the incentive constraints (46) hold and $w(\theta|\cdot) \in \mathcal{V}$ for all θ . Define an allocation \mathbf{u} as follows. Let $u_1(\theta) = u(\theta)$. For each $\theta \in \Theta$, since $w(\theta|\cdot) \in \mathcal{V}$ there exists some allocation $\tilde{\mathbf{u}}(\theta)$ such that $v_{\tilde{\mathbf{u}}(\theta)}(\theta_{(j)}) = w(\theta|\theta_{(j)})$ for all $j \in \{1, \dots, |\Theta|\}$. Define $\mathbf{u}_2^\infty(\theta) = \tilde{\mathbf{u}}(\theta)$. The allocation \mathbf{u} constructed in this way is in \mathcal{U} , satisfies the incentive constraints (50), and delivers $v_{\mathbf{u}}(\theta_{(j)}) = v(\theta_{(j)})$. Thus, $\vec{v} \in \mathcal{V}$ and hence $\mathcal{AV} \subset \mathcal{V}$. \square

2.5.1 Continuum of shocks and the first order approach

The previous section provides a general way to characterize recursively the solution to the optimal insurance problem when shocks are Markovian. One practical difficulty in using the Bellman equation (51) in applications is that the dimensionality of the state space grows with the number of shocks. As the number of shocks becomes large, solving problem (51) becomes intractable. To keep the problem manageable, it is useful to have a method that keeps the number of state variables small.

One approach is to guess that only some of the incentive constraints (38) bind at the optimum. In this case all the non-binding constraints can be dropped, which also eliminates the need to keep track of the state variables associated with that state. The natural candidate for binding constraints are the local constraints, that ensure that a type θ does not want to mimic the types closest to his. In this section we describe how to construct this relaxed problem and provide sufficient conditions that can be verified ex-post to make sure that the dropped incentive constraints are satisfied.¹²

This analysis can be done with a discrete number of shocks, but it becomes particularly simple if instead we allow for a continuum of shocks. In this case applying the envelope theorem to the incentive compatibility condition gives a simple and tractable way to derive the Bellman equation. This problem has been analyzed in detail by Kapička (2013) and Pavan et al. (2014), and here we follow the exposition of the former.

Let the taste shocks θ_t in each period belong to an interval $\Theta = (\underline{\theta}, \bar{\theta}) \subset \mathbb{R}_+^*$, with $\bar{\theta} < \infty$. We assume that the stochastic process for the shocks θ_t is Markov with continuous density $\pi(\theta_t|\theta_{t-1})$. We use $\pi_s(\cdot|\theta_t)$ to denote the p.d.f. of histories θ_{t+1}^{t+s} given that the shock θ_t

¹²It is important to keep in mind that there is a large class of incentive problems with persistent shocks in which non-local incentive constraints bind (Battaglini and Lamba (2015)) and thus the relaxed problem may not satisfy the sufficient conditions.

occured in period t , that is,

$$\pi_s(\theta_{t+1}^{t+s} | \theta_t) = \pi(\theta_{t+s} | \theta_{t+s-1}) \times \dots \times \pi(\theta_{t+1} | \theta_t).$$

Assume as in the previous section that these probabilities are generated from the seed value $\theta_0 = \theta_{(1)}$. We make the following assumptions:

Assumption 4. *Assume that the density $\pi(\theta | \cdot)$ is uniformly Lipschitz-continuous for all θ , and that the derivatives $\hat{\pi}(\theta | \theta_-) \equiv \frac{\partial \pi(\theta | \theta_-)}{\partial \theta_-}$ exist and are uniformly bounded.*

These assumptions can be substantially relaxed (see Kapička (2013), Pavan et al. (2014) for more general treatments of the problem), but they considerably simplify our analysis. To simplify the integrability conditions, we further assume in this section that the utility function is bounded.

Having a continuum of shocks does not change the arguments leading to the recursive characterization of the constraints (41), (40), with the only difference that the sum over a finite number of shock realizations in equation (41) is now replaced by integrals. Constraints (41) and (40) can be written as

$$v_t(\theta^{t-1}, \hat{\theta}_t | \theta_t) = \int_{\Theta} [\theta' u_{t+1}(\theta^{t-1}, \hat{\theta}_t, \theta') + \beta v_{t+1}(\theta^{t-1}, \hat{\theta}_t, \theta' | \theta_t)] \pi(\theta' | \theta_t) d\theta', \quad \forall \theta^t, \hat{\theta}_t, \quad (52)$$

and

$$\theta_t u_t(\theta^{t-1}, \theta_t) + \beta v_t(\theta^{t-1}, \theta_t | \theta_t) = \max_{\hat{\theta} \in \Theta} \left\{ \theta_t u_t(\theta^{t-1}, \hat{\theta}) + \beta v_t(\theta^{t-1}, \hat{\theta} | \theta_t) \right\}, \quad \forall \theta^{t-1}, \theta_t. \quad (53)$$

Lemma 4. *Suppose Assumption 4 is satisfied and the utility function is bounded. Then the function $v_t(\theta^{t-1}, \theta_t | \cdot)$ is differentiable with respect to the realized period- t type θ for each history of reports $\theta^t = (\theta^{t-1}, \theta_t)$ and its derivative evaluated at θ_t is given by*

$$\hat{v}_t(\theta^t) = \int_{\Theta} [\theta' u_{t+1}(\theta^t, \theta') + \beta v_{t+1}(\theta^t, \theta' | \theta^t)] \hat{\pi}(\theta' | \theta_t) d\theta'. \quad (54)$$

Moreover, if an allocation is incentive compatible, then for all $t \geq 1$, $\theta^t \in \Theta^t$,

$$\theta_t u_t(\theta^{t-1}, \theta_t) + \beta v_t(\theta^{t-1}, \theta_t | \theta_t) = \int_{\underline{\theta}}^{\theta_t} \{u_t(\theta^{t-1}, \theta) + \beta \hat{v}_t(\theta^{t-1}, \theta)\} d\theta + \nu(\theta^{t-1}), \quad (55)$$

where $\nu(\theta^{t-1}) = \lim_{\theta \rightarrow \underline{\theta}} \{\theta u_t(\theta^{t-1}, \theta) + \beta v_t(\theta^{t-1}, \theta | \theta)\}$.

Proof. Let $S_{t+1}(\theta^t, \theta') = \theta' u_{t+1}(\theta^t, \theta') + \beta v_{t+1}(\theta^t, \theta' | \theta')$. Then

$$\begin{aligned} \frac{v_t(\theta^{t-1}, \theta_t | \theta + \Delta\theta) - v_t(\theta^{t-1}, \theta_t | \theta)}{\Delta\theta} &= \int_{\Theta} S_{t+1}(\theta^t, \theta') \frac{\pi(\theta' | \theta + \Delta\theta) - \pi(\theta' | \theta)}{\Delta\theta} d\theta' \\ &\xrightarrow{\Delta\theta \rightarrow 0} \int_{\Theta} S_{t+1}(\theta^t, \theta') \hat{\pi}(\theta' | \theta) d\theta', \end{aligned}$$

where the last step follows from the dominated convergence theorem, noting that $S_{t+1}(\theta^t, \theta')$ and $\left| \frac{\pi(\theta' | \theta + \Delta\theta) - \pi(\theta' | \theta)}{\Delta\theta} \right|$ are bounded by the uniform Lipschitz continuity of $\pi(\theta' | \cdot)$. Hence $v_t(\theta^{t-1}, \theta_t | \cdot)$ is differentiable, and it is Lipschitz continuous on $[\eta, \bar{\theta}]$ for all $\eta > \underline{\theta}$ since $\hat{\pi}(\theta' | \theta)$ is uniformly bounded.

Let $\hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t) = \theta_t u_t(\theta^{t-1}, \hat{\theta}_t) + \beta v_t(\theta^{t-1}, \hat{\theta}_t | \theta_t)$. Then \hat{S}_t is differentiable in θ_t on $(\eta, \bar{\theta})$ (denote by $\hat{S}_{\theta,t}$ its derivative) and Lipschitz in θ_t on $[\eta, \bar{\theta}]$. Hence it is absolutely continuous and has a bounded derivative with respect to θ_t on $(\eta, \bar{\theta})$. Since the allocation is incentive compatible, $\hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t)$ is maximized at $\hat{\theta}_t = \theta_t$. By Theorem 2 in Milgrom and Segal (2002), $\hat{S}_t(\theta^{t-1}, \theta_t | \theta_t)$ can be represented as an integral of its derivative:

$$\begin{aligned} \hat{S}_t(\theta^{t-1}, \theta_t | \theta_t) &= \int_{\eta}^{\theta_t} \hat{S}_{\theta,t}(\theta^{t-1}, \theta, \theta) d\theta + \hat{S}_t(\theta^{t-1}, \eta | \eta) \\ &= \int_{\eta}^{\theta_t} \{u_t(\theta^{t-1}, \theta) + \beta v_{\theta,t}(\theta^{t-1}, \theta | \theta)\} d\theta + \hat{S}_t(\theta^{t-1}, \eta | \eta). \end{aligned}$$

Take the limit as $\eta \rightarrow \underline{\theta}$ to get expression (55). \square

We can now define a relaxed problem by replacing the temporary incentive compatibility constraints (53) by the envelope condition (55) for all histories. This substantially simplifies the analysis, as the latter constraint only depends on the lifetime utility and marginal lifetime utility of the truthteller, rather than the continuation utility of all possible types as in Section 2.5. In the recursive formulation of the planner's problem, the choice variables are the current utility $u(\theta)$, the continuation utility of the truthtelling agent $w(\theta)$, and the marginal change in the continuation utility of the truthtelling agent $\hat{w}(\theta)$. The state variables are the reported taste shock realization θ_- in the previous period, the promised utility v of an agent who truthfully announced θ_- last period, and the marginal promised utility \hat{v} of an agent who truthfully announced θ_- last period.

We now show the recursive formulation of the relaxed problem in an infinite-period economy. Let $\hat{\mathcal{V}}(\theta_-)$ denote the set of lifetime utility and marginal lifetime utility pairs $(v, \hat{v}) \in \mathbb{R}^2$ for which there exist values $\{u(\theta), w(\theta), \hat{w}(\theta)\}_{\theta \in \Theta}$ such that the following conditions hold:

(i) the envelope condition:

$$\theta u(\theta) + \beta w(\theta) = \int_{\underline{\theta}}^{\theta} \{u(\theta') + \beta \hat{w}(\theta')\} d\theta' + \lim_{\theta \rightarrow \underline{\theta}} \{\theta u(\theta) + \beta w(\theta)\}, \quad \forall \theta \in \Theta, \quad (56)$$

(ii) the promise-keeping constraint:

$$v = \int_{\Theta} \{\theta' u(\theta') + \beta w(\theta')\} \pi(\theta' | \theta_-) d\theta', \quad (57)$$

(iii) the marginal promise-keeping constraint:

$$\hat{v} = \int_{\Theta} \{\theta' u(\theta') + \beta w(\theta')\} \hat{\pi}(\theta' | \theta_-) d\theta', \quad (58)$$

(iv) $(w(\theta), \hat{w}(\theta)) \in \hat{\mathcal{V}}(\theta)$ for all $\theta \in \Theta$.

Note that in general $\hat{\mathcal{V}}(\theta)$ depends on the realized value of θ . It can be characterized along the lines of Proposition 8.

For any $\theta_- \in \Theta$ and pair $(v, \hat{v}) \in \hat{\mathcal{V}}(\theta_-)$, the Bellman equation writes

$$K(v, \hat{v}, \theta_-) = \sup_{(\bar{u}, \bar{w}, \bar{w})} \int_{\Theta} \{-C(u(\theta)) + \beta K(w(\theta), \hat{w}(\theta), \theta)\} \pi(\theta | \theta_-) d\theta \quad (59)$$

subject to (56), (57), (58), and $u(\theta) \in \mathbb{U}$, $(w(\theta), \hat{w}(\theta)) \in \hat{\mathcal{V}}(\theta)$ for all $\theta \in \Theta$.

We finally discuss when the relaxed problem gives the solution to the original problem. The envelope condition (55) is necessary but not sufficient for an allocation to be temporarily incentive compatible. A simple sufficient condition is given in the next proposition:

Proposition 9. *Suppose that an allocation \mathbf{u} satisfies the envelope condition (55) and, in addition,*

$$u_t(\theta^{t-1}, \hat{\theta}_t) + \beta \hat{v}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t) \quad (60)$$

is increasing in $\hat{\theta}_t$ for all t, θ^{t-1} and almost all θ_t , where $v_t(\theta^{t-1}, \hat{\theta}_t | \theta_t) \equiv \frac{\partial}{\partial \theta} v_t(\theta^{t-1}, \hat{\theta}_t | \theta_t)$. Then \mathbf{u} is temporary incentive compatible.

Proof. Fix t, θ^{t-1} , and let $S_t(\theta^{t-1}, \theta_t) \equiv \hat{S}_t(\theta^{t-1}, \theta_t | \theta_t)$. An allocation is temporarily incentive compatible if $S_t(\theta^{t-1}, \theta_t) \geq \hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t)$ for all $\theta^{t-1}, \theta_t, \hat{\theta}_t$. Equation (55) shows that $S_t(\theta^{t-1}, \cdot)$ is differentiable for almost all $\theta \in \Theta$ with

$$\frac{\partial}{\partial \theta} S_t(\theta^{t-1}, \theta) = u_t(\theta^{t-1}, \theta) + \beta \hat{v}_t(\theta^{t-1}, \theta).$$

We thus have

$$\begin{aligned} & S_t(\theta^{t-1}, \theta_t) - S_t(\theta^{t-1}, \hat{\theta}_t) \\ &= \int_{\hat{\theta}_t}^{\theta_t} \frac{\partial}{\partial \theta'} S_t(\theta^{t-1}, \theta') d\theta' = \int_{\hat{\theta}_t}^{\theta_t} \{u_t(\theta^{t-1}, \theta') + \beta \hat{v}_t(\theta^{t-1}, \theta' | \theta_t)\} d\theta' \\ &\geq \int_{\hat{\theta}_t}^{\theta_t} \{u_t(\theta^{t-1}, \hat{\theta}_t) + \beta \hat{v}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t)\} d\theta' = \hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta_t) - S_t(\theta^{t-1}, \hat{\theta}_t), \end{aligned}$$

where the inequality follows from the monotonicity of (60), and the last equality follows from the differentiability of $\hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta)$, with $\frac{\partial}{\partial \theta} \hat{S}_t(\theta^{t-1}, \hat{\theta}_t | \theta) = u_t(\theta^{t-1}, \hat{\theta}_t) + \beta \hat{v}_t(\theta^{t-1}, \hat{\theta}_t | \theta)$. We obtain that \mathbf{u} is temporarily incentive compatible. \square

If the shocks are i.i.d., the second term in expression (60) drops out and the proposition is equivalent to a simple requirement that $u_t(\theta^t)$ is increasing in θ_t , and one can show that this requirement is necessary as well. In the static setting, u satisfies the Spence-Mirrlees condition and this sufficient condition reduces to the familiar *necessary and sufficient* condition that *allocations* are monotonic, see Myerson (1981). Unfortunately, in the dynamic model with persistent shocks, the monotonicity condition on (60) is not necessary, and moreover there is no one-to-one mapping between marginal lifetime utilities and allocations. Moreover, in practice condition (60) is difficult to verify directly, and we have to check ex post (possibly numerically) in specific applications (see, e.g., Section 4.1) whether the solution to the relaxed problem is indeed an optimal allocation.

2.6 Other Models

The techniques that we introduced in the previous sections in the context of the taste shock model can be easily applied to many more environments. First, Green (1987) and Thomas and Worrall (1990) study a model closely related to the one we analyzed above, in which the agent receives privately observed i.i.d. or persistent endowment (or income) shocks $\theta_t \in \Theta$: in each period $t \geq 1$, the agent observes his income shock θ_t and reports its realization to the planner, who then provides a transfer $\tau_t(\theta^t)$ to the agent. Second, Spear and Srivastava (1987) and Phelan and Townsend (1991) study a moral hazard model in which agents exert a privately observed effort level $\theta_t \in \Theta$ in each period. The output produced from that effort is stochastic and observable to the planner. The case where current effort affects only current output corresponds to the i.i.d. assumption 3 in the taste shock model, while the case where current effort also affects future output corresponds to the taste shock model with persistent types. Third, Thomas and Worrall (1988), Kocherlakota (1996), Ligon et al. (2002) show that models of limited commitment, in which there is no asymmetry of information but one or both parties are free to walk away to the insurance contract, can be analyzed using similar recursive techniques using promised utility as state variable; we discuss applications of these models in Section 4.

Here we describe briefly how to apply our recursive techniques to a model of moral hazard. Agents exert an effort level $\theta_t \in \Theta = [0, \infty)$ in each period. The planner does not observe the agent's effort, but only the (random) output produced from that effort, $y_t \in Y = \{y_{(1)}, y_{(2)}\}$, with $0 = y_{(1)} < y_{(2)}$. The flow utility at time t is $U(c_t) - h(\theta_t)$, where the utility from consumption $U(\cdot) : \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing and strictly concave, and the disutility of effort $h(\cdot) : \mathbb{R}_+ \rightarrow \mathbb{R}$ is twice differentiable, strictly increasing and strictly convex with $h(0) = 0$ and $h'(0) \geq 0$.

We assume that the probability of output $y_t \in Y$ in period t depends only on the effort $\theta_t \in \Theta$ exerted by the agent in the *current* period.¹³ We denote it by $\pi(y_t | \theta_t)$, with $0 = \pi(y_{(2)} | 0) < \pi(y_{(2)} | \theta) < 1$ for all $\theta \in (0, \infty)$, and $\pi(y_{(2)} | \cdot)$ is twice differentiable with $\pi_\theta(y_{(2)} | \cdot) > 0$. An allocation in this model is a pair of sequences of functions $\boldsymbol{\theta} = \{\theta_t(y^{t-1})\}_{t \geq 1}$ (with $y^0 = \emptyset$) describing the effort recommended by the planner to the agent given the observed history at the beginning of each period t , and a sequence of utility payments $\mathbf{u} = \{u_t(y^t)\}_{t \geq 1}$ given the observed history of output at the end of each period t . The planner chooses the incentive compatible allocation $\{\mathbf{c}, \boldsymbol{\theta}\}$ that minimizes the cost of delivering lifetime utility v_0 , that is, letting $C \equiv U^{-1}$,

$$\begin{aligned}
K(v_0) \equiv \quad & \max_{\boldsymbol{\theta}, \mathbf{u}} \quad \mathbb{E}^{\boldsymbol{\theta}} \left[\sum_{t=1}^{\infty} \beta^{t-1} \{y_t - C(u_t(y^t))\} \right] \\
\text{subject to} \quad & \mathbb{E}^{\boldsymbol{\theta}} \left[\sum_{t=1}^{\infty} \beta^{t-1} \{u_t(y^t) - h(\theta_t(y^{t-1}))\} \right] = v_0, \\
& \mathbb{E}^{\hat{\boldsymbol{\theta}}} \left[\sum_{t=1}^{\infty} \beta^{t-1} \{u_t(y^t) - h(\hat{\theta}_t(y^{t-1}))\} \right] \leq v_0, \quad \forall \hat{\boldsymbol{\theta}},
\end{aligned} \tag{61}$$

where the superscripts over expectations $\mathbb{E}^{\boldsymbol{\theta}}$ and $\mathbb{E}^{\hat{\boldsymbol{\theta}}}$ indicate that the probability distributions over the paths of output $\{y^t\}_{t \geq 1}$ inside the brackets depend on the agent's corresponding strategies $\boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}}$ (respectively), that is, for any t and random variable $X_t(y^t)$ we let $\mathbb{E}^{\boldsymbol{\theta}}[X_t] \equiv \sum_{y^t \in Y^t} \pi_t(y^t | \boldsymbol{\theta}^t) X_t(y^t)$. Thus, each expectation in the incentive constraint depends on the agent's effort directly through the cost of effort $h(\theta_t)$, and indirectly through its effect on the probability distribution $\pi_t(y^t | \boldsymbol{\theta})$ over the paths of y^t .

Defining the continuation utility of an obedient agent up to and after date t as in (11), we can rewrite this problem recursively:

$$\begin{aligned}
K(v) = \quad & \max_{\theta \geq 0, \bar{u} \geq u, \bar{w} \geq v} \sum_{y \in Y} \pi(y | \theta) [y - C(u(y)) + \beta K(w(y))] \\
\text{s.t.} \quad & v = \sum_{y \in Y} \pi(y | \theta) [u(y) - h(\theta) + \beta w(y)], \\
& \pi_\theta(y_{(2)} | \theta) [(u(y_{(2)}) - u(y_{(1)})) + \beta (w(y_{(2)}) - w(y_{(1)}))] - h'(\theta) \leq 0 \\
& \text{with equality if } \theta > 0,
\end{aligned} \tag{62}$$

where the incentive compatibility constraint is replaced by a first-order condition, assuming for

¹³The analysis of the case where current effort also affects future output is slightly more involved than that of Section 2.5. This is because there is a form of non-separability of the agent's lifetime utility (incentives in a given period depend no longer only on his current true type and past reports, but also on his past true types) which implies that truthful revelation does not necessarily hold off the equilibrium path (i.e., after the agent has deviated from the recommended action in the past); see Example of Section S.5. in Pavan et al. (2014). Thus, after a deviation an agent may prefer to engage in a strategy of infinite deviations. Fernandes and Phelan (2000) nevertheless show how to write a recursive formulation of this problem.

simplicity that the first order approach is valid.

Following the steps leading to Proposition 5 gives the following characterization of the solution to the planner’s problem. For any interior v , the optimal contract $(\theta_v, \vec{u}_v, \vec{w}_v)$ satisfies the following martingale property (with respect to the probability measure associated with the optimum effort strategy \mathbb{P}^θ):

$$K'(v) = \mathbb{E}^{\theta_v} [-C'(u_v)] = \mathbb{E}^{\theta_v} [K'(w_v)]. \quad (63)$$

This can be re-written as:

$$\frac{1}{u'(c_t(y^t))} = \sum_{y_{t+1} \in Y} \pi(y_{t+1} | \theta_{t+1}(y^t)) \frac{1}{u'(c_{t+1}(y^t, y_{t+1}))}, \quad (64)$$

and is known in the literature as the *Inverse Euler Equation* (see Diamond and Mirrlees (1978), Rogerson (1985), Spear and Srivastava (1987), Golosov et al. (2003)). We derive implications of this equation in Section 4.1 and show by comparing it to the individual’s Euler equation that agents’ savings are constrained in the optimal insurance arrangement.

We can analyze this problem (62) along the lines of the proof of Proposition 5. A utility-effort pair (v, θ_v) is absorbing if and only if $\theta_v = 0$ and $(u_v(y), w_v(y)) = ((1 - \beta)v, v)$. The recommended effort θ_v is strictly positive as long as the promised utility is small enough, $v < \bar{v}$. If $h'(0) = 0$, we find $\bar{v} = \infty$, and the Martingale Convergence Theorem implies that immiseration occurs: $v_t(\theta^t) \rightarrow \underline{v}$ as $t \rightarrow \infty$ with probability 1. If instead $h'(0) > 0$, the principal may eventually “retire” the agent (i.e., recommend effort $\theta_t(y^t) = 0$ and provide constant consumption $c_t(y^t) = c$) when $v_t(\theta^t) \geq \bar{v}$, as for a large enough promised utility the benefit of inducing him to work outweighs the cost of providing the necessary incentives and compensating him for the higher effort.

3 Advanced Topics

In this Section we discuss three additional topics that significantly expand the applicability of the recursive contracts. In Section 3.1 we overview the theory of Lagrange multipliers and show how it can help solve many dynamic incentive problems recursively even if they do not fit into the canonical setup described in Section 2. Section 3.2 shows how to extend the analysis to settings in which the ability of the principal to commit is imperfect. Finally, in Section 3.3, we describe the analysis of dynamic contracting problems in continuous time using martingale methods. Throughout this section we do not aim at the same level of rigor as in Section 2; we omit several technical details and refer to the relevant papers for the complete proofs.

3.1 Lagrange multipliers

The key feature that allowed us to analyze the dynamic contracting problem (17) is that we could write the incentive constraints in a simple recursive form. In many applications, however, the optimal contracting problem often has additional constraints that cannot easily be written recursively. For example, if we replaced the present value budget constraint (2) with a requirement that the total consumption of all agents should be equal to the total endowment in *each* period, the previous method could not be applied directly. In this section we describe a simple approach that allows us to extend our analysis to such problems. The main idea behind this approach is to assign Lagrange multipliers to all the constraints that do not have a straightforward recursive representation, and to apply the techniques developed in the previous sections to the resulting Lagrangian.

We start in Section 3.1.1 by giving a general theoretical background about the properties of Lagrange multipliers in infinite dimensional spaces. Infinite dimensional spaces are common in macroeconomic applications but the Lagrangian techniques are more subtle in such spaces than in finite dimensions. The main results of this section are, first, Theorems 3 and 4, which provide conditions under which the Lagrangian exists and characterize the solution to the constrained optimization problem, and second, Theorem 5, which provides sufficient conditions that ensure that the Lagrangian can be written as an infinite sum, allowing us to apply the standard techniques familiar from finite-dimensional optimization theory. Sections 3.1.2 through 3.1.4 give several examples of applications of these techniques. The reader only interested in practical applications can skip Section 3.1.1 in the first reading.

3.1.1 Main theoretical results

The classical reference about using Lagrange multipliers to solve optimization problems is Luenberger (1969). Here we state two main results from this book, adapting them to our setting. To use this approach, we need to set our problem in abstract linear spaces.¹⁴ Before starting our analysis we introduce the notions of convex cones and mappings, dual spaces, and l_p spaces.

First, let P be a convex cone in a vector space \mathcal{V} , that is, P satisfies $\alpha x + \beta y \in P$ for all $x, y \in P$ and $\alpha, \beta > 0$. This convex cone defines a partial order \leq on \mathcal{V} , such that $x \geq y$ if $x - y \in P$. By definition, P is the positive cone with respect to this partial order, i.e., the subset $\mathcal{V}^+ = \{x \in \mathcal{V} : x \geq 0\}$. We write $x > 0$ if x is an interior point of the positive cone P . By introducing a cone defining the positive vectors in the vector space \mathcal{V} , we thus define an ordering relation \leq and make it possible to consider inequality problems in the abstract vector space \mathcal{V} . (Often the positive cones of the vector spaces we consider are constructed naturally, e.g., the positive orthant of \mathbb{R}^n , or the non-negative continuous functions of $\mathcal{C}([a, b])$.) A mapping $G : \mathcal{V}_1 \rightarrow \mathcal{V}_2$ from the vector space \mathcal{V}_1 to a vector space \mathcal{V}_2 having a cone P defined as the positive cone is said to be *convex* if the domain Ω of G is a convex set and if

¹⁴For a review of basic functional analysis, see Luenberger (1969), or Chapters 3 and 15 in Stokey et al. (1989).

$G(\alpha x_1 + (1 - \alpha)x_2) \leq \alpha G(x_1) + (1 - \alpha)G(x_2)$ for all $x_1, x_2 \in \Omega$ and all $\alpha \in (0, 1)$.

Second, the dual \mathcal{V}^* of a normed vector space \mathcal{V} is the space of all bounded linear functionals on \mathcal{V} , i.e., $f : \mathcal{V} \rightarrow \mathbb{R}$. The norm of an element $f \in \mathcal{V}^*$ is $\|f\| = \sup_{\|x\| \leq 1} |f(x)|$. The value of the linear functional $x^* \in \mathcal{V}^*$ at the point $x \in \mathcal{V}$, that is $x^*(x)$, is denoted by $\langle x, x^* \rangle$. For $1 \leq p < \infty$, the space l_p consists of all sequences of scalars $\{u_1, u_2, \dots\}$ for which $\sum_{n=1}^{\infty} |u_n|^p < \infty$, and the space l_∞ consists of the bounded sequences. The norm of an element $u = \{u_n\}_{n \geq 1} \in l_p$ is defined as $\|u\|_p = (\sum_{i=1}^{\infty} |u_i|^p)^{1/p}$ for $p < \infty$, and as $\|u\|_p = \sup_n |u_n|$ for $p = \infty$. Then for every $p \in [1, \infty)$, the dual space of l_p is l_q , where $q = (1 - p^{-1})^{-1}$. This is because every bounded linear functional f on l_p , $1 \leq p < \infty$ can be represented uniquely in the form $f(u) = \sum_{n=1}^{\infty} v_n u_n$, where $v = \{v_n\}_{n \in \mathbb{N}^*}$ is an element of l_q ; specifically, for all $n \geq 1$, $v_n \equiv f(e_n)$, where $e_n \in l_p$ is the sequence that is identically zero except for a 1 in the n^{th} component. The dual of l_∞ *strictly contains* l_1 . Finally, given a normed space \mathcal{V} together with a positive convex cone $P \subset \mathcal{V}$, it is natural to define a corresponding positive convex cone P^* in the dual space \mathcal{V}^* by $P^* = \{x^* \in \mathcal{V}^* : \forall x \in P, \langle x, x^* \rangle \geq 0\}$.

