Implications of Uncertainty for Optimal Policy*

Maxim Troshkin†

July 2, 2018

Abstract

We study the implications of a broader view of uncertainty for social insurance and redistribution in conventional macro public finance models - with asymmetric information, with exogenous lack of commitment, and with a combination of the two. We show that uncertainty manifests as endogenous lack of commitment, implying the optimality of periodically reformed policies. Optimal periodic reforms imply simpler policies, often leading to loss of full history dependence. Optimal policies also can be characterized without full backward induction to compute promised utilities when time horizon is finite. However, linear policies can be far from optimal. We show that decentralized economies with uncertainty may not be efficient unless markets are complete, implying a meaningful role for government provision of insurance.

Keywords: uncertainty, robustness, social insurance, redistribution, optimal policy, efficiency, risk sharing, private insurance, crowding out

1 Introduction

A sizable and growing literature shares the following approach to social insurance, redistribution, and classical normative questions more generally: Start with a friction, commonly asymmetric information about idiosyncratic shocks, or an exogenous inability to commit on the part of the policy designer or the agents, and characterize friction-constrained allocations that maximize a social objective, usually social welfare. The optimal policies are 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

*Early versions were circulated as "On the Optimality of Periodically Reformed, Simple Fiscal Policies", which benefited from comments by Marco Battaglini, Steve Coate, Lars Hansen, and Narayana Kocherlakota. I thank John Zhu for early discussions that led to a more productive approach and key results. Todd Lensman provided absolutely indispensable research assistance.

†Cornell University, troshkin@cornell.edu.

dependence in dynamic contexts. In addition, because the policy designer and agents are typically assumed to have perfect knowledge of the data-generating process, the optimal policies are designed once and maintained forever.[1]

At the same time, a growing body of evidence that has benefited more recently from administrative data suggests that idiosyncratic shock distributions change significantly and quite often.[2] One implication is that certainty about the data-generating processes in the economy is a strong assumption when designing policies. This can be further interpreted as suggesting that social contracts that are potentially incomplete and periodically renegotiated could be better suited for understanding optimal policies. Casual empiricism also points to real world policies that are at least somewhat incomplete, history independent, and periodically reformed.

This paper moves away from the assumption of certainty about the future distributions of idiosyncratic shocks, and characterizes optimal social insurance and redistribution that are robust with respect to incomplete or inaccurate knowledge of stochastic properties of the economy. In order 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 in various literatures as models, priors, or beliefs). The approach we take to modeling risk and uncertainty, with aversion to both, follows the approaches in the macroeconomic and finance literatures (e.g., Hansen and Sargent (2001), Epstein and Schneider (2003)).[3] We therefore study mutual implications between dynamic optimal policy and macroeconomic and finance approaches to uncertainty.

The economy we consider has a unit measure of agents who experience idiosyncratic skill shocks over time. The data-generating process for skills is arbitrary and unknown to anyone in the economy. Each agent has a set of distributions that he believes may represent the data-generating process for skills in the next period. We consider both the baseline case

[1]For reviews of applications to dynamic optimal fiscal policies see, e.g., Golosov, Tsyvinski, and Werning (2006), Kocherlakota (2010), Golosov and Tsyvinski (2015). For a textbook treatment of theses approaches to social insurance more generally see, e.g., Ljungqvist and Sargent (2012, ch.20, 21).

[2]In particular, recent administrative data evidence suggests that the distribution of pre-tax earnings has been going through dramatic changes in many developed countries. These significant changes appear quite often and at irregular intervals, and they are apparently unanticipated by the governments (see, e.g., Piketty, Saez, Zucman (2017)). A related literature also suggests that people hold inconsistent and at times distribution-incompatible beliefs about the distributions of their future productivities, their future tax liabilities, etc. (for a review see, e.g., Rees-Jones and Taubinsky (2017); for a recent example see, e.g., Aghion et al. (2017)).

[3]See also Hansen, Sargent, Turmuhambetova, and Williams (2006). A number of recent studies in those literatures have focused on showing that such broader notion of uncertainty can help explain the behavior of economic aggregates to a surprising degree, e.g., Bianchi, Ilut, and Schneider (2017), Hansen and Sargent (2017), Ilut, Kehrig, and Schneider (2017).

of public information and the cases of private beliefs, private skills, and exogenous outside options. The government seeks to provide social insurance and a degree of redistribution, and it is constrained by the same lack of certainty about the distribution of future skills. In other words, the government is not an abstract entity with perfect knowledge of the data-generating processes infinitely far into the future. 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 individuals' beliefs and keep the environment virtually agnostic about learning that may map histories of individual actions 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 requirements that their current beliefs are consistent with what has been observed in the past and that they are sufficiently uncertain about future beliefs. To aid technical intuition, we demonstrate arguments using a common parametric formulation of repeated maxmin expected utility with arbitrary belief updating rules.[4]

We derive three main results. First, we show that with uncertainty as a friction, 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 reform the allocations as needed after the economy realizes a new distribution of shocks. This follows from the observation that uncertainty 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, subsequently reforming them as necessary. We show that this is the case even when information about realized skills and about beliefs is public and even when the government and the agents have the ability to fully commit to allocations designed ex ante. This optimality of periodic reforms extends to private beliefs, private skills, and exogenous lack of commitment. When the time horizon is finite, this implies that efficient allocations can be characterized without the full backward induction ordinarily required to compute promised utilities.

Second, we show that optimal policies themselves are also simplified: Optimal allocations are generally independent of full history of shocks because they can be periodically reformed

[4] For dynamically consistent representation of repeated maxmin expected utility formulation see Epstein and Schneider (2003). Our results are straightforwardly generalizable to, e.g., dynamic variational preferences (Maccheroni, Marinacci and Rustichini (2006b)), and more generally would extend to a dynamic version of uncertainty-averse preferences (Cerreia-Vioglio, Maccheroni, Marinacci, and Montrucchio (2011)) that unify many other specifications, including in particular multiplier preferences as in, e.g., Hansen and Sargent (2001).

at any date. However, linear and affine policies can be far from optimal even when agents are uncertain about the future distribution of idiosyncratic shocks. In particular, we argue that quite 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 economies with uncertainty, where agents contract with firms to provide labor and capital in exchange for consumption, competitive equilibria may not be efficient unless markets are complete. This is because firms that are uncertain about the next period's distribution of shocks would need to purchase securities that insure against this uncertainty in order to provide as much insurance as the government. As a result, markets that do not permit trade in securities for each contingency - including contingencies that describe distributions of beliefs - lead to less than efficient provision of insurance. This implies that a broader view of uncertainty creates a potentially meaningful role for the government provision of insurance. The decentralized economies nevertheless maintain the features of simplified, period-by-period insurance that is periodically reformed, not unlike what is commonly observed in reality.

In the next section, we consider a baseline environment, in which we introduce uncertainty about future shocks as the only friction and show that periodic reforms and history independence are optimal. In Section 3, we extend the characterization of constrained efficient policies to private beliefs. We prove that a version of the Revelation Principle holds for private belief revelation, show that incentive constraints are often non-binding, and suggest simple ways of implementing constrained efficient allocations in information-constrained economies. We then derive extensions to asymmetric information about skills and exogenous lack of commitment. Section 4 analyzes the decentralized versions of these economies. In Section 5, to further examine the simplicity of optimal policies implied by the periodic reforms, we characterize conditions under which affine and linear policies can be optimal.

2 Uncertainty as the Friction

To focus on uncertainty, we first consider a baseline environment that is essentially a conventional dynamic heterogeneous-agents economy, with public information and full commitment. The only unconventional element - and the only friction - is uncertainty about 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 prove that the government can achieve efficiency with policies that are periodically reformed and do not need to depend 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 unit measure of agents.[5] At the beginning of $t = 0$, nature draws a sequence of idiosyncratic shocks $s^T \equiv (s_0, .., s_T)$ for each agent, and reveals each s_t to the agent at the beginning of t. 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 once a shock is revealed to an agent it becomes public information. Consequently, the distribution of shock histories up to period t is public information at the beginning of t. Let $\varphi_t \in \Delta(s^t)$ denote any possible distribution of histories s^t in period t.

An idiosyncratic shock has two components: a skill and a set of subjective beliefs, $s_t \equiv (\theta_t, \Pi_{t+1})$.[6] First, θ_t denotes an agent's idiosyncratic skill (productivity), so that if the amount of labor the agent exerts is $l_t \in [0, L]$, $L < \infty$, then the effective labor supplied is $z_t \equiv \theta_t l_t$. Assume that θ_t has compact 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 $\theta^t \equiv (\theta_0, \ldots, \theta_t)$ denote a period t skill history.

Second, each shock contains a set of subjective beliefs, i.e., the next period distributions of idiosyncratic histories that the agent believes to be possible. An agent's beliefs at t about $t + 1$ form a non-empty set $\Pi_{t+1} \subset \Delta(s^{t+1})$ of probability measures π_{t+1} over history types s^{t+1}. The beliefs are consistent with what has been observed so far: If φ_t is the distribution realized at t, then it is equal to the marginal distribution of s^t under belief π_{t+1}, for any $\pi_{t+1} \in \Pi_{t+1}$. Let $\Pi^{t+1} \equiv (\Pi_0, \ldots, \Pi_{t+1})$ denote a $t + 1$ history of beliefs.[7]

To analyze properties common to social insurance and redistribution in general, without restricting attention to specific sets of policy tools, we organize our analysis around

[5] For simplicity, in the main text we maintain finite countability of time periods and discuss the technical details of extending to the infinite horizon case in Appendix.

[6] Alternatively, we could specify idiosyncratic skill shocks θ_t and an updating (and potentially learning) process by which beliefs Π_{t+1} are formed. Rather than introducing additional notation to describe arbitrary updating, the exposition is significantly simplified by including beliefs in the shock and imposing assumptions on the beliefs directly.

[7] One example of this broader view of uncertainty that is commonly used in macroeconomics (e.g., Hansen and Sargent (2001)) supposes that in any period, an agent has a particular statistical model in mind, i.e., a distribution $\pi^*_{t+1} \in \Delta(s^{t+1})$ that he thinks may be the true distribution of $t + 1$ types. This distribution is often described as the agent's "benchmark" or "approximating" model. However, the agent distrusts this model and considers other models π_{t+1} that are "close to" π^*_{t+1} in the sense of distance d on $\Delta(s^{t+1})$, 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 $\Pi_{t+1} \equiv \{\pi_{t+1} | \, d(\pi_{t+1}, \pi^*_{t+1}) \leq \epsilon\}$.

allocations that a policy would deliver to agents. An *allocation*

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

is a sequence of consumption, effective labor, and capital functions that depend on the idiosyncratic type s^t. Since all agents are uncertain about the data-generating process, each of the variables in an allocation also depends on the type distribution φ_t in a potentially nontrivial way.

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 k_t and effective labor z_t. Given an initial distribution φ_0, an allocation C is *feasible* if for every t and every $\varphi_t \geq \varphi_0$,[8]

$$\int c_t\left(s^t,\varphi_t\right) d\varphi_t\left(s^t\right) + K_{t+1}\left(\varphi_t\right) \leq f\left(K_t\left(\varphi_{t-1}\right), Z_t\left(\varphi_t\right)\right),$$

where $K_t\left(\varphi_{t-1}\right) \equiv \int k_t\left(s^{t-1},\varphi_{t-1}\right) d\varphi_{t-1}\left(s^{t-1}\right)$ and $Z_t\left(\varphi_t\right) \equiv \int z_t\left(s^t,\varphi_t\right) d\varphi_t\left(s^t\right)$ are the aggregate capital and effective labor functions respectively. This aggregate ex post feasibility constraint must hold for any distribution φ_t that could follow φ_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

$$U_t\left(C,s^t,\varphi_t\right) \equiv u\left(c_t\left(s^t,\varphi_t\right), \frac{z_t\left(s^t,\varphi_t\right)}{\theta_t}\right) + \beta \inf_{\pi_{t+1}\in\Pi_{t+1}} \mathbb{E}_{\pi_{t+1}}\left[U_{t+1}\left(C,s^{t+1},\pi_{t+1}\right)\middle| s^t\right],$$

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 $\mathbb{E}_{\pi_{t+1}}$ denotes an expectation with respect to the belief $\pi_{t+1} \in \Pi_{t+1}$. When $t = T$, a U_{t+1} term does not appear on the right hand side. Agents are thus averse to both risk originating from stochastic skills and uncertainty captured by multiple beliefs. The latter can be 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 seeking a degree of redistribution while providing social insurance. The government's problem is to maximize a weighted average of the agents'

[8]For any two distributions φ_t and φ_τ with $t \geq \tau$, we say that φ_t *follows* φ_τ, written $\varphi_t \geq \varphi_\tau$, if $\varphi_\tau\left(s^\tau\right) = \int_{\{s^t \geq s^\tau\}} d\varphi_t\left(s^t\right)$.

[9]For an axiomatic treatment and a characterization of the existence of a recursive representation of this form see Epstein and Schneider (2003). We further discus dynamic consistency below in Section 2.4.

utilities subject to feasibility and non-negativity of allocations, where the weighting captures the redistribution motive and is given by a non-negative measure η on $\Theta \times 2^{\Delta(s^1)}$.[10] Given social welfare weights η and an initial type distribution φ_0, an *efficient* allocation $C^*(\varphi_0)$ is given by

$$C^*(\varphi_0) \in \arg\max_C \int U_0(C, s_0, \varphi_0)\, d\eta(s_0) \tag{1}$$

subject to feasibility and non-negativity. Note that instead of an abstract entity with perfect knowledge of the data-generating process, the government quite concretely possesses the same information about realized types and the future distributions as do all of the agents combined. Note also that to focus 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.

2.2 Optimality of periodic reforms

We first show that under a mild condition on agents' beliefs, the government in this economy can achieve efficiency with simplified, short-term policies that 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 short term as well as on an endogenous outside option that remains feasible even when the actual path of the economy falls outside of the belief overlap. One obvious example of such overlap is when all agents agree on the worst path for the economy. We consider this case to illustrate the result and then discuss its generality.

