Implications of Uncertainty for Optimal Policies*

Todd Lensman† Maxim Troshkin‡

September 9, 2018

Abstract

This paper studies the implications of a broader view of uncertainty for social insurance and redistribution in otherwise conventional macro public finance environments, with heterogeneous agents and with private idiosyncratic shocks. We show that uncertainty manifests as endogenous lack of commitment on the part of the government, leading to the optimality of periodically-reformed policies. Periodic reforms imply simplified policies that are not fully state-contingent and at times lose full history dependence. Simplified policies can be characterized without complete backward induction to compute promised utilities when the time horizon is finite. However, linear policies can be far from optimal. We show that equilibria in decentralized versions of these economies are not generally efficient, implying a meaningful role for government provision of insurance, unlike in conventional environments without uncertainty.

Keywords: social insurance, redistribution, efficiency, competitive equilibria, risk sharing, insurance, uncertainty, ambiguity, model misspecification, robustness

JEL: H21, H11, E62, D8

1 Introduction

A sizable and growing literature shares the following approach to social insurance, redistribution, and normative questions more generally: Start with a friction, typically private information about idiosyncratic shocks, and characterize friction-constrained allocations that maximize an ex ante social objective, usually social welfare. The optimal policies are

*We benefited significantly from a discussion by Stefanie Stantcheva and from comments by Marco Battaglini, Marcus Berliant, Steve Coate, Lars Hansen, Narayana Kocherlakota, Alessandro Pavan, Aleh Tsyvinski, Florian Scheuer, and Iván Werning, as well as from the audiences at the NBER SI MPF meeting, the 4^{th} Taxation Theory Conference, LAEF at UC Santa Barbara, Cornell, Georgetown, Maryland, Montréal, Penn State, Rochester. We particularly thank John Zhu for early discussions that led to a more productive approach and key results. Early versions were circulated as "On the Optimality of Periodically Reformed, Simple Fiscal Policies".

†Cornell University, tal84@cornell.edu

‡Cornell University, troshkin@cornell.edu

then the ones that implement constrained-optimal allocations. A notable property of such policies is that they are generically complex, featuring significant nonlinearities and history dependence in dynamic contexts. In addition, the policy designer and agents are commonly assumed to know the data-generating processes with certainty, so optimal policies are designed once and maintained forever.[1]

Casual empiricism, however, points to real world policies that are at least somewhat incomplete, history independent, and periodically reformed. More systematically, growing evidence suggests that idiosyncratic shock distributions change significantly and often.[2] One implication is that certainty about the data-generating processes in the economy is a strong assumption when designing policies. It further suggests that social contracts that are potentially incomplete (not fully state-contingent) and periodically renegotiated could be better suited for understanding optimal policies.

This paper moves away from the assumption of certainty about future distributions of idiosyncratic shocks, and characterizes general properties of optimal policies that are robust with respect to incomplete knowledge of stochastic elements of the economy. To do this, we allow a broader view of uncertainty. The agents in the economy face both risk in the conventional sense of stochastic, heterogeneous skills, as well as (Knightian or model) uncertainty in the sense that agents entertain multiple possible distributions of future skills, commonly referred to as beliefs, models, or priors. The approach we take to modeling risk and uncertainty, with aversion to both, follows the approaches in macroeconomics and finance (e.g., Hansen and Sargent 2001, Epstein and Schneider 2003).[3]

The economy we consider is otherwise conventional, populated by heterogeneous agents who experience idiosyncratic and potentially persistent skill shocks over time. The data-generating process for skills is arbitrary, and it is not known to anyone in the economy. As a result, each agent forms a set of distributions that he believes may represent the data-generating process for skills in the next period. The economy has a government that seeks to provide social insurance against risk and against uncertainty, as well as a degree

[1]For reviews of applications to optimal dynamic fiscal policies see, e.g., Golosov, Tsyvinski, and Werning (2006), Kocherlakota (2010). For an example of this approach to dynamic social insurance more broadly see, e.g., Williams (2011).

[2]Recent administrative data evidence suggests that the distribution of pre-tax earnings in many developed countries has been going through dramatic changes that appear irregular, frequent, and apparently unanticipated by the governments (see, e.g., Piketty, Saez, and Zucman 2018). Related evidence suggests that people hold inconsistent and at times distribution-incompatible beliefs about their future productivities, tax liabilities, etc. (see, e.g., Aghion, Akcigit, Lequien, and Stantcheva 2017). In addition, a large empirical literature documents substantial uncertainty about both macro and micro variables (see, e.g., Bloom 2014).

[3]See also Hansen, Sargent, Turmuhambetova, and Williams (2006). Recent studies in those literatures have focused on showing that uncertainty can help explain the behavior of economic aggregates to a surprising degree, e.g., Bhandari, Borovička, and Ho (2017).

of redistribution. It is constrained by the same lack of certainty about the distribution of future skills, so the government is not an abstract entity with perfect knowledge of the data-generating processes. Rather, the government is interpreted concretely as having at best the information that all of the agents in the economy have combined.

We impose minimalistic assumptions on beliefs, and we keep the environment virtually agnostic about any learning that may map histories of observations and current distributions into updated beliefs about the future. That is, agents are allowed to be arbitrarily uncertain or certain about the next period's distribution of skills, with only the requirement that beliefs have sufficient overlap. To aid technical intuition, we demonstrate arguments using a common parametric formulation of repeated maxmin expected utility with arbitrary belief updating rules.[4]

To make it transparent that the main force behind the results is uncertainty, we first develop our results in a baseline finite-horizon environment with a finite number of agents, in which agents' heterogeneous skills and beliefs are publicly observable. We then prove a version of the revelation principle when skills and beliefs are private, and we re-establish our results when the government designs policies that incentivize truthful revelation of skills and beliefs by the agents. We also demonstrate that our results persist with an infinite time horizon and a continuum of agents.

We first show that it is optimal to periodically reform policies. Periodically-reformed policies are those for which the government finds it optimal to design allocations period by period and subsequently redesign them as needed after the economy realizes a new set of shocks. Specifically, suppose that agents' beliefs have sufficient overlap so that everyone agrees on at least some paths that are possible for the economy to follow. If this set of paths is sufficiently large, it will be optimal for the government to design a policy that assumes the economy will follow paths in this belief overlap. In subsequent periods, if the economy does not follow such a path, the government will find it optimal to reform the policy. In this sense, uncertainty with a sufficient belief overlap essentially presents itself in the government's problem as endogenous lack of commitment: Even though the government has the ability to fully commit, it may choose not to and instead design optimal allocations period by period, reforming them as necessary.

We then show that, unlike with exogenous lack of commitment, optimal policies can

[4]For a dynamically consistent representation of repeated maxmin expected utility preferences, see Epstein and Schneider (2003). Our results are readily generalizable to, e.g., dynamic Variational Preferences, and more generally would extend to a dynamic version of the Uncertainty Averse Preferences that unify many other specifications, including in particular the Multiplier Preferences of Hansen and Sargent (2001). See, e.g., Machina and Siniscalchi (2014) for a review of the interconnections among these preference representations.

be constructed in a simplified way and are simplified themselves, in particular by losing dependence on the full history of shocks. Because they are periodically reformed, optimal allocations can be constructed by solving reform problems we derive for each period, without first solving for a complete sequence of fully state-contingent constrained-efficient allocations. When a reform provides strict improvement to the previous policy, optimal allocations lose dependence on the full history of shocks. When the time horizon is finite, reform problems also imply that efficient allocations can be characterized without the full backward induction ordinarily required to compute promised utilities. Despite these simplifications, we argue that restrictive assumptions are required for linear or even affine policies to be optimal, e.g., independence of shock distributions and inelastic labor supply.

Finally, we show that in decentralized versions of the private information economy, competitive equilibria are not generally efficient. In other words, unlike in conventional environments, a broader view of uncertainty in our environment creates a potentially meaningful role for the government provision of insurance. In these decentralizations, agents contract with competitive firms to provide labor and capital in exchange for consumption, and both agents and firms are uncertain about the data-generating process. If firms form heterogeneous beliefs about the data-generating process, they will seek to trade state-contingent securities with other firms that have different beliefs. However, competitive markets will not permit trade in every state-contingent security because of private information, leading to less than efficient provision of insurance. That is, this additional heterogeneity naturally leads to market outcomes that can be strictly improved upon by the government. Nevertheless, we show that any insurance that decentralized economies do provide is again simplified and periodically reformed, not unlike what is observed in reality.

We discuss related literature next. Section 2 establishes our results in the baseline environment with uncertainty as the only friction. Section 3 considers private information. Section 4 shows that decentralized versions of the private information economy may not be efficient. Section 5 further examines the simplicity of optimal policies, characterizing conditions under which linear policies can be optimal. Throughout, we show the key steps of the arguments, while the Online Appendix contains the remaining details of the proofs.

Related Literature

We contribute most directly to a sizable and growing literature studying the design of social insurance and redistribution in dynamic economies (see, e.g., reviews in Golosov, Tsyvinski, and Werning 2006 and Kocherlakota 2010). It generally assumes perfect certainty of the data-generating process to characterize optimal policies when the government has few or no direct constraints on policy tools, but potentially faces informational or commitment frictions. A key lesson without uncertainty is that optimal dynamic policies are generically complex. For example, Farhi and Werning (2013) use the first-order approach to characterize complex dynamics of optimal income taxes over the life cycle. Recent contributions also introduce additional frictions or permit additional heterogeneity on the part of agents (e.g., Scheuer and Werning 2016, Stantcheva 2017, and Makris and Pavan 2018). However, many of these contributions also compute the welfare losses from restrictions on policy tools to argue that simpler policies can deliver welfare that closely approaches that of the optimal dynamic policies. Our results provide a theoretical foundation for simpler policies by showing that such policies can in fact be optimal when there is uncertainty about the data-generating process.

Parts of this broader literature focus on a government that has an inability to commit to its own policies and may seek to implement reforms. In recent examples, Farhi, Sleet, Werning, and Yeltekin (2012) and Golosov and Iovino (2016) approach this by constraining the government ex ante to choosing policies that it will not seek to reform later. We instead study conditions under which periodic reforms are optimal when the government has the ability to commit to a policy.

Closely related to the implementation of optimal policies, an important observation is that decentralized competitive equilibria result in (constrained) efficient allocations in many cases, even when agents have private information and markets are incomplete (e.g., Golosov and Tsyvinski 2007, Acemoglu and Simsek 2012). A common interpretation is that the only result of a government's social insurance policy is to crowd out insurance provided by private markets. Our results show that removing the assumption of certainty about the data-generating process and permitting firms and agents to react heterogeneously to this uncertainty can overturn this conclusion, creating a potentially meaningful role for the government provision of insurance.

The literature has also started to consider optimal policies when either the government or the agents are uncertain about the data-generating process. Kocherlakota and Phelan (2009) consider an endowment shock economy in which the government, but not the agents, is uncertain about the data-generating process. They derive conditions under which the government cannot improve on the competitive equilibrium allocation. We consider optimal policies in a dynamic production economy in which no one is certain, with a potentially beneficial role for government intervention. Bhandari (2015) studies properties of optimal risk-sharing arrangements between uncertain agents, using the multiplier preferences of Hansen and Sargent (2001) to characterize the optimal consumption path and long-run inequality. Kocherlakota and Song (2018) examine mechanisms for the provision of a public good in economies where agents are uncertain about the distribution of private valuations. Similarly to our results, they find that uncertainty can lead to simple implementations of efficient policies.

Methodologically, our formulation of the agents' preferences closely follows the recursive multiple-priors utility axiomatized by Epstein and Schneider (2003), and it includes a variation of Hansen and Sargent's (2001) constraint preferences as a special case. Our approach to characterizing optimal allocations for uncertainty averse agents follows Zhu (2016), who studies incomplete and affine contracts in the context of financial contracting. The formalism we use is distinct from but in the spirit of robust mechanism design (see, e.g., Bergemann and Morris 2013), in the sense that we seek to characterize policies that are robust with respect to misspecification of the environment and determine if such policies display any inherent simplicity. The particular assumption of common knowledge that we remove is that of the data-generating process. However, the social choice functions we focus on depend on agents' beliefs so that the government is not paternalistic, evaluating the policies in line with how the agents do. The economies we study are also dynamic and allow for private information.

2 Uncertainty as the Friction

To avoid confounded effects that cannot be attributed to different causes, we first consider a conventional dynamic heterogeneous-agents economy with public information. The only unconventional element - and the only friction in this baseline setup - is the assumption that no one in the economy knows with certainty the future distributions of idiosyncratic shocks. We define an efficient allocation as a solution to the problem of a government seeking to provide social insurance and a degree of redistribution. We then show that the government can achieve efficiency with policies that are both simplified and that can be obtained in a simplified way. Such policies are not fully state contingent, are periodically reformed, and lose dependence on the full history of shocks.

2.1 Baseline setup

The economy exists in discrete time, $t = 0$ to $t = T \leq \infty$, and is populated by a finite number of agents, indexed by $i \in \{1, \dots, N\}$.[5] At the beginning of $t = 0$, nature draws a sequence of idiosyncratic shocks $s_i^T \equiv (s_{i,0}, .., s_{i,T})$ for each agent, and reveals $s_{i,t}$ to agent

[5] We focus in the main text on finite countability of time periods as well as finite numbers of agents and possible shocks. Our results readily extend to infinite cases as we discuss in the Online Appendix.

i at the beginning of t. The set of possible shocks $s_{i,t}$ that could be realized at each t is finite. The data-generating process for the idiosyncratic shocks is arbitrary and is not known with certainty to anyone in the economy. In the baseline economy of this section, this lack of certainty is the only friction, so a shock becomes public information once it is revealed to an agent. Consequently, the vector of all agents' shock histories up to period t, $s^t \equiv (s_1^t, \ldots, s_N^t)$, is public information at the beginning of t. We call s^t the period-t *state* of the economy.

Each idiosyncratic shock has two components: a skill and a set of subjective beliefs, $s_{i,t} \equiv (\theta_{i,t}, \Pi_{i,t+1})$.[6] First, $\theta_{i,t}$ denotes agent i's idiosyncratic *skill* (productivity), so that if the amount of labor the agent exerts is $l_{i,t} \in [0, L]$, $L < \infty$, then the effective labor supplied is $z_{i,t} \equiv \theta_{i,t} l_{i,t}$. Assume that $\theta_{i,t}$ has finite support $\Theta \subset \mathbb{R}_+ - \{0\}$ for all t, and let $\underline{\theta}$ and $\bar{\theta}$ denote the minimum and maximum elements of Θ respectively. Let $\theta_i^t \equiv (\theta_{i,0}, \ldots, \theta_{i,t})$ denote a period-t history of skills.

Second, each shock contains a set of subjective *beliefs*, i.e., distributions over the next period's state of the economy that the agent thinks could represent the data-generating process. An agent's beliefs at t about the $t+1$ state are a non-empty set $\Pi_{i,t+1} \subset \Delta\left(s^{t+1}\right)$ of probability distributions $\pi_{i,t+1}$, where $\Delta\left(s^{t+1}\right)$ denotes the set of all probability distributions over states s^{t+1}. Let $\Pi_i^{t+1} \equiv (\Pi_{i,0}, \ldots, \Pi_{i,t+1})$ denote a period-$t+1$ history of beliefs.[7]

To examine properties common to social insurance and redistribution in general, without restricting attention to specific collections of policy tools, we focus on the allocations that a policy would deliver to agents: An *allocation*

$$C \equiv \left\{c_t\left(s^t\right), z_t\left(s^t\right), k_{t+1}\left(s^t\right)\right\}_{t=0}^{T}$$

is a sequence of N-vector valued consumption, effective labor, and capital functions that depend on the state s^t. For example, the period-t consumption of agent i is given by $c_{i,t}\left(s^t\right)$, the i^{th} entry of $c_t\left(s^t\right)$.

---

[6]Alternatively, we could specify idiosyncratic skill shocks $\theta_{i,t}$ and an arbitrary updating (and potentially learning) process by which beliefs $\Pi_{i,t+1}$ are formed. The exposition is significantly simplified by treating beliefs as part of the shock and imposing assumptions directly on the beliefs.

[7]A common example of this broader uncertainty (the model uncertainty of Hansen and Sargent 2001) supposes that in any period, any agent i has a particular statistical model in mind, i.e., a distribution $\pi^*_{i,t+1} \in \Delta\left(s^{t+1}\right)$ that he thinks may be the true distribution of $t+1$ types, often described as the "benchmark" or "approximating" model. The agent distrusts this model and considers other models $\pi_{i,t+1}$ that are "close to" $\pi^*_{i,t+1}$ in the sense of distance d on $\Delta\left(s^{t+1}\right)$, commonly taken to be relative entropy (expected log likelihood ratio), total variation, etc. Given a parameter $\epsilon \geq 0$ that governs how uncertain the agent is, his set of beliefs is $\Pi_{i,t+1} \equiv \left\{\pi_{i,t+1} \middle| d\left(\pi_{i,t+1}, \pi^*_{i,t+1}\right) \leq \epsilon\right\}$.

This formalism of $\Pi_{i,t+1}$ is intuitive but somewhat circular, given that $\Pi_{i,t+1}$ is a component of the state s^{t+1} and is also a set of distributions over s^{t+1}. Beliefs can be defined rigorously, e.g., using the construction of type spaces typically used in robust mechanism design (see Bergemann and Morris 2013).

Each agent is initially endowed with capital $k_0 > 0$. Output is produced using a constant returns to scale production function $f : \mathbb{R}_+^2 \to \mathbb{R}_+$ increasing in capital and effective labor. Given an initial state s_0, an allocation C is *feasible* if for every $t = 0$ and every $s^t \geq s_0$,[8]

$$\sum_i c_{i,t}\left(s^t\right) + K_{t+1}\left(s^t\right) \leq f\left(K_t\left(s^{t-1}\right), Z_t\left(s^t\right)\right),$$

where $K_t\left(s^{t-1}\right) \equiv \sum_i k_{i,t}\left(s^{t-1}\right)$ and $Z_t\left(s^t\right) \equiv \sum_i z_{i,t}\left(s^t\right)$ are the aggregate capital and effective labor functions, respectively. This aggregate ex post feasibility constraint must hold for any state s^t that could follow s_0 in order to account for idiosyncratic uncertainty about the data-generating process.

An agent's preferences are assumed to have a recursive representation with continuation utility

$$U_{i,t}\left(C\left|s^t\right.\right) \equiv u\left(c_{i,t}\left(s^t\right), \frac{z_{i,t}\left(s^t\right)}{\theta_{i,t}}\right) + \beta \inf_{\pi_{i,t+1}\in\Pi_{i,t+1}} \mathbb{E}_{\pi_{i,t+1}}\left[U_{i,t+1}\left(C\left|s^{t+1}\right.\right)\middle| s^t\right],$$

where $\beta \in (0,1)$ is the subjective discount factor, the von Neumann-Morgenstern utility function u is C^2 with $-u_c, u_l < 0$ and $u_{cc}, u_{ll} \leq 0$, and $\mathbb{E}_{\pi_{i,t+1}}$ denotes an expectation with respect to the belief $\pi_{i,t+1} \in \Pi_{i,t+1}$. When $t = T$, a $U_{i,t+1}$ term does not appear on the right side. Agents are thus averse to both risk originating from stochastic skills and uncertainty captured by multiple beliefs. The latter is usually interpreted as seeking to make choices that are robust with respect to the shock distribution: Instead of choosing what works best in a particular future scenario, agents choose what works decently in any scenario, which entails choosing what works best in the worst scenario.[9]

The economy has a government that seeks a degree of redistribution while providing social insurance. The government's problem is to maximize a weighted average of the agents' utilities, subject to feasibility and non-negativity of policy functions, where the weighting captures the redistribution motive and is given by a non-negative vector $\eta \in \mathbb{R}_+^N$. To focus on fiscal policies and more broadly on social insurance and redistribution, the agents are not allowed to leave the social contract designed by the government, i.e., they do not have access to an exogenous outside option. Given social welfare weights η and an initial state

[8]For any two states s^t and s^τ with $t \geq \tau$, we say that s^t *follows* s^τ, written $s^t \geq s^\tau$, if the first $\tau + 1$ components of s^t are s^τ. We use this ex post notion of feasibility for simplicity, but the results below readily generalize to settings with other feasibility constraints.

[9]For an axiomatic treatment and a characterization of the existence of a recursive representation of this form, see Epstein and Schneider (2003).

s_0, an allocation $C^*(s_0)$ is *efficient* if

$$C^*(s_0) \in \arg\max_C \sum_i \eta_i U_{i,0}(C\,|s_0) \tag{1}$$

subject to feasibility and non-negativity. Note that instead of an abstract entity with perfect knowledge of the data-generating process, the government concretely possesses the same information about realized types and the future state distribution as do all of the agents combined. Note also that C^* could be quite complex and fully history dependent, and it is typically designed once and maintained forever. We assume that the government has full commitment power in the sense that if it considers a new allocation in period $t \geq 0$, it must always deliver the period-t continuation utility promised to agents at t, regardless of any considered subsequent reforms.

2.2 Optimality of periodic reforms

We first show that under a mild condition on agents' beliefs, the government can achieve efficiency with simplified policies that are not fully contingent on future state realizations and are periodically reformed. The condition on agents' beliefs is, loosely speaking, that they have a sufficient overlap. The overlap is needed for agents to agree on a simplified policy as well as a fallback option that remains feasible even when the actual path of the economy falls outside of the belief overlap. One intuitive case is when all agents agree on the worst path for the economy. We first illustrate the result in this case and then discuss its generality.