We can now introduce the theory of Lagrange multipliers. Consider a problem

$$\begin{aligned} \min_{\mathbf{x}} \quad & \varphi(\mathbf{x}) \\ \text{subject to} \quad & \Phi(\mathbf{x}) \leq \mathbf{0}, \quad \mathbf{x} \in \Gamma, \end{aligned} \tag{65}$$

where Γ is a convex subset of a vector space X , $\varphi : \Gamma \rightarrow \mathbb{R}$ is a convex functional, $\Phi : \Gamma \rightarrow Z$ is a convex mapping to a normed vector space Z that has positive cone P . Let Z^* be the dual space of Z and Z_+^* be its positive orthant (i.e., all $\mathbf{z}^* \in Z^*$ such that $\mathbf{z}^* \geq \mathbf{0}$). We assume throughout this section that the minimum of problem (65) is attained. This assumption is not necessary but it simplifies the statement of the theorems, and we will see in our context (Proposition 10) that it can often be verified directly. Theorems 1 p. 217 and Corollary 1 p. 219 in Luenberger (1969) give the main results for solving the minimization problem (65) using Lagrange multipliers.

Theorem 3. *Assume that the minimum in (65) is achieved at $\hat{\mathbf{x}}$. Suppose that P contains an interior point, and that there exists $\mathbf{x}' \in \Gamma$ such that $\Phi(\mathbf{x}') < \mathbf{0}$. Then there is $\hat{\mathbf{z}}^* \in Z_+^*$ such that the Lagrangian*

$$\mathcal{L}(\mathbf{x}, \mathbf{z}^*) = \varphi(\mathbf{x}) + \langle \Phi(\mathbf{x}), \hat{\mathbf{z}}^* \rangle$$

has a saddle point at $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$, i.e.,

$$\mathcal{L}(\hat{\mathbf{x}}, \mathbf{z}^*) \leq \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*) \leq \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*), \quad \forall \mathbf{x} \in \Gamma, \hat{\mathbf{z}}^* \in Z_+^*. \tag{66}$$

Moreover,

$$\langle \Phi(\hat{\mathbf{x}}), \hat{\mathbf{z}}^* \rangle = 0.$$

Theorem 3 establishes that for convex problems there generally exists a Lagrangian such that the solution to the original constrained minimization problem is also a solution to the

minimization of the unconstrained Lagrangian. The next result (Theorem 2 p. 221 in Luenberger (1969)) ensures the sufficiency:

Theorem 4. *Let $X, Z, \Gamma, P, \varphi, \Phi$ be as above and assume that the positive cone $P \subset Z$ is closed. Suppose that there exist $\hat{\mathbf{z}}^* \in Z_+^*$ and an $\hat{\mathbf{x}} \in \Gamma$ such that the Lagrangian $\mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ has a saddle point at $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$. Then $\hat{\mathbf{x}}$ is a solution to (65).*

Thus, if φ and Φ are convex, the positive cone $P \subset Z$ is closed and has non-empty interior, and the regularity condition $\Phi(\mathbf{x}') < 0$ is satisfied, then the saddle point condition is necessary and sufficient for optimality of $\hat{\mathbf{x}}$.

One way to find a saddle point of \mathcal{L} is to use the following result (see Bertsekas et al. (2003)).

Corollary 1. *$(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$ is a saddle point of \mathcal{L} if and only if the equality*

$$\inf_{\mathbf{x} \in \Gamma} \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}, \mathbf{z}^*) = \sup_{\mathbf{z}^* \in Z_+^*} \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*) \quad (67)$$

is satisfied, and

$$\begin{aligned} \hat{\mathbf{x}} &= \arg \min_{\mathbf{x} \in \Gamma} \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}, \mathbf{z}^*), \\ \hat{\mathbf{z}}^* &= \arg \max_{\mathbf{z}^* \in Z_+^*} \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*). \end{aligned} \quad (68)$$

In particular, suppose that the conditions of Theorem 3 hold, so that $\mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ has a saddle point at $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$. Suppose moreover that $\arg \min_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ exists for each $\mathbf{z}^* \in Z_+^*$ and is unique for $\mathbf{z}^* = \hat{\mathbf{z}}^*$. Then $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$ is the solution to $\max_{\mathbf{z}^* \in Z_+^*} \min_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*)$.

Proof. Suppose $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$ is a saddle point. Then

$$\inf_{\mathbf{x} \in \Gamma} \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}, \mathbf{z}^*) \leq \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\hat{\mathbf{x}}, \mathbf{z}^*) = \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*) = \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*) \leq \sup_{\mathbf{z}^* \in Z_+^*} \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*).$$

By the max-min inequality, $\inf_{\mathbf{x} \in \Gamma} \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}, \mathbf{z}^*) \geq \sup_{\mathbf{z}^* \in Z_+^*} \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*)$, establishing that all these inequalities hold with equality, and hence (67) and (68) is satisfied.

Conversely, suppose that (67) and (68) hold. Then

$$\sup_{\mathbf{z}^* \in Z_+^*} \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*) = \inf_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*) \leq \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*) \leq \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\hat{\mathbf{x}}, \mathbf{z}^*) = \inf_{\mathbf{x} \in \Gamma} \sup_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}, \mathbf{z}^*).$$

Using (67) implies that $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$ is a saddle point.

Finally suppose that the conditions of Theorem 3 are satisfied, so that $\mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ has a saddle point at $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$, and that $\mathbf{x}(\mathbf{z}^*) \equiv \arg \min_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ exists for each $\mathbf{z}^* \in Z_+^*$ and is unique for $\mathbf{z}^* = \hat{\mathbf{z}}^*$. Then, by (68) we have $\hat{\mathbf{z}}^* = \arg \max_{\mathbf{z}^* \in Z_+^*} \mathcal{L}(\mathbf{x}(\mathbf{z}^*), \mathbf{z}^*)$. By the uniqueness assumption we have $\mathcal{L}(\mathbf{x}(\hat{\mathbf{z}}^*), \hat{\mathbf{z}}^*) < \mathcal{L}(\mathbf{x}, \hat{\mathbf{z}}^*)$ for all $\mathbf{x} \neq \mathbf{x}(\hat{\mathbf{z}}^*)$, so that the saddle point (66) can only be achieved at $(\mathbf{x}(\hat{\mathbf{z}}^*), \hat{\mathbf{z}}^*)$, establishing that $\hat{\mathbf{x}} = \mathbf{x}(\hat{\mathbf{z}}^*)$. We obtain that the solution to $\max_{\mathbf{z}^* \in Z_+^*} \min_{\mathbf{x} \in \Gamma} \mathcal{L}(\mathbf{x}, \mathbf{z}^*)$ is $(\hat{\mathbf{x}}, \hat{\mathbf{z}}^*)$. \square

The max min problem in Corollary 1 provides a simple way to find the solution to the minimization problem together with the corresponding Lagrangian. The uniqueness qualifier is important for that result; without it there may exist solutions to the max min problem that are not saddle points, i.e., that are not a solution to the original optimization problem (see, e.g., Messner and Pavoni (2004)).

In economic applications, Φ often represents per-period constraints and it can be written as $\Phi = \{\Phi_1, \Phi_2, \dots\}$. The most natural vector space to choose in such situations is the space of bounded sequences, l_∞ . In this case we define the positive cone P of l_∞ as the positive orthant, i.e., the subset of non-negative sequences of l_∞ . Exercise 15.7 in Stokey et al. (1989) shows that l_∞ is the only l_p -space that has a positive orthant with a non-empty interior, which is a requirement needed to apply the theorems above.

A limitation of the space l_∞ is that its dual is complicated. It contains the space of summable sequences l_1 , but it also includes other sequences which are not summable. This makes the analysis difficult because the linear operator $\langle \Phi(\mathbf{x}), \mathbf{z} \rangle$ may take a complicated form. The analysis simplifies if it can be ensured that the mappings φ and Φ are not affected by how \mathbf{x} behaves “at infinity”, in which case we provide an l_1 representation of the Lagrange multipliers and each constraint $\Phi_n(\mathbf{x})$ will have a scalar multiplier λ_n associated with it. For any $\mathbf{x}, \mathbf{y} \in l_\infty$, define an operator $x^T(\mathbf{x}, \mathbf{y})$ as $x^T(\mathbf{x}, \mathbf{y}) = x_t$ if $t \leq T$ and $x^T(\mathbf{x}, \mathbf{y}) = y_t$ if $t > T$. We use the notation $x_t^T(\mathbf{x}, \mathbf{y})$ to denote the t -th element of this operator.

Assumption 5. Let $X, Z = l_\infty$, $\Psi = \{\mathbf{x} \in \Gamma : \varphi(\mathbf{x}) < \infty\}$. Suppose that:

(i) If $(\mathbf{x}, \mathbf{y}) \in \Psi \times l_\infty$ satisfy $x^T(\mathbf{x}, \mathbf{y}) \in \Psi$ for all T large enough, then $\varphi(x^T(\mathbf{x}, \mathbf{y})) \rightarrow \varphi(\mathbf{x})$ as $T \rightarrow \infty$.

(ii) If $\mathbf{x}, \mathbf{y} \in \Gamma$ and $x^T(\mathbf{x}, \mathbf{y}) \in \Gamma$ for all T large enough, then:

- (a) $\forall t, \lim_{T \rightarrow \infty} \Phi_t(x^T(\mathbf{x}, \mathbf{y})) = \Phi_t(\mathbf{x})$,
- (b) $\exists M$ s.t. $\forall T$ large enough, $\|\Phi(x^T(\mathbf{x}, \mathbf{y}))\| \leq M$,
- (c) $\forall T$ large enough, $\lim_{t \rightarrow \infty} [\Phi_t(x^T(\mathbf{x}, \mathbf{y})) - \Phi_t(\mathbf{y})] = 0$.

Le Van and Saglam (2004) prove that under these assumptions the Lagrangian can be written as an infinite sum:¹⁵

Theorem 5. Let $\hat{\mathbf{x}}$ be a solution to (65). Suppose that for all $\mathbf{x} \in \Gamma$, we have $\Phi(\mathbf{x}) \in l_\infty$. Assume that there exists $\mathbf{x}' \in \Gamma$ such that $\Phi(\mathbf{x}') < 0$, that is, $\sup_t \Phi_t(\mathbf{x}') < 0$ (Slater condition). Assume finally that Assumption 5 is satisfied and that $x^T(\hat{\mathbf{x}}, \mathbf{x}') \in \Gamma \cap \Psi$ for all T large enough. Then there exists $\hat{\mathbf{z}}^* \in l_1$ with $\hat{\mathbf{z}}^* \geq \mathbf{0}$ such that

$$\sum_{t=1}^{\infty} \hat{z}_t^* \Phi_t(\hat{\mathbf{x}}) = 0,$$

¹⁵See also Rustichini (1998) who provides an alternative set of sufficient conditions ensuring the summability of the Lagrange multipliers.

and

$$\varphi(\mathbf{x}) + \sum_{t=1}^{\infty} \hat{z}_t^* \Phi_t(\mathbf{x}) \geq \varphi(\hat{\mathbf{x}}) + \sum_{t=1}^{\infty} \hat{z}_t^* \Phi_t(\hat{\mathbf{x}}), \quad \forall \mathbf{x} \in \Gamma.$$

In the next sections, we apply this theory to dynamic contracting problems.

3.1.2 Application: recursive contracts in general equilibrium

Consider a simple modification of the setup in Section 2.3, in which the planner can no longer freely borrow and lend at an exogenous interest rate. Instead we require that the economy-wide feasibility constraint hold period by period, i.e.,

$$\sum_{\theta^t \in \Theta^t} \pi_t(\theta^t) C(u_t(\theta^t)) \leq e, \quad \forall t \geq 1. \quad (69)$$

This problem is analyzed by Atkeson and Lucas (1992). For simplicity we assume $|\Theta| = 2$ and i.i.d. shocks to parallel our discussion in Sections 2.3 and 2.4. Thus we study the problem

$$\max_{\mathbf{u}} \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t(\theta^t) \right] \quad (70)$$

subject to

$$\mathbb{E}_0 [C(u_t(\theta^t))] \leq e, \quad \forall t \geq 1, \quad (71)$$

and

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t \{u_t(\theta^t) - u_t(\sigma^t(\theta^t))\} \right] \geq 0, \quad \forall \sigma. \quad (72)$$

Assume for now that the maximum in the problem (70) is attained for all $e > 0$; we will show this formally below.

Let Γ be the set of sequences $\mathbf{u} = \{u_t(\theta^t)\}_{t \geq 1, \theta^t \in \Theta^t}$, indexed by (t, θ^t) , such that \mathbf{u} satisfies the period-0 incentive constraint (72) and the sequence $\{\mathbb{E}_0[C(u_t)] - e\}_{t=1}^{\infty}$ is bounded in sup-norm. The set Γ is convex and has an interior point, e.g., $u_t(\theta^t) = \varepsilon$ for all t, θ^t and $\varepsilon > 0$ sufficiently small.

We start with the sufficient conditions first. Let X be the space of all infinite sequences, $Z = l_{\infty}$ and $\Phi = \{\Phi_1, \Phi_2, \dots\}$ where $\Phi_t : \Gamma \rightarrow \mathbb{R}$ is defined by $\Phi_t(\mathbf{u}) = \mathbb{E}_0[C(u_t)] - e$. Suppose that we can find a non-negative sequence $\boldsymbol{\lambda} = \{\lambda_t\}_{t=1}^{\infty}$ such that the problem¹⁶

$$\max_{\mathbf{u} \in \Gamma} \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t \right] - \sum_{t=1}^{\infty} \lambda_t \{\mathbb{E}_0[C(u_t)] - e\} \quad (73)$$

has a maximum $\hat{\mathbf{u}}$ and $\mathbb{E}_0[C(\hat{u}_t)] = e$ for all t . To verify that $(\hat{\mathbf{u}}, \boldsymbol{\lambda})$ is a saddle point, observe

¹⁶To be consistent with discussion in Section 3.1.1, we use the fact that minimizing φ is equivalent to maximizing $-\varphi$.

that for any $\mathbf{z}^* \in Z_+^*$, we have $\langle \Phi(\hat{\mathbf{u}}), \mathbf{z}^* \rangle \leq 0 = \langle \Phi(\hat{\mathbf{u}}), \boldsymbol{\lambda} \rangle$. Moreover, the regularity condition $\Phi(\mathbf{u}') < 0$ holds for some $\mathbf{u}' \in \Gamma$, i.e. \mathbf{u}' satisfies incentive compatibility (take $u'_t(\theta^t) = \varepsilon$). Therefore $\hat{\mathbf{u}}$ is a solution to the original problem (70) by Theorem 4. Note that we impose no boundedness assumption on the utility function.

To illustrate an application of this result, consider an example with logarithmic preferences. We argue that the Lagrange multiplier $\boldsymbol{\lambda}$ has the form $\lambda_t = \lambda_1 \beta^t$ for some λ_1 . Following the same steps as in Section 2.3, replace the period-0 incentive constraints with a sequence of one-shot constraints. Moreover, we consider an auxiliary planner's problem that has a recursive structure, by augmenting the set of constraints with the promise-keeping condition

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t(\theta^t) \right] = v_0. \quad (74)$$

The constraint set is then the set $\Gamma(v_0)$ defined in (16). We can re-write the problem (73)-(74) as

$$\max_{(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)} v_0 - \sum_{t=1}^{\infty} \lambda_t \{ \mathbb{E}_0 [C(u_t)] - e \}.$$

The solution to this problem coincides, given our guess $\lambda_t = \lambda_1 \beta^t$, with the solution to the problem

$$\max_{(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)} - \sum_{t=1}^{\infty} \beta^{t-1} \mathbb{E}_0 [C(u_t)],$$

which is, of course, the same problem as the one we analyzed in Section 2.3. Therefore, if we can show that the solution to that problem satisfies the feasibility constraint (69) for each period t , we found the solution to our new problem. We recover the solution to the original problem by maximizing the auxiliary problem over v_0 .

We now check that this is the case. Let (\mathbf{u}, \mathbf{v}) be the allocation generated by the policy functions to the Bellman equation (23) for some v_0 . The optimality conditions (29) imply

$$K'(v_0) = \mathbb{E}_0 [-C'(u_1)] = \mathbb{E}_0 [K'(v_1)] = \mathbb{E}_0 [\mathbb{E}_1 [-C'(u_2)]] = \mathbb{E}_0 [-C'(u_2)].$$

When preferences are logarithmic, $C = C' = \exp$, thus forward induction implies $\mathbb{E}_0 [C(u_1)] = \mathbb{E}_0 [C(u_t)]$ for all t . Since v_0 must satisfy $K(v_0) = \frac{1}{1-\beta} e$, this implies that $\mathbb{E}_0 [C(u_t)] = e$ for all t , establishing our result (and justifying our guess for λ_t).

When we set up the maximization problem (73) we assumed the existence of a summable sequence $\boldsymbol{\lambda}$ such that the feasibility constraints are satisfied with equality in all periods at the optimum. We subsequently showed how to explicitly construct such a sequence of multipliers in an example with logarithmic preferences. We now conclude this section by discussing sufficient conditions ensuring the existence of a summable sequence of Lagrange multipliers. Note that

without any further assumptions, the maximization problem

$$\max_{\mathbf{u} \in \Gamma, \Phi(\mathbf{u}) \leq \mathbf{0}} \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t \right]$$

satisfies all the conditions of Theorem 3, so that a Lagrangian exists. To show that it is a summable sequence we verify conditions of Theorem 5. It is the easiest to do in the case of bounded utility.¹⁷ In this case any sequence \mathbf{u} lies in l_∞ . Assumption 5.i holds following the same arguments as those used in the proof of Proposition 10 (see below). Since the constraint (69) holds for each t , we immediately have $\mathbb{E}_0 [C(u_t)] = \mathbb{E}_0 [C(x_t^T(\mathbf{u}, \mathbf{v}))]$ for T sufficiently large holding t fixed, and $\mathbb{E}_0 [C(x_t^T(\mathbf{u}, \mathbf{v}))] = \mathbb{E}_0 [C(v_t)]$ for t sufficiently large holding T fixed, which verifies Assumptions 5.ii.a and 5.ii.c. Assumption 5.ii.b holds by definition of Γ . Therefore Theorem 5 establishes that the Lagrange multipliers form a summable sequence.

Existence of a maximum

We finally show the existence of the maximum in problem (70). Note that we already showed the existence in Section 2.3.2 using the (finite-dimensional) Bellman formulation of the problem. Here we do so directly, using techniques that can be applied to other contexts where the previous approach is not readily available. It is not obvious a priori that the maximum in this problem exists. In finite dimensional spaces, the continuity of the objective function and the compactness of the constraint sets are easily obtained, implying directly the existence of a maximum. These properties are more difficult to obtain in infinite period economies. The next proposition guarantees that the infinite-horizon planner's problem is a well-defined maximization problem, i.e., there exist feasible $(\mathbf{u}^*, \mathbf{v}^*)$ for which the supremum is achieved. The reader interested mostly in the applications can skip this section.

Proposition 10. *The maximum in the problem (70) is attained for all $e > 0$.*

Proof. One of the easiest ways to show the existence of the minimum in the planner's problem is to truncate the economy at any finite period T , show the existence of the solution for this truncated economy, and finally show that the limit of this solution achieves the supremum of the original problem as $T \rightarrow \infty$. To show these we adapt the arguments of Ekeland and Scheinkman (1986).

We first restrict allocations in each period to compact sets as follows. Fix $e > 0$. For any $t \geq 1$, define $\bar{u}_t \in (\underline{u}, \bar{u})$ such that $\pi_t(\theta^t) C(\bar{u}_t) \geq e$ for all θ^t , that is,

$$\bar{u}_t = C^{-1} \left(\frac{e}{\min_{\theta \in \Theta} (\pi(\theta))^t} \right).$$

If $u_t(\theta^t) > \bar{u}_t$ for any history $\theta^t \in \Theta^t$, then $\mathbb{E}_0 [C(u_t)] > e$, and the allocation is not feasible.

¹⁷See Rustichini (1998) for existence arguments when utility is not bounded.

This gives us an upper bound $u_t(\theta^t) \leq \bar{u}_t$, and similarly, letting

$$\bar{v}_t = \mathbb{E}_t \left[\sum_{s=1}^{\infty} \beta^{s-1} \bar{u}_{t+s} \right] < \bar{v},$$

we have $v_t(\theta^t) \leq \bar{v}_t$. Moreover, $\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t(\theta^t) \right]$ goes to $-\infty$ if $u_s(\theta^s) \rightarrow -\infty$ for some (s, θ^s) , because of the upper bounds \bar{u}_t for all $t \geq 0$. But this allocation is dominated by $\tilde{u}_t(\theta^t) = C^{-1}(e)$ for all (t, θ^t) . Thus for each t we must have $u_t(\theta^t) \geq \underline{u}_t(\theta^t)$, which then gives us lower bounds $u_t(\theta^t) \geq \underline{u}_t \equiv \min_{\Theta^t} \underline{u}_t(\theta^t)$ and, similarly, $v_t(\theta^t) \geq \underline{v}_t$. Therefore we can impose the additional constraints

$$\begin{aligned} \underline{u}_t &\leq u_t(\theta^t) \leq \bar{u}_t, \\ \underline{v}_t &\leq v_t(\theta^t) \leq \bar{v}_t. \end{aligned}$$

Next, we truncate the economy to $T < \infty$ periods and allow the planner to provide incentives in the last period “for free”. That is, we define

$$V^T(e) = \sup_{\substack{u_t(\theta^t) \in [\underline{u}_t, \bar{u}_t] \\ v_t(\theta^t) \in [\underline{v}_t, \bar{v}_t]}} \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t u_t(\theta^t) \right]$$

subject to the promise keeping constraints

$$v_t(\theta^t) = \sum_{\theta \in \Theta} \pi(\theta) [\theta u_{t+1}(\theta^t, \theta) + \beta v_{t+1}(\theta^t, \theta)], \quad \forall t \leq T-1,$$

the incentive compatibility constraints

$$\theta u_t(\theta^{t-1}, \theta) + \beta v_t(\theta^t) \geq \theta u_t(\theta^{t-1}, \hat{\theta}) + \beta v_t(\theta^{t-1}, \hat{\theta}), \quad \forall t \leq T,$$

and the feasibility constraints

$$\mathbb{E}_0 [C(u_t)] \leq e, \quad \forall t \leq T.$$

Note that the last incentive constraint is:

$$\theta u_T(\theta^{T-1}, \theta) + \beta v_T(\theta^T) \geq \theta u_T(\theta^{T-1}, \hat{\theta}) + \beta v_T(\theta^{T-1}, \hat{\theta}) \quad \text{for all } \theta^{T-1}, \hat{\theta},$$

and the last two promise-keeping constraints are:

$$v_{T-1}(\theta^{T-1}) = \sum_{\theta \in \Theta} \pi(\theta) [\theta u_T(\theta^{T-1}, \theta) + \beta v_T(\theta^{T-1}, \theta)] \quad \text{and} \quad \underline{v}_T \leq v_T(\theta^T) \leq \bar{v}_T,$$

that is, the promise in period T has no resource cost. In the truncated problem we maximize a continuous function over a compact set, namely $\prod_{\substack{1 \leq t \leq T \\ \theta^t \in \Theta^t}} [\underline{u}_t, \bar{u}_t] \times [\underline{v}_t, \bar{v}_t]$, so a maximum

exists (which is, in fact, unique, since the objective is strictly convex). Call this maximum $(\mathbf{u}^T, \mathbf{v}^T) = \{u_t^T(\theta^t); v_t^T(\theta^t)\}_{t, \theta^t}$.

We now show that $\lim_{T \rightarrow \infty} (\mathbf{u}^T, \mathbf{v}^T)$ achieves the maximum of the original problem. By definition of a supremum, for any $\varepsilon > 0$ we can find an incentive compatible and feasible allocation $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$ for the original problem such that

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t \tilde{u}_t(\theta^t) \right] > V(e) - \varepsilon.$$

(Note that the r.h.s. is finite.) The truncation at T periods satisfies all the constraints of the truncated economy, so

$$V^T(e) \geq \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} \theta_t \tilde{u}_t(\theta^t) \right], \quad \forall T \geq 1.$$

Hence

$$\liminf_{T \rightarrow \infty} V^T(e) \geq \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t \tilde{u}_t(\theta^t) \right].$$

Since ε is arbitrary,

$$\liminf_{T \rightarrow \infty} V^T(e) \geq V(e).$$

To show the reverse inequality, fix $t \geq 1$. For all $T \geq t$, $(u_t^T(\theta^t), v_t^T(\theta^t)) \in [\underline{u}_t, \bar{u}_t] \times [\underline{v}_t, \bar{v}_t]$. Thus, the sequences $\{u_t^T(\theta^t)\}_{T \geq t}$ and $\{v_t^T(\theta^t)\}_{T \geq t}$ must have convergent subsequences as $T \rightarrow \infty$. We can then use a diagonal procedure to obtain an incentive compatible and feasible allocation $\{(u_t^\infty(\theta^t), v_t^\infty(\theta^t))\}_{t \geq 1, \theta^t \in \Theta^t}$, as follows. Arrange states as

$$\mathcal{R} = \{\theta_{(1)}, \dots, \theta_{(|\Theta|)}, (\theta_{(1)}, \theta_{(1)}), \dots, (\theta_{(1)}, \theta_{(|\Theta|)}), \dots\}.$$

Choose a subsequence of u^T, v^T so that the first element converges, i.e.,

$$\lim_{T \rightarrow \infty} (u_1^T(\theta_{(1)}), v_1^T(\theta_{(1)})) = (u_1^\infty(\theta_{(1)}), v_1^\infty(\theta_{(1)})).$$

From that subsequence choose another subsequence so that the second element converges, i.e.,

$$\lim_{T \rightarrow \infty} (u_1^T(\theta_{(2)}), v_1^T(\theta_{(2)})) = (u_1^\infty(\theta_{(2)}), v_1^\infty(\theta_{(2)})).$$

Repeat the procedure so that we get (u^∞, v^∞) , and call the final subsequence $\{T_n\}_{n \geq 0}$. Since for each $t \leq T, \theta^t \in \Theta^t$, $(u_t^T(\theta^t), v_t^T(\theta^t))$ lie in a closed set defined by the incentive constraints, $(u_t^\infty(\theta^t), v_t^\infty(\theta^t))$ also lie in the same set, i.e. they are incentive compatible. Since $C(u_t(\theta^t))$

is continuous on $\left[\underline{u}_t, \bar{u}_t\right]$, $C^T(\theta^t) \equiv C(u_t^T(\theta^t))$ (and $C^T(\theta^t) = 0$ for $t \geq T$) converges pointwise,

$$\lim_{T_n \rightarrow \infty} C^{T_n}(\theta^t) = C^\infty(\theta^t) \in \left[C(\underline{u}_t), C(\bar{u}_t)\right].$$

Now we can think of $\{\pi_1(\theta_{(1)}), \dots, \pi_1(\theta_{(|\Theta|)}) , \beta\pi_2(\theta_{(1)}, \theta_{(1)}), \dots\}$ as a measure on \mathcal{R} . For all $t \geq 1$, $\{\theta_t u_t^{T_n}(\theta^t)\}_{n \geq 1}$ is a sequence of positive measurable functions on that space that converges pointwise to $\theta_t u_t^\infty(\theta^t)$ as $n \rightarrow \infty$. By Fatou's lemma (Lemma 7.9 in Stokey et al. (1989)) $\theta_t u_t^\infty(\theta^t)$ is also measurable, and

$$\limsup_{n \rightarrow \infty} \sum_{t=1}^{T_n} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta_t) \theta_t u_t^{T_n}(\theta^t) \leq \sum_{t=1}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^{t-1} \pi_t(\theta_t) \theta_t u_t^\infty(\theta^t) \leq V(e),$$

where the last inequality follows from the fact that $\{u_t^\infty(\theta^t)\}$ satisfies the constraints of problem (9), but may not maximize the objective. Therefore, we obtain

$$\limsup_{n \rightarrow \infty} V^{T_n}(e) \leq V(e).$$

We therefore showed that $\lim_{n \rightarrow \infty} V^{T_n}(e)$ exists and

$$V(e) = \lim_{n \rightarrow \infty} V^{T_n}(e).$$

Moreover, we showed that a maximum of $V(e)$ is achieved by the limit of the sequence $(\mathbf{u}^{T_n}, \mathbf{v}^{T_n})$. This concludes the proof. \square

3.1.3 Application: Sustainability constraints

Suppose in addition to constraint (69) we further impose a constraint that social welfare in any period cannot drop below a threshold \underline{U} ,

$$\mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} u_{t+s} \right] \geq \underline{U}, \quad \forall t \geq 1. \quad (75)$$

Such constraints naturally arise in various settings with imperfect commitment, participation constraints, etc. We discuss an example of those in Section 4.4 in the context of an international finance model, where they capture the need to provide incentives for the agents to stick to the contract rather than defaulting and reverting to their outside option (in that case, the value of autarky).¹⁸ We add this constraint (75) to problem (70) and assume that the utility function is bounded.

¹⁸Such constraints would also appear in models of political economy in which a government is tempted to re-optimize, see e.g. Acemoglu et al. (2008), Sleet and Yeltekin (2008) and Farhi et al. (2012).