To construct the "worst" beliefs that an agent could have at any date, let $\underline{\Pi}_T = \{\underline{\pi}_T\}$ denote the unique belief set with $\underline{\pi}_T$ placing unit weight on everyone having skills $\underline{\theta}$ 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 unique belief set with $\underline{\pi}_{t+1}$ placing unit weight on $t+1$ shocks of the form $s_{t+1} = \left(\underline{\theta}, \underline{\Pi}_{t+2}\right)$ 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 pessimistic beliefs will persist, so he is certain that everyone will get $\underline{\theta}$ at $t+2, \ldots, T$. In other words, any agent who realizes beliefs $\underline{\Pi}_{t+1}$ is maximally pessimistic about the future path of the economy.

Given any belief $\pi_{t+1} \in \Pi_{t+1}$, it will be convenient to let $\pi_{t+1}(\cdot | s^t)|_\theta$ denote the marginal distribution of θ_{t+1}, conditional on a type s^t. Similarly, let $\pi_{t+1}(\cdot | s^t)|_\Pi$ denote the conditional marginal distribution of Π_{t+2}. The following is then one simple way to ensure that

[10] For instance, a utilitarian motive implies $\eta = \varphi_0$.

all agents agree on the worst path for the economy:

Assumption 1. *For any* $t = 0, \dots, T-1$, *any* Π_{t+1} *realized by an agent, and any belief* $\pi_{t+1} \in \Pi_{t+1}$, *there exists another belief* $\pi^-_{t+1} \in \Pi_{t+1}$ *such that*

$$\pi_{t+1}\left(\cdot|\, s^t\right)\Big|_\theta = \pi^-_{t+1}\left(\cdot|\, s^t\right)\Big|_\theta \quad \forall s^t$$

but $\pi^-_{t+1}\left(\cdot|\, s^t\right)\big|_\Pi$ *places unit weight on* $\underline{\Pi}_{t+2}$.

In other words, regardless of what an agent believes at any t about future skills, he believes it possible that everyone at $t+1$ could be maximally pessimistic with 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 $s_\tau = \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ in $\tau = t+2, \dots, T$.[11]

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' heterogeneous beliefs about the next-period skill distribution. Besides the basic consistency with what has been observed so far, this assumption is the only condition on agents' beliefs and belief updating. The following establishes that under these conditions, the efficient allocations can be simplified and periodically reformed:

Proposition 1. *The efficient allocation* C^* *can be implemented by a sequence of two-period allocations* $\{C^t\}_{t=0}^T$, *where* $C^t = \left\{c^t_\tau, z^t_\tau, k^t_{\tau+1}\right\}_{\tau=t}^{t+1}$.

The implementation here is in the sense that every agent's $t = 0$ utility under the two-period allocation C^0 is equal to that under the complete efficient allocation C^*:

$$U_0\left(C^0, s_0, \varphi_0\right) = U_0\left(C^*, s_0, \varphi_0\right) \quad \forall s_0.$$

In essence, the result comes from agents being sufficiently uncertain about future beliefs and from a sufficient overlap in the sets of distributions entertained by the agents. While the technical details of the proof are provided in Appendix, the construction of the two-period allocations is intuitive. At $t = 0$, each agent believes it is possible to realize a $t = 1$ distribution wherein everyone has the worst belief $\underline{\Pi}_2$ about the future path of the economy, i.e., each agent would be certain that everyone will realize the worst skill $\underline{\theta}$ at all $t \geq 2$. The

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

two-period 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 distributions are realized at $t = 1$. Start with $t = 0$ 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 $t = 1$, construct 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_1^*, z_1^*, and k_2^* prescribe in the case when every agent has beliefs $\underline{\Pi}_2$. In other words, C^0 takes for granted that agents will realize the worst beliefs $\underline{\Pi}_2$. It is not a completely state-contingent allocation, but instead it depends only on the $t = 0$ type $s_0 = (\theta_0, \Pi_1)$, the $t = 1$ skill θ_1, and the marginal distribution of $((\theta_0, \Pi_1), \theta_1)$ in the economy.[12]

To see that at $t = 0$ all agents are indifferent between C^* and C^0, let $\Pi_1^- \subset \Pi_1$ denote the subset of beliefs π_1^- that place unit weight on the worst next-period belief set $\underline{\Pi}_2$. Since agents are averse to uncertainty, the lowest expected utility over Π_1^- places a weak upper bound on the lowest expected utility over Π_1:

$$\inf_{\Pi_1} \mathbb{E}_{\pi_1} \left[U_1 \left(C^*, s_0, \pi_1\right) \middle| s_0 \right] \leq \inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-} \left[U_1 \left(C^*, s_0, \pi_1^-\right) \middle| s_0 \right].$$

By construction, C^0 and C^* coincide when the distribution realized at $t = 1$ is actually of the same form as the belief π_1^-, so we have

$$\inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-} \left[U_1 \left(C^*, s_0, \pi_1^-\right) \middle| s_0 \right] = \inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-} \left[U_1 \left(C^0, s_0, \pi_1^-\right) \middle| s_0 \right].$$

Assumption 1 guarantees that Π_1^- is not only non-empty, but for any belief $\pi_1 \in \Pi_1$ there is another belief $\pi_1^- \in \Pi_1$ with the same marginal distribution over $((\theta_0, \Pi_1), \theta_1)$. An agent not only considers distributions that place unit weight on $\underline{\Pi}_2$, but he thinks it is possible for everyone to realize the belief set $\underline{\Pi}_2$ 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 $t = 1$,

$$\inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-} \left[U_1 \left(C^0, s_0, \pi_1^-\right) \middle| s_0 \right] = \inf_{\Pi_1} \mathbb{E}_{\pi_1} \left[U_1 \left(C^0, s_0, \pi_1\right) \middle| s_0 \right].$$

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

[12] We abuse notation here since as we explain below C^0 effectively specifies an *endogenous outside option* 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 $t \geq 2$.

indifferent between C^* and the simpler, not fully state-contingent, two-period 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 Pareto optimal anymore. In that case, at $t = 1$ the government can reform C^0 to a new allocation C^1. In general, given a two-period allocation C^{t-1}, the government can reform the allocation to the new optimal two-period allocation C^{*t} given by

$$C^{*t}\left(\varphi_t, C^{t-1}\right) \in \arg\max_{C^t} \int U_t\left(C^t, s^t, \varphi_t\right) d\eta\left(s_0\right) \tag{2}$$

$$\text{subject to}$$

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

$$U_t\left(C^t, s^t, \varphi_t\right) \geq U_t\left(C^{t-1}, s^t, \varphi_t\right),$$

and non-negativity, where the first constraint is feasibility that must hold for $\tau = t, t+1$ and $\forall \varphi_\tau \geq \varphi_t$ and the second constraint is a form of self-enforcement that must hold $\forall s^t$ (with no such constraint at $t = 0$). The component of the C^{*t-1} allocation after $t-1$ serves as an endogenous outside option for C^{*t}. In this sense, the underlying uncertainty can be seen as leading to a form of endogenous lack of commitment: Even though the government has the ability to fully commit, it may choose not to and instead design optimal policies period by period, subsequently reforming them as necessary.

2.3 History independence

The allocations C^{*t} defined above have limited dependence on future shocks and shock distributions. In particular, they do not depend on beliefs Π_{t+2} or type s_τ for $\tau \geq t+2$. We next show that full history dependence is lost whenever periodic reforms provide an improvement to previously designed policies. That is, whenever a reform results in the government's self-enforcement constraints being slack at t, then the optimal allocation C^{*t} is independent of $s_1^{t-1} \equiv (s_1, \ldots, s_{t-1})$ and the distribution of s_1^{t-1} in the economy.

Proposition 2. *For any t at which the self-enforcement constraints in the government's problem (2) are slack, the optimal allocation C^{*t} is independent of full history.*

Proof. When the self-enforcement constraints in problem (2) do not bind at t, the government must simply maximize an η-weighted average of agents' t continuation utilities, subject to feasibility at t and $t+1$. The continuation utility $U_t\left(C^t, s^t, \varphi_t\right)$ does not depend on s_1^{t-1}

or on the distribution of s_1^{t-1} other than through C^t, and similarly s_1^{t-1} and the distribution of s_1^{t-1} are only relevant to the feasibility constraint through date t capital $K_t^{t-1}(\varphi_{t-1})$, which is fixed at the beginning of t. Thus to maximize its objective, the government will choose the optimal reformed allocation C^{*t} so that it does not depend on s_1^{t-1} or on the distribution of s_1^{t-1} in the economy, given fixed $K_t^{t-1}(\varphi_{t-1})$. □

As a specific example, consider the case in which agents' beliefs are homogeneous and C^{*t} is computed by backward induction for all t.[13] Then C^{*t} will never depend on s_1^{t-1} or on the distribution of s_1^{t-1} for $t \geq 2$. To see this, let $t = 0$ 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 (s_0, θ_1) and the distribution of (s_0, θ_1) in the economy, and we have

$$C_1^{*0}\left(\underline{\varphi}_1, k_1^0\right) \in \arg\max_{C_1^0} \int U_1\left(C_1^0, s^1, \underline{\varphi}_1\right) d\eta\left(s_0\right)$$

subject to

$$\int c_1^0\left(s^1, \underline{\varphi}_1\right) d\underline{\varphi}_1\left(s^1\right) + K_2^0\left(\underline{\varphi}_1\right) \leq f\left(K_1^0\left(\varphi_0\right), Z_t^0\left(\underline{\varphi}_1\right)\right),$$

and non-negativity, where $\underline{\varphi}_1$ is a $t = 1$ type distribution such that every agent realizes the belief set $\underline{\Pi}_2$ at $t = 1$. The government solves this problem for each $\underline{\varphi}_1 \geq \varphi_0$ and each 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 type distribution of the form $\underline{\varphi}_1$ is realized, then by construction the continuation allocation C_1^{*0} is optimal. In this case, the self-enforcement constraints in the $t = 1$ government's problem are obviously slack. If the $t = 1$ type distribution φ_1 is not of the form $\underline{\varphi}_1$, then there exist types s^1 in the support of φ_1 with $\Pi_2 \neq \underline{\Pi}_2$. Under a belief $\pi_2 \notin \underline{\Pi}_2$, the agent believes that at some period $t \geq 2$, a set of agents with positive measure will realize skills $\theta_t > \underline{\theta}$. In this case, the government will be able to have strictly greater output than if all agents realized $\theta_t = \underline{\theta}$, while still providing t continuation utility $U_t\left(C_t^{*0}, s^t, \varphi_t'\right)$ to all agents. In effect, the government's feasibility constraints at $t \geq 2$ are looser under a distribution $\pi_2 \notin \underline{\Pi}_2$ than under $\underline{\pi}_2$. Since the government would not reduce agents' continuation utilities in response to loosened feasibility constraints, it follows that any re-optimized C^{*1} would trivially satisfy the self-enforcement constraints. In particular, if we require that the optimal two-period C^{*t} be constructed by backward induction for all t, then this argument shows that self-enforcement constraints will never bind. The resulting

[13]The assumption of homogeneous beliefs implies that the allocation C^{*0} can be computed through backward induction.

optimal C^{*t} will then be independent of s_1^{t-1} and the distribution of s_1^{t-1} in the economy for $t \geq 2$.

2.4 Discussion

We further discuss two key elements behind the results above. First, our results make use of a condition on agents' beliefs that requires them to be sufficiently similar. Assumption 1 is a convenient and intuitive way to guarantee this condition. That assumption can be weakened in various ways while maintaining the results, and we describe one example of such weakening and discuss others.

Second, it is important to understand the extent to which the results may be driven by a form of dynamic inconsistency, because taking a broader view of uncertainty (in particular, relaxing Savage's Sure-Thing Principle) creates the potential for dynamic inconsistency regardless of the particular preference representation or updating rule used (see, e.g., Machina and Siniscalchi (2014)). We show that agents' preferences satisfy dynamic consistency in the sense commonly used in the literature.

Weaker belief conditions

At 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 $[\underline{\theta}, \theta_{t,T}]$. 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 $s^{\tau-1}$, $\pi_{t,\tau}\left(\cdot\middle| s^{\tau-1}\right)\big|_{\theta}$ is supported in $[\underline{\theta}, \theta_{t,\tau}]$ and $\pi_{t,\tau}\left(\cdot\middle| s^{\tau-1}\right)\big|_{\Pi}$ places unit weight on the belief set $\underline{\Pi}_{t,\tau+1}$. With this definition, any agent who realizes the belief set $\underline{\Pi}_{t,\tau}$ at $\tau \geq t+2$ 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 Π_{t+1} realized by an agent, and any $\pi_{t+1} \in \Pi_{t+1}$, there exists $\pi'_{t+1} \in \Pi_{t+1}$ such that

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

but $\pi'_{t+1}\left(\cdot\middle| s^{t}\right)\big|_{\Pi}$ places unit weight on the belief set $\underline{\Pi}_{t,t+2}$. With this condition, at t any agent thinks it is possible that at $t+1$, every agent 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 completely state-contingent at $t+1$.

In particular, 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 $s_t = \left(\theta_t, \underline{\Pi}_{0,t+1}\right)$ for $t \geq 2$, where $\theta_t \in [\underline{\theta}, \theta_{0,t}]$. 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 an endogenous outside option 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 $\left(\theta_\tau, \underline{\Pi}_{0,\tau+1}\right)$ at $\tau \leq t$ with $\theta_\tau \in [\underline{\theta}, \theta_{0,\tau}]$. If an agent actually realizes a τ skill $\theta_\tau \notin [\underline{\theta}, \theta_{0,\tau}]$, 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 outside option constructed in this example will remain feasible if any agent realizes a type $\theta_t \notin [\underline{\theta}, \theta_{t,\tau}]$ 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 two-period allocation C^1, but it must additionally ensure that each agent receives weakly greater $t = 1$ continuation utility under C^1 than under C^0. By reforming the allocation at each date, the government can find a sequence of two-period allocation $\{C^t\}_{t=0}^T$ that implements the (potentially completely 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 two-period allocations $\{C^t\}_{t=0}^T$ display loss of full history dependence.