To construct the "worst" beliefs that an agent could have in any period, let $\underline{\Pi}_T = \{\underline{\pi}_T\}$ denote the belief set containing only $\underline{\pi}_T$, a distribution that places unit weight on everyone realizing the skill $\underline{\theta}$ at T. Having this belief simply means the agent is certain at $T-1$ that everyone will realize the worst skill $\underline{\theta}$ at T. For $t \leq T-2$, let $\underline{\Pi}_{t+1} = \left\{\underline{\pi}_{t+1}\right\}$ denote the belief set containing only $\underline{\pi}_{t+1}$, a distribution that places unit weight the $t+1$ shock $\left(\underline{\theta}, \underline{\Pi}_{t+2}\right)$ for all agents. If an agent has beliefs $\underline{\Pi}_{t+1}$ at t, he is certain that everyone in the economy will realize the worst skill $\underline{\theta}$ at $t+1$. Moreover, the agent is certain that his pessimism will persist, so he is certain that everyone will realize $\underline{\theta}$ at $t+2, \ldots, T$.

Given any belief $\pi_{i,t+1} \in \Pi_{i,t+1}$, it will be convenient to let $\pi_{i,t+1}(\cdot|\, s^t)|_\theta$ denote the marginal distribution of $t+1$ skills $\theta_{t+1} \equiv (\theta_{1,t+1}, \ldots, \theta_{N,t+1})$, conditional on the state s^t. Similarly, let $\pi_{i,t+1}(\cdot|\, s^t)|_\Pi$ denote the conditional marginal distribution of $t+2$ beliefs $\Pi_{t+2} \equiv (\Pi_{1,t+2}, \ldots, \Pi_{N,t+2})$. The following is then one intuitive way to ensure that all agents agree on the worst path for the economy:

Assumption 1. *For any $t = 0, \ldots, T-1$, any i, any s^t, and any belief $\pi_{i,t+1} \in \Pi_{i,t+1}$, there exists another belief $\pi'_{i,t+1} \in \Pi_{i,t+1}$ such that*

$$\pi_{i,t+1}\left(\cdot|\, s^t\right)\big|_\theta = \pi'_{i,t+1}\left(\cdot|\, s^t\right)\big|_\theta$$

but $\pi'_{i,t+1}\left(\cdot|\, s^t\right)\big|_\Pi$ places unit weight on the event $\Pi_{j,t+2} = \underline{\Pi}_{t+2}$ for all j.[10]

That is, regardless of what an agent believes at any t about future skills, he believes it possible that everyone at $t+1$ could have the worst beliefs $\underline{\Pi}_{t+2}$. This implies that all of the agents agree on at least one particular path that the economy could take, i.e., where everyone realizes the worst shock $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ in $\tau = t+2, \ldots, T$. Note, however, that this does not require the data-generating process to actually place weight on the beliefs $\underline{\Pi}_{\tau+1}$. This also places no restrictions on agents' potentially heterogeneous beliefs about the distribution of the next period's skills. This assumption is the only condition we place on agents' beliefs and belief updating.[11]

Proposition 1. *An efficient allocation C^* can be implemented by a sequence of allocations $\{C^t\}_{t=0}^T$, where $C^t = \left\{c_\tau^t, z_\tau^t, k_{\tau+1}^t\right\}_{\tau=t}^{t+1}$.*

The implementation here is in the sense that every agent's $t = 0$ utility under the simplified allocation C^0 is equal to that under the potentially fully state-contingent allocation C^*: For all i,

$$U_{i,0}\left(C^0\, |s_0\right) = U_{i,0}\left(C^*\, |s_0\right).$$

While complete details of the proof are provided in the Online Appendix, the construction of the simplified allocations and the reasoning behind why they work are intuitive enough to provide here. At $t = 0$, each agent believes it is possible to realize a $t = 1$ state wherein everyone has the worst belief $\underline{\Pi}_2$ about the future path of the economy. In such states, each agent is certain that everyone will realize the worst skill $\underline{\theta}$ at all $t \geq 2$. The simplified allocation $C^0 \equiv \left\{c_t^0, z_t^0, k_{t+1}^0\right\}_{t=0}^1$ can be constructed based on what C^* prescribes when such states are realized at $t = 1$. Start with $t = 0$ by setting $\{c_0^0, z_0^0, k_1^0\} \equiv \{c_0^*, z_0^*, k_1^*\}$, so that C^0 and C^* coincide at $t = 0$. For $t = 1$, construct c_1^0, z_1^0, and k_2^0 so that regardless

[10] With a finite number of agents, the Law of Large Numbers does not imply that the actual data-generating distribution for the t state s^t is observable after shocks are realized in period t. As a result, the distribution $\underline{\pi}_{t+1}$ and the belief set $\underline{\Pi}_{t+1}$ are not necessarily unique. Assumption 1 assumes a particular sequence $\left\{\underline{\Pi}_{\tau+1}\right\}_{\tau=1}^T$.

[11] This is also satisfied by common specifications of agents' beliefs. For example, with the model uncertainty reinterpretation, it is sufficient that at any t the agent finds it too difficult to have models $\pi_{i,t+1}$ predict future models $\pi_{i,t+2}$ and so on, so the agent does not rule out any possible distribution over $t+2$ models. As such, for any $\pi_{i,t+1} \in \Delta\left(s^{t+1}\right)$ the distance function d takes into account only the marginal distribution over the physical part of the state, θ_{t+1}.

of their actual beliefs Π_2, each agent is allocated consumption, labor, and capital based on what c_1^*, z_1^*, and k_2^* prescribe when every agent has beliefs $\underline{\Pi}_2$. In other words, C^0 takes for granted that agents will realize the worst beliefs $\underline{\Pi}_2$. Thus C^0 is not a fully state-contingent allocation, but instead it depends only on the $t = 0$ state s_0 and the $t = 1$ skills θ_1.[12]

To see that at $t = 0$ all agents are indifferent between C^* and C^0, let $\Pi'_{i,1} \subset \Pi_{i,1}$ denote the set of beliefs $\pi'_{i,1}$ that place unit weight on the event in which all agents realize the worst beliefs $\underline{\Pi}_2$ at $t = 1$. Since agents are averse to uncertainty, the lowest expected utility over $\Pi'_{i,1}$ places a weak upper bound on the lowest expected utility over $\Pi_{i,1}$:

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[U_{i,1} \left(C^* \left| s^1 \right) \right| s_0 \right] \leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^* \left| s^1 \right) \right| s_0 \right].$$

By construction, C^0 and C^* coincide when all agents realize beliefs $\underline{\Pi}_2$, so we have

$$\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^* \left| s^1 \right) \right| s_0 \right] = \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right) \right| s_0 \right].$$

Assumption 1 guarantees that $\Pi'_{i,1}$ is not only non-empty, but for any belief $\pi_{i,1} \in \Pi_{i,1}$ there is another belief $\pi'_{i,1} \in \Pi_{i,1}$ with the same marginal distribution over $t = 1$ skills θ_1. Thus an agent not only considers distributions that place unit weight on the belief set $\underline{\Pi}_2$, but he thinks it is possible for everyone to realize $\underline{\Pi}_2$ regardless of the distribution of skills θ_1. In this sense, an agent's uncertainty about beliefs has some independence from his uncertainty about skills. Because of this and the fact that C^0 takes for granted the worst beliefs at $t = 1$,

$$\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right) \right| s_0 \right] = \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right) \right| s_0 \right].$$

Together with the fact that C^0 and C^* coincide at $t = 0$, this implies that every agent weakly prefers C^0 to C^*. But C^* is an efficient allocation, so it must be that all agents are indifferent between C^* and the simplified allocation C^0.

Of course, if at $t = 1$ all agents do not realize beliefs $\underline{\Pi}_2$, the continuation allocation $\{c_1^0, z_1^0, k_2^0\}$ may not be optimal anymore. In that case, the government will seek to reform C^0 to a new allocation C^1 in order to raise $t = 1$ social welfare while continuing to deliver the $t = 0$ continuation utility promised to each agent under C^0. The same arguments as above imply that the reform C^1 can be simplified in the same sense as C^0, so by solving for optimal reform allocations in each period, the government can construct a sequence of

[12] We abuse notation here since C^0 also specifies a fallback allocation where for $t \geq 2$, c_t^0, z_t^0, and k_{t+1}^0 allocate based on what c_t^*, z_t^*, and k_{t+1}^* prescribe when every agent realizes $\underline{\Pi}_2$ at $t = 1$ and then $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$ at $t \geq 2$.

simplified allocations $\{C^t\}_{t=0}^T$ that implements C^*.

Constructing Optimal Reforms

This argument that the government can implement the efficient allocation with a sequence of simplified allocations also suggests a way to construct the simplified allocations $\{C^t\}_{t=0}^T$, without first solving for the efficient allocation C^*. In general, given a simplified allocation C^{t-1} that is not fully state-contingent, it can be reformed to the optimal simplified allocation C^{*t} given by

$$C^{*t}\left(s^t, C^{t-1}\right) \in \arg\max_{C^t} \sum_i \eta_i U_{i,t}\left(C^t \left| s^t\right.\right) \tag{2}$$

subject to non-negativity and

$$\sum_i c_{i,\tau}^t\left(s^\tau\right) + K_{\tau+1}^t\left(s^\tau\right) \leq f\left(K_\tau^{\tau-1}\left(s^{\tau-1}\right), Z_\tau^t\left(s^\tau\right)\right),$$

$$U_{i,t-1}\left(C_{t-1}^{t-1}, \left(C_\tau^t\right)_{\tau=t}^T \left| s^{t-1}\right.\right) \geq U_{i,t-1}\left(C^{t-1}\left| s^{t-1}\right.\right),$$

where the first constraint is feasibility that must hold for $\tau = t, t+1$ and for all $s^\tau \geq s^t$, and the second constraint is a form of promise-keeping that must hold for all i (with no such constraint at $t = 0$). Here $\left(C_{t-1}^{t-1}, \left(C_\tau^t\right)_{\tau=t}^T\right)$ denotes the allocation that uses the C^{t-1} policy functions in period $t-1$ and the C^t policy functions in periods $\tau = t, .., T$. The promise-keeping constraint appears because when the government chooses a policy at $t-1$, it commits to delivering at least the $t-1$ continuation utility promised to each agent, even after subsequent reforms.

Note that the component of the C^{*t-1} allocation after $t-1$ serves as a fallback option (i.e., an endogenous outside option) in the case that the government cannot construct a better reform allocation at t. In this sense, the friction created by uncertainty leads to a form of endogenous lack of commitment: Even though the government has the ability to fully commit to maintaining an allocation forever, it may choose not to and instead design optimal simplified policies period by period, subsequently reforming them as necessary.

While Proposition 1 guarantees the existence of simplified, periodically-reformed allocations, its full power for optimal policy applications is more transparent in this reinterpretation as an algorithm that permits the construction of simplified allocations without first constructing complete, fully state-contingent efficient allocations.

2.3 History independence

The allocations C^{*t} defined above are simplified because they are not fully contingent on future shocks. In particular, their construction makes it clear that they do not depend on beliefs $\Pi_{i,t+2}$ or type $s_{i,\tau}$ for $\tau \geq t+2$. We next show that they are simplified also in the sense that they lose full history dependence when reforms provide an improvement to previously designed government policies. Specifically, whenever agents' beliefs are Markov and a reform results in the government's promise-keeping constraints being slack at t, the optimal C^{*t} is independent of the $t-1$ state s^{t-1}.

Proposition 2. *For any t at which agents' beliefs are Markov and the promise-keeping constraints in the reform problem (2) are slack, the optimal C^{*t} is independent of full history.*

Proof. When the promise-keeping constraints in problem (2) do not bind at t, the government must maximize an η-weighted average of agents' t continuation utilities, subject to feasibility at t and $t+1$. If all agents' belief distributions $\pi_{i,t+1}$ are Markov, the continuation utility $U_{i,t}\left(C^t \left| s^t\right.\right)$ does not depend on s^{t-1} other than through C^t, and similarly s^{t-1} is only relevant to the feasibility constraint through period-t capital $K_t^{t-1}\left(s^{t-1}\right)$, which is fixed at the beginning of t. Thus to maximize its objective, the government will choose the optimal reformed C^{*t} so that it does not depend on s^{t-1}, given fixed $K_t^{t-1}\left(s^{t-1}\right)$. □

This naturally generalizes to the case in which agents' conditional belief distributions at t are not Markov but also depend on shocks from period $\tau \leq t$ to period t, $(s_\tau, \ldots, s_t)$. In this case, the optimal C^{*t} is independent of the $t-\tau-1$ state $s^{t-\tau-1}$ whenever a reform leads to the government's promise-keeping constraints being slack at t.

It is also useful to consider an example in which the allocation C^{*t} will never depend on the $t-1$ state s^{t-1} for $t \geq 2$. In particular, suppose C^{*t} can be computed by backward induction for all t. Fix $t=0$ and start by solving for the optimal $t=1$ continuation allocation $C_1^{*0} \equiv \{c_1^{*0}, z_1^{*0}, k_2^{*0}\}$, assuming that all belief realizations at $t=1$ will be $\underline{\Pi}_2$. Then each of the policy functions in the continuation allocation C_1^{*0} depend only on the $t=0$ state s_0 and the $t=1$ skills θ_1, and we have

$$C_1^{*0}\left(\underline{s}^1, k_1^0\right) \in \arg\max_{C_1^0} \sum_i \eta_i U_{i,1}\left(C_1^0 \left| \underline{s}^1\right.\right)$$

subject to non-negativity and

$$\sum_i c_{i,1}^0\left(\underline{s}^1\right)\eta_i + K_2^0\left(\underline{s}^1\right) \leq f\left(K_1^0\left(s_0\right), Z_1^0\left(\underline{s}^1\right)\right),$$

where $\underline{s}^1$ is a $t=1$ state such that every agent realizes the belief set $\underline{\Pi}_2$ at $t=1$. The government solves this problem for each state $\underline{s}^1 \geq s_0$ and each capital function k_1^0, and the remaining policy functions $c_0^{*0}, z_0^{*0}, k_1^{*0}$ are then found by a similar optimization problem, taking C_1^{*0} as given.

At $t = 1$, if a state of the form $\underline{s}^1$ is realized, then by construction the continuation allocation C_1^{*0} is optimal. Thus the government can choose the optimal reform allocation C^{*1} so that it coincides with C_1^{*0} when a state of the form $\underline{s}^1$ is realized. If the $t = 1$ state s^1 is not of the form $\underline{s}^1$, then some agent realized a belief set $\Pi_{i,2} \neq \underline{\Pi}_2$. Under a belief $\pi_{i,2} \notin \underline{\Pi}_2$, the agent believes that at some period $t \geq 2$, some agent j will realize skill $\theta_{j,t} > \underline{\theta}$. In this case, the government will be able to produce strictly greater output than if all agents realized skill $\underline{\theta}$, while still providing t continuation utility $U_{m,t}\left(C_t^{*0}, \underline{s}^t\right)$ to all agents m. In effect, the government's feasibility constraints at $t \geq 2$ are looser under a distribution $\pi_{i,2} \notin \underline{\Pi}_2$ than under $\underline{\pi}_2$. Since the government would not reduce agents' continuation utilities in response to loosened feasibility constraints, the optimal reform allocation C^{*1} must deliver weakly greater $t = 1$ continuation utility to every agent when a state s^1 not of the form $\underline{s}^1$ is realized.

These arguments imply that C^{*1} delivers weakly greater $t = 1$ continuation utility than C_1^{*0} to every agent regardless of the realized state s^1, so any re-optimized C^{*1} must trivially satisfy the promise-keeping constraint. In particular, if the optimal simplified allocation C^{*t} can be constructed by backward induction for all t, then promise-keeping constraints will never bind. The resulting C^{*t} will then be independent of s^{t-1} for $t \geq 1$.

2.4 Discussion

Two points closely related to our results above warrant further discussion. First, we make use of a condition on agents' beliefs that requires them to be sufficiently similar. Assumption 1 is an intuitive way to guarantee this condition, but it can be weakened in various ways while maintaining the results. We describe one example of such weakening below. Second, uncertainty (relaxing Savage's Sure-Thing Principle) creates the potential for dynamic inconsistency regardless of the particular preference representation or updating rules (see, e.g., Machina and Siniscalchi 2014). We discuss the extent to which the results can be attributed to a form of dynamic inconsistency. In particular, we show that agents' preferences satisfy dynamic consistency in the sense commonly used in the literature.

Weaker Belief Conditions

Suppose now that in each period t, everyone agrees it is possible that everyone will be certain at $t+1$ that skills will fall within some time-dependent interval from then on. Suppose also that agents think it is possible to have these beliefs regardless of the distribution of $t+1$ skills, and with these beliefs at $t+1$ everyone is completely uncertain about the particular distribution of future skills. Then any efficient allocation can again be implemented by a sequence of simplified policies with periodic reforms, though each policy in the sequence features greater state dependence than before.

Formally, for each t, let $\{\theta_{t,\tau}\}_{\tau=t+2}^{T}$ be a set of skill values, and inductively define the belief set $\underline{\Pi}_{t,\tau}$ for $\tau \geq t+2$ as follows: Let $\underline{\Pi}_{t,T}$ be the set of all beliefs $\pi_{t,T}$ such that for any s^{T-1}, $\pi_{t,T}\left(\cdot\middle| s^{T-1}\right)\big|_{\theta}$ is supported on $[\underline{\theta}, \theta_{t,T}]^N$. Then for any $\tau \geq t+2$, let $\underline{\Pi}_{t,\tau}$ be the set of all beliefs $\pi_{t,\tau}$ such that for any $s^{\tau-1}$, $\pi_{t,\tau}\left(\cdot| s^{\tau-1}\right)|_{\theta}$ is supported in $[\underline{\theta}, \theta_{t,\tau}]^N$ and $\pi_{t,\tau}\left(\cdot| s^{\tau-1}\right)|_{\Pi}$ places unit weight on all agents realizing belief set $\underline{\Pi}_{t,\tau+1}$. With this definition, any agent who realizes the belief set $\underline{\Pi}_{t,\tau}$ at $\tau-1 \geq t+1$ believes that every agent will realize a skill in the set $[\underline{\theta}, \theta_{t,\tau}]$ for $\tau \geq t+2$, but he is completely uncertain about the particular distribution of skills. When $\theta_{t,\tau} = \underline{\theta}$ for all t and τ, $\underline{\Pi}_{t,\tau}$ becomes the belief set $\underline{\Pi}_{\tau}$ of Assumption 1.

We relax Assumption 1 in the following way: For any $t = 0, \ldots, T-1$, any i, any s^t, and any $\pi_{i,t+1} \in \Pi_{i,t+1}$, there exists $\pi'_{i,t+1} \in \Pi_{i,t+1}$ such that

$$\pi_{i,t+1}\left(\cdot\middle| s^{t}\right)\big|_{\theta} = \pi'_{i,t+1}\left(\cdot\middle| s^{t}\right)\big|_{\theta}$$

but $\pi'_{i,t+1}\left(\cdot| s^{t}\right)\big|_{\Pi}$ places unit weight on all agents realizing belief set $\underline{\Pi}_{t,t+2}$. With this condition, at t any agent thinks it is possible that at $t+1$, all agents will believe that everyone will realize a skill in the set $[\underline{\theta}, \theta_{t,\tau}]$ for $\tau \geq t+2$. This is a direct weakening of Assumption 1, and it similarly guarantees that agents' beliefs have sufficient overlap so that they can each agree to a policy that is not fully state-contingent at $t+1$.

As before, let C^* be an efficient allocation. Similarly to the construction in Section 2.2, C^0 is defined based off of what C^* prescribes when every agent realizes belief set $\underline{\Pi}_{0,2}$ at $t=1$ and shocks $\left(\theta, \underline{\Pi}_{0,t+1}\right)$ for $t \geq 2$, where $\theta \in [\underline{\theta}, \theta_{0,t}]$. Begin by setting

$$\left\{c_0^0, z_0^0, k_1^0\right\} \equiv \left\{c_0^*, z_0^*, k_1^*\right\},$$

so that C^0 and C^* coincide at $t=0$. For $t=1$, define c_1^0, z_1^0, and k_2^0 so that regardless of their beliefs Π_2, each agent is allocated consumption, labor, and capital based on what C^* would allocate if everyone realized the belief set $\underline{\Pi}_{0,2}$. As in Section 2.2, C^0 also specifies a fallback allocation at $t \geq 2$, where c_t^0, z_t^0, and k_{t+1}^0 are defined by what C^* prescribes when every agent realizes a shock of the form $\left(\theta, \underline{\Pi}_{0,\tau+1}\right)$ at $\tau \leq t$ with $\theta \in [\underline{\theta}, \theta_{0,\tau}]$. If an agent actually realizes a τ skill $\theta \notin [\underline{\theta}, \theta_{0,\tau}]$, then he is allocated consumption and effective labor as if he realized the shock $\left(\theta_{0,\tau}, \underline{\Pi}_{0,\tau+1}\right)$. Note that the sets $[\underline{\theta}, \theta_{t,\tau}]$ were chosen so that the fallback allocation constructed in this example will remain feasible if an agent realizes a skill $\theta_{i,t} \notin [\underline{\theta}, \theta_{t,\tau}]$ at some t.

The government may wish to reform the allocation C^0 at $t = 1$ if all agents do not realize the belief set $\underline{\Pi}_{0,2}$. In this case, we can apply the same argument as above to conclude that the government can construct a simplified allocation C^1, but it must additionally ensure that it delivers weakly greater $t = 0$ continuation utility after the reform than under C^0. By reforming the allocation in each period, the government can find a sequence of simplified allocations $\{C^t\}_{t=0}^{T}$ that implements the (potentially fully state-contingent) efficient allocation C^*.

