

Institution Building without Commitment*

Marco Bassetto
Federal Reserve Bank of Minneapolis

Zhen Huo
Yale University

José-Víctor Ríos-Rull
University of Pennsylvania
CAERP, CEPR, NBER, UCL

Thursday 4th January, 2024

Abstract

We propose a theory for the slow implementation of “good” policies, as carbon taxes or free trade. The resulting policies are a way out of the generally “poor” (i.e. Markovian) policies that arise in environments with time inconsistency. The policies arise as the equilibrium of environments where agents are allowed both to wait for future agents to start the equilibrium and to restart the equilibrium by ignoring the history. The equilibrium allocation gradually transits towards one that weighs both short- and long-term concerns, stopping short of the Ramsey outcome. We document the properties of the implied equilibrium concept and we use it to account for the slow emergence of climate policies and for the gradual reduction of global tariff rates.

*For useful comments, we thank Gadi Barlevy, Jess Benhabib, Anmol Bhandari, V.V. Chari, Harold Cole, Meredith Crowley, Mariacristina De Nardi, Per Krusell, Nicola Pavoni, Facundo Piguillem, Larry Samuelson, Tom Sargent, Aleh Tsyvinski, Pierre Yared and participants at several conferences and seminars. An older version of this paper circulated under the title “Organizational Equilibrium with Capital”. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

Europe will not be made all at once, or according to a single plan. It will be built through concrete achievements which first create a de facto solidarity.

Robert Schumann, May 9th 1950

1 Introduction

Why are many obvious arguably “good” policies (against climate change, to expand free trade, or even the extension of civil rights or the expansion of the European Union) adopted slowly? We claim that it is often due to an incentive of the current policy makers to delegate to future policy makers the implementation because of the initial costs involved in carrying out a fast adoption.¹ To make the case, we postulate environments where policy makers display time inconsistency, either due to having a shorter horizon than the population at large, or to the presence of an incentive problem. We propose an equilibrium concept—organizational equilibrium—where decision makers are allowed to wait and avoid addressing the issue as well as to replicate the behavior of previous ones. As a result, good institutions have to arise slowly but eventually converge to reasonable outcomes—the best time-invariant policy from the point of view of any agent. We apply our ideas to the setting of carbon taxes and to the expanding of free trade, but it is much more general, embracing many environments including those that display preferences with heterogeneous discounting ([Jackson and Yariv, 2014](#)).

The gist of our approach is to model explicitly aggregate economies where policy makers have to address an issue where short term costs are at odds with long term benefits and a time inconsistency problems arise. A fast implementation of good policies with delayed benefits gives incentives to wait, but agreement over what good policies are in the distant future opens the door for their gradual implementation. We then use organizational equilibrium to analyze the policy outcomes. The implied allocations are vastly superior to those predicted by Markov equilibria, yet they do not need to be supported by trigger-strategy reversion to dominated outcomes. An attractive feature of our equilibrium is that it involves a gradual evolution of policy and institutions. We show how the equilibrium is simple to study: although we provide game theoretic foundations, it can be computed recursively by looking just at the equilibrium path, as is the case in competitive equilibria that are the workhorse of macroeconomic models. This property makes it much easier to apply to specific macroeconomic problems, particularly where the blend of strategic and competitive elements would

¹As an example, though China started to take an active role in carbon reduction more recently, the government was initially reluctant to implement a stringent environmental policy. The former president Hu Jintao stated that “China is a developing country in the process of industrialization and modernization... China’s central task now is to develop the economy and make life better for the people.” at 2008 G8 Outreach Session.

otherwise require an intricate description of the agents' strategies.² It is also easily amenable to comparative statics analysis.

The substantive concern that we address in this paper is the slowness in the implementation of policies that are generally considered desirable. We highlight two such sets of policies that are central issues of our time with relative agreement among experts: that carbon emissions should be reduced, and reduced fast, to mitigate global warming and that the reduction of trade barriers yield more efficient allocations. Yet governments are not embracing these policies, at least to the degree that is typically desired, but they are timidly moving in the right direction: there is a gradual increase of the carbon tax³ and it has taken several decades to reduce the global tariff rate to its current level. While there can be other reasons for this slowness, we do believe that a satisfactory explanation has to come from models that are explicit about the intertemporal conflicts associated with those policies.

The discussion around policies on climate change has often been cast in terms of fairness and institutional agreements across generations that support better policies. Taking this concern to heart, we extend the climate change model in [Goloso et al. \(2014\)](#) to accommodate possible conflicts of interest between current and future generations: current generations bear the cost of reduced carbon emissions and future generations benefit from a cooler planet. Organizational equilibrium resolves this tension by a graduate increase in carbon taxation as in the data while neither the Markov equilibrium nor the Ramsey outcome yields predictions that are consistent with the observed qualitative features. In similar fashion, we see the implementation of free trade as a conflict between the import substitute producers and long-run growth ([Autor et al., 2013](#); [Acemoglu et al., 2016](#); [Pierce and Schott, 2016](#); [Caliendo et al., 2019](#)). Accordingly, we pose a two-country model with international trade where the optimal tariff rate strikes a balance between the short-run redistribution and the long-run growth. A lower tariff rate improves allocation efficiency and facilitates long-run growth, but it hurts the short-run welfare of workers in the importing sector. Therefore, the current government always prefers future governments to implement the reduction of tariff rates. Again, organizational equilibrium resolves the tension with a graduate policy change, reducing tariffs, which is consistent with the actual globalization process.

The theoretical foundation of organizational equilibrium is based on the logic that there is nothing special in the initial period (as Ramsey type solutions assume), yet some intertemporal collaboration is possible (which Markov solutions neglect). Specifically, we argue that equilibria should satisfy three conditions in environments with a sequence of decision makers that see themselves in a similar spot

²[Bassetto \(2005\)](#) describes some of the challenges that such hybrid environments entail.

³Based on the World Bank Carbon Pricing Dashboard, we document that (1) the carbon taxes rise gradually over the last 30 years in Scandinavian countries that first initiated the carbon tax system; (2) across different regions, the level of current carbon tax is increasing in the number of years since the carbon tax implementation, controlling the level of per capita GDP.

—a form of stationarity even if there are state variables. The first such condition, a *no-restarting condition*, is that any outcome should have the property that no decision maker would rather become an earlier member of the decision-making sequence. This limits the use of trigger strategies as a future punishment. A second condition, a *no-delay condition*, prevents free riding at the start of the process: no agent can do better by sitting out the system (playing Markov) and waiting for future agents to start on a given equilibrium path and eliminates the intrinsic advantage that is often associated with the first agent. This condition prevents jumping to desirable allocations fast. We interpret the implications of this condition as the need for institutions to *slowly earn good will*, like earning a reputation for good behavior without need of unobserved types or triggers. Finally, the third condition is an *an optimality requirement* within the class of allocations that satisfy the previous two requirements. Note that these conditions are defined over properties of the equilibrium path which drastically simplifies the analysis of any application.

Environments with time consistency issues (hyperbolic discounting consumers but also carbon tax and tariff rates policy settings) have typically been modeled as a specific game that has a sequence of decision makers sometimes described as the *future selves* or *future governments*. In these games the time zero agent has a special position. We pose our equilibrium concept as a particular refinement of the set of subgame perfect Nash equilibria of this game that tries to deemphasize the preeminent role of the time zero agent. But we also pose an alternative game to model the same class of environments, where any agent has the ability to hide the past history. This alternative game eliminates the specificity of the time zero agent and conveys a recursive spirit to the passage of time. We show how the no-delay condition becomes a necessary condition for any symmetric, Pareto optimal subgame perfect equilibrium of this game.

Under mild conditions, we prove the existence of an organizational equilibrium. The equilibrium converges to a stationary allocation that we refer to as a steady state.⁴ From this steady state, the entire transition path then can be solved recursively.⁵ Crucially, due to the no-delay condition, agents' actions converge only gradually. The stationary allocation to which our equilibrium converges weighs the concerns over immediate events associated to time inconsistent environments with those later events. It has a larger weight into the future than the allocation of the Markov equilibrium but not so much as that implied by the Ramsey solution that completely ignores any short term considerations. It is also easy to calculate and characterize. In fact it is the best *constant* action

⁴We call the state vector of the long-run stationary allocation a steady state, but note that it does not have the property that, if the economy starts there, it will remain there for ever, as sometimes steady state is understood. This is because the equilibrium sequence of actions converges to the constant action that supports this steady state only asymptotically.

⁵On the technical side, while the approach in [Abreu et al. \(1986, 1990\)](#) cannot be adopted to construct the set of continuation values to compute organizational equilibria or reconsideration-proof equilibria, our method exploits some similar ideas in its quest for a recursive representation, which we think may have independent value.

from the point of view of the initial decision maker. Notice that the policy outcomes need not be particularly “good” only that they will converge to the ones that are “good”. In this sense our applied analysis predicts that the actual policies about climate change or international trade eventually will hit the right level, even if we will end up with a hotter planet than we would like.

Our paper is related to the literature that studies macroeconomic environments with time-inconsistency features typically characterized in terms of their Markov equilibria (e.g., [Cohen and Michel \(1988\)](#), [Currie and Levine \(1993\)](#), [Krusell and Ríos-Rull \(1996\)](#), [Klein and Ríos-Rull \(2003\)](#), [Klein et al. \(2005\)](#), [Bassetto and Sargent \(2006\)](#), [Klein et al. \(2008\)](#), [Bassetto \(2008\)](#), [Krusell et al. \(2010\)](#), [Martin \(2011\)](#), [Azzimonti \(2011\)](#)). It also addresses the type of environments previously studied by posing trigger strategies ([Chari and Kehoe \(1990\)](#), [Phelan and Stacchetti \(2001\)](#)). Our workhorse example builds upon the quasi-geometric discounting growth model analyzed by [Strotz \(1956\)](#), [Phelps and Pollak \(1968\)](#), [Laibson \(1997\)](#), [Krusell and Smith \(2003\)](#), [Bernheim et al. \(2015\)](#), [Chatterjee and Eyigungor \(2016\)](#), [Cao and Werning \(2018\)](#), [Halac and Yared \(2017\)](#), among others. Finally, we build on the literature on refinements of subgame perfect equilibrium, particularly in relation to renegotiation proofness ([Farrell and Maskin \(1989\)](#), [Kocherlakota \(1996\)](#), [Asheim \(1997\)](#), [Ales and Sleet \(2014\)](#)).

Other papers that have analyzed dynamic institution building applied to macroeconomic problems include [Acemoglu and Robinson \(2000\)](#), [Acemoglu et al. \(2012\)](#), and [Acemoglu et al. \(2015\)](#). These papers emphasize the role of changing the distribution of power within groups in the context of Markov equilibria as the mechanism that generates slow institutional buildup. [Piguillem and Riboni \(2015, 2020\)](#) explore the role of institutional arrangements where the status quo plays a special role as a way of disciplining time-inconsistent policymakers in the presence of heterogeneity.

Our notion of equilibrium is related to Reconsideration Proofness in [Kocherlakota \(1996\)](#) which is based on a notion of symmetry in the values obtained by all agents along the equilibrium path. It is also related to the equilibrium concept for overlapping-generations economies in [Prescott and Rios-Rull \(2005\)](#) where the symmetry included the passive first generation, but were the notion of no-delay was implicitly stated if one were to interpret the second generation as the first. In this paper we make the symmetry notion operational via the no restarting condition and we make explicit the no-delay condition. We also make the definition of organizational equilibrium to be compatible with state variables.⁶ Our use of state variables is defined for environments that display a weak separability property: preferences can be decomposed between a set of actions that we label “re-scaled actions”

⁶Kocherlakota defines a “state” in his work, but his state only depends on the expectation about current and future actions and is thus purely forward looking. In our case, we define a state as arising from past actions (including possibly past actions of nature, if randomness is present). This is in line with the literature on optimal control and dynamic programming.

and the state of the economy.⁷ Separability allows us to compare actions and outcomes across periods in which state variables are different.

For economies that do not have the separability property we propose to approximate them using separable environments as approximating economies. The approximation, while being first order, is not necessarily linear, but can take any functional form of our choice. We can look for functional forms that yield solutions to the Markov and Ramsey allocations that are close to those of the original economy that can be found with arbitrary accuracy using standard methods. Having similar Markov and Ramsey allocations gives a rationale to believe that the organizational equilibrium allocation of the approximating economy is also close to our object of interest.

Two other papers are of special relevance. Like us, [Nozawa \(2018\)](#) extends the notion of a reconsideration-proof equilibrium to economies with state variables. However, his extension imposes too strict requirements and leads to nonexistence of an equilibrium in many applications. By relying on weak separability, our approach allows us to define “state-free” notions of the economic environment and to establish existence. [Brendon and Ellison \(2018\)](#) analyze optimal policy in the Ramsey tradition, but they restrict the planner to choosing policies that satisfy a recursive Pareto criterion: this criterion disallows sequences that benefit policymakers in the early periods but are dominated for all policymakers from a given time onward. Like them, we also reject policies that allow early decision makers to dictate future paths that lead to early benefits purely at the expense of future decision makers. Rather than developing an optimality criterion, we propose a solution concept aimed at positive analysis, where implicit cooperation across policymakers at different times builds over time. Because of this different motivation, our “no-restarting condition” is imposed on a period-by-period basis. The presence of state variables causes problems in their environment as well, and our approach based on weak separability could be fruitfully applied there too.⁸ Interestingly the allocation that they propose coincides with the steady state to which our equilibrium converges to, the one that maintains behavior constant and is the best one among those.

Concerning the specific applications, [Matsuyama \(1990\)](#) analyzed a setting where trade liberalization creates a time inconsistency problem. He studies subgame-perfect equilibria, finding cyclical behavior, but argues that these equilibria featuring temporary protectionism fail renegotiation proofness. The short-run motive for protectionism in our environment is different from his, as we emphasize the adverse distributional impact accruing to workers in the sector exposed to international competition rather than to the firms that invest in that sector; more importantly, by applying our equilibrium

⁷The same type of separability property is also explored in [Halac and Yared \(2014\)](#) and [Halac and Yared \(2017\)](#), but they focus on different SPE refinements where agents’ type is private information.

⁸Our approach encompasses the more specific cases introduced by [Brendon and Ellison](#) in their latest version to account for state variables.

notion we find equilibria that gradually build institutions supporting free trade, and can make predictions about the level at which the process stops. [Maggi and Rodríguez-Clare \(2007\)](#) also build a political-economy model that entails features similar to our environment, but they focus on a one-time agreement that sets a dynamic path of liberalization. They also find that an agreement features gradual liberalization, since the adverse effect on sectors exposed to international competition gradually vanish as factors of production (capital, in their case) become more mobile across sectors over time. In a setting in which trade agreements can evolve over time, their environment also generates time inconsistency. Using our equilibrium, we show how this time inconsistency is resolved and how liberalization can take place in a succession of agreements, rather than a single one.