As before, we have for all t ,

$$\lim_{T \rightarrow \infty} \mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} x_t^T(\mathbf{u}, \mathbf{v}) \right] = \mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} u_{t+s} \right],$$

since the utility is bounded; and for t sufficiently large holding T fixed, we have

$$\mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} x_t^T(\mathbf{u}, \mathbf{v}) \right] = \mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} v_{t+s} \right],$$

which verifies Assumptions 5.ii.a and 5.ii.c for the constraints (75). The other parts of Assumption 5 are verified as before. As long as \underline{U} is not too high, we can find an interior point \mathbf{x}' that satisfies $\Phi(\mathbf{x}') < 0$. Theorem 5 thus establishes that there exists a non-negative summable sequence of Lagrange multipliers $\{\mu_t\}_{t=1}^{\infty}$, such that the solution to our problem is also a solution to

$$\max_{\mathbf{u} \in \Gamma} \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t \right] + \sum_{t=1}^{\infty} \mu_t \mathbb{E}_0 \left[\sum_{s=1}^{\infty} \beta^{s-1} \theta_{t+s} u_{t+s} \right] - \sum_{t=1}^{\infty} \tilde{\lambda}_t \mathbb{E}_0 [C(u_t)].$$

Since $\{\mu_t\}_{t=1}^{\infty}$ is summable, we can rewrite the equation above as

$$\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t + \sum_{t=1}^{\infty} \sum_{s=1}^{\infty} \beta^{s-1} \mu_t \theta_{t+s} u_{t+s} - \sum_{t=1}^{\infty} \tilde{\lambda}_t C(u_t) \right] = \mathbb{E}_0 \left[\sum_{t=1}^{\infty} \bar{\beta}_t \{ \theta_t u_t - \lambda_t C(u_t) \} \right], \quad (76)$$

where $\bar{\beta}_t = \beta^{t-1} + \mu_1 \beta^{t-2} + \dots + \mu_{t-2} \beta + \mu_{t-1}$, letting $\mu_0 = 0$, with $\sum_{t=1}^{\infty} \bar{\beta}_t < \infty$ and $\lambda_t = \tilde{\lambda}_t / \bar{\beta}_t$. This problem can be solved using our usual techniques. Augmenting the problem with a promise-keeping constraint, we can replace $\mathbf{u} \in \Gamma$ with $(\mathbf{u}, \mathbf{v}) \in \Gamma(v_0)$ and observe that the maximum can be written recursively, letting $\hat{\beta}_{t+1} = \bar{\beta}_{t+1} / \bar{\beta}_t$, as

$$k_t(v) = \max_{\{u(\theta), w(\theta)\}_{\theta \in \Theta}} \mathbb{E} \left[\theta u - \lambda_t C(u) + \hat{\beta}_{t+1} k_{t+1}(w) \right]$$

subject to (19), (20).

We can extend Lemma 1 directly to k_t , with the only exception that k_t is not strictly decreasing but rather inversely U -shaped. (We can easily show that $v \mapsto k_t(v)$ is concave, continuous, and satisfies $k_t(v) \xrightarrow{v \rightarrow \bar{v}} -\infty$ and $k_t(v) \xrightarrow{v \rightarrow \underline{v}} \underline{v}$.) Much of the analysis in Section 2.4 continues to hold but the condition (29) now becomes

$$k'_t(v) = \frac{\hat{\beta}_{t+1}}{\beta} \mathbb{E} [k'_{t+1}(w_v)]. \quad (77)$$

Observe that $\hat{\beta}_{t+1} \geq \beta$ with strict inequality if $\mu_t > 0$. Therefore the marginal cost k'_t is no longer a martingale if constraint (75) binds, which implies a form of mean-reversion. To see this, observe that since k is inversely U -shaped, k'_t is positive for low v and negative for

high v . Therefore equation (77) shows that the marginal cost decreases in expectation if v is low (because then $\mathbb{E} [k'_{t+1}(w_v)] = \frac{\beta}{\beta_{t+1}} k'_t(v) \leq k'_t(v)$) and increases if v is high (because then $\mathbb{E} [k'_{t+1}(w_v)] = \frac{\beta}{\beta_{t+1}} k'_t(v) \geq k'_t(v)$).

3.1.4 Using Lagrange multipliers instead of promised utilities

In our discussion so far we have used the following technique to solve dynamic incentive problems: we formed a Lagrangian using all the constraints except the incentive constraints, and then introduced the promised utilities in order to write the incentive constraints recursively. In principle, there is nothing special about the incentive constraints per se: we could extend the Lagrangian to those constraints also, eliminating the need to use promised utilities at all. Here we describe how this can be done, using a version of our benchmark partial equilibrium model of Section 2.3.2 as an example.

Consider the maximization problem

$$K(v_0) \equiv \max_{\mathbf{u}} \mathbb{E} \left[- \sum_{t=1}^{\infty} \beta^{t-1} C(u_t) \right] \quad (78)$$

subject to (10) $\forall t \geq 1$, and (14),

that we analyzed in Section 2.3.2. Since the objective function is strictly concave and the constraint set is convex, its solution $\hat{\mathbf{u}}$ is unique. Define $W : \mathbb{R} \rightarrow \mathbb{R}$ as

$$W(\alpha) \equiv \max_{\mathbf{u}} \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} (\alpha \theta_t u_t - C(u_t)) \right] \quad (79)$$

subject to (10) $\forall t \geq 1$.

If $\hat{\alpha}$ is the Lagrange multiplier on constraint (14), then $W(\hat{\alpha})$ is simply the Lagrangian associated with problem (78), whose unique maximum (the objective is strictly concave and the constraints are linear) is attained at $\hat{\mathbf{u}}$ by Theorem 4.¹⁹

We now show how applying Corollary 1 leads to a recursive characterization of this problem, using different techniques than those described in Section 2.3.2. We then discuss the strengths and weaknesses of these two alternative approaches. For simplicity we assume that $|\Theta| = 2$; the analysis extends straightforwardly to any number of shocks.

When $|\Theta| = 2$ there are two incentive constraints (10) in period 1 corresponding to shocks $\theta_{(1)}$ and $\theta_{(2)}$. Adapting the arguments of Proposition 5 shows that the constraint of type $\theta_{(2)}$ is slack. Let $\hat{\xi}(\theta_{(1)})$ be the Lagrange multiplier on the first-period incentive constraint of type $\theta_{(1)}$ in problem (79). By Corollary 1 (written for a maximization rather than minimization

¹⁹ $\hat{\alpha}$ can be found from a maxmin problem as in Corollary 1 by observing that by monotonicity we can replace the equality constraint (14) in problem (78) with the inequality constraint $\mathbb{E} [\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t (\theta^t)] \geq v_0$.

problem),²⁰ $\hat{\xi}(\theta_{(1)})$ and the solution to (79) are also the solution to

$$W(\alpha) \equiv \min_{\xi \geq 0} \max_{\mathbf{u}} \mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} (\alpha \theta_t u_t - C(u_t)) \right] \\ + \xi \left[\left\{ \theta_{(1)} u_1(\theta_{(1)}) + \mathbb{E} \left[\sum_{t=2}^{\infty} \beta^{t-1} \theta_t u_t \mid \theta_1 = \theta_{(1)} \right] \right\} \right. \\ \left. - \left\{ \theta_{(1)} u_1(\theta_{(2)}) + \mathbb{E} \left[\sum_{t=2}^{\infty} \beta^{t-1} \theta_t u_t \mid \theta_1 = \theta_{(2)} \right] \right\} \right]$$

subject to (10) $\forall t \geq 2$.

Re-arrange these terms and use the definition of W to obtain

$$W(\alpha) \equiv \min_{\xi \geq 0} \max_{u(\theta_{(1)}), u(\theta_{(2)})} \\ \pi(\theta_{(1)}) \left[\left(\alpha \theta_{(1)} + \frac{\xi \theta_{(1)}}{\pi(\theta_{(1)})} \right) u_1(\theta_{(1)}) - C(u_1(\theta_{(1)})) + \beta W \left(\alpha + \frac{\xi}{\pi(\theta_{(1)})} \right) \right] \\ + \pi(\theta_{(2)}) \left[\left(\alpha \theta_{(2)} - \frac{\xi \theta_{(1)}}{\pi(\theta_{(2)})} \right) u_1(\theta_{(2)}) - C(u_1(\theta_{(2)})) + \beta W \left(\alpha - \frac{\xi}{\pi(\theta_{(2)})} \right) \right]. \quad (80)$$

Problem (80) is an alternative way to characterize the solution to the maximization problem (78). The function W can be found using standard contraction mapping techniques (see Marcat and Marimon (2015) for proofs). The policy functions to this Bellman equation can then be used to generate the solution $\hat{\mathbf{u}}$ the same way we did in Section 2.3.2.

We conclude this section by comparing the two alternative recursive formulations (23) and (80). On the one hand, the max operator in (23) is simpler to handle than the min max operator in (80).²¹ This makes (23) easier to use in many simple applications. On the other hand, the function W is defined over an a priori known domain, \mathbb{R} , while the domain of K is endogenous. We could easily characterize the latter in the setup of Section 2.3.2, see footnote 7. In more general settings, however (with Markov shocks, additional constraints, etc.), characterizing the state space is more difficult and requires using the techniques of Abreu et al. (1990) (see Propo-

²⁰We can verify the sufficient conditions allowing us to apply Corollary 1 using the same steps as in Section 3.1.2.

²¹Many insights can be obtained from problem (80) without considering the min part. Since $\hat{\xi}(\theta_{(1)}) \geq 0$ (in fact, with strict inequality from the discussion in Proposition 5), problem (80) immediately shows that the weight $\alpha(\theta_{(1)}) \equiv \hat{\alpha} + \frac{\hat{\xi}(\theta_{(1)})}{\pi(\theta_{(1)})}$ increases for the agent who reports $\theta_{(1)}$, i.e. $\alpha(\theta_{(1)}) \geq \hat{\alpha}$, while the weight $\alpha(\theta_{(2)}) \equiv \hat{\alpha} - \frac{\hat{\xi}(\theta_{(1)})}{\pi(\theta_{(2)})}$ decreases, i.e. $\alpha(\theta_{(2)}) \leq \hat{\alpha}$. To see the implication of this fact, observe that the solution to (79), $\hat{\mathbf{u}}^\alpha$, has a property that $\mathbb{E} \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t u_t^\alpha \right]$ is increasing in α , so that higher weights correspond to higher lifetime utilities. Therefore we showed, without explicitly considering the min operator, that the expected lifetime utility starting from next period increases for an agent who reports $\theta_{(1)}$, and decreases for the agent who reports $\theta_{(2)}$. See Acemoglu et al. (2011) for another application of this technique.

sition 8), so that using the tools described in this section can be simpler. An in-depth discussion of this approach is outside the scope of our work and we refer the interested reader to the papers that describe it in more detail. The pionnering work that first developed this approach is Marcat and Marimon (2015). The more recent applications are Messner et al. (2012, 2014), Cole and Kubler (2012), Espino et al. (2013).

3.2 Mechanism Design Without Commitment

In our discussion so far we assumed that the principal, who provides insurance to agents, has perfect commitment: it implicitly promises a menu of allocations for infinitely many periods and never entertains the possibility of renegeing on those promises as time goes by. This assumption was critical in proving the Revelation Principle in Theorem 1. This assumption is not innocuous. For example, we saw in Proposition 6 that long-run immiseration is a common feature of the optimal insurance contracts. While such a contract is optimal ex-ante in period 0, it provides the worst possible allocation in the long run. Any benevolent principal would like to re-optimize at that point. Thus the assumption that the principal has perfect commitment is very strong in many applications.

In this section we discuss several approaches to analyze dynamic contracting problems in environments where the principal cannot commit. We start with the set up of Section 2.2 with two modifications. First, we assume that insurance in that economy is provided by a benevolent principal, which we call “the government”, that cannot commit ex-ante, in period 0 to its future actions. Second, in order to focus on information revelation and various generalizations of Theorem 1 we abstract from borrowing and lending and assume that the total consumption of all agents should be equal to the total endowment e in each period, as in Section 3.1.2.

Since the government cannot commit, we formally describe the environment as an infinitely repeated game between one large player (the government) and a continuum of atomistic agents.²² Each period of the game is divided into two stages. In the first stage agents report information about the realization of their idiosyncratic shock to the government, and in the second stage the government chooses allocations.²³ As in Section 2.2, it is helpful to start by describing communication between the agent and the principal using a general message space M .

Agents’ reporting strategies in period t are maps $\tilde{\sigma}_t : M^{t-1} \times \Theta^t \times H^t \rightarrow \Delta(M)$, and the government’s strategy is a map $\tilde{c}_t : M^t \times \check{H}^t \rightarrow \Delta(\mathbb{R}_+)$, where M^{t-1} and Θ^t are the histories of reports and the realizations of shocks for each agent, and H^t and \check{H}^t (described below) are the aggregate histories of the game. To avoid complicating our discussion with measure-theoretic apparatus, we assume that $\Delta(M)$ and $\Delta(\mathbb{R}_+)$ assign finite probabilities to finitely many elements. The assumption of a continuum of agents simplifies the analysis. By the law

²²See Chari and Kehoe (1990, 1993) for classic references.

²³Although this set up appears a bit stylized, many of its features emerge naturally in many more sophisticated models of political economy. For example, models in which policies are chosen via probabilistic voting à la Lindbeck and Weibull (1987) in each period often reduce to our set up with a benevolent government that cannot commit to its future actions. See Farhi et al. (2012), Scheuer and Wolitzky (2014) or Dovic et al. (2015) for applications.

of the large numbers, $\tilde{\sigma}$ generates the aggregate distribution of reports that the government receives from the agents, and \tilde{c} generates the distribution of consumption allocations provided by the government. Moreover, these distributions are not affected if an individual agent (who is of measure zero) deviates from his equilibrium strategy. Assuming that these aggregate distributions are observed, history H^t consists of the aggregate distributions generated by $\tilde{\sigma}$ and \tilde{c} up to period $t - 1$, while \check{H}^t consists of H^t and the distribution of aggregate reports generated by $\tilde{\sigma}_t$.

We describe how to characterize the perfect Bayesian equilibrium (PBE) of this game that delivers that highest ex-ante utility to agents. Well-known arguments (see Chari and Kehoe (1990) or a textbook treatment in Chapter 23 of Ljungqvist and Sargent (2012) for details) imply that to characterize such an equilibrium it is sufficient to focus only on a subset of the histories of the game. Namely, it is sufficient to characterize the reporting done by the agents and allocations provided by the government “on the equilibrium path”. If the government ever deviates from the equilibrium path distribution of allocations (up to a measure zero) then in subsequent histories agents and the government switch to the worst PBE. With a slight abuse of notation, we use $(\tilde{\sigma}, \tilde{c})$ to describe the behavior of agents and the government “on the equilibrium path”, i.e., mappings $\tilde{\sigma}_t : M^{t-1} \times \Theta^t \rightarrow \Delta(M)$ and $\tilde{c}_t : M^t \rightarrow \Delta(\mathbb{R}_+)$ which no longer have the aggregate histories in H^t, \check{H}^t as arguments. We use $\tilde{\sigma}_t(m | m^{t-1}, \theta^t)$ to denote the probability that an agent with history (m^{t-1}, θ^t) reports the message m in period t .

The pair $(\tilde{\sigma}, \tilde{c})$ must satisfy three constraints. First, in equilibrium each individual agent finds it optimal to stick to his reporting strategy $\tilde{\sigma}$ rather than deviating to any other reporting strategy $\tilde{\sigma}'$, so that the constraint (4) is satisfied. Note that to write this constraint we implicitly used the assumption of a continuum of agents. If an individual agent chooses $\tilde{\sigma}'$ rather than $\tilde{\sigma}$, the aggregate distribution of reports to the government remains unchanged and therefore the equilibrium allocations remain the same. Thus the same \tilde{c} appears on both sides of the incentive constraint. Second, any allocation that the government chooses must also be feasible, i.e. satisfy

$$\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}} [c_t] \leq e, \quad \forall t. \quad (81)$$

Third, the government should not find it optimal to deviate from its equilibrium play at any point of time. This constraint can be written as

$$\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}} \left[\sum_{s=t}^{\infty} \beta^{s-t} \theta_s U(c_s) \right] \geq \tilde{W}_t(\{\tilde{\sigma}_s\}_{s=1}^t) + \frac{\beta}{1-\beta} U(e), \quad \forall t. \quad (82)$$

The left hand side of this constraint is the government’s payoff from continuing to play its equilibrium strategy in period t . The right hand side consists of two parts: the value of the best one-time deviation \tilde{W}_t (to be defined below) followed by the value of the worst PBE starting from the next period. Since we assumed that shocks are i.i.d., it is easy to show that the worst PBE is such that agents reveal no information to the government and receive forever the same

per capita allocation e independently of the shock. The expected value of this allocation is $\frac{1}{1-\beta}U(e)$.

We now derive the value of deviation \tilde{W}_t . Let $\mu_t(m^t)$ denote the measure of agents who report history m^t . It is defined recursively as $\mu_{-1} = 1$ and

$$\mu_t(m^t) = \mu_{t-1}(m^{t-1}) \sum_{\theta^t \in \Theta^t} \pi(\theta^t) \tilde{\sigma}_t(m_t | m^{t-1}, \theta^t).$$

The measure μ_t depends on the entire history of reports up to period t , $\{\tilde{\sigma}_s\}_{s=1}^t$. We use $\mathbb{E}^{\tilde{\sigma}}[\theta | m^t]$ to denote the government's posterior expectation of an agent's type being θ , conditional on the history of reports m^t . The best deviation solves

$$\tilde{W}_t(\{\tilde{\sigma}_s\}_{s=1}^t) = \max_{\{c^w(m^t)\}_{m^t \in M^t}} \sum_{m^t \in M^t} \mu_t(m^t) [\mathbb{E}^{\tilde{\sigma}}[\theta | m^t] U(c^w(m^t))] \quad (83)$$

subject to the feasibility constraint

$$\sum_{m^t \in M^t} \mu_t(m^t) c^w(m^t) \leq e. \quad (84)$$

At this stage of our discussion it is useful to compare our set up to that with commitment in Section 2.2. Relative to the environment in that section, we have one additional constraint, (82). The important feature of this constraint is that posterior beliefs appear on both sides of this constraint. This complicates the analysis. In particular, note that the proof of Theorem 1 does not need to go through when constraint (82) is imposed. If we replace $\tilde{\Gamma} = (M, \tilde{c} \circ \tilde{\sigma})$ with a direct truthful mechanism $(\Theta, c \circ \sigma^{truth})$, we still obtain feasible and incentive compatible allocations for all agents as in the proof of Theorem 1. However we have $\tilde{W}_t(\{\sigma_s^{truth}\}_{s=1}^t) \geq \tilde{W}_t(\{\tilde{\sigma}_s\}_{s=1}^t)$, generally with a strict inequality, since a direct mechanism reveals more precise information to the government and increases its incentives to deviate. Since by construction we have $\mathbb{E}^{\tilde{c} \circ \tilde{\sigma}}[\sum_{s=t}^{\infty} \beta^{s-t} \theta_s U(c_s)] = \mathbb{E}^{c \circ \sigma^{truth}}[\sum_{s=t}^{\infty} \beta^{s-t} \theta_s U(c_s)]$, the direct truthtelling mechanism tightens the sustainability constraint of the government. Intuitively, this mechanism always reveals more information to the government than any other communication mode, increasing the gains for the government from ex-post reoptimization and lowering ex-ante welfare.

The discussion in the previous paragraph implies that it is generally *not* without loss of generality to restrict attention to mechanisms in which agents report their type directly to the government, as we did in Section 2.2, and that one needs to work with more general message spaces to characterize the optimal insurance in this setting. Here we outline how it can be done. Our discussion is based on Golosov and Iovino (2014); for more detailed discussion and proofs we refer the reader to that paper.²⁴

²⁴Formally, Golosov and Iovino (2014) study a slightly more general game that allows agents' and government's strategies to depend on the realization of payoff irrelevant variables. This convexifies the set of equilibrium payoffs

To find the optimal insurance without commitment, the best PBE solves

$$\max_{\tilde{c}, \tilde{\sigma}} \mathbb{E}^{\tilde{c}, \tilde{\sigma}} \left[\sum_{t=1}^{\infty} \beta^{t-1} \theta_t U(c_t) \right] \quad (85)$$

subject to (4), (81), (82). Under some technical conditions this problem can be significantly simplified. In particular with i.i.d. shocks the history of past realization of shocks is irrelevant and we can simply restrict attention to reporting strategies of the form $\tilde{\sigma}_t : M^{t-1} \times \Theta_t \rightarrow \Delta(M)$. Similarly, one can also show the analogue of Proposition 1 that stochastic allocations of consumption are suboptimal, so that we can assume without loss of generality that $\tilde{c}_t : M^t \rightarrow \mathbb{R}_+$. Finally without loss of generality we can restrict M to a finite set.²⁵

We now show how to write this problem recursively. As in Section 2, it is more convenient to change variables and optimize with respect to $u_t = U(\tilde{c}_t)$, and constraint (4) simplifies if we use a one shot deviation principle. Using the same arguments as those leading to equation (10), we can rewrite (4) as: for all $m^t \in M^t$,

$$v_t(m^t) = \sum_{(\theta, m) \in \Theta \times M} \pi(\theta) \tilde{\sigma}_{t+1}(m | m^t, \theta) [\theta u_{t+1}(m^t, m) + \beta v_{t+1}(m^t, m)], \quad (86)$$

and for all $(m^t, \theta) \in M^t \times \Theta$, for all $m \in M$ and some $m_\theta \in M$,

$$\theta u_{t+1}(m^t, m_\theta) + \beta v_{t+1}(m^t, m_\theta) \geq \theta u_{t+1}(m^t, m) + \beta v_{t+1}(m^t, m) \quad (87)$$

with for all $m \in M$,

$$\tilde{\sigma}_{t+1}(m | m^t, \theta) [\{\theta u_{t+1}(m^t, m_\theta) + \beta v_{t+1}(m^t, m_\theta)\} - \{\theta u_{t+1}(m^t, m) + \beta v_{t+1}(m^t, m)\}] = 0. \quad (88)$$

Equation (86) is simply a generalization of (12) to the setting in which agents reveal noisy information to the government. The next two equations form the incentive compatibility conditions. Equation (87) says that there must be some message m_θ that agent θ prefers to all others given the past history of messages m^t and shock realization θ . Equation (88) says that if an agent with current shock realization θ reports any message m other than m_θ with positive probability $\tilde{\sigma}_{t+1}(m | m^t, \theta)$, then he must be indifferent between reporting m and m_θ , since any report he sends must give him the highest utility. Equations (87) and (88) are a generalization of (13) and have recursive structure, with $v_t(m^t)$ playing the role of the state variable.

We now show how to write the problem of maximizing (85) subject to (81), (82), (86)-(88) recursively using the Lagrangian techniques introduced in Section 3.1.2. Let $\tilde{\lambda} = \{\tilde{\lambda}_t\}_{t=1}^{\infty}$ and $\tilde{\chi} = \{\tilde{\chi}_t\}_{t=1}^{\infty}$ be sequences of multipliers on the constraints (5) and (82), respectively. Assuming

and ensures that some technical conditions simplifying the analysis hold. To ease the exposition we simply assume that those conditions are satisfied.

²⁵Specifically, the cardinality of M can be taken to be $2|\Theta| - 1$.

that these sequences are summable (see Section 3.1.1), we can write the Lagrangian, using Abel's formula, as²⁶

$$\max_{\mathbf{u}, \tilde{\sigma}} \mathbb{E}^{\tilde{\sigma}} \sum_{t=1}^{\infty} \bar{\beta}_t \left[\theta_t u_t - \lambda_t C(u_t) - \chi_t \tilde{W}_t \right] \quad (89)$$

subject to (86)-(88), where $\bar{\beta}_t = \beta^{t-1} + \sum_{s=1}^t \beta^{t-s} \tilde{\chi}_s$, $\lambda_t = \tilde{\lambda}_t / \bar{\beta}_t$, $\chi_t = \tilde{\chi}_t / \bar{\beta}_t$. Note that this problem is very similar to the problem considered in Section 3.1.3, except that now we choose the optimal amount of information that is revealed to the government, $\tilde{\sigma}$, and the costs of information revelation are captured by the terms $\chi_t \tilde{W}_t$.

This problem still does *not* have a natural recursive form. Our recursive characterization in Section 3.1.2 relied on the fact that the linearity of the objective function allowed us to separately solve for the optimal allocations after any history θ^t (which without commitment becomes a history of reports m^t) without paying attention to the other histories. The key difficulty now is that \tilde{W}_t depends on the distribution of reports that are sent by all agents. We show here how to write a recursive formulation under the assumption that preferences are logarithmic. Golosov and Iovino (2014) use the techniques of Section 3.1.1 to obtain the same characterization for arbitrary concave utility functions.

When preferences are logarithmic, \tilde{W}_t is easy to simplify. The first order conditions of problem (89) give

$$\lambda_t^w C'(u_t^w(m^t)) = \mathbb{E}^{\tilde{\sigma}} [\theta | m^t] = \frac{\tilde{\sigma}_t(m_t | m^{t-1}, \theta)}{\sum_{\theta' \in \Theta} \pi(\theta') \tilde{\sigma}_t(m_t | m^{t-1}, \theta')},$$

where $u_t^w \equiv U(u_t^w)$ and λ_t^w is the Lagrange multiplier on constraint (84). With logarithmic preferences, $C' = C = \exp$. Using this fact together with (84) we can easily find that $\lambda_t^w = 1/e$. The key property is that this multiplier does not depend on particular values of $\{\tilde{\sigma}_t\}_{m^t, \theta^t}$, and therefore \tilde{W}_t can be written as

$$\tilde{W}_t(\{\tilde{\sigma}_s\}_{s=1}^t) = \sum_{m^{t-1}} \mu_{t-1}(m^{t-1}) W_t(\{\tilde{\sigma}_t(m | m^{t-1}, \theta)\}_{(m, \theta) \in M \times \Theta}),$$

where

$$W_t(\{\tilde{\sigma}_t(m | \cdot, \theta)\}_{(m, \theta) \in M \times \Theta}) = \max_{\{u^w(m)\}_{m \in M}} \sum_{(m, \theta) \in M \times \Theta} \pi(\theta) \tilde{\sigma}_t(m | \cdot, \theta) [\theta u^w(m) - \lambda_t^w C(u^w(m))].$$

If we substitute this equation into (89), we can easily write the problem recursively, letting

²⁶Since consumption allocations are deterministic, we write $\mathbb{E}^{\tilde{\sigma}^*}$ rather than $\mathbb{E}^{\mathbf{u} \circ \tilde{\sigma}^*}$.

$\hat{\beta}_{t+1} \equiv \bar{\beta}_{t+1}/\bar{\beta}_t$, as

$$k_t(v) = \max_{\{u(m), w(m), \sigma(m|\theta)\}_{(m,\theta) \in M \times \Theta}, \sigma(\cdot|\theta) \in \Delta(M)} \mathbb{E}^\sigma \left[\theta u - \lambda_t C(u) + \hat{\beta}_{t+1} k_{t+1}(w) \right] - \chi_t W_t \left(\{\sigma(m|\theta)\}_{m,\theta} \right) \quad (90)$$

subject to: for all θ ,

$$v = \sum_{(\theta, m) \in \Theta \times M} \pi(\theta) \sigma(m|\theta) [\theta u(m) + \beta w(m)],$$

for all m and some m_θ ,

$$\theta u(m_\theta) + \beta w(m_\theta) \geq \theta u(m) + \beta w(m),$$

and for all m ,

$$\sigma(m|\theta) [\{\theta u(m_\theta) + \beta w(m_\theta)\} - \{\theta u(m) + \beta w(m)\}] = 0.$$

Note that problem (90) is very similar to problem (18) in Section 2, with two modifications. First, the objective function has an additional term $-\chi_t W_t$ which captures the additional cost of information revelation off the equilibrium path. Second, agents generally play mixed strategies over the message space M rather than a pure reporting strategy over the set Θ .

Golosov and Iovino (2014) analyze this problem and show that the optimal amount of information that each agent reveals depends on the promised utility v . The key insight of their paper is that the agents who should reveal more information to the government are those for whom such revelation saves the most resources to the government *on the equilibrium path*. In particular, in the set up discussed above, the government loses relatively little resources if it delivers a low value of v without knowing the realization of θ , while information revelation by agents with higher v leads the government to save more resources. Golosov and Iovino (2014) show that for all v sufficiently small agents reveal no information to the government and play the same reporting strategy independently of the realization of their shock; on the other extreme, agents with sufficiently high promise v reveal full information to the government (at least as long as U exhibits decreasing absolute risk aversion) just as in Section 2.2. Golosov and Iovino (2014) further generalize their analysis by considering Markov shocks and obtaining a recursive characterization along the lines of Section 2.5.²⁷ They show that the government's participation constraints imply the existence of an endogenous lower bound below which agents' promised utility never falls, preventing the emergence of long-run immiseration which was obtained in Section 2.4.

²⁷The problem of information revelation with persistent shocks is related to the literature on the ratchet effect, see Freixas et al. (1985) and Laffont and Tirole (1988).

3.2.1 Optimal insurance with a mediator

In the game described in the previous section we assumed a particular communication protocol between the agents and the government: agents first report some information to the government, then the government takes some action. In settings where the government could commit, as in Section 2.5, restricting attention to such communication protocols was without loss of generality due to Theorem 1. As we saw, Theorem 1 fails when the government cannot commit. One may wonder if better outcomes can be attained if richer ways to communicate between agents and the government are available. The answer to this question turns out to be yes. Here we describe what the optimal communication devices are and how to characterize the optimal contracts in such settings.

Suppose agents and the government can communicate indirectly, using a third party called a “mediator”. The mediator can be a trusted third person with no stake in the outcome of the game, or simply a machine that takes reports from the agents and recommends the action to the government as a function of those reports using a pre-determined rule. Thus, the game is essentially the same as in the previous section, with the following modification. In each period, the agents first send reports $\tilde{\sigma}_t : M^{t-1} \times \Theta^t \rightarrow \Delta(M)$ to the mediator, then the mediator makes recommendations $\tilde{\sigma}_t^{med} : M^t \rightarrow \Delta(\mathbb{R}_+)$ to the government about which consumption allocation the government should pick. The government is then free to make any choice it wants.