It is easy to see that the allocation C^0 displays limited state dependence at $t \geq 1$: The policy functions do not depend on beliefs Π_{t+1} for $t \geq 1$, and they do not differentiate between any two types $\theta, \theta' \in \left[\theta_{0,t}, \overline{\theta}\right]$ for $t \geq 2$. The degree to which C^0 is state-dependent is increasing in $\theta_{0,t}$ for each $t \geq 2$, and setting $\theta_{0,t} \equiv \underline{\theta}$ for $t \geq 2$ leads to the setting addressed in Section 2.2. Moreover, the same argument as in Section 2.3 implies that the allocations $\{C^t\}_{t=0}^{T}$ display loss of full history dependence when beliefs are Markov and reforms lead to an improvement on previous policies.

Dynamic Consistency

We show that agents' recursive maxmin preferences satisfy dynamic consistency in a natural sense that is commonly used in the literature on dynamic preferences with uncertainty.

Lemma 1. *Fix any $t < T$ and any s^t. If C and $\tilde{C}$ coincide at t and*

$$U_{i,t+1}\left(C \left| s^{t+1}\right.\right) \leq U_{i,t+1}\left(\tilde{C} \left| s^{t+1}\right.\right) \tag{3}$$

for all $s^{t+1} \geq s^t$, then

$$U_{i,t}\left(C \left| s^{t}\right.\right) \leq U_{i,t}\left(\tilde{C} \left| s^{t}\right.\right).$$

According to the lemma, if an agent weakly prefers the allocation $\tilde{C}$ to C at $t+1$ for any state s^{t+1} consistent with the realized state s^t, and the two allocations do not differ at t, then the agent will weakly prefer $\tilde{C}$ to C at t. In this sense, we say that the preferences described in Section 2.1 are *dynamically consistent.* The proof of Lemma 1 makes straightforward use of the recursive form of $U_{i,t}$: Since C and $\tilde{C}$ coincide at t,

$$u\left(c_{i,t}\left(s^t\right), \frac{z_{i,t}\left(s^t\right)}{\theta_{i,t}}\right) = u\left(\tilde{c}_{i,t}\left(s^t\right), \frac{\tilde{z}_{i,t}\left(s^t\right)}{\theta_{i,t}}\right).$$

Allocation $\tilde{C}$ is weakly preferred at $t+1$ for all $s^{t+1} \geq s^t$, and the agent's beliefs $\Pi_{i,t+1}$ are not allocation-dependent, so we also have

$$\inf_{\Pi_{i,t+1}} \mathbb{E}_{\pi_{i,t+1}} \left[U_{i,t+1} \left(C \left| s^{t+1} \right. \right) \right| s^t \right] \leq \inf_{\Pi_{i,t+1}} \mathbb{E}_{\pi_{i,t+1}} \left[U_{i,t+1} \left(\tilde{C} \left| s^{t+1} \right. \right) \right| s^t \right]. \tag{4}$$

By the definition of the utility function $U_{i,t}$, we have the result.

Note that for Lemma 1 to hold, it is crucial that the agent's beliefs $\Pi_{i,t+1}$ are not allocation-dependent. To see this, let $\Pi_{i,t+1}(C)$ denote the agent's beliefs given allocation C, and define $\Pi_{i,t+1}\left(\tilde{C}\right)$ similarly. In the lemma, we suppose that inequality (3) holds pointwise (i.e., for each $s^{t+1} \geq s^t$). But, as a simple example, there may exist belief distributions $\pi_{i,t+1}$ and $\tilde{\pi}_{i,t+1}$ such that

$$U_{i,t+1} \left(\tilde{C} \left| \tilde{s}^{t+1} \right. \right) < U_{i,t+1} \left(C \left| s^{t+1} \right. \right)$$

for all $\tilde{s}^{t+1}$ in the support of $\tilde{\pi}_{i,t+1}$ and all s^{t+1} in the support of $\pi_{i,t+1}$. If $\Pi_{i,t+1}\left(\tilde{C}\right) = \{\tilde{\pi}_{i,t+1}\}$ and $\Pi_{i,t+1}(C) = \{\pi_{i,t+1}\}$ then inequality (4) will be strictly reversed. More generally, this problem arises because inequality (3) does not imply any restrictions about the values of $U_{i,t+1}(C\,|\cdot)$ and $U_{i,t+1}\left(\tilde{C}\,|\cdot\right)$ when evaluated at different points in their domains.

The dynamic consistency property described above features prominently in the literature on dynamic preferences with uncertainty. For example, the axiomatization of recursive Variational Preferences implies the same definition of dynamic consistency as Lemma 1. Similar assumptions are made in the axiomatizations of recursive Maxmin Preferences and recursive Smooth Ambiguity Preferences. Epstein and Schneider (2003) additionally show that dynamic consistency with recursive Maxmin Preferences is equivalent to a "rectangularity" condition on agents' prior distributions (see their Theorem 3.2), and it can be shown that the sets of priors induced by our belief sets $\Pi_{i,t+1}$ satisfy that condition as well.

Lemma 1 implies that agents' preferences also satisfy a slightly weaker property, but one that is arguably more relevant in macroeconomics and public finance (Hansen and Sargent 2006): Dynamic preferences are sufficiently consistent over time if a solution to a dynamic choice problem computed by backward induction is also optimal ex ante. This condition is strictly weaker than the notion of dynamic consistency above because it only implies consistency in preference orderings involving the ex ante optimal solution. For example, the constraint preferences of Hansen and Sargent (2001) do not satisfy the definition of dynamic consistency above, but they do satisfy the weaker property (see, e.g., Epstein and Schneider 2003, Section 5).

3 Private Skills and Private Beliefs

For some applications, it is clearly too demanding to presume that lack of certainty about the data-generating process is the only friction in the design of optimal policies. This section is devoted to relaxing the assumption of publicly observable skills and beliefs. We consider the consequences of private shocks and demonstrate that the main results of Section 2 persist in economies constrained by uncertainty as well as private information.

3.1 Private information setup

Consider the following modification to the baseline setup in Section 2: Agents are privately informed about their shocks $s_{i,t}$ in each period, and after receiving their shocks at the beginning of the period, they can make reports. Let $\hat{s}_{i,t}$ denote a period-t *reported shock*, and let $\hat{s}_i^t \equiv \left(\hat{\theta}_i^t, \hat{\Pi}_i^{t+1}\right)$ denote a history of reported shocks up to period t, called a period-t *reported type*. Also let $\hat{s}^t \equiv (\hat{s}_1^t, \ldots, \hat{s}_N^t)$ denote a vector of reported types for every agent in the economy, called a period-t *reported state*.

A *reporting strategy* is $\sigma_i \equiv \{\sigma_{i,t}\}_{t=0}^T$, where $\sigma_{i,t}$ maps a reported state $\hat{s}^{t-1}$ and a type s_i^t to a reported shock $\hat{s}_{i,t}$. Let Σ be the set of possible reporting strategies, and let $\sigma \equiv \{\sigma_i\}_{i=1}^N \in \Sigma^N$ be a strategy profile. The truth-telling strategy is denoted σ_i^*, and similarly $\sigma^* \equiv \{\sigma_i^*\}_{i=1}^N$ denotes the strategy profile in which all agents use the truth-telling strategy.

Given an allocation $C = \{c_t(\hat{s}^t), z_t(\hat{s}^t), k_{t+1}(\hat{s}^t)\}_{t=0}^T$, a reported state $\hat{s}^{t-1}$, a type s_i^t, and a strategy profile σ, agent i's period-t continuation utility is[13]

$$U_{i,t}\big(C \,\big|\hat{s}^{t-1}, s_i^t\big)(\sigma) \equiv \inf_{\Pi_{i,t+1}} \mathbb{E}_{\pi_{i,t+1}}\left[W_{i,t}\big(C\,\big|\hat{s}^{t-1}, s^t\big)(\sigma)\,\big|\,\hat{s}^{t-1}, s_i^t\right], \tag{5}$$

where

$$W_{i,t}\big(C\,\big|\hat{s}^{t-1}, s^t\big)(\sigma) \equiv u\left(c_{i,t}\left(\hat{s}^{t-1}, \sigma_t\left(\hat{s}^{t-1}, s^t\right)\right), \frac{z_{i,t}\left(\hat{s}^{t-1}, \sigma_t\left(\hat{s}^{t-1}, s^t\right)\right)}{\theta_{i,t}}\right) + \beta U_{i,t+1}\big(C\,\big|\left(\hat{s}^{t-1}, \sigma_t\left(\hat{s}^{t-1}, s^t\right)\right), s_i^{t+1}\big)(\sigma). \tag{6}$$

In (6) we slightly abuse notation letting $\sigma_t(\hat{s}^{t-1}, s^t) \equiv (\sigma_{1,t}(\hat{s}^{t-1}, s_1^t), \ldots, \sigma_{N,t}(\hat{s}^{t-1}, s_N^t))$ denote the vector of all agents' t reports. Thus $(\hat{s}^{t-1}, \sigma_t(\hat{s}^{t-1}, s^t))$ denotes the reported state $\hat{s}^t$, given that the true state realized at t is s^t and that the agents follow the strategy profile σ. With definitions (5) and (6), $W_{i,t}$ is the payoff of agent i after types are reported in

[13] The expectation in (5) is implicitly conditioned on the strategy profile up to period $t-1$, $(\sigma_\tau)_{\tau=0}^{t-1}$. In what follows, it should be understood that conditioning on a $t-1$ reported state $\hat{s}^{t-1}$ implicitly includes conditioning on the strategy profile $(\sigma_\tau)_{\tau=0}^{t-1}$ that the agents followed through period $t-1$.

period t, while $U_{i,t}$ gives the infimum of the expected payoff before reporting. In contrast to Section 2, period-t instantaneous utility u now appears inside of the expectation $\mathbb{E}_{\pi_{i,t+1}}$ (and inside of the infimum) because the other agents' types s^t_{-i} are not publicly observable. As a result, agent i is potentially uncertain about the distribution of the other agents' current types as well as the distribution of future shocks, and he is averse to this uncertainty.

Given the information structure, we define an *equilibrium strategy profile* as a strategy profile $\sigma^e \in \Sigma^N$ such that

$$U_{i,0}\left(C\left|\hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^e) \geq U_{i,0}\left(C\left|\hat{s}^{-1}, s_{i,0}\right.\right)\left(\sigma^e_{-i}, \sigma_i\right)$$

for all i, all $\sigma_i \in \Sigma$, and all $s_{i,0}$.[14] Here $\left(\sigma^e_{-i}, \sigma_i\right)$ denotes the strategy profile in which agent i uses strategy σ_i and any agent $j \neq i$ uses his respective equilibrium strategy σ^e_j. Given the agents' preferences, this definition of an equilibrium strategy profile is natural: It implies that at $t = 0$, any agent i prefers to follow his equilibrium strategy σ^e_i in all future periods, given that all other agents $j \neq i$ also follow their equilibrium strategies σ^e_j.[15]

3.2 Constrained efficiency

We first show that a form of the revelation principle extends to this environment with uncertainty. This justifies focusing on *incentive-compatible* allocations, i.e., allocations for which σ^* is an equilibrium strategy profile. The proof is standard and the details are provided in the Online Appendix.

Lemma 2. *For any allocation C and any equilibrium strategy profile σ^e, there exists an allocation $\tilde{C}$ such that σ^* is an equilibrium strategy profile, with*

$$U_{i,0}\big(\tilde{C}\left|\hat{s}^{-1}, s_{i,0}\right.\big)\big(\sigma^*\big) = U_{i,0}\big(C\left|\hat{s}^{-1}, s_{i,0}\right.\big)\big(\sigma^e\big)$$

for all i and all $s_{i,0}$.[16]

---

[14] Note that $\hat{s}_{-1}$ is an empty vector because agents make no reports before $t = 0$, and it is included only to maintain consistent notation.

[15] For any C, by Nash's existence theorem there exists an equilibrium profile in mixed strategies. This holds because there are finitely many agents and the set of types $s_{i,t}$ that can be realized at any t is finite, so the set of all reporting strategies Σ is finite. The existence of an equilibrium profile for any allocation implies that it is sufficient for agents to form beliefs only about other agents' true types, i.e., each allocation comes with a proposed equilibrium strategy profile, and we make the standard assumption that the agents follow the proposed equilibrium strategies. The focus on pure strategy equilibria is justified by the revelation principle discussed next.

[16] If preferences satisfy "no-hedging", i.e., the expectation for any mixed strategy appears outside of the infimum in $U_{i,0}$, a far more general statement holds: Given an arbitrary mechanism and an equilibrium in mixed strategies, there exists an incentive-compatible direct mechanism that implements the equilibrium

With private shocks, the government can no longer observe the $t = 0$ state s_0 before designing a policy. As a result, it cannot evaluate the objective in problem (1) unless it is allowed to form beliefs in each period about agents' types, which it can then use to take expectations over the unobservable state. For example, each vector of t shocks s_t could be expanded to include a belief set $\Pi_{g,t+1}$ for the government that it would use to evaluate social welfare in each period. This requires specifying how the government's beliefs are formed. One natural option is to choose one of the agents to serve as the government and use her beliefs to take expectations over the unobservable state in each period. We use this approach below, but the results generalize to any setting in which the government's objective consists of aggregating agents' utilities $U_{i,0}$.

At $t = 0$, an agent is chosen uniformly at random to serve as the government and design a social insurance policy. Denote the index of the governing agent by $g \in \{1, \ldots, N\}$. Given social welfare weights η and the governing agent's type $s_{g,0}$, an allocation $C^*(s_{g,0})$ is *constrained-efficient* if

$$C^*(s_{g,0}) \in \arg\max_{C} \inf_{\Pi_{g,1}} \mathbb{E}_{\pi_{g,1}} \left[\sum_i \eta_i U_{i,0} \left(C \left| \hat{s}_{-1}, s_{i,0} \right.\right) (\sigma^*) \middle| s_{g,0} \right] \tag{7}$$

$$\text{subject to non-negativity and}$$

$$\sum_i c_{i,t}\left(\hat{s}^t\right) + K_{t+1}\left(\hat{s}^t\right) \le f\left(K_t\left(\hat{s}^{t-1}\right), Z_t\left(\hat{s}^t\right)\right),$$

$$U_{i,0}\left(C \left| \hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^*) \ge U_{i,0}\left(C \left| \hat{s}^{-1}, s_{i,0}\right.\right)\left(\sigma^*_{-i}, \sigma_i\right),$$

where the feasibility constraint holds for $t = 0, \ldots, T$ and for all $\hat{s}^t$, and the incentive-compatibility constraint holds for all i, all $\sigma_i \in \Sigma$, and all $s_{i,0}$, reflecting the additional friction imposed by private shocks.

Similarly to Section 2, we assume that the government has full commitment power in the following sense: Once the government chooses an allocation at $t \ge 0$, it commits to delivering at least the t continuation utility promised to any truthfully-reporting agent, regardless of any possible future reforms; the government also commits to the incentives for truthful revelation, so that under any future reform, an agent who deviates from the truth-telling strategy at t cannot receive higher t continuation utility as a result of the reform.

outcome. This is proven in the same manner as Lemma 2, and we omit the proof. It requires, however, that we allow the policy functions of an allocation C to map a reported state $\hat{s}^t$ to non-degenerate lotteries over real numbers. The results below are unaffected with this interpretation of an allocation.

3.3 Periodic reforms with private shocks

We next demonstrate that an analog of Proposition 1 holds. As with public information, the result requires a condition on agents' beliefs in each period that ensures they have sufficient overlap. To simplify the exposition, we again focus on an intuitive version of the condition where all agents agree on the worst path for the economy, which guarantees the sufficient overlap. This worst path is defined in the same way as before: Given any period t and any state s^t, the worst path occurs when all agents realize the skill shock $\underline{\theta}$ at $\tau \geq t+2$. Then an analog of Assumption 1 is that at any t, each agent believes it is possible for all agents to be certain at $t+1$ that the economy will follow the worst path.

Sufficient Belief Overlap

For $t \geq 0$ let $\underline{\Pi}_{t+1} = \left\{\underline{\pi}_{t+1}\right\}$ be the belief set defined in Section 2.2, but modified as follows: For $\underline{\pi}_{t+1} \in \underline{\Pi}_{t+1}$, the marginal distribution of the t shocks places unit weight on $(\theta_{i,t}, \Pi_{i,t+1}) = \left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$ for all i. As with public information, any agent who realizes the belief set $\underline{\Pi}_{t+1}$ at t is certain that all agents will receive the worst skill shock $\underline{\theta}$ for $\tau \geq t+1$. The modification implies that he also believes that all agents received the worst skill shock $\underline{\theta}$ at t and that the other agents share his worst beliefs $\underline{\Pi}_{t+1}$. Unlike in the public information economy, here it is necessary to specify the marginal distribution of the t shocks because agents now use their belief sets $\Pi_{i,t+1}$ to evaluate expectations over other agents' t types.

Note that if all agents realize the belief set $\underline{\Pi}_{t+1}$ at t, by definition they have homogeneous beliefs about the distribution of t skills θ_t in the economy. To allow for heterogeneous beliefs about t skills while maintaining the possibility of the worst beliefs about the future, the belief set $\underline{\Pi}_{t+1}$ must be generalized slightly. In particular, we must define a class of belief sets $\underline{\underline{\Pi}}_{i,t+1}$ such that any agent i who realizes a belief set $\underline{\underline{\Pi}}_{i,t+1}$ is certain that all agents will realize the skill $\underline{\theta}$ at $\tau \geq t+1$, but he has potentially nontrivial beliefs about t shocks θ_t. As before, under $\underline{\underline{\Pi}}_{i,t+1}$ agent i is certain that all other agents share his pessimism, i.e., he believes that the other agents also realized belief sets of the form $\underline{\underline{\Pi}}_{j,t+1}$.

Formally, for any $t \geq 1$, let $\underline{\underline{\Delta}}\left(s^{t+1}\right) \subset \Delta\left(s^{t+1}\right)$ denote the subset of belief distributions $\underline{\underline{\pi}}_{i,t+1}$ that place unit weight on $t+1$ shocks $(\theta_{j,t+1}, \Pi_{j,t+2}) = \left(\underline{\theta}, \underline{\Pi}_{t+2}\right)$ and $t+1$ belief sets $\underline{\underline{\Pi}}_{j,t+1} \subset \underline{\underline{\Delta}}\left(s^{t+1}\right)$ for all j. Any agent who realizes a belief set $\underline{\underline{\Pi}}_{i,t+1} \subset \underline{\underline{\Delta}}\left(s^{t+1}\right)$ at t is thus certain that all agents will realize the worst skill shock $\underline{\theta}$ at $\tau \geq t+1$, and he is also certain that all agents are similarly pessimistic about the future. However, a belief distribution $\underline{\underline{\pi}}_{i,t+1} \in \underline{\underline{\Delta}}\left(s^{t+1}\right)$ can have any marginal distribution over the $t-1$ state s^{t-1} and t skills θ_t, so $\underline{\underline{\Delta}}\left(s^{t+1}\right)$ contains many more distributions than simply the distribution $\underline{\pi}_{t+1}$ defined before. In fact, in the Online Appendix we show that for any $t+1$ distribution $\pi_{i,t+1}$, there exists a natural corresponding distribution $\underline{\underline{\pi}}_{i,t+1} \in \underline{\underline{\Delta}}\left(s^{t+1}\right)$ such that $\pi_{i,t+1}$ and $\underline{\underline{\pi}}_{i,t+1}$ have the same marginal distribution over s^{t-1} and θ_t.[17] Assumption 1 is then modified as follows:

Assumption 2. *For any $t = 0, \dots, T-1$, any i, any $\hat{s}^{t-1}$, any s_i^t, and any belief $\pi_{i,t+1} \in \Pi_{i,t+1}$, there exists another belief $\pi'_{i,t+1} \in \Pi_{i,t+1}$ such that*

$$\pi_{i,t+1}\left(\cdot \middle| \hat{s}^{t-1}, s_i^t\right)\Big|_{s_{-i,t},\theta_{t+1}} = \pi'_{i,t+1}\left(\cdot \middle| \hat{s}^{t-1}, s_i^t\right)\Big|_{s_{-i,t},\theta_{t+1}}$$

but $\pi'_{i,t+1}\left(\cdot \middle| \hat{s}^{t-1}, s_i^t\right)\Big|_{\Pi_{t+2}}$ places unit weight on the event $\Pi_{j,t+2} \subset \underline{\underline{\Delta}}\left(s^{t+2}\right) \; \forall j$.

Similarly to Assumption 1, this requires that at the beginning of t, regardless of his type s_i^t or the reported state $\hat{s}^{t-1}$, any agent i believes that it is possible for everyone at $t+1$ to have the worst beliefs of the form $\underline{\underline{\Pi}}_{j,t+2}$. The agent believes that this is possible under any beliefs about the other agents' shocks $s_{-i,t}$ or skills θ_{t+1}. As a result, at t all agents agree on at least one possible path for the economy.

Notice that in the public information economy, this worst path in terms of feasibility coincides with every agent's subjective worst path in terms of continuation utility because each agent's continuation utility is monotonically increasing in the amount of resources available to the government in each period. This monotonicity may not hold when the government must also satisfy incentive compatibility. To ensure then that all agents agree on the worst path in terms of both feasibility and subjective utility, we assume a weak version of monotonicity: An allocation C is *weakly monotonic at t* if for any agent i, any type s_i^t, any reported state $\hat{s}^{t-1}$, and any strategy $\sigma_i \in \Sigma$, the agent's expected payoff under any belief $\pi_{i,t+1} \in \Pi_{i,t+1}$,

$$\mathbb{E}_{\pi_{i,t+1}}\left[W_{i,t}\left(C \middle| \hat{s}^{t-1}, s^t\right)\left(\sigma^*_{-i}, \sigma_i\right) \middle| \hat{s}^{t-1}, s_i^t\right],$$

is weakly greater than when the other agents are certain to report worst beliefs $\underline{\underline{\hat{\Pi}}}_{j,t+2}$ at $t+1$ and shock $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ at $\tau \geq t+2$. Intuitively, weak monotonicity implies that any agent evaluates his utility of a deviation from truth-telling based on how the deviating strategy σ_i performs along the worst path.