We start by posing the issues with time-inconsistent preferences in the context of the well understood quasi-geometric discounting growth model with log utility and full depreciation in [Section 2](#). We define organizational equilibrium for separable economies in [Section 3](#), where we also describe the connections to game theory and the strategy to study non-separable economies by using approximating separable economies. We also show how to tackle economies where a large decision maker such as the a policymaker interacts with a continuum of competitive agents, which requires that we adapt our concept to hybrid settings of competitive and strategic behavior. We then turn to the two main substantive issues addressed in this paper. [Section 4](#) has the analysis of how the slow implementation of policies to address climate change is both a fair description of the policies chosen by current governments as well as the predictions of our equilibrium notion because of the time inconsistency of the environment. We then turn to study the implementation of free trade policies, again documenting the slow implementation of tariff reductions as well as how a suitable multi-country model generates them as an organizational equilibrium in [Section 5](#). [Section 6](#) concludes.

2 A Motivating Example

Our equilibrium concept can be heuristically described as the best among those requiring that:

- in equilibrium no period- t agent can do worse than any period- τ agent for $\tau < t$ because then it could become a τ agent;
- in equilibrium no period- t agent can do worse than by staying out of the plan and letting the equilibrium unfold as if the economy started in the following period.

To provide the basic intuition, we revisit the canonical growth model with quasi-geometric discounting, log utility and full depreciation and compare what our equilibrium notion implies relative to other standard equilibrium concepts. The quasi-geometric discounting could be interpreted more

broadly as the consequence of aggregating different parties in the population with different time preferences (Jackson and Yariv, 2014, 2015).⁹

This example is easy to characterize (it has some closed form solutions), and it allows us to ignore any consideration related to a competitive equilibrium emerging from the interaction with other agents, a case that we analyze in Section 3.6. More importantly it displays a form of separability that allows us to decompose the rewards of any feasible allocation as a separable function of the initial capital and the subsequent sequence of saving rates. We will exploit this decomposition to provide a way of comparing rewards across agents who may be endowed with different levels of capital.

Assume that the production function is

$$f(k_t) = k_t^\alpha,$$

and the agent's period utility function is

$$u(c_t) = \log c_t.$$

The relevant state of the economy is capital k_t with law of motion

$$k_{t+1} = f(k_t) - c_t.$$

The lifetime utility for the agent at period t is

$$u(c_t) + \delta \sum_{\tau=1}^{\infty} \beta^\tau u(c_{t+\tau}).$$

It is easy to see that the agent will disagree with itself in the next period if $\delta \neq 1$.

To see the separability property, it is useful to work with saving rates defined as $s_t = k_{t+1}/k_t^\alpha$. Any sequence of saving rates $\{s_j\}_{j=0}^{\infty}$, together with an initial capital stock k_0 , implies a sequence of capital levels $k_t = k_0^\alpha \prod_{j=0}^{t-1} s_j^{\alpha^{t-j-1}}$. The corresponding lifetime utility for the agent in period 0 is

$$U(k_0, s_0, s_1, \dots) = \log[(1 - s_0)k_0^\alpha] + \delta \sum_{j=1}^{\infty} \beta^j \log[(1 - s_j)k_j^\alpha] = \frac{\alpha(1 - \alpha\beta + \delta\alpha\beta)}{1 - \alpha\beta} \log k_0 + V(s_0, s_1, \dots),$$

⁹Jackson and Yariv (2014, 2015) discuss how present bias emerges naturally in collective decision problems when agents have heterogeneous discount factors. One interpretation of this setting is that it is the problem of a planner facing such households. The functional form that we assume follows Laibson (1997) and corresponds to the one implied by a utilitarian planner aggregating the preferences of a two-person household with heterogeneous discount factors, in the limiting case in which one member of the household only cares about present consumption. We show how the equilibrium concept can be applied away from the limiting case in Appendix C.

where

$$V(s_0, s_1, \dots) \equiv \log(1 - s_0) + \frac{\delta\alpha\beta}{1 - \alpha\beta} \log(s_0) + \delta \sum_{j=1}^{\infty} \beta^j \left(\log(1 - s_j) + \frac{\alpha\beta}{1 - \alpha\beta} \log(s_j) \right). \quad (1)$$

The same logic follows for the period- t agent, its lifetime utility is the sum of a term that depends on the period- t capital and a term that depends only on saving rates of periods t and after. We write it compactly as

$$\underbrace{U(k_t, s_t, s_{t+1}, \dots)}_{\text{total payoff}} = \underbrace{\frac{\alpha(1 - \alpha\beta + \delta\alpha\beta)}{1 - \alpha\beta} \log k_t}_{\text{capital payoff}} + \underbrace{V(s_t, s_{t+1}, \dots)}_{\text{action payoff}}. \quad (2)$$

Two relevant implications of separability are that, as of period t , the relative preferences over a sequence of saving rates $\{s_t, s_{t+1}, \dots\}$ are independent of the initial level of capital, and also the set of feasible sequences is the same no matter what initial capital is.¹⁰

Before we discuss our proposed notion of equilibrium, we first characterize the allocations implied by some commonly used equilibrium concepts, including the Ramsey outcome, the (differentiable) Markov equilibrium,¹¹ and the best allocation supported by a constant saving rate.

Ramsey outcome The assumption that the period-0 agent is able to commit to a particular sequence of saving rates $\{s_\tau\}_{\tau=0}^{\infty}$ chosen at time 0, gives us a useful benchmark. The problem is

$$\max_{\{s_t\}_{t=0}^{\infty}} u(c_0) + \delta \sum_{t=1}^{\infty} \beta^t u(c_t),$$

$$\text{subject to} \quad k_{t+1} = s_t k_t^\alpha, \quad c_t = (1 - s_t) k_t^\alpha, \quad k_0 \text{ given.}$$

The solution to the Ramsey problem can be summarized as

$$s_t = \begin{cases} s_0^R = \frac{\alpha\delta\beta}{1 - \alpha\beta + \delta\alpha\beta}, & t = 0, \\ s^R = \alpha\beta, & t > 0. \end{cases}$$

The initial agent discounts period 1 at a higher rate than future periods, so she chooses a lower saving rate in period 0 than in the future, $s_0^R < \alpha\beta$.

¹⁰While this example satisfies additive separability, weak separability is sufficient for our results. When the production function is linear in capital, as is the case for an individual who takes the interest rate as given, separability holds for all CRRA utility functions.

¹¹More precisely, we look at the Markov equilibrium that is the limit of finite economies. See [Krusell and Smith \(2003\)](#) for details of the trigger-strategy equilibria that can be represented via non-differentiable Markov perfect equilibrium.

Markov equilibrium Another useful equilibrium concept is the Markov equilibrium. Its implied allocation satisfies a generalized Euler equation (GEE). Let $g(k)$ denote the policy function for tomorrow's capital k' , the GEE is

$$u_c(f(k) - g(k)) = \beta u_c\left(f[g(k)] - g[g(k)]\right) \left[\delta f_k[g(k)] + (1 - \delta) g_k[g(k)] \right],$$

which yields a closed form solution for the policy function, $g(k) = \frac{\alpha\delta\beta}{1-\alpha\beta+\delta\alpha\beta}k^\alpha$, and a constant saving rate

$$s^M = \frac{\alpha\delta\beta}{1 - \alpha\beta + \delta\alpha\beta}. \quad (3)$$

Note that the saving rate in the Markov equilibrium is the same as the first period's saving rate in the Ramsey outcome, and $s^M < \alpha\beta$. Note also that this saving rate is independent of the level of capital: hence, the current player cannot influence the future saving rates, and makes a choice taking those future rates as given. This is a general consequence of separability.

Best constant savings rate The Markov equilibrium features a particular constant saving rate, but it may not yield the best payoff compared with other constant saving rates. Suppose that agents are restricted to choose a constant saving rate for themselves and for all future agents; then, the best constant saving rate solves

$$\max_s u(c_0) + \delta \sum_{t=1}^{\infty} \beta^t u(c_t)$$

$$\text{subject to} \quad k_{t+1} = s k_t^\alpha, \quad c_t = (1 - s) k_t^\alpha, \quad k_0 \text{ given.}$$

The solution to this problem is given by

$$s^B = \frac{\delta\alpha\beta}{(1 - \beta + \delta\beta)(1 - \alpha\beta) + \delta\alpha\beta}, \quad (4)$$

and it satisfies $s^B \in (s^M, s^R)$ when $\delta \in (0, 1)$.

Towards Organizational Equilibrium Separability makes it easy to discuss the properties of the allocations implied by these equilibrium concepts, and in particular whether any time- t agent would prefer the sequence of savings rates given to another agent. We next explore how the two criteria mentioned at the beginning of this section relate to the previous equilibrium concepts.

In the Ramsey outcome, the initial agent achieves a higher action payoff than any subsequent agent, as $V(s^M, s^R, s^R, \dots)$ is higher than $V(s^R, s^R, s^R, \dots)$ and it is feasible for future agents as well. If a time- t agent were to be able to become the initial agent, it would always do so, violating the first criterion in our wish list. Our notion of organizational equilibrium excludes such an allocation as

an equilibrium. A similar property would hold in the best subgame perfect equilibrium that can be supported under the threat of reverting to Markov after a deviation; there too the initial agent receives a more favorable treatment than subsequent agents, in the sense that its action payoff is higher.

The best constant saving rate s^B does satisfy our first criterion for an organizational equilibrium, that no agent can do better by switching to an earlier allocation. However, it does not satisfy the second, since $V(s^M, s^B, s^B, \dots) > V(s^B, s^B, \dots)$: the initial agent would rather choose s^M and let every subsequent agent choose s^B , which amounts to sitting out and letting the equilibrium unfold from the following period.

The Markov equilibrium avoids these issues. First, the Markov equilibrium does not favor the initial agent, as the equilibrium path features a constant saving rate and the action payoff is the same for all agents. Second, when staying out and letting equilibrium unfold the following period, the current agent will choose the Markov saving rate itself, which yields exactly the same payoff. The question is then whether anything else can be better than the Markov equilibrium and still satisfy the two criteria.

The answer is yes. Putting the two criteria together, the organizational equilibrium points to an allocation implied by a sequence of saving rates $\{s_0^*, s_1^*, s_2^*, \dots\}$, such that

$$\bar{V} = V(s_t^*, s_{t+1}^*, s_{t+2}^*, \dots) \text{ for all } t, \quad (5)$$

and also

$$V(s_t^*, s_{t+1}^*, s_{t+2}^*, \dots) \geq V(s^M, s_0^*, s_1^*, s_2^*, \dots). \quad (6)$$

To see how we obtain condition (5), recall that the no-restarting condition requires that

$$V(s_t^*, s_{t+1}^*, s_{t+2}^*, \dots) \leq V(s_{t+1}^*, s_{t+2}^*, s_{t+3}^*, \dots) \text{ for all } t,$$

as period $t + 1$ agent can become period- t agent. Meanwhile, if the sequence did not yield a constant value, agent 0 could improve by skipping the periods up to the point at which the value becomes constant.¹² We will show in Section 3.4 that the maximum constant value of \bar{V} that can be attained is given by $V(s^B, s^B, s^B, \dots)$, and the equilibrium path of saving rates increases gradually such that choosing s^M followed by $\{s^*\}_{t=0}^\infty$ does not yield a higher utility than \bar{V} . On the other hand, condition (6) makes sure the no-delay condition is satisfied by construction. In fact, it also implies

¹²Our compactness assumptions ensure that, if such a sequence remained strictly increasing even as time goes to infinity, then there would exist a sequence of constant value attaining the limit value (higher than all of the elements of the strictly increasing sequence).

that restarting from s_0 amounts to a sufficient (although often stronger than necessary) punishment to deter any deviation along the transition.

To construct the sequence of saving rates s_t^* , recall that in our example economy, the action payoff function $V(\cdot)$ is given by equation (1). Together with condition (5), it implies a recursive relationship between s_{t+1}^* and s_t^*

$$\beta(1 - \delta) \log(1 - s_{t+1}^*) = \frac{\delta\alpha\beta}{1 - \alpha\beta} \log s_t^* + \log(1 - s_t^*) - (1 - \beta)\bar{V}, \quad (7)$$

Following this relationship, there are a continuum of paths that converge to the stationary point, which is given by the best constant saving rate s^B . By setting the starting point s_0^* sufficiently low, the initial agent will not find it attractive to stay out and wait for the equilibrium to unfold. An equilibrium path has the following properties: first, it displays a gradual transition from a relatively low saving rate to a high saving rate s^B , as if a good reputation is built over time. Second, along the transition path, the action payoff stays the same as that implied by the best constant saving rate, and is larger than that in the Markov equilibrium.

In the next section, we formalize these notions as the result of a game-theoretical refinement, and we show that the features of this motivating example are general properties of an organizational equilibrium.

3 Organizational Equilibrium

This section contains the core of our equilibrium analysis. After we describe the economic environment (preferences and technology), we setup a standard game in which there is a specific period 0 and the entire history is recorded. In such a game, we define the organizational equilibrium as a subgame perfect equilibrium with a particular refinement. In Section 3.2, we provide an alternative representation of the same economic environment as a different game with incomplete information that eliminates the special role of period 0. In this game, we show that the no-delay condition central to the notion of organizational equilibrium is a requirement for a sequential equilibrium rather than a refinement. In Section 3.3, we prove the general properties of an organizational equilibrium and a recursive method that directly constructs the equilibrium outcome.

Consider a generic environment of sequential decision makers (typically those that have a time-consistency problem) where there is a physical state variable $k \in K$. Specifically, given the current level of k , the agent making a decision will choose an action a from a set A . The state evolves according to $k_{t+1} = F(k_t, a_t)$. Preferences for an agent making decisions in period t are given by

$U(k_t, a_t, a_{t+1}, a_{t+2}, \dots)$. The first assumption is that functions U and F are independent of calendar time, which allows meaningful welfare comparisons across decision makers.

We restrict the environments that we study to those in which the utility is weakly separable between the state and the sequence of actions, such that the preference ordering over sequences is independent of the initial state. Formally:

Assumption 1. 1. *At any point in time t , the set of feasible actions A is independent of the state k_t ;*

2. *U is weakly separable in k and in $\{a_s\}_{s=t}^\infty$, i.e., there exist functions $v : K \times \mathbb{R} \rightarrow \mathbb{R}$ and $V : A^\infty \rightarrow \mathbb{R}$ such that*

$$U(k, a_t, a_{t+1}, a_{t+2}, \dots) \equiv v(k, V(a_t, a_{t+1}, a_{t+2}, \dots)). \quad (8)$$

and such that v is strictly increasing in its second argument.