Studying equilibria in this communication game using a mediator is interesting for the following reason. First, without loss of generality we can restrict attention to direct truth-telling strategies σ^{truth} for the agents, as defined in Section 2.2 (and hence we can assume that $\tilde{\sigma}_t^{med}$ is a mapping from Θ^t to $\Delta(\mathbb{R}_+)$). Moreover, with a mediator we can replicate the outcome of any perfect Bayesian equilibrium with any other communication device. Thus, we get a version of Theorem 1 for Bayesian Nash equilibria (see Myerson (1982, 1986) and Mas-Colell et al. (1995) (Sec. 23.D)). Therefore, the equilibrium with a mediator provides an upper bound on what can be achieved using any other communicating device.

We want to make two observations about games with a mediator. First, while without loss of generality we can assume that agents report their types truthfully to the mediator, the mediator generally randomizes to garble the information that the government receives – otherwise we government would be able to learn information perfectly about the agent’s type and this mechanism would be equivalent to the direct truth-telling mechanism discussed in the previous section. Second, while any PBE (using arbitrary communicating devices) can be implemented as a PBE in a game with a mediator, the converse is not true. Thus, whether the equilibrium with a mediator provides a reasonable description of the optimal insurance arrangement often depends on the context. For example, many negotiations of the resolutions of conflicts between countries already use mediators, so that this approach may be natural. On the other hand, in many political settings it seems often difficult to introduce an uninterested third party outside of the politician’s control, and the approach we described in the previous section may be preferable.

To see how this approach alters the incentive constraints, we consider the analogue of the recursive problem (90). The mediator generally needs to randomize between different allocations that it recommends to the government. For simplicity we assume that the mediator offers finitely many recommendations m_1, \dots, m_I to the government for each agent's report. The reporting strategies of the mediator are now simply $\tilde{\sigma}_t^{med}(m | m^{t-1}, \theta^t)$ for $m \in M \equiv \{m_1, \dots, m_I\}$. Assuming the one-shot deviation principle and that the dependence of period- t strategies on θ^{t-1} is redundant, we can write the agents' incentive constraint as (86) and

$$\begin{aligned} & \sum_m \tilde{\sigma}_{t+1}^{med}(m | m^t, \theta) [\theta u(m^{t+1}) + \beta v_{t+1}(m^{t+1})] \\ & \geq \sum_m \tilde{\sigma}_{t+1}^{med}(m | m^t, \theta') [\theta u(m^{t+1}) + \beta v_{t+1}(m^{t+1})], \quad \forall \theta'. \end{aligned} \tag{91}$$

Constraint (91) is weaker than constraints (87) and (88), so that more allocations are incentive compatible when a mediator is used. One way to understand the intuition is as follows. When an agent communicates using a mediator, he has no control over which recommendation the mediator makes to the government. Thus his incentive constraint (91) should hold in expectation, over all the recommendations that the mediator may make. When an agent communicates with the government directly, he would never send any message to the government which is dominated by another message. Therefore his incentive constraint (87), (88) should be satisfied for all the messages sent to the government.

It remains to describe how the government forms posterior beliefs based on the mediator's recommendations. The government's behavior is formally identical to that in (83) except that the value of the best deviation is now simply $\tilde{W}_t \left(\{\tilde{\sigma}_s^{med}\}_{s=1}^t \right)$ rather than $\tilde{W}_t \left(\{\tilde{\sigma}_s\}_{s=1}^t \right)$, so that the government uses the mediator's recommendations described by $\tilde{\sigma}^{med}$ rather than agents' reports $\tilde{\sigma}$ to form its posterior beliefs. Nevertheless the mathematical structure of the two problems is identical and we obtain a similar recursive representation as in (90), except that the incentive constraints are replaced by: for all θ' ,

$$\sum_m \sigma^{med}(m | \theta) [\theta u(m) + \beta w(m)] \geq \sum_m \sigma^{med}(m | \theta') [\theta u(m) + \beta w(m)].$$

A caution. We conclude this section with a caution about the usage of the term ‘‘Revelation Principle’’ in the literature. Some authors reserve this term only for principal-agent models and Theorem 1. Since Theorem 1 does not hold if the principal cannot commit, those authors often say that ‘‘the Revelation Principle fails without commitment’’ (see, e.g. Laffont and Tirole (1988) or Bester and Strausz (2001)). Other authors use this term more broadly as in Section 3.2.1, when the mechanism designer is thought not as a principal per se but rather as a mechanical randomizing device. In such settings truthful direct revelation holds both when the agents and the principal can and cannot commit (see, e.g. Myerson (1982, 1986) and Mas-Colell et al. (1995)), and one often hears that ‘‘the Revelation Principle always holds’’. While it may be

confusing, there is no disagreement about the mathematical facts, and one just needs to be careful about which version of the Revelation Principle one refers to.

3.3 Martingale Methods in Continuous Time

We now show how dynamic contracting problems can be conveniently analyzed in continuous time frameworks. We only briefly touch on this literature here. Sannikov (2008, 2014) analyzed a continuous-time dynamic moral hazard problem where observable output follows a Brownian motion whose drift is given by the agent's unobservable effort. Williams (2009, 2011) uses the stochastic Pontryagin principle based on the work of Bismut (1973, 1978) to analyze a continuous-time version of the Thomas and Worrall (1990) endowment shock model, and Cvitanic and Zhang (2012) apply the same techniques to moral hazard and adverse selection problems. Zhang (2009) considers a dynamic contracting problem with a finite Markov chain for the types. Miao and Zhang (2014) extend the Lagrangian techniques introduced in Section 3.1 to a model of limited commitment in continuous time.

Here we follow Sannikov (2008) who uses the dynamic programming principle in continuous time to analyze the moral hazard model described in the discrete time setting in Section 2.6. We start with a short primer on the mathematical techniques that allow us to solve this problem. A fully rigorous exposition of these techniques is beyond the scope of this paper, but we present the main tools that allow us to describe Sannikov (2008)'s model in a self-contained way.

3.3.1 Mathematical Background

For the basics of Brownian motion and stochastic processes, see e.g. Revuz and Yor (1999), Øksendal (2003), or Karatzas and Shreve (2012). For an exposition of the theory of stochastic optimal control, see e.g. Yong and Zhou (1999). In this section, after briefly introducing the basics of stochastic processes, we simply state the three fundamental theorems that will be important in the analysis of the continuous-time dynamic contracting model below, namely Itô's lemma, the Martingale Representation Theorem, and Girsanov's theorem. We also describe heuristically the dynamic programming principle in continuous time.

A stochastic process X is a family of random variables $\{X_t\}_{t \geq 0}$ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Define the filtration $\{\mathcal{F}_t\}_{t \geq 0}$ such that for all t , $\mathcal{F}_t = \sigma(\{X_s\}_{s \leq t})$ is the σ -algebra generated by X from time 0 to time t . We say that the process X is Markovian if $\mathbb{P}(X_t \in A | \mathcal{F}_s) = \mathbb{P}(X_t \in A | X_s)$ for all $t > s$ and all Borel sets A . The process X is a martingale (resp., submartingale) if $\mathbb{E}[|X_t|] < \infty$ for all $t \geq 0$ and $\mathbb{E}[X_t | \mathcal{F}_s] = X_s$ (resp. $\mathbb{E}[X_t | \mathcal{F}_s] \geq X_s$) for all $t > s$. A stochastic process $Z = \{Z_t\}_{t \geq 0}$ on $(\Omega, \mathcal{F}, \mathbb{P})$ is a *Brownian motion* if it satisfies (see Section 37 in Billingsley):

- (i) The process starts at 0: $\mathbb{P}(Z_0 = 0) = 1$;

(ii) The increments are independent: if $0 \leq t_0 \leq \dots \leq t_n$, then

$$\mathbb{P}(\mathcal{Z}_{t_k} - \mathcal{Z}_{t_{k-1}} \in A_k, \forall k \leq n) = \prod_{k=1}^n \mathbb{P}(\mathcal{Z}_{t_k} - \mathcal{Z}_{t_{k-1}} \in A_k);$$

(iii) For $0 \leq s < t$ the increment $\mathcal{Z}_t - \mathcal{Z}_s$ is normally distributed with mean 0 and variance $t - s$:

$$\mathbb{P}(\mathcal{Z}_t - \mathcal{Z}_s \in A) = \frac{1}{\sqrt{2\pi(t-s)}} \int_A e^{-x^2/2(t-s)} dx;$$

(iv) The sample paths are continuous: for each $\omega \in \Omega$, the function $t \mapsto \mathcal{Z}_t(\omega)$ is continuous. In particular, a Brownian motion is a martingale.

We now define the concept of quadratic variation of a martingale. Consider a martingale M that has continuous sample paths. Consider a partition $\Pi_t = \{t_0, \dots, t_n\}$ of the interval $[0, t]$ with $0 = t_0 < t_1 < \dots < t_n = t$, and denote its mesh by $\|\Pi_t\| \equiv \max_{1 \leq k \leq n} (t_k - t_{k-1})$. Denoting by $\mathbb{P} \lim$ the limit of a process in the sense of the convergence in probability, we can show that

$$\mathbb{P} \lim_{\|\Pi_t\| \rightarrow 0} \sum_{k=1}^n (M_{t_k} - M_{t_{k-1}})^2 = \langle M \rangle_t,$$

where $\langle M \rangle$ is an adapted process with continuous and non-decreasing sample paths, called the *quadratic variation* of the martingale M . In particular, in the case where M is a Brownian motion, $\langle M \rangle$ is the deterministic process $\langle M \rangle_t = t$, and the convergence holds almost surely. Since $\langle M \rangle$ has non-decreasing sample paths ω , we can define the (path-by-path) Lebesgue-Stieljes integral $\int_0^t X_s(\omega) d\langle M \rangle_s(\omega)$ for each ω of a stochastic process X on an interval $[0, T]$ with $T < \infty$ (in the case where M is a Brownian motion, $d\langle M \rangle_s = ds$ is simply the Lebesgue measure).

We refer to Revuz and Yor (1999) for the rigorous construction of the stochastic integral $\int_0^t X_s dM_s$ of a process X with respect to a martingale M that has continuous sample paths (e.g., a Brownian motion). For such a martingale M , let $L^2(M)$ denote the (Hilbert) space of processes X such that for all $t \geq 0$ the map $(\omega, s) \mapsto X_s(\omega)$ defined on $\Omega \times [0, t]$ is measurable with respect to $\mathcal{F}_t \otimes \mathcal{B}([0, t])$, and $\mathbb{E} \left[\int_0^T X_s^2 d\langle M \rangle_s \right] < \infty$. The construction of the stochastic integral involves several steps. Suppose first that X is a “simple” process, in the sense that there exists a partition $0 = t_0 < t_1 < \dots < t_n = T$ of $[0, T]$ such that $X_s = \xi_j$ for all $s \in (t_j, t_{j+1}]$, where ξ_j is a bounded \mathcal{F}_{t_j} -measurable random variable. That is, X can be written as

$$X_s(\omega) = \sum_{j=0}^{n-1} \xi_j(\omega) \mathbb{I}_{(t_j, t_{j+1}]}(s).$$

We can then define, for $t_k < t \leq t_{k+1}$,

$$I_t(X) = \int_0^t X_s dM_s \equiv \sum_{j=0}^{k-1} \xi_j (M_{t_{j+1}} - M_{t_j}) + \xi_k (M_t - M_{t_k}).$$

The integral $I(X)$ is then a square integrable continuous martingale with quadratic variation given by $\langle I(X) \rangle_t = \int_0^t X_s^2 d\langle M \rangle_s$. Next, any process $X \in L^2(M)$ can be approximated by a sequence of simple processes $\{X^n\}_{n \geq 0}$ in the sense that $\mathbb{E} \left[\int_0^T (X_s^n - X_s)^2 d\langle M \rangle_s \right] \rightarrow 0$. We can then show that the sequence of integrals $I(X^n)$ is a Cauchy sequence in the complete space $L^2(M)$. Its limit defines the stochastic integral. It satisfies $\mathbb{E}[I_t(X)] = 0$ and is a martingale.

We now state the three main theorems which we use in our analysis. The first, *Itô's lemma*, is an extension of the chain rule from standard calculus:

Theorem 6. (*Itô's lemma.*) *Let f be a deterministic C^2 function and M a squared integrable martingale. We have:*

$$f(M_t) = f(M_0) + \int_0^t f'(M_s) dM_s + \frac{1}{2} \int_0^t f''(M_s) d\langle M \rangle_s.$$

The second important result is the *Martingale Representation Theorem*. If M is a martingale, define the *exponential martingale*

$$\mathcal{E}(M)_t = \exp \left(M_t - \frac{1}{2} \langle M \rangle_t \right). \quad (92)$$

We can then show that $\mathcal{E}(M)_t$ is a supermartingale, and it is a martingale if in addition $\mathbb{E} \left[\exp \left(\frac{1}{2} \langle M \rangle_T \right) \right] < \infty$. In particular, if M_t is defined as a stochastic integral with respect to a Brownian motion \mathcal{Z} , i.e. $M_t = \int_0^t \mu_s d\mathcal{Z}_s$ with $\int_0^t \mu_s^2 ds < \infty$ a.s., then

$$\mathcal{E}(M)_t = \exp \left(\int_0^t \mu_s d\mathcal{Z}_s - \frac{1}{2} \int_0^t \mu_s^2 ds \right) \quad (93)$$

is a martingale if $\mathbb{E} \left[\exp \left(\frac{1}{2} \int_0^T \mu_s^2 ds \right) \right] < \infty$.

Theorem 7. (*Martingale Representation theorem.*) *Let \mathcal{Z} be a given Brownian motion. Every square integrable continuous martingale M adapted to the filtration $\mathcal{F}^{\mathcal{Z}}$ generated by \mathcal{Z} admits a unique representation*

$$M_t = M_0 + \int_0^t \beta_s d\mathcal{Z}_s$$

for some process β adapted to $\mathcal{F}^{\mathcal{Z}}$ that satisfies $\mathbb{E} \left[\exp \left(\int_0^T \beta_s^2 ds \right) \right] < \infty$.

Finally, the third important result that we will use is *Girsanov's theorem*, which concerns the changes of measures.

Theorem 8. (Girsanov theorem.) Let \mathcal{Z} be a Brownian motion and μ be an adapted process with $\int_0^t \mu_s^2 ds < \infty$ a.s. Let $\mathcal{E}(M)_t$ be defined by (93). If $\mathbb{E}[\mathcal{E}(M)_T] = 1$ (that is, if $\mathcal{E}(M)$ is a martingale) then, under

$$\tilde{\mathbb{P}}(d\omega) = \mathcal{E}(M)_T(\omega) \times \mathbb{P}(d\omega),$$

the process

$$\tilde{\mathcal{Z}} = \mathcal{Z} - \int_0^t \mu_s ds$$

is a Brownian motion.

Finally we describe heuristically the dynamic programming principle in continuous time. We skip many of the technicalities and refer to Yong and Zhou (1999) for a rigorous exposition. Consider a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, on which a Brownian motion \mathcal{Z} is defined, and let $T \in (0, \infty)$ and $A \subset \mathbb{R}$ be a given Borel set. The state of a system at time t is described by a stochastic process $X_t \in \mathbb{R}$ that evolves according to

$$X_{t'} = x + \int_t^{t'} b(s, X_s, u_s) ds + \int_t^{t'} \sigma(s, X_s, u_s) d\mathcal{Z}_s, \quad 0 \leq t \leq t' \leq T, \quad (94)$$

where $\mathbf{u} : [0, T] \times \Omega \rightarrow A$ is the *control process*, and $b, \sigma : [0, T] \times \mathbb{R} \times A \rightarrow \mathbb{R}$. The goal is to choose \mathbf{u} to maximize the functional

$$J(\mathbf{u}) \equiv \mathbb{E} \left[\int_t^T f(s, X_s, u_s) ds + g(X_T) \right], \quad (95)$$

where $f : [0, T] \times \mathbb{R} \times A \rightarrow \mathbb{R}$ and $g : \mathbb{R} \rightarrow \mathbb{R}$.²⁸ We assume that the functions b, σ, f, g , satisfy suitable conditions ensuring that there exists a unique solution X to (94) for any t, x, \mathbf{u} and that the functional $J(\mathbf{u})$ in (95) is well defined (see Definition 1.6.15. and Conditions (S1)' and (S2)' in Yong and Zhou (1999)). The control process \mathbf{u} is *admissible* if: (i) u_t is $\{\mathcal{F}_t\}_{t \geq 0}$ -adapted; (ii) X is the unique solution of equation (94); (iii) the functions $s \mapsto f(s, X_s, u_s)$ and $s \mapsto g(X_T)$ are in $L^1_{\mathcal{F}}([0, T], \mathbb{R})$ and $L^1_{\mathcal{F}_T}(\Omega, \mathbb{R})$, respectively. The value function of the stochastic control problem that we consider is

$$V(t, x) = \sup_{\mathbf{u}} J(\mathbf{u}),$$

where the supremum is over all admissible controls \mathbf{u} .²⁹

²⁸Here the control problem ends at a fixed duration T . In our analysis of the moral hazard problem we will deal instead with random horizons T optimally chosen by the principal (“retirement”), that is, where T is the stopping time $T \equiv \inf \{t \geq 0 : x_t \notin \mathcal{O}\}$ for some open set $\mathcal{O} \subset \mathbb{R}$. The dynamic programming principle can be extended to this case, see e.g. Section 2.7. in Yong and Zhou (1999) and Chapter 4 in Øksendal and Sulem (2005).

²⁹Rigorously, it is often natural and necessary to consider a *weak formulation* of the problem, in which the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ and the Brownian motion \mathcal{Z} are not fixed, but parts of the control (see Sections 2.4.2. and 4.3.1. in Yong and Zhou (1999)). This is because the objective of the stochastic control problem is to minimize the expectation of a random variable that depends only on the distribution of the processes involved. We ignore this distinction in the sequel.

Theorem 9. (Dynamic Programming Principle.) For any stopping time τ with values in $[0, T]$, the value function $V(t, x)$ is equal to

$$V(t, x) = \sup_{\mathbf{u}} \mathbb{E} \left[\int_t^\tau f(s, X_s, u_s) ds + V(\tau, X_\tau) \right].$$

Moreover, for all admissible controls \mathbf{u} ,

$$M_{t'} \equiv \int_t^{t'} f(s, X_s, u_s) ds + V(t', X_{t'})$$

is a supermartingale (i.e., $-M_{t'}$ is a submartingale), and it is a martingale if and only if \mathbf{u} is optimal. Suppose the value function $V \in \mathcal{C}^{1,2}([0, T] \times \mathbb{R})$. Then V is a solution to the following second-order Hamilton-Jacobi-Bellman partial differential equation:

$$\begin{cases} -\frac{\partial V}{\partial t} + \sup_{u \in A} \left[f(t, x, u) + b(t, x, u) \frac{\partial V}{\partial x} + \frac{1}{2} \sigma^2(t, x, u) \frac{\partial^2 V}{\partial x^2} \right] = 0, & \forall (t, x) \in [0, T] \times \mathbb{R}, \\ V(T, x) = g(x), & \forall x \in \mathbb{R}. \end{cases}$$

Note that the last statement assumes smoothness conditions about the value function V , which is endogenous.³⁰

3.3.2 Continuous-Time Model: Transitory effects of effort on output

We now analyze the moral hazard problem in a continuous time framework (see Section 2.6 for the discrete time version of the model), following Sannikov (2008)'s exposition. Our aim is to derive and explain the main results with the minimum of technicalities. Therefore we omit many technical details and refer to Sannikov's work for the fully rigorous proofs.

We analyze a model where the agent's current effort affects only current output. The agent derives utility $U(c_t) - h(\theta_t)$ from consumption $c_t \geq 0$ and effort $\theta_t \in [0, \bar{\theta}]$ at time t , where U is twice continuously differentiable, increasing and concave with $U(0) = 0$ and $\lim_{c \rightarrow \infty} U'(c) = 0$, and h is differentiable, increasing and convex with $h(0) = 0$ and $h'(0) > 0$.

Fix a reference probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a standard Brownian motion \mathcal{Z} under \mathbb{P} . If the agent works according to the effort process $\boldsymbol{\theta} = \{\theta_t\}_{t \in [0, \infty)}$ with $0 \leq \theta_t \leq \bar{\theta}$ for all t , he generates an output y_t given by $y_t = \int_0^t \theta_s ds + \sigma \mathcal{Z}_t$, i.e.

$$dy_t = \theta_t dt + \sigma d\mathcal{Z}_t,$$

where $\sigma > 0$ is a constant. The principal observes y_t , but not θ_t nor \mathcal{Z}_t , and compensates the agent with a consumption process $\mathbf{c} = \{c_t\}_{t \geq 0}$ with $c_t \geq 0$ for all t . Denoting by \mathcal{F}_t^y the filtration generated by y_t , we impose that the process c_t is \mathcal{F}_t^y -adapted, i.e., the agent's compensation c_t

³⁰There exist other notions of solutions to stochastic differential equations, called viscosity solutions, which avoid making such assumptions, see e.g. Section 4.5. in Yong and Zhou (1999).

is conditional on past output $\{y_s\}_{s \leq t}$.

Rather than fixing the underlying Brownian motion \mathcal{Z} and solving for the agent's effort choice θ as a function of \mathcal{Z} , we can instead view the agent as choosing a probability measure \mathbb{P}^θ on the output space.³¹ That is, for each effort process θ we can define a process $\mathcal{Z}_t^\theta = \sigma^{-1} \left(y_t - \int_0^t \theta_s ds \right)$. By Girsanov's theorem, \mathcal{Z}_t^θ is a Brownian motion under the measure \mathbb{P}^θ , where

$$\mathbb{P}^\theta(d\omega) = \mathcal{E}(\mathcal{Z})_t \mathbb{P}(d\omega) = e^{\int_0^t \theta_s d\mathcal{Z}_s - \frac{1}{2} \int_0^t \theta_s^2 ds} \mathbb{P}(d\omega).$$

A change of measure from \mathbb{P}^θ to $\mathbb{P}^{\hat{\theta}}$ on the space of output paths corresponds to a change in the drift of the output process from θ to $\hat{\theta}$.

Planner's problem. If he receives consumption $\mathbf{c} = \{c_t\}_{t \geq 0}$ and provides effort $\theta = \{\theta_t\}_{t \geq 0}$, the agent gets the expected utility

$$U(\mathbf{c}, \theta) = \mathbb{E}^\theta \left[\int_0^\infty e^{-rt} (U(c_t) - h(\theta_t)) dt \right], \quad (96)$$

where \mathbb{E}^θ denotes the expectation under the probability measure \mathbb{P}^θ induced by the strategy θ , as defined above. The superscript θ over the expectation \mathbb{E}^θ highlights that the agent's strategy affects the probability distribution over paths of output, and thus over compensation realizations. Thus, the utility depends on the agent's effort directly, as it enters the cost of effort $h(\theta_t)$, and indirectly through its effect on the probability distribution over the paths of y_t .

The principal gets expected profit

$$\mathbb{E}^\theta \left[\int_0^\infty e^{-rt} (dy_t - c_t dt) \right] = \mathbb{E}^\theta \left[\int_0^\infty e^{-rt} (\theta_t - c_t) dt \right]. \quad (97)$$

A contract (\mathbf{c}, θ) is *incentive compatible* if the agent finds it optimal to exert the contractual effort θ_t at every t , i.e., if $\{\theta_t\}_{t \geq 0}$ maximizes his expected utility $U(\mathbf{c}, \theta)$ given $\{c_t\}_{t \geq 0}$:

$$\mathbb{E}^\theta \left[\int_0^\infty e^{-rt} (U(c_t) - h(\theta_t)) dt \right] \geq \mathbb{E}^{\hat{\theta}} \left[\int_0^\infty e^{-rt} (U(c_t) - h(\hat{\theta}_t)) dt \right], \quad \forall \hat{\theta}. \quad (98)$$

The contract must deliver initial promised utility \hat{v}_0 , i.e.

$$\mathbb{E}^\theta \left[\int_0^\infty e^{-rt} (U(c_t) - h(\theta_t)) dt \right] \geq \hat{v}_0. \quad (99)$$

The principal's problem consists of choosing the contract (\mathbf{c}, θ) that maximizes his expected profit (97) among all the contracts that satisfy the incentive compatibility (98) and promise

³¹Similarly, in the standard static moral hazard problem, we can view the agent as choosing the probability distribution $\mathbb{P}(y|\theta)$ over output values y generated by his effort θ .

keeping (99) constraints, that is,

$$\begin{aligned} & \max_{\mathbf{c}, \boldsymbol{\theta}} \mathbb{E}^{\boldsymbol{\theta}} \left[\int_0^{\infty} e^{-rt} (\theta_t - c_t) dt \right] \\ & \text{subject to (98), (99).} \end{aligned}$$

The principal can commit to the contract he offers.

Reducing the planner's problem to an optimal stochastic control problem. The planner's problem can be solved by reducing it to an optimal stochastic control problem. As in the discrete time framework, we use the agent's continuation utility v_t (defined formally below) as state variable. The key simplification of the planner's problem comes again from the (continuous-time equivalent of the) one-shot deviation principle, which substantially reduces the set of incentive constraints: the agent's incentive constraints hold for all alternative strategies $\hat{\boldsymbol{\theta}} = \{\hat{\theta}_t\}_{t \geq 0}$ if they hold just for strategies that differ from $\boldsymbol{\theta} = \{\theta_t\}_{t \geq 0}$ for an instant. The Martingale Representation Theorem then allows us to express the instantaneous incentive constraints in terms of v_t .

Fix an arbitrary consumption process $\mathbf{c} = \{c_t\}_{t \geq 0}$ and an effort strategy $\boldsymbol{\theta} = \{\theta_t\}_{t \geq 0}$ (not necessarily optimal for the agent given \mathbf{c}). The agent's *continuation value* $v_t(\mathbf{c}, \boldsymbol{\theta})$, defined as his expected future payoff from $(\mathbf{c}, \boldsymbol{\theta})$ after time t (i.e., after a given history of output $\{y_s\}_{s \leq t}$), is given by

$$v_t(\mathbf{c}, \boldsymbol{\theta}) = \mathbb{E}^{\boldsymbol{\theta}} \left[\int_t^{\infty} e^{-r(s-t)} (U(c_s) - h(\theta_s)) ds \mid \mathcal{F}_t \right]. \quad (100)$$

Throughout this section, for a given time t and contract $(\mathbf{c}, \boldsymbol{\theta})$, we also define the agent's *total* expected payoff from the contract $(\mathbf{c}, \boldsymbol{\theta})$ given the information at time t as:³²

$$V_t^{\mathbf{c}, \boldsymbol{\theta}} = \mathbb{E}^{\boldsymbol{\theta}} \left[\int_0^{\infty} e^{-rs} (U(c_s) - h(\theta_s)) ds \mid \mathcal{F}_t \right] = \int_0^t e^{-rs} (U(c_s) - h(\theta_s)) ds + e^{-rt} v_t(\mathbf{c}, \boldsymbol{\theta}). \quad (101)$$

We first derive the law of motion of $v_t(\mathbf{c}, \boldsymbol{\theta})$ by applying the Martingale Representation Theorem.

Proposition 11. *Fix a contract $(\mathbf{c}, \boldsymbol{\theta})$ with finite expected payoff to the agent. An adapted process v_t is the continuation value process (as defined in (100)) associated with the contract $(\mathbf{c}, \boldsymbol{\theta})$ if and only if there exists an \mathcal{F}_t -adapted process $\boldsymbol{\beta} = \{\beta_t\}_{t \geq 0}$ with $\mathbb{E} \left[\int_0^t \beta_s^2 ds \right] < \infty$ for all t such that, for all $t \geq 0$,*

$$dv_t = (rv_t - U(c_t) + h(\theta_t)) dt + \beta_t (dy_t - \theta_t dt) \quad (102)$$

and the transversality condition $\lim_{t \rightarrow \infty} \mathbb{E}^{\boldsymbol{\theta}} [e^{-rt} v_{t_0+t} \mid \mathcal{F}_{t_0}] = 0$ holds almost everywhere.

Proof. Fix a contract $(\mathbf{c}, \boldsymbol{\theta})$. The process $V_t^{\mathbf{c}, \boldsymbol{\theta}}$ defined in (101) is a martingale under the prob-

³²See Theorem 9 above.

ability measure \mathbb{P}^θ . Hence by the Martingale Representation Theorem there exists an adapted process β_t such that

$$V_t^{\mathbf{c},\theta} = V_0^{\mathbf{c},\theta} + \int_0^t e^{-rs} \beta_s \sigma dZ_s^\theta, \quad 0 \leq t < \infty,$$

where $Z_t^\theta = \sigma^{-1} \left(y_t - \int_0^t \theta_s ds \right)$ is a Brownian motion under \mathbb{P}^θ . Differentiating both expressions for $V_t^{\mathbf{c},\theta}$ with respect to t and equating them implies that $v_t(\mathbf{c}, \theta)$ satisfies (102). The transversality condition (for simplicity with $t_0 = 0$) follows from

$$\lim_{t \rightarrow \infty} \mathbb{E}^\theta \left[\int_0^t e^{-rs} (U(c_s) - h(\theta_s)) ds \right] = \mathbb{E}^\theta \left[\int_0^\infty e^{-rs} (U(c_s) - h(\theta_s)) ds \right],$$

by the Dominated Convergence Theorem using that θ_s , and thus $\int_0^t e^{-rs} (U(c_s) - h(\theta_s)) ds$, is bounded. A similar argument shows that $\lim_{t \rightarrow \infty} \mathbb{E}^\theta [e^{-rt} v_{t_0+t} | \mathcal{F}_{t_0}] = 0$ for all times $t_0 \geq 0$.