Optimality of Periodic Reforms

Simplified, periodically-reformed policies are optimal given the sufficient belief overlap in Assumption 2 and weak monotonicity of reforms:

[17] The formalism of $\underline{\underline{\Delta}}\left(s^{t+1}\right)$ again involves a degree of circularity, which can be eliminated by considering a belief-closed subset of an appropriate type space. Restricting types to such a subset would imply common knowledge about the worst path beliefs at t.

Proposition 3. *A constrained-efficient allocation* C^* *can be implemented by a sequence of allocations* $\{C^t\}_{t=0}^T$*, where* $C^t = \left\{c_\tau^t, z_\tau^t, k_{\tau+1}^t\right\}_{\tau=t}^{t+1}$.

As with Proposition 1, the implementation here is in the sense that every agent's $t = 0$ utility under the simplified allocation C^0 is equal to that under C^*: For any i and any s_0,

$$U_{i,0}\left(C^0\,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*) = U_{i,0}\left(C\,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*)\,.$$

We again leave a detailed proof in the Online Appendix and here demonstrate the key steps of the argument. Under Assumption 2, at $t = 0$ everyone agrees that it is possible for the economy to follow the worst path. The $t = 0$ simplified allocation C^0 is constructed such that it coincides with C^* at $t = 0$, and at $t = 1$ it is defined based on how C^* allocates when the worst path attains, i.e., when all agents report beliefs $\underline{\underline{\hat{\Pi}}}_{i,2} \subset \underline{\underline{\Delta}}\left(s^2\right)$. If an agent reports a belief set $\hat{\Pi}_{i,2}$ that is not contained in $\underline{\underline{\Delta}}\left(s^2\right)$, then C^0 allocates as if the agent reported a belief set $\underline{\underline{\hat{\Pi}}}_{i,2}$ that is contained in $\underline{\underline{\Delta}}\left(s^2\right)$ but otherwise corresponds to $\hat{\Pi}_{i,2}$.[18]

Under the truth-telling strategy profile σ^*, the same argument as with Proposition 1 shows that every agent weakly prefers C^0 to C^*: For all i and all $s_{i,0}$,

$$U_{i,0}\left(C^*\,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*) \leq U_{i,0}\left(C^0\,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*) \tag{8}$$

By construction, the allocation C^0 also satisfies non-negativity and feasibility. Unlike in the public information case, we must also verify that C^0 is incentive-compatible to ensure that the truth-telling strategy profile σ^* is actually followed by the agents.

Intuitively, C^0 remains incentive-compatible because under either C^0 or C^*, the only distributions relevant to an agent when choosing a strategy are the "worst beliefs" $\pi'_{i,1}$ defined in Assumption 2. This holds trivially for C^0 because this allocation presumes that agents will report belief sets $\underline{\underline{\hat{\Pi}}}_{j,2}$ at $t = 1$ and the shock $\left(\underline{\theta}, \underline{\Pi}_{i,t+1}\right)$ at $t \geq 2$, and the beliefs $\pi'_{i,1}$ place unit weight precisely on such belief and shock realizations. For C^*, weak monotonicity at $t = 0$ implies that the worst case for agent i occurs when the other agents report belief sets $\underline{\underline{\hat{\Pi}}}_{j,2}$ at $t = 1$ and the shock $\left(\underline{\theta}, \underline{\Pi}_{i,t+1}\right)$ at $t \geq 2$. Given truth-telling by others, this happens when shocks are distributed according to a belief of the form $\pi'_{i,1}$.

Now any agent i must only consider the beliefs $\pi'_{i,1}$ when choosing a strategy for C^0 or C^*. Truthful reporting by others then amounts to assuming that the $-i$ agents report beliefs $\underline{\underline{\hat{\Pi}}}_{j,2}$ at $t = 1$ and the shock $\left(\underline{\theta}, \underline{\Pi}_{i,t+1}\right)$ at $t \geq 2$. By construction, C^0 and C^* coincide

[18] As in Section 2.2, allocation C^0 also specifies a fallback option $\left\{c_t^0, z_t^0, k_{t+1}^0\right\}_{t=2}^T$ that is defined by what C^* prescribes when all agents realize a belief set contained in $\underline{\underline{\Delta}}\left(s^2\right)$ at $t = 1$ and then the worst shock $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$ at $t \geq 2$.

when all agents make such reports, so we can bound the value of agent i's problem under C^0 by the value of his problem under C^*:

$$\begin{aligned}&\max_{\sigma_i \in \Sigma} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big) \big|\, s_{i,0}\right] \\ &\leq \max_{\sigma_i \in \Sigma} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big) \big|\, s_{i,0}\right].\end{aligned}$$

The inequality is weak because C^* does not necessarily "assume" that agent i will report belief set $\hat{\underline{\underline{\Pi}}}_{i,2}$ at $t = 1$ and the shock $\left(\underline{\theta}, \underline{\Pi}_{i,t+1}\right)$ at $t \geq 2$, so agent i can effectively deviate in a larger number of ways under C^* than under C^0. But C^* is incentive-compatible by definition, so the right side is equal to the agent's $t = 0$ utility under C^* with truth-telling, $U_{i,0}\left(C^* \,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*)$. By (8), every agent weakly prefers C^0 to C^* under the truth-telling strategy profile, so we can bound the value of the agent's problem under C^0 by the utility with truth-telling:

$$\begin{aligned}&\max_{\sigma_i \in \Sigma} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big) \big|\, s_{i,0}\right] \\ &\leq U_{i,0}\left(C^0 \,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*).\end{aligned}$$

Clearly this inequality must bind, so we conclude that C^0 is incentive-compatible. This additionally implies that all agents are indifferent between C^0 and C^*, so C^0 must be a solution to the government's problem (7).

Finally, to describe how the government iterates forward to C^1 and on, note that since the continuation allocation $\{c_1^0, z_1^0, k_2^0\}$ is not fully state-contingent or necessarily $t = 1$ incentive-compatible, it may not be optimal at $t = 1$ if the economy does not follow the worst path. Thus the government may wish to construct a reform allocation C^1 that raises $t = 1$ social welfare relative to C^0. As in problem (7), the reform C^1 must satisfy feasibility and non-negativity constraints, and it must be incentive-compatible with respect to the agents' $t = 1$ utility functions $U_{i,1}$. In addition, the government remains committed to the utility promises it made at $t = 0$, as well as the incentives it provided for truthful revelation at $t = 0$. These commitments respectively take the form of promise-keeping and threat-keeping constraints that the allocation C^1 must satisfy. The *promise-keeping* constraint requires that the government deliver at least as much $t = 0$ utility to every agent when the reform C^1 is substituted for the $t = 0$ allocation C^0 in periods $t \geq 1$: For all i and for all $s_{i,0}$,

$$U_{i,0}\big(C_0^0, \big(C_\tau^1\big)_{\tau=1}^T \,\big|\hat{s}^{-1}, s_{i,0}\big)\big(\sigma^*\big) \geq U_{i,0}\big(C^0 \,\big|\hat{s}^{-1}, s_{i,0}\big)\big(\sigma^*\big).$$

This is equivalent to the promise-keeping constraint in the public information reform problem (2).

The *threat-keeping* constraint requires that the government punish $t = 0$ deviations from truth-telling at least as much with the reform C^1 as with the allocation C^0: For all i, all $s_{i,0}$, and all $\sigma_{i,0}$,

$$\begin{aligned} U_{i,0}\big(C_0^0, \big(C_\tau^1\big)_{\tau=1}^T \,\big|\hat{s}^{-1}, s_{i,0}\big)\big(\sigma_{-i,0}^*, \sigma_{i,0}, \big(\sigma_\tau^*\big)_{\tau=1}^T\big) \\ \leq U_{i,0}\big(C^0 \,\big|\hat{s}^{-1}, s_{i,0}\big)\big(\sigma_{-i,0}^*, \sigma_{i,0}, \big(\sigma_\tau^*\big)_{\tau=1}^T\big). \end{aligned}$$

This constraint has no analog in the public information reform problem (2) because it pertains directly to how the government incentivizes truthful revelation. Along with the $t = 1$ and $t = 0$ incentive-compatibility constraints, the threat-keeping constraint ensures that every agent weakly prefers the truth-telling equilibrium at $t = 0$, even if C^1 is substituted for C^0 in the future.

If there exists an allocation that satisfies non-negativity, feasibility, incentive-compatibility, promise-keeping, and threat-keeping at $t = 1$, then the government chooses such an allocation, call it C^{*1}, to maximize $t = 1$ social welfare

$$\inf_{\Pi_{g,2}} \mathbb{E}_{\pi_{g,2}} \left[\sum_i \eta_i U_{i,1}\big(C^1 \,\big|s_0, s_i^1\big)\big(\sigma^*\big) \middle| s_0, s_g^1 \right],$$

where s_0 is the truthfully-reported $t = 0$ state, and s_g^1 is the $t = 1$ type of the governing agent. Arguments analogous to those above imply that the government can then choose a simplified allocation C^1 to solve this $t = 1$ reform problem. If the $t = 1$ constraint set is empty, the government keeps the $t = 0$ allocation C^0 and picks the welfare-maximizing equilibrium σ^e of the allocation C^0 at $t = 1$. In subsequent periods, the government again attempts to reform the previous period's allocation.

Constructing Optimal Reforms

As in the public information case, the arguments above suggest it is not necessary to solve first for the constrained-efficient allocation C^* in order to find a sequence of simplified allocations $\{C^t\}_{t=0}^T$ that implements it. Rather, the sequence can be constructed by solving for an optimal reform in each period. In general, at any $t \geq 0$ the government seeks to design a reform that maximizes t social welfare. Suppose the allocation was last reformed in period $r < t$. Given the reported state $\tilde{s}^{t-1}$ and the governing agent's type s_g^t, the optimal reform $C^{*t}\left(\tilde{s}^{t-1}, s_g^t, C^r\right)$ is a solution to the reform problem

$$\max_{C^t} \inf_{\Pi_{g,t+1}} \mathbb{E}_{\pi_{g,t+1}} \left[\sum_i \eta_i U_{i,t}\big(C^t \left| \tilde{s}^{t-1}, s_i^t \right.\big)\left(\sigma^*\right) \middle| \tilde{s}^{t-1}, s_g^t \right] \tag{9}$$

$$\text{subject to non-negativity and}$$

$$\sum_i c_{i,\tau}^t\left(\hat{s}^\tau\right) + K_{\tau+1}^t\left(\hat{s}^\tau\right) \leq f\left(K_\tau^t\left(\hat{s}^{\tau-1}\right), Z_\tau^t\left(\hat{s}^\tau\right)\right),$$

$$U_{i,t}\big(C^t \left| \hat{s}^{t-1}, s_i^t \right.\big)\left(\sigma^*\right) \geq U_{i,t}\big(C^t \left| \hat{s}^{t-1}, s_i^t \right.\big)\left(\sigma_{-i}^*, \sigma_i\right),$$

$$U_{i,r}\big(\left(C_\tau^r\right)_{\tau=r}^{t-1}, \left(C_\tau^t\right)_{\tau=t}^{T} \left| \tilde{s}^{r-1}, s_i^r \right.\big)\left(\sigma^*\right) \geq U_{i,r}\big(C^r \left| \tilde{s}^{r-1}, s_i^r \right.\big)\left(\sigma^*\right)$$

$$U_{i,r}\big(\left(C_\tau^r\right)_{\tau=r}^{t-1}, \left(C_\tau^t\right)_{\tau=t}^{T} \left| \hat{s}^{r-1}, s_i^r \right.\big)\big(\left(\sigma_{-i,\tau}^*, \sigma_{i,\tau}\right)_{\tau=r}^{t-1}, \left(\sigma_\tau^*\right)_{\tau=r}^{T}\big)$$

$$\leq U_{i,r}\big(C^r \left| \hat{s}^{r-1}, s_i^r \right.\big)\big(\left(\sigma_{-i,\tau}^*, \sigma_{i,\tau}\right)_{\tau=r}^{t-1}, \left(\sigma_\tau^*\right)_{\tau=r}^{T}\big),$$

where the first constraint is feasibility for all $\tau \geq t$ and $\hat{s}^\tau$; the second constraint is the t incentive-compatibility that holds for all i, $\hat{s}^{t-1}$, s_i^t, and $\sigma_i \in \Sigma$; the third constraint is the promise-keeping for all i and s_i^r; and the last constraint is the threat-keeping for all i, $\hat{s}^{r-1}$, s_i^r, and $\left(\sigma_{i,\tau}\right)_{\tau=r}^{t-1}$. When $t = 0$, there are no promise-keeping or threat-keeping constraints.

If at any $t > 0$ the constraint set is empty, then the government sets $C^t \equiv C^{t-1}$ and picks the equilibrium σ^e of C^t with respect to the period-t utility functions $U_{i,t}$ that maximizes period-t social welfare

$$\inf_{\Pi_{g,t+1}} \mathbb{E}_{\pi_{g,t+1}} \left[\sum_i \eta_i U_{i,t}\big(C^t \left| \tilde{s}^{t-1}, s_i^t \right.\big)\left(\sigma^e\right) \middle| \tilde{s}^{t-1}, s_g^t \right].$$

If the constraint set in the reform problem (9) is non-empty at $t \geq 0$, the same argument as in the $t = 1$ case above implies that the government can choose a simplified allocation C^t as the solution. Following this process then constructs a sequence of simplified allocations $\{C^t\}_{t=0}^T$ that implements the constrained-efficient allocation C^*, without first solving for C^*.

As with public information, the existence of simplified, periodically-reformed allocations is guaranteed by Proposition 3. The utility of this result for optimal policy applications is demonstrated by this algorithmic interpretation that permits the construction of simplified allocations without first constructing complete, fully state-contingent constrained-efficient allocations.

History Independence

Similarly to the public information case in Section 2.3, the structure of the reform problem (9) characterizes conditions under which optimal policies may lose full dependence on the history of reports $\hat{s}^{t-1}$. In particular, suppose that at t, agents' beliefs $\pi_{i,t+1}$ are such that the shocks s_t and s_{t+1} are drawn from Markov distributions, and suppose that agents followed the truth-telling strategy profile σ^* at $t-1$. If a reform at t can provide an improvement to previously designed government policies, i.e., if the promise- and threat-keeping constraints in the reform problem (9) do not bind at t, then the optimal reform C^{*t} will not depend on the $t-2$ reported state $\hat{s}^{t-2}$. The reasoning for this is analogous to that in Section 2.3: With Markov beliefs and truthful reporting at $t-1$, the objective function and the incentive-compatibility constraint depend on $\hat{s}^{t-2}$ only through C^t, and the feasibility constraint depends on $\hat{s}^{t-2}$ only through C^t and period-t capital $K_t^{t-1}(\hat{s}^{t-1})$, which is fixed at the beginning of t. Because of this, the government will choose C^{*t} so that it does not depend on $\hat{s}^{t-2}$, conditional on fixed capital $K_t^{t-1}(\hat{s}^{t-1})$.

With private information, such improvement and consequent loss of full history dependence require that the threat-keeping constraint be non-binding at t in addition to the promise-keeping constraint. Since the threat-keeping constraint does not appear in the public information reform problem (2), policies may lose history dependence less frequently when agents have private information about shocks $s_{i,t}$. This is because the government must provide incentives for truthful revelation in the current period, so it will generally condition an agent's future consumption and effective labor on his current report. The government also commits to maintaining these incentives even after subsequent reforms, and this may require that reforms maintain some dependence on the history of shocks. In contrast, with no uncertainty, loss of history dependence is generically suboptimal and can confer large welfare losses over policies with full history dependence (see, e.g., Kapička 2017).

4 Inefficiency of Competitive Equilibria

In conventional economies without broader uncertainty, competitive equilibria result in (constrained) efficient allocations even when types are privately known to agents and markets are incomplete. A common interpretation is that the only result of a social insurance policy the government can put in place is to crowd out insurance provided by the decentralized, private markets (e.g., Golosov and Tsyvinski 2007, Acemoglu and Simsek 2012). With uncertainty and private information, we show that competitive equilibria may not be efficient. This holds because agents and firms may form heterogeneous beliefs in response to uncertainty, and when combined with private information, this additional heterogeneity naturally leads to market outcomes that can be strictly improved upon by the government. In other words, a broader view of uncertainty creates a potentially meaningful role for the government provision of insurance.

Nevertheless, we also show that the results of previous sections persist in that insurance in a decentralized economy can be provided with sequences of simplified allocations that are periodically reformed.

4.1 Decentralization

Consider once again the economy of Section 3 in which agents are privately informed about their shocks $s_{i,t}$ in each period. To decentralize this economy, introduce N firms that are owned equally by the agents and compete by offering contracts to agents on a one-to-one basis. Firms purchase capital k_0 from agents (equivalently, rent capital period by period) and employ the agents to supply effective labor $z_{i,t}$. In return, the agents receive consumption $c_{i,t}$. The firms produce output using the same deterministic, constant returns to scale production function $f : \mathbb{R}_+^2 \to \mathbb{R}$ as above and seek to maximize the net present value of dividends. It is easy to show that a revelation principle akin to Lemma 2 holds once again, so we assume without loss of generality that the firms offer contracts that incentivize truthful revelation by the agents. The firms face the same uncertainty as the agents about the data-generating process, so they also form potentially heterogeneous beliefs about the current and future states of the economy.

Suppose the firms had access to complete markets where they could trade fully-contingent Arrow-Debreu securities among themselves to insure against their uncertainty as well as risk from agents' idiosyncratic skills. That is, at the end of each period t, firms trade securities that are contingent on the idiosyncratic shock $\hat{s}_{i,t+1}$ reported by an agent at $t+1$. In equilibrium, however, securities of this kind that are contingent on idiosyncratic shocks are not traded:[19]

Lemma 3. *Securities that are contingent on idiosyncratic type reports are not traded in competitive equilibria with uncertainty.*

To see this, reinterpret the equilibrium as agents having direct access to the production function f and the ability to trade Arrow-Debreu securities among themselves. Fix t, and let $a\left(\hat{s}_{i,t+1}\right)$ denote a security that pays one unit of consumption after reporting at $t+1$ if agent i reports shock $\hat{s}_{i,t+1}$.[20] Let $q\left(\hat{s}_{i,t+1} \middle| \hat{s}^t, s^t\right)$ denote the equilibrium price of this security, and let $Q\left(\hat{s}^t, s^t\right)$ denote the price of a risk-free bond that pays one unit of consumption after reporting at $t+1$. In equilibrium, we must have $q\left(\hat{s}_{i,t+1} \middle| \hat{s}^t, s^t\right) \leq Q\left(\hat{s}^t, s^t\right)$ because $a\left(\hat{s}_{i,t+1}\right)$

[19] This is an extension of a straightforward result that holds even when agents are Bayesian and have rational expectations (e.g., Golosov and Tsyvinski 2007).

[20] The security $a\left(\hat{s}_{i,t+1}\right)$ may be contingent on other uncertain properties of the economy, but we abuse notation and write only agent i's idiosyncratic shock report to keep the exposition intuitive.

pays a unit of consumption only when agent i reports shock $\hat{s}_{i,t+1}$, while a risk-free bond pays one unit of consumption in all $t+1$ realizations. If the inequality were strict, agent i would purchase arbitrarily many $a(\hat{s}_{i,t+1})$ securities and sell a corresponding number of risk-free bonds. By reporting shock $\hat{s}_{i,t+1}$ at $t+1$ regardless of his actual realization, agent i could ensure that he would net arbitrarily high profits from these trades. Since sellers of the securities $a(\hat{s}_{i,t+1})$ would be guaranteed to lose, we cannot have $q(\hat{s}_{i,t+1} | \hat{s}^t, s^t) < Q(\hat{s}^t, s^t)$. Hence equality must obtain, and this immediately implies that securities contingent on idiosyncratic shocks are not traded in equilibrium.

The only securities potentially traded in equilibrium are then risk-free bonds. In this case, with uncertainty there exist economies in which competitive equilibria are not efficient:

Proposition 4. *Competitive equilibria with uncertainty may not be efficient.*

Proof. It suffices to construct a simple example economy in which a competitive equilibrium is not efficient. Fix $t=0$, and suppose that there are only two agents, $i \in \{1,2\}$. At $t=0$, agent 1 can only realize the belief sets $\{\bar{\pi}_{1,1}\}$ and $\{\underline{\pi}_{1,1}, \bar{\pi}_{1,1}\}$, and agent 2 can only realize the belief sets $\{\bar{\pi}_{2,1}\}$ and $\{\underline{\pi}_{2,1}, \bar{\pi}_{2,1}\}$. Here $\underline{\pi}_{1,1}$ is the distribution such that both agents are certain to realize the worst skill $\underline{\theta}$ at $t \geq 1$, and the marginal distribution of agent 2's beliefs $\underline{\Pi}_{2,1}$ under $\underline{\pi}_{1,1}$ places unit weight on $\{\underline{\pi}_{2,1}, \bar{\pi}_{2,1}\}$. Similarly, $\bar{\pi}_{1,1}$ is the distribution such that both agents are certain to realize the best skill $\bar{\theta}$ at $t \geq 1$, and the marginal distribution of $\underline{\Pi}_{2,1}$ under $\bar{\pi}_{1,1}$ places unit weight on $\{\underline{\pi}_{2,1}, \bar{\pi}_{2,1}\}$. The distributions $\underline{\pi}_{2,1}$ and $\bar{\pi}_{2,1}$ are defined symmetrically for agent 2.[21]

Suppose that agent 1 realized beliefs $\{\underline{\pi}_{1,1}, \bar{\pi}_{1,1}\}$, while agent 2 realized beliefs $\{\bar{\pi}_{2,1}\}$. Suppose also that agent 2 is the governing agent, with any strictly positive Pareto weights $\eta_1, \eta_2 > 0$. Then the government seeks to maximize

$$\mathbb{E}_{\bar{\pi}_{2,1}}\left[\sum_{i=1}^{2} \eta_i U_{i,0}\left(C\,|\hat{s}_{-1}, s_{i,0}\right)(\sigma^*)\middle| s_{2,0}\right],$$

subject to feasibility and period-0 incentive-compatibility. Let C^{aut} denote the allocation such that both agents produce and consume output in autarky, with no social insurance or income redistribution. This allocation satisfies feasibility, and it is incentive-compatible by construction. With C^{aut}, the infimum in agent 1's continuation utility $U_{1,0}$ occurs at the worst-case distribution $\underline{\pi}_{1,1}$, which places unit weight on all agents realizing the skill $\underline{\theta}$ at

[21]The agents' beliefs in this example do not satisfy Assumption 2, but they can be easily modified to do so by adding a single belief into all belief sets such that all agents receive $\bar{\theta}$ at $t=1$ and $\underline{\theta}$ at $t \geq 2$. This shows that competitive equilibria may not be efficient regardless of whether the sufficient overlap condition holds.