Sometimes the original problem does not satisfy Assumption 1, but it is possible to re-scale actions in such a way that it does.¹³ As an example, the original specification of the saving problem with quasi-geometric discounting does not satisfy Assumption 1 if we define the action to be consumption: the feasible set of consumption levels depends on initial capital.¹⁴ Formally, suppose that the set of feasible actions at any capital level k is $\tilde{A}(k) \subseteq \tilde{A}$ and that preferences are given by $\tilde{U}(k_t, \tilde{a}_t, \tilde{a}_{t+1}, \tilde{a}_{t+2}, \dots)$. Our construction still applies as long as it is possible to find a set of actions A and a function γ such that $\tilde{a} = \gamma(a, k)$ and that Assumption 1 holds for A , where

$$U(k, a_t, a_{t+1}, a_{t+2}, \dots) \equiv \tilde{U}(k, \tilde{a}_t, \tilde{a}_{t+1}, \tilde{a}_{t+2}, \dots),$$

and where for $t \geq 0$, \tilde{a}_t is computed recursively as

$$\begin{aligned} \tilde{a}_t &= \gamma(a_t, k_t), \\ k_{t+1} &= F(k_t, \tilde{a}_t). \end{aligned} \quad (9)$$

¹³Notice also that any one-to-one transformation of a will preserve weak separability, so the action space is only defined up to such transformations. As an example, for the case analyzed in Section 2, any monotone transformation of the saving rate would be an equally valid action, yielding the same equilibrium outcomes.

¹⁴Note that weak separability automatically fails if certain actions are only feasible for some levels of capital, since, holding actions fixed, the left-hand side of (8) would then be well defined for some values of k and not for others.

3.1 The Standard Game

The general setup that we have described is typically modeled as the following game. There is an infinity of players indexed by the time at which they act, $\{0, 1, \dots\}$, each of whom has preferences given by (8). At each time t , the history of play is given by $h^t := (a_0, a_1, \dots, a_{t-1})$, with $h^0 := \emptyset$. A strategy σ_t for player t is a mapping from the set of time- t histories, H^t , to the set of actions A . A strategy profile is a sequence of strategies, one for each player: $\sigma := (\sigma_0, \sigma_1, \dots)$. As usual, it is also convenient to define a continuation strategy after history h^t , $\sigma|_{h^t}$, represented by the restriction of $(\sigma_t, \sigma_{t+1}, \dots)$ to the histories following h^t . From any history h^t , a strategy profile induces a sequence of future actions, which we denote through the following short-hand notation:

$$a_{t+1, \sigma|_{h^t}} := \sigma_{t+1}(h^t, \sigma(h^t)), \quad a_{t+s, \sigma|_{h^t}} := \sigma_{t+s}(h^t, a_{t+1, \sigma|_{h^t}}, \dots, a_{t+s-1, \sigma|_{h^t}}).$$

Starting from the set of subgame-perfect equilibria of this game, we limit the equilibria in our analysis by first imposing the following refinements:

Requirement 1 (State Independence). *We limit attention to equilibria in which the strategies followed by all players are independent of the state k .*

Requirement 2 (No-restarting and optimality). *We limit attention to equilibria such that:*

- *they are symmetric, in that the action payoff*

$$V(a_{t+1, \sigma|_{h^t}}, a_{t+2, \sigma|_{h^t}}, \dots)$$

is the same after any history of play;

- *No other symmetric state-independent equilibrium exists that attains a higher payoff.*

In the absence of a state variable (when preferences are independent of k), Requirement R2 corresponds to Kocherlakota's (1996) definition of reconsideration-proof equilibrium. In particular, no restarting is the natural adaptation of Kocherlakota's symmetry requirement: factoring out the effect of the state variable, the utility that an agent receives from its actions and those of its successors is independent of the past history of play.

As a first step, we extend the notion of reconsideration proofness to dynamic games, rather than purely repeated games. In the presence of a state variable, we assume that players coordinate on strategies that only depend on the history of play h^t and not on the physical state. Weak

separability plays an important role in this selection criterion: it ensures that the same actions can indeed be played independently of the current value of the state, and that each player’s preferences over current and future actions are independent of the state. This ensures that, if $\sigma|_{h^t}$ is a subgame-perfect equilibrium for a history that attains some level of the state \bar{k} , it is also a subgame-perfect equilibrium for any other history, even though the initial level of capital may be different from \bar{k} . In this case, an organizational equilibrium imposes symmetry only in that the payoff of the subutility V is independent of the history of play, but the payoff of each time- t player is still different across histories which lead to different levels of the state. Intuitively, a different state implies a different set of possible utility levels going forward, so we should expect it to affect payoffs in the subgames going forward. However, this dependence of utility from the state takes a simple form under weak separability, and there is a natural mapping across histories with different levels of capital: the same sequences of actions are possible under any level of capital, and the preferences of player t over the sequences from date t on are also represented by the subutility V , independent of k_t . For this reason, imposing reconsideration proofness on preferences represented by V alone is appealing.

It is useful to compare our notion to previous attempts at dealing with state variables in this context. An extension of reconsideration proofness to environments with state variables was proposed by [Nozawa \(2018\)](#). Nozawa requires weakly reconsideration-proof equilibria to be such that the equilibria of all subgames share the same payoff *function* $\Psi(k)$, which depends on the state; in the absence of the state, this reduces to Kocherlakota’s (1996) symmetry requirement. A strong reconsideration-proof equilibrium is then an equilibrium in which $\Psi(k)$ is undominated by any other equilibrium *point by point*. This is often too strong a requirement, and hence existence may fail.¹⁵ Our approach avoids this problem because symmetry is defined by a single utility level \bar{V} , namely the action payoff attained by each agent, rather than a function. This is possible because weak separability allows us to extend this single level to the complete payoff (which remains a function of the state) by setting it equal to $v(k, \bar{V})$.

An alternative approach adopted in the past is revision proofness, which was introduced by [Asheim \(1997\)](#) and made explicit as a game in [Ales and Sleet \(2014\)](#). In their papers, a larger class of credible punishments is allowable. Specifically, under reconsideration proofness, if Σ is the set of equilibrium strategies of the game, each player at any time t is allowed to coordinate current and future play to its favorite element of Σ . Under revision proofness, player t ’s coordination power is limited because it is required to propose deviations from the equilibrium path of play that benefit *all* future players. The resulting equilibrium set is much larger. For the case of quasi-geometric discounting with linear preferences, [Ales and Sleet \(2014\)](#) show that all subgame-perfect paths better than the Markov equilibrium are revision proof. In environments with state variables, a limitation of revision

¹⁵As an example, no reconsideration-proof equilibrium would exist in the example of Section 2.

proofness is that it is unclear how a future player could “block” a revision proposal when it would inherit a different state under the revision proposal and would thus not be able to continue with the original strategy. Our notion of organizational equilibrium retains the unilateral aspect of deviations from reconsideration proofness, but it relies on weak separability to define and impose symmetry across different levels of capital.

Proposition 6 in Appendix A proves existence of an equilibrium satisfying Requirements 1 and 2 under the following technical assumption:

Assumption 2. 1. A is a convex compact subset of a locally convex topological linear space with topology ρ_x .

2. V is quasiconcave over A^∞ .

3. V is continuous over A^∞ with respect to the product topology ρ_x^∞ .

There can be many such equilibria. In Section 2, Requirements 1 and 2 imply that saving rates must satisfy the difference equation (7), but they do not provide an initial condition s_0 . Our next step is to push our selection further, and capture the idea that time inconsistency is not overcome all at once out of the blue, but rather that intertemporal cooperation takes time to build. If players coordinated to start from a high degree of cooperation from the beginning, there might be an incentive for a player to defect and hope that such a coordination takes place in the future. We approach this reasoning in a formal way by altering the game in Section 3.2, but for now we simply add as a third requirement that the first player has no incentive to deviate from the equilibrium path if the threat is that the same equilibrium will be played starting from the next period, as if the initial period had not taken place. Formally:

Requirement 3 (No Delay). Let σ be a subgame-perfect equilibrium strategy profile satisfying Requirements 1 and 2. Then, for each possible action $a \in A$,

$$V(a_{0,\sigma}, a_{1,\sigma}, a_{2,\sigma}, \dots) \geq V(a, a_{0,\sigma}, a_{1,\sigma}, \dots).$$

Our no-delay condition is related to Prescott and Rios-Rull (2005), although the form it takes here is different. In their overlapping-generations economy, there is an implicit form of no-delay motivated by the presence of an initial elderly cohort which is assumed to have played their best one-shot response before time zero. Here, we argue that explicitly imposing no-delay becomes desirable because it ensures that the coordination that gives rise to the initial equilibrium is not as generous as to tempt the first player to sit it out, play its best one-shot action, and count on the same coordinating mechanism to arise in the future.

We are now ready to define an organizational equilibrium:

Definition 1. *An organizational equilibrium is the outcome of any subgame-perfect equilibrium of the game described above that satisfies Requirements 1, 2, and 3.*

We choose to define an organizational equilibrium as the *outcome* of an equilibrium in a game-theoretic sense. In Proposition 3 and Corollary 1 below, we provide a way of characterizing this outcome directly, without describing the underlying strategies that support it. This makes it easier to apply our notion in a macroeconomic context: just as a competitive equilibrium can be defined as a single path, without reference to the underlying strategies, an organizational equilibrium also involves only the description of the path. At the same time, our paper provides the full strategic foundations that support such an equilibrium path, and the interested reader can then apply our proofs to construct them, if desired.

In order to prove existence of an organizational equilibrium, we use the following additional sufficient condition:

Assumption 3. *V is weakly separable in a_t and $\{a_s\}_{s=t+1}^\infty$, i.e., there exist functions $\tilde{V} : A \times \mathbb{R} \rightarrow \mathbb{R}$ and $\hat{V} : A^\infty \rightarrow \mathbb{R}$ such that, for all sequences $(a_t, a_{t+1}, a_{t+2}, \dots) \in A^\infty$,*

$$V(a_t, a_{t+1}, a_{t+2}, \dots) = \tilde{V}(a_t, \hat{V}(a_{t+1}, a_{t+2}, \dots)), \quad (10)$$

with \tilde{V} strictly increasing in its second argument.

Assumption 3 implies that a player’s preference ordering over the actions of future players is independent of its own choice: what player 0 views as “desirable” or “undesirable” future actions does not depend on its own choice. This is only a sufficient and not a necessary condition for existence of an organizational equilibrium; Appendix C shows an example where this Assumption fails and an organizational equilibrium can nonetheless be found. Assumption 3 implies that there is a clear way of defining what “cooperation” means, because the preferences of past players over the actions of the current and future players are not tied to the choices that those past players made. This makes our Requirement 3 particularly salient. In our hyperbolic discounting, past players would like future players to adopt a saving rate which is above the Markovian saving rate and closer to the long-run Ramsey outcome, and this is independent of what they themselves chose. It is in this context that the vague notion of “gradual development of cooperation” can be given formal meaning.

Proposition 1. *Under Assumptions 1, 2 and 3, an organizational equilibrium exists.*

Proof. See Appendix A □

The proof of Proposition 1 supports the equilibrium outcome by relying on a trigger strategy that reverts to restarting the path of play from the equilibrium action a_0 following any deviation, but in most applications (including all of those that we discuss below) this punishment is unnecessarily harsh. As an example, suppose that player t deviates (being the first to deviate) and chooses a_s , the equilibrium action that player $s < t$ took. A (weakly) sufficient deterrent is for future players to continue by playing a_{s+1}, a_{s+2}, \dots . We find it appealing that the punishment can be often made proportional to the size of the deviation.

While generalizing our equilibrium notion to a stochastic environment is beyond the goal of our current paper, we would regard exploiting this feature as an important ingredient in such a generalization. Intuitively, suppose that players are subject to impatience shocks, that lead them sometimes to backtrack on the way to intertemporal cooperation. As a concrete example, using our international trade application of Section 5, suppose that a political shock leads a government to reintroduce some trade barriers. Rather than restarting from square one, it is natural to assume that future players would react by resuming the slow march towards cooperation from this new, lower level.

3.2 An Alternative Game where Period 0 is not Special

We wish to go one step further and formalize the notion that intertemporal cooperation is fostered by the emergence of “good institutions,” or “good norms.” To do so, we build upon the game above, but we modify it so as to make sharper predictions about the start of play. In this alternative game, every agent has the ability to erase history and become the agent in period “minus one,” effectively letting the agent in the following period become the period zero agent. This game is another representation of environments with time consistency problems, that are our objects of interest, provided that we are willing to entertain that agents can actually erase history, or at least, provide a clean separation from their past.

More generally, our game pins down what do we mean by “the first period,” an issue generally ignored in the literature. We accomplish that by giving any agent the option to either go along with whatever time index it has from the past or to become the (or better, a) time zero agent. This approach has a recursive flavor in the sense there is nothing special to the timing of birth of any particular agent and we formalize it below. We think that it has various attractive features and an unattractive one: it allows us to use the powerful tools of dynamic games, while at the same time preventing any specific agent to be the special time-zero agent; it provides a natural justification for the no-delay condition embedded in Requirement 3, which emerges naturally as a requirement for a sequential equilibrium; finally, it provides a rationale for the unorthodox name (organizational) of the equilibrium concept, since an organization consists of an ongoing, uninterrupted set of agents

that choose to go along with a plan rather than set up their own organization. The unattractive feature of our interpretation of the initial period is that it requires the make-believe assumption that subsequent agents can forget the previous history whenever one of them chooses to become the type-zero agent. Such assumption may not sound appealing in a literal interpretation of the problem of an agent with time-inconsistent preferences, but it is more so when we think of collections of governments and their possible explicit choices of breaking with the past, claiming that they do not share anything with previous governments, which we take as becoming agent 0 or, simply, restarting history.

We choose to present this alternative game in a separate section because this interpretation or design of the strategic interaction is not strictly necessary to develop our equilibrium notion, but it makes the no-delay condition not yet one more refinement of subgame perfection, but a necessary property of a sequential equilibrium.

Formally, the game of Section 3.1 is modified as follows. We now assume that the actions of past players remain unobservable to the current player until an “organization” is set up to record past play. The opportunity to set up an organization arrives at a stochastic point in time t , where the probability distribution over the time of arrival is unrestricted, except that it is assumed to have full support over \mathbb{N} ; the precise time is unobserved by the players, who can only know whether setting up an organization is possible when they are called to play. In sum, let \hat{t} be the time at which the opportunity to set up record keeping emerges. For $t < \hat{t}$, players do not observe past play and choose an action a_t that cannot be conditioned on (a_0, \dots, a_{t-1}) .¹⁶ In each subsequent period $t \geq \hat{t}$, if no record-keeping organization is in place, player t can start one, so that player $t + 1$ will be able to condition its actions on player t 's choice a_t . This choice is taken without knowing whether the opportunity was available in the past, or whether it newly arrived in period t . If record-keeping has been in place since a period $\tilde{t} \in [\hat{t}, t)$, player t can choose to continue the current organization, so that player $t + 1$ can condition its actions on $(a_{\tilde{t}}, \dots, a_t)$, or it can start a new organization, in which case only a_t is known to player $t + 1$, or it can discontinue the current organization without replacing it, in which case player $t + 1$ cannot condition its actions on any of the past actions (a_0, \dots, a_{t-1}) .¹⁷

With the limitations on record-keeping described above, the game unfolds otherwise as in Section 3.1, with each player at time t choosing an action $a_t \in A$ (after making a record-keeping choice, if a choice is available). The preferences and the evolution of the state are the same as in Section 3.1. To quickly distinguish between the two games, we will from now on refer to the game of Section 3.1

¹⁶We continue to only consider equilibria in which strategies are independent of the state variable, which rules out inferring past play through this indirect channel.