Dynamic consistency

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

Lemma 1. *Fix any $t < T$, any φ_t, and any s^t with $\varphi_t\left(s^t\right) > 0$. If $C, \tilde{C}$ coincide at t and*

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

for all $s^{t+1} \geq s^t$ *and all* $\varphi_{t+1} \geq \varphi_t$*, then*

$$U_t\left(C, s^t, \varphi_t\right) \leq U_t\left(\tilde{C}, s^t, \varphi_t\right).$$

According to the lemma, if an agent weakly prefers the allocation $\tilde{C}$ to C at $t+1$ for any s^{t+1} and φ_{t+1} consistent with the realized type s^t and type distribution φ_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_t: Since C and $\tilde{C}$ coincide at t,

$$u\left(c_t\left(s^t, \varphi_t\right), \frac{z_t\left(s^t, \varphi_t\right)}{\theta_t}\right) = u\left(\tilde{c}_t\left(s^t, \varphi_t\right), \frac{\tilde{z}_t\left(s^t, \varphi_t\right)}{\theta_t}\right).$$

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

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

By the recursive definition of the utility function U_t, we have the result.

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

$$U_{t+1}\left(\tilde{C}, s^{t+1}, \tilde{\varphi}_{t+1}\right) < U_{t+1}\left(C, s^{t+1}, \varphi_{t+1}\right)$$

for all $s^{t+1} \geq s^t$ and all $\varphi_{t+1} \neq \tilde{\varphi}_{t+1}$. If $\Pi_{t+1}\left(\tilde{C}\right) = \{\tilde{\varphi}_{t+1}\}$ and $\tilde{\varphi}_{t+1} \notin \Pi_{t+1}(C)$ then inequality (4) will be strictly reversed. More generally, this problem arises because inequality (3) does not imply any restrictions about the values of $U_{t+1}(C, \cdot, \cdot)$ and $U_{t+1}\left(\tilde{C}, \cdot, \cdot\right)$ 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, Axiom 4 in the axiomatization of recursive variational preferences in Maccheroni, et al. (2006b) implies the same definition of dynamic consistency as Lemma 1. Similar assumptions are made by Epstein and Schneider (2003) and Klibanoff, et al. (2009) in their axiomatizations of recursive maxmin preferences and recursive smooth ambiguity preferences, respectively. 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 Π_{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).[14]

3 Private Information and Lack of Commitment

In this section, we relax the assumptions of public information and full commitment. We consider the consequences of agents having private beliefs, as well as further extensions in which current skills are also private information and agents have exogenous outside options. We show that a version of the revelation principle holds for the environment with uncertainty and private beliefs. Our results below indicate that incentive constraints for belief revelation are often non-binding, and they suggest simple ways of implementing efficient policies in economies constrained by uncertainty, asymmetric information, and a lack of full commitment.

3.1 Private beliefs

Consider the following modifications to the baseline setup in Section 2: Agents are privately informed about their belief sets Π_{t+1}, so they must choose reporting strategies and have beliefs about other agents' reports. Let $\hat{\Pi}_{t+1}$ denote a $t+1$ belief set reported at t. Any belief must now be a joint distribution over $t+1$ types s^{t+1} and $t+1$ reported belief histories

[14]Gul and Pesendorfer (2017) define an alternative notion of "weak dynamic consistency" whereby an agent's expected continuation utilities conditional on any event cannot be uniformly lower than his unconditional expected continuation utility (i.e., the resolution of uncertainty cannot make an agent uniformly worse-off). Our Lemma 1 implies that our preferences satisfy this property given the information structure described in Section 2.1. Gul and Pesendorfer (2017) provide an axiomatization of dynamic uncertainty-averse preferences that satisfy this property (among others) for any information structure.

$\hat{\Pi}^{t+1}$, so we have $\pi_{t+1} \in \Pi_{t+1} \subset \Delta\left(s^{t+1}, \hat{\Pi}^{t+1}\right)$. Similarly, now $\varphi_{t+1} \in \Delta\left(s^{t+1}, \hat{\Pi}^{t+1}\right)$. It will be also convenient to let $s^{t,\tau} = \left(\theta^t, \hat{\Pi}^\tau\right)$ denote a t history of skills and a τ history of reported beliefs, with $\varphi_{t,\tau} \in \Delta\left(s^{t,\tau}\right)$ denoting a distribution over $s^{t,\tau}$.

A *reporting strategy* is $\sigma = \{\sigma_t\}_{t=0}^T$, where σ_t maps a type s^t and a distribution $\varphi_{t,t}$ to a reported type $s^{t,t+1}$. Note that skill realizations are still public information. Let Σ denote the set of possible reporting strategies, and let σ^* denote the truth-telling strategy.

Given an allocation $C = \{c_t\left(s^{t,t+1}, \varphi_{t,t+1}\right), z_t\left(s^{t,t+1}, \varphi_{t,t+1}\right), k_{t+1}\left(s^{t,t+1}, \varphi_{t,t+1}\right)\}_{t=0}^T$ and a strategy $\sigma \in \Sigma$, an agent's t continuation utility is defined as

$$U_t\left(C, \sigma, s^t, \varphi_{t,t}\right) \equiv \inf_{\Pi_{t+1}} \left\{ u\left(c_t\left(\sigma_t\left(s^t, \varphi_{t,t}\right), \pi_{t,t+1}\right), \frac{z_t\left(\sigma_t\left(s^t, \varphi_{t,t}\right), \pi_{t,t+1}\right)}{\theta_t}\right) + \mathbb{E}_{\pi_{t+1}}\left[\beta U_{t+1}\left(C, \sigma, s^{t+1}, \pi_{t+1,t+1}\right) \middle| s^t\right] \right\}.$$

Note that the distribution $\varphi_{t,t}$ is known to the agents when forming their $t+1$ beliefs Π_{t+1}, so we impose the consistency condition $\pi_{t+1} \in \Pi_{t+1} \Rightarrow \pi_{t,t} = \varphi_{t,t}$. We also remark that with this specification, the agent does not evaluate an expectation of current-period utility with respect to a distribution over other agents' reports. Rather, to evaluate his total t utility given a distribution $\pi_{t+1} \in \Pi_{t+1}$, the agent simply substitutes the marginal distribution $\pi_{t,t+1}$ into the current-period allocation. This is without loss of generality: We could instead define each $\pi_{t+1} \in \Pi_{t+1}$ as a second-order distribution over the set $\Delta\left(s^{t+1}, \hat{\Pi}^{t+1}\right)$, and the agent would then be forced to evaluate an expectation of current-period utility. Since results below go through with only minor modifications, we adopt the simpler definition of the agents' utility given above.

With this characterization of the information structure and the agents' preferences, we define an *equilibrium strategy* as a strategy $\sigma^e \in \Sigma$ such that

$$U_t\left(C, \sigma^e, s^t, \varphi_{t,t}\right) \geq U_t\left(C, \sigma, s^t, \varphi_{t,t}\right)$$

for all $\sigma \in \Sigma$, all $t = 0, \ldots, T$, all s^t, and all $\varphi_{t,t} \in \Delta\left(s^{t,t}\right)$.

3.2 Revelation principle for private beliefs

We first show that a standard revelation principle extends to the environment with private beliefs.

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

$$U_t\left(\tilde{C},\sigma^*,s^t,\varphi_{t,t}\right)=U_t\left(C,\sigma^e,s^t,\varphi_{t,t}\right)$$

for all $t=0,\ldots,T$*, all* s^t*, and all* $\varphi_{t,t}$*.*

The proof is standard and is provided in Appendix. (Moreover, note that a far more general statement also holds: Given an arbitrary mechanism and an equilibrium strategy, there exists an incentive-compatible direct mechanism that implements the equilibrium outcome. This statement is proven in the same manner as Lemma 2 and we omit the proof.) This justifies focusing on incentive-compatible allocations: an allocation C is *incentive-compatible* if

$$U_t\left(C,\sigma^*,s^t,\varphi_{t,t}\right)\geq U_t\left(C,\sigma,s^t,\varphi_{t,t}\right)$$

for all $\sigma\in\Sigma$, all $t=0,\ldots,T$, all s^t, and all $\varphi_{t,t}$.

Given an initial distribution $\varphi_{0,0}\in\Delta\left(\theta_0\right)$, an allocation C is *Pareto optimal* if there does not exist another allocation $\tilde{C}$ such that[15]

$$U_0\left(\tilde{C},\sigma^*,s_0,\varphi_{0,0}\right)\geq U_0\left(C,\sigma^*,s_0,\varphi_{0,0}\right)$$

for all s_0, where the inequality is strict for some type s_0. An allocation C is *renegotiation-proof* if for every t and every distribution $\varphi_{t,t}$, there does not exist another allocation $\tilde{C}$ such that

$$U_t\left(\tilde{C},\sigma^*,s^t,\varphi_{t,t}\right)\geq U_t\left(C,\sigma^*,s^t,\varphi_{t,t}\right)$$

for all s^t, where the inequality is strict for some type s^t. Assumption 1 can then be modified as follows:

Assumption 2. *Fix any* $t=0,\ldots,T$*, any* $\varphi_{t,t}$*, and any* Π_{t+1}*. Then for any* $\pi_{t+1}\in\Pi_{t+1}$*, and any* $\pi'_{t+1}\geq\varphi_{t,t}$ *such that* π_{t+1} *and* π'_{t+1} *have the same marginal distribution over* s^{t+1}*, we must have* $\pi'_{t+1}\in\Pi_{t+1}$*.*

One interpretation of this form of the condition on beliefs is that agents are completely uncertain about the distribution of $t+1$ reported beliefs $\hat{\Pi}_{t+1}$, and they are willing to consider any such distribution given a marginal distribution over $t+1$ types s^{t+1}. When that is the case, incentive compatibility does not bind:

Proposition 3. *Every renegotiation-proof allocation is incentive-compatible.*

[15]Note that $\varphi_{0,0}$ is simply a distribution over θ_0 because agents have no beliefs to report before $t=0$.

The intuition behind this result is straightforward:[16] Consider a renegotiation-proof allocation C, and suppose that it is not incentive-compatible. Then there exist a date τ, a type $\tilde{s}^{\tau}$, a distribution $\tilde{\varphi}_{\tau,\tau}$, and a strategy σ such that

$$U_{\tau}\left(C,\sigma^{*},\tilde{s}^{\tau},\tilde{\varphi}_{\tau,\tau}\right) < U_{\tau}\left(C,\sigma,\tilde{s}^{\tau},\tilde{\varphi}_{\tau,\tau}\right).$$

We define a new allocation $\bar{C}$ that is equal to C before τ and for all distributions $\varphi_{t,t+1}$, $t \geq \tau$, that do not follow $\tilde{\varphi}_{\tau,\tau}$. For any distribution $\varphi_{t,t+1}$ that follows $\tilde{\varphi}_{\tau,\tau}$, $\bar{C}$ is defined so that agents with type $s^t \geq \tilde{s}^{\tau}$ will receive the same utility by following the truth-telling strategy σ^* as they would under C by following the deviating strategy σ. These agents are obviously strictly better off under $\bar{C}$ than under C, given the truth-telling strategy σ^*. For agents with types $s^t \ngeq \tilde{s}^{\tau}$, $\bar{C}$ is defined so that they receive the same utility by following the truth-telling strategy σ^*, given that agents with type $s^t \geq \tilde{s}^{\tau}$ also follow σ^*, as they would under C, given that agents with type $s^t \geq \tilde{s}^{\tau}$ follow the deviating strategy σ. For example, the consumption function $\bar{c}_t$ in $\bar{C}$ is given by

$$\bar{c}_t\left(s^t,\varphi_{t,t+1}\right) = \begin{cases} c_t\left(\sigma_t\left(s^t,\varphi_{t,t}\right),\bar{\varphi}_{t,t+1}\right) & \text{if } t \geq \tau, s^t \geq \tilde{s}^{\tau}, \varphi_{t,t+1} \geq \tilde{\varphi}_{\tau,\tau}, \\ c_t\left(s^t,\bar{\varphi}_{t,t+1}\right) & \text{if } t \geq \tau, s^t \ngeq \tilde{s}^{\tau}, \varphi_{t,t+1} \geq \tilde{\varphi}_{\tau,\tau}, \\ c_t\left(s^t,\varphi_{t,t+1}\right) & \text{else,} \end{cases}$$

where $\bar{\varphi}_{t,t+1}$ is the distribution such that $\bar{\varphi}_{t,t} = \varphi_{t,t}$, but $\bar{\varphi}_{t,t+1}$ reflects the fact that agents with type $s^t \geq \tilde{s}^{\tau}$ report $\sigma_t\left(s^t,\varphi_{t,t}\right)$ instead of s^t.

Now fix a history $s^t \ngeq \tilde{s}^{\tau}$ and a distribution $\varphi_{t,t} \geq \tilde{\varphi}_{\tau,\tau}$. By Assumption 2, we know that a particular belief distribution π_{t+1} is in the belief set Π_{t+1} if and only if there exists another distribution $\pi'_{t+1} \in \Pi_{t+1}$ with the same marginal distribution over $t+1$ types s^{t+1}. Intuitively, this means that the agent believes any t distribution $\varphi_{t,t+1} \geq \varphi_{t,t}$ to be possible, so he considers all such distributions independently of his consideration of marginal distributions over $t+1$ types s^{t+1}.