$t \geq 1$. By contrast, the infimum in agent 2's continuation utility $U_{2,0}$ trivially occurs at the best-case distribution $\bar{\pi}_{2,1}$, which places unit weight on all agents realizing the skill $\bar{\theta}$ at $t \geq 1$. As such, the government can strictly improve upon C^{aut} by shifting a small amount of consumption from agent 2 to agent 1 at $t = 1$ in the event that all agents report the skill $\underline{\theta}$. This change leaves agent 2's continuation utility unchanged, but it strictly raises agent 1's continuation utility. In addition, the new allocation remains incentive-compatible because agent 2 does not believe that any agent will realize the skill $\underline{\theta}$ at $t = 1$.

A competitive equilibrium cannot achieve the same $t = 0$ social welfare as the government, because Lemma 3 implies that report-contingent securities are not traded in equilibrium. As a result, to insure himself against realizing the worst skill $\underline{\theta}$ in future periods, agent 1 will seek to purchase risk-free bonds from agent 2. But this lowers agent 1's $t = 0$ instantaneous utility, whereas the adjustment made by the government does not. As a result, any competitive equilibrium cannot be constrained-efficient. □

Note that the result persists in the reinterpretation of competitive equilibria where firms trade securities among themselves, as long as firms have heterogeneous beliefs that are not all exactly the same as those of the government. To an extent, the result persists even when there is a continuum of agents in the economy and the Law of Large Numbers ensures that there is no aggregate risk. In particular, whenever firms or agents cannot trade securities contingent on the aggregate state of the economy, competitive equilibria with uncertainty may not be efficient.

4.2 Periodic reforms in equilibrium

Even with the potential for an equilibrium that is not efficient, agents may still obtain some degree of insurance in a decentralized economy. We next show that any insurance provided by a competitive equilibrium can be obtained with simplified allocations that are periodically reformed.

Lemma 3's arguments justify the focus here on the public information version of the economy. In the decentralized economy introduced above, at $t = 0$ each agent solves for a fully contingent allocation

$$C_i = \left\{c_{i,t}\left(s^t\right), z_{i,t}\left(s^t\right), k_{i,t+1}\left(s^t\right), b_{i,t+1}\left(s^t\right)\right\}_{t=0}^{T},$$

taking prices $Q\left(s^t\right)$ as given. Here $b_{i,t+1}$ denotes the agent's holdings of risk-free bonds that pay one unit of consumption after types are realized at $t + 1$. To see the result, it is helpful to again assume that agents' beliefs satisfy the intuitive sufficient overlap condition in Assumption 1. In this case, a decentralized version of Proposition 1 holds:

Proposition 5. *For any set of prices* $\{Q(s^t)\}_{t=0}^{T-1}$ *and any competitive equilibrium allocation* C*, there exists a sequence of allocations* $\{C^t\}_{t=0}^{T}$ *such that*

$$U_{i,0}(C|s_0) = U_{i,0}\left(C^0|s_0\right) \quad \forall s_0, \forall i.$$

Here an allocation C^t is a set of consumption, effective labor, capital, and bond functions that specify an agent's allocative actions at t and $t+1$. As before, each simplified allocation C^t also effectively provides a fallback option, i.e., a contingency plan for $\tau \geq t+2$ in the case that the agents do not choose to reform the allocation at $t+1$. We omit the proof since it follows essentially identical steps of the proof of Proposition 1. As in the government's problem, the key principle is that at period t, agents simply plan their allocations for t and $t+1$ under the assumption that all agents will receive the worst shock $\underline{\theta}$ at $\tau \geq t+2$.

5 Taking Simplicity Further

Finally, we consider whether even simpler, linear or affine policies can be optimal.[22] In particular, we ask when an optimal simplified allocation C^{*t} can be implemented with fiscal policies that are affine in individual income, $f(k_t, z_t)$. We argue that this is not generally the case: One must place strong assumptions on agents' beliefs and on allocations for affine policy functions to be optimal, and there are substantive limitations to generalizations. As an example, risk aversion provides a significant roadblock, as does elastically supplied labor. When linearity does persist, we show that it is with respect to an agent's skill shock, which is generally not equivalent to linearity in income.

It suffices to consider again the baseline economy of Section 2, and a strengthening of Assumption 1 as follows:[23]

[22] This linearity conjecture is motivated in part by the findings that, under some conditions, uncertainty can lead to linearity in financial contracting (see, e.g., Carroll 2015, Zhu 2016). A further, applied motivation is strongly suggested by a vast Ramsey optimal taxation literature that typically starts with a restriction to linear or affine policy instruments.

[23] As before, this assumption is stronger than necessary but more intuitive. The arguments below only require that under each $\pi_{i,t+1} \in \Pi_{i,t+1}$, $t+1$ skill shocks are independent and there exists $\pi'_{i,t+1} \in \Pi_{i,t+1}$ such that

1. $\pi_{i,t+1}$ and $\pi'_{i,t+1}$ have the same marginal distribution over $\theta_{-i,t+1}$ conditional on s^t,
2. the conditional marginal distribution of $\theta_{i,t+1}$ under $\pi'_{i,t+1}$ places weight only on $\{\underline{\theta}, \bar{\theta}\}$ to satisfy

$$\mathbb{E}_{\pi'_{i,t+1}}\left[\theta_{i,t+1}|s^t\right] = \mathbb{E}_{\pi_{i,t+1}}\left[\theta_{i,t+1}|s^t\right].$$

Assumption 3. *For any t, any i, any j, and any s^t, there exists $\theta_{j|i,t+1}(s^t)$ in the convex hull of Θ such that $\pi_{i,t+1} \in \Pi_{i,t+1}$ if and only if $\mathbb{E}_{\pi_{i,t+1}}[\theta_{j,t+1}| s^t] = \theta_{j|i,t+1}(s^t)$ for all j and the $t+1$ skill shocks $\theta_{1,t+1}, \ldots, \theta_{N,t+1}$ are independent under $\pi_{i,t+1}$.*

That is, at t agent i forms an expectation $\theta_{j|i,t+1}(s^t)$ about every agent j's $t+1$ skill $\theta_{j,t+1}$, and he considers any belief $\pi_{i,t+1}$ under which skills are independently distributed and the expected value of $\theta_{j,t+1}$ is $\theta_{j|i,t+1}(s^t)$ for all j. We require the independence to ensure that agent i's beliefs about other agents' skills do not change how he evaluates expectations with respect to his own skill. This is essential to the arguments, because the linearity conjecture concerns policies or continuation utilities that are linear with respect to agent i's skill, holding other agents' skills fixed.

The characterization of agents' beliefs in Assumption 3 satisfies Assumption 1, so an analogue of Proposition 1 holds: The efficient allocation C^* can be implemented by a sequence of simplified allocations $\{C^t\}_{t=0}^T$. The proof is almost identical to that of Proposition 1, so it is omitted. However, it is crucial to note that in the period-t allocation C^t, the consumption and effective labor functions $\left\{c_{i,\tau}^t, z_{i,\tau}^t\right\}_{\tau=t}^T$ for each agent i depend only on (s^t, θ_{t+1}). This holds because C^t can be constructed by assuming that all agents will realize beliefs $\underline{\Pi}_{t+2}$ and type shock $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ at $\tau \geq t+2$. We can then assume without loss of generality that the government implements C^* by designing a sequence of allocations $\{C^{*t}\}_{t=0}^T$, where each C^{*t} specifies a fallback option for subsequent periods and solves the government's period-t reform problem (2).

We focus on conditions under which the solution to the period-t reform problem features linear (affine) policy functions or linearity in agents' utilities with respect to their skills. We address the latter in detail, and at the end of the section we describe the additional assumptions needed for linear policy functions to be optimal. To state our result, we begin by examining the solution $C^{*0}(s_0)$ to the government's $t=0$ problem and considering how to modify it so that agent i's $t=1$ continuation utility is affine in his skill shock $\theta_{i,1}$. To simplify notation, we will suppress the $t=0$ superscript on all policy functions.

Notice that agent i's $t=1$ instantaneous utility

$$u\left(c_{i,1}^*(s_0, \theta_1), \frac{z_{i,1}^*(s_0, \theta_1)}{\theta_{i,1}}\right)$$

is not generally an affine function of his own skill $\theta_{i,1}$, holding fixed $(s_0, \theta_{-i,1})$. As such, this function typically does not coincide with its secant line from $\theta_{i,1} = \underline{\theta}$ to $\theta_{i,1} = \bar{\theta}$. Define new policy functions $\hat{c}_{i,1}, \hat{z}_{i,1}$ such that for any fixed $(s_0, \theta_{-i,1})$, $u(\hat{c}_{i,1}, \hat{z}_{i,1}/\theta_{i,1})$ is equal to $u\left(c_{i,1}^*, z_{i,1}^*/\theta_{i,1}\right)$ for $\theta_{i,1} \in \left\{\underline{\theta}, \bar{\theta}\right\}$ but is affine over Θ. Then $u(\hat{c}_{i,1}, \hat{z}_{i,1}/\theta_{i,1})$ is the function of the secant line of $u\left(c_{i,1}^*, z_{i,1}^*/\theta_{i,1}\right)$ from $\theta_{i,1} = \underline{\theta}$ to $\theta_{i,1} = \bar{\theta}$.

If feasible, the government would like to deliver utility $u\left(\hat{c}_{i,1}, \hat{z}_{i,1}/\theta_{i,1}\right)$ wherever $u\left(c^*_{i,1}, z^*_{i,1}/\theta_{i,1}\right)$ falls below its secant line. In particular, the government would like to use the policy functions $\tilde{c}_{i,1}, \tilde{z}_{i,1}$ defined by

$$\left(\tilde{c}_{i,1}\left(s_0,\theta_1\right), \tilde{z}_{i,1}\left(s_0,\theta_1\right)\right) = \begin{cases} \left(\hat{c}_{i,1}\left(s_0,\theta_1\right), \hat{z}_{i,1}\left(s_0,\theta_1\right)\right) & \text{if } u\left(\hat{c}_{i,1}\left(s_0,\theta_1\right), \frac{\hat{z}_{i,1}\left(s_0,\theta_1\right)}{\theta_{i,1}}\right) \\ & \geq u\left(c^*_{i,1}\left(s_0,\theta_1\right), \frac{z^*_{i,1}\left(s_0,\theta_1\right)}{\theta_{i,1}}\right), \\ \left(c^*_{i,1}\left(s_0,\theta_1\right), z^*_{i,1}\left(s_0,\theta_1\right)\right) & \text{else.} \end{cases}$$

Thus $\tilde{c}_{i,1}, \tilde{z}_{i,1}$ give the agent the maximum of the utilities $u\left(\hat{c}_{i,1}, \hat{z}_{i,1}/\theta_{i,1}\right)$ and $u\left(c^*_{i,1}, z^*_{i,1}/\theta_{i,1}\right)$.

For $t > 1$, we define $\hat{c}_{i,t}, \hat{z}_{i,t}$ and $\tilde{c}_{i,t}, \tilde{z}_{i,t}$ similarly to that above: Agent i's period t instantaneous utility under allocation C^{*0} is given by

$$u\left(c^*_{i,t}\left(s_0,\theta_1\right), \frac{z^*_{i,t}\left(s_0,\theta_1\right)}{\underline{\theta}}\right),$$

so we let $\hat{c}_{i,t}, \hat{z}_{i,t}$ be such that for fixed $\left(s_0, \theta_{-i,1}\right)$, $u\left(\hat{c}_{i,1}, \hat{z}_{i,1}/\underline{\theta}\right)$ is the function of the secant line of $u\left(c^*_{i,1}, z^*_{i,1}/\underline{\theta}\right)$ from $\theta_{i,1} = \underline{\theta}$ to $\theta_{i,1} = \bar{\theta}$. We then define $\tilde{c}_{i,t}, \tilde{z}_{i,t}$ analogously to how $\tilde{c}_{i,1}, \tilde{z}_{i,1}$ are defined. Finally, at $t = 0$ let

$$\left\{\tilde{c}_{i,0}, \tilde{z}_{i,0}\right\} = \left\{\hat{c}_{i,0}, \hat{z}_{i,0}\right\} = \left\{c^*_{i,0}, z^*_{i,0}\right\}.$$

Feasibility constraints may preclude the government from using the policy functions $\left\{\tilde{c}_{i,t}, \tilde{z}_{i,t}\right\}_{t=0}^{T}$, but these are clearly preferable to both $\left\{c^*_{i,t}, z^*_{i,t}\right\}_{t=0}^{T}$ and $\left\{\hat{c}_{i,t}, \hat{z}_{i,t}\right\}_{t=0}^{T}$.

We next show that the government is actually indifferent between the policy functions $\left\{\tilde{c}_{i,t}, \tilde{z}_{i,t}\right\}_{t=0}^{T}$ and $\left\{\hat{c}_{i,t}, \hat{z}_{i,t}\right\}_{t=0}^{T}$.

Proposition 6. *Let* $\hat{C}^0_i \equiv \left\{\hat{c}_{i,t}, \hat{z}_{i,t}, k^*_{i,t+1}\right\}_{t=0}^{T}$, $\tilde{C}^0_i \equiv \left\{\tilde{c}_{i,t}, \tilde{z}_{i,t}, k^*_{i,t+1}\right\}_{t=0}^{T}$. *Then*

$$U_{i,0}\left(\hat{C}^0_i \,|s_0\right) = U_{i,0}\left(\tilde{C}^0_i \,|s_0\right).$$

To see this, note first that to evaluate his $t = 0$ utility with allocation $\tilde{C}^0_i$, agent i considers the expected $t = 1$ continuation utility he would receive under each of his belief distributions $\pi_{i,1} \in \Pi_{i,1}$. When considered as a function of $\theta_{i,1}$ with fixed $\left(s_0, \theta_{-i,1}\right)$, his instantaneous utility at $t \geq 1$ lies weakly above its secant line from from $\theta_{i,1} = \underline{\theta}$ to $\theta_{i,1} = \bar{\theta}$, so the "worst" belief distributions $\pi'_{i,1} \in \Pi_{i,1}$ are those such that $\pi'_{i,1}\left(\cdot | \, s_0, \theta_{-i,1}\right)\big|_{\theta_i}$ places weight only on $\left\{\underline{\theta}, \bar{\theta}\right\}$ so as to satisfy

$$\mathbb{E}_{\pi'_{i,1}}\left[\theta_{i,1} | \, s_0\right] = \theta_{i|i,1}\left(s_0\right).$$

However, this only holds because of the independence of $\theta_{i,1}$ and $\theta_{-i,1}$ under each $\pi_{i,1} \in \Pi_{i,1}$. Without this independence, the distribution $\pi_{i,1}(\cdot|\, s_0, \theta_{-i,1})|_{\theta_i}$ may change depending on the other agents' $t = 1$ skill realizations $\theta_{-i,1}$, and it may not hold that agent i's expected $t = 1$ utility with the distribution $\pi_{i,1}(\cdot|\, s_0, \theta_{-i,1})|_{\theta_i}$ is weakly higher than his expected $t = 1$ utility with a distribution of the form $\pi'_{i,1}$. Intuitively, the independence of $\theta_{i,1}$ and $\theta_{-i,1}$ implies that $\theta_{-i,1}$ does not affect how agent i evaluates expectations with respect to his own skill.

Since agent i only considers distributions of the form $\pi'_{i,1} \in \Pi_{i,1}$ when evaluating his expected $t = 1$ utility, it is easy to see that he is indifferent between $\tilde{C}_i^0$ and $\hat{C}_i^0$. Under $\hat{C}_i^0$, the agent's $t \geq 1$ instantaneous utility is affine in $\theta_{i,1}$, and it coincides with the instantaneous utility under $\tilde{C}_i^0$ when $\theta_{i,1} \in \{\underline{\theta}, \bar{\theta}\}$. Distributions of the form $\pi'_{i,1}$ place weight only on events with $\theta_{i,1} \in \{\underline{\theta}, \bar{\theta}\}$, so agent i's expected $t = 1$ utility is the same under $\tilde{C}_i^0$ and $\hat{C}_i^0$. These allocations also give him the same $t = 0$ instantaneous utility, so the agent is indifferent between $\tilde{C}_i^0$ and $\hat{C}_i^0$.

Proposition 6 implies that the government weakly prefers the allocation $\hat{C}_i^0$, for which agent i's $t \geq 1$ instantaneous utility is affine in $\theta_{i,1}$, to C_i^{*0}. The government will always seek to design $t = 0$ allocations in which each agent's instantaneous utility at $t \geq 1$ is affine in his own $t = 1$ skill. The significance of this is that it demonstrates an interesting property of uncertainty aversion. However, the assumptions needed to prove this result are so strong that affine policies cannot be expected to be optimal in practice. For example, the proof of Proposition 6 makes heavy use of the independence condition in Assumption 3, but it is not clear why an uncertain agent would restrict his beliefs to product distributions. The use of this condition makes the result difficult to imagine in a setting without aggregate risk. In particular, the analogue of the independence condition with a continuum of agents would require a measureless agent to view himself as distinct from other agents with the same type, and it is difficult to justify this property. Another serious issue regards feasibility: We can clearly assume without loss of generality that the feasibility constraints in the $t = 0$ government's problem will bind, so the intermediate allocation $\tilde{C}_i^0$ constructed above is almost certain to be infeasible. In this case, the allocation $\hat{C}_i^0$ may also be infeasible.

Moreover, to obtain a more concrete property we must make the additional assumption that labor supply is inelastic. In particular, suppose that at $t = 0$, the government is constrained so that at each $t \geq 1$, agent i will exert some fixed amount of labor $\bar{l}_{i,t}(s_0)$. If this is the case, then agent i's $t = 1$ skill $\theta_{i,1}$ affects his $t \geq 1$ instantaneous utility only through consumption, and we can use the same methods as above to show that the government weakly prefers consumption functions that are affine in an agent's own skill. The following corollary collects these observations:

Corollary 1. *Fix i. If $t \geq 1$ labor supply is inelastic, agent i weakly prefers $\left\{c_{i,0}^{*}, \hat{c}_{i,t}\right\}_{t=1}^{T}$ to $\left\{c_{i,t}^{*}\right\}_{t=0}^{T}$, where $\hat{c}_{i,t}$ $t \geq 1$ is the affine consumption function given by*

$$\hat{c}_{i,t}\left(s_0, \theta_1\right) \equiv \frac{\bar{\theta} - \theta_{i,1}}{\bar{\theta} - \underline{\theta}} c_{i,t}^{*}\left(s_0, \left(\underline{\theta}, \theta_{-i,1}\right)\right) + \frac{\theta_{i,1} - \underline{\theta}}{\bar{\theta} - \underline{\theta}} c_{i,t}^{*}\left(s_0, \left(\underline{\theta}, \theta_{-i,1}\right)\right).$$

6 Concluding Remarks

This paper studied the optimal policy implications of moving away from the assumption of certainty about future distributions of idiosyncratic shocks, and characterized general properties of policies that are robust with respect to incomplete knowledge of stochastic elements of the economy. In otherwise conventional dynamic environments, we developed a simpler way to characterize general properties of optimal policies than under a narrower view of uncertainty limited to the standard notion of risk. Optimal policies themselves are also simplified, losing at times dependence on the full history of idiosyncratic shocks, and reformed periodically, consistent with what is commonly observed in reality. We argued, however, that restrictive assumptions are required for linear or even affine policies to be optimal. As a result of uncertainty, decentralized versions of these economies are not generally efficient, implying a potentially meaningful role for the government provision of insurance, in contrast with conventional environments without uncertainty.

While this paper focused on social insurance and fiscal policies, the above insights are applicable to risk-sharing environments more broadly. We believe that the paper also opens a number of interesting questions for future research. One significant possible extension is to environments where exogenous lack of commitment on the part of the agents is a friction that is more relevant or as relevant as private information about idiosyncratic attributes. It can still be considered an informational problem: Exogenous lack of commitment typically takes place when it is difficult or impossible for third parties to acquire information about contracting parties, so the enforcement of contracts by third parties is limited or absent. Applications include wage contracting between a firm and its workers who are free to walk away from long-term contracts, or informal insurance arrangements in village economies as well as other contexts in development economics. State-run (rather than federally-run) risk-sharing programs are a further example of broader risk-sharing environments where enforcement of contracts is difficult given relatively low costs of moving to a different state. In such applications, however, there is likely also private information about crucial individual economic attributes.