¹⁷Even in this case, player t does not know if the opportunity to set up an organization first appeared in period \tilde{t} , or was available in earlier periods but was not taken up, or it was taken up but discontinued by earlier players.

as the game where record-keeping starts at time 0, and the game just described here as the game where history can be hidden. The introduction of incomplete information requires switching from considering subgame-perfect equilibria to sequential equilibria.¹⁸ We relegate the notation defining histories, information sets, and strategies to Appendix B.

Since this new game involves uncertainty, in order to find equilibria where strategies are independent of the state, we need to strengthen Assumption 1:

Assumption 4. 1. *At any point in time t , the set of feasible actions A is independent of the state k_t ;*

2. *There exist functions $\bar{v} : K \times \mathbb{R}_{++}$, $\bar{v} : K \times \mathbb{R}$, and $V : A^\infty \rightarrow \mathbb{R}$, such that*

$$U(k, a_t, a_{t+1}, a_{t+2}, \dots) \equiv \bar{v}(k)V(a_t, a_{t+1}, a_{t+2}, \dots) + \bar{v}(k). \quad (11)$$

This assumption is satisfied throughout all of our examples.

Proposition 2. *Consider a state-independent sequential equilibrium that satisfies Requirement 2 from period \hat{t} on. Such an equilibrium exists under Assumptions 2 and 4. Let \hat{t} be the realization of the (random) first time in which record keeping is possible, and let $(a_{\hat{t}}, a_{\hat{t}+1}, a_{\hat{t}+2}, \dots)$ be the path implied by the equilibrium, conditional on \hat{t} . Then $(a_{\hat{t}}, a_{\hat{t}+1}, a_{\hat{t}+2}, \dots)$ is an organizational equilibrium.*

Proof. See Appendix B. □

Remarks on no-delay and no-restarting conditions Notice that we do not redefine an organizational equilibrium in this section; rather, we prove that the same definition that we used in Section 3.1 also describes the equilibrium path of this new game from the point at which record keeping becomes possible. The no-delay condition is now a requirement for a sequential equilibrium. Suppose it were violated. Then, the player that moves at the first instance in which record keeping is possible would have an incentive to “pass the buck,” pretend that nothing has happened, let next period’s agent think that she is the first one to have access to record keeping, and play the best response to that. In their own way, both the no-restarting condition and the no-delay condition are designed to downplay the special role that period 0 has under time inconsistency. No-restarting looks forward and formalizes the notion that the time-0 player cannot impose an equilibrium that treats

¹⁸Note however that we do not need to keep track of beliefs within an information set. This is for two reasons: first, no future player will be able to distinguish between histories that are in the same information set at time t , which means that their actions will be the same independently of the specific node within the current information set; second, player t ’s payoff conditional on current and future actions and on the state is also independent of the specific node within the information set.

herself more favorably than the agents that follow. No-delay looks backward instead: it captures the idea that there is no special period 0 in which the world starts, but rather there always is a history that precedes any particular policy that we wish to analyze. In this context, a player at a given point in time will only find it optimal to start a process (or an organization) to overcome time inconsistency only if it is (weakly) in her best interest not to hope that the next player will do so instead.

3.3 Equilibrium Properties

We are interested in macroeconomic applications of the notion of organizational equilibrium. For these applications, keeping track of strategies is cumbersome. It would be desirable to characterize an organizational equilibrium directly in terms of the equilibrium sequence, without fully specifying the supporting strategies; this is similar to the way in which competitive equilibria are defined in macroeconomic models, which is also purely in terms of sequences of actions (and prices). In this section, we develop sufficient conditions that allow us to do this, and that provide further properties of the equilibrium allocation. Appendix C discusses the role of these conditions in greater details and shows how to apply similar ideas to characterize organizational equilibria in examples where these specific assumptions fail.

The following proposition proved in Appendix D provides a first step:

Proposition 3. *Let Assumption 1 hold. A sequence $\{\bar{a}_t\}_{t=0}^{\infty}$ that satisfies the following properties is an organizational equilibrium:*

1. *No-restarting:*

$$V(\bar{a}_t, \bar{a}_{t+1}, \bar{a}_{t+2}, \dots) = \bar{V} \quad \forall t \geq 0;$$

2. *Optimality: No other sequence satisfying no-restarting achieves a higher constant value;*

3. *No-delay:*

$$V(\bar{a}_0, \bar{a}_1, \bar{a}_2, \dots) \geq \max_a V(a, \bar{a}_0, \bar{a}_1, \dots).$$

Furthermore, if a sequence satisfying the three properties above exists, then all the organizational equilibria satisfy the same conditions.

When a sequence that satisfies the three properties of Proposition 3 can be found, we have a way of characterizing organizational equilibria directly in terms of sequences. We are also interested in

establishing the converse: specifically, that any organizational equilibrium is a sequence that satisfies the three properties of Proposition 3 (or, equivalently, that a sequence satisfying the 3 properties exists). To do so, we require further assumptions. The same assumptions are also instrumental in establishing that some of the properties that we observed in the simple example of Section 2 are true more generally.

In Section 3.1, we introduced some structure on the action payoff function V through Assumption 3 as a sufficient condition for the existence of an organizational equilibrium. Assumption 3 implies that the current action payoff is determined by the current action and a scalar sufficient statistic for the sequence of future actions. Our next assumption imposes a recursive structure on this sufficient statistic:

Assumption 5. *Let \widehat{V} be defined as in Assumption 3. There exists a function $W : A \times \mathbb{R} \rightarrow \mathbb{R}$, increasing in the second argument, such that, given any sequence $\{a_s\}_{s=t}^\infty \in A^\infty$,*

$$\widehat{V}(a_t, a_{t+1}, a_{t+2}, \dots) \equiv W\left(a_t, \widehat{V}(a_{t+1}, a_{t+2}, \dots)\right). \quad (12)$$

This assumption of course holds in the example of Section 2 and in many other applications of economic interest.¹⁹ As is often the case, a recursive structure is instrumental in constructing equilibria in infinite-horizon economies in which backward induction cannot be applied. Note however that the recursion applies only on preferences about *future* actions: a_t enjoys a special role in the preferences of player t , which is the essence of the time-consistency problem.

We then obtain (again proved in Appendix D):

Proposition 4. *Under Assumptions 2, 3, and 5, there exists an organizational equilibrium $\{a_t\}_{t=0}^\infty$ which is recursive in the value $\widehat{V}(a_t, a_{t+1}, a_{t+2}, \dots)$: that is, there exists a function $g : \mathbb{R} \rightarrow A \times \mathbb{R}$ such that $(a_t, v_{t+1}) = g(v_t)$, and $v_t = \widehat{V}(a_t, a_{t+1}, a_{t+2}, \dots)$ for all $t = 0, 1, \dots$*

Proposition 4 uses values as a state variable in ways similar to Abreu et al. (1986, 1990) (APS). However, as the proof shows, constructing the set of possible values is considerably more involved than in the case of APS. In APS, the set of equilibrium values can be obtained by starting from a

¹⁹The key difference between Assumption 3 and Assumption 5 is that the same function \widehat{V} appears both on the left and right-hand side in Assumption 5. The function V represents the utility that an agent at t receives from a sequence of actions starting at t . The function \widehat{V} represents the continuation utility that an agent at t derives from a sequence of actions starting at $t + 1$. Assumption 5 allows us to employ recursive methods. When 5 fails, computing an organizational equilibrium may be more challenging, but often there are alternative ways of making the problem recursive. As an example, Appendix C computes an organizational equilibrium for a planner that aggregates the preferences of agents with heterogeneous discount factors, as in Jackson and Yariv (2014, 2015). While Assumption 5 fails in that context, an alternative recursive representation is possible, and the nature of the equilibrium is similar.

large (convex) set of possible continuation values, and iterating backward through a monotonically shrinking sequence of sets until convergence. In our case, the symmetry property of organizational equilibria imposes that continuation equilibria must share the same value as the equilibrium of the game from time 0. This equality constraint breaks the APS procedure and requires a more complicated argument.

The following proposition and corollaries provide the main characterization of an organizational equilibrium as we use it throughout all of our applications.

Proposition 5. *Under Assumptions 2, 3, and 5, the following properties hold:*

1. *The action payoff of an organizational equilibrium coincides with the maximum value that $V(a, a, a, \dots)$ can attain under a constant action.*
2. *The payoff of an organizational equilibrium is weakly below that of the Ramsey outcome and weakly better than that of the best (state-independent) Markov equilibrium. The inequality is strict, except in the case in which the Ramsey outcome is attained by a constant allocation as in a Markov equilibrium.*
3. *$\widehat{V}(a_t, a_{t+1}, a_{t+2}, \dots)$ is increasing over time for any organizational equilibrium, and it converges to the value associated with the constant action profile that maximizes $V(a, a, a, \dots)$.*
4. *If \widehat{V} is a strictly quasiconcave function and the steady state that maximizes $V(a, a, a, \dots)$ is not a Markov equilibrium, then the convergence is not immediate: the initial value $\widehat{V}(a_0, a_1, a_2, \dots)$ is strictly below $\widehat{V}(a, a, a, \dots)$.*

Corollary 1. *Under Assumptions 2, 3, and 5, there exists a sequence $\{\bar{a}_t\}_{t=0}^{\infty}$ that satisfies Properties 1, 2, and 3 of Proposition 3.*

Part 1 of Proposition 5 implies that an organizational equilibrium attains the same payoff as the allocation associated with the best constant action. In the presence of time inconsistency, this constant action necessarily trades off the short run costs and long-run benefits of deviating from a Markov equilibrium; Part 2 accordingly shows that an organizational equilibrium always has an intermediate payoff between the Ramsey outcome and the best Markov equilibrium. The three notions yield the same payoff only when the Ramsey outcome can be attained by a Markov equilibrium: in this case, there is no time consistency problem, because player t has an incentive to choose the Ramsey allocation if it believes that future players will also do the same. Whenever time consistency has bite, an organizational equilibrium falls short of the Ramsey outcome. Part 3 and 4 in turn shows that convergence takes time. Whether this is interpreted as a gradual build-up of trust or a willingness

by the current institutions to own up to the actions of previous decision makers, this is at the heart of our notion of *organizational* equilibrium.

Finally, Corollary 1 shows that the procedure we used to compute organizational equilibria in Section 2 applies more generally. Specifically, we first compute a constant action profile that maximizes $V(a, a, a, \dots)$: this is what would be chosen by the decision maker at time 0 if future players were committed to take the same action. This maximization yields a value V^* , which must remain constant along the path, i.e., $V(a_t, a_{t+1}, a_{t+2}, \dots) = V^*$. The proof of Proposition 5 shows how to exploit the recursive structure implied by Assumption 5 to construct other sequences that attain the value V^* , and how to pick the initial action a_0 ; the resulting sequence converges to the constant action profile (a, a, a, \dots) , but not immediately (except for the degenerate case of no time-inconsistency).

3.4 Application to the Growth Model

We now apply the results that we derived for a generic case to the specific application of Section 2. By inspecting Equations 1 and 2, we can verify that our preferences satisfy Assumptions 1, 3, 4, and 5. Assumption 2 is satisfied if we pick an arbitrarily small $\epsilon > 0$ and we impose a minimum saving rate ϵ and a maximum saving rate $1 - \epsilon$.²⁰ Following Proposition 3 and Corollary 1, we compute an organizational equilibrium by directly looking at sequences that satisfy the properties of no-restarting, optimality, and no-delay of Proposition 3.

First, in order to attain a constant action value, as implied by the no-delay condition, the sequence of saving rates must satisfy the difference equation (7). From Proposition 4, we know that the solution has a recursive structure. In this case, instead of writing the recursion in terms of continuation values, it is more convenient to write it directly in terms of the saving rates that will be undertaken. Accordingly, using (7), we define

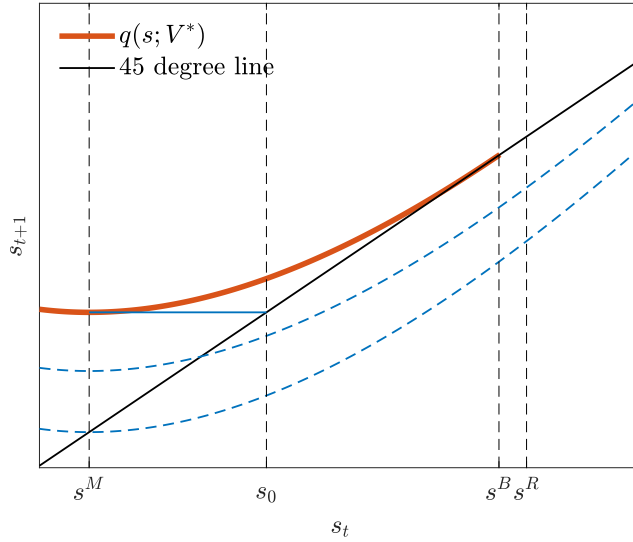
$$q(s; \bar{V}) := 1 - \exp \left\{ \frac{-(1 - \beta)\bar{V} + \frac{\delta\alpha\beta}{1-\alpha\beta} \log s + \log(1 - s)}{\beta(1 - \delta)} \right\}, \quad (13)$$

so that equation (7) can be rewritten as $s_{t+1} = q(s_t; \bar{V})$. The blue lines and the red line in Figure 1 represent this difference equation under different values for \bar{V} .

The optimality property is satisfied by the solution that achieves the highest possible action payoff;

²⁰We could prove that these bounds do not affect our results when ϵ is sufficiently close to zero. Intuitively, while household preferences are not time consistent, the degree of time inconsistency is limited, so that the player moving at time t would not want to consume all of its endowment and at the same time it would not want future players to save all of theirs. We omit a formal analysis of these bounds for brevity.

FIGURE 1: Evolution of the Saving Rate for Sequences that Attain a Constant Value



by Proposition 5 this is the payoff attained by the action s that maximizes $V(s, s, s, \dots)$, which is given by (4). This is represented by the red line in Figure 1 and is associated with a value that we define V^* .²¹

Even after we pin down $\bar{V} = V^*$, the difference equation (7) admits a continuum of solutions, indexed by the initial condition. From the perspective of the agent who makes the proposal, all of these sequences of saving rates yield the same payoff; however, they yield different payoffs from the perspective of future agents. As an example, consider the following two proposals: the first one is $s_t = s^B$ for all t , and the second one starts from $s_0 = s^M$ and subsequently follows a sequence dictated by the difference equation (7). Both of these sequences imply the same action payoff $V(s_0, s_1, \dots) = V(s^B, s^B, \dots) = V^*$ for the agent who makes the proposal, as a higher saving rate today is rewarded with higher saving rates in the future. The no-delay condition restricts the range of initial conditions that correspond to an organizational equilibrium; s_0 must be such that $V(s^M, s_0, s_1, \dots) \leq V^*$, so that the initial agent would not gain from waiting for the sequence to start in the next period.²² To guarantee the no-delay condition is satisfied, the sequence of saving

²¹For points to the right of the steady-state saving rate, the difference equation implies a sequence of saving rates converging to 1. This would violate the upper bound $1 - \epsilon$, so eventually no solution would be possible. Even if we ignore the bound, this solution violates the transversality condition and yields infinitely negative utility rather than V^* . When a higher value than V^* is used in the difference equation, there is no fixed point and all the solutions of the difference equation violate the transversality condition.