When the agent is evaluating his t utility, he takes each possible distribution $\pi_{t+1} \in \Pi_{t+1}$ and calculates the sum of his current-period utility and expected continuation utility, supposing that $\pi_{t+1,t+1}$ will be the realized distribution of $\left(\theta^{t+1},\hat{\Pi}^{t+1}\right)$. After making these calculations, he takes as his utility the infimum value of these sums over all $\pi_{t+1} \in \Pi_{t+1}$. Now the allocation $\bar{C}$ is constructed so that its consumption, effective labor, and capital functions are identical on distribution pairs of the form $\varphi_{t,t+1}$ and $\bar{\varphi}_{t,t+1}$, so the set of

[16]See Appendix for the proof, which is a dynamic generalization of an argument in de Castro and Yannelis (2010).

current-period and expected continuation utility sums under $\bar{C}$ must be a subset of those under C. More specifically,

$$\left\{ u\left(\bar{c}_t\left(s^t,\pi_{t,t+1}\right),\frac{\bar{z}_t\left(s^t,\pi_{t,t+1}\right)}{\theta_t}\right)+\mathbb{E}_{\pi_{t+1}}\left[\beta U_{t+1}\left(\bar{C},\sigma^*,s^{t+1},\pi_{t+1,t+1}\right)\middle|\, s^t\right]\middle|\, \pi_{t+1}\in\Pi_{t+1}\right\}$$
$$\subseteq\left\{ u\left(c_t\left(s^t,\pi_{t,t+1}\right),\frac{z_t\left(s^t,\pi_{t,t+1}\right)}{\theta_t}\right)+\mathbb{E}_{\pi_{t+1}}\left[\beta U_{t+1}\left(C,\sigma^*,s^{t+1},\pi_{t+1,t+1}\right)\middle|\, s^t\right]\middle|\, \pi_{t+1}\in\Pi_{t+1}\right\}.$$

It then follows immediately that agents with types $s^t \not\geq \tilde{s}^\tau$ are weakly better off under $\bar{C}$ than under C, given the truth-telling strategy σ^*. Since agents with τ type $\tilde{s}^\tau$ are strictly better off at τ, this implies that C is not renegotiation-proof, a contradiction.

3.3 Periodic reforms with private beliefs

The utility of the result in Proposition 3 above is apparent when we consider constrained efficient allocations in economies with asymmetric information about beliefs. For a specific illustration, suppose Assumption 2 still holds and hence the government solves the following problem:

$$\arg\max_{C}\int U_0\left(C,\sigma^*,s_0,\varphi_{0,0}\right)d\eta\left(s_0\right)$$

subject to non-negativity and

$$\int c_t\left(s^t,\varphi_t\right)d\varphi_t\left(s^t\right)+K_{t+1}\left(\varphi_t\right)\leq f\left(K_t\left(\varphi_{t-1}\right),Z_t\left(\varphi_t\right)\right),$$
$$U_t\left(C,\sigma^*,s^t,\varphi_{t,t}\right)\geq U_t\left(C,\sigma,s^t,\varphi_{t,t}\right)$$

Here the feasibility constraint holds for $t=0,\ldots,T$ and $\forall\varphi_t\geq\varphi_0$, and the incentive-compatibility constraint holds for $t=0,\ldots,T$, $\forall\sigma\in\Sigma$, $\forall s^t$, and $\forall\varphi_{t,t}\geq\varphi_{0,0}$. It is clear that without loss of generality, we can restrict our choice set to renegotiation-proof allocations. By Proposition 3, the incentive-compatibility constraint is non-binding, so it now suffices to solve the symmetric information government problem (2) considered in Section 2 for each initial distribution φ_0. If agents' beliefs satisfy Assumption 1, we can immediately apply Proposition 1 to conclude that the government can solve this problem via a sequence of two-period allocations $\{C^t\}_{t=0}^T$ with endogenous outside options. This implies as before the properties of periodic reforms as well as loss of full history dependence.

3.4 Private skills and lack of commitment

We next show that the method used to prove Proposition 3 can be used in a straightforward way to simplify incentive constraints when agents also have private information about their skill shocks θ_t. Moreover, we can also apply the analysis from Section 2 when the agents lack exogenous commitment, i.e., when agents have access to an exogenous outside option and have to be compelled to participate in the government's allocation. In that case self-enforcement constraints appear in the government's problem (2) but Proposition 1 applies essentially unchanged. That is, an efficient allocation can still be implemented with a sequence of two-period allocations but with a modified self-enforcement constraint at each date, once again with the same implications for history independence as above.

Private skills

Suppose now agents are privately informed about both their beliefs Π_{t+1} and their skills θ_t at t. In addition to reporting beliefs $\hat{\Pi}_{t+1}$ at t, agents must also report a skill $\hat{\theta}_t$. Let $\hat{\theta}_t$ denote a skill reported at t. Then $s^{t,t+1} = \left(\hat{\theta}^t, \hat{\Pi}^{t+1}\right)$ denotes a t history of reported skills and a $t+1$ history of reported beliefs, called a t *reported type*. Any belief must now be a joint distribution over $t+1$ types s^{t+1} and t reported types $s^{t,t+1}$, so $\pi_{t+1} \in \Pi_{t+1} \subset \Delta\left(s^{t+1}, s^{t,t+1}\right)$. Similarly, $\varphi_{t+1} \in \Delta\left(s^{t+1}, s^{t,t+1}\right)$ now denotes a distribution over $t+1$ types and t reported types, and $\varphi_{t,t+1} \in \Delta\left(s^{t,t+1}\right)$ denotes the marginal distribution of φ_{t+1} over t reported types.

Any report function σ_t in a reporting strategy $\sigma = \{\sigma_t\}_{t=0}^{T}$ maps a type s^t and a distribution $\varphi_{t-1,t}$ to a reported type $s^{t,t+1}$. The set of all reporting strategies is again denoted Σ, and σ^* is the truth-telling strategy. Given a strategy $\sigma \in \Sigma$ and an allocation $C = \left\{c_t\left(s^{t,t+1}, \varphi_{t,t+1}\right), z_t\left(s^{t,t+1}, \varphi_{t,t+1}\right), k_{t+1}\left(s^{t,t+1}, \varphi_{t,t+1}\right)\right\}_{t=0}^{T}$, an agent's t continuation utility is

$$U_t\left(C, \sigma, s^t, \varphi_{t-1,t}\right) \equiv \inf_{\Pi_{t+1}} \left\{ u\left(c_t\left(\sigma_t\left(s^t, \varphi_{t-1,t}\right), \pi_{t,t+1}\right), \frac{z_t\left(\sigma_t\left(s^t, \varphi_{t-1,t}\right), \pi_{t,t+1}\right)}{\theta_t}\right) + \mathbb{E}_{\pi_{t+1}}\left[\beta U_{t+1}\left(C, \sigma, s^{t+1}, \pi_{t,t+1}\right)\middle| s^t\right] \right\}.$$

Equilibrium, incentive-compatibility, Pareto optimality, and renegotiation-proofness are all defined as in Sections 3 and 3.1, and a revelation principle analogous to Lemma 2 holds. To demonstrate that all renegotiation-proof allocations are incentive-compatible, we will assume the analogue of Assumption 2 in this economy: For any $\varphi_{t-1,t}$, any $\pi_{t+1} \in \Pi_{t+1}$, and any $\pi'_{t+1} \geq \varphi_{t-1,t}$ such that π_{t+1} and π'_{t+1} have the same marginal distribution over s^{t+1}, we must have $\pi'_{t+1} \in \Pi_{t+1}$. As with Assumption 2, this condition implies that agents are completely uncertain about the distribution of t reported types $s^{t,t+1}$, and they are willing to consider any such distribution given a marginal distribution over $t+1$ types s^{t+1}.

Given this characterization of beliefs, incentive-compatibility constraints are again satisfied for all renegotiation-proof allocations: Following the proof of Proposition 3, take an allocation C, and suppose that it is renegotiation-proof but not incentive-compatible. Then there exist a lowest date τ, a type $\tilde{s}^{\tau}$, a $\tau-1$ reported type distribution $\tilde{\varphi}_{\tau-1,\tau}$, and a strategy σ such that an agent with type $\tilde{s}^{\tau}$ strictly prefers σ to σ^*. Define a new allocation $\bar{C}$ that coincides with C before τ and after τ in the case that the distribution $\tilde{\varphi}_{\tau-1,\tau}$ is not realized at $\tau-1$. If $\tilde{\varphi}_{\tau-1,\tau}$ is realized at $\tau-1$, $\bar{C}$ is adjusted at $t\geq\tau$ as follows: Any agent with type $s^t\geq\tilde{s}^{\tau}$ receives the same allocation in $\bar{C}$ with σ^* as he would have received in C with σ. Clearly any such agent strictly prefers $\bar{C}$ to C when evaluating both allocations with the truth-telling strategy σ^*. Any agent with type $s^t\not\geq\tilde{s}^{\tau}$ receives the same allocation in $\bar{C}$ with σ^* as he would have received in C with σ^* if all agents with type $s^t\geq\tilde{s}^{\tau}$ used the deviating strategy σ. Such an agent is actually indifferent between C and $\bar{C}$ because agents are completely uncertain about the t distribution of reported types. This contradicts the assumption that C is renegotiation-proof, so we have the result.

This fact significantly simplifies the implementation of optimal allocations when agents have asymmetric information over shocks s_t. For example, suppose that the government solves the following problem:

$$\max_{C}\int U_0\left(C,\sigma^*,s_0\right)d\eta\left(s_0\right)$$

subject to non-negativity and

$$\int c_t\left(s^t,\varphi_t\right)d\varphi_t\left(s^t\right)+K_{t+1}\left(\varphi_t\right)\leq f\left(K_t\left(\varphi_{t-1}\right),Z_t\left(\varphi_t\right)\right),$$

$$U_t\left(C,\sigma^*,s^t,\varphi_{t-1,t}\right)\geq U_t\left(C,\sigma,s^t,\varphi_{t-1,t}\right)$$

Here the feasibility constraint holds for $t=0,\ldots,T$ and $\forall\varphi_t$, and the incentive-compatibility constraint holds for $t=0,\ldots,T$, $\forall\sigma\in\Sigma$, $\forall s^t$, and $\forall\varphi_{t-1,t}$. It is clear that a renegotiation-proof allocation is in the solution set to this problem, so without loss of generality we can restrict the government to renegotiation-proof allocations. In this case, we know that the incentive-compatibility constraint is non-binding, so the government's problem above reduces essentially to problem (2) addressed in Section 2.1:

$$\max_{C} \int U_0\left(C, s_0, \varphi_0\right) d\eta\left(s_0\right)$$

subject to non-negativity and

$$\int c_t\left(s^t, \varphi_t\right) d\varphi_t\left(s^t\right) + K_{t+1}\left(\varphi_t\right) \leq f\left(K_t\left(\varphi_{t-1}\right), Z_t\left(\varphi_t\right)\right)$$

Here $U_0\left(C, s_0, \varphi_0\right)$ is defined as in Section 2.1. The sole difference between this problem and the one described in Section 2.1 regards the government's information about the initial distribution φ_0. In this problem, we do not assume that the government can observe the distribution φ_0 when it is realized at the beginning of $t = 0$, so a solution to this problem must provide a contingency for each possible distribution φ_0. Hence the feasibility constraint holds for $t = 0, \ldots, T$ and $\forall \varphi_t$, and we denote a solution to this government's problem by $C^* \equiv \left\{C^*\left(\varphi_0\right)\right\}_{\varphi_0}$.

If we assume that agents' beliefs additionally satisfy Assumption 1, then we can apply the arguments from Section 2.2 to further simplify the government's problem. In this case, any efficient allocation can be implemented by a sequence of two-period allocations with periodic reforms at each date. Since the arguments detailed in this section imply that the government's problem with asymmetric information about shocks is essentially the same as the symmetric information problem, we can simply apply the argument of Section 2.2 to each component $C^*\left(\varphi_0\right)$ of the efficient allocation C^*.

Lack of commitment

Consider now agents that have an exogenous inability to commit in addition to the private information about types introduced above. Suppose that in addition to private information over shocks s_t agents now have access to an exogenous outside option that yields type-contingent utility $\underline{u}_t\left(s^t\right)$ at t. After receiving their type shocks at t and making their reports, agents have the opportunity to leave any allocation C and take their outside option at $t, \ldots, T$. We assume that agents cannot return to an allocation once they have left. We recursively define the t continuation utility under the outside option by

$$\underline{U}_t\left(s^t\right) \equiv \underline{u}_t\left(s^t\right) + \inf_{\Pi_{t+1}} \mathbb{E}_{\pi_{t+1}}\left[\underline{U}_{t+1}\left(s^{t+1}\right) \middle| s^t\right].$$

An allocation C is *self-enforcing* if

$$U_t\left(C, \sigma, s^t, \varphi_{t-1,t}\right) \geq \underline{U}_t\left(s^t\right)$$

for $t = 0, \ldots, T$, $\forall s^t$, $\forall \sigma$, and $\forall \varphi_{t-1,t}$.

At $t = 0$, a government wishes to implement an allocation that maximizes an η-weighted average of agents' $t = 0$ continuation utilities. However, the government must now design the allocation so that agents choose to truthfully reveal their types and weakly prefer staying in the allocation to their outside option. Thus the government must solve the problem

$$\max_C \int U_0 \left(C, \sigma^*, s_0\right) d\eta \left(s_0\right)$$

subject to non-negativity and

$$\int c_t \left(s^t, \varphi_t\right) d\varphi_t \left(s^t\right) + K_{t+1} \left(\varphi_t\right) \leq f \left(K_t \left(\varphi_{t-1}\right), Z_t \left(\varphi_t\right)\right),$$

$$U_t \left(C, \sigma^*, s^t, \varphi_{t-1,t}\right) \geq U_t \left(C, \sigma, s^t, \varphi_{t-1,t}\right),$$

$$U_t \left(C, \sigma^*, s^t, \varphi_{t-1,t}\right) \geq \underline{U}_t \left(s^t\right)$$

Here the feasibility, incentive-compatibility, and self-enforcement constraints hold for $t = 0, \ldots, T$, $\forall s^t$, $\forall \sigma$, $\forall \varphi_t$, and $\forall \varphi_{t-1,t}$.[17] Given that Assumption 2 holds and agents are completely uncertain about the distribution of reports, we can follow the same analysis given in the private types economy to conclude that any renegotiation-proof allocation is incentive-compatible. We can obviously assume without loss of generality that the solution to the government's problem is renegotiation-proof, so we can simplify the government's problem by restricting the choice set to renegotiation-proof allocations and assuming that incentive-constraints are non-binding. Thus the government must solve the symmetric information problem given in Section 2.1 for each type distribution φ_0, but with the additional self-enforcement constraint

$$U_t \left(C, s^t, \varphi_t\right) \geq \underline{U}_t \left(s^t\right)$$

for each $t = 0, \ldots, T$, $\forall s^t$, and $\forall \varphi_t$. We let $C^{**} \equiv \left\{C^{**} \left(\varphi_0\right)\right\}_{\varphi_0}$ denote the constrained-efficient, renegotiation-proof solution to this problem.