Conversely, suppose that v_t is a process that satisfies (102) (for some starting value v_0 and some volatility process β_t) and the transversality condition. Define V_t as

$$V_t = \int_0^t e^{-rs} (U(c_s) - h(\theta_s)) ds + e^{-rt} v_t.$$

Differentiating V_t implies that it is a martingale when the agent is following the effort strategy θ , i.e., under the probability measure \mathbb{P}^θ . Therefore

$$v_0 = V_0 = \mathbb{E}^\theta [V_t | \mathcal{F}_0] = \mathbb{E}^\theta \left[\int_0^t e^{-rs} (U(c_s) - h(\theta_s)) ds | \mathcal{F}_0 \right] + \mathbb{E}^\theta [e^{-rt} v_t | \mathcal{F}_0].$$

Since the transversality condition is satisfied (for $t_0 = 0$), taking limits as $t \rightarrow \infty$ in the previous equation implies that $v_0 = v_0(\mathbf{c}, \theta)$. A similar argument shows that v_t is the continuation value process $v_t(\mathbf{c}, \theta)$ defined by (100) at any time $t \geq 0$. \square

The law of motion (102) of the continuation utility has the following interpretation. Since $dy_t - \theta_t dt = \sigma dZ_t^\theta$ has mean zero when the agent takes the recommended effort level θ , $[rv_t(\mathbf{c}, \theta) - (U(c_t) - h(\theta_t))]$ is the drift of the agent's continuation value. The value that the principal owes to the agent (future expected payoff), $v_t(\mathbf{c}, \theta)$, grows at the rate of interest r , and falls due to the flow of repayments $(U(c_t) - h(\theta_t))$. The transversality condition has to hold if the debt is eventually repaid. Since the agent's compensation and recommended effort are determined by output y_t , his continuation payoff $v_t(\mathbf{c}, \theta)$ is also determined by output, and the process $r\beta_t$ then expresses the sensitivity of the agent's continuation value to output at a given time, which will be the key to affect the agent's incentives.

The previous lemma is useful because it allows us to simplify the set of incentive constraints with a version of the one-shot deviation principle (Proposition 12 below), which shows that the agent's incentive constraints hold for all alternative strategies $\hat{\theta}$ if they hold for all strategies which differ from θ for an infinitesimally small amount of time. Heuristically, suppose that

the agent has conformed to the contract (c_s, θ_s) for $s \leq t$ and cheats by performing effort $\hat{\theta}$ in the interval $[t, t + dt]$ and reverting to $\{\theta_s\}$ for $s \geq t + dt$. His immediate consumption c_t is unaffected, his cost on $[t, t + dt]$ is $h(\hat{\theta}) dt$, and his expected benefit on $[0, \infty)$, i.e. the expected impact of effort on his continuation value, is $\mathbb{E}^{\hat{\theta}}[\beta_t dy_t] = \beta_t \hat{\theta} dt$. Hence for the contract to be incentive compatible we must have

$$\beta_t \theta_t - h(\theta_t) = \max_{\hat{\theta} \geq 0} \left\{ -h(\hat{\theta}) + \beta_t \hat{\theta} \right\},$$

almost everywhere.³³ This argument can be made rigorous, and in addition the condition is not only necessary but also sufficient: if this one-shot condition holds at each instant t , then any dynamic deviation strategy $\hat{\theta} = \{\hat{\theta}_s\}_{s \geq 0}$ is suboptimal.

Proposition 12. *Let (\mathbf{c}, θ) be a contract with agent's continuation value $v_t(\mathbf{c}, \theta)$ and let β_t be the process from Lemma 11 that represents $v_t(\mathbf{c}, \theta)$. Then (\mathbf{c}, θ) is incentive compatible if and only if $\forall \hat{\theta} \in [0, \bar{\theta}], \forall t \geq 0$,*

$$\theta_t \in \arg \max_{\hat{\theta} \geq 0} \left\{ \beta_t \hat{\theta} - h(\hat{\theta}) \right\}, \quad a.e. \quad (103)$$

Proof. Suppose that (103) is satisfied. Suppose that an agent follows the alternative effort process $\hat{\theta} = \{\hat{\theta}_s\}_{s \geq 0}$ until time t and reverts back to θ thereafter; denote by $\hat{\theta}^t$ this strategy.

The time- t expectation of his total payoff is given by $V_t^{\mathbf{c}, \hat{\theta}^t}$ defined in (101),

$$V_t^{\mathbf{c}, \hat{\theta}^t} = \int_0^t e^{-rs} \left(U(c_s) - h(\hat{\theta}_s) \right) ds + e^{-rt} v_t(\mathbf{c}, \theta).$$

Differentiating $V_t^{\mathbf{c}, \hat{\theta}^t}$ and using equation (102) to compute $d[e^{-rt} v_t(\mathbf{c}, \theta)]$, we find³⁴

$$dV_t^{\mathbf{c}, \hat{\theta}^t} = e^{-rt} \left\{ \left(\beta_t \hat{\theta}_t - h(\hat{\theta}_t) \right) - \left(\beta_t \theta_t - h(\theta_t) \right) \right\} dt + e^{-rt} \beta_t \left(dy_t - \hat{\theta}_t dt \right).$$

Thus, since $(dy_t - \hat{\theta}_t dt)$ is a Brownian motion under $\mathbb{P}^{\hat{\theta}}$, if (103) holds the drift of $V_t^{\mathbf{c}, \hat{\theta}^t}$ under the probability measure $\mathbb{P}^{\hat{\theta}}$ is non-positive and thus $V_t^{\mathbf{c}, \hat{\theta}^t}$ is a $\mathbb{P}^{\hat{\theta}}$ -supermartingale. Hence we have

$$\mathbb{E}^{\hat{\theta}} \left[\int_0^t e^{-rs} \left(U(c_s) - h(\hat{\theta}_s) \right) ds \right] + \mathbb{E}^{\hat{\theta}} \left[e^{-rt} v_t(\mathbf{c}, \theta) \right] = \mathbb{E}^{\hat{\theta}} \left[V_t^{\mathbf{c}, \hat{\theta}^t} | \mathcal{F}_0 \right] \leq V_0^{\mathbf{c}, \hat{\theta}^0} = v_0(\mathbf{c}, \theta).$$

³³Note the fixed point nature of the argument: θ_t generates $v_t(\mathbf{c}, \theta)$ which yields β_t ; in turn, the incentives have to be satisfied given this process β_t .

³⁴This equation evaluates the incremental change in the agent's utility from pursuing the alternative effort strategy $\hat{\theta}$ for an additional unit of time during $[t, t + dt]$, and shows that in expectation such an incremental deviation hurts the agent. The next equation then uses a supermartingale argument to obtain inductively that the whole deviation strategy $\{\hat{\theta}_t\}_{t \geq 0}$ is worse than $\{\theta_t\}_{t \geq 0}$.

Taking the limit as $t \rightarrow \infty$ using the fact that $\mathbb{E}^{\hat{\theta}} [e^{-rt} v_t(\mathbf{c}, \theta)] \geq -e^{-rt} h(\bar{\theta})$, we obtain $v_0(\mathbf{c}, \hat{\theta}) \leq v_0(\mathbf{c}, \theta)$.

Conversely, if (103) does not hold on a set of times and sample paths with positive measure, then pick a deviation $\hat{\theta}$ defined as $\hat{\theta}_t = \arg \max_{\hat{\theta}} (-h(\hat{\theta}) + \beta_t \hat{\theta})$ everywhere. The drift of $V_t^{\mathbf{c}, \hat{\theta}^t}$ under $\mathbb{P}^{\hat{\theta}}$ is non-negative and positive on a set of positive measure, so that for t large enough the time-0 expected payoff from following $\hat{\theta}$ until time t and switching to θ thereafter is $v_0(\mathbf{c}, \hat{\theta}^t) = \mathbb{E}^{\hat{\theta}} [V_t^{\mathbf{c}, \hat{\theta}^t} | \mathcal{F}_0] > V_0^{\mathbf{c}, \hat{\theta}^0} = v_0(\mathbf{c}, \theta)$. Thus the strategy θ is suboptimal. \square

For a given effort θ , denote by $\beta(\theta)$ the sensitivity that maximizes $(-h(\theta) + \beta\theta)$, namely $\beta(\theta) = h'(\theta)$ if $\theta > 0$ and take $\beta(\theta) = 0$ if $\theta = 0$.

Propositions 11 and 12 imply that there is a one-to-one correspondence between incentive-compatible contracts (\mathbf{c}, θ) and controlled processes v_t that satisfy the one-shot deviation conditions (103), so that $\beta(\theta_t)$ replaces β_t in (102). We are now ready to reformulate the planner's problem as a stochastic control problem, using the continuation value v_t as the single state variable.

Solution to the optimal stochastic control problem. The planner maximizes his expected profit (97) over incentive compatible contracts (\mathbf{c}, θ) subject to the law of motion of v_t , the transversality conditions, and delivering initial promised utility \hat{v}_0 . We consider a relaxed problem without the transversality condition (to be checked ex-post). Before we analyze this problem, note that as in the discrete time setting, the principal has the option of “retiring” the agent at a given time τ by allocating a constant consumption $c_t = c$ and recommending zero effort $\theta_t = 0$ for all $t \geq \tau$. The continuation value at retirement time τ is then $v_\tau = r^{-1}U(c)$, so that $c = U^{-1}(rv_\tau)$. The retirement time τ must be specified in the contract, so it is a stopping time with respect to the filtration \mathcal{F}_t generated by the output process y . We can thus write the principal's value of the optimal contract as

$$K(\hat{v}_0) = \max_{\mathbf{c}, \theta, \tau} \mathbb{E}^{\theta} \left[\int_0^{\tau} e^{-rt} \left((\theta_t - c_t) dt + \sigma d\mathcal{Z}_t^{\theta} \right) - \frac{e^{-r\tau}}{r} U^{-1}(rv_\tau) \right] \quad (104)$$

$$\text{subject to} \quad dv_t = (rv_t - U(c_t) + h(\theta_t)) dt + \beta(\theta_t) \sigma d\mathcal{Z}_t^{\theta} \quad (105)$$

$$v_0 = \hat{v}_0. \quad (106)$$

Note in particular that the incentive constraints (98) are automatically satisfied if the constraint (105) holds. The function $K(v)$ can be found using standard optimal control and optimal stopping techniques, where the control variables are θ_t, c_t, τ and the state variable is v_t . The principal's problem can be solved in two steps: first, guess an optimal contract using the appropriate Bellman equation; second, verify ex post that this contract is indeed optimal.

We start by conjecturing the optimal contract. The function K is continuous on $[0, \infty)$ with $K(v) \geq -r^{-1}U^{-1}(rv)$ for all v . It satisfies the following Hamilton-Jacobi-Bellman (HJB)

equation.³⁵

$$rK(v) = \max \left\{ -U^{-1}(rv); \right. \\ \left. \max_{\substack{0 \leq \theta \leq \bar{\theta} \\ c \geq 0}} (\theta - c) + (rv - U(c) + h(\theta)) K'(v) + \frac{1}{2} \sigma^2 (\beta(\theta))^2 K''(v) \right\} \quad (107)$$

with the three boundary conditions

$$K(0) = 0, \quad K(\bar{v}) = -r^{-1}U^{-1}(r\bar{v}), \quad K'(\bar{v}) = -U^{-1'}(r\bar{v}), \quad (108)$$

for some $\bar{v} \geq 0$. Intuitively, (107) means that the principal maximizes the expected current flow of profit $(\theta - c)$ plus the expected change of future profit due to the drift and volatility of the agent's continuation value, until the stopping time τ at which the principal either retires the agent (if $v_\tau = \bar{v}$) or fires him (if $v_\tau = 0$). The second and third boundary conditions in (108) mean that the optimal retirement time occurs at the continuation value \bar{v} where the value-matching condition (which equates the value of retiring the agent with that of continuing with positive effort) and the smooth-pasting condition (which equates the marginal values of retiring and continuing) are satisfied. We can show that there exists a unique function K that satisfies the HJB equation (107) with the three boundary conditions (108). The stopping time $\tau = \inf \{t \geq 0 : K(v) \leq -r^{-1}U^{-1}(rv)\}$ satisfies $\tau < \infty$ a.s.,³⁶ and the function K is concave.

Define, for an arbitrary control policy $(\mathbf{c}, \boldsymbol{\theta})$, the process

$$G_t^{\mathbf{c}, \boldsymbol{\theta}} = \int_0^t e^{-rs} (\theta_s - c_s) ds + e^{-rt} K(v_t). \quad (109)$$

The following proposition conjectures the optimal contract from the solution to (107), (108) and then verifies that it is indeed optimal using martingale techniques:

Proposition 13. *Denote by $\theta(v), c(v)$ the maximizers in the right hand side of the HJB equation.³⁷ Consider the unique solution $K(v) \geq -r^{-1}U^{-1}(rv)$ to the HJB equation (107) that satisfies the conditions (108) for some $\bar{v} \geq 0$. For any $\hat{v}_0 \in [0, \bar{v}]$, define the process v_t by*

³⁵In fact, this is a Hamilton-Jacobi-Bellman Variational Inequality (HJBVI), where the HJB comes from the optimal stochastic control problem and the VI comes from the optimal stopping problem. See Chapter 4 in Øksendal and Sulem (2005).

³⁶If $h'(0) = 0$, the retirement point \bar{v} may not be finite, so that $K(v), c(v), \theta(v)$ asymptote to $K_0(v), \infty, 0$ as $v \rightarrow \infty$.

³⁷The optimal effort maximizes the difference between the expected flow of output θ , and the costs of compensating the agent for his effort, $-h(\theta)K'(v)$, and of exposing him to income uncertainty to provide incentives, $-\frac{\sigma^2}{2}\beta(\theta)^2K''(v)$. The optimal consumption is 0 for v small enough (i.e., for $K'(v) \geq -1/u'(0)$), and it is increasing in v according to $K'(v) = -1/U'(c)$ otherwise, where $1/U'(c)$ and $-K'(v)$ are the marginal costs of giving the agent value through current consumption and through his continuation payoff, respectively.

$v_0 = \hat{v}_0$ and

$$dv_t = r(v_t - u(c(v_t)) + h(\theta(v_t))) dt + r\beta(\theta(v_t))(dy_t - \theta(v_t) dt) \quad (110)$$

until the stopping time τ when v_τ hits 0 or \bar{v} . Define the contract $(\mathbf{c}, \boldsymbol{\theta})$ with payments $c_t = c(v_t)$ and recommended effort $\theta_t = \theta(v_t)$ for $t < \tau$, and $c_t = U^{-1}(rv_\tau)$ and $\theta_t = 0$ for $t \geq \tau$. Then $(\mathbf{c}, \boldsymbol{\theta})$ is incentive compatible and it has value $\hat{v}_0 = v_0(\mathbf{c}, \boldsymbol{\theta})$ to the agent and profit $K(\hat{v}_0)$ to the principal. Moreover, consider a concave solution K of the HJB equation (107). Any incentive compatible contract $(\mathbf{c}, \boldsymbol{\theta})$ yields to the principal a profit less than or equal to $K(v_0(\mathbf{c}, \boldsymbol{\theta}))$.

Proof. Let v_t be given by the stochastic differential equation (110) for $t \leq \tau$ and $v_t = v_\tau$ for $t > \tau$ (note in particular that $v \in [0, \bar{v}]$ is bounded). We show that $v_t = v_t(\mathbf{c}, \boldsymbol{\theta})$ for all $t \geq 0$, where $v_t(\mathbf{c}, \boldsymbol{\theta})$ is the agent's true continuation value in the contract $(\mathbf{c}, \boldsymbol{\theta})$ constructed above. This will imply in particular that the agent gets value $v_0(\mathbf{c}, \boldsymbol{\theta}) = \hat{v}_0$ from the contract. From the representation of $v_t(\mathbf{c}, \boldsymbol{\theta})$ in Proposition 11, we have

$$d(v_t(\mathbf{c}, \boldsymbol{\theta}) - v_t) = r(v_t(\mathbf{c}, \boldsymbol{\theta}) - v_t) dt + (\beta_t - \beta(\theta(v_t))) \sigma dZ_t^\theta,$$

hence for all $s \geq 0$, $\mathbb{E}^\theta[v_{t+s}(\mathbf{c}, \boldsymbol{\theta}) - v_{t+s}] = e^{rs}(v_t(\mathbf{c}, \boldsymbol{\theta}) - v_t)$. But $\mathbb{E}^\theta[v_{t+s}(\mathbf{c}, \boldsymbol{\theta}) - v_{t+s}]$ is bounded, hence $v_t = v_t(\mathbf{c}, \boldsymbol{\theta})$. Moreover, the contract $(\mathbf{c}, \boldsymbol{\theta})$ is incentive compatible by construction since the process from Proposition 11 that represents $v_t(\mathbf{c}, \boldsymbol{\theta})$ is $\beta_t = \beta(\theta_t)$.

Next we show that the principal gets expected profit $K(\hat{v}_0)$ from the contract. Differentiating expression (109) and applying Itô's lemma to $K(v_t)$ yields that the drift of $G_t^{\mathbf{c}, \boldsymbol{\theta}}$ under \mathbb{P}^θ is

$$e^{-rt} \left\{ (\theta_t - c_t - rK(v_t)) + (rv_t - U(c_t) + h(\theta_t)) K'(v_t) + \frac{1}{2} \sigma^2 (\beta(\theta_t))^2 K''(v_t) \right\}.$$

Thus, when $c_t = c(v_t)$ and $\theta_t = \theta(v_t)$, the drift of $G_t^{\mathbf{c}, \boldsymbol{\theta}}$ under \mathbb{P}^θ is equal to zero before time τ , so that $G_t^{\mathbf{c}, \boldsymbol{\theta}}$ is a martingale. By the Optional Stopping Theorem, we thus obtain that the principal's profit from the contract is

$$\mathbb{E}^\theta \left[\int_0^\tau e^{-rs} (\theta_s - c_s) ds \right] + \mathbb{E}^\theta [e^{-r\tau} K(v_\tau)] = \mathbb{E} [G_\tau^{\mathbf{c}, \boldsymbol{\theta}}] = G_0^{\mathbf{c}, \boldsymbol{\theta}} = K(v_0(\mathbf{c}, \boldsymbol{\theta})).$$

Finally, consider an alternative incentive compatible contract $(\mathbf{c}, \boldsymbol{\theta})$. Then (107) implies that the drift of $G_t^{\mathbf{c}, \boldsymbol{\theta}}$ under \mathbb{P}^θ is smaller than zero, so that $G_t^{\mathbf{c}, \boldsymbol{\theta}}$ is a bounded supermartingale. By the Optional Stopping Theorem, we obtain that the principal's expected profit at time 0 is less than or equal to $G_0^{\mathbf{c}, \boldsymbol{\theta}} = K(v_0(\mathbf{c}, \boldsymbol{\theta}))$. We refer to Sannikov (2008) for the technical details omitted in this sketch of proof. \square

For any $v_0 > \bar{v}$, the function $K(v)$ is negative and is an upper bound on the principal's value function; thus, there is no profitable contract with positive profit to the principal in that range. In the range $(0, \bar{v})$, replacing the optimal consumption $c(v)$ and effort $\theta(v)$ into the

HJB equation, we obtain a non-linear second order differential equation for $K(v)$ which can be solved numerically. Finally, note that the envelope theorem applied to the HJB equation before retirement implies

$$(rv - U(c) + h(\theta)) K''(v) + \frac{1}{2} \sigma^2 (\beta(\theta))^2 K'''(v) = 0.$$

By Itô's lemma, the left hand side is the drift of $K'(v_t) = -1/U'(c_t)$ on the interval $[\underline{v}, \bar{v}]$. Thus the inverse of the agent's marginal utility is a martingale when the agent's consumption is positive, a result that parallels the Inverse Euler Equation (64) found in the discrete-time model.

Sannikov (2014) extends the analysis of the moral hazard model in continuous time to the case where current actions affect not only current output, but also future output. The solution to this problem is more involved than that of Sannikov (2008), but the steps and the martingale techniques (using the Martingale Representation Theorem to simplify the set of incentive constraints and reduce the problem to a stochastic control problem) are similar.

We conclude this section with a brief discussion of the benefits of using a continuous-time rather than discrete-time framework to analyze dynamic contracting problems. First, the Hamilton-Jacobi-Bellman equation is more tractable analytically than the discrete-time Bellman equation (23). In particular De Marzo and Sannikov (2006) show how differentiating the HJB equation and its boundary conditions that characterize the optimal contract allows us to derive comparative statics results analytically. Here we illustrate their method on a simple example. Suppose we are interested in the effect of the volatility σ^2 on the principal's profit $K(v)$. Differentiating (107) yields, for $v \in (0, \bar{v})$,

$$r \frac{\partial K(v)}{\partial \sigma^2} = \frac{1}{2} (\beta(\theta))^2 K''(v) + (rv - U(c) + h(\theta)) \left(\frac{\partial K(v)}{\partial \sigma^2} \right)' + \frac{1}{2} \sigma^2 (\beta(\theta))^2 \left(\frac{\partial K(v)}{\partial \sigma^2} \right)'' ,$$

with the following boundary condition, obtained by differentiating the value-matching condition (108) at \bar{v} :

$$\frac{\partial K}{\partial \sigma^2}(\bar{v}) = - (K'(\bar{v}) + U^{-1'}(r\bar{v})) \frac{\partial \bar{v}}{\partial \sigma^2} = 0.$$

A generalization of the Feynman-Kac formula (see De Marzo and Sannikov (2006) for the technical details) implies that the solution to this differential equation can be written as a conditional expectation:

$$\frac{\partial K(v)}{\partial \sigma^2} = \frac{1}{2} \mathbb{E}^\theta \left[\int_0^\tau e^{-rt} \beta^2(\theta_t) K''(v_t) dt + e^{-r\tau} \frac{\partial K}{\partial \sigma^2}(\bar{v}) | v_0 = v \right] < 0,$$

where v_t evolves according to (105), and where the inequality follows from the strict concavity of the profit function K . Intuitively, the right hand side of this equation sums the profit gains and losses along the path of v_t due to an increase in σ^2 . This shows that a higher volatility

σ^2 reduces the principal's profit. We can similarly evaluate the effects of all the parameters of the model on the principal's profit, the agent's time-0 utility (by differentiating the optimality condition $K'(v_0) = 0$), or the value at retirement \bar{v} (by differentiating the boundary conditions (108)).

Finally, another advantage of the continuous-time problem is that it is also more tractable computationally. In particular, the continuous-time formulation (107) can be computed more easily as the solution to an ordinary differential equation with a free boundary, while computing the solution to the discrete time Bellman equation (23) is more involved.

4 Applications

In this section we discuss several applications of the theory of recursive contracts. The methods developed in the previous sections can be used to analyze questions in public economics, corporate finance, development, international finance, and political economy. Our goal is not to provide a comprehensive overview of those fields. Rather we want to show how several general principles emphasized above can be used to obtain rich insights in very different areas and relate those insights to empirical observations.

4.1 Public Finance

Individuals are subject to a variety of idiosyncratic shocks. Illness, disability, job loss, structural changes in the economy that diminish the value of human capital, unexpected promotions and demotions, success and failure in business ventures, all significantly affect individuals' incomes. It has been recognized at least since the work of Vickrey (1947) that the tax and transfer system can provide insurance against such shocks and help individuals smooth their consumption across different dates and states. A natural question is then how to design the optimal social insurance system that provides the best insurance given the distortions imposed by those programs.

Diamond and Mirrlees (1978), Diamond et al. (1980) and Diamond and Mirrlees (1986) were some of the first papers to systematically study this question. At the same time, solving these problems either analytically or computationally is very difficult even in relatively simple dynamic settings. The advances in the theory of recursive contracts in the late 1980s and 1990s delivered a set of tools that that allowed to overcome many of the difficulties. The New Dynamic Public Finance literature applied those tools to the study of dynamic optimal taxation: see, e.g., Golosov et al. (2003), Albanesi and Sleet (2006), Golosov et al. (2006), Golosov and Tsyvinski (2006, 2007), Farhi and Werning (2013, 2012, 2007), Werning (2009), Kocherlakota (2010), Golosov et al. (2016), Stantcheva (2014). In what follows, we describe a model that illustrates some of the main results of this literature.

We focus on a partial equilibrium model in which individuals are subject to idiosyncratic

shocks to labor productivity.³⁸ The economy lasts T periods, where T can be finite or infinite. Each agent's preferences are described by a time separable utility function over consumption $c_t \geq 0$ and labor supply $l_t \geq 0$,

$$\mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} U(c_t, l_t) \right], \quad (111)$$

where $\beta \in (0, 1)$ is a discount factor, \mathbb{E}_0 is a period-0 expectation operator, conditional on the shock at date $t = 0$, and $U : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ is differentiable, strictly increasing and concave in consumption, and decreasing and concave in labor supply. The partial derivatives of the utility function are denoted by U_c and U_l .

Agents draw their initial type (skill) θ_1 from a distribution $\pi_1(\cdot)$ in period 1. From then on skills follow a Markov process $\pi_t(\theta_t | \theta_{t-1})$, where θ_{t-1} is the agent's skill realization in period $t - 1$. We denote the probability density function of period- t types conditional on θ_{t-1} by $\pi_t(\cdot | \theta_{t-1})$. Skills are non-negative: $\theta_t \in \Theta \subset \mathbb{R}_+$ for all t . At this stage we are agnostic about the dimensionality of Θ and allow Θ to be discrete or continuous. The set of possible histories up to period t is denoted by Θ^t . An agent of type θ_t who supplies l_t units of labor produces $y_t = \theta_t l_t$ units of output.

In this partial equilibrium economy, y_t also denotes the labor income of individuals. Individuals can freely borrow and lend at an exogenous interest rate R . We assume that there is no insurance available to individuals except self-insurance through borrowing and lending and through taxes and transfers provided by the government. We are interested in understanding how the government can design the optimal tax system $\mathcal{T}_t(\cdot)$ as a function of the information it has about individuals. We are thinking of the function \mathcal{T}_t in very general terms: it is a combination of all taxes and transfers that individuals pay to or receive from the government. We are seeking a function \mathcal{T}_t that maximizes welfare given by (111) in a competitive equilibrium.³⁹

If individuals' skills are observable, the optimal tax function is very simple: \mathcal{T}_t should depend on the realization of the shocks θ^t and prescribe positive or negative transfers without distorting either labor supply or savings decisions. In reality idiosyncratic shocks are difficult to observe. Even disability insurance programs which extensively employ medical examinations to determine whether an applicant is subject to medical conditions that make a person unable to work are subject to substantial moral hazard problems and asymmetric information (see Golosov and Tsyvinski (2006) and references therein). Therefore we make the assumption that the realizations of θ_t are not observed by the government and the only observable choices are labor income, consumption and capital.

We study the optimal taxes using a two-step procedure. In the first step, we invoke the Revelation Principle (see Section 2.2) and write the problem as a mechanism design program

³⁸See Albanesi (2011), Shourideh (2010), Abraham and Pavoni (2008) for applications of recursive contracting tools to taxation with shocks to savings, Stantcheva (2014) for human capital accumulation, Hosseini et al. (2013) for fertility choices.

³⁹It is straightforward to extend this analysis and allow other welfare criteria or expenditures on public goods (see, e.g. Golosov et al. (2003)).

whose solution can be characterized using recursive techniques. In the second step, we back out a tax function \mathcal{T}_t that can implement that solution in a competitive equilibrium.

The mechanism design problem is as follows. Let reports be given by $\sigma_t : \Theta^t \rightarrow \Theta$ and allocations by $c_t : \Theta^t \rightarrow \mathbb{R}_+$, $y_t : \Theta^t \rightarrow \mathbb{R}_+$, for all $t \geq 1$. The incentive constraint (8) writes

$$\begin{aligned} & \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} U \left(c_t(\theta^t), \frac{y_t(\theta^t)}{\theta_t} \right) \right] \\ & \geq \mathbb{E}_0 \left[\sum_{t=1}^T \beta^{t-1} U \left(c_t(\sigma^t(\theta^t)), \frac{y_t(\sigma^t(\theta^t))}{\theta_t} \right) \right], \quad \forall \sigma^T \in \Sigma^T, \end{aligned} \quad (112)$$

and the feasibility constraint (2) becomes

$$\mathbb{E}_0 \left[\sum_{t=1}^T R^{1-t} c_t(\theta^t) \right] \leq \mathbb{E}_0 \left[\sum_{t=1}^T R^{1-t} y_t(\theta^t) \right]. \quad (113)$$

The planner maximizes the ex-ante expected utility (111) of the agents, i.e., provides optimal ex-ante insurance. This problem is thus similar to that analyzed in Section 2, see equations (9) or (61). Solving this problem directly is difficult. There are prohibitively many incentive constraints (112) either for analytical or numerical analysis in most applications. In the next sections we overcome this problem using recursive techniques. For concreteness, we assume separable isoelastic preferences

$$U(c, l) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \frac{l^{1+\varepsilon}}{1+\varepsilon}. \quad (114)$$

While these preferences are not needed for most of the insights, they simplify the exposition of the main results.

4.1.1 Analysis with i.i.d. shocks

We start the analysis by assuming that shocks are independent and identically distributed over time, so that the probability of realization of any $\theta \in \Theta$ in any period can be written as $\pi(\theta)$. This assumption, although unrealistic, allows us to illustrate many insights very transparently.