²²In general, the no-delay condition requires $\max_s V(s, s_0, s_1, \dots) \leq V(s_0, s_1, \dots)$. The fact that this maximum is always attained by s^M independent of the sequence is due to the specific functional-form assumptions of our example.

rates has to start with a low enough point to prevent the initial agent from sitting out; in Figure 1, this is no higher than the point s_0 , which corresponds to $q(s^M; V^*)$. This condition excludes the possibility of jumping to the steady-state saving rate immediately, and a gradual transition has to take place.

While there are many organizational equilibria, all of which give the same utility to the first generation, the equilibrium in which

$$s_0 = q(s^M; V^*)$$

yields the highest total utility for the subsequent generation by delivering the most capital. We select this one because it is the natural outcome if there is an arbitrarily small amount of altruism involved.²³

In the organizational equilibrium, time inconsistency is gradually overcome through time: at least from period 1 on, the saving rate exceeds that of the Markov equilibrium, and a virtuous cycle is started, with a monotonic increase which converges to s^B . Initial saving is limited by the temptation to let the next generation start the virtuous cycle. This temptation diminishes in subsequent periods, since restarting the virtuous cycle from scratch implies giving up on the accumulated effect of previous increases in s_t . Note that s^B is below the long-term savings rate of the Ramsey outcome, no matter how close to 1 δ is (as long as it is strictly less than 1): while the equilibrium path converges to the Ramsey outcome as $\delta \rightarrow 1$, it never coincides with it. This is in contrast with the best subgame-perfect equilibrium, which (by the folk theorem) coincides with the Ramsey outcome for all values of β sufficiently close to 1.

In this example, we can further establish the following properties of the converging dynamics.

Corollary 2. *The slope of the transition function $s_{t+1} = q(s_t; V^*)$ is positive when $s_t \in (s^M, s^B]$, equals to 0 when $s = s^M$, and equals to 1 when $s = s^B$.*

This result implies that not only the saving rate is monotonically converging to s^B , but also the speed of the convergence slows down when approaching the steady state. The transition begins with a comparatively low savings rate and will require a considerable amount of time to ultimately stabilize.

Comparison with other equilibria We now compare the properties of the sequence of capital stocks and the lifetime utilities in the organizational equilibrium with those in the Ramsey outcome

²³For the same reason, a fortiori we neglect saving rates below the Markov saving rate. Technically, such saving rates would also represent organizational equilibria, up to a lower bound which is the second solution to the equation $V(s^M, s_0, s_1, \dots) = V^*$, with the sequence (s_0, s_1, \dots) satisfying (7).

and the Markov equilibrium.²⁴

We first turn to the transition paths of different equilibria. We assume that the initial capital stock is the steady state capital stock in the Markov equilibrium, i.e., $k_0 = k^M$. The first row of Figure 2 displays the transition paths for the saving rate s_t and capital k_t . In the Markov equilibrium, the capital stock remains unchanged at the steady-state level that we assumed as a starting point. The Ramsey outcome features the same saving rate as the Markov equilibrium in the first period, so that the capital stock remains the same at the beginning of the second period. From the second period onwards, the saving rate increases to s^R permanently. The sequence of saving rates in the organizational equilibrium is induced by the transition function $s_{t+1} = q(s_t; V^*)$. Particularly, the saving rate in the first period is $s_0 = q(s^M; V^*) > s^M$, and the capital is initially higher than in the Ramsey allocation. Over time, the saving rates increase gradually and converge to $s^* < s^R$. Asymptotically, capital in the organizational equilibrium settles between the Ramsey outcome and the Markov equilibrium.

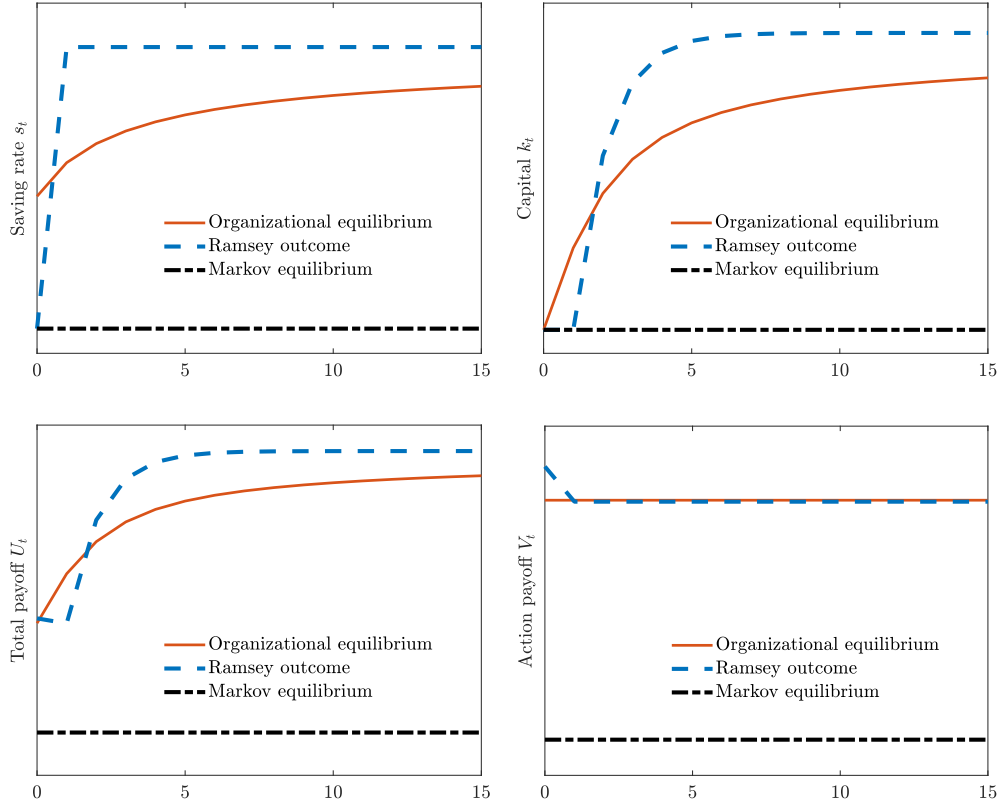
Now we turn to the welfare comparison. Given a particular sequence of saving rates $\{s_\tau\}_{\tau=0}^\infty$, based on the analysis in the last section, the lifetime utility for generation t can be written as

$$U_t(k_t, \{s_\tau\}_{\tau=t}^\infty) = \frac{\alpha(1 - \alpha\beta + \delta\alpha\beta)}{1 - \alpha\beta} \log k_t + V_t.$$

The total payoff U_t and the action payoff V_t are depicted in the second row of Figure 2. The total payoff in the Markov equilibrium is the lowest during the entire transition, which is the result of both the lowest capital stock and action payoff. The comparison between the Ramsey outcome and the organizational equilibrium is more subtle. In the first period, the total payoff in the Ramsey outcome is higher than that in the organizational equilibrium: this has to happen by definition, since the Ramsey outcome maximizes the total payoff from the perspective of period 0. In the following period, the comparison reverses, and the total payoff in the organizational equilibrium is actually higher than the Ramsey outcome. This happens both because the initial generation accumulates additional capital, and because the organizational equilibrium does not impose as high a saving rate, allowing for some indulgence for the short-run impatience that arises in the second period. Our notion of organizational equilibrium treats initial capital as a bygone, factoring it out of the payoff that is relevant in computing the equilibrium itself; however, it captures the notion that the initial agent is not privileged compared to future decision makers and cannot impose on them sacrifices that it has not undertaken. For this reason, when we focus on V_t , an organizational equilibrium redistributes from the initial agent to all future decision makers. When comparing the total payoff, after period 0, early decision makers benefit both from a higher capital level and a higher action

²⁴For the numerical computation, we use the following parameter values: $\alpha = 0.4$, $\beta = 0.96$, and $\delta = 0.9$.

FIGURE 2: Transition Paths in Various Equilibria



payoff, while eventually capital falls below the Ramsey outcome and late generations lose from this.

3.5 Equilibrium of Approximating Economies

The assumption of weakly separable utility is restrictive. In this section, we propose a strategy to study organizational equilibrium for economies where this assumption is not satisfied. Our approach is to look at an economy that is weakly separable and similar in a particular metric to the original one and then study organizational equilibrium in this alternative economy. Two considerations make this a fruitful approach. One is that the approximation, while being first order, is not necessarily linear, but takes any functional form of our choice. The other is that we can solve with arbitrary accuracy using standard methods for the Markov and Ramsey allocations of the original economy. This allows us to explore, if needed, what particular functional forms of the approximating economy works best in the sense of yielding sufficiently close Markov and Ramsey allocations to those of the original economy. Having similar Markov and Ramsey allocations gives a rationale to believe that

the organizational equilibrium allocation of the approximating economy is also close to our object of interest. This strategy has a strong tradition in macroeconomics where little (if anything) is known about recursive equilibrium in distorted economies that do not have a particular functional form. Consequently, the equilibrium is computed for a similar economy in a certain sense (see [Kubler and Schmedders \(2005\)](#) and [Kubler \(2007\)](#) for related discussions).

For brevity, we limit our analysis to a more specific class of economies, in which the action and the state are univariate and a recursive structure is present from the outset. Specifically, we assume the preferences for the agent in period t have the following recursive formulation

$$U(k_t, a_t, a_{t+1}, \dots) = P(k_t, a_t) + Q(k_{t+1}, a_{t+1}) + \beta U(k_{t+1}, a_{t+1}, a_{t+2}, \dots),$$

where U is strictly concave, and, following our previous notation, k_t is the period- t state, a_t is the period- t action, and $k_{t+1} = F(k_t, a_t)$. We also assume that F is weakly concave, and that, for each action a , $F(a, k) - k > (<) 0$ in a neighborhood of $k = 0$ ($k = \infty$). The time inconsistency issue arises as the period utility function depends on future agents' action a_{t+1} . In general, the life-time utility U may not be separable between k_t and the sequence of actions $\{a_{t+\tau}\}_{\tau=0}^{\infty}$.

Consider any stationary point of the mapping F , that is, any point (\bar{k}, \bar{a}) such that $\bar{k} = F(\bar{k}, \bar{a})$.²⁵ We propose a relatively flexible first-order approximation of our original economy around any such stationary point:

$$\begin{aligned} P(k_t, a_t) &= \bar{P} + \frac{\bar{P}_k}{\bar{h}_k} (h(k_t) - h(\bar{k})) + \frac{\bar{P}_a}{\bar{m}_a} (m(a_t) - m(\bar{a})), \\ Q(k_t, a_t) &= \bar{Q} + \frac{\bar{Q}_k}{\bar{h}_k} (h(k_t) - h(\bar{k})) + \frac{\bar{Q}_a}{\bar{m}_a} (m(a_t) - m(\bar{a})), \\ F(k_t, a_t) &= \bar{F} + \frac{\bar{F}_k}{\bar{h}_k} (h(k_t) - h(\bar{k})) + \frac{\bar{F}_a}{\bar{m}_a} (g(a_t) - g(\bar{a})). \end{aligned}$$

That is, we allow flexible function choices when approximating the actions a_t in the preference and the technology, respectively.²⁶ The requirement is that the functions $h(\cdot)$, $m(\cdot)$, and $g(\cdot)$ are monotonic functions. When these functions are linear, we return to the standard linear approximation.

The lifetime utility can therefore be approximated by

$$\tilde{U}(k_t, a_t, a_{t+1}, \dots) = \bar{U} + \frac{\bar{U}_k}{\bar{h}_k} h(k_t) + \sum_{j=0}^{\infty} \frac{\bar{U}_{m,j}}{\bar{m}_a} m(a_{t+j}) + \sum_{j=0}^{\infty} \frac{\bar{U}_{g,j}}{\bar{g}_a} g(a_{t+j}), \quad (14)$$

²⁵Our assumptions about F imply that, for each action a , there exists a unique point $k(a)$ such that $F(k(a), a) = k(a)$.

²⁶The choice of the approximating function for the state k_t is irrelevant as eventually the equilibrium is about the action a_t .

where $\bar{U}_k = \frac{\bar{P}_k + \bar{Q}_k \bar{F}_k}{1 - \beta \bar{F}_k}$ and

$$\begin{aligned}\bar{U}_{m,0} &= \bar{P}_a, & \bar{U}_{g,0} &= \bar{F}_a(\bar{Q}_k + \beta \bar{U}_k) \\ \bar{U}_{m,1} &= \bar{Q}_a + \beta \bar{U}_{m,0}, & \bar{U}_{g,1} &= \beta \bar{U}_{g,0}, \\ \bar{U}_{m,j} &= \beta \bar{U}_{m,j-1}, & \bar{U}_{g,j} &= \beta \bar{U}_{g,j-1}, \quad j > 1.\end{aligned}$$

Clearly, this approximating economy is weakly separable between the state variable k and the sequence of actions, and our notion of organizational equilibrium can be applied. The action payoff can be written as

$$\begin{aligned}V(a_t, a_{t+1}, \dots) &= \sum_{j=0}^{\infty} \frac{\bar{U}_{m,j}}{\bar{m}_a} m(a_{t+j}) + \sum_{j=0}^{\infty} \frac{\bar{U}_{g,j}}{\bar{g}_a} g(a_{t+j}), \\ &= \frac{\bar{U}_{m,0}}{\bar{m}_a} m(a_t) + \frac{\bar{U}_{g,0}}{\bar{g}_a} g(a_t) + \frac{\bar{U}_{m,1} - \beta \bar{U}_{m,0}}{\bar{m}_a} m(a_{t+1}) + \beta V(a_{t+1}, a_{t+2}, \dots).\end{aligned}$$

The remaining question is how to determine the point around which the approximating economy is constructed. The approximating economy satisfies Assumptions 2, 3, and 5, so that Proposition 5 applies. Let a^* denote the choice of action that maximizes the stationary payoff $V(a, a, \dots)$. A natural requirement for the choice of \bar{a} is therefore that it coincides with a^* . Note that

$$(1 - \beta)V(a, a, a, \dots) = \frac{\bar{U}_{m,0} + \bar{U}_{m,1} - \beta \bar{U}_{m,0}}{\bar{m}_a} m(a) + \frac{\bar{U}_{g,0}}{\bar{g}_a} g(a).$$

An interior solution for a^* is only possible if \bar{a} is chosen at the point that satisfies²⁷

$$\bar{P}_a + \bar{Q}_a + \bar{F}_a \left(\frac{\bar{Q}_k + \beta \bar{P}_k}{1 - \beta \bar{F}_k} \right) = 0. \quad (15)$$

Utilizing the recursive structure of the action payoff, the transition path can be then constructed via the following difference equation

$$(1 - \beta)V(\bar{a}, \bar{a}, \dots) = \frac{\bar{U}_{m,0}}{\bar{m}_a} m(a_t) + \frac{\bar{U}_{m,1} - \beta \bar{U}_{m,0}}{\bar{m}_a} m(a_{t+1}) + \frac{\bar{U}_{g,0}}{\bar{g}_a} g(a_t).$$

In Appendix F, we further illustrate this strategy in a non-separable growth economy with quasi-geometric discounting, partial capital depreciation, and CRRA preference. In this example, we set $h(\cdot)$ and $g(\cdot)$ to be log functions and $m(\cdot)$ to be the power function with the same curvature

²⁷Such a point is unique because our original utility function is assumed to be strictly concave. Since we need a compact action set, it is in principle possible that no solution to (15) exists if the \bar{a} is at the lower or upper bound, where the appropriate inequality would need to be imposed. We only consider interior equilibria here for brevity.

as the period utility function. We compare the Markov and the Ramsey solutions obtained in this economy to those obtained in the approximating economy according to our strategy. Relative to the true solution, the capital and the saving rate paths in the approximating economy are quite close, showing that there is no need to search for better behaved approximating functions. The predictions of the organizational equilibrium in this approximating economy are consistent with our aforementioned equilibrium properties.