If agents' beliefs additionally satisfy Assumption 1, then the analogue of Proposition 1 holds in this economy: The constrained-efficient allocation C^{**} can be implemented by a sequence of two-period allocations $\{C^t\}_{t=0}^T$. This is proven by applying the argument in the proof of Proposition 1 to each component $C^{**} \left(\varphi_0\right)$ of the constrained-efficient allocation C^{**}: Given $C^{**} \left(\varphi_0\right)$, we define a new allocation $C^0 = \left\{C_t^0 \left(\varphi_t\right)\right\}_{t=0}^T$ that coincides with $C^{**} \left(\varphi_0\right)$ at $t = 0$, but allocates consumption, effective labor, and capital at $t \geq 1$ as if every agent realized the belief set $\underline{\Pi}_2$ and date $t \geq 2$ shock $\left(\underline{\theta}, \underline{\Pi}_{t+1}\right)$. Hence the policy functions in C^0

[17] In the analysis below, we assume that the outside option is such that the government's constraint set is non-empty.

depend only on (s_0, θ_1) and the distribution of (s_0, θ_1) in the economy. Under Assumption 1, agents weakly prefer the two-period allocation C^0 to $C^{**}(\varphi_0)$.

At $t=1$, the two-period allocation C^0 may not be Pareto optimal with respect to agents' $t=1$ continuation utilities, and it may not even be preferable to the outside option since C^0 is not fully state-contingent. For these reasons, the government may wish to reform the allocation after $t=1$ types are realized. By the same argument as above, it can simply design a new two-period allocation C^1 that delivers each agent weakly greater $t=1$ continuation utility than C^0 or the outside option. In this sense, the $t=1$ continuation allocation in C^0 serves as an endogenous outside option that provides a lower bound on agents' $t=1$ continuation utilities.

In general, given a two-period allocation C^{t-1} designed at $t-1$ and a t type distribution φ_t, the government with social welfare weights η designs the optimal two-period allocation $C^{**t}(\varphi_t, C^{t-1})$ as the solution to the problem

$$\max_{C^t} \int U_t\left(C^t, s^t, \varphi_t\right) d\eta\left(s_0\right)$$

subject to non-negativity and

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

$$U_t\left(C^t, s^t, \varphi_t\right) \geq \max\left\{U_t\left(C^{t-1}, s^t, \varphi_t\right), \underline{U}_t\left(s^t\right)\right\}$$

Here the feasibility constraint holds for $\tau = t, t+1$ and $\forall \varphi_\tau \geq \varphi_t$, and the self-enforcement constraint holds $\forall s^t$. This problem also applies when $t=0$, but only the outside option utility $\underline{U}_0(s_0)$ appears in the self-enforcement constraint. By solving the problem at $t = 0, \ldots, T$, we find a sequence of two-period allocations $\{C^{*t}\}_{t=0}^T$ that implements the efficient allocation C^{**}.

4 Decentralization

In conventional economies without uncertainty, competitive equilibrium results in efficient allocations even when types are privately known to agents and markets are incomplete (Prescott and Townsend (1984), Atkeson and Lucas (1992)). 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)). With uncertainty, we prove that competitive equilibrium may be inefficient when types are privately known to agents and markets are not complete. This implies that government insurance does not necessarily crowd out private insurance unless there exists a *full* (allowing for the broader view of uncertainty) set of Arrow-Debreu securities. Another implication is that this is also the case with public information. 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 competitive equilibrium can be provided with sequences of simplified, short-term allocations that are periodically reformed.

4.1 Inefficiency of competitive equilibrium

Consider once again the economy of Section 3.4 in which agents are privately informed about their shocks s_t at each date. To decentralize this economy, introduce a continuum of 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 agents to supply effective labor z_t. In return, the agents receive consumption c_t. 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 type revelation by the agents. Each firm then adopts the beliefs of the agent with whom it contracts, i.e., the firms face the same uncertainty as the agents about the data-generating process.

Consider first complete markets where firms can trade fully contingent Arrow-Debreu securities among themselves to insure against risk from agents' idiosyncratic skills and against uncertainty. That is, at the beginning of t, firms can trade securities that are contingent on the idiosyncratic shock $s_{t,t+1}$ reported by an agent at t and on the t distribution of reported types $\varphi_{t,t+1}$. In equilibrium, however, securities contingent on idiosyncratic shocks will not be traded:[18]

Lemma 3. *Securities that are contingent on idiosyncratic shocks are not traded in a competitive equilibrium with uncertainty.*

To see this, it is convenient to reinterpret the equilibrium as agents having direct access to the production function f and the ability to trade Arrow-Debreu securities among

[18]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)).

themselves. Fix t, and let $a_i(s_{t,t+1})$ denote a security that pays one unit of consumption after reporting at t if agent i reports shock $s_{t,t+1}$.[19] Let $q_{i,t}(s_{t,t+1}|\varphi_t, \varphi_{t-1,t})$ denote the equilibrium price of this security, and let $Q_t(\varphi_t, \varphi_{t-1,t})$ denote the price of a risk-free bond that pays one unit of consumption after reporting at t. In equilibrium, we must have $q_{i,t}(s_{t,t+1}|\varphi_t, \varphi_{t-1,t}) \leq Q_t(\varphi_t, \varphi_{t-1,t})$ because $a_i(s_{t,t+1})$ pays a unit of consumption only when agent i reports shock $s_{t,t+1}$, while a risk-free bond pays one unit of consumption in all t realizations. If the inequality were strict, agent i could purchase arbitrarily many $a_i(s_{t,t+1})$ securities and sell a corresponding number of risk-free bonds. By reporting type $s_{t,t+1}$ at t regardless of his actual type realization, agent i could ensure that he would net arbitrarily high profits from these trades. Since sellers of the securities $a_i(s_{t,t+1})$ would be guaranteed to lose, we cannot have $q_{i,t}(s_{t,t+1}|\varphi_t, \varphi_{t-1,t}) < Q_t(\varphi_t, \varphi_{t-1,t})$. 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 and securities that are contingent on the t distribution of reported types $\varphi_{t,t+1}$. Let $a(\varphi_{t,t+1})$ denote a security that pays one unit of consumption if the distribution of reported types is $\varphi_{t,t+1}$, and let $q_t(\varphi_{t,t+1}|\varphi_t, \varphi_{t-1,t})$ be its price in equilibrium. The following proposition shows that if a market for a security of the form $a(\varphi_{t,t+1})$ is not available, there exist economies in which competitive equilibria are inefficient.

Proposition 4. *Competitive equilibria with uncertainty may be inefficient if markets are not complete.*

Proof. It suffices to construct an example in which inefficiencies result from markets that are not complete. Note first that without securities that are contingent on idiosyncratic shocks, agents are indifferent between all reporting strategies $\sigma \in \Sigma$. Hence we can assume without loss of generality that agents adopt the truth-telling strategy σ^*. This is then equivalent to a decentralized version of the public information economy of Section 2 in which agents can trade securities that are contingent on the observable t type distribution $\varphi_t \in \Delta(s^t)$.[20]

Fix $t < T$, and let φ_t be the realized distribution of types s^t. Let φ_t place positive measure on a type $\tilde{s}^t$ such that $\tilde{\theta}_t = \bar{\theta}$ and $\tilde{\Pi}_{t+1} = \left\{\underline{\pi}_{t+1}, \bar{\pi}_{t+1}\right\}$. Here $\underline{\pi}_{t+1}$ is the type distribution described in Section 2.2 such that every agent realizes the $t+1$ shock $\left(\underline{\theta}, \underline{\Pi}_{t+2}\right)$ and believes that all agents in the economy will realize skill $\underline{\theta}$ at date $\tau \geq t+2$. Similarly, $\bar{\pi}_{t+1}$ is the type distribution such that every agent realizes skill $\bar{\theta}$ at $t+1$, and moreover

[19]The security $a_i(s_{t,t+1})$ may be contingent on other uncertain properties of the economy, but we abuse notation and write only agent i's idiosyncratic shock to simplify the explanation.

[20]As a result, the proposition also applies to economies in which idiosyncratic shocks are publicly observable.

every agent believes that everyone will realize skill $\bar{\theta}$ at date $\tau \geq t+2$. A type $\tilde{s}^t$ agent chooses his own allocation of consumption, effective labor, capital investment, and securities purchases for $\tau \geq t$ to maximize his t continuation utility U_t, so to insure himself against his uncertainty about the data-generating process, at the beginning of $t+1$ he will seek to purchase the security $a\left(\underline{\pi}_{t+1}\right)$ while selling the security $a\left(\bar{\pi}_{t+1}\right)$. Risk-free securities are an imperfect substitute for $a\left(\underline{\pi}_{t+1}\right)$ securities, so the t continuation utility of a type $\tilde{s}^t$ agent is strictly lower when agents cannot trade in a market for the security $a\left(\underline{\pi}_{t+1}\right)$. The Weak Axiom of Revealed Preference implies that agents with other types also have weakly lower t continuation utility without access to the security $a\left(\underline{\pi}_{t+1}\right)$, so the equilibrium allocation is not Pareto optimal. By a standard argument an allocation is Pareto optimal if and only if it solves the government's problem (1) for a set of Pareto weights η, hence the competitive equilibrium is inefficient. □

More generally, if some but not all distribution-contingent securities are available (i.e., if agents do not have access to securities that are contingent on the future type distributions φ_{t+1} that they think are possible), there may be room for a strict Pareto improvement on the competitive equilibrium allocation. This implies that uncertainty creates a role for government in providing insurance when markets are not complete.

4.2 Periodic reforms in equilibrium

Even with the potential for inefficiency in equilibrium, agents may still obtain some degree of insurance in a decentralized economy. In this subsection, we show that any insurance provided by a competitive equilibrium can be obtained with simplified, short-term allocations that are periodically reformed.

In the decentralized economy introduced above, at $t=0$ each agent solves for a fully contingent allocation

$$C=\left\{c_t\left(s^t,\varphi_t\right),z_t\left(s^t,\varphi_t\right),k_{t+1}\left(s^t,\varphi_t\right),a_t\left(\varphi_t\middle|\,s^{t-1},\varphi_{t-1}\right),b_t\left(s^{t-1},\varphi_{t-1}\right)\right\}_{t=0}^{T},$$

taking prices $q_t\left(\varphi_t\middle|\,\varphi_{t-1}\right)$ and $Q_t\left(\varphi_{t-1}\right)$ as given. Here $a_t\left(\varphi_t\middle|\,s^{t-1},\varphi_{t-1}\right)$ denotes the agent's holdings of securities that pay one unit of consumption at t if the type distribution φ_t is realized, and b_t similarly denotes the agent's holdings of risk-free bonds that pay one unit of consumption after types are realized at t.[21] To see the result, it is helpful to again assume that agents' beliefs satisfy the intuitive Assumption 1. In this case, a decentralized version

[21] Lemma 3 and subsequent arguments justify the focus here on the public information version of the economy.

of Proposition 1 holds:

Proposition 5. *For any set of prices* $\{q_t(\varphi_t|\varphi_{t-1}), Q_t(\varphi_{t-1})\}_{t=0}^{T-1}$ *and any allocation* C*, there exists a sequence of two-period allocations* $\{C^t\}_{t=0}^{T}$ *such that*

$$U_0(C, s_0, \varphi_0) = U_0\left(C^0, s_0, \varphi_0\right) \quad \forall s_0, \forall \varphi_0.$$

Here a two-period allocation C^t is a set of consumption, effective labor, capital, and securities functions that specify an agent's allocative actions at t and $t+1$. As before, each two-period allocation C^t also effectively provides an endogenous outside option, i.e., a contingency plan for $\tau \geq t+2$ in the case that the agent does not choose to adjust the allocation at $t+1$ or $t+2$. We omit the proof since it follows essentially identical steps of the proof of Proposition 1. Once again, the key principle is that at date 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: Linearity

Finally, we consider whether even simpler linear or affine policies can be optimal. In particular, we ask if an optimal two-period allocation C^{*t} can be implemented with policy functions that are affine in individual income, $f(k_t, z_t)$. Our results below indicate 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 generalizing these results. For example, risk-aversion and elastically supplied labor provide significant roadblocks, as does a continuum of measureless agents. Where 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.

For simplicity, we present the arguments in the baseline economy of Section 2 and with a finite number of agents indexed by $i \in \{1, \ldots, N\}$. The period t shock of agent i is denoted $s_{i,t} \equiv (\theta_{i,t}, \Pi_{i,t+1})$, his period t type is denoted $s_i^t \equiv (s_{i,0}, \ldots, s_{i,t})$, and as in Section 2 types are public information. We let $h^t \equiv (s_1^t, \ldots, s_N^t)$ be the vector of the agents' period t types, which we call the date t *state.* A belief set $\Pi_{i,t+1}$ is now a set of distributions $\pi_{i,t+1}$ over h^{t+1}.