We follow the steps familiar from the analysis in Sections 2.3 and 2.4. To ensure convexity, we rewrite our maximization problem in terms of utils of consumption and leisure rather than c and l . To this end we define the functions $C(u) = [1 + (1 - \sigma)u]^{1/(1-\sigma)}$ and $Y(h) = [(1 + \varepsilon)h]^{1/(1+\varepsilon)}$. We apply the one-shot deviation result from Propositions 2 and 3 to write the incentive compatibility and promise keeping constraints as:

$$u_t(\theta^t) - \theta_t^{-(1+\varepsilon)} h(\theta^t) + \beta v_t(\theta^t) \geq u_t(\theta^{t-1}, \hat{\theta}) - \theta_t^{-(1+\varepsilon)} h(\theta^{t-1}, \hat{\theta}) + \beta v_t(\theta^{t-1}, \hat{\theta})$$

for all $\theta^{t-1} \in \Theta^{t-1}$, $\hat{\theta} \in \Theta$, with

$$v_{t-1}(\theta^{t-1}) = \int_{\Theta} \left[u_t(\theta^{t-1}, \theta) - \theta^{-(1+\varepsilon)} h(\theta^{t-1}, \theta) + \beta v_t(\theta^{t-1}, \theta) \right] d\pi(\theta).$$

Following the same steps as in Section 2.3 we write the Bellman equation as

$$K_t(v) = \min_{\{u(\theta), h(\theta), w(\theta)\}_{\theta \in \Theta}} \int_{\Theta} \left[C(u(\theta)) - Y(h(\theta)) + R^{-1} K_{t+1}(w(\theta)) \right] d\pi(\theta) \quad (115)$$

subject to the incentive constraints: for all $\theta, \hat{\theta} \in \Theta$,

$$u(\theta) - \theta^{-(1+\varepsilon)} h(\theta) + \beta w(\theta) \geq u(\hat{\theta}) - \theta^{-(1+\varepsilon)} h(\hat{\theta}) + \beta w(\hat{\theta}),$$

the promise-keeping constraint:

$$v = \int_{\Theta} \left[u(\theta) - \theta^{-(1+\varepsilon)} h(\theta) + \beta w(\theta) \right] d\pi(\theta),$$

and $K_{T+1}(w) = 0$ for all w if T is finite. When T is infinite, the subscript t drops out of the Bellman equation above.

Many of the qualitative properties of this model can be obtained along the lines of Proposition 5. For example, using steps analogous to those of Section 2.4, it is easy to show the analogue of equations (29) and (63):⁴⁰

$$K'_t(v) = \mathbb{E} [C'(u_{v,t})] = (\beta R)^{-1} \mathbb{E} [K'_{t+1}(w_{v,t}) | v]. \quad (116)$$

Moreover, optimality also requires: for all θ, t, v ,

$$C'(u_{v,t}(\theta)) = (\beta R)^{-1} K'_{t+1}(w_{v,t}(\theta)). \quad (117)$$

The intuition for this result is simple. The planner can provide incentives to reveal information either intratemporally, by giving an agent higher contemporaneous utility, or intertemporally, by giving higher future promises. Condition (117) implies that it is optimal to equalize the marginal costs of the two ways of providing incentives.

These conditions have some immediate but unexpected implications for taxation. Note that $C' = \frac{1}{U_c}$, where U_c is the marginal utility of consumption, and hence (116) can be re-written as

$$\frac{\beta R}{U_c(c_{v,t}(\theta))} = \mathbb{E} \left[\frac{1}{U_c(c_{v,t+1})} \middle| v \right], \quad \forall v, t, \theta.$$

⁴⁰This condition is particularly easy to derive if Θ is finite, in which case it can be obtained by simple manipulation of the Lagrangians on the incentive constraints.

The policy functions generate the constrained-optimal stochastic processes $\{c_t^*, y_t^*\}_{t=1}^T$ that satisfy the Inverse Euler equation (see our discussion in Section 2.6 as well as Golosov et al. (2003)):

$$\frac{\beta R}{U_c(c_t^*)} = \mathbb{E}_t \left[\frac{1}{U_c(c_{t+1}^*)} \right]. \quad (118)$$

By Jensen's inequality, we have $\mathbb{E} \left[\frac{1}{X} \right] \geq \frac{1}{\mathbb{E}[X]}$ for any random variable X , with strict inequality if X is non-deterministic. Therefore this equation implies that at the optimum,

$$U_c(c_t^*) \leq \beta R \mathbb{E}_t [U_c(c_{t+1}^*)],$$

with strict inequality if future consumption is uncertain. Therefore, it follows that the optimal tax system must introduce positive savings distortions in this economy. One useful way to summarize the distortions introduced by the tax system is to define the *savings wedge* as

$$1 - \tau_t^s(\theta^t) \equiv \frac{1}{\beta R} \frac{U_c(c_t^*(\theta^t), y_t^*(\theta^t)/\theta_t)}{\mathbb{E}_t [U_c(c_{t+1}^*(\theta^{t+1}), y_{t+1}^*(\theta^{t+1})/\theta_{t+1})]}. \quad (119)$$

Optimality implies that $\tau_t^s(\theta^t) \geq 0$ for all θ^t , with strict equality if consumption in $t + 1$ is uncertain.

Decentralization We now describe how the government can design a tax system \mathcal{T}_t such that agents optimally choose consumption and income $\{c_t^*, y_t^*\}_{t=1}^T$ given a budget constraint

$$c_t + k_{t+1} \leq y_t + Rk_t - \mathcal{T}_t.$$

That is, this tax function \mathcal{T}_t is an *implementation* or *decentralization* of the constrained optimum. We want to understand what arguments $\mathcal{T}_t(\cdot)$ should depend on and how to construct it.

In general, there are many tax systems that implement the same allocation.⁴¹ Here we consider a particularly simple implementation that arises naturally from the recursive problem. Observe that to find the optimal allocations in period t in the Bellman equation (115), we did not need to know the whole past history θ^t . It was sufficient to know the summary statistics $v_{t-1}(\theta^{t-1})$ together with the current period shock θ_t . A natural analogue of the promised utility in competitive equilibrium is the agent's savings. Albanesi and Sleet (2006) use this insight to show that when types are i.i.d. and the utility function is separable between consumption and labor supply we can construct an optimal tax system in which taxes in period t depend only on labor income y_t and on savings k_t at the beginning of that period.

Proposition 14. *Assume that shocks are i.i.d. and preferences are separable in consumption*

⁴¹For example, an extreme tax system $\mathcal{T}_t(\{y_s\}_{s=1}^t)$ defined as $\mathcal{T}_t(\{y_s\}_{s=1}^t) = y_t^*(\theta^t) - c_t^*(\theta^t)$ if $y_s = y_s^*(\theta^s)$ for all $\theta^s \leq \theta^t$ and $\mathcal{T}_t(\{y_s\}_{s=1}^t) = \infty$ otherwise, ensures that the only feasible choices for a consumer are $\{c_t^*, y_t^*\}_{t=1}^T$. Then the incentive compatibility constraint ensures that $\mathcal{T}_t(\{y_s\}_{s=1}^t)$ implements $\{c_t^*, y_t^*\}_{t=1}^T$.

and labor. The optimal allocations can be implemented by a tax system $\mathcal{T}_t(k_t, y_t)$.

Proof. We show this result in a two-period economy. Let $K_2(w_2)$ denote planner's minimized cost function (115) in period 2, and $u_{w_2}^*(\theta), h_{w_2}^*(\theta)$ denote the policy functions that solve the second-period planner's problem.

In period 2, consider an individual who enters the period with savings k_2 and chooses labor income y_2 . Suppose that $k_2 = K_2(w_2)$ for some promised utility w_2 , and $y_2 = Y(h_{w_2}^*(\theta))$ for some $\theta \in \Theta = [\underline{\theta}, \bar{\theta}]$. We then define the tax function $T_2(k_2, y_2)$ as⁴²

$$T_2(K_2(w_2), Y(h_{w_2}^*(\theta))) = K_2(w_2) + Y(h_{w_2}^*(\theta)) - C(u_{w_2}^*(\theta)).$$

By incentive compatibility, an agent with savings $k_2 = K_2(w_2)$ and type θ in period 2 chooses labor supply and consumption $(y_2, c_2) = (Y(h_{w_2}^*(\theta)), C(u_{w_2}^*(\theta)))$, that is, the levels optimally assigned to his promised utility-type pair (w_2, θ) .

In period 1, consider an individual who enters the period with savings k_1 and chooses labor income y_1' (which may or may not be optimal given his first-period type θ). Denote by $(c_1', R^{-1}k_2')$ his optimal consumption-savings choice given (k_1, y_1') , and by $\tilde{u}' = U(c_1') + \beta E[V_2(k_2', \theta_2)]$ (where V_2 is the maximized objective of the agent in period 2) the utility that he achieves with this combination, gross of the disutility of labor. The key is that the cost-minimizing way for the planner to deliver utility \tilde{u}' to the agent is to offer the pair $(u_1', w_2') = (U(c_1'), K_2^{-1}(k_2'))$, and the corresponding cost is $C_{\tilde{u}'} = c_1' + R^{-1}k_2'$.

Now suppose that $y_1' = y_1^*(k_1, \theta')$ for some $\theta' \in \Theta = [\underline{\theta}, \bar{\theta}]$, where $y_1^*(k_1, \theta')$ denotes the first-period income optimally allocated to type θ' in the solution to the planner's problem. Define the tax function $T_1(k_1, y_1')$ as

$$T_1(k_1, y_1^*(k_1, \theta')) = k_1 + y_1^*(k_1, \theta') - C_{\tilde{u}'}$$

If the individual's true type is $\theta \neq \theta'$, by lying he reaches utility $\tilde{u}' - \theta^{-(1+\varepsilon)}h(y_1^*(k_1, \theta'))$. But by incentive compatibility this is smaller than the utility he gets by reporting his true type, namely $\tilde{u} - \theta^{-(1+\varepsilon)}h(y_1^*(k_1, \theta))$. Thus under this tax function the agent finds it optimal to choose the income that corresponds to his true type in period 1, and his choice of savings will be exactly equal to $k_2 = K_2(w_2(\theta))$, since his net income is $C_{\tilde{u}}$. \square

This proposition shows simultaneously that optimal allocations can be implemented by a joint tax on current period savings and labor income, and provides a method of constructing this tax.

When thinking about the relationship between this tax \mathcal{T}_t and taxes in the data, it is important to keep in mind that \mathcal{T}_t in the model corresponds to the sum of all taxes and transfers in the data. The marginal distortions with respect to capital and labor income, $\frac{\partial \mathcal{T}_t}{\partial k_t}$ and $\frac{\partial \mathcal{T}_t}{\partial y_t}$,

⁴²The tax function can be easily extended to deter any move $y_2 > Y(h_{w_2}^*(\underline{\theta}))$ and $y < Y(h_{w_2}^*(\bar{\theta}))$.

correspond to the effective marginal tax rates in the data, which are a sum of statutory tax rates and the rates of phasing out of transfers in capital and labor income respectively. Because of the phasing out of transfers, there is no reason to expect a priori that marginal taxes in the model and effective marginal taxes in the data are progressive.⁴³ For example, if individuals with more wealth receive less insurance against labor income shocks (e.g., if they are not eligible to some welfare programs because of means-testing), we should expect the marginal labor taxes to be decreasing in capital.

4.1.2 Persistent shocks

An important limitation of the previous discussion is the assumption that shocks are i.i.d. The empirical labor literature has emphasized that idiosyncratic shocks are highly persistent (for example, Storesletten et al. (2004) or Guvenen et al. (2015)). In this section we discuss how to extend our analysis to persistent (Markov) shocks.

It is useful to assume, both for analytical tractability and for connecting the analysis to the empirical literature, that shocks are drawn from a continuous distribution. We focus on a family of stochastic processes frequently used in the applied labor and public finance literatures.⁴⁴

Assumption 6. *Suppose that shocks θ_t evolve according to*

$$\ln \theta_t = b_t + \rho \ln \theta_{t-1} + \eta_t,$$

where η_t is drawn from one of the following three distributions:

- (a) *lognormal*: $\eta_t \sim \mathcal{N}(0, \nu)$;
- (b) *Pareto-lognormal*: $\eta_t \sim \mathcal{NE}(\mu, \nu, a)$, where \mathcal{NE} is a normal-exponential distribution;
- (c) *mixture of lognormals*: $\eta_t \sim \mathcal{N}(\mu_i, \nu_i)$ with probability p_i for $i = 1, \dots, I$. Let $\nu = \max_i \nu_i$.

We can write the planner's problem recursively by applying the first-order approach discussed in Section 2.5. Under these assumptions the Bellman equation writes:

$$K_t(v, \hat{v}, \theta_-) = \max_{(\vec{u}, \vec{h}, \vec{w}, \vec{\hat{w}})} \int_0^\infty (Y(h(\theta)) - C(u(\theta)) + R^{-1}K_{t+1}(w(\theta), \hat{w}(\theta), \theta)) \pi_t(\theta | \theta_-) d\theta \quad (120)$$

⁴³In the U.S. there is significant heterogeneity in the shapes of the effective tax rates as a function of income as they vary by state, family status, age, type of residence a person lives in, etc. Some typical patterns of the effective marginal rates in the U.S. data are increasing, U-shaped, and inverted S-shaped (see CBO (2007) and Maag et al. (2012)).

⁴⁴For example, Storesletten et al. (2004) and Farhi and Werning (2013) use lognormal distributions, Badel and Huggett (2014) and Lockwood et al. (2014) use Pareto-lognormal distributions, Geweke and Keane (2000) and Guvenen et al. (2015) use mixtures of lognormals.

subject to the promise-keeping and marginal promise-keeping constraints

$$v = \int_0^\infty \varpi(\theta) \pi_t(\theta | \theta_-) d\theta, \quad (121)$$

$$\hat{v} = \int_0^\infty \varpi(\theta) \hat{\pi}_t(\theta | \theta_-) d\theta, \quad (122)$$

$$\varpi(\theta) = u(\theta) - \theta^{-(1+\varepsilon)} h(\theta) + \beta w(\theta), \quad (123)$$

and the envelope condition

$$\dot{\varpi}(\theta) = (1 + \varepsilon) \theta^{-(2+\varepsilon)} h(\theta) + \beta \hat{w}(\theta). \quad (124)$$

This problem can then be analyzed using optimal control techniques (see Golosov et al. (2016)).

The analysis of savings distortions remains unchanged. In particular, the Inverse Euler Equation (118) continues to hold in this economy. The same arguments as in the previous section immediately imply the optimality of savings distortions.

We now turn to the analysis of labor distortions. We define the *labor wedge* as

$$1 - \tau_t^y(\theta^t) \equiv \frac{-U_l(c_t^*(\theta^t), y_t^*(\theta^t) / \theta_t)}{\theta_t U_c(c_t^*(\theta^t), y_t^*(\theta^t) / \theta_t)}. \quad (125)$$

To simplify the notations, for any history $\theta^t = (\theta^{t-1}, \theta)$ and random variable x_t , we use the short-hand notations $x_t(\theta)$ to denote $x_t(\theta^{t-1}, \theta)$ and x_{t-1} to denote $x_{t-1}(\theta^{t-1})$. Manipulating the first-order conditions we obtain

$$\begin{aligned} \frac{\tau_t^y(\theta)}{1 - \tau_t^y(\theta)} &= (1 + \varepsilon) \frac{\int_\theta^\infty \pi_t(x') dx'}{\theta \pi_t(\theta)} \int_\theta^\infty \frac{U_{c,t}(\theta)}{U_{c,t}(x)} \left(1 - \int_0^\infty \frac{U_{c,t}(x)}{U_{c,t}(x')} \pi_t(x') dx' \right) \frac{\pi_t(x) dx}{\int_\theta^\infty \pi_t(x') dx'} \\ &\quad + \rho \beta R \frac{\tau_{t-1}^y}{1 - \tau_{t-1}^y} \frac{U_{c,t}(\theta)}{U_{c,t-1}}. \end{aligned} \quad (126)$$

Equation (126) shows that the optimal labor distortion is the sum of two terms. The first (“intra-temporal”) term on the right hand side captures the costs and benefits of labor distortions in providing insurance against period- t shocks. A labor distortion for type θ discourages that type’s labor supply, as captured by the Frisch elasticity of labor supply ε . This lowers total output in proportion to $\theta \pi_t(\theta)$ but allows the planner to relax the incentive constraints for all types above θ , a trade-off summarized by the hazard ratio (of period- t shocks conditional on a given history θ^{t-1}), $\frac{\int_\theta^\infty \pi_t(x') dx'}{\theta \pi_t(\theta)}$. Finally, the relaxed incentive constraints allow the planner to extract more resources from individuals with skills above θ and transfer them to all agents. The social value of this transfer is captured by the integral term on the r.h.s., which depends on the marginal utilities of consumption of agents with skills above θ , weighted by the average marginal utility. The second term (“inter-temporal”) on the right hand side captures how the

planner uses distortions in the current period t to provide incentives for information revelation in earlier periods. It depends on the information that the period- t shock carries about θ^{t-1} , summarized by the coefficient ρ , and on the ratio $\frac{U_{c,t}(\theta)}{U_{c,t-1}}$ which captures the fact that it is cheaper to provide incentives in those states in which the marginal utility of consumption is high.

We can also use the decomposition (126) to obtain insights about the time series properties of the optimal labor distortions, as studied by Farhi and Werning (2013). Multiplying the expression above by $\frac{1}{U_{c,t}}\pi_t(\theta)$ and integrating by parts yields

$$\mathbb{E}_{t-1} \left[\frac{\tau_t^y}{1 - \tau_t^y} \frac{1}{U_{c,t}} \right] = \rho\beta R \frac{\tau_{t-1}^y}{1 - \tau_{t-1}^y} \frac{1}{U_{c,t-1}} + (1 + \varepsilon) \text{Cov}_{t-1} \left(\ln \theta, \frac{1}{U_{c,t}} \right). \quad (127)$$

Equation (127) shows that the marginal utility-adjusted labor distortions follow an AR(1) process with a drift. The persistence of that process is determined by the persistence of the shock process ρ , and its drift is strictly positive since we should generally expect that $\text{Cov}_{t-1} \left(\ln \theta, \frac{1}{U_{c,t}} \right) > 0$. Farhi and Werning (2013) conclude that the optimal labor distortions should increase with age.

Golosov et al. (2016) use condition (126) to characterize the dependence of labor wedges on the realization of the shock θ . In particular they show the asymptotic laws of motion⁴⁵

$$\frac{\tau_t^y(\theta)}{1 - \tau_t^y(\theta)} \underset{\theta \rightarrow \infty}{\sim} \begin{cases} \left(\frac{a}{1+\varepsilon} - \frac{\sigma}{\sigma+\varepsilon} \right)^{-1}, & \text{if } \eta_t \text{ is Pareto-lognormal, } \frac{a}{1+\varepsilon} - \frac{\sigma}{\sigma+\varepsilon} > 0 \\ \left(\frac{\ln \theta}{\nu^2} \frac{1}{1+\varepsilon} \right)^{-1}, & \text{if } \eta_t \text{ is lognormal or a mixture,} \end{cases} \quad (128)$$

and

$$\frac{\tau_t^y(\theta)}{1 - \tau_t^y(\theta)} \underset{\theta \rightarrow 0}{\sim} \rho\beta R \frac{\tau_{t-1}^y}{1 - \tau_{t-1}^y} \left(\frac{c_t(0)}{c_{t-1}} \right)^{-\sigma}. \quad (129)$$

Given the fact that $(\ln \theta)^{-1}$ is very slowly moving, equation (128) implies that the labor distortions are approximately flat for high realizations of θ_t for all three classes of distributions (although in the cases of lognormal and mixture of lognormal distributions they eventually converge to zero), they do not depend on the past history of shocks, and they are given by relatively simple closed-form expressions. Equation (129) shows that the labor distortions for low shocks depend on the persistence, the past history, and the growth rate of consumption, and are generally increasing in age.

Another implication of these equations is that the higher moments, such as the kurtosis, play an important qualitative and quantitative role for the size of the labor distortions. Some of the best estimates of those moments are obtained by Guvenen et al. (2014) and Guvenen et al. (2015) who use U.S. administrative data on a random sample of 10% of the U.S. male taxpayers to estimate the stochastic process for labor earnings. Golosov et al. (2016) use that finding to

⁴⁵For any functions h, g and $c \in \bar{\mathbb{R}}$, $h(x) \underset{x \rightarrow c}{\sim} g(x)$ if $\lim_{x \rightarrow c} h(x)/g(x) = 1$.

calibrate their model using newly available estimates of idiosyncratic shocks. The optimal labor distortions are U-shaped while savings distortions are increasing in current earnings. Welfare in the constrained optimum is 2 to 4 percent higher than in the equilibrium with affine taxes. These findings (both the U-shape and the relatively high welfare gains from non-linear, history-dependent taxation) are largely driven by the high kurtosis found in the labor earnings process in the data. This suggests that a system of progressive taxes and history-dependent transfers that are being phased out relatively quickly with income can capture most of the welfare gains in this economy.

4.2 Corporate Finance

In this section we describe some applications of the recursive contract theory to corporate finance. We show how financing frictions arise endogenously from agency problems, leading to implications for firms' capital structure and dynamics. To cite only a few papers in this literature, such models have been analyzed by Albuquerque and Hopenhayn (2004), Clementi and Hopenhayn (2006), DeMarzo and Fishman (2007a,b) in discrete-time environments, and by De Marzo and Sannikov (2006), Biais et al. (2007), DeMarzo et al. (2012), Biais et al. (2010), He (2009) in continuous-time environments.

Endogenous financing frictions and firm dynamics. A large empirical literature (see, e.g., Caves (1998) for a survey) describes the properties of firm dynamics, e.g., the characteristics and evolution of their size, growth rates, and survival probabilities. In particular, as firms get older, their size and survival probability increase, the mean and variance of their growth rates decrease, and the hazard rates for exit first increase and then decrease. Moreover, starting with the work of Fazzari et al. (1988), many authors have found that firms' investment responds positively to innovations in the cash-flow process (after controlling for Tobin's q), suggesting the importance of borrowing constraints, and that the investment-cash flow sensitivity decreases with the firm's age and size. Clementi and Hopenhayn (2006) analyze a dynamic moral hazard model where such features arise *endogenously* in the optimal contract between a borrower (a firm, or agent) and a lender (bank, or principal) who cannot observe the outcome of the project. They describe the optimal contract and show that the model yields rich testable predictions about firm dynamics that are in line with the evidence presented above.

Their model is as follows. The borrower and the lender are both risk-neutral, have the same discount factor β , and have the ability to commit to contracts. The agent's project requires a fixed initial investment $I_0 > 0$ and a per-period investment of capital, denoted k_t . Revenues are stochastic (i.i.d.) and increase with the amount of capital advanced by the bank. Specifically, in each period t , with probability π the project is successful and yields revenue $R(k_t)$, where R is continuous, bounded and concave, whereas with probability $(1 - \pi)$ it yields zero revenues. Denote the outcome of the project in period t by $\theta_t \in \Theta = \{\theta_{(1)}, \theta_{(2)}\} = \{0, 1\}$, where $\theta_{(1)} = 0$ is

failure and $\theta_{(2)} = 1$ is success, and histories up to period t by θ^t . The project can be liquidated at the beginning of each period, with scrap value $S \geq 0$.

Suppose first that revenues are observable. The efficient amount of capital is advanced in every period, $k^* = \arg \max (\pi R(k_t) - k_t)$, and running the project is efficient if $W^* \equiv \frac{1}{1-\beta} (\pi R(k^*) - k^*) > I_0$. Assume that $W^* > S$. Thus, in the benchmark complete-information version of the model, the firm neither grows, shrinks, or exits: its size k^* is constant. This feature allows us to cleanly analyze the implications of informational frictions for the firm's dynamics.

Now suppose that revenues are private information to the borrower. Thus, in each period the borrower observes the outcome of the project and sends a report to the lender according to a strategy $\sigma = \{\sigma_t(\theta^t)\}_{t \geq 1}$. The timing of events is as follows. At the beginning of each period t , the bank decides whether (and if so, with which probability) to liquidate the firm, in which case it gets the scrap value S and compensates the agent with a transfer $Q_t \geq 0$. Denote by $\alpha = \{\alpha_t(\sigma^t(\theta^t))\}_{t \geq 1}$ the liquidation probabilities and by $\mathbf{Q} = \{Q_t(\sigma^t(\theta^t))\}_{t \geq 1}$ the transfers from the lender to the borrower in case of liquidation. Then, if the firm is not liquidated, the bank chooses the amount of capital k_t it lends to the firm, and the borrower's repayment τ_t if the project is successful. Denote by $\mathbf{k} = \{k_t(\sigma^{t-1}(\theta^{t-1}))\}_{t \geq 1}$ the capital advancements and by $\boldsymbol{\tau} = \{\tau_t(\sigma^{t-1}(\theta^{t-1}), \sigma_t(\theta_t))\}_{t \geq 1}$ the contingent payments from the borrower to the lender in case of success (there is no transfer in case of failure). The firm is restricted at all times to have a nonnegative cash flow, i.e., the following limited-liability constraint must be satisfied: $\tau_t(\sigma^{t-1}(\theta^{t-1}), \theta_{(2)}) \leq R(k_t(\sigma^{t-1}(\theta^{t-1})))$ for all t, θ^{t-1} . The outcome θ_t of the project is then realized and privately observed by the borrower, who sends a report $\sigma_t(\theta^t)$ and transfers $\tau_t(\sigma^t(\theta^t))$ to the bank in case of a (truthfully reported) success.

We define "equity" as the entrepreneur's share of the total firm's value, and "debt" as the lender's share. That is, equity (denoted by $V_t(\{\mathbf{k}, \boldsymbol{\tau}, \boldsymbol{\alpha}, \mathbf{Q}, \boldsymbol{\sigma}\}, \sigma^{t-1}(\theta^{t-1}))$) and debt (denoted by $B_t(\{\mathbf{k}, \boldsymbol{\tau}, \boldsymbol{\alpha}, \mathbf{Q}, \boldsymbol{\sigma}\}, \sigma^{t-1}(\theta^{t-1}))$) are the expected discounted cash flows (or continuation values) accruing to the borrower and the lender, respectively, under the contract $\{\mathbf{k}, \boldsymbol{\tau}, \boldsymbol{\alpha}, \mathbf{Q}, \boldsymbol{\sigma}\}$ and reporting strategy $\boldsymbol{\sigma}$, given the history of reports $\sigma^{t-1}(\theta^{t-1})$. In particular, the value of equity corresponds to the promised utility variable (11) in the taste shock model of Section 2.

This setup is formally similar to that in Section 2.3 and can be analyzed using the same recursive techniques. Specifically, we write the problem in recursive form using the value of equity v as the state variable. We can show that the set of continuation values v that can be supported by a feasible contract (see footnote 7) is $[0, \infty)$. A constrained-efficient contract maximizes the value obtained by the lender, $B(v_0)$, in the space of incentive-compatible and feasible contracts, subject to delivering some utility $v_0 \geq 0$ to the entrepreneur. The pair $(v_0, B(v_0))$ defines the capital structure of the firm (equity and debt) and implies a total value for the firm $W(v_0) = v_0 + B(v_0)$.

Denote by $W(\cdot)$ the value of the firm prior to the liquidation decision, and by $\hat{W}(\cdot)$ the value of the firm conditional on not being liquidated. Following the steps of Proposition 3, we

obtain that the latter is given by the following Bellman equation:

$$\hat{W}(\hat{v}) = \max_{\{k, \tau, \bar{w}\} \geq 0} (\pi R(k) - k) + \beta [(1 - \pi) W(w(\theta_{(1)})) + \pi W(w(\theta_{(2)}))]$$

subject to the promise-keeping constraint:

$$\hat{v} = \pi (R(k) - \tau) + \beta [(1 - \pi) w(\theta_{(1)}) + \pi w(\theta_{(2)})], \quad (130)$$

the incentive compatibility constraint in the high state:

$$R(k) - \tau + \beta w(\theta_{(2)}) \geq R(k) + \beta w(\theta_{(1)}), \quad (131)$$

and the limited liability constraint

$$\tau \leq R(k). \quad (132)$$

The liquidation decision of the firm can be formalized as follows. At the beginning of the period, the firm is liquidated with probability α , and the borrower receives Q , and it is kept in operation with probability $1 - \alpha$, in which case the borrower receives the continuation value \hat{v} . The value function $W(\cdot)$ then solves the following Bellman equation:

$$W(v) = \max_{\alpha \in [0, 1], \{Q, \hat{v}\} \geq 0} \alpha S + (1 - \alpha) \hat{W}(\hat{v})$$

subject to

$$v = \alpha Q + (1 - \alpha) \hat{v}.$$

Clementi and Hopenhayn (2006) show that the solution to this problem is as follows. If the equity v is large enough, then the policy of providing the unconstrained efficient level of capital k^* in every period is both feasible and incentive compatible. The minimum value v^* for which this is the case is given by the solution to the problem:

$$\begin{aligned} v^* &\equiv \min_{\tau, \{w(\theta)\}_{\theta \in \Theta}} (\pi R(k^*) - \tau) + \beta [(1 - \pi) w(\theta_{(1)}) + \pi w(\theta_{(2)})] \\ &\text{subject to } R(k^*) - \tau + \beta w(\theta_{(2)}) \geq R(k^*) + \beta w(\theta_{(1)}), \\ &\tau \leq R(k^*), \quad w(\theta_{(1)}) \geq v^*, \quad w(\theta_{(2)}) \geq v^*. \end{aligned}$$

We get $v^* \equiv \frac{1}{1-\beta} \pi R(k^*)$. Next, we can show that there exists a value $v_* \in (0, v^*)$ such that:

(i) when $v \geq v^*$, the firm's value $W(v)$ is equal to W^* . Letting $k_v = k^*$ at any future date, with $\tau_v = 0$ and $w_v(\theta_{(j)}) = v^*$ for all $j \in \{1, 2\}$, is optimal.