3.6 Agents as Policy Makers

Up to now, we analyzed environments in which at each point in time t a single player takes an action. Most macroeconomic policy problems are different, in that we often are interested in situations in which there is a single large player (the “government”), but also a continuum of competitive private agents (the “consumers”). This is the case in our applications of Sections 4 and 5. Typically, when policy is endogenously determined, the economy is described as a hybrid of a game and a competitive environment. In this section, we show how we can adapt our equilibrium concept to such an environment.²⁸

We now assume that, given the current level of $k \in K$, the government chooses an action a from a set A , and the consumers choose an action s from the set $s(k) \subseteq S$. The state evolves according to $k' = F(k, a, s)$. Let the preferences for the government in period t be given by $\Psi(k_t, a_t, s_t, a_{t+1}, s_{t+1}, a_{t+2}, s_{t+2}, \dots)$.²⁹

Assumption 6. *Given a sequence of government actions $\mathbf{a} := \{a_t\}_{t=0}^\infty$, there exists a unique competitive equilibrium $\mathbf{s}(\mathbf{a}) := \{s_t(\mathbf{a})\}_{t=0}^\infty$, where the sequence $\mathbf{s}(\mathbf{a})$ is independent of the state k_0 .*

Assumption 6 plays two roles. First, the uniqueness allows us to define government preferences directly over the sequence of government actions, taking as given that households will play the associated competitive equilibrium. Second, the fact that \mathbf{s} is independent of the initial state extends the weak separability requirement that is at the heart of our method. We can then define the government’s preferences over sequences of actions as

$$U(k, a_t, a_{t+1}, a_{t+2}, \dots) := \Psi(k, a_t, s_t(\mathbf{a}), a_{t+1}, s_{t+1}(\mathbf{a}), a_{t+2}, s_{t+2}(\mathbf{a}), \dots), \quad (16)$$

²⁸The setup here uses ideas from Stokey (1991).

²⁹As in Section 3, sometimes it may be necessary to transform the original government action so that it is feasible independently of the choices of the private sector and the current level of the physical state, and so that the desired separability property of preferences emerges. A similar re-scaling may be needed for the household choices.

where for $t \geq 0$, k_t is computed recursively as

$$k_{t+1}(k) = F(k_t(k), a_t, s_t(\mathbf{a})). \quad (17)$$

These preferences now take the same form as in the one-agent case, so once again we impose Assumptions 1 and 2, and we define an organizational equilibrium as in Definition 1. While the definition of an organizational equilibrium is the same in terms of sequences of actions, its connection to symmetric subgame-perfect equilibria of an underlying game is slightly different, due to the presence of competitive households that act in anticipation of the government’s future actions. In Appendix G, we adapt the analysis of Section 3 to this new environment and we provide an explicit game that describes the strategic interaction between the government and the households over time.³⁰

The interaction of private and government decisions is such that Assumptions 3 and 5 fail in our applications. Fortunately, Proposition 3 provides an alternative way of proving existence, that we adopt in what follows.

In Section 3.1 and 3.2, an organizational equilibrium represents a refinement of a subgame perfect equilibrium based on specific beliefs that the single player at each stage entertains about future play. In the richer environment considered here, coordination of beliefs involves both the government and a continuum of private players. It is natural for this coordination to take the form of institutions and laws, which is why we call ours an “organizational equilibrium.” Nonetheless, it is important to contrast this role of institutions as purely coordinating expectations from an alternative, in which they represent forms of commitment. As in Prescott and Rios-Rull (2005), we take the view here that laws can be freely changed ex post and that government agencies can be reformed, so that they do not represent effective forms of commitment, and show how cooperation across different players over time can still be sustained, even when the self-interest of future players rules out the usual grim-trigger strategies.

4 Climate Change Mitigation

The discussion around policies on climate change has often been cast in terms of intergenerational fairness and institutional agreements that support better policies. Moreover, such policy formation has taken a gradual approach of successive rounds of international negotiation that led to increasingly ambitious goals over time, from the UN Framework Convention on Climate Change in 1992, to the Kyoto Protocol in 2005 and the Paris agreement in 2015 and even to the recent agreement to end fossil

³⁰In the application that we will describe shortly, we assume that the government is a first mover within each period, so that households react contemporaneously to a government deviation. The definition could be adapted to environments where the opposite timing prevails, as in Ortigueira (2006).

fuels in The United Nations Climate Change Conference (COP28) of December 2023. Undoubtedly there are many other factors that shape the evolution of the carbon tax. Nevertheless, we think that our theory captures well such features and delivers a natural way of thinking about the reason agreements are gradually ratcheting up emission standards rather than achieving an ultimate goal in a big-bang fashion.

Environment We follow [Goloso et al. \(2014\)](#) when specifying the preferences, the production technology, and the carbon evolution dynamics, with the exception that we allow quasi-geometric discounting which captures the idea that short-sighted politicians may have a bias in favor of the current generations. Assume that the objective of the policy maker at period t is given by

$$\log(c_t) + \delta \sum_{j=1}^{\infty} \beta^j \log(c_{t+j}),$$

where $\delta \in (0, 1]$ accommodates the idea that the policy maker is facing the conflict of interest between different generations as in [Jackson and Yariv \(2014, 2015\)](#).

The production of the final goods requires energy input e_t in addition to standard inputs capital k_t and labor n_t .

$$c_t + k_{t+1} = y_t = \exp(-\gamma q_t) k_t^\alpha n_t^{1-\alpha-\nu} e_t^\nu.$$

The use of the energy input generates carbon in the atmosphere, which negatively affects production efficiency. We denote q_t as the amount of carbon accumulated in the atmosphere. The parameters α and ν control the relative importance of various production inputs, and γ determines the damage of carbon on the economy.

Energy is produced through labor alone:

$$e_t = A(1 - n_t),$$

where A stands for the productivity in the energy sector, and we have normalized the total labor supply to be 1. We interpret the fossil fuel energy as coal. As discussed in [Goloso et al. \(2014\)](#), coal is the most important type of fossil fuel energy that calls for government intervention, and it is practically in unlimited supply.

The evolution of the stock of carbon is driven by the sum of a permanent component q_{1t} and a

persistent component q_{2t} : $q_t = q_{1t} + q_{2t}$, where

$$q_{1t} = q_{1t-1} + \varphi_L e_t, \quad \text{and} \quad q_{2t} = \varphi q_{2t-1} + (1 - \varphi_L) \varphi_0 e_t.$$

That is, a fraction φ_L of the newly emitted carbon e_t stays in the atmosphere permanently, a fraction $\varphi_L(1 - \varphi_0)$ of e_t exits the atmosphere immediately, and the rest stays in the atmosphere but decays at a geometric rate φ .

Consider the following cost-benefit analysis on the allocation of labor between the final goods and energy production:

$$\frac{\nu}{A} \frac{y_t}{1 - n_t} = \frac{1 - \alpha - \nu}{A} \frac{y_t}{n_t} + y_t \Lambda_t.$$

The left-hand side corresponds to the marginal product of e_t after a marginal increase of labor in the energy sector. The first-term on the right-hand side is the marginal cost due to the reduction of labor in the final goods sector, and the second term represents the associated environmental cost expressed in terms of current output. The variable Λ_t can therefore be interpreted as the carbon tax in a decentralized economy.

A notable feature of this environment is that it is separable between the actions (saving rate s_t and labor allocation to final goods production n_t) and the state variables (capital stock k_t and carbon in the atmosphere q_{1t}, q_{2t}). Given a sequence of saving rates $\{s_0, s_1, \dots\}$, a sequence of labor allocation $\{n_0, n_1, \dots\}$, and the initial state $(k_0, q_{1,-1}, q_{2,-1})$, the payoff of the policy maker at time 0 can be expressed as

$$U_0 = G(k_0, q_{1,-1}, q_{2,-1}) + W(s_0, s_1, \dots) + V(n_0, n_1, \dots),$$

where the part involving the allocation of labor on energy production can be expressed as³¹

$$\begin{aligned} V(n_0, n_1, \dots) = & \gamma A \left(1 + \frac{\delta \chi \beta}{1 - \chi \beta} \right) (1 - n_0) + (1 - \alpha - \nu) \log n_0 + \nu \log(1 - n_0) \\ & - \delta \beta \left(\frac{\gamma A}{1 - \beta \chi} \sum_{j=0}^{\infty} \beta^j (1 - n_{j+1}) + (1 - \alpha - \nu) \sum_{j=0}^{\infty} \beta^j \log n_{j+1} + \nu \sum_{j=0}^{\infty} \beta^j \log(1 - n_{j+1}) \right) \end{aligned}$$

³¹The expression for W is given by

$$W(s_0, s_1, \dots) = \log(1 - s_0) + \frac{\delta \alpha \beta}{1 - \alpha \beta} \log(s_0) + \delta \sum_{j=1}^{\infty} \beta^j \left(\log(1 - s_j) + \frac{\alpha \beta}{1 - \alpha \beta} \log(s_j) \right).$$

Note that the optimal carbon tax is independent of the sequence of saving rates.

Note that once the sequence of labor is specified, the implied carbon tax Λ_t is also determined. Therefore, to understand the dynamics of the carbon tax Λ_t , it is sufficient to characterize the equilibrium allocation of labor.

Steady-state analysis Thanks to the separability property, it is straightforward to apply our theoretical results to derive the allocation in the organizational equilibrium. First, we compare carbon taxes in the long-run steady state. These are given by

$$\left\{ \begin{array}{l} \text{Ramsey outcome:} \\ \text{Markov equilibrium:} \\ \text{Organizational equilibrium:} \end{array} \right. \Lambda^R = \gamma \left(\frac{\varphi_L}{1-\beta} + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta} \right)$$

$$\Lambda^M = \gamma \left(\frac{\varphi_L}{1-\beta} (1-\beta + \delta\beta) + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta} (1-\varphi\beta + \delta\varphi\beta) \right)$$

$$\Lambda^O = \gamma \left(\frac{\varphi_L}{1-\beta} \left(1-\beta + \frac{\delta\beta}{1-\beta(1-\delta)} \right) + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta} \left(1-\varphi\beta + \frac{\delta\varphi\beta}{1-\beta(1-\delta)} \right) \right),$$

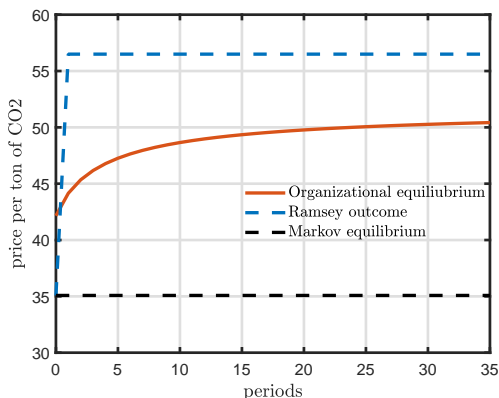
with $\Lambda^R > \Lambda^O > \Lambda^M$ when $\delta \in (0, 1)$.

Let us unpack these expressions. Starting from the Ramsey outcome, $\gamma \frac{\varphi_L}{1-\beta} + \gamma \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta}$ simply represents the present value of the damage after one unit of carbon emission, taking into account the time path of the carbon diffusion. This formula is identical to that derived in [Goloso et al. \(2014\)](#). In the Markov economy, the additional discounting resulting from $\delta < 1$ induces the policy maker to behave in a more myopic way, as captured by the term $1 - \beta + \delta\beta < 1$ and $1 - \varphi\beta + \delta\varphi\beta < 1$, respectively. Finally, in the organizational equilibrium, each policy maker directly takes into account the welfare of all their subsequent policy makers as their action payoff needs to be equalized, which makes them more patient than in the Markov equilibrium. It follows that the steady state carbon tax for the organizational equilibrium lies in the middle of the Ramsey outcome and the Markov equilibrium.

Transition dynamics In addition to the difference in terms of steady-state carbon tax, the organizational equilibrium also predicts an interesting dynamic pattern that sheds light on the evolution of the carbon tax. To proceed, we set the parameters related to production technology and environmental damages as in [Goloso et al. \(2014\)](#), and set $\delta = 0.95$ at annual frequency to allow the time inconsistency to arise.

The solid line in [Figure 3](#) displays the transition path of the implied carbon tax Λ_t under the organizational equilibrium, which features a gradually increasing path. A carbon tax acts as an investment, as it reduces current productivity, but it increases future productivity. As in our benchmark consumption-saving problem, in an organizational equilibrium policymakers overcome the temptation to ignore the carbon externality only slowly and partially. In contrast, such implicit coordination

FIGURE 3: Transition Path of Carbon Tax



Notes: The calibration of $\{\beta, \gamma, \varphi, \varphi_L, \varphi_0\}$ follows Golosov et al. (2014). The additional discounting δ at annual frequency correspond to 0.95. The prices are expressed in terms of 2010 world nominal GDP.

is absent in the Markov equilibrium, and the carbon tax remains constant at a lower level. In the Ramsey outcome, the initial policy maker enjoys the privilege and sets a low tax level initially, but a sudden jump to a high level follows immediately after that.