One approach to such situations is to characterize contracts that are self enforcing, i.e., providing incentives to stay within the contract. The algorithmic interpretation of our optimal reform problems naturally lends itself to this approach. In an optimal reform problem, a previous simplified allocation effectively provides an endogenous outside option, so self-enforcement constraints with exogenously specified outside options can be naturally accounted for as well. For example, by solving reform problems with exogenous self-enforcement constraints in each period, a sequence of simplified self-enforcing allocations can be characterized without the need to construct complete constrained-efficient allocations. Similarly, the optimal self-enforcing reform contracts are likely to feature significant simplifications. This is likely to provide novel insights about the optimal self-enforcing contracts.

References

[1] Acemoglu, Daron and Alp Simsek, "Moral Hazard and Efficiency in General Equilibrium with Anonymous Trading," working paper, August 2012.

[2] Aghion, Philippe, Ufuk Akcigit, Matthieu Lequien, and Stefanie Stantcheva, "Tax Simplicity and Heterogeneous Learning," NBER Working Paper, No. 24049, November 2017.

[3] Bergemann, Dirk, and Stephen Morris, "An Introduction to Robust Mechanism Design," Foundations and Trends in Microeconomics, 8 (3): 169-230, 2013.

[4] Bhandari, Anmol, "Doubts, Asymmetries, and Insurance," working paper, September 2015.

[5] Bhandari, Anmol, Jaroslav Borovička, and Paul Ho, "Survey data and subjective beliefs in business cycle models," working paper, November 2017.

[6] Bloom, Nicholas, "Fluctuations in Uncertainty," Journal of Economic Perspectives 28 (2):153-176, 2014.

[7] Carroll, Gabriel, "Robustness and Linear Contracts," American Economic Review, 105 (2), 536-63, 2015.

[8] Epstein, Larry, and Martin Schneider, "Recursive Multiple Priors," Journal of Economic Theory, 113 (1): 32-50, 2003.

[9] Farhi, Emmanual, Christopher Sleet, Iván Werning, and Sevin Yeltekin, “Non-linear Capital Taxation without Commitment,” Review of Economic Studies 79: 1469–1493, 2012.

[10] Farhi, Emmanuel and Iván Werning, “Insurance and Taxation over the Life Cycle,” Review of Economic Studies 80 (2): 596-635, 2013.

[11] Golosov, Mikhail and Luigi Iovino, “Social Insurance, Information Revelation, and Lack of Commitment,” working paper, March 2016.

[12] Golosov, Mikhail and Aleh Tsyvinski, “Optimal Taxation with Endogenous Insurance Markets,” Quarterly Journal of Economics 122 (2): 487-534, 2007.

[13] Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning, "New Dynamic Public Finance: A User’s Guide," NBER Macroeconomics Annual 21 (2006): 317-363, 2006.

[14] Hansen, Lars P. and Thomas J. Sargent, “Robust Control and Model Uncertainty,” American Economic Review 91 (2):60-66, 2001.

[15] Hansen, Lars P. and Thomas J. Sargent, “Time Inconsistency of Robust Control?” November 2006.

[16] Hansen, Lars P., Thomas J. Sargent, Gauhar A. Turmuhambetova, and Noah Williams, “Robust Control and Model Misspecification,” Journal of Economic Theory 128 (1):45-90, 2006.

[17] Harsanyi, John C., “Games with Incomplete Information Played by Bayesian Players,” Management Science, 14: 159-182, 320-334, 486-502, 1967-1968.

[18] Kapička, Marek, “Quantifying the Welfare Gains from History Dependent Income Taxation,” working paper, August 2017.

[19] Kocherlakota, Narayana R., “The New Dynamic Public Finance,” Princeton University Press, 2010.

[20] Kocherlakota, Narayana R. and Christopher Phelan, “On the Robustness of Laissez-Faire,” Journal of Economic Theory 144: 2372-2387, 2009.

[21] Kocherlakota, Narayana R. and Yangwei Song, “Public Goods with Ambiguity in Large Economies,” working paper, February 2018.

[22] Machina, Mark J. and Marciano Siniscalchi, "Ambiguity and Ambiguity Aversion," Handbook of the Economics of Risk and Uncertainty, Vol. 1, North-Holland, 729-807, 2014.

[23] Makris, Miltiadis and Alessandro Pavan, “Taxation under Learning-by-Doing,” working paper, March 2018.

[24] Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States," Quarterly Journal of Economics 133 (2): 553-609, 2018.

[25] Scheuer, Florian and Iván Werning, “Mirrlees meets Diamond-Mirrlees: Simplifying Nonlinear Income Taxation,” working paper, September 2016.

[26] Stantcheva, Stefanie, “Optimal Taxation and Human Capital Policies over the Life Cycle,” Journal of Political Economy 125 (6): 1931-1990, 2017.

[27] Williams, Noah, “Persistent Private Information,” Econometrica 79 (4): 1233-1275, 2011.

[28] Zhu, John Y., “Renegotiation of Dynamically Incomplete Contracts,” working paper, November 2016.

Online Appendix

A Proof Details Omitted in the Main Text

A.1 Proofs for Section 2

Proof. **(Proposition 1)** We begin by defining the period-0 allocation $C^0 = \{c_t^0, z_t^0, k_t^0\}_{t=0}^T$. We will show that each agent is indifferent between C^* and C^0, and that C^0 is simplified in the sense that it is not fully contingent on future states. First define

$$\left\{c_0^0, z_0^0, k_1^0\right\} \equiv \left\{c_0^*, z_0^*, k_1^*\right\}.$$

Thus C^0 and C^* coincide at $t = 0$.

For $t \geq 1$, define

$$\begin{aligned}
c_t^0\left(s^t\right) &\equiv c_t^*\left(\left(\theta_{i,0}, \Pi_{i,1}, \theta_{i,1}, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right), \\
z_t^0\left(s^t\right) &\equiv z_t^*\left(\left(\theta_{i,0}, \Pi_{i,1}, \theta_{i,1}, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right), \\
k_{t+1}^0\left(s^t\right) &\equiv k_{t+1}^*\left(\left(\theta_{i,0}, \Pi_{i,1}, \theta_{i,1}, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right).
\end{aligned}$$

At $t \geq 1$, regardless of agents' realized $t = 2$ beliefs $\Pi_{i,2}$ and $\tau \geq 2$ shocks $s_{i,\tau}$, C^0 allocates consumption, effective labor, and capital as if every agent realized the beliefs $\underline{\Pi}_2$ and period $\tau \geq 2$ shocks $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$.

To see that the agents are indifferent between C^0 and C^*, fix an initial state s_0 and let $\Pi'_{i,1} \subseteq \Pi_{i,1}$ be all distributions $\pi'_{i,1}$ supported on the set $\left\{s^1 \middle| \Pi_{j,2} = \underline{\Pi}_2 \, \forall j\right\}$. Since $\Pi_{i,1} \neq \emptyset$, Assumption 1 implies that $\Pi'_{i,1}$ is non-empty. We have

$$\begin{aligned}
U_{i,0}\left(C^* \middle| s_0\right) &= u\left(c_{i,0}^*\left(s_0\right), \frac{z_{i,0}^*\left(s_0\right)}{\theta_{i,0}}\right) + \beta \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[U_{i,1}\left(C^* \middle| s^1\right) \middle| s_0\right] \\
&\leq u\left(c_{i,0}^*\left(s_0\right), \frac{z_{i,0}^*\left(s_0\right)}{\theta_{i,0}}\right) + \beta \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[U_{i,1}\left(C^* \middle| s^1\right) \middle| s_0\right] \\
&= u\left(c_{i,0}^0\left(s_0\right), \frac{z_{i,0}^0\left(s_0\right)}{\theta_{i,0}}\right) + \beta \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[U_{i,1}\left(C^0 \middle| s^1\right) \middle| s_0\right] \\
&\leq u\left(c_{i,0}^0\left(s_0\right), \frac{z_{i,0}^0\left(s_0\right)}{\theta_{i,0}}\right) + \beta \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[U_{i,1}\left(C^0 \middle| s^1\right) \middle| s_0\right] \\
&= U_{i,0}\left(C^0 \middle| s_0\right).
\end{aligned}$$

The key step is the fourth line, which makes substantial use of Assumption 1: By definition, the $t \geq 1$ policy functions $\left\{c_t^0, z_t^0, k_{t+1}^0\right\}_{t=1}^T$ in C^0 depend only on the $t = 0$ state s_0 and $t = 1$ skills θ_1. By Assumption 1, any belief $\pi_{i,1} \in \Pi_{i,1}$ and its corresponding belief $\pi'_{i,1} \in \Pi'_{i,1}$ have the same marginal distribution over θ_1 conditional on s_0, so the distribution of future consumption and effective labor allocated the agent is the same under $\pi_{i,1}$ and $\pi'_{i,1}$. However, under the distribution $\pi_{i,1}$, it is possible that agent i could realize a skill $\theta_{i,t} \neq \underline{\theta}$ at some $t \geq 2$. With this realization, the agent's period-t continuation utility would be strictly higher than if he realized the skill $\underline{\theta}$. Given the recursive definition of $U_{i,1}$, this implies that the expected $t = 1$ continuation utility under $\pi'_{i,1}$ is weakly lower than the expected $t = 1$ continuation utility under $\pi_{i,1}$:

$$\mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right. \right) \right| s_0 \left. \right] \leq \mathbb{E}_{\pi_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right. \right) \right| s_0 \left. \right].$$

Assumption 1 implies that for each belief $\pi_{i,1} \in \Pi_{i,1}$, there exists a corresponding belief $\pi'_{i,1} \in \Pi'_{i,1}$, so we can take infima on both sides to find

$$\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right. \right) \right| s_0 \left. \right] \leq \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[U_{i,1} \left(C^0 \left| s^1 \right. \right) \right| s_0 \left. \right],$$

and the fourth line then follows.

Since C^* is a solution to the government's problem, the above inequality must be an equality for all i. Thus C^0 is also a solution to the government's problem. However, C^0 is not fully state-contingent, so the government may wish to modify the allocation if all agents do not realize the belief set $\underline{\Pi}_2$ at $t = 1$. The government remains fully committed to the $t = 0$ continuation utilities promised to each agent under C^0, so it can only reform if doing so does not reduce $t = 0$ continuation utilities. In this sense, the $t = 1$ continuation allocation $\left\{c_t^0, z_t^0, k_{t+1}^0, \right\}_{t=1}^T$ serves as a fallback allocation that the government can use if it cannot design a reform that raises $t = 1$ social welfare while fulfilling $t = 0$ promises. Note that we construct this fallback allocation based on what C^* prescribes when every agent realizes the worst productivity shock at each $t \geq 2$, so we are assured that it is feasible in the case that some agent j realizes a skill shock $\theta_{j,t} \neq \underline{\theta}$ at some $t \geq 2$.

In general, given a simplified allocation C^{t-1} designed at $t - 1$ and a period-t state s^t, the government with social welfare weights η designs the optimal simplified allocation

$C^{*t}\left(s^t, C^{t-1}\right)$ as the solution to the problem

$$\max_{C^t} \sum_i \eta_i U_{i,t}\left(C^t \left| s^t\right.\right)$$

subject to non-negativity and

$$\sum_i c^t_{i,\tau}\left(s^\tau\right) + K^t_{\tau+1}\left(s^\tau\right) \leq f\left(K^t_\tau\left(s^{\tau-1}\right), Z^t_\tau\left(s^\tau\right)\right),$$

$$U_{i,t-1}\left(C^{t-1}_{t-1}, \left(C^t_\tau\right)^T_{\tau=t} \left| s^{t-1}\right.\right) \geq U_{i,t-1}\left(C^{t-1} \left| s^{t-1}\right.\right)$$

Here the feasibility constraint holds for $\tau = t, \ldots, T$ and for all $s^\tau \geq s^t$, and the promise-keeping constraint holds for all i. This problem also applies when $t = 0$, but there is no promise-keeping constraint. By solving the problem at $t = 0, \ldots, T$, we can find a sequence of simplified allocations $\{C^{*t}\}^T_{t=0}$ that implements the efficient allocation C^*. □

A.2 Proofs for Section 3

Proof. **(Lemma 2)** For every t and every $\hat{s}^t$, define

$$\tilde{c}_t\left(\hat{s}^t\right) \equiv c_t\left(\hat{s}^{t-1}, \sigma^e_t\left(\hat{s}^{t-1}, \hat{s}^t\right)\right).$$

Define $\tilde{z}_t$ and $\tilde{k}_{t+1}$ analogously, and let $\tilde{C} = \left\{\tilde{c}_t, \tilde{z}_t, \tilde{k}_{t+1}\right\}^T_{t=0}$. With this definition, we have

$$U_{i,t}\big(\tilde{C} \left| \hat{s}^{t-1}, s^t_i\right.\big)\left(\sigma^*\right) = U_{i,t}\big(C \left| \hat{s}^{t-1}, s^t_i\right.\big)\left(\sigma^e\right)$$

for all i, all $\hat{s}^{t-1}$, and all s^t_i. Suppose that σ^* is not an equilibrium of $\tilde{C}$. Then there exist i, $\sigma_i \in \Sigma$, and $s_{i,0}$ such that

$$U_{i,0}\big(\tilde{C} \left| \hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma_i, \sigma^*_{-i}\right) > U_{i,0}\big(\tilde{C} \left| \hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma^*\right).$$

By the definition of $\tilde{C}$, this implies

$$U_{i,0}\big(C \left| \hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma^e_i \circ \sigma_i, \sigma^e_{-i}\right) > U_{i,0}\big(C \left| \hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma^e\right),$$

where $\sigma^e_i \circ \sigma_i \equiv \left\{\sigma^e_{i,t} \circ \sigma_{i,t}\right\}^T_{t=0}$ is the strategy formed by first applying the deviating strategy $\sigma_{i,t}$ to $\left(\hat{s}^{t-1}, s^t_i\right)$ and then applying the equilibrium strategy $\sigma^e_{i,t}$ to $\left(\hat{s}^{t-1}, \left(\hat{s}^{t-1}_i, \sigma_{i,t}\left(\hat{s}^{t-1}, s^t_i\right)\right)\right)$ at each t. This contradicts the assumption that σ^e is an equilibrium of C, so σ^* must be an equilibrium of $\tilde{C}$. □

Details of Proposition 3 Assumptions

It is useful to begin by describing an operator $I : \Delta(s^{t+1}) \to \underline{\underline{\Delta}}(s^{t+1})$ that will be used to define the simplified allocations $\{C^t\}_{t=0}^T$. For any belief distribution $\pi_{i,t+1}$, $I\pi_{i,t+1}$ is the distribution with the same marginal distribution over (s^{t-1}, θ_t), but which places unit weight on all agents receiving shock $(\underline{\theta}, \underline{\Pi}_{t+2})$ at $t+1$ and all agents realizing belief sets contained in $\underline{\underline{\Delta}}(s^{t+1})$ at t. More specifically, we have that for any $\pi_{i,t+1}$-measurable set E,

$$(I\pi_{i,t+1})(E) = \pi_{i,t+1}\left(\left\{ s^{t+1} \middle| \left(s_j^{t-1}, (\theta_{j,t}, I\Pi_{j,t+1}), (\underline{\theta}, \underline{\Pi}_{t+2})\right)_{j=1}^N \in E \right\}\right).$$

Here we have written $I\Pi_{j,t+1} \equiv \{I\pi_{j,t+1} | \pi_{j,t+1} \in \Pi_{j,t+1}\}$. This description of I is circular, but the operator can be defined precisely by viewing each $\pi_{i,t+1}$ as representing an infinite hierarchy of beliefs (see, e.g., Harsanyi 1967-1968). I then operates on hierarchies by shifting beliefs of all orders so that under any distribution $I\pi_{i,t+1}$, it is common knowledge that all agents believe everyone will realize the shock $(\underline{\theta}, \underline{\Pi}_{t+2})$ at $t+1$. Significantly, this definition implies that I is also the identity transformation on $\underline{\underline{\Delta}}(s^{t+1})$.

With the operator I, we can state our Assumption 2 more precisely as follows: For any $t = 0, \ldots, T-1$, any i, any $\hat{s}^{t-1}$, any s_i^t, and any belief $\pi_{i,t+1} \in \Pi_{i,t+1}$, there exists another belief $\pi'_{i,t+1} \in \Pi_{i,t+1}$ such that for any $\pi_{i,t+1}$-measurable set E,

$$\pi'_{i,t+1}\left(E \middle| \hat{s}^{t-1}, s_i^t\right) = \pi_{i,t+1}\left(\left\{ s^{t+1} \middle| \left(s^t, (\theta_{j,t+1}, I\Pi_{j,t+2})_{j=1}^N\right) \in E \right\} \middle| \hat{s}^{t-1}, s_i^t\right). \tag{10}$$

Thus $\pi'_{i,t+1}$ is essentially the same as the distribution $\pi_{i,t+1}$ conditional on $(\hat{s}^{t-1}, s_i^t)$, but shifted using I to place unit weight on all agents realizing worst beliefs $I\Pi_{j,t+2} \subset \underline{\underline{\Delta}}(s^{t+2})$ at $t+1$.

It will also be useful to give a technical description of weak monotonicity. Given a strategy $\sigma_i \in \Sigma$, let σ_i^t denote the strategy defined by

$$\sigma_{i,\tau}^t\left(\hat{s}^{\tau-1}, s_i^\tau\right) \equiv \begin{cases} \sigma_{i,\tau}\left(\hat{s}^{\tau-1}, s_i^\tau\right) & \tau \le t \\ \sigma_{i,\tau}\left(\hat{s}^{\tau-1}, \left(s_i^{\tau-1}, (\theta_{i,\tau}, I\Pi_{i,\tau+1})\right)\right) & \tau = t+1 \\ \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right) & \tau \ge t+2 \end{cases} \tag{11}$$

In the above definition, $s_i^{\tau-1}$ is the vector of shocks up to period $\tau - 1$ for type s_i^τ, and $(\theta_{i,\tau}, \Pi_{i,\tau+1})$ is the period-τ shock for type s_i^τ. Thus σ_i^t coincides with σ_i through period t, and at period $\tau \ge t+2$, the strategy always reports shock $(\underline{\theta}, \underline{\Pi}_{\tau+1})$. At $t+1$, an agent using reporting strategy σ_i^t reports as if he realized belief set $I\Pi_{i,t+2}$ instead of $\Pi_{i,t+2}$ and if he were using the reporting strategy σ_i. Intuitively, σ_i^t is formed from σ_i by shifting the shocks $(\theta_{i,t}, \Pi_{i,t+1})$ for each t in precisely the way that they are shifted to define $\pi'_{i,t+1}$ from $\pi_{i,t+1}$.

With this definition, an allocation C is weakly monotonic at t if for any agent i, any $t-1$ reported state $\hat{s}^{t-1}$, any t type s_i^t, any strategy $\sigma_i \in \Sigma$, and any belief distribution $\pi_{i,t+1} \in \Pi_{i,t+1}$,

$$\begin{aligned} &\mathbb{E}_{\pi_{i,t+1}} \left[W_{i,t}\big(C \left| \hat{s}^{t-1}, s^t \right.\big)\big(\sigma^*_{-i}, \sigma_i\big) \big| \, \hat{s}^{t-1}, s_i^t \right] \\ &\geq \mathbb{E}_{\pi_{i,t+1}} \left[W_{i,t}\big(C \left| \hat{s}^{t-1}, s^t \right.\big)\big(\sigma^{*t}_{-i}, \sigma_i\big) \big| \, \hat{s}^{t-1}, s_i^t \right]. \end{aligned}$$

This condition implies that if the $-i$ agents report truthfully, then regardless of the reporting strategy used by agent i, his expected utility at $t+1$ under any belief distribution $\pi_{i,t+1}$ is weakly greater than his expected utility when the $-i$ agents report belief sets $I\Pi_{j,t+2} \subset \underline{\underline{\Delta}}\left(s^{t+2}\right)$ at $t+1$ and the worst shock $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ at $\tau \geq t+2$.

Proof. **(Proposition 3)**

To prove the Proposition, we will construct a $t = 0$ simplified allocation C^0 such that each agent is indifferent between C^0 and C^*. We will prove that this new allocation satisfies the constraints of the government's problem (7) and then explain how to iterate this process in each period to find a sequence of simplified allocations $\{C^t\}_{t=0}^T$ that implements C^*.

We start by defining $C^0 = \{c_t^0, z_t^0, k_t^0\}_{t=0}^T$. Let

$$\left\{c_0^0, z_0^0, k_1^0\right\} \equiv \left\{c_0^*, z_0^*, k_1^*\right\},$$

so that C^0 and C^* coincide at $t = 0$. For $t \geq 1$, define

$$\begin{aligned} c_t^0\left(\hat{s}^t\right) &\equiv c_t^*\left(\left(\hat{\theta}_{i,0}, \hat{\Pi}_{i,1}, \hat{\theta}_{i,1}, I\hat{\Pi}_{i,2}, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right), \\ z_t^0\left(\hat{s}^t\right) &\equiv z_t^*\left(\left(\hat{\theta}_{i,0}, \hat{\Pi}_{i,1}, \hat{\theta}_{i,1}, I\hat{\Pi}_{i,2}, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right), \\ k_{t+1}^0\left(\hat{s}^t\right) &\equiv k_{t+1}^*\left(\left(\hat{\theta}_{i,0}, \hat{\Pi}_{i,1}, \hat{\theta}_{i,1}, I\hat{\Pi}_{i,2}, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right)_{i=1}^N\right). \end{aligned}$$

Thus at $t \geq 1$, regardless of agents' $t = 2$ reported beliefs $\hat{\Pi}_{i,2}$ and $\tau \geq 2$ reported shocks $\hat{s}_{i,\tau}$, C^0 allocates consumption, effective labor, and capital as if agents reported belief sets $I\hat{\Pi}_{i,2} \subset \underline{\underline{\Delta}}\left(s^2\right)$ and period $\tau \geq 2$ shocks $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$.