(ii) when $v \in [v_*, v^*)$, the policy function is as follows: $\alpha(v) = 0$, $k_v < k^*$, $\tau_v = R(k_v)$, and $w_v(\theta_{(1)}) < v < w_v(\theta_{(2)})$. The values $w_v(\theta_{(1)})$, $w_v(\theta_{(2)})$ are given as a function of k_v by the promise-keeping and incentive compatibility constraints (130), (131), which both hold with

equality.⁴⁶ Moreover, k_v is increasing in v for v close enough to v^* , k_v is decreasing in v for v close enough to v_* , $w_v(\theta_{(1)})$, $w_v(\theta_{(2)})$ are increasing in v , and equity is a submartingale, i.e. $v < \mathbb{E}[w_v]$. The value function $W(v)$ is strictly increasing and concave.

(iii) when $v < v_*$, the firm is liquidated with positive probability $\alpha(v) = 1 - v/v_*$ and transfer $Q = 0$, and continues at value $\hat{v} = v_*$ with probability $(1 - \alpha(v))$. The firm's value is equal to $W(v) = \alpha(v)S + (1 - \alpha(v))\hat{W}(v_*)$.

This characterization of the optimal contract has the following interpretation and implications. The contract determines stochastic processes for the firm size k_t , equity v_t , and debt $B(v_t) = W(v_t) - v_t$. Specifically, consider an entrepreneur who starts with equity $v_0 \in (v_*, v^*)$. Starting from this region, a good shock raises the value of equity to $w_v(\theta_{(2)}) > v$, and a bad shock reduces it. The submartingale property (which follows from equation (130)) implies that the equity v_t of surviving firms on average increases over time, and the monotonicity of the functions $w_v(\theta)$ implies the this process v_t displays persistence. Eventually, equity reaches either the lower threshold v_* (after a series of negative shocks), leading to the region where it is optimal to liquidate the firm with positive probability, or the upper threshold v^* (after a series of positive shocks), at which the incentive constraints no longer bind and the unconstrained efficient level of capital k^* is advanced from then on. There are therefore two absorbing states: either the firm is liquidated, or it attains its efficient size. In the transition, the transfer τ_v in the event of a good shock is set equal to the maximum possible amount $R(k_v)$. This is because the bank and the firm are both risk-neutral, so that it is optimal to backload the distribution of dividends to the borrower (by choosing the highest possible value of transfers τ and raising $w_v(\theta_{(j)})$ accordingly) in order to allow the equity to reach v^* as fast as possible. Finally, when v^* is attained, the firm's future cash flows are $v^* = W(v^*) + \frac{k^*}{1-\beta}$, and the lender's continuation value is $B(v^*) = -\frac{k^*}{1-\beta}$. This means that the entrepreneur has accumulated assets at the bank (at the interest rate r such that $\beta = \frac{1}{1+r}$) up to the positive balance $k^*/(1-\beta)$ while his payments were being postponed and all the cash flows were received by the lender; this balance is exactly enough to finance the project at the efficient scale from then on.

Next, the optimal contract shows that when equity is below the threshold, the amount of capital advanced by the bank is strictly smaller than the unconstrained efficient level: $k_v < k^*$. We can interpret this result as an (endogenous) borrowing constraint to which the entrepreneur is subject. Moreover, if v is close enough to v^* , higher equity relaxes the borrowing constraint and allows the entrepreneur to finance the project on a more efficient scale, as k_v is increasing in v .⁴⁷ Such financing frictions arise endogenously in the optimal contract due to moral hazard. To provide incentives for the successful entrepreneur to truthfully report the (good) outcome of his project, the optimal contract requires the borrower's compensation to be sensitive to reported

⁴⁶If these values are such that $w_v(\theta_{(2)}) > v^*$, then other values for the transfer τ_v (along with $w_v(\theta_{(2)})$) are also optimal.

⁴⁷The authors also run numerical simulations and find that for all v , $k_{w_v(\theta_{(1)})} < k_v$ and $k_{w_v(\theta_{(2)})} > k_v$, so that the amount of capital advanced by the lender increases after a success, and decreases after a decline.

output, which necessitates a spread $w_v(\theta_{(2)}) - w_v(\theta_{(1)})$ between the future equity values in the successful vs. unsuccessful states. Moreover, advancing more capital today tightens the incentive constraint (as the borrower will have to repay more in case of success, since $\tau = R(k)$) and thus requires a larger spread between future continuation values. But this spread is costly, because the marginal revenue is decreasing in capital and hence the firm's total value $W(\cdot)$ is concave. Therefore, the trade-off between higher capital and profits today against a lower firm's value in future periods implies an inefficient level of financing $k_v < k^*$ in the optimal contract.

These results imply that revenue shocks affect the financial structure $(v, B(v))$ of the firm, and yield rich implications for firm dynamics (size, growth and survival probability). Defining the firm's size as the level of capital k_t invested in the project, and investment as $k_t - k_{t-1}$, and simulating a calibrated version of the model, the authors obtain the following testable predictions. First, firm age and size are positively correlated. Second, the mean and variance of growth decrease with size and age. Third, the survival probability $\mathbb{P}(T > t | v)$, where T is the stopping time for exit, increases with the value of equity v and thus with age. The hazard rates for exit follow an inverted U-shaped function of age, as it takes a few periods for young firms to reach the liquidation region from their initial value v_0 , and a selection effect implies that older (surviving) firms have on average higher values and hence lower hazard rates. All these properties are consistent with the empirical evidence on firm dynamics (see the references in Clementi and Hopenhayn (2006) for a survey of the empirical literature). Finally, the authors argue that simulated data generated using the policy functions of the model would reproduce the empirical prediction that investment responds positively to innovations in the cash-flow process, and that the sensitivity of investment to cash-flows decreases with the age and size of the firm. Importantly, in the model, the financing frictions (borrowing constraints) arise *endogenously* as a feature of the optimal contract.

Optimal capital structure. We now describe another application of recursive contracts to corporate finance in a continuous-time framework using the techniques described in Section 3.3.2, following a simple version of De Marzo and Sannikov (2006).⁴⁸ In their model the agent (firm) can unobservably divert cash flows for its private benefit; investors control its wage and choose when to liquidate the project. While Clementi and Hopenhayn (2006) focused on the importance of informational frictions for firm investment and growth as a function of the the history of profit realizations (so that the scale, i.e. the capital, of the firm is an endogenous part of the optimal incentive contract), the closely related framework of De Marzo and Sannikov (2006) assume instead that the firm has a fixed size and are concerned with the optimal choice of the firm's capital structure.⁴⁹ Specifically, they propose an implementation of the optimal contract using simple financial instruments. This implementation is composed of a combination of long term debt with a constant coupon, a credit line, and equity. In this implementation the firm is

⁴⁸A discrete time version of this problem has been analyzed by DeMarzo and Fishman (2007b).

⁴⁹DeMarzo et al. (2012) extend this model to include investment and non-constant firm size.

compensated by holding a fraction of the equity, and defaults if debt service payments are not made or the credit line is overdrawn; dividends are paid when cash flows exceed debt payments and the credit line is paid off. This analysis can therefore help understand the choice between various forms of borrowing for firms, in particular the characteristics of credit line contracts, an empirically important component of firm financing. Finally, as we saw in Section 3.3.2, setting the model in continuous time allows the authors to obtain both a clean characterization of the optimal contract through an ordinary differential equation, as well as analytical comparative statics of the optimal contract with respect to the parameters of the model.

We now turn to a formal description of the model. An agent manages a project that generates stochastic cash flows given by:

$$d\hat{y}_t = (\mu - \theta_t) dt + \sigma dZ_t,$$

where Z_t is a standard Brownian motion, and $\theta_t \geq 0$ is the agent's private action, which can be interpreted as cash flow diversion. This unobserved diversion generates private benefit to the agent at rate $\lambda\theta_t$, with $\lambda \in (0, 1]$. The principal observes only the reported cash flows $\{\hat{y}_t\}_{t \geq 0}$.⁵⁰ The principal and the agent are risk-neutral and discount the future at rate r and γ respectively, with $\gamma > r$. The project requires external capital of $I_0 \geq 0$ to be started. The principal offers a contract (\mathbf{c}, τ) that specifies the agent's compensation $dc_t \geq 0$ for all t and a termination date τ , as functions of the histories $\{\hat{y}_s\}_{s \leq t}$. In the event of termination, the agent gets his outside option $R \geq 0$ and the principal receives the liquidation payoff $L \geq 0$.

The optimal contract maximizes the principal's expected profit subject to delivering expected utility \hat{v}_0 to the agent and the incentive compatibility constraints. We can show that in the optimal contract we have $\theta_t = 0$ for all $t \geq 0$. The problem is similar to that analyzed in Section 3.3.2 and can be expressed as:

$$\max_{\mathbf{c}, \tau} \mathbb{E}^{\theta=0} \left[\int_0^\tau e^{-rt} (d\hat{y}_t - dc_t) + e^{-r\tau} L \right]$$

subject to the promise-keeping constraint:

$$\hat{v}_0 = \mathbb{E}^{\theta=0} \left[\int_0^\tau e^{-\gamma t} dc_t + e^{-\gamma\tau} R \right]$$

and the incentive compatibility constraints:

$$\hat{v}_0 \geq \mathbb{E}^{\hat{\theta}} \left[\int_0^\tau e^{-\gamma t} (dc_t + \lambda \hat{\theta}_t dt) + e^{-\gamma\tau} R \right],$$

for any deviation strategy $\hat{\theta}$.

Following identical steps as in Section 3.3.2 (see in particular Proposition 11), we find that

⁵⁰De Marzo and Sannikov (2006) consider a more general model in which the agent can secretly save and thus over-report, i.e. $\theta_t < 0$, but show that in the optimal contract the agent always chooses to maintain zero savings.

there is a one-to-one correspondence between incentive-compatible contracts (\mathbf{c}, τ) and controlled processes (with controls (c_t, β_t))

$$dv_t = \gamma v_t dt - dc_t + \beta_t (d\hat{y}_t - \mu dt), \quad (133)$$

where the sensitivity of the agent's promised value to his report satisfies $\beta_t \geq \lambda$ for all $t \leq \tau$. The termination time τ is the earliest time that the agent's promised value v_t reaches R . The one-shot incentive constraint (Proposition 12) here says that truth-telling is incentive compatible if and only if $\beta_t \geq \lambda$ for all t , since the agent has incentives not to steal cash flows if he gets at least λ of promised value for each reported dollar.

De Marzo and Sannikov (2006) characterize the optimal contract as follows. Denote by $K(v)$ the principal's value function. It is easy to see that the optimal contract must satisfy $K'(v) \geq -1$ for all v . This is because the principal can always give to the agent with current promised utility v a lump-sum transfer $dc > 0$ and then revert to the optimal contract with utility $v - dc$, so that $K(v) \geq K(v - dc) - dc$. Defining \bar{v} as the lowest value such that $K'(\bar{v}) = -1$, it is optimal to keep the agent's promised utility in the range $[R, \bar{v}]$ and to set $dc_t = (v - \bar{v}) \mathbb{I}_{\{v \geq \bar{v}\}}$. The function $K(v)$ can then be characterized recursively as in Section 3.3.2. The Hamilton-Jacobi-Bellman equation is

$$rK(v) = \max_{\beta \geq \lambda} \mu + \gamma v K'(v) + \frac{1}{2} \beta^2 \sigma^2 K''(v),$$

$$\text{with } K(v) = K(\bar{v}) - (v - \bar{v}) \text{ for } v > \bar{v},$$

with the following value-matching, smooth-pasting, and super-contact conditions

$$K(R) = L, \quad K'(\bar{v}) = -1, \quad K''(\bar{v}) = 0.$$

The function $K(\cdot)$ is concave so that it is optimal to set $\beta_t = \lambda$ for all t . The optimal contract (with $\hat{v}_0 \in [R, \bar{v}]$) is such that v_t evolves according to (133) with $dc_t = 0$ when $v_t \in [R, \bar{v}]$. If $v_t = \bar{v}$, payments dc_t cause v_t to reflect at \bar{v} . The contract is terminated at time τ when v_t reaches R .⁵¹

De Marzo and Sannikov (2006) propose an implementation of the optimal contract using equity, long-term debt D , and a credit line C^L . If the agent defaults on a debt coupon payment or his credit balance exceeds C^L , the project is terminated. The idea behind this implementation is to map the interval of continuation values $[R, \bar{v}]$ into a credit line, with point \bar{v} corresponding to balance 0. From (133), we can write the evolution of the credit balance $\lambda^{-1}(\bar{v} - v_t)$ (where

⁵¹In the discrete-time setting described the previous section (based on the work of Clementi and Hopenhayn (2006)), allowing for randomization over the decision to terminate the project could improve the contract. De Marzo and Sannikov (2006) show that in the continuous-time framework, such randomization is not necessary: without loss of generality the termination time τ is based only on the firm's (reported) past performance.

λ is simply a normalization) as

$$d\left(\frac{\bar{v} - v_t}{\lambda}\right) = -d\hat{y}_t + \left\{ \gamma \left(\frac{\bar{v} - v_t}{\lambda}\right) dt + \left(\mu - \frac{\gamma}{\lambda}\bar{v}\right) dt + \frac{dc_t}{\lambda} \right\}.$$

The first term in the right hand side of this expression, $-d\hat{y}_t$, is the credit balance reduction due to the cash flows, where each dollar of cash flow subtracts exactly one dollar from the credit line balance. The next three terms in the right hand side (inside the brackets) are the three components that compose the implementation of the contract. The first term inside the brackets is the interest charged on the credit balance $\lambda^{-1}(\bar{v} - v_t)$, so that the implementation of the optimal contract has a credit line $C^L = \lambda^{-1}(\bar{v} - R)$, up to which credit is available to the firm at interest rate γ . The second term inside the brackets is the coupon rD on long-term debt, so that the face value of the debt is $D = r^{-1}(\mu - \gamma\bar{v}/\lambda)$. Finally the third term inside the brackets consists of the dividend payments made by the firm, i.e. the equity. The agent gets a fraction λ of the dividends dc_t , while outside investors hold the remaining firm's equity, debt and credit line. Cash flows in excess of the debt coupon payments are issued as dividends once the credit line is fully repaid. Termination occurs when the credit line balance reaches the credit limit C^L . Observe that the balance on the credit line fluctuates with the past performance of the firm, in particular leverage decreases with its profitability since the firm pays off the credit line when it makes profits.

De Marzo and Sannikov (2006) further analyze this optimal capital structure, i.e. how the amount of long-term debt and the size of the credit line depend on the parameters of the model, by deriving analytical comparative statics using the techniques described in Section 3.3.2. We refer the reader to the original paper for an in-depth analysis of these questions.

4.3 Development Economics

There is a large literature that studies informal insurance arrangements in the context of village economies. An early work of Townsend (1994), for example, showed that in rural India idiosyncratic variation in consumption is systematically related to idiosyncratic variation in income, implying that households can only achieve partial insurance against their idiosyncratic risks. Models of limited commitment developed by Thomas and Worrall (1988), Kehoe and Levine (1993), Kocherlakota (1996), Alvarez and Jermann (2000), Ligon et al. (2000), Ligon et al. (2002) can potentially explain these observations. In these models, all the information is public (there is no *information* friction); instead there is an *enforcement* friction: agents are free to walk away from the insurance contract at any time. Nevertheless, these models can be analyzed using the same recursive techniques as those described in Section 2. Analogous to the asymmetric information models we analyzed, the state variable is the utility promised to the agent. The only formal difference is that the incentive compatibility constraints (8) are replaced by participation constraints that we formally define in equation (134).

Here we describe the two-sided limited commitment framework analyzed by Ligon et al. (2002). The (observable) period- t state of nature $\theta_t \in \Theta = \{\theta_{(1)}, \dots, \theta_{(|\Theta|)}\}$ is stochastic and follows a Markov process with transition probability $\pi(\theta_{(i)} | \theta_{(j)}) > 0$ for all i, j . There are two agents with period- t utilities $U^1(c_t^1), U^2(c_t^2)$ and exogenous non-storable endowments (y_t^1, y_t^2) determined by θ_t . At least one of the two households is risk averse, and they both discount the future at rate β . A risk-sharing contract τ specifies for every date t and history θ^t a (possibly negative) transfer $\tau_t(\theta^t)$ from household 1 to household 2. A first-best, or full risk-pooling, contract τ is such that for all dates and states, the ratio of marginal utilities $\frac{U^{2'}(y_t^2(\theta_t) + \tau_t(\theta^t))}{U^{1'}(y_t^1(\theta_t) - \tau_t(\theta^t))}$ is constant across all histories and dates, so that each individual's consumption is only a function of the aggregate endowment.

The key friction of the model is that agents can walk away from the insurance contract, after which both households consume at autarky levels forever after, i.e. $\tau_t(\theta^t) = 0$ for all t, θ^t . Household $j \in \{1, 2\}$ has no incentive to break the contract if the following *sustainability constraint* holds: for all $\theta^t \in \Theta^t$,

$$U^j(c_t^j(\theta^t)) + \mathbb{E}_t \left[\sum_{s=1}^{\infty} \beta^s U^j(c_{t+s}^j(\theta^{t+s})) \right] \geq U^j(y_t^j(\theta_t)) + \mathbb{E}_t \left[\sum_{s=1}^{\infty} \beta^s U^j(y_{t+s}^j(\theta_{t+s})) \right], \quad (134)$$

where $c_s^1(\theta^s) = y_s^1(\theta_s) - \tau_s(\theta^s)$ and $c_s^2(\theta^s) = y_s^2(\theta_s) + \tau_s(\theta^s)$ for all s , and where \mathbb{E}_t is the expectation conditional on θ^t . As in Section 3.2 (where the government was unable to commit), it is useful to describe the present environment with bilateral lack of commitment as a repeated game between the two agents. Since reversion to autarky is the worst subgame-perfect punishment, there is a one-to-one relationship between sustainable contracts and subgame perfect equilibria (see Abreu (1988)). We now show how to characterize the set of constrained-efficient sustainable contracts, using recursive arguments formally similar to those we used in Section 2. By the Markov structure of the income process and the forward-looking nature of the constraints, the set of sustainable continuation contracts (i.e., for which conditions (134) are satisfied) depends only on the current state θ and not on the history leading to that state. The constrained efficient allocations maximize the expected lifetime utility of agent 2 subject to both sustainability constraints (134), and to delivering at least a given utility level v^1 to the agent 1, given that the current state is θ . Before we formally write this problem, we describe the space of discounted expected utilities v^1, v^2 for each agent (defined as in (11)) for which there exists a sustainable contract that delivers those values, given that the current state is θ (see footnote 7). We can show that this set is an interval of the form $[\underline{v}^j(\theta), \bar{v}^j(\theta)]$ for each agent $j \in \{1, 2\}$, where the minimum sustainable utilities when the current state is θ are

$$\underline{v}^j(\theta) = U^j(y^j(\theta)) + \mathbb{E} \left[\sum_{s=1}^{\infty} \beta^s U^j(y_s^j(\theta_s)) \mid \theta \right]$$

for $j \in \{1, 2\}$, that is, the value of autarky for agent j from state θ onwards.

The ex-post efficiency frontier, calculated once the current state θ is known, can then be characterized in recursive form as follows: for $v^1 \in [\underline{v}^1(\theta), \bar{v}^1(\theta)]$,

$$V(v^1, \theta) = \max_{\tau(\theta), \{w^1(\theta')\}_{\theta' \in \Theta}} U^2(y^2(\theta) + \tau(\theta)) + \beta \sum_{\theta' \in \Theta} \pi(\theta' | \theta) V(w^1(\theta'), \theta')$$

subject to to the promise-keeping constraint

$$U^1(y^1(\theta) - \tau(\theta)) + \beta \sum_{\theta' \in \Theta} \pi(\theta' | \theta) w^1(\theta') = v^1, \quad (135)$$

the sustainability constraints

$$w^1(\theta') \geq \underline{v}^1(\theta'), \quad \forall \theta', \quad (136)$$

$$V(w^1(\theta'), \theta') \geq \underline{v}^2(\theta'), \quad \forall \theta', \quad (137)$$

(the constraint $w^1(\theta') \leq \bar{v}^1(\theta')$ is equivalent to (137)), and the non-negativity constraints

$$y^1(\theta) - \tau(\theta) \geq 0, \quad \text{and} \quad y^2(\theta) + \tau(\theta) \geq 0. \quad (138)$$

The Lagrange multiplier λ associated with the constraint (135) is the key variable in the analysis of optimal insurance contracts. The first-order conditions and envelope condition of the problem imply that λ is related to the ratio of the marginal utilities of consumption by

$$\lambda = -\frac{\partial}{\partial v} V(v^1, \theta) = \frac{U^{2'}(y^2(\theta) + \tau(\theta))}{U^{1'}(y^1(\theta) - \tau(\theta))} + \frac{\psi_2 - \psi_1}{U^{1'}(y^1(\theta) - \tau(\theta))}, \quad (139)$$

where ψ_1, ψ_2 are the Lagrange multipliers associated with the non-negativity constraints (138). and other multipliers for the evolution. Suppose that the value of λ is known. If λ is in the set of marginal utility ratios which can be generated by feasible transfers in state θ (i.e., by $\tau(\theta) \in [-y_2(\theta), y_1(\theta)]$), then there is a unique interior solution and the value of the transfer $\tau(\theta)$ is pinned down by equation (139) with $\psi_1 = \psi_2 = 0$. Otherwise, there is a corner solution with all income going to one of the households, i.e., $\tau(\theta) \in \{-y_2(\theta), y_1(\theta)\}$ (with a positive multiplier ψ_2 or ψ_1). Therefore the constrained efficient contracts can be fully characterized by the evolution of the multiplier $\lambda(\theta^t)$ along with an initial value $\lambda_0 \in [\min_{\theta} \{\underline{\lambda}_{\theta}\}, \max_{\theta} \{\bar{\lambda}_{\theta}\}]$ (varying the value of λ_0 traces out the Pareto frontier). This can be easily done using the first-order conditions with respect to $w^1(\theta')$, which writes, for all $\theta' \in \Theta$, $-\frac{\partial}{\partial v} V(w^1(\theta'), \theta') = \frac{\lambda + \chi_1(\theta')}{1 + \chi_2(\theta')}$, where $\beta\pi(\theta' | \theta)\chi_1(\theta')$ and $\beta\pi(\theta' | \theta)\chi_2(\theta')$ are the multipliers associated with the constraints (136), (137). Ligon et al. (2002) show that there exist $|\Theta|$ intervals $[\underline{\lambda}_{\theta}, \bar{\lambda}_{\theta}]$, such that $\lambda(\theta^t)$ evolves

according to

$$\lambda(\theta^t, \theta_{t+1}) = \begin{cases} \underline{\lambda}_{\theta_{t+1}}, & \text{if } \lambda(\theta^t) < \underline{\lambda}_{\theta_{t+1}}, \\ \lambda(\theta^t), & \text{if } \lambda(\theta^t) \in [\underline{\lambda}_{\theta_{t+1}}, \bar{\lambda}_{\theta_{t+1}}], \\ \bar{\lambda}_{\theta_{t+1}}, & \text{if } \lambda(\theta^t) > \bar{\lambda}_{\theta_{t+1}}. \end{cases} \quad (140)$$

Specifically, $\underline{\lambda}_\theta \equiv -\frac{\partial}{\partial v} V(\underline{v}^1(\theta), \theta)$ and $\bar{\lambda}_\theta \equiv -\frac{\partial}{\partial v} V(\bar{v}^1(\theta), \theta)$ where $\bar{v}^1(\theta)$ is the maximum feasible expected value for agent 1, which satisfies $V(\bar{v}^1(\theta), \theta) = \underline{v}^2(\theta)$.

To understand the intuition underlying this result, suppose for simplicity that the non-negativity constraints on consumption (138) never bind, i.e., $\psi_1 = \psi_2 = 0$. We already argued that in a full risk-pooling contract, the current transfers in every period are chosen such that the ratio of the two households' marginal utilities (139) is constant. Now consider a constrained-efficient contract, where the evolution of this ratio is given by equation (140). Suppose that the marginal utility ratio last period was $\lambda(\theta^t)$, and that the current state is $\theta_{t+1} = \theta'$, which defines an interval of possible marginal utility ratios $[\underline{\lambda}_{\theta'}, \bar{\lambda}_{\theta'}]$. If $\lambda(\theta^t) \in [\underline{\lambda}_{\theta'}, \bar{\lambda}_{\theta'}]$, then we choose $\tau(\theta^t)$ so that $\lambda(\theta^t, \theta') = \lambda(\theta^t)$. If instead $\lambda(\theta^t) < \underline{\lambda}_{\theta'}$ (respectively, if $\lambda(\theta^t) > \bar{\lambda}_{\theta'}$), household 1 (resp., household 2) would want to break the contract if the ratio of marginal utilities remained constant, as the short-term costs of making the corresponding transfer in the current period would exceed the long-term insurance benefits coming from promises of future reciprocation. Hence full risk-pooling, which would occur with complete markets, is not feasible in this case. We then choose $\lambda(\theta^t, \theta') = \underline{\lambda}_{\theta'}$ (resp., $\lambda(\theta^t, \theta') = \bar{\lambda}_{\theta'}$). The value $\lambda = \underline{\lambda}_{\theta'}$ (respectively, $\lambda = \bar{\lambda}_{\theta'}$) corresponds to household 1 receiving its minimum possible sustainable surplus $\underline{v}^1(\theta')$ in state θ' (resp., its maximum surplus $\bar{v}^1(\theta')$), or equivalently household 2 getting $\bar{v}^2(\theta')$ (resp., $\underline{v}^2(\theta')$). In other words, if full risk-sharing is not possible the ratio of marginal utilities must change to an endpoint (i.e., by the minimum possible amount) so that one of the households is just indifferent between staying in the contract and renegeing.⁵²

Ligon et al. (2002) then test the model on the data for three Indian villages, using the model to predict consumption allocations (by estimating empirically the initial ratio of marginal utilities and values for the model's parameters that provide the best fit to the data), and measuring the difference between these predictions and the actual data. They find that the dynamic limited commitment model does a substantially better job at explaining the dynamic response of consumption to income than do models of full insurance, static limited commitment, or autarky.

In models of limited commitment, the key to the amount of informal insurance that can be provided in the optimal contract depends on how costly renegeing is for the households. That is, the value of autarky is the most important determinant of the extent of insurance. Recent

⁵²We can show that for a sufficiently high discount factor $\beta \geq \beta^* \in [0, 1)$ the λ -intervals overlap and thus there is some first-best contract which is sustainable, whereas if the households are sufficiently impatient, i.e. $\beta \leq \beta^* \in (0, 1)$, then no non-autarkic contract exists. In the former case, irrespective of initial value of λ_0 , and hence of the initial division of the surplus from the contract, the contract converges with probability one to a first-best contract. Thus, if people are sufficiently patient, absence of commitment cannot justify the observed lack of diversification in individual consumption as being efficient.

work by Morten (2013) studies a model of risk sharing with endogenous commitment in which temporary migration is possible. The possibility of migration has the unintended consequence of improving self insurance of individuals and the value of autarky, and worsening the risk sharing in the economy. She studies the joint determination of risk sharing and migration decisions and decomposes the welfare effects of migration between changes in income and changes in the endogenous structure of insurance. Morten (2013) further structurally estimates the model on a panel from rural India and argues that the possibility of migration may significantly reduce risk sharing.

There is by now a large literature studying the predictions of models with contracting frictions in the development economics context. For example, Karaivanov and Townsend (2014) is a comprehensive study comparing exogenously incomplete markets to markets which are endogenously incomplete due to contractual frictions. Their focus is on consumption, income, investment and asset behavior of small businesses in Thailand. They conclude that the exogenously incomplete market model has the best fit for their rural sample while the dynamic moral hazard model is more appropriate for urban households. A recent paper by Kinnan (????) develops a test to distinguish barriers to informal insurance in Thai villages for three types of models: limited commitment, moral hazard, and hidden income, based on the theoretical prediction (see, e.g., equation (64)) that a single lag of inverse marginal utility is sufficient to forecast current inverse marginal utility, which is satisfied by the first two models but not the latter. She concludes that hidden income is more likely to be the cause of barriers to insurance.

4.4 International Finance

In this section, we describe an application of the recursive contract models to the international finance context, based on Kehoe and Perri (2002). The benchmark model is one of limited commitment similar to that studied in the previous section, but we now analyze it using the duality theory described in Section 3.1.4. Models of limited commitment are useful to analyze questions related to sovereign debt default as they provide a framework that can explain the mechanisms by which countries are induced to participate in contracts involving transfers backed only by promises of future repayment, i.e., without a legal authority enforcing them. In such models, countries are free to renege on their debts; the only threat is exclusion from future participation in the financial market.

Standard international business cycle models with either complete or exogenously incomplete markets typically deliver predictions that are at odds with the data (see Backus et al. (1992)), for instance, that cross-country correlations of consumption are much higher than those for output, and that both employment and investment in different countries comove negatively. Moreover, net exports and investment are much more volatile in these models than in the data. Kehoe and Perri (2002) show that introducing endogenously incomplete markets due to limited loan enforcement frictions in an otherwise standard international business cycle model can resolve

these puzzles. This feature allows the model to reproduce the data's positive cross-country comovements of factors of production, consumption and output.