An allocation $C \equiv \{c_t(h^t), z_t(h^t), k_t(h^{t-1})\}_{t=0}^{T}$ is a sequence of N-vector-valued consumption, effective labor, and capital functions. The t consumption of agent i, for example, is given by $c_{i,t}(h^t)$, the i^{th} entry of $c_t(h^t)$. Let $C_i \equiv \{c_{i,t}(h^t), z_{i,t}(h^t), k_{i,t+1}(h^t)\}_{t=0}^{T}$ denote the i^{th} component of the allocation C. Define the aggregate capital and effective labor functions

$$K_t\left(h^{t-1}\right) \equiv \sum_{i=1}^{N} k_{i,t}\left(h^{t-1}\right),$$

$$Z_t\left(h^{t}\right) \equiv \sum_{i=1}^{N} z_{i,t}\left(h^{t}\right).$$

For an initial state h_0, an allocation C is *feasible* if for every t and every $h^t \geq h_0$,

$$\sum_{i=1}^{N} c_{i,t}\left(h^{t}\right) + K_{t+1}\left(h^{t}\right) \leq f\left(K_t\left(h^{t-1}\right), Z_t\left(h^{t}\right)\right).$$

Given an allocation C, agent i's period t continuation utility is defined by

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

The government designs an allocation C for the economy after the $t = 0$ state h_0 is realized. Given h_0 and a non-negative vector of social welfare weights $\eta \in \mathbb{R}^N_+$, the *efficient* allocation $C^*\left(h_0\right)$ is defined by

$$C^*\left(h_0\right) \in \arg\max_{C} \sum_{i=1}^{N} U_{i,0}\left(C_i, h_0\right)\eta_i,$$

subject to feasibility and non-negativity.

We modify Assumption 1 as follows:

Assumption 3. *For any t, i, j, and h^t, there exists $\bar{\theta}_{j|i,t+1}\left(h^t\right) \in \Theta$ such that $\pi_{i,t+1} \in \Pi_{i,t+1}$ if and only if $\mathbb{E}_{\pi_{i,t+1}}\left[\theta_{j,t+1}\middle|\, h^t\right] = \bar{\theta}_{j|i,t+1}\left(h^t\right)$ and all $t+1$ skill shocks $\theta_{1,t+1}, \ldots, \theta_{N,t+1}$ are independent.*[22]

Under Assumption 3, at t each agent i forms an expectation $\bar{\theta}_{j|i,t+1}\left(h^t\right)$ about agent j's $t+1$ skill $\theta_{j,t+1}$. Agent i believes that the data-generating process is such that the expected

[22]As before, this assumption is stronger than necessary but more convenient. The arguments below only require that for 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 h^t,

2. $\pi'_{i,t+1}$ places weight only on $\left\{\underline{\theta}, \bar{\theta}\right\}$ to satisfy

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

$t+1$ skill of agent j is $\bar{\theta}_{j|i,t+1}(h^t)$, and he is willing to consider any distribution $\pi_{i,t+1}$ over the $t+1$ state h^{t+1} such that skills are independently distributed and the expected value of $\theta_{j,t+1}$ under $\pi_{i,t+1}$ is $\bar{\theta}_{j|i,t+1}(h^t)$ for all $j=1,\ldots,N$. We require the independence condition 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 results concern policies or continuation utilities that are linear with respect to agent i's skill shock, holding other agents' skill shocks fixed.

The characterization of agents' beliefs in Assumption 3 satisfies Assumption 1, so a finite-agents analogue of Proposition 1 holds: The efficient allocation C^* can be implemented by a sequence of two-period allocations $\{C^t\}_{t=0}^T$, where each C^t specifies an endogenous outside option. The proof of this proposition is almost identical to that of Proposition 1, so it is omitted. However, it is crucial to note that in the period t two-period 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 (h^t, θ_{t+1}), where $\theta_{t+1} \equiv (\theta_{1,t+1}, \ldots, \theta_{N,t+1})$ denotes the N-vector of t skill realizations. This holds because C^t can be constructed by assuming that all agents will realize maximally pessimistic beliefs $\underline{\Pi}_{i,t+2}$ at $t+1$ and type shock $\left(\underline{\theta}, \underline{\Pi}_{i,\tau+1}\right)$ at $\tau \geq t+2$.[23]

With this implementation result, we can assume without loss of generality that the government implements C^* by designing a sequence of two-period allocations $\{C^{*t}\}_{t=0}^T$, where each C^{*t} specifies an endogenous outside option and satisfies

$$
\begin{aligned}
C^{*t}\left(h_0, C^{*t-1}\right) \in \arg\max_{C^t} & \sum_{i=1}^N U_{i,t}\left(C_i^t, h^t\right)\eta_i \\
& \text{subject to} \\
& \sum_{i=1}^N c_{i,\tau}^t\left(h^\tau\right) + K_{\tau+1}^t\left(h^\tau\right) \leq f\left(K_\tau^t\left(h^{\tau-1}\right), Z_\tau^t\left(h^\tau\right)\right), \\
& U_{i,t}\left(C_i^t, h^t\right) \geq U_{i,t}\left(C_i^{t-1}, h^t\right),
\end{aligned}
$$

and non-negativity, where the feasibility constraint holds for $\tau = t, \ldots, T$ and $\forall h^\tau \geq h^t$, and the self-enforcement constraint holds for $i = 1, \ldots, N$.

We seek to determine conditions under which the solution to the date t government's 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 $C^{*0}(h_0)$ to the government's $t=0$ problem and

[23]See Section 2.2 for details.

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.

Agent i's $t=1$ instantaneous utility

$$u\left(c_{i,1}^{*}\left(h_0,\theta_1\right),\frac{z_{i,1}^{*}\left(h_0,\theta_1\right)}{\theta_{i,1}}\right)$$

is not generally an affine function of his own skill $\theta_{i,1}$, holding fixed $(h_0,\theta_{-i,1})$. 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 $(h_0,\theta_{-i,1})$, $u\left(\hat{c}_{i,1},\hat{z}_{i,1}/\theta_{i,1}\right)$ 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 $u\left(\hat{c}_{i,1},\hat{z}_{i,1}/\theta_{i,1}\right)$ 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

$$\left(\tilde{c}_{i,1}\left(h_0,\theta_1\right),\tilde{z}_{i,1}\left(h_0,\theta_1\right)\right)=\begin{cases}\left(\hat{c}_{i,1}\left(h_0,\theta_1\right),\hat{z}_{i,1}\left(h_0,\theta_1\right)\right) & \text{if } u\left(\hat{c}_{i,1}\left(h_0,\theta_1\right),\frac{\hat{z}_{i,1}\left(h_0,\theta_1\right)}{\theta_{i,1}}\right)\\ & \quad\geq u\left(c_{i,1}^{*}\left(h_0,\theta_1\right),\frac{z_{i,1}^{*}\left(h_0,\theta_1\right)}{\theta_{i,1}}\right),\\ \left(c_{i,1}^{*}\left(h_0,\theta_1\right),z_{i,1}^{*}\left(h_0,\theta_1\right)\right) & \text{else.}\end{cases}$$

Hence $\tilde{c}_{i,1},\tilde{z}_{i,1}$ 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 date t instantaneous utility under allocation C^{*0} is given by

$$u\left(c_{i,t}^{*}\left(h_0,\theta_1\right),\frac{z_{i,t}^{*}\left(h_0,\theta_1\right)}{\underline{\theta}}\right),$$

so we let $\hat{c}_{i,t},\hat{z}_{i,t}$ be such that for fixed $(h_0,\theta_{-i,1})$, $u\left(\hat{c}_{i,1},\hat{z}_{i,1}/\underline{\theta}\right)$ 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 $\left\{c_{i,t}^{*},z_{i,t}^{*}\right\}_{t=0}^{T}$.

As the next proposition shows, 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* $\hat{C}_i^0 \equiv \left\{\hat{c}_{i,t}, \hat{z}_{i,t}, k_{i,t}^*\right\}_{t=0}^T$, $\tilde{C}_i^0 \equiv \left\{\tilde{c}_{i,t}, \tilde{z}_{i,t}, k_{i,t}^*\right\}_{t=0}^T$. *Then*

$$U_{i,0}\left(\hat{C}_i^0, h_0\right) = U_{i,0}\left(\tilde{C}_i^0, h_0\right).$$

The proposition holds by the following argument: To evaluate his $t = 0$ utility with allocation $\tilde{C}_i^0$, 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 $(h_0, \theta_{-i,1})$, 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 $\pi'_{i,1}\left(\cdot \mid h_0, \theta_{-i,1}\right)\big|_{\theta_i}$ places weight only on $\left\{\underline{\theta}, \bar{\theta}\right\}$ so as to satisfy

$$\mathbb{E}_{\pi_{i,1}}\left[\theta_{i,1} \mid h_0\right] = \bar{\theta}_{i|i,1}\left(h_0\right).$$

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 $\pi_{i,1}\left(\cdot \mid h_0, \theta_{-i,1}\right)\big|_{\theta_i}$ 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 $\pi_{i,1}\left(\cdot \mid h_0, \theta_{-i,1}\right)\big|_{\theta_i}$ 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 $\theta_{i,1} \in \left\{\underline{\theta}, \bar{\theta}\right\}$. Distributions of the form $\pi'_{i,1}$ place weight only on events with $\theta_{i,1} \in \left\{\underline{\theta}, \bar{\theta}\right\}$, 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 we find that 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}. In particular, the government will always seek to implement $t = 0$ allocations in which each agent's instantaneous utility at $t \geq 1$ is affine in his own $t = 1$ skill. This result is significant in that it demonstrates an interesting property of the uncertainty-averse preferences we have adopted. However, we argue that the assumptions needed to prove this result are so strong that affine policies cannot be expected to be optimal in practical environments.

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. Moreover, the use of this condition makes the result difficult to generalize to a setting with a continuum of agents. 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 in such a continuous model. 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.

In order to prove a more concrete property about affine consumption functions, 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 $\bar{l}_{i,t}(h_0)$. 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 type. We collect this observation in the following corollary:

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*

$$\hat{c}_{i,t}(h_0, \theta_1) \equiv \frac{\bar{\theta} - \theta_{i,1}}{\bar{\theta} - \underline{\theta}} c_{i,t}^*(h_0, \theta_{-i,1}, \underline{\theta}) + \frac{\theta_{i,1} - \underline{\theta}}{\bar{\theta} - \underline{\theta}} c_{i,t}^*\left(h_0, \theta_{-i,1}, \bar{\theta}\right).$$

6 Conclusion

This paper argued that allowing a broader view of uncertainty drastically changes the conclusions about what types of policies are optimal in providing social insurance and redistribution. In otherwise conventional dynamic macroeconomic environments, efficiency is simpler to characterize than under a narrower view of uncertainty limited to the standard notion of risk. Efficient allocations themselves also can be simpler, independent of the full history of idiosyncratic shocks, and reformed periodically, akin to what is commonly observed in reality. However, linear policies can be far from optimal. Decentralized economies with uncertainty as a friction may not efficient unless markets are complete, implying a meaningful role for the government provision of insurance.

References

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

[2] Atkeson, Andrew and Lucas, Robert, E. Jr. (1992), “On Efficient Distribution With Private Information”, Review of Economic Studies, 59, issue 3, p. 427-453.

[3] Bianchi, Francesco, Cosmin Ilut, and Martin Schneider, “Uncertainty shocks, asset supply and pricing over the business cycle”, forthcoming Review of Economic Studies, March 2017.

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

[5] Cerreia-Vioglio, S., Maccheroni, F., Marinacci, M., Montrucchio, L. (2011). “Uncertainty Averse Preferences”. Journal of Economic Theory 146, 1275-1330.

[6] de Castro, Luciano I., and Nicholas C. Yannelis, “Uncertainty, efficiency and incentive compatibility,” Journal of Economic Theory, forthcoming, March 2018.

[7] Epstein, Larry, and Martin Schneider, “Recursive multiple priors”, Journal of Economic Theory, 113(1), 32-50 (2003)

[8] Golosov, Mikhail, and Aleh Tsyvinski, “Policy Implications of Dynamic Public Finance”, Annual Review of Economics (2015) 7:1, 147-171.

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

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

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

[12] Hansen, Lars P. and Thomas J. Sargent, 2017, “The Price of Macroeconomic Uncertainty with Tenuous Beliefs”, November, 2017.

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

[14] Ilut, Cosmin, Matthias Kehrig, and Martin Schneider, “Slow to hire, quick to fire: employment dynamics with asymmetric responses to news”, forthcoming Journal of Political Economy, April 2017.

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

[16] Maccheroni, F., Marinacci, M., and Rustichini, A. (2006b). “Dynamic Variational Preferences”. Journal of Economic Theory 128, 4-44.

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

[18] Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, "Distributional National Accounts: Methods and Estimates for the United States", forthcoming Quarterly Journal of Economics, September 2017

[19] Prescott, Edward C. and Robert M. Townsend, “Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard”. Econometrica 52 (1):21-45.

[20] Rees-Jones, Alex, and Dmitry Taubinsky, “Taxing Humans: Pitfalls of the Mechanism Design Approach and Potential Resolutions,” NBER Working Paper, No. 23980, October 2017.

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

Appendix A: Proofs omitted in the main text

Proofs for Section 2

Proof. **(Proposition 1)** We begin by defining the date 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 effectively a two period allocation with an endogenous outside option. 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 each possible $t \geq 1$ type distribution $\varphi_t \in \Delta\left(s^t\right)$, let $\underline{\varphi}_t$ be the distribution defined as follows: For any $\underline{\varphi}_t$-measurable set E, let

$$\underline{\varphi}_t\left(E\right) \equiv \varphi_t\left(\left\{ s^t \middle| \left(\theta_0, \Pi_1, \theta_1, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right) \in E\right\}\right).$$

Hence $\underline{\varphi}_t$ is formed from the distribution φ_t by shifting probabilities so that all agents realize belief set $\underline{\Pi}_2$ at date 2 and shocks $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$ at $\tau = 2, \ldots, t$. Distributions of the form $\underline{\varphi}_t$ are the "worst case" t type distributions, from the perspective of an agent at date 0.