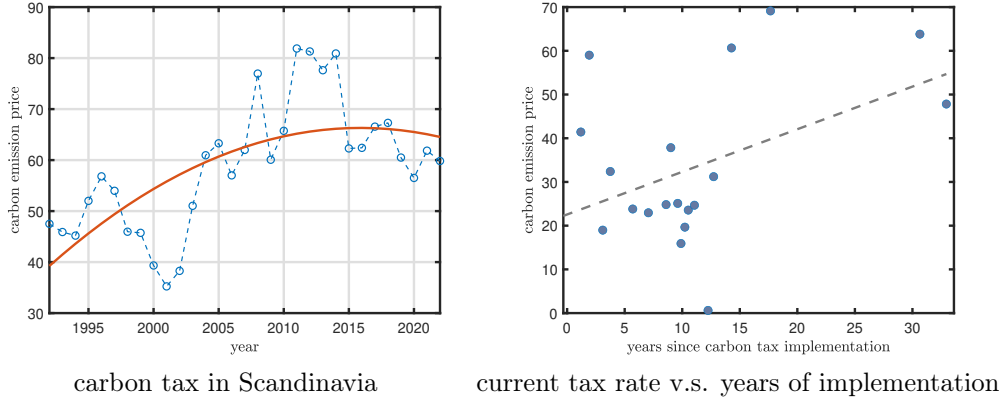
The gradualism resulting from the trade-off between current and future policy makers in the organizational equilibrium speaks to the observed evolution in the data.³² The left panel of Figure 4 presents the average carbon tax rates in Scandinavian countries over time, which are the first group of countries that implemented the carbon tax in the world. Instead of featuring a steady level or a disruptive jump, the carbon tax rates in these countries display a gradual increasing trend, a pattern that is broadly consistent with the predictions of the organizational equilibrium.

We also explore how the current carbon tax rates across countries vary with the number of years since the implementation of the carbon tax. *Ceteris paribus*, the carbon tax rate is increasing in the years of implementation in an organizational equilibrium. After controlling for the quantity of carbon emission covered by the tax and the GDP per capita, we find that the carbon tax rates are positively correlated with the years of implementation, which lends additional credibility to our theory. The right panel of Figure 4 displays the corresponding binscatter plot.

Another important case is the gradual change of the climate policy in China, which has the most GHG emissions and just started Emissions Trading System in 2021. The trade-off between current versus future generations is directly revealed from speeches by their successive presidents. In 2008,

³²The main data source for the carbon tax is the Carbon Pricing Dashboard constructed by the World Bank.

FIGURE 4: Evolution of Carbon Taxes



Notes: In the left panel, the blue dots correspond to the average carbon price in Denmark, Finland, Norway, and Sweden. The original nominal carbon prices are converted to real prices using country specific GDP deflator with 2010 as the base year prices. The red solid line repents a quadratic fit of the data over time. In the right panel, we consider the following regression: $\tau_j = \beta_0 + \beta_1 t_j + \text{controls} + \text{residuals}$, where τ_j is country/region j 's carbon tax rate in 2023, t_j is the number of years since implementation, and we include the carbon tax covered CO₂ emission and the GDP per capita as controls. The estimated coefficient β_1 is 1.11 with p value 0.017.

the former president Hu Jintao stated that “China is a developing country in the process of industrialization and modernization... China’s central task now is to develop the economy and make life better for the people.” In 2020, the current president Xi Jinping stated that “China will increase its nationally determined contributions, adopt more powerful policies and measures, strive to reach the peak of carbon dioxide emissions by 2030, and strive to achieve carbon neutrality by 2060.”³³ These statements directly reflect the conflict of interest between current and future generations and signals the incentive of delaying the process. Our theory provides a natural solution that balances the interests across generations via a gradual reform.

5 Tariff Policies

Similar to the case of climate change, international institutions to foster free trade have evolved gradually. The GATT (General Agreement on Tariffs and Trade) evolved through several rounds of negotiations, and eventually became the World Trade Organization. Through each of these rounds, the number of goods covered by agreements was gradually expanded, and agreements covered an increasingly broad set of protectionist policies. We describe here a simple model that captures these features well. In this model, policymakers trade off a short-run cost of free trade, arising from distributional losses across different types of workers, with a long-run benefit arising from faster

³³These are quoted from Hu Jintao’s remarks at G8 Outreach Session on July 9, 2008, and from Xi Jinping’s speech at the United Nations General Assembly on September 22, 2020, respectively.

growth.

Environment The world consists of two symmetric countries, each of which produces two types of tradable intermediate goods. We mainly describe the environment in the home country, and we focus on the symmetric equilibrium.

The problem of a representative firm that produces intermediate good i is

$$\max_{\ell_{it}, k_{it}} p_{it} A_i k_t^{1-\alpha} \ell_{it}^{1-\alpha} k_{it}^\alpha - w_{it} \ell_{it} - r_t k_{it},$$

where ℓ_{it} and k_{it} are capital and labor used by firm i , and k_t is aggregate capital. The production efficiency in sector 1 is higher than that in sector 2, $A_1 = A > A_2 = 1$, while in the foreign country, the opposite prevails. That is, country 1 has comparative advantage in producing good 1, and country 2 has comparative advantage in producing good 2.

We assume that there is a measure one of workers in each sector and there is no labor mobility across sectors. This assumption of fixed labor supply captures the idea that it is difficult for workers to reallocate across sectors.³⁴ As a consequence, wage rates in the two sectors may differ. Capital is free to flow across sectors.³⁵ Aggregate capital provides a positive externality on the production efficiency, captured by $k_t^{1-\alpha}$, generating the potential for endogenous growth.

The final good that is used for investment and consumption is produced by aggregating the intermediate goods via a CES function with associated aggregate price index p_t

$$y_t = [0.5^{1-\rho} m_{1t}^\rho + 0.5^{1-\rho} m_{2t}^\rho]^{\frac{\rho-1}{\rho}}, \quad p_t = \left[0.5 p_{1t}^{\frac{\rho}{\rho-1}} + 0.5 p_{2t}^{\frac{\rho}{\rho-1}} \right]^{\frac{\rho-1}{\rho}}.$$

We normalize the price of goods 2 in the home country to be 1, $p_{2t} = 1$. In a free trade economy without tariffs, symmetry implies $p_{1t} = p_{2t} = 1$, with country 1 exporting some of good 1 in exchange for some of country 2's production of good 2. When the government sets the tariff rate to be τ_t on imported goods, the price of goods 1 in the home country becomes $p_{1t} = \frac{1}{1+\tau_t}$. A higher tariff lowers the relative price of the good 1 in country 1, since it prevents it from fully exploiting its comparative advantage.

Given the total capital k_t and the tariff τ_t in both home and foreign countries, it is straightforward to solve for the static allocation and prices. In the static symmetric equilibrium, the wage rates and

³⁴We could relax this assumption and state that workers are only immobile for one period, generating even stronger time inconsistency.

³⁵This assumption is also not essential.

the interest rate are given by

$$w_{1t} = (1 - \alpha)(1 + \tau_t)^{-1}A(1 - \phi_t)^\alpha k_t, \quad w_{2t} = (1 - \alpha)\phi_t^\alpha k_t, \quad r_t = \alpha\phi_t^{\alpha-1}, \quad (18)$$

where ϕ_t is the fraction of capital allocated in sector 2, $\phi_t = \left(1 + \left(\frac{A}{1+\tau_t}\right)^{\frac{1}{1-\alpha}}\right)^{-1}$.

With a reduction of the tariff rate (both in the home and foreign countries), the exporting sector expands, resulting a lower ϕ_t . Accordingly, the wage rate in sector 1 is higher, and the wage rate in sector 2 is lower. The reduction in importing sector's wage rate provides a rationale for protectionism. Overall, a lower tariff rate improves production efficiency, and the return to capital, r_t , is higher.

Capitalists are responsible for the capital accumulation. Their problem is

$$\max \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$c_t + k_{t+1} = \left(\frac{r_t}{p_t} + 1 - \delta\right) k_t.$$

The capitalists prefer a low tariff as the implied interest rate will be higher. Let s_t denote capitalists' saving rate. The Euler condition can be expressed as

$$\left(\frac{1-s_t}{s_t}\right)^{-\sigma} = \beta \left(\frac{r_{t+1}}{p_{t+1}} + 1 - \delta\right)^{1-\sigma} (1-s_{t+1})^{-\sigma} \quad (19)$$

We focus on the case in which saving is increasing in the interest rate (and thus decreasing in the tariff); this happens when the intertemporal elasticity of substitution (σ^{-1}) is larger than unity. Note that the real interest rate is only a function of the tariff rate.

The policy maker in period t attempts to maximize an average worker's welfare in the economy, which is given by

$$U_t \equiv \sum_{k=0}^{\infty} \beta^k (\log c_{1t} + \log c_{2t}).$$

Workers live hand to mouth, consuming their labor income w_{it} and the tariff revenues that are rebated to them. In the static equilibrium, the consumption levels are proportional to the capital shock k_t

$$\log c_{1t} + \log c_{2t} = \chi(\tau_t) \log k_t,$$

where $\chi(\tau_t)$ only depends on the tariff and the formula is specified in [Appendix I](#).

From the policymaker's perspective, a tariff benefits workers in sector 2 in the short run, shielding them from the foreign competition. The utilitarian welfare function makes a tariff desirable in the short run, since workers in sector 2 earn a lower wage under free trade. For the future, the expectation of a low tariff encourages capital accumulation, benefiting all workers. As a result, under commitment a policymaker would always prefer setting relatively higher tariffs initially and lower tariffs down the road. This is the root of time inconsistency.

Equilibrium allocation The economy is separable between capital stock k_t and the tariff rate τ_t . Substituting the wage rates from the static equilibrium, we obtain that, given an initial capital stock k_0 , the payoff to the policymaker from a sequence of tariff rates $\{\tau_t\}$ is

$$U_0 = \frac{1}{1-\beta} \log k_0 + \sum_{t=0}^{\infty} \beta^t \chi(\tau_t) + \frac{\beta}{1-\beta} \sum_{t=0}^{\infty} \beta^t \left(\log s_t + \log \left(\frac{r_t}{p_t} + 1 - \delta \right) \right).$$


The second term related to $\chi(\tau_t)$ captures the short-run consequence of tariffs which protect the importing sector. The third term related to s_t and r_t captures the long-run consequence of tariffs which discourage investments by capitalists.

Note that the current saving rate is shaped by future tariff policies according to the Euler condition [\(19\)](#). In a Markov equilibrium, the policy maker will take tariff rates in the future as given. When setting the current tariff rate, the dependence of the saving rate on the tariff rate will be ignored. As a result, the policy maker fails to internalize the benefit of lowering tariffs and inducing more capital in the future, and leans towards protectionism. In contrast, in the organizational equilibrium, the action payoffs of policy makers in different periods are equalized. The forward-looking behavior allows the policy maker to partially internalize the long-run benefit of free trade. As a result, the steady-state tariff rate in the organizational equilibrium is lower than that in the Markov equilibrium.

Transitional dynamics Now we turn to the dynamic pattern of tariffs. It is well documented that there is a decline of world tariff over time in the process of globalization. [Figure 5](#) displays the world average tariff from 1950 to present, which gradually decreases from around 16% to 4%. There is a large literature that speaks to this gradual change. For example, in [Bond and Park \(2002\)](#), gradualism comes from the asymmetry between countries: the efficient allocation specifies a rising payoff for the country that has an initial binding incentive constraint. [Chisik \(2003\)](#) considers an environment with specialization and capacity irreversibility in the development of partner-specific capital which increase the benefit of further liberalization and the penalty of defection. This dynamic decreases the lowest self-enforcing tariff over time. Instead of the contract enforcement issue, [Zissimos](#)

(2007) takes a more applied view, and shows directly how specific rules in the GATT (General Agreement of Tariffs and Trade) created strategic incentives for gradual liberalization.

FIGURE 5: World Average Tariff



Notes: This tariff series is the unweighted world average tariff from Figure 8 in [Antràs \(2020\)](#).

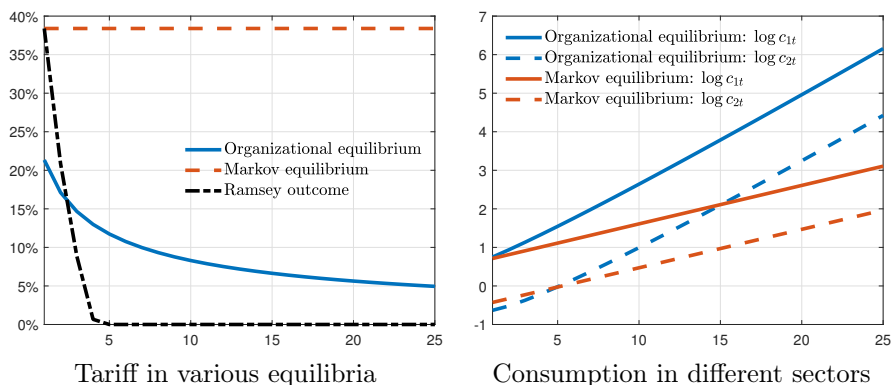
Our theory provides an alternative rationale for the observed gradual decline, which is based on the intertemporal trade-off between the short-run welfare of workers that calls for protectionism and the long-run welfare of workers that benefits from development resulting from open trade. Different equilibrium notions we have considered put different weights on this intertemporal trade-off. To illustrate the behavior of the model, we set the duration of a model period to 7 years and choose $\alpha = 0.33$, $\beta = 0.7$, $\delta = 0.65$. The elasticity of substitution is set to be 1.1, that lies in the middle of various empirical estimates.³⁶ Finally, we set $A = 3.5$ so that the average level of tariff in the organizational equilibrium is comparable to that in the data.

The left panel of Figure 6 displays the transition paths of tariff. Starting from the Markov equilibrium, it features a constant high level of tariff rate as policy makers behave in the most myopic way. In the Ramsey outcome, the policy maker sets a high level of tariff initially but quickly eliminates it afterwards. With organizational equilibrium, the tariff starts from a level in the middle of the previous two cases and gradually converges to a low level of tariff rate, which is broadly consistent with the observed pattern in the data.

The right panel of Figure 6 further illustrates the aforementioned trade-off by comparing the consumption paths that start with the same initial capital stock. It is interesting to inspect the consumption of workers in sector 2 who are subject to the competition with imported goods. In the

³⁶See [Boehm et al. \(2023\)](#) for a more detailed discussion.

FIGURE 6: Transition Paths of Tariff and Consumption



Notes: The left panel displays the transition paths of tariff rates in the three different equilibria. The right panel compares the consumption paths for workers in the exporting sector (c_{1t}) and importing sectors (c_{2t}) under the organizational equilibrium and the Markov equilibrium, respectively.

Markov equilibrium, workers enjoys more consumption initially than the organizational equilibrium thanks to the higher level of tariff. However, in the long run, a lower tariff in the organizational equilibrium induces a higher allocation efficiency and encourages a faster capital growth rate. The consumption level eventually exceeds that of the Markov equilibrium.