We now formally describe the model. There are two countries $i = 1, 2$ that produce using domestic labor and capital inputs and face exogenous idiosyncratic Markov technology shocks $A_i(\theta^t)$. Output in country i after a history of shocks θ^t is given by $F(k_i(\theta^{t-1}), A_i(\theta^t) l_i(\theta^t))$. The social planner's problem consists of choosing allocations $\{c_i(\theta^t), l_i(\theta^t), k_i(\theta^{t-1})\}_{i,t}$ to maximize a weighted (with weights λ_i) sum of utilities of the representative consumers in each country:

$$\max_{\mathbf{c}, \mathbf{l}, \mathbf{k}} \sum_{i=1,2} \lambda_i \left\{ \sum_{t=0}^{\infty} \sum_{\theta^t \in \Theta^t} \beta^t \pi(\theta^t) U(c_i(\theta^t), l_i(\theta^t)) \right\} \quad (141)$$

subject to the feasibility constraint

$$\sum_{i=1,2} (c_i(\theta^t) + k_i(\theta^t)) = \sum_{i=1,2} [F(k_i(\theta^{t-1}), A_i(\theta^t) l_i(\theta^t)) + (1 - \delta) k_i(\theta^{t-1})],$$

and the enforcement constraints (similar to (134)): for $i = 1, 2$,

$$\sum_{s=t}^{\infty} \sum_{\theta^s \geq \theta^t} \beta^{s-t} \pi(\theta^s | \theta^t) U(c_i(\theta^s), l_i(\theta^s)) \geq \underline{V}_i(k_i(\theta^{t-1}), \theta^t), \quad (142)$$

where $\underline{V}_i(k_i(\theta^{t-1}), \theta^t)$ denotes country i 's value of autarky from θ^t onward, given by

$$\underline{V}_i(k_i(\theta^{t-1}), \theta^t) = \max_{\mathbf{c}, \mathbf{l}, \mathbf{k}} \sum_{s=t}^{\infty} \sum_{\theta^s} \beta^{s-t} \pi(\theta^s | \theta^t) U(c_i(\theta^s), l_i(\theta^s)) \quad (143)$$

$$\text{subject to } c_i(\theta^s) + k_i(\theta^s) \leq F(k_i(\theta^{s-1}), A_i(\theta^s) l_i(\theta^s)) + (1 - \delta) k_i(\theta^{s-1}).$$

The enforcement constraints are formally derived from arguments similar to those we used to obtain (82) in Section 3.2. They ensure that it is the best response for each country to stick to their equilibrium strategies.

We can re-write this problem recursively using the Marcet and Marimon (2015) approach (see Section 3.1.4). Letting $\beta^t \pi(\theta^t) \mu_i(\theta^t)$ denote the multipliers on the enforcement constraints (142), we can write the Lagrangian of the social planner's problem as (similar derivation as equation (76))

$$\sum_{t=0}^{\infty} \sum_{\theta^t \in \Theta^t} \sum_{i=1,2} \beta^t \pi(\theta^t) \left\{ M_i(\theta^{t-1}) U(c_i(\theta^t), l_i(\theta^t)) + \mu_i(\theta^t) [U(c_i(\theta^t), l_i(\theta^t)) - \underline{V}_i(k_i(\theta^{t-1}), \theta^t)] \right\} \quad (144)$$

subject to the feasibility constraint, where $M_i(\theta^t)$ is a cumulative Lagrange multiplier defined

recursively as

$$M_i(\theta^t) = M_i(\theta^{t-1}) + \mu_i(\theta^t), \quad (145)$$

for $t \geq 0$, with $M_i(\theta^{-1}) = \lambda_i$. Thus the cumulative multiplier $M_i(\theta^t)$ is equal to the original planning weight λ_i at time 0, and to the sum of the initial weight and the past multipliers on the enforcement constraints at time $t \geq 1$ and history θ^t . Using the techniques described in Section 3.1.4 and denoting by $z(\theta^t) = \frac{M_2(\theta^t)}{M_1(\theta^t)}$ the relative weight on country 2, this problem can be written recursively and its solution is stationary in the state space that consists of the current shock, the current capital stocks, and the relative weight, i.e. $x_t = (\theta_t, k_1(\theta^{t-1}), k_2(\theta^{t-1}), z(\theta^{t-1}))$.

It is instructive to compare this objective (144) with the unconstrained objective (141). The enforcement constraints introduce three key differences. First, starting at the beginning of the period, the cumulative Lagrange multiplier $M_i(\theta^{t-1})$ shifts the (relative) weights of each agent. Second, the current Lagrange multiplier $\mu_i(\theta^t)$ on the sustainability constraint further changes the weight on current consumption (as well as on future consumption by affecting the future cumulative multiplier $M_i(\theta^t)$). These two forces translate in the first-order conditions into a distortion of the relative marginal utilities of consumption:

$$\frac{U_{1c}(\theta^t)}{U_{2c}(\theta^t)} = \frac{M_2(\theta^{t-1}) + \mu_2(\theta^t)}{M_1(\theta^{t-1}) + \mu_1(\theta^t)}. \quad (146)$$

Third, accumulating more capital $k_i(\theta^{t-1})$ tightens the enforcement constraint by increasing the value of autarky. As a result, the Euler equation (and capital accumulation) is distorted as follows:

$$U_{ic}(\theta^t) = \beta \sum_{\theta_{t+1}} \pi(\theta_{t+1} | \theta_t) \left[\frac{M_i(\theta^{t+1})}{M_i(\theta^t)} U_{ic}(\theta^{t+1}) (F_{ik}(\theta^{t+1}) + 1 - \delta) - \frac{\mu_i(\theta^{t+1})}{M_i(\theta^t)} V_{ik}(\theta^{t+1}) \right]. \quad (147)$$

The last first-order condition writes $\frac{U_{il}(\theta^t)}{U_{ic}(\theta^t)} = F_{il}(\theta^t)$: there is no distortion in the consumption-labor decision, since this margin does not affect the enforcement constraint.⁵³ These first-order conditions along with the transition law for $z_2(\theta^t)$ can be straightforwardly rewritten as functions of $z_2(\theta^{t-1})$ and the normalized multipliers $\tilde{\mu}_i(\theta^t) \equiv \frac{\mu_i(\theta^t)}{M_i(\theta^{t-1})}$. The solution to this problem can then be characterized by allocations of the form $(c_i(x_t), l_i(x_t), k_i(x_t))$, where the state vector is $x_t = (\theta_t, k_1(\theta^{t-1}), k_2(\theta^{t-1}), z(\theta^{t-1}))$. These policy functions satisfy the first-order conditions above, the feasibility and enforcement constraints, and the complementary slackness conditions on the multipliers.

The model has the following implications. Suppose that the home country (say, country $i = 1$) is hit in history (t, θ^t) with a positive and persistent productivity shock $A_1(\theta^t) > 0$. Equation (143) shows that such a shock increases the home country's value of autarky, and thus

⁵³Kehoe and Perri (2004) show how to decentralize the constrained efficient allocation as a competitive equilibrium using a tax on capital income to replicate the wedge in the Euler equation (147) generated by the enforcement constraint.

tightens its enforcement constraint (142). This may lead the enforcement constraint to bind, which translates into a positive multiplier $\mu_1(\theta^t)$ in the first-order condition (146). This in turn implies that the planner increases the relative weight to the home country in its objective and allocates it higher consumption $c_1(\theta^t)$ (i.e., lower marginal utility $U_{1c}(\theta^t)$) to prevent it from defaulting. Moreover, this increase in consumption is persistent, because the productivity shock is persistent and the positive multiplier $\mu_1(\theta^t)$ raises the cumulative multiplier $M_1(\theta^s)$ of the home country (defined in (145)) in all future periods $s \geq t$. In contrast, consumption in the foreign country does not vary much, as risk-sharing in this economy is limited. Finally, the planner optimally restricts the investment flow into country 1 in order to reduce the home country's future value of autarky in equation (147) and relax the enforcement constraint. It also increases labor effort and investment in the foreign country to raise country 1's value of participating into the contract, leading to positive cross-country correlations of investment and employment and to a trade surplus (positive net exports) in the home country.

Now compare these effects with those that would occur in an economy without enforcement frictions, i.e., with complete markets. In response to a positive productivity shock in the home country, and hence a higher productivity of capital and labor, the planner optimally increases the domestic labor effort and the capital stock, both by saving more and increasing investment flowing from abroad. In contrast, foreign labor effort and investment decrease. Moreover, because of risk sharing, the domestic economy shares its consumption gains, leading to an increase in the consumption of the foreign country. The responses are qualitatively similar but muted in a model where markets are exogenously incomplete (only bonds are allowed). In such models, therefore, output is less correlated across countries than is consumption, the cross-country correlations of investment and employment are negative, and a positive productivity shock leads to a trade deficit in the home country (due to the net inflow of investment).

Kehoe and Perri (2002) calibrate the economy and analyze numerically these implications of the model with endogenously incomplete markets. They find that it matches the data's positive cross-country comovements of factors of production (employment, investment) and the cross-country comovements of consumption and output. This resolves several of the puzzles arising in standard (complete or exogenously incomplete market-)models of international finance described in the first paragraph of this section.

There is a large literature that analyzes questions of international debt and sovereign default using models of (one- or two-sided) limited commitment. The seminal paper is Eaton and Gersovitz (1981), and this literature has been comprehensively reviewed by Aguiar and Amador (2013). In particular, Aguiar et al. (2009) analyze the behavior of sovereign debt and foreign direct investment in a small open economy (rather than in a two-country general equilibrium environment as analyzed in the previous paragraphs) where the government lacks commitment (leading to potential default and expropriation of capital) and is more impatient than the market. While the standard one-sided limited commitment model (see Thomas and Worrall (1994)) predicts that the government will eventually accumulate enough assets to overcome its commit-

ment problem,⁵⁴ the additional assumption of a higher degree of impatience (and hence, the combination of front loading due to impatience and back loading due to limited commitment) leads to cycles in both sovereign debt and foreign direct investment, as well as a “debt overhang” effect whereby investment is distorted by more in recessions than in booms.

5 Conclusion

The theory of recursive contracts underpins a variety of applications in a range of fields, from public finance to development economics. A unifying feature of these applications is that they feature frictions such as unobservability of shocks or actions or non-enforceability of contracts that endogenously limit the amount of risk sharing and insurance that can be achieved. This chapter provides a self-contained treatment of the fundamental techniques and the more advanced topics of recursive contracts. We also survey a number of applications through the lens of this unified theoretical treatment that illustrate the versatility of the theoretical apparatus.

References

- ABRAHAM, A. AND N. PAVONI (2008): “Efficient allocations with moral hazard and hidden borrowing and lending: A recursive formulation,” *Review of Economic Dynamics*, 11, 781–803.
- ABREU, D. (1988): “On the theory of infinitely repeated games with discounting,” *Econometrica*, 383–396.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a theory of discounted repeated games with imperfect monitoring,” *Econometrica*, 1041–1063.
- ACEMOGLU, D., M. GOLOSOV, AND A. TSYVINSKI (2008): “Political economy of mechanisms,” *Econometrica*, 76, 619–641.
- (2011): “Power fluctuations and political economy,” *Journal of Economic Theory*, 146, 1009–1041.

⁵⁴One-sided limited commitment models generally imply that the optimal contract features a form *back loading*: the profile of consumption is shifted towards the future. The intuition is as follows. Additional consumption in a particular period helps ensure the agent’s participation in the contract. Moreover, it also helps satisfy the enforcement constraints in all previous periods as well, since the left-hand side of the enforcement constraint (e.g., (142)) is forward-looking. At the margin, therefore, consumption in the future is preferable as it relaxes all the preceding participation constraints. As a result the relevant Euler equation includes the cumulative sums of Lagrange multipliers that take into account all of the binding constraints in the previous periods. When the government and the market have the same degree of impatience, the economy will eventually achieve perfect risk sharing with constant consumption, so that a country has an incentive to save to grow out of the enforcement constraints if it is patient enough. Ray (2002) shows that the backloading result and eventual reaching of the unconstrained allocations apply in very general settings.

- AGUIAR, M. AND M. AMADOR (2013): “Sovereign debt,” *chapter prepared for the Handbook of International Economics*, 4.
- AGUIAR, M., M. AMADOR, AND G. GOPINATH (2009): “Investment cycles and sovereign debt overhang,” *The Review of economic studies*, 76, 1–31.
- AIYAGARI, S. R., A. MARCET, T. J. SARGENT, AND J. SEPPÄLÄ (2002): “Optimal taxation without state-contingent debt,” *Journal of Political Economy*, 110, 1220–1254.
- ALBANESI, S. (2011): “Optimal taxation of entrepreneurial capital with private information,” *Working paper*.
- ALBANESI, S. AND C. SLEET (2006): “Dynamic optimal taxation with private information,” *Review of Economic Studies*, 73, 1–30.
- ALBUQUERQUE, R. AND H. HOPENHAYN (2004): “Optimal lending contracts and firm dynamics,” *Review of Economic Studies*, 71, 285–315.
- ALVAREZ, F. AND U. J. JERMANN (2000): “Efficiency, equilibrium, and asset pricing with risk of default,” *Econometrica*, 775–797.
- ATKESON, A. AND R. E. LUCAS (1992): “On efficient distribution with private information,” *Review of Economic Studies*, 59, 427–453.
- (1995): “Efficiency and equality in a simple model of efficient unemployment insurance,” *Journal of Economic Theory*, 66, 64–88.
- BACKUS, D. K., P. J. KEHOE, AND F. E. KYDLAND (1992): “International real business cycles,” *Journal of political Economy*, 745–775.
- BADEL, A. AND M. HUGGETT (2014): “Taxing top earners: a human capital perspective,” *Working paper*.
- BARRO, R. J. (1979): “On the determination of the public debt,” *The Journal of Political Economy*, 940–971.
- BATTAGLINI, M. AND R. LAMBA (2015): “Optimal dynamic contracting: the first-order approach and beyond,” *Working paper*.
- BENVENISTE, L. M. AND J. A. SCHEINKMAN (1979): “On the differentiability of the value function in dynamic models of economics,” *Econometrica*, 47, 727–32.
- BERTSEKAS, D. P., A. NEDI, AND A. E. OZDAGLAR (2003): *Convex analysis and optimization*, Athena Scientific.

- BESTER, H. AND R. STRAUZ (2001): “Contracting with imperfect commitment and the revelation principle: The single agent case,” *Econometrica*, 69, 1077–98.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J.-C. ROCHET (2007): “Dynamic security design: Convergence to continuous time and asset pricing implications,” *The Review of Economic Studies*, 74, 345–390.
- BIAIS, B., T. MARIOTTI, J.-C. ROCHET, AND S. VILLENEUVE (2010): “Large risks, limited liability, and dynamic moral hazard,” *Econometrica*, 73–118.
- BILLINGSLEY, P. (1995): *Probability and measure*, Wiley Series in Probability and Statistics, Wiley.
- BISMUT, J.-M. (1973): “Conjugate convex functions in optimal stochastic control,” *Journal of Mathematical Analysis and Applications*, 44, 384–404.
- (1978): “An introductory approach to duality in optimal stochastic control,” *SIAM review*, 20, 62–78.
- CAVES, R. E. (1998): “Industrial organization and new findings on the turnover and mobility of firms,” *Journal of economic literature*, 1947–1982.
- CBO (2007): “Historical effective federal tax rates, 1979 to 2005,” *Congressional Budget Office*.
- CHAMBERLAIN, G. AND C. A. WILSON (2000): “Optimal intertemporal consumption under uncertainty,” *Review of Economic dynamics*, 3, 365–395.
- CHARI, V. V. AND P. J. KEHOE (1990): “Sustainable plans,” *Journal of Political Economy*, 98, 783–802.
- (1993): “Sustainable plans and debt,” *Journal of Economic Theory*, 61, 230–261.
- CLEMENTI, G. L. AND H. A. HOPENHAYN (2006): “A theory of financing constraints and firm dynamics,” *The Quarterly Journal of Economics*, 121, 229–265.
- COLE, H. AND F. KUBLER (2012): “Recursive contracts, lotteries and weakly concave pareto sets,” *Review of Economic Dynamics*, 15, 479–500.
- CVITANIC, J. AND J. ZHANG (2012): *Contract theory in continuous-time models*, Springer Science & Business Media.
- DE MARZO, P. AND Y. SANNIKOV (2006): “Optimal security design and dynamic capital structure in a continuous-time agency model,” *The Journal of Finance*, 61, 2681–2724.
- DEMARZO, P. M. AND M. J. FISHMAN (2007a): “Agency and optimal investment dynamics,” *Review of Financial Studies*, 20, 151–188.

- (2007b): “Optimal long-term financial contracting,” *Review of Financial Studies*, 20, 2079–2128.
- DEMARZO, P. M., M. J. FISHMAN, Z. HE, AND N. WANG (2012): “Dynamic agency and the q theory of investment,” *The Journal of Finance*, 67, 2295–2340.
- DIAMOND, P. AND J. MIRRLEES (1978): “A model of social insurance with variable retirement,” *Journal of Public Economics*, 10, 295–336.
- (1986): “Payroll-tax financed social insurance with variable retirement,” *The Scandinavian Journal of Economics*, 88, 25–50.
- DIAMOND, P. A., L. J. HELMS, AND J. A. MIRRLEES (1980): “Optimal taxation in a stochastic economy: A Cobb–Douglas example,” *Journal of Public Economics*, 14, 1–29.
- DOVIS, A., M. GOLOSOV, AND A. SHOURIDEH (2015): “Political economy of sovereign debt: Cycles of debt crisis and inequality overhang,” *Working paper*.
- EATON, J. AND M. GERSOVITZ (1981): “Debt with potential repudiation: Theoretical and empirical analysis,” *The Review of Economic Studies*, 289–309.
- EKELAND, I. AND J. A. SCHEINKMAN (1986): “Transversality conditions for some infinite horizon discrete time optimization problems,” *Mathematics of operations research*, 11, 216–229.
- ESPINO, E., J. KOZLOWSKI, AND J. M. SANCHEZ (2013): “Too big to cheat: efficiency and investment in partnerships,” *FRB of St. Louis Working Paper No 2013-001C*.
- FARHI, E., C. SLEET, I. WERNING, AND S. YELTEKIN (2012): “Non-linear capital taxation without commitment,” *Review of Economic Studies*, 79, 1469–1493.
- FARHI, E. AND I. WERNING (2007): “Inequality and social discounting,” *Journal of Political Economy*, 115, 365–402.
- (2012): “Capital taxation: Quantitative explorations of the inverse Euler equation,” *Journal of Political Economy*, 120, 398–445.
- (2013): “Insurance and taxation over the life cycle,” *Review of Economic Studies*, 80, 596–635.
- FAZZARI, S. M., R. G. HUBBARD, B. C. PETERSEN, A. S. BLINDER, AND J. M. POTERBA (1988): “Financing constraints and corporate investment,” *Brookings Papers on Economic Activity*, 141–206.
- FERNANDES, A. AND C. PHELAN (2000): “A recursive formulation for repeated agency with history dependence,” *Journal of Economic Theory*, 91, 223–247.

- FREIXAS, X., R. GUESNERIE, AND J. TIROLE (1985): “Planning under incomplete information and the ratchet effect,” *Review of Economic Studies*, 52, 173–91.
- FRIEDMAN, M. (1957): *A theory of the consumption function*, National Bureau of Economic Research, Inc.
- FUDENBERG, D. AND J. TIROLE (1991): *Game theory*, MIT Press.
- GEWEKE, J. AND M. KEANE (2000): “An empirical analysis of earnings dynamics among men in the PSID: 1968-1989,” *Journal of Econometrics*, 96, 293–356.
- GOLOSOV, M. AND L. IOVINO (2014): “Social insurance, information revelation, and lack of commitment,” *NBER Working paper No. w20633*.
- GOLOSOV, M., N. KOCHERLAKOTA, AND A. TSYVINSKI (2003): “Optimal indirect and capital taxation,” *Review of Economic Studies*, 70, 569–587.
- GOLOSOV, M., M. TROSHKIN, AND A. TSYVINSKI (2016): “Redistribution and social insurance,” *American Economic Review*, *forthcoming*.
- GOLOSOV, M. AND A. TSYVINSKI (2006): “Designing optimal disability insurance: A case for asset testing,” *Journal of Political Economy*, 114, 257–279.
- (2007): “Optimal taxation with endogenous insurance markets,” *The Quarterly Journal of Economics*, 122, 487–534.
- GOLOSOV, M., A. TSYVINSKI, AND I. WERNING (2006): “New dynamic public finance: A user’s guide,” *NBER Macroeconomics Annual*, 21, 317–363.
- GREEN, E. J. (1987): “Lending and the smoothing of uninsurable income,” in *Contractual Arrangements for Intertemporal Trade*, ed. by E. C. Prescott and N. Wallace, Minneapolis: University of Minnesota Press.
- GUVENEN, F., S. OZKAN, AND J. SONG (2014): “The nature of countercyclical income risk,” *Journal of Political Economy*, 122, 621–660.
- GUVENEN, F., J. SONG, S. OZKAN, AND F. KARAHAN (2015): “What do data on millions of US workers reveal about life-cycle earnings risk?” *NBER Working Paper No. w20913*.
- HALL, R. E. (1978): “Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence,” *The Journal of Political Economy*, 86, 971–987.
- HE, Z. (2009): “Optimal executive compensation when firm size follows geometric brownian motion,” *Review of Financial Studies*, 22, 859–892.

- HOSSEINI, R., L. E. JONES, AND A. SHOURIDEH (2013): “Optimal contracting with dynastic altruism: family size and per capita consumption,” *Journal of Economic Theory*, 148, 1806–1840.
- HURWICZ, L. (1960): “Optimality and informational efficiency in resource allocation processes,” in *Mathematical methods in the social sciences, 1959: proceedings of the first Stanford symposium*, Stanford University Press, 27.
- (1972): “On informationally decentralized systems,” in *Decision and organization: a volume in honor of Jacob Marschak*, North-Holland.
- KAPIČKA, M. (2013): “Efficient allocations in dynamic private information economies with persistent shocks: A first order approach,” *Review of Economic Studies*, 80, 1027–1054.
- KARAIVANOV, A. AND R. M. TOWNSEND (2014): “Dynamic financial constraints: Distinguishing mechanism design from exogenously incomplete regimes,” *Econometrica*, 82, 887–959.
- KARATZAS, I. AND S. SHREVE (2012): *Brownian motion and stochastic calculus*, vol. 113, Springer Science & Business Media.
- KEHOE, P. J. AND F. PERRI (2002): “International business cycles with endogenous incomplete markets,” *Econometrica*, 70, 907–928.
- (2004): “Competitive equilibria with limited enforcement,” *Journal of Economic Theory*, 119, 184–206.
- KEHOE, T. J. AND D. K. LEVINE (1993): “Debt-constrained asset markets,” *The Review of Economic Studies*, 865–888.
- KINNAN, C. (????): “Distinguishing barriers to insurance in Thai villages,” *Working paper*.
- KOCHERLAKOTA, N. (2010): *The New Dynamic Public Finance*, Princeton University Press, USA.
- KOCHERLAKOTA, N. R. (1996): “Implications of efficient risk sharing without commitment,” *The Review of Economic Studies*, 63, 595–609.
- LAFFONT, J.-J. AND J. TIROLE (1988): “The dynamics of incentive contracts,” *Econometrica*, 56, 1153–75.
- LE VAN, C. AND H. C. SAGLAM (2004): “Optimal growth models and the Lagrange multiplier,” *Journal of Mathematical Economics*, 40, 393–410.
- LIGON, E., J. P. THOMAS, AND T. WORRALL (2000): “Mutual insurance, individual savings, and limited commitment,” *Review of Economic Dynamics*, 3, 216–246.

- (2002): “Informal insurance arrangements with limited commitment: Theory and evidence from village economies,” *The Review of Economic Studies*, 69, 209–244.
- LINDBECK, A. AND J. W. WEIBULL (1987): “Balanced-budget redistribution as the outcome of political competition,” *Public choice*, 52, 273–297.
- LJUNGQVIST, L. AND T. SARGENT (2012): *Recursive Macroeconomic Theory*, MIT Press.
- LOCKWOOD, B. B., C. G. NATHANSON, AND E. G. WEYL (2014): “Taxation and the allocation of talent,” *Working paper*.
- LUENBERGER, D. (1969): *Optimization by vector space methods*, Wiley-Interscience.
- MAAG, E., C. E. STEUERLE, R. CHAKRAVARTI, AND C. QUAKENBUSH (2012): “How marginal tax rates affect families at various levels of poverty,” *National Tax Journal*, 65, 759–82.
- MARCET, A. AND R. MARIMON (2015): “Recursive contracts,” *Mimeo, European University Institute*.
- MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic theory*, Oxford University Press, New York.
- MESSNER, M. AND N. PAVONI (2004): “On the recursive saddle point method,” *Dynamic Games and Applications*, 1–13.
- MESSNER, M., N. PAVONI, AND C. SLEET (2012): “Recursive methods for incentive problems,” *Review of Economic Dynamics*, 15, 501–525.
- (2014): “The dual approach to recursive optimization: Theory and examples,” in *2014 Meeting Papers*, Society for Economic Dynamics, 1267.
- MIAO, J. AND Y. ZHANG (2014): “A duality approach to continuous-time contracting problems with limited commitment,” *Journal of Economic Theory*, 159, 929–928.
- MILGROM, P. AND I. SEGAL (2002): “Envelope theorems for arbitrary choice sets,” *Econometrica*, 70, 583–601.
- MORTEN, M. (2013): “Temporary migration and endogenous risk sharing in village india,” *Working paper*.
- MYERSON, R. B. (1981): “Optimal auction design,” *Mathematics of operations research*, 6, 58–73.
- (1982): “Optimal coordination mechanisms in generalized principal-agent problems,” *Journal of Mathematical Economics*, 10, 67–81.
- (1986): “Multistage games with communication,” *Econometrica*, 323–358.

- ØKSENDAL, B. (2003): *Stochastic differential equations*, Springer.
- ØKSENDAL, B. K. AND A. SULEM (2005): *Applied stochastic control of jump diffusions*, vol. 498, Springer.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic mechanism design: A Myersonian approach,” *Econometrica*, 82, 601–653.
- PHELAN, C. (1995): “Repeated moral hazard and one-sided commitment,” *Journal of Economic Theory*, 66, 488–506.
- PHELAN, C. AND R. M. TOWNSEND (1991): “Computing multi-period, information-constrained optima,” *Review of Economic Studies*, 58, 853–81.
- RAY, D. (2002): “The time structure of self-enforcing agreements,” *Econometrica*, 70, 547–582.
- REVUZ, D. AND M. YOR (1999): *Continuous martingales and Brownian motion*, vol. 293, Springer Science & Business Media.
- ROGERSON, W. P. (1985): “Repeated moral hazard,” *Econometrica*, 53, 69–76.
- RUDIN, W. (1976): *Principles of mathematical analysis*, New York: McGraw-Hill Book Co., third ed.
- RUSTICHINI, A. (1998): “Lagrange multipliers in incentive-constrained problems,” *Journal of Mathematical Economics*, 29, 365–380.
- SANNIKOV, Y. (2008): “A continuous-time version of the principal-agent problem,” *The Review of Economic Studies*, 75, 957–984.
- (2014): “Moral hazard and long-run incentives,” .
- SCHEUER, F. AND A. WOLITZKY (2014): “Capital taxation under political constraints,” .
- SHOURIDEH, A. (2010): “Optimal taxation of capital income: A Mirrleesian approach to capital accumulation,” *Working paper*.
- SLEET, C. AND S. YELTEKIN (2008): “Politically credible social insurance,” *Journal of Monetary Economics*, 55, 129–151.
- SPEAR, S. AND S. SRIVASTAVA (1987): “On repeated moral hazard with discounting,” *The Review of Economic Studies*, 54, 599–617.
- STANTCHEVA, S. (2014): “Optimal taxation and human capital policies over the life cycle,” *Working paper*.

- STOKEY, N. L., R. E. LUCAS, AND E. C. PRESCOTT (1989): *Recursive methods in economic dynamics*, Cambridge, MA: Harvard University Press.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): “Consumption and risk sharing over the life cycle,” *Journal of Monetary Economics*, 51, 609–633.
- SUN, Y. (2006): “The exact law of large numbers via Fubini extension and characterization of insurable risks,” *Journal of Economic Theory*, 126, 31–69.
- THOMAS, J. AND T. WORRALL (1988): “Self-enforcing wage contracts,” *The Review of Economic Studies*, 55, 541–554.
- (1990): “Income fluctuation and asymmetric information: An example of a repeated principal-agent problem,” *Journal of Economic Theory*, 51, 367–390.
- (1994): “Foreign direct investment and the risk of expropriation,” *The Review of Economic Studies*, 81–108.
- TOWNSEND, R. M. (1994): “Risk and insurance in village India,” *Econometrica*, 539–591.
- UHLIG, H. (1996): “A law of large numbers for large economies,” *Economic Theory*, 8, 41–50.
- VICKREY, W. (1947): *Agenda for progressive taxation*, The Ronald Press Company.
- WERNING, I. (2009): “Nonlinear capital taxation,” *MIT working paper*.
- WILLIAMS, N. (2009): “On dynamic principal-agent problems in continuous time,” *Working paper*.
- (2011): “Persistent Private Information,” *Econometrica*, 79, 1233–1275.
- YONG, J. AND X. Y. ZHOU (1999): *Stochastic controls: Hamiltonian systems and HJB equations*, vol. 43, Springer Science & Business Media.
- ZHANG, Y. (2009): “Dynamic contracting with persistent shocks,” *Journal of Economic Theory*, 144, 635–675.