For $t \geq 1$, define

$$c_t^0\left(s^t, \varphi_t\right) \equiv c_t^*\left(\left(\theta_0, \Pi_1, \theta_1, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right), \underline{\varphi}_t\right),$$
$$z_t^0\left(s^t, \varphi_t\right) \equiv z_t^*\left(\left(\theta_0, \Pi_1, \theta_1, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right), \underline{\varphi}_t\right),$$
$$k_{t+1}^0\left(s^t, \varphi_t\right) \equiv k_{t+1}^*\left(\left(\theta_0, \Pi_1, \theta_1, \underline{\Pi}_2, \left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)_{\tau=2}^t\right), \underline{\varphi}_t\right).$$

At $t \geq 1$, regardless of the actual distribution of $t = 2$ beliefs Π_2 and $\tau \geq 2$ shocks s_τ, C^0 allocates consumption, effective labor, and capital as if every agent realized the beliefs $\underline{\Pi}_2$ and date $\tau \geq 2$ shocks $\left(\underline{\theta}, \underline{\Pi}_{\tau+1}\right)$.

To see that the agents are indifferent between C^0 and C^*, fix an initial $s_0 = (\theta_0, \Pi_1)$ and let $\Pi_1^- \subseteq \Pi_1$ be all distributions π_1^- of the form $\underline{\varphi}_1$ described above, i.e., beliefs supported on the set $\left\{s^1 \middle| \Pi_2 = \underline{\Pi}_2\right\}$. Since $\Pi_1 \neq \emptyset$, Assumption 1 implies that Π_1^- is non-empty. We have

$$
\begin{aligned}
U_0\left(C^*, s_0, \varphi_0\right) &= u\left(c_0^*\left(s_0, \varphi_0\right), \frac{z_0^*\left(s_0, \varphi_0\right)}{\theta_0}\right) + \beta \inf_{\Pi_1} \mathbb{E}_{\pi_1}\left[U_1\left(C^*, s^1, \pi_1\right) \middle| s_0\right] \\
&\leq u\left(c_0^*\left(s_0, \varphi_0\right), \frac{z_0^*\left(s_0, \varphi_0\right)}{\theta_0}\right) + \beta \inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-}\left[U_1\left(C^*, s^1, \pi_1^-\right) \middle| s_0\right] \\
&= u\left(c_0^0\left(s_0, \varphi_0\right), \frac{z_0^0\left(s_0, \varphi_0\right)}{\theta_0}\right) + \beta \inf_{\Pi_1^-} \mathbb{E}_{\pi_1^-}\left[U_1\left(C^0, s^1, \pi_1^-\right) \middle| s_0\right] \\
&= u\left(c_0^0\left(s_0, \varphi_0\right), \frac{z_0^0\left(s_0, \varphi_0\right)}{\theta_0}\right) + \beta \inf_{\Pi_1} \mathbb{E}_{\pi_1}\left[U_1\left(C^0, s^1, \pi_1\right) \middle| s_0\right] \\
&= U_0\left(C^0, s_0, \varphi_0\right).
\end{aligned}
$$

The key step is the fourth line, which makes substantial use of Assumption 1: By definition, the date $t \geq 1$ functions $\{c_1^0, z_1^0, k_2^0, \ldots\}$ in C^0 depend only on an agent's date 0 type s_0 and date 1 skill θ_1, as well as the distribution of (s_0, θ_1) in the economy. Hence

$$
\mathbb{E}_{\pi_{t+1}^-}\left[U_1\left(C^0, s^1, \pi_{t+1}^-\right) \middle| s_0\right]
$$

depends on π_1^- only through its marginal distribution over (s_0, θ_1). But by Assumption 1, $\pi_1^- \in \Pi_1^-$ if and only if there exists $\pi_1 \in \Pi_1$ such that π_1^- and π_1 have the same marginal distribution over (s_0, θ_1), so the fourth line follows.

Since C^* is a solution to the government's problem, the above inequality must be an equality for all s_0. 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 ex post if the $t = 1$ distribution φ_1 is not of the form $\underline{\varphi}_1$. The government can only do this if the updated allocation delivers weakly greater $t = 1$ continuation utility to all agents than C^0. In this sense, the $t = 1$ continuation allocation $\{c_1^0, z_1^0, k_2^0, \ldots\}$ gives the agents an endogenous outside option that they may take if the government cannot design a weakly preferable allocation. For this reason and the fact that C^0 is completely uncontingent on t shocks for $t \geq 2$, we say that C^0 is a two-period allocation with an endogenous outside option. Note that we construct the outside option based off of what C^* prescribes when every agent realizes the worst productivity shock at each $t \geq 2$, so we are assured that the outside option is feasible in the case that a type distribution of the form $\underline{\varphi}_t$ is not realized at some $t \geq 2$.

In general, given a two-period allocation C^{t-1} designed at $t-1$ and a t type distribution φ_t, the government with social welfare weights η designs the optimal two-period allocation

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

$$\max_{C^t} \int U_t\left(C^t, s^t, \varphi_t\right) d\eta\left(s_0\right)$$

subject to non-negativity and

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

$$U_t\left(C^t, s^t, \varphi_t\right) \geq U_t\left(C^{t-1}, s^t, \varphi_t\right)$$

Here the feasibility constraint holds for $\tau = t, \ldots, T$ and $\forall \varphi_\tau \geq \varphi_t$, and the self-enforcement constraint holds $\forall s^t$. This problem also applies when $t = 0$, but there is no self-enforcement constraint. By solving the problem at $t = 0, \ldots, T$, we can find a sequence of two-period allocations $\{C^{*t}\}_{t=0}^T$ that implements the efficient allocation C^*. □

Proofs for Section 3

Proof. **(Lemma 2)** For each t, each s^t, and each $\varphi_{t,t+1} \in \Delta\left(s^{t,t+1}\right)$, define

$$\tilde{c}_t\left(s^t, \varphi_{t,t+1}\right) \equiv c_t\left(\sigma_t^e\left(s^t, \varphi_{t,t}\right), \varphi_{t,t+1}\right).$$

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

$$U_t\left(\tilde{C}, \sigma^*, s^t, \varphi_{t,t}\right) = U_t\left(C, \sigma^e, s^t, \varphi_{t,t}\right)$$

$\forall s^t$ and $\forall \varphi_{t,t} \in \Delta\left(s^{t,t}\right)$. Suppose that σ^* is not an equilibrium of $\tilde{C}$. Then there exist $\sigma \in \Sigma$, s^t, and $\varphi_{t,t}$ such that

$$U_t\left(\tilde{C}, \sigma, s^t, \varphi_{t,t}\right) > U_t\left(\tilde{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

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

$$U_t\left(C, \sigma \circ \sigma^e, s^t, \varphi_{t,t}\right) > U_t\left(C, \sigma^e, s^t, \varphi_{t,t}\right),$$

where $\sigma \circ \sigma^e \equiv \{\sigma_t \circ \sigma_t^e\}_{t=0}^T$. This contradicts the assumption that σ^e is an equilibrium of C, so σ^* must be an equilibrium of $\tilde{C}$. □

Proof. **(Proposition 3)** Let C be a renegotiation-proof allocation, and suppose that it is not incentive-compatible. Then there exist a date τ, a history $\tilde{s}^\tau$, a distribution $\tilde{\varphi}_{\tau,\tau}$, and a strategy σ such that

$$U_\tau\left(C, \sigma^*, \tilde{s}^\tau, \tilde{\varphi}_{\tau,\tau}\right) < U_\tau\left(C, \sigma, \tilde{s}^\tau, \tilde{\varphi}_{\tau,\tau}\right).$$

Without loss of generality, we can assume that τ is such that

$$\sigma_t = \sigma_t^* \quad \forall t < \tau.$$

We can also take σ to be an optimal deviation in the sense that

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(C, \sigma, s^t, \varphi_{t,t}\right)$$

for all $t = 0, \ldots, T$, all s^t, and all $\varphi_{t,t}$.

For any $t \geq \tau$ and any $\varphi_{t,t+1} \in \Delta\left(s^{t,t+1}\right)$, let $\bar{\varphi}_{t,t+1}$ be the distribution defined as follows: For any $\varphi_{t,t+1}$-measurable set E, let

$$\begin{aligned}\bar{\varphi}_{t,t+1}(E) = \varphi_{t,t+1}&\left(\left\{ s^{t,t+1} \middle| \, s^{t,t+1} \in E,\ s^{t,t+1} \ngeq \tilde{s}^\tau \right\}\right) \\ &+ \varphi_{t,t+1}\left(\left\{ s^t \middle| \, s^t \geq \tilde{s}^\tau,\ \sigma_t\left(s^t, \varphi_{t,t}\right) \in E\right\}\right).\end{aligned}$$

The first term measures all types $s^{t,t+1}$ (truthfully) reported by agents who do not follow the history $\tilde{s}^\tau$, and the second term measures the deviating type $s^t \geq \tilde{s}^\tau$. Thus $\bar{\varphi}_{t,t+1}$ is the distribution such that $\bar{\varphi}_{t,t} = \varphi_{t,t}$, but $\bar{\varphi}_{t,t+1}$ reflects the fact that agents with history $s^t \geq \tilde{s}^\tau$ report $\sigma_t\left(s^t, \varphi_{t,t}\right)$ instead of s^t.[24] For each $t = 0, \ldots, T$, define

$$\bar{c}_t\left(s^t, \varphi_{t,t+1}\right) = \begin{cases} c_t\left(\sigma_t\left(s^t, \varphi_{t,t}\right), \bar{\varphi}_{t,t+1}\right) & \text{if } t \geq \tau, s^t \geq \tilde{s}^\tau, \varphi_{t,t+1} \geq \tilde{\varphi}_{\tau,\tau}, \\ c_t\left(s^t, \bar{\varphi}_{t,t+1}\right) & \text{if } t \geq \tau, s^t \ngeq \tilde{s}^\tau, \varphi_{t,t+1} \geq \tilde{\varphi}_{\tau,\tau}, \\ c_t\left(s^t, \varphi_{t,t+1}\right) & \text{else.} \end{cases}$$

Define $\bar{z}_t$ and $\bar{k}_{t+1}$ similarly, and set $\bar{C} \equiv \left\{\bar{c}_t, \bar{z}_t, \bar{k}_t\right\}_{t=0}^T$. Now $\bar{C}$ is feasible because C must be feasible at any date t when the distribution of reported states is given by $\bar{\varphi}_t$. By construction, we also have

$$U_\tau\left(C, \sigma^*, \tilde{s}^\tau, \tilde{\varphi}_{\tau,\tau}\right) < U_\tau\left(\bar{C}, \sigma^*, \tilde{s}^\tau, \tilde{\varphi}_{\tau,\tau}\right).$$

We claim that for any $t = 0, \ldots, T$, any s^t and any $\varphi_{t,t}$, we have

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

This would imply that C is not renegotiation-proof, the desired contradiction.

[24]This definition implies that the government treats agents' types as names, where the actual measure of agents of a particular name is known at each date. This can be made consistent with the information structure by assuming a finite number of agents in the economy.

Fix any $t = \tau, \ldots, T$, any s^t, and any $\varphi_{t,t}$. If $\varphi_{t,t} \ngeq \tilde{\varphi}_{\tau,\tau}$ we have

$$\{c_r, z_r\}_{r=t}^{T} = \{\bar{c}_r, \bar{z}_r\}_{r=t}^{T},$$

and this obviously implies

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) = U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

Now suppose $\varphi_{t,t} \geq \tilde{\varphi}_{\tau,\tau}$ and $s^t \geq \tilde{s}^\tau$. Then

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(C, \sigma, s^t, \varphi_{t,t}\right) = U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

For the case in which $\varphi_{t,t} \geq \tilde{\varphi}_{\tau,\tau}$ but $s^t \ngeq \tilde{s}^\tau$, first let $\varphi_{t+1} \in \Delta\left(s^{t+1}, \hat{\Pi}^{t+1}\right)$ denote any distribution over $\left(s^{t+1}, \hat{\Pi}^{t+1}\right)$, and write $\varphi_{t+1} \sim \varphi'_{t+1}$ if φ_{t+1} and φ'_{t+1} have the same marginal distribution over s^{t+1}. Write $\varphi_{t+1} \sim \Pi_{t+1}$ if there exists $\pi_{t+1} \in \Pi_{t+1}$ such that $\varphi_{t+1} \sim \Pi_{t+1}$. Define the set

$$\mathcal{W}_t\left(C, s^t, \varphi_{t,t}\right) \equiv \left\{u\left(c_t\left(s^t, \varphi_{t,t+1}\right), \frac{z_t\left(s^t, \varphi_{t,t+1}\right)}{\theta_t}\right) + \beta \mathbb{E}_{\varphi_{t+1}}\left[U_{t+1}\left(C, \sigma^*, s^{t+1}, \varphi_{t+1,t+1}\right) \middle| s^t\right] \right.$$
$$\left. | \varphi_{t+1} \geq \varphi_{t,t}, \varphi_{t+1} \sim \Pi_{t+1}\right\}$$

$\mathcal{W}_t\left(C, s^t, \varphi_{t,t}\right)$ is the set of all possible t utility levels that could be realized by the agent for some $\varphi_{t+1} \geq \varphi_{t,t}$ with $\varphi_{t+1} \sim \Pi_{t+1}$. By the definition of $\bar{C}$,

$$\begin{aligned}
\mathcal{W}_t\left(\bar{C}, s^t, \varphi_{t,t}\right) &= \left\{u\left(\bar{c}_t\left(s^t, \varphi_{t,t+1}\right), \frac{\bar{z}_t\left(s^t, \varphi_{t,t+1}\right)}{\theta_t}\right) + \beta \mathbb{E}_{\varphi_{t+1}}\left[U_{t+1}\left(\bar{C}, \sigma^*, s^{t+1}, \varphi_{t+1,t+1}\right) \middle| s^t\right] \right. \\
&\qquad \left. | \varphi_{t+1} \geq \varphi_{t,t}, \varphi_{t+1} \sim \Pi_{t+1}\right\} \\
&= \left\{u\left(c_t\left(s^t, \bar{\varphi}_{t,t+1}\right), \frac{z_t\left(s^t, \bar{\varphi}_{t,t+1}\right)}{\theta_t}\right) + \beta \mathbb{E}_{\bar{\varphi}_{t+1}}\left[U_{t+1}\left(C, \sigma^*, s^{t+1}, \bar{\varphi}_{t+1,t+1}\right) \middle| s^t\right] \right. \\
&\qquad \left. | \varphi_{t+1} \geq \varphi_{t,t}, \varphi_{t+1} \sim \Pi_{t+1}\right\} \\
&\subseteq \left\{u\left(c_t\left(s^t, \varphi_{t,t+1}\right), \frac{z_t\left(s^t, \varphi_{t,t+1}\right)}{\theta_t}\right) + \beta \mathbb{E}_{\varphi_{t+1}}\left[U_{t+1}\left(C, \sigma^*, s^{t+1}, \varphi_{t+1,t+1}\right) \middle| s^t\right] \right. \\
&\qquad \left. | \varphi_{t+1} \geq \varphi_{t,t}, \varphi_{t+1} \sim \Pi_{t+1}\right\} \\
&= \mathcal{W}_t\left(C, s^t, \varphi_{t,t}\right)
\end{aligned}$$