Under the truth-telling equilibrium profile σ^*, all agents weakly prefer C^0 to C^*. To see this, fix any initial state s_0 and any i. Let $\Pi'_{i,1} \subset \Pi_{i,1}$ be the set of all shifted distributions

$\pi'_{i,1}$ defined by (10). Since $\Pi_{i,1} \neq \emptyset$, Assumption 2 implies that $\Pi'_{i,1}$ is non-empty. We have

$$\begin{aligned} U_{i,0}\big(C^* \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big) &= \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right] \\ &\leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right] \\ &= \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right] \\ &\leq \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right] \\ &= U_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big). \end{aligned} \tag{12}$$

The second line holds because $\Pi'_{i,1} \subset \Pi_{i,1}$, and the third line holds because C^0 and C^* coincide when agents report beliefs $\underline{\underline{\hat{\Pi}}}_{j,2}$ at $t = 1$ and the worst shock $\big(\underline{\theta}, \underline{\Pi}_{t+1}\big)$ at $t \geq 2$. The fourth line follows from Assumption 2 by the following argument: Fix $\pi_{i,1} \in \Pi_{i,1}$. Because agents have private information over shocks, agent i will use his next-period belief set $\Pi_{i,2}$ to evaluate expectations with respect to the $-i$ agents' $t = 1$ shocks $s_{-i,1}$. As a result, the expectation

$$\mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right]$$

depends on $\pi_{i,1}$ only through its marginal distribution over $(\theta_0, \Pi_1, \theta_{i,1}, \Pi_{i,2})$, conditional on $s_{i,0}$. The policy functions in C^0 are defined such that regardless of any agent j's reports at $t \geq 1$, C^0 allocates as if he reported beliefs $I\hat{\Pi}_{j,2} \subset \underline{\underline{\Delta}}\left(s^2\right)$ at $t = 1$ and shock $\big(\underline{\theta}, \underline{\Pi}_{t+1}\big)$ at $t \geq 2$, so the distribution of future consumption and effective labor allocated to agent i is the same under $\pi_{i,1}$ and $\pi'_{i,1}$. (Note that this uses the fact that I leaves belief sets $\underline{\underline{\hat{\Pi}}}_{j,2} \subset \underline{\underline{\Delta}}\left(s^2\right)$ unchanged.) However, under the distribution $\pi_{i,1}$, it is possible that agent i could realize a skill $\theta_{i,t} \neq \underline{\theta}$ at some $t \geq 2$. In this case, the agent's period-t continuation utility is weakly higher than if he realized skill $\underline{\theta}$. By the recursive definition of $W_{i,0}$, this implies that the above expectation is weakly greater than the expectation with respect to the shifted distribution $\pi'_{i,1}$:

$$\mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right].$$

Assumption 2 requires that $\Pi'_{i,1} \subset \Pi_{i,1}$ contains a shifted distribution $\pi'_{i,t+1}$ for each distribution $\pi_{i,t+1} \in \Pi_{i,t+1}$, so we have

$$\begin{aligned} &\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right] \\ &\leq \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*\big) \middle| \, s_{i,0} \right], \end{aligned}$$

and the fourth line in (12) follows.

We next show that the simplified allocation C^0 satisfies the constraints of the government's problem (7) and that all agents are actually indifferent between C^0 and C^*. It is clear that C^0 satisfies non-negativity and feasibility, and we now demonstrate that C^0 is incentive-compatible. To see this, consider agent i's problem at $t = 0$:

$$\max_{\sigma_i \in \Sigma} U_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*_{-i}, \sigma_i\big). \tag{13}$$

The allocation C^0 is defined such that it does not depend on reports $\hat{s}_{i,t}$ for $t \geq 2$, and it assumes that the agent will report beliefs $\hat{\underline{\underline{\Pi}}}_{i,2} \subset \underline{\underline{\Delta}}\left(s^2\right)$ at $t = 1$. As a result, we can without loss of generality restrict the agent to the subset of strategies $\Sigma' \subset \Sigma$ such that for any $\sigma_i \in \Sigma'$, the image of $\sigma_{i,t}$ is $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$ for $t \geq 2$ and the image of $\sigma_{i,1}$ is contained in the set $\left\{ \left(\hat{\theta}_{i,1}, \hat{\underline{\underline{\Pi}}}_{i,2}\right) \Big| \hat{\underline{\underline{\Pi}}}_{i,2} \subset \underline{\underline{\Delta}}\left(s^2\right)\right\}$. Substituting the definition of $U_{i,0}$, (13) is then equivalent to

$$\max_{\sigma_i \in \Sigma'} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big)\big|\, s_{i,0}\right]. \tag{14}$$

Since $\Pi'_{i,1} \subset \Pi_{i,1}$, we can replace $\Pi_{i,1}$ by $\Pi'_{i,1}$ in the infimum to see that (14) is bounded above by

$$\max_{\sigma_i \in \Sigma'} \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big)\big|\, s_{i,0}\right]. \tag{15}$$

Recall that C^0 is defined such that it coincides with C^* at $t = 0$, and it allocates as C^* does at $t \geq 1$, but assuming that all agents report beliefs contained in $\underline{\underline{\Delta}}\left(s^2\right)$ at $t = 1$ and type shock $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$ at $t \geq 2$. We can replace the C^0 policy functions with the C^* policy functions since the support of any belief in $\Pi'_{i,1}$ and the range of any reporting strategy in Σ' are restricted to $t = 1$ belief sets contained in $\underline{\underline{\Delta}}\left(s^2\right)$ and the $t \geq 2$ type shock $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$. Then (15) is equal to

$$\max_{\sigma_i \in \Sigma'} \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big)\big|\, s_{i,0}\right]. \tag{16}$$

Since $\Sigma' \subset \Sigma$, we can replace Σ' by Σ in the maximum to find that (16) is bounded above by

$$\max_{\sigma_i \in \Sigma} \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big)\big|\, s_{i,0}\right]. \tag{17}$$

Below, we show that because C^* is weakly monotonic at $t = 0$, we can replace $\Pi'_{i,1}$ with $\Pi_{i,1}$ in the infimum to find that (17) is equal to

$$\begin{aligned} &\max_{\sigma_i \in \Sigma} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma_i\big) \big|\, s_{i,0}\right] \\ &= \max_{\sigma_i \in \Sigma} U_{i,0}\big(C^* \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*_{-i}, \sigma_i\big) \\ &= U_{i,0}\big(C^* \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big), \end{aligned} \tag{18}$$

where the last line holds by the incentive-compatibility of C^*. By (12),

$$U_{i,0}\big(C^* \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big) \leq U_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big), \tag{19}$$

and we can then conclude

$$\max_{\sigma_i \in \Sigma} U_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*_{-i}, \sigma_i\big) \leq U_{i,0}\big(C^0 \,|\hat{s}_{-1}, s_{i,0}\big)\big(\sigma^*\big), \tag{20}$$

so C^0 is incentive-compatible. This additionally implies that (19) must hold with equality because (20) holds with equality, so we also have that all agents are indifferent between C^0 and C^*.

To justify the equality between (17) and (18), note immediately that (17) is weakly greater than (18) because $\Pi'_{i,1} \subset \Pi_{i,1}$. To establish the reverse inequality, fix $\sigma_i \in \Sigma$ and $\pi_{i,1} \in \Pi_{i,1}$. If σ^1_i and σ^{*1}_{-i} are defined as in (11), then the weak monotonicity of C^* at $t = 0$ implies that

$$\mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^*_{-i}, \sigma^1_i\big) \big|\, s_{i,0}\right] \tag{21}$$

is weakly bounded below by

$$\mathbb{E}_{\pi_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^{*1}_{-i}, \sigma^1_i\big) \big|\, s_{i,0}\right]. \tag{22}$$

Let $\pi'_{i,1} \in \Pi_{i,1}$ be the distribution that corresponds to $\pi_{i,1}$ in the sense of (10). We claim that (22) is weakly bounded below by

$$\mathbb{E}_{\pi'_{i,1}} \left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\big(\sigma^{*1}_{-i}, \sigma^1_i\big) \big|\, s_{i,0}\right]. \tag{23}$$

To see this, note that for any t, the distribution of reported histories $\hat{s}^t$ implied by $\pi_{i,1}$ and $\big(\sigma^{*1}_{-i}, \sigma^1_i\big)$ is the same as with $\pi'_{i,1}$. This holds because the strategy profile $\big(\sigma^{*1}_{-i}, \sigma^1_i\big)$ is formed by shifting the profile $\big(\sigma^*_{-i}, \sigma_i\big)$ in precisely the same way that $\pi_{i,1}$ is shifted to form $\pi'_{i,1}$. Thus agent i's beliefs about the consumption and effective labor he will be allocated in each period are the same under $\pi_{i,1}$ and $\pi'_{i,1}$. However, under the distribution $\pi_{i,1}$, it is possible that agent i believes he could realize a skill $\theta_{i,t} \neq \underline{\theta}$ in some period $t \geq 2$. In this case, agent i's period-t continuation payoff $W_{i,t}$ is strictly greater than if he realized skill $\underline{\theta}$. The distribution of agent i's $t = 1$ skill $\theta_{i,1}$ is the same under $\pi_{i,1}$ and $\pi'_{i,1}$, so we must have that (22) is weakly bounded below by (23).

Now since $\pi'_{i,1}$ incorporates the shift used to define $\left(\sigma^{*1}_{-i}, \sigma^1_i\right)$ from $\left(\sigma^*_{-i}, \sigma_i\right)$, we can replace the strategy profile in (23) to see that (23) is equal to

$$\mathbb{E}_{\pi'_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma_i\right)\big|\, s_{i,0}\right]. \tag{24}$$

With expressions (21)-(24), we have the inequality

$$\begin{aligned}&\mathbb{E}_{\pi_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma^1_i\right)\big|\, s_{i,0}\right]\\ &\geq \mathbb{E}_{\pi'_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma_i\right)\big|\, s_{i,0}\right].\end{aligned}$$

But since $\pi_{i,1} \in \Pi_{i,1}$ was arbitrary and $\pi'_{i,1} \in \Pi'_{i,1}$ by Assumption 2, we can take infima on both sides to find

$$\begin{aligned}&\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma^1_i\right)\big|\, s_{i,0}\right]\\ &\geq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma_i\right)\big|\, s_{i,0}\right],\end{aligned}$$

and since $\sigma_i \in \Sigma$ was arbitrary, $\sigma^1_i \in \Sigma$ implies that we can maximize both sides over Σ to find

$$\begin{aligned}&\max_{\sigma_i \in \Sigma} \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma_i\right)\big|\, s_{i,0}\right]\\ &\geq \max_{\sigma_i \in \Sigma} \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[W_{i,0}\big(C^* \,|\hat{s}_{-1}, s_0\big)\left(\sigma^*_{-i}, \sigma_i\right)\big|\, s_{i,0}\right].\end{aligned}$$

This is precisely the inequality needed to conclude that (17) and (18) are equal. We have thus shown that the simplified allocation C^0 satisfies the constraints of the government's problem (7). Since all agents are indifferent between C^0 and C^*, we conclude that C^0 is also a solution to the government's problem (7).

Now C^0 is not fully state-contingent, so it may no longer be optimal when the government receives new beliefs $\Pi_{g,2}$ at $t = 1$. The government will then seek to construct a reform allocation C^1 to raise $t = 1$ social welfare, while remaining fully committed to the utility promised to agents at $t = 0$ as well as the incentives it provided for truthful revelation.

Because of this commitment, the reform C^1 must not reduce agents' $t=0$ utility relative to C^0, and it must enforce the threats made by the government at $t=0$ to lower the utility of any agent who reported untruthfully. In particular, given the simplified allocation C^0 designed at $t=0$, the truthfully-reported $t=0$ state s_0, and the $t=1$ type s_g^1 of the governing agent, the government seeks to solve the problem

$$\max_{C^1}\inf_{\Pi_{g,2}}\mathbb{E}_{\pi_{g,2}}\left[\sum_i \eta_i U_{i,1}\big(C^1\,\big|s_0,s_i^1\big)\big(\sigma^*\big)\bigg|\,s_0,s_g^1\right] \tag{25}$$

subject to non-negativity and

$$\sum_i c_{i,t}^1\left(\hat{s}^t\right)+K_{t+1}^1\left(\hat{s}^t\right)\leq f\left(K_t^1\left(\hat{s}^{t-1}\right),Z_t^1\left(\hat{s}^t\right)\right),$$

$$U_{i,1}\big(C^1\,\big|\hat{s}_0,s_i^1\big)\big(\sigma^*\big)\geq U_{i,1}\big(C^1\,\big|\hat{s}_0,s_i^1\big)\big(\sigma_{-i}^*,\sigma_i\big),$$

$$U_{i,0}\big(C_0^0,\big(C_\tau^1\big)_{\tau=1}^T\,\big|\hat{s}^{-1},s_{i,0}\big)\big(\sigma^*\big)\geq U_{i,0}\big(C^0\,\big|\hat{s}^{-1},s_{i,0}\big)\big(\sigma^*\big)$$

$$U_{i,0}\big(C_0^0,\big(C_\tau^1\big)_{\tau=1}^T\,\big|\hat{s}^{-1},s_{i,0}\big)\big(\sigma_{-i,0}^*,\sigma_{i,0},\big(\sigma_\tau^*\big)_{\tau=1}^T\big)$$
$$\leq U_{i,0}\big(C^0\,\big|\hat{s}^{-1},s_{i,0}\big)\big(\sigma_{-i,0}^*,\sigma_{i,0},\big(\sigma_\tau^*\big)_{\tau=1}^T\big).$$

The feasibility constraint holds for $t\geq 1$ and for all $\hat{s}^t$, and the $t=1$ incentive-compatibility constraint holds for all i, all $\sigma_i\in\Sigma$, all $\hat{s}_0$, and all s^1. The promise-keeping constraint holds for all i and all s_0, and the threat-keeping constraint holds for all i, all s_0, and all $\sigma_{i,0}$. If the constraint set is non-empty, then the government chooses the allocation that solves problem (25). In this case, we can use arguments similar to those above to show that a simplified allocation is in the solution set.

In particular, let C^{*1} solve problem (25). Define the simplified allocation $C^1=\{c_t^1,z_t^1,k_t^1\}_{t=1}^T$ analogously to how C^0 was defined from C^*: Let

$$\left\{c_1^1,z_1^1,k_2^1\right\}\equiv\left\{c_1^{*1},z_1^{*1},k_2^{*1}\right\},$$

so that C^1 and C^{*1} coincide at $t=1$. For $t\geq 2$, define

$$c_t^1\left(\hat{s}^t\right)\equiv c_t^{*1}\left(\left(\left(\hat{\theta}_{i,\tau},\hat{\Pi}_{i,\tau+1}\right)_{\tau=0}^1,\hat{\theta}_{i,2},I\hat{\Pi}_{i,3},\left(\underline{\theta},\underline{\Pi}_{\tau+1}\right)_{\tau=3}^t\right)_{i=1}^N\right),$$

$$z_t^1\left(\hat{s}^t\right)\equiv z_t^{*1}\left(\left(\left(\hat{\theta}_{i,\tau},\hat{\Pi}_{i,\tau+1}\right)_{\tau=0}^1,\hat{\theta}_{i,2},I\hat{\Pi}_{i,3},\left(\underline{\theta},\underline{\Pi}_{\tau+1}\right)_{\tau=3}^t\right)_{i=1}^N\right),$$

$$k_{t+1}^1\left(\hat{s}^t\right)\equiv k_{t+1}^{*1}\left(\left(\left(\hat{\theta}_{i,\tau},\hat{\Pi}_{i,\tau+1}\right)_{\tau=0}^1,\hat{\theta}_{i,2},I\hat{\Pi}_{i,3},\left(\underline{\theta},\underline{\Pi}_{\tau+1}\right)_{\tau=3}^t\right)_{i=1}^N\right).$$

Thus at $t \geq 2$, regardless of agents' $t = 3$ reported beliefs $\hat{\Pi}_{i,3}$ and $\tau \geq 3$ reported shocks $\hat{s}_{i,\tau}$, C^1 allocates consumption, effective labor, and capital as if agents reported belief sets $I\hat{\Pi}_{i,3} \subset \underline{\underline{\Delta}}(s^3)$ and period $\tau \geq 3$ shocks $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$.

The simplified allocation C^1 clearly satisfies non-negativity and feasibility, and since we assume that C^{*1} is weakly monotonic at $t = 1$, the same argument as for C^0 implies that C^1 is incentive-compatible at $t = 1$. This argument also implies that all agents are indifferent between C^{*1} and C^1 at $t = 1$ under the truth-telling strategy profile. To see that C^1 satisfies the promise-keeping constraint (the third constraint in problem 25), note first that the $t = 1$ analogue of (12) implies that for all i, all $\hat{s}_0$, and all s_i^1,

$$U_{i,1}\left(C^{*1}\left|\hat{s}_0, s_i^1\right.\right)(\sigma^*) \leq U_{i,1}\left(C^1\left|\hat{s}_0, s_i^1\right.\right)(\sigma^*). \tag{26}$$

Given the allocation $\left(C_0^0, (C_t^1)_{t=1}^T\right)$, the strategy profile σ^*, and the $t = 0$ type $s_{i,0}$, consider the $t = 0$ utility of agent i:

$$\begin{aligned} &U_{i,0}\left(C_0^0, \left(C_\tau^1\right)_{\tau=1}^T\left|\hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^*) \\ &= \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[u\left(c_0^0(s_0), \frac{z_0^0(s_0)}{\theta_{i,0}}\right) + \beta U_{i,1}\left(C^1\left|s_0, s_i^1\right.\right)(\sigma^*)\middle| s_{i,0}\right]. \end{aligned} \tag{27}$$

By setting $\hat{s}_0 = s_0$, inequality (26) implies that (27) is bounded below by

$$\begin{aligned} &\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[u\left(c_0^0(s_0), \frac{z_0^0(s_0)}{\theta_{i,0}}\right) + \beta U_{i,1}\left(C^{*1}\left|s_0, s_i^1\right.\right)(\sigma^*)\middle| s_{i,0}\right] \\ &= U_{i,0}\left(C_0^0, \left(C_\tau^{*1}\right)_{\tau=1}^T\left|\hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^*). \end{aligned} \tag{28}$$

But C^{*1} satisfies the promise-keeping constraint since it is a solution to problem (25), so

$$U_{i,0}\left(C_0^0, \left(C_\tau^{*1}\right)_{\tau=1}^T\left|\hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^*) \geq U_{i,0}\left(C^0\left|\hat{s}^{-1}, s_{i,0}\right.\right)(\sigma^*). \tag{29}$$

Expressions (27)-(29) then imply that C^1 satisfies the promise-keeping constraint.

Finally, we show that the simplified allocation C^1 also satisfies the threat-keeping constraint (the fourth constraint in problem 25). Given the allocation $\left(C_0^0, (C_t^1)_{t=1}^T\right)$, the strategy profile $\left(\sigma_{-i,0}^*, \sigma_{i,0}, (\sigma_\tau^*)_{\tau=1}^T\right)$, and the $t = 0$ type $s_{i,0}$, consider the $t = 0$ utility of agent

i:

$$U_{i,0}\big(C_0^0, \left(C_\tau^1\right)_{\tau=1}^T \left|\hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma_{-i,0}^*, \sigma_{i,0}, \left(\sigma_\tau^*\right)_{\tau=1}^T\right) \tag{30}$$
$$= \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[u\left(c_0^0\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right), \frac{z_0^0\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right)}{\theta_{i,0}}\right)\right.$$
$$\left. + \ \beta U_{i,1}\big(C^1 \left|\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right), s_i^1\right.\big)\left(\sigma^*\right) \right| s_{i,0}\Big].$$

The argument for the incentive-compatibility of C^{*1} implies that inequality (26) holds with equality, so we can replace C^1 with C^{*1} in the argument of $U_{i,1}$ to find that (30) is equal to

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}} \left[u\left(c_0^0\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right), \frac{z_0^0\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right)}{\theta_{i,0}}\right)\right. \tag{31}$$
$$\left. + \ \beta U_{i,1}\big(C^{*1} \left|\left(s_{-i,0}, \sigma_{i,0}\left(\hat{s}^{-1}, s_{i,0}\right)\right), s_i^1\right.\big)\left(\sigma^*\right) \right| s_{i,0}\Big].$$
$$= U_{i,0}\big(C_0^0, \left(C_\tau^{*1}\right)_{\tau=1}^T \left|\hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma_{-i,0}^*, \sigma_{i,0}, \left(\sigma_\tau^*\right)_{\tau=1}^T\right).$$

But C^{*1} satisfies the threat-keeping constraint since it is a solution to problem (25), so

$$U_{i,0}\big(C_0^0, \left(C_\tau^{*1}\right)_{\tau=1}^T \left|\hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma_{-i,0}^*, \sigma_{i,0}, \left(\sigma_\tau^*\right)_{\tau=1}^T\right) \tag{32}$$
$$\leq U_{i,0}\big(C^0 \left|\hat{s}^{-1}, s_{i,0}\right.\big)\left(\sigma_{-i,0}^*, \sigma_{i,0}, \left(\sigma_\tau^*\right)_{\tau=1}^T\right).$$

Expressions (30)-(32) then imply that C^1 also satisfies the threat-keeping constraint. We have shown that the simplified allocation C^1 is in the constraint set of problem (25), so inequality (26) implies that it is also a solution to problem (25). The incentive-compatibility argument implies that (26) is actually satisfied with equality, so we have

$$U_{i,1}\big(C^{*1} \left|\hat{s}_0, s_i^1\right.\big)\left(\sigma^*\right) = U_{i,1}\big(C^1 \left|\hat{s}_0, s_i^1\right.\big)\left(\sigma^*\right)$$

for all i, all $\hat{s}_0$, and all s_i^1.