The set inclusion holds by the following argument: Assumption 2 implies that Π_{t+1} contains all distributions $\varphi_{t+1} \sim \Pi_{t+1}$ with $\varphi_{t+1} \geq \varphi_{t,t}$, and it is obvious from the definition of

$\bar{\varphi}_{t+1}$ that $\bar{\varphi}_{t+1} \sim \varphi_{t+1}$. Thus $\bar{\varphi}_{t+1} \sim \Pi_{t+1}$, but then the containment is immediate because distributions of the form $\bar{\varphi}_{t+1}$ form a subset of all distributions $\varphi_{t+1} \sim \Pi_{t+1}$ with $\varphi_{t+1} \geq \varphi_{t,t}$. Again by Assumption 2, we find

$$\begin{aligned} U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) &= \inf \mathcal{W}_t\left(C, s^t, \varphi_{t,t}\right) \\ &\leq \inf \mathcal{W}_t\left(\bar{C}, s^t, \varphi_{t,t}\right) \\ &= U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right). \end{aligned}$$

Hence for any s^t and any $\varphi_{t,t}$,

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

For $t = 0, \ldots, \tau - 1$, it follows immediately from the above argument and the fact that

$$\{c_r, z_r\}_{r=0}^{\tau-1} = \{\bar{c}_r, \bar{z}_r\}_{r=0}^{\tau-1},$$

that for any s^t and any $\varphi_{t,t}$,

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

The above analysis shows that for any $t = 0, \ldots, T$, any s^t and any $\varphi_{t,t}$,

$$U_t\left(C, \sigma^*, s^t, \varphi_{t,t}\right) \leq U_t\left(\bar{C}, \sigma^*, s^t, \varphi_{t,t}\right).$$

Moreover, the inequality is strict when $t = \tau$, $s^t = \tilde{s}^\tau$, and $\varphi_{t,t} = \tilde{\varphi}_{\tau,\tau}$. This implies that the original allocation C is not renegotiation-proof, a contradiction. □

Proofs for Section 5

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

$$\begin{aligned} U_{i,0}\left(\tilde{C}_i^0, h_0\right) = u\left(\tilde{c}_{i,0}\left(h_0\right), \frac{\tilde{z}_{i,0}\left(h_0\right)}{\theta_{i,0}}\right) \\ + \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\Bigg[\beta u\left(\tilde{c}_{i,1}\left(h_0, \theta_1\right), \frac{\tilde{z}_{i,1}\left(h_0, \theta_1\right)}{\theta_{i,1}}\right) \\ + \sum_{t=2}^{T} \beta^t u\left(\tilde{c}_{i,t}\left(h_0, \theta_1\right), \frac{\tilde{z}_{i,t}\left(h_0, \theta_1\right)}{\underline{\theta}}\right) \Bigg| h_0\Bigg] \end{aligned}$$

For notational simplicity, define

$$\tilde{v}\left(h_0,\theta_1\right) \equiv \beta u\left(\tilde{c}_{i,1}\left(h_0,\theta_1\right), \frac{\tilde{z}_{i,1}\left(h_0,\theta_1\right)}{\theta_{i,1}}\right) + \sum_{t=2}^{T} \beta^t u\left(\tilde{c}_{i,t}\left(h_0,\theta_1\right), \frac{\tilde{z}_{i,t}\left(h_0,\theta_1\right)}{\underline{\theta}}\right).$$

Define $\hat{v}\left(h_0,\theta_1\right)$ 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 h_0, but $\left.\pi'_{i,1}\left(\cdot|\, h_0\right)\right|_{\theta_i}$ places weight only on $\underline{\theta}$ and $\bar{\theta}$ so as to satisfy

$$\mathbb{E}_{\pi_{i,1}}\left[\theta_{i,1}|\, h_0\right] = \bar{\theta}_{i|i,1}\left(h_0\right).$$

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:

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right] \leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right],$$
$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\hat{v}\left(h_0,\theta_1\right)|\, h_0\right] \leq \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}\left(h_0,\theta_1\right)|\, h_0\right].$$

To see that the opposite inequalities also hold, let $p\left(h_0\right) \in [0,1]$ be such that

$$p\left(h_0\right)\underline{\theta} + \left(1 - p\left(h_0\right)\right)\bar{\theta} = \bar{\theta}_{i|i,1}\left(h_0\right).$$

Now by the definition of $\tilde{C}_i^0$, $\tilde{v}\left(h_0,\theta_{-i,1},\cdot\right)$ 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

$$\begin{aligned}
\mathbb{E}_{\pi_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right] &= \mathbb{E}_{\pi_{i,1}}\left[\mathbb{E}_{\pi_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0,\theta_{-j,1}\right]\Big|\, h_0\right] \\
&\geq \mathbb{E}_{\pi_{i,1}}\left[p\left(h_0\right)\tilde{v}\left(h_0,\theta_{-i,1},\underline{\theta}\right)\right. \\
&\qquad \left. + \left(1-p\left(h_0\right)\right)\tilde{v}\left(h_0,\theta_{-i,1},\bar{\theta}\right)\Big|\, h_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[p\left(h_0\right)\tilde{v}\left(h_0,\theta_{-i,1},\underline{\theta}\right)\right. \\
&\qquad \left. + \left(1-p\left(h_0\right)\right)\tilde{v}\left(h_0,\theta_{-i,1},\bar{\theta}\right)\Big|\, h_0\right] \\
&= \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right].
\end{aligned}$$

A similar set of calculations applies to $\hat{v}\left(h_0,\theta_1\right)$, so we must have

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right] = \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}\left(h_0,\theta_1\right)|\, h_0\right],$$
$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\hat{v}\left(h_0,\theta_1\right)|\, h_0\right] = \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}\left(h_0,\theta_1\right)|\, h_0\right].$$

Now since $\tilde{v}(h_0, \theta_{-i,1}, \cdot)$ and $\hat{v}(h_0, \theta_{-i,1}, \cdot)$ coincide on the endpoints of Θ, the independence of $\theta_{i,1}$ and $\theta_{-i,1}$ imply

$$\begin{aligned}\mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}(h_0,\theta_1)|\, h_0\right] &= \mathbb{E}_{\pi'_{i,1}}\left[p(h_0)\,\tilde{v}(h_0,\theta_{-i,1},\underline{\theta})\right.\\ &\qquad\left.+(1-p(h_0))\,\tilde{v}\left(h_0,\theta_{-i,1},\bar{\theta}\right)\big|\, h_0\right]\\ &= \mathbb{E}_{\pi'_{i,1}}\left[p(h_0)\,\hat{v}(h_0,\theta_{-i,1},\underline{\theta})\right.\\ &\qquad\left.+(1-p(h_0))\,\hat{v}\left(h_0,\theta_{-i,1},\bar{\theta}\right)\big|\, h_0\right]\\ &= \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}(h_0,\theta_1)|\, h_0\right].\end{aligned}$$

Hence

$$\inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\tilde{v}(h_0,\theta_1)|\, h_0\right] = \inf_{\Pi'_{i,1}} \mathbb{E}_{\pi'_{i,1}}\left[\hat{v}(h_0,\theta_1)|\, h_0\right].$$

With the above equations, we can conclude

$$\inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\tilde{v}(h_0,\theta_1)|\, h_0\right] = \inf_{\Pi_{i,1}} \mathbb{E}_{\pi_{i,1}}\left[\hat{v}(h_0,\theta_1)|\, h_0\right].$$

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

Appendix B: Infinite time horizon

This section describes a possible extension of the model in Section 2 to the $T = \infty$ case. The model formulation will be sequential, and we assume a finite number of agents $i \in \{1, \ldots, N\}$ for simplicity. Let $(\Omega, \mathcal{F})$ be a measurable space, let $\Theta \subset \mathbb{R} - \{0\}$ be compact, and let $\theta_i^\infty : \Omega \to \Theta^{\mathbb{N}}$ be $\mathcal{F}$-measurable, where $\Theta^{\mathbb{R}}$ is endowed with the sigma algebra generated by finite rectangles. (For concreteness, we can take $\mathcal{F} := \sigma(\theta_i^\infty, i \in \{1, \ldots, N\})$.) Here θ_i^∞ represents an infinite sequence of skill shocks for a single agent. Define θ_i^t as the obvious projection of θ_i^∞ to Θ^t, and let $\mathcal{F}_t := \sigma(\theta^t)$, where $\theta^t := (\theta_1^t, \ldots, \theta_N^t)$ denotes the t histories of skill shocks for every agent in the economy. Then $(\mathcal{F}_t)_{t=0}^\infty$ is a filtration on Ω to which the shock history process $(\theta^t)_{t=0}^\infty$ is adapted.

Let $C := \{c_t(\theta^t), z_t(\theta^t), k_{t+1}(\theta^t)\}_{t=0}^\infty$ 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

$$\sum_{i=1}^N \left[c_{i,t}\left(\theta^t\right) + k_{i,t+1}\left(\theta^t\right)\right] \le f\left(\sum_{i=1}^N k_{i,t}\left(\theta^{t-1}\right), \sum_{i=1}^N z_{i,t}\left(\theta^t\right)\right) \quad \forall t, \forall \theta^t.$$

Let $P_i \subset \Delta(\Omega, \mathcal{F})$ be a non-empty set of prior distributions on $(\Omega, \mathcal{F})$ that represent agent i's beliefs. Assume that agents update their beliefs using Bayes's Theorem prior-by-prior. Given an allocation C, an agent's t continuation utility is given by

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

where $\beta \in [0,1)$. With social welfare weights $\eta \in \mathbb{R}_+^N$ and an initial state θ_0, an *efficient* allocation $C^*(\theta_0)$ is given by

$$C^*(\theta_0) \in \arg\max_C \sum_{i=1}^{N} U_{i,0}(C, \theta_0)\, \eta_i$$

subject to feasibility and non-negativity.

The analog of Assumption 1 in this setup is

Assumption 4. *$\forall t$, $\forall i$, $\forall \theta^t$, and $\forall p_i \in P_i$, $\exists p_i^- \in P_i$ such that*

$$p_i^-\left(\theta^{t+1}\middle|\, \theta^t\right) = p_i\left(\theta^{t+1}\middle|\, \theta^t\right) \quad \forall \theta^{t+1},$$

but

$$p_i^-\left(\cap_{j=1}^N \left(\theta_j^\infty\right)^{-1}\left(\left\{\ell \in \Theta^{\mathbb{N}} : \ell_\tau = \underline{\theta},\ \tau \geq t+2\right\}\right)\right) = 1.$$

With agents' t utilities defined as above, a version of Proposition 1 holds:

Proposition 7. *The efficient allocation C^* can be implemented by a sequence of two-period allocations $\{C^t\}_{t=0}^\infty$, where each C^t specifies an endogenous outside option.*

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

$$c_t^0\left(\theta^t\right) := c_t^*\left(\theta^1, \left(\underline{\theta}, \overset{N \text{ times}}{\cdots}, \underline{\theta}\right)^{t-2}\right).$$

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

$\forall i,$

$$\begin{aligned}U_{i,0}\left(C^{*},\theta^{0}\right) &= \inf_{p_i\in P_i}\mathbb{E}_{p_i}\left[\sum_{t=0}^{\infty}\beta^{t}u\left(c_t^{*}\left(\theta^{t}\right),\frac{z_t^{*}\left(\theta^{t}\right)}{\theta_{i,t}}\right)\middle|\,\theta^{0}\right]\\ &\leq \inf_{p_i^{-}\in P_i^{-}}\mathbb{E}_{p_i^{-}}\left[\sum_{t=0}^{\infty}\beta^{t}u\left(c_t^{*}\left(\theta^{t}\right),\frac{z_t^{*}\left(\theta^{t}\right)}{\theta_{i,t}}\right)\middle|\,\theta^{0}\right]\\ &= \inf_{p_i^{-}\in P_i^{-}}\mathbb{E}_{p_i^{-}}\left[\sum_{t=0}^{\infty}\beta^{t}u\left(c_t^{0}\left(\theta^{t}\right),\frac{z_t^{0}\left(\theta^{t}\right)}{\theta_{i,t}}\right)\middle|\,\theta^{0}\right]\\ &= \inf_{p_i\in P_i}\mathbb{E}_{p_i}\left[\sum_{t=0}^{\infty}\beta^{t}u\left(c_t^{0}\left(\theta^{t}\right),\frac{z_t^{0}\left(\theta^{t}\right)}{\theta_{i,t}}\right)\middle|\,\theta^{0}\right].\end{aligned}$$

The last line holds because c_t^0 and z_t^0 only depend on θ^1 for $t \geq 2$. Hence all agents weakly prefer C^0 to C^*, and the feasibility of C^0 follows from that of C^*. By iterating this process at each date, we have the result. □

With Proposition 7, we find that even when the time horizon is infinite, the efficient allocation can be implemented by a sequence of simplified, two-period allocations $\{C^t\}_{t=0}^{\infty}$ that are reformed after each date. As in Section 2.2, each two-period allocation C^t displays limited state dependence at $\tau \geq t+2$, and an argument analogous to that in Section 2.3 implies that if the government constructs these allocations via backward induction, then they are often history independent. This is true in particular when agents' beliefs are homogeneous.