Now C^0 is not necessarily incentive-compatible with respect to the $t = 1$ utility functions $U_{i,1}$, so the constraint set for problem (25) may be empty. In this case, the government is not able to reform the allocation C^0, so we set $C^1 \equiv C^0$. The government then recommends that the agents play the welfare-maximizing equilibrium σ^{*e} of C^1, i.e., the equilibrium σ^e that maximizes

$$\inf_{\Pi_{g,2}} \mathbb{E}_{\pi_{g,2}} \left[\sum_i \eta_i U_{i,1}\big(C^1 \left|s_0, s_i^1\right.\big)\left(\sigma^e\right) \right| s_0, s_g^1\Bigg].$$

In all subsequent periods, the government follows essentially the same reform process, thereby constructing a sequence of simplified allocations $\{C^t\}_{t=0}^T$ that implements C^*: Whenever possible, the government designs a reform allocation C^t that raises t social welfare relative to the previous allocation. If the government cannot design such an allocation while maintaining full commitment to previously promised utilities and incentives, then the government leaves the old allocation in place and suggests an equilibrium for the agents to play at t. □

Optimal Reform Problems ($t \geq 0$)

In general, at $t \geq 0$ the government seeks to design a reform allocation that maximizes t continuation social welfare. Suppose the allocation was last reformed in period $r < t$. Let $\tilde{s}^{t-1}$ denote the $t-1$ reported state, and let s_g^t be the t type of the governing agent. Then at t, the government seeks to solve the following problem:

$$\max_{C^t} \inf_{\Pi_{g,t+1}} \mathbb{E}_{\pi_{g,t+1}} \left[\sum_i \eta_i U_{i,t}\big(C^t \,\big|\tilde{s}^{t-1}, s_i^t\big)\big(\sigma^*\big) \middle| \tilde{s}^{t-1}, s_g^t \right] \tag{33}$$

$$\begin{aligned}
&\text{subject to non-negativity and}\\
&\sum_i c_{i,\tau}^t\left(\hat{s}^\tau\right) + K_{\tau+1}^t\left(\hat{s}^\tau\right) \leq f\left(K_\tau^t\left(\hat{s}^{\tau-1}\right), Z_\tau^t\left(\hat{s}^\tau\right)\right),\\
&U_{i,t}\big(C^t \,\big|\hat{s}^{t-1}, s_i^t\big)\big(\sigma^*\big) \geq U_{i,t}\big(C^t \,\big|\hat{s}^{t-1}, s_i^t\big)\big(\sigma_{-i}^*, \sigma_i\big),\\
&U_{i,r}\big(\big(C_\tau^r\big)_{\tau=r}^{t-1}, \big(C_\tau^t\big)_{\tau=t}^{T} \,\big|\tilde{s}^{r-1}, s_i^r\big)\big(\sigma^*\big) \geq U_{i,r}\big(C^r \,\big|\tilde{s}^{r-1}, s_i^r\big)\big(\sigma^*\big)\\
&U_{i,r}\big(\big(C_\tau^r\big)_{\tau=r}^{t-1}, \big(C_\tau^t\big)_{\tau=t}^{T} \,\big|\hat{s}^{r-1}, s_i^r\big)\big(\big(\sigma_{-i,\tau}^*, \sigma_{i,\tau}\big)_{\tau=r}^{t-1}, \big(\sigma_\tau^*\big)_{\tau=r}^{T}\big)\\
&\qquad \leq U_{i,r}\big(C^r \,\big|\hat{s}^{r-1}, s_i^r\big)\big(\big(\sigma_{-i,\tau}^*, \sigma_{i,\tau}\big)_{\tau=r}^{t-1}, \big(\sigma_\tau^*\big)_{\tau=r}^{T}\big).
\end{aligned}$$

The feasibility constraint holds for $\tau \geq t$ and for all $\hat{s}^\tau$, and the t incentive-compatibility constraint holds for all i, all $\sigma_i \in \Sigma$, all $\hat{s}^{t-1}$, and all s^t. The promise-keeping constraint holds for all i and all s^r, and the threat-keeping constraint holds for all i, all $\hat{s}^{r-1}$, all s^r, and all $\big(\sigma_{i,\tau}\big)_{\tau=r}^{t-1}$. Note that the threat-keeping constraint in this problem involves punishments for multi-period deviations from truth-telling because the allocation may not have been reformed in the previous period. If $t = 0$, the promise-keeping and threat-keeping constraints do not appear in the problem. When the constraint set is non-empty, we assume that the maximizing allocation C^{*t} is weakly monotonic at t. The same arguments as in the $t = 1$ case then imply that the government can solve problem (33) with a simplified allocation C^t.

If the constraint set in the reform problem (33) is empty for some $t > 0$, then the government sets $C^t \equiv C^{t-1}$ and recommends that the agents play the equilibrium σ^e of C^t

that maximizes period-t social welfare

$$\inf_{\Pi_{g,t+1}} \mathbb{E}_{\pi_{g,t+1}} \left[\sum_i \eta_i U_{i,t} \big(C^t \,\big|\tilde{s}^{t-1}, s_i^t \big) \big(\sigma^e \big) \middle| \tilde{s}^{t-1}, s_g^t \right].$$

By solving the period-t reform problem (33) in each period, retaining the previous allocation when the constraint set of (33) is empty, the government can find a sequence of simplified allocations $\{C^t\}_{t=0}^T$ that implements C^*.

We remark that under Assumption 2, if the government designs a simplified allocation C^t at t, then the government will be able to reform the allocation at either $t+1$ or $t+2$. In particular, if the government fails to reform the allocation at $t+1$, then the constraint set at $t+2$ is non-empty. To see this, note that at $t+2$, the allocation C^t trivially satisfies the non-negativity, feasibility, promise-keeping, and threat-keeping constraints. In addition, the construction described in the proof of Proposition 3 implies that the C^t policy functions do not depend on agents' reports at $\tau \geq t+2$. Thus C^t also satisfies the $t+2$ incentive-compatibility constraint, so the constraint set in the reform problem is non-empty.

The incentive-compatibility of C^t at $t+2$ additionally makes use of how we characterize reporting. In particular, each agent reports a shock $\hat{s}_{i,t}$ in each period, rather than an entire history of shocks $\hat{s}_i^t$. This is not a meaningful distinction when agents truthfully report in every period, but during the course of constructing reform allocations, the government may occasionally recommend that the agents play an untruthful equilibrium σ^e. Since agents only report shocks in each period, the incentive-compatibility constraint in the reform problem (33) does not require that agents correct their untruthful reports from previous periods. If agents were instead required to report histories of shocks $\hat{s}_i^t$ in each period, the incentive-compatibility constraint in (33) would be strictly stronger, and C^t may fail to be incentive-compatible at $t+2$.

A.3 Proofs for Section 5

Proof. **(Proposition 6)** Given allocation $\tilde{C}_i^0$, agent i's $t = 0$ continuation utility is given by

$$U_{i,0}\left(\tilde{C}_i^0 \left| s_0\right.\right) = u\left(\tilde{c}_{i,0}\left(s_0\right), \frac{\tilde{z}_{i,0}\left(s_0\right)}{\theta_{i,0}}\right) + \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\beta u\left(\tilde{c}_{i,1}\left(s_0,\theta_1\right), \frac{\tilde{z}_{i,1}\left(s_0,\theta_1\right)}{\theta_{i,1}}\right) + \sum_{t=2}^{T} \beta^t u\left(\tilde{c}_{i,t}\left(s_0,\theta_1\right), \frac{\tilde{z}_{i,t}\left(s_0,\theta_1\right)}{\underline{\theta}}\right)\right| s_0\right]$$

For notational simplicity, define

$$\tilde{v}\left(s_0,\theta_1\right) \equiv \beta u\left(\tilde{c}_{i,1}\left(s_0,\theta_1\right), \frac{\tilde{z}_{i,1}\left(s_0,\theta_1\right)}{\theta_{i,1}}\right) + \sum_{t=2}^{T} \beta^t u\left(\tilde{c}_{i,t}\left(s_0,\theta_1\right), \frac{\tilde{z}_{i,t}\left(s_0,\theta_1\right)}{\underline{\theta}}\right).$$

Define $\hat{v}\left(s_0,\theta_1\right)$ similarly. By Assumption 3, we know that for each $\pi_{i,1} \in \Pi_{i,1}$, there exists $\pi'_{i,1} \in \Pi_{i,1}$ such that $\pi_{i,1}$ and $\pi'_{i,1}$ have the same marginal distribution over $\theta_{-j,1}$ conditional on s_0, but $\pi'_{i,1}\left(\cdot\right| s_0)\big|_{\theta_i}$ places weight only on $\underline{\theta}$ and $\bar{\theta}$ so as to satisfy

$$\mathbb{E}_{\pi_{i,1}}\left[\theta_{i,1}\right| s_0] = \bar{\theta}_{i|i,1}\left(s_0\right).$$

Let $\Pi'_{i,1} \subset \Pi_{i,1}$ denote the subset of all distributions of the form $\pi'_{i,1}$. By set inclusion, the following inequalities are immediate:

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\tilde{v}\left(s_0,\theta_1\right)\right| s_0] \leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}\left(s_0,\theta_1\right)\right| s_0], \tag{34}$$

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\hat{v}\left(s_0,\theta_1\right)\right| s_0] \leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}\left(s_0,\theta_1\right)\right| s_0]. \tag{35}$$

To see that the opposite inequalities also hold, let $p\left(s_0\right) \in [0,1]$ be such that

$$p\left(s_0\right)\underline{\theta} + \left(1 - p\left(s_0\right)\right)\bar{\theta} = \bar{\theta}_{i|i,1}\left(s_0\right).$$

By the definition of $\tilde{C}_i^0$, $\tilde{v}(s_0, \theta_{-i,1}, \cdot)$ lies weakly above its secant line from $\underline{\theta}$ to $\bar{\theta}$, so by the independence of $\theta_{-i,1}$ and $\theta_{i,1}$ under $\pi_{i,1}$ and $\pi'_{i,1}$, we have

$$\begin{aligned}
\mathbb{E}_{\pi_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0\right] &= \mathbb{E}_{\pi_{i,1}}\left[\mathbb{E}_{\pi_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0, \theta_{-j,1}\right]\middle| s_0\right] \\
&\geq \mathbb{E}_{\pi_{i,1}}\left[p(s_0)\,\tilde{v}(s_0, \theta_{-i,1}, \underline{\theta})\right. \\
&\qquad\qquad \left.+ (1 - p(s_0))\,\tilde{v}\left(s_0, \theta_{-i,1}, \bar{\theta}\right)\middle| s_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[p(s_0)\,\tilde{v}(s_0, \theta_{-i,1}, \underline{\theta})\right. \\
&\qquad\qquad \left.+ (1 - p(s_0))\,\tilde{v}\left(s_0, \theta_{-i,1}, \bar{\theta}\right)\middle| s_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0\right].
\end{aligned}$$

A similar set of calculations applies to $\hat{v}(s_0, \theta_1)$. By taking infima on both sides, we find that equality must hold in (34) and (35).

Now since $\tilde{v}(s_0, \theta_{-i,1}, \cdot)$ and $\hat{v}(s_0, \theta_{-i,1}, \cdot)$ coincide on the endpoints of Θ, the independence of $\theta_{i,1}$ and $\theta_{-i,1}$ imply

$$\begin{aligned}
\mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0\right] &= \mathbb{E}_{\pi'_{i,1}}\left[p(s_0)\,\tilde{v}(s_0, \theta_{-i,1}, \underline{\theta})\right. \\
&\qquad\qquad \left.+ (1 - p(s_0))\,\tilde{v}\left(s_0, \theta_{-i,1}, \bar{\theta}\right)\middle| s_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[p(s_0)\,\hat{v}(s_0, \theta_{-i,1}, \underline{\theta})\right. \\
&\qquad\qquad \left.+ (1 - p(s_0))\,\hat{v}\left(s_0, \theta_{-i,1}, \bar{\theta}\right)\middle| s_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}(s_0, \theta_1)\middle| s_0\right].
\end{aligned}$$

Taking infima over $\Pi'_{i,1}$ on both sides, we find

$$\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0\right] = \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}(s_0, \theta_1)\middle| s_0\right]. \tag{36}$$

Since equality holds in (34) and (35), (36) then implies

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\tilde{v}(s_0, \theta_1)\middle| s_0\right] = \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\hat{v}(s_0, \theta_1)\middle| s_0\right].$$

Thus agent i is indifferent between $\tilde{C}_i^0$ and $\hat{C}_i^0$. □

B Infinite Time Horizon, Continuum of Agents

This section describes an extension of the model in Section 2 to the case in which the time horizon is infinite ($T = \infty$) and there is a continuum of agents. To avoid difficulties in using backward induction to define certain belief sets, the model formulation will be sequential.

Identify each agent with a real number in the interval $I_1 = [0,1]$, and equip I_1 with the Borel sigma algebra $\mathcal{B}$ and Lebesgue measure λ. Let $\Theta \subset \mathbb{R}_+ - \{0\}$ be the set of skill shocks that are possible in each period, which we will suppose is compact and connected for simplicity. As before, let $\underline{\theta}$ and $\bar{\theta}$ denote the minimum and maximum elements of Θ, respectively. Give $\Theta^{\mathbb{N}_0}$ the sigma algebra generated by finite rectangles, i.e., sets of the form

$$(a_1, b_1] \times (a_2, b_2] \times \ldots \times (a_n, b_n] \times \Theta \times \Theta \times \ldots$$

with $\underline{\theta} \leq a_k \leq b_k \leq \bar{\theta}$, $k = 1, \ldots, n$, and $n \in \mathbb{N}$.

Let $(\Omega, \mathcal{F}, \mu)$ be a measure space, where $\omega \in \Omega$ denotes a possible state of the economy and μ is the distribution of states. Give $I_1 \times \Omega$ the product sigma algebra $\mathcal{B} \otimes \mathcal{F}$ with the product measure $\varphi \equiv \lambda \otimes \mu$. Let $\theta^\infty : I_1 \times \Omega \to \Theta^{\mathbb{N}_0}$ be $\mathcal{B} \otimes \mathcal{F}$-measurable. For any $i \in I_1$, $\theta^\infty(i, \omega)$ denotes agent i's infinite sequence of skill shocks if the state is ω. For notational simplicity, we will write $\theta_i^\infty \equiv \theta^\infty(i, \cdot)$. Define θ^t as the obvious projection of θ^∞ to Θ^{t+1}, and let $\mathcal{F}_t \equiv \sigma(\theta^t) \subset \mathcal{B} \otimes \mathcal{F}$ be the sigma algebra generated by θ^t. Then $(\mathcal{F}_t)_{t=0}^\infty$ is a filtration on $\mathcal{B} \otimes \mathcal{F}$ to which the shock history process $(\theta^t)_{t=0}^\infty$ is adapted.

Let $C \equiv \{c_{i,t}(\theta^t), z_{i,t}(\theta^t), k_{i,t+1}(\theta^t)\}_{t=0, i \in I_1}^\infty$ denote an allocation, and let $f : \mathbb{R}_+^2 \to \mathbb{R}_+$ be a constant returns-to-scale production function that is increasing in capital and effective labor. We say that C is *feasible* if for all t and all θ^t,

$$\int_{I_1} c_{i,t}(\theta^t) + k_{i,t+1}(\theta^t)\, d\lambda \leq \int_{I_1} f\left(k_{i,t}(\theta^{t-1}), z_{i,t}(\theta^t)\right) d\lambda.$$

For each agent $i \in I_1$, Let $P_i \subset \Delta(I_1 \times \Omega, \mathcal{B} \otimes \mathcal{F})$ be a non-empty set of prior distributions on $(I_1 \times \Omega, \mathcal{B} \otimes \mathcal{F})$ that represent agent i's beliefs. We assume that each distribution $p_i \in P_i$ is a product measure of the form $\lambda \otimes \mu_i$ for some $\mu_i \in \Delta(\Omega, \mathcal{F})$. Unlike in the main text, for simplicity we will assume here that that agents update their beliefs using Bayes's Theorem prior-by-prior. Given an allocation C, and a state θ^t, agent i's t continuation utility is given by

$$U_{i,t}\left(C \,\middle|\, \theta^t\right) := \inf_{p_i \in P_i} \mathbb{E}_{p_i}\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u\left(c_{i,\tau}(\theta^\tau), \frac{z_{i,\tau}(\theta^\tau)}{\theta_{i,\tau}}\right) \middle|\, \theta^t\right],$$

where $\beta \in (0,1)$. With social welfare weights $\eta \in \Delta(I_1, \mathcal{B})$ and an initial state θ_0, an *efficient* allocation $C^*(\theta_0)$ is given by

$$C^*(\theta_0) \in \arg\max_C \int_{I_1} U_{i,0}(C\,|\theta_0)\,d\eta,$$

subject to feasibility and non-negativity.

The analog of Assumption 1 in this setup is

Assumption 4. *For all t, all i, all θ^t, and all $p_i \in P_i$, there exists $p_i' \in P_i$ such that*

$$p_i'\left(\cdot|\,\theta^t\right)\big|_{\mathcal{F}_{t+1}} = p_i\left(\cdot|\,\theta^t\right)\big|_{\mathcal{F}_{t+1}},$$

but

$$p_i'\left((\theta^\infty)^{-1}\left(\left\{\ell \in \Theta^{\mathbb{N}_0} : \ell_\tau = \underline{\theta},\ \tau \geq t+2\right\}\right)\right) = 1.$$

This assumption implies that in any period t and for any belief $p_i \in P_i$, there exists another belief $p_i' \in P_i$ such that p_i and p_i' imply the same distribution of the $t+1$ state θ^{+1} conditional on θ^t. However, under p_i', almost surely almost all agents will realize the shock $\underline{\theta}$ at $\tau \geq t+2$.

To prove an analogue of Proposition 1, we must make an additional assumption. Since the policy functions in an allocation depend fully on the state θ^t, it is possible that changing the shock histories of agents in a set of λ-measure zero could alter an agent i's allocation. This property is inconsistent with Assumption 4, which ensures only that a full measure of agents realize the shock $\underline{\theta}$ at $\tau \geq t+2$ under the distribution p_i'. We will thus assume that the policy functions do not distinguish between any two t states θ^t and $\tilde{\theta}^t$ such that $\theta_i^t \neq \tilde{\theta}_i^t$ for i in a set of λ-measure zero.[24]

Given the assumptions above, periodically-reformed policies are optimal:

Proposition 7. *The efficient allocation C^* can be implemented by a sequence of simplified allocations $\{C^t\}_{t=0}^\infty$.*

Proof. Define C^0 to coincide with C^* at $t=0$ and at $t=1$. Then let

$$c_t^0\left(\theta^t\right) \equiv c_t^*\left(\theta^1, (\underline{\theta}_\tau)_{\tau=2}^{t-2}\right).$$

[24]This assumption is without loss of generality when the social welfare measure η is absolutely continuous with respect to Lebesgue measure. Alternatively, a natural way to enforce this assumption is to constrain policy functions to depend only on an agent's shock history θ_i^t as well as the distribution of shock histories in the economy. To define the information structure appropriately, random measures should be used to model the agent's beliefs about the distribution of shock histories that will be observed in subsequent periods.

Here $\underline{\theta}_\tau : I_1 \times \Omega \to \Theta$ satisfies $\underline{\theta}_\tau (i, \omega) = \underline{\theta}$ for all $(i, \omega) \in I_1 \times \Omega$. Define z_t and k_{t+1} similarly. Let $P_i' \subseteq P_i$ denote all distributions of the form $p_i \in P_i$. Then for all i,

$$\begin{aligned}
U_{i,0}\left(C^* \left|\theta^0\right.\right) &= \inf_{P_i} \mathbb{E}_{p_i}\left[\sum_{t=0}^{\infty} \beta^t u\left(c_{i,t}^*\left(\theta^t\right), \frac{z_{i,t}^*\left(\theta^t\right)}{\theta_{i,t}}\right)\middle|\theta^0\right] \\
&\leq \inf_{P_i'} \mathbb{E}_{p_i'}\left[\sum_{t=0}^{\infty} \beta^t u\left(c_{i,t}^*\left(\theta^t\right), \frac{z_{i,t}^*\left(\theta^t\right)}{\theta_{i,t}}\right)\middle|\theta^0\right] \\
&= \inf_{P_i'} \mathbb{E}_{p_i'}\left[\sum_{t=0}^{\infty} \beta^t u\left(c_{i,t}^0\left(\theta^t\right), \frac{z_{i,t}^0\left(\theta^t\right)}{\theta_{i,t}}\right)\middle|\theta^0\right] \\
&\leq \inf_{P_i} \mathbb{E}_{p_i}\left[\sum_{t=0}^{\infty} \beta^t u\left(c_{i,t}^0\left(\theta^t\right), \frac{z_{i,t}^0\left(\theta^t\right)}{\theta_{i,t}}\right)\middle|\theta^0\right] \\
&= U_{i,0}\left(C^0 \left|\theta^0\right.\right).
\end{aligned}$$

The third line holds because C^0 and C^* coincide when almost every agent realizes the shock $\underline{\theta}$ at $t \geq 2$. The last line holds because c_t^0 and z_t^0 only depend on θ^1 for $t \geq 2$, and agent i can potentially realize a shock $\theta_{i,t} \neq \underline{\theta}$ at some $t \geq 2$ under a distribution $p_i \in P_i - P_i'$. Hence all agents weakly prefer C^0 to C^*, and the feasibility of C^0 follows from that of C^*. By iterating this process in each period, we have the result. □

With Proposition 7, we find that even when the time horizon is infinite and there is a continuum of measureless agents, the efficient allocation can be implemented by a sequence of simplified allocations $\{C^t\}_{t=0}^{\infty}$ that are reformed after each period. As in Section 2.2, each allocation C^t displays limited state dependence at $\tau \geq t+2$, and an argument analogous to that in Section 2.3 implies that they are history independent whenever the government's promise-keeping constraints are slack.