While introducing shocks is beyond the scope of this paper, we conjecture that our setup could be well suited to analyze what happens following periods of temporary disruption in international cooperation. To explain why, we consider a tremble, in which for some reason in period t governments unexpectedly raise tariffs to some value $\tilde{\tau}$ rather than progressing along the path of liberalization. In an organizational equilibrium, there is no need for a grim-trigger strategy to punish such a deviation; rather, it is sufficient to revert to the path that starts from $\tilde{\tau}$ on. We view this as a potentially useful way of studying the current environment in international trade, where the reimposition of tariffs have not led to all-out trade war, but have set cooperation back, with negotiations restarting to undo the newly introduced barriers.

6 Conclusion

Rome was not built in a day. Likewise, many institutions that underpin solid policy evolved slowly. [Sargent \(2017\)](#), and [Hall and Sargent \(2014; 2015; 2018\)](#) describe the way the United States acquired a reputation for honoring its debt over time. The current environment of relatively low inflation emerged after the tumultuous 70s, during which governments gradually learned how to manage

monetary policy without resorting to the anchor of a commodity standard. Social security started as a narrow and limited program and only later grew in size and scope.

In this paper, we provided an equilibrium concept that is well suited to analyzing such situations. The equilibrium eschews abrupt transitions and is not (or at least not necessarily) supported by grim triggers, but rather cooperation for the common long-term good evolves slowly and would potentially also erode slowly. While the constraints that we impose on equilibrium strategies appear very restrictive, what is most interesting to us is that in our computed examples, they still permit very good outcomes, bringing substantial improvement to the dismal predictions of Markov equilibria. At the same time, the notion of organizational equilibrium allows for clear-cut comparative statics exercises and does not suffer from the pervasive multiplicity of subgame-perfect equilibria, which is implied by the folk theorem. Because of this, it is more readily amenable to empirical analysis.

For future research, our equilibrium concept can be applied to understand certain sociopolitical phenomena. Two fitting examples are the extension of civil rights in the U.S. after the civil war and the gradual creation of the European union. In these situations, a common good (such as the institution) has to be provided but different states or parties may have different time horizons. The aggregated social welfare function is necessarily time inconsistent ([Jackson and Yariv, 2014](#)). Through the lens of our equilibrium, a good institution will gradually emerge (unlike in a Markov equilibrium), but will remain imperfect (worse than Ramsey).

References

- Abreu, Dilip, David Pearce, and Ennio Stacchetti**, “Optimal Cartel Equilibria with Imperfect Monitoring,” *Journal of Economic Theory*, 1986, *39*, 251–269.
- , —, and —, “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 1990, *58* (5), 1041–1063.
- Acemoglu, Daron and James A. Robinson**, “Why Did the West Extend the Franchise? Democracy, Inequality and Growth in Historical Perspective,” *Quarterly Journal of Economics*, 2000, *115* (4), 1167–1199.
- , **David Autor, David Dorn, Gordon H Hanson, and Brendan Price**, “Import competition and the great US employment sag of the 2000s,” *Journal of Labor Economics*, 2016, *34* (S1), S141–S198.
- , **Georgy Egorov, and Konstantin Sonin**, “Dynamics and Stability of Constitutions, Coalitions, and Clubs,” *American Economic Review*, 2012, *102* (4), 1446–76.
- , —, and —, “Political Economy in a Changing World,” *Journal of Political Economy*, 2015, *123* (5), 1038–1086.
- Ales, Laurence and Christopher Sleet**, “Revision proofness,” *Journal of Economic Theory*, 2014, *152* (C), 324–355.
- Antràs, Pol**, “De-globalisation? Global value chains in the post-COVID-19 age,” Technical Report, National Bureau of Economic Research 2020.
- Asheim, Geir B.**, “Individual and Collective Time-consistency,” *Review of Economic Studies*, 1997, *64* (3), 427–443.
- Autor, David H, David Dorn, and Gordon H Hanson**, “The China syndrome: Local labor market effects of import competition in the United States,” *American economic review*, 2013, *103* (6), 2121–2168.
- Azzimonti, Marina**, “Barriers to Investment in Polarized Societies,” *American Economic Review*, August 2011, *101* (5), 2182–2204.
- Bassetto, Marco**, “Equilibrium and Government Commitment,” *Journal of Economic Theory*, 2005, *124* (1), 79–105.
- , “Political Economy of Taxation in an Overlapping-Generations Economy,” *Review of Economic Dynamics*, 2008, *11* (1), 18–43.
- and **Thomas J. Sargent**, “Politics and Efficiency of Separating Capital and Ordinary Government Budgets,” *The Quarterly Journal of Economics*, 2006, *121* (4), 1167–1210.
- Bernheim, B Douglas, Debraj Ray, and Şevin Yeltekin**, “Poverty and Self-Control,” *Econometrica*, 2015, *83* (5), 1877–1911.

- Boehm, Christoph E, Andrei A Levchenko, and Nitya Pandalai-Nayar**, “The Long and Short (Run) of Trade Elasticities,” *American Economic Review*, 2023, 113 (4), 861–905.
- Bond, Eric W and Jee-Hyeong Park**, “Gradualism in trade agreements with asymmetric countries,” *The Review of Economic Studies*, 2002, 69 (2), 379–406.
- Brendon, Charles and Martin Ellison**, “Time-Consistently Undominated Policies,” 2018. Unpublished Manuscript, Oxford University.
- Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro**, “Trade and labor market dynamics: General equilibrium analysis of the china trade shock,” *Econometrica*, 2019, 87 (3), 741–835.
- Cao, Dan and Ivàn Werning**, “Saving and Dissaving with Hyperbolic Discounting,” *Econometrica*, 2018, 86 (3), 805–857.
- Chari, V. V. and P. J. Kehoe**, “Sustainable Plans,” *Journal of Political Economy*, 1990, 98 (4), 784–802.
- Chatterjee, Satyajit and Burcu Eyigungor**, “Continuous Markov equilibria with quasi-geometric discounting,” *Journal of Economic Theory*, 2016, 163 (C), 467–494.
- Chisik, Richard**, “Gradualism in free trade agreements: a theoretical justification,” *Journal of International Economics*, 2003, 59 (2), 367–397.
- Cohen, Daniel and Philippe Michel**, “How Should Control Theory Be Used to Calculate Time-Consistent Government Policy?,” *Review of Economic Studies*, 1988, 55 (2), 263–274.
- Currie, David A. and Paul L. Levine**, *Rules, Reputation and Macroeconomic Contracts*, Cambridge University Press, Cambridge, 1993.
- Farrell, Joseph and Eric Maskin**, “Renegotiation in repeated games,” *Games and Economic Behavior*, 1989, 1 (4), 327–360.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski**, “Optimal taxes on fossil fuel in general equilibrium,” *Econometrica*, 2014, 82 (1), 41–88.
- Halac, Marina and Pierre Yared**, “Fiscal rules and discretion under persistent shocks,” *Econometrica*, 2014, 82 (5), 1557–1614.
- and —, “Fiscal Rules and Discretion under Self-Enforcement,” Working Paper 23919, NBER 2017.
- Hall, George J. and Thomas J. Sargent**, “Fiscal Discrimination in Three Wars,” *Journal of Monetary Economics*, 2014, 61, 148–166.
- and —, “A History of U.S. Debt Limits,” Working Paper 21799, NBER 2015.
- and —, “Brief History of US Debt Limits Before 1939,” *Proceedings of the National Academy of Sciences*, 2018, 115 (12), 2942–2945.

- Jackson, Matthew O. and Leeat Yariv**, “Present Bias and Collective Dynamic Choice in the Lab,” *American Economic Review*, 2014, 104 (12), 4184–4204.
- and —, “Collective Dynamic Choice: The Necessity of Time Inconsistency,” *American Economic Journal: Microeconomics*, 2015, 7 (4), 150–78.
- Jovanovic, Boyan and Robert W. Rosenthal**, “Anonymous Sequential Games,” *Journal of Mathematical Economics*, 1988, 17 (1), 77–87.
- Klein, Paul and José-Víctor Ríos-Rull**, “Time-Consistent Optimal Fiscal Policy,” *International Economic Review*, 2003, 44 (4), 1217–1246.
- , **Per Krusell**, and **José-Víctor Ríos-Rull**, “Time-Consistent Public Policy,” *Review of Economic Studies*, 2008, 75 (3), 789–808.
- , **Vincenzo Quadrini**, and **José-Víctor Ríos-Rull**, “Optimal Time-Consistent Taxation with International Mobility Of Capital,” *Advances in Macroeconomics*, 2005, 5 (1). Article 2.
- Kocherlakota, Narayana R.**, “Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency,” *Games and Economic Behavior*, July 1996, 15 (1), 33–54.
- Krusell, P. and José-Víctor Ríos-Rull**, “Vested Interests in a Positive Theory of Stagnation and Growth,” *The Review of Economic Studies*, 1996, 63 (215), 301–331.
- Krusell, Per and Anthony A. Smith**, “Consumption-Savings Decisions with Quasi-Geometric Discounting,” *Econometrica*, 2003, 71 (1), 365–375.
- , **Burhanettin Kuruşçu**, and **Anthony A. Smith**, “Temptation and Taxation,” *Econometrica*, November 2010, 78 (6), 2063–2094.
- Kubler, Felix**, “Approximate Generalizations and Computational Experiments,” *Econometrica*, 07 2007, 75 (4), 967–992.
- and **Karl Schmedders**, “Approximate versus Exact Equilibria in Dynamic Economies,” *Econometrica*, 2005, 73 (4), 1205–1235.
- Laibson, David**, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, pp. 443–477.
- Maggi, Giovanni and Andrés Rodríguez-Clare**, “A Political-Economy Theory of Trade Agreements,” *American Economic Review*, 2007, 97 (4), 1374–1406.
- Martin, Fernando M.**, “On the joint determination of fiscal and monetary policy,” *Journal of Monetary Economics*, March 2011, 58 (2), 132–145.
- Matsuyama, Kiminori**, “Perfect Equilibria in a Trade Liberalization Game,” *The American Economic Review*, 1990, 80 (3), 480–492.

- Nozawa, Wataru**, “On Nonexistence of Reconsideration-proof Equilibrium with State Variables,” *Theory and Decision*, 2018, *85* (2), 253–273.
- Ortigueira, Salvador**, “Markov-Perfect Optimal Taxation,” *Review of Economic Dynamics*, January 2006, *9* (1), 153–178.
- Phelan, Christopher and Ennio Stacchetti**, “Sequential Equilibria in a Ramsey Taxes Model,” *Econometrica*, 2001, *69* (6), 1491–1518.
- Phelps, Edmund S. and Robert A. Pollak**, “On Second-best National Saving and Game-equilibrium Growth,” *The Review of Economic Studies*, 1968, *35* (185-199).
- Pierce, Justin R and Peter K Schott**, “The surprisingly swift decline of US manufacturing employment,” *American Economic Review*, 2016, *106* (7), 1632–1662.
- Pigullem, Facundo and Alessandro Riboni**, “Spending-Biased Legislators: Discipline Through Disagreement,” *Quarterly Journal of Economics*, 02 2015, *130* (2), 901–949.
- and —, “Fiscal Rules as Bargaining Chips,” *Review of Economic Studies*, 2020, *88* (5), 2439–2478.
- Prescott, Edward C. and Jose-Victor Rios-Rull**, “On Equilibrium for Overlapping Generations Organizations*,” *International Economic Review*, November 2005, *46* (4), 1065–1080. doi: 10.1111/j.1468-2354.2005.00360.x.
- Sargent, Thomas J.**, “Honoring Public Debts,” http://www.tomsargent.com/research/Sargent_debt_BFI.pdf 2017. Mimeo, New York University.
- Stokey, Nancy L.**, “Credible Public Policy,” *Journal of Economic Dynamics and Control*, 1991, *15*, 627–656.
- Strotz, Robert H.**, “Myopia and Inconsistency in Dynamic Utility Maximization,” *The Review of Economic Studies*, 1956, *23*, 165–80.
- Zissimos, Ben**, “The GATT and gradualism,” *Journal of International Economics*, 2007, *71* (2), 410–433.

Appendix for Online Publication

A Proofs of Section 3.1	2
A.1 Proposition 6	2
A.2 Proof of Proposition 1	2
B Proofs and Details in Section 3.2	3
C Further Discussion of Assumptions 1, 3, 4, and 5.	9
D Proofs of Section 3.3.	12
D.1 Proof of Proposition 3	12
D.2 Proof of Proposition 4.	13
D.3 Proof of Proposition 5	16
D.4 Proof of Corollary 1	19
E Proofs for Section 3.4	19
F Example of Approximating Strategy	20
G Organizational Equilibrium in Policy Problems	22
H Details for Section 4	25
I Details for Section 5	27

A Proofs of Section 3.1

A.1 Proposition 6

Proposition 6. *Under Assumptions 1 and 2, there exists a subgame-perfect equilibrium of the game that satisfies Requirements 1 and 2.*

Proof. Assumption 2 follows Kocherlakota (1996), who uses it in Proposition 4 to prove that a reconsideration-proof equilibrium exists for the game whose period- t payoff is $V(a_t, a_{t+1}, a_{t+2}, \dots)$. The strategies of such a game represent a subgame-perfect equilibrium of our game with a state variable: weak separability implies that the state does not affect the preference ordering of each player over the sequence of future actions. Moreover, these strategies satisfy Requirements 1 and 2 by the definition of a reconsideration-proof equilibrium. \square

A.2 Proof of Proposition 1

Proof. Let (a_0^E, a_1^E, \dots) be the outcome of a reconsideration-proof equilibrium for the game whose period- t payoff is $V(a_t, a_{t+1}, a_{t+2}, \dots)$, and let \bar{V} be its associated value. This means that, for any period t and any actions $a \in A$, there exists a continuation sequence $(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots)$ which is also the outcome of a reconsideration-proof equilibrium and is such that

$$V(a_t^E, a_{t+1}^E, a_{t+2}^E, \dots) \geq V(a, a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots). \quad (20)$$

We then have

$$V(a, a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots) = \tilde{V}(a, \hat{V}(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots)).$$

Acknowledging that the sequence $(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots)$ is potentially a function of the deviation a (as well as of time t , which we can hold fixed), define

$$\underline{V} := \inf_{a \in A} \hat{V}(a_{t+1}^{\bar{E}}, a_{t+2}^{\bar{E}}, \dots). \quad (21)$$

By the compactness of A , Tychonoff's theorem, and continuity of \hat{V} , we can find a sequence of actions a_0^*, a_1^*, \dots that attains the infimum in equation (21) above. Exploiting Assumption 3, this sequence ensures subgame perfection and satisfies the no-restarting condition (Requirement 3):

$$V(a_0^*, a_1^*, a_2^*, \dots) \geq V(a, a_0^*, a_1^*, \dots).$$

This path attains the value \bar{V} , so that it continues to satisfy the optimality condition of Requirement 2. Hence, playing (a_0^*, a_1^*, \dots) followed by a restart after any deviation is an equilibrium that satisfies Requirements 1, 2, and 3, and therefore (a_0^*, a_1^*, \dots) is an organizational equilibrium. \square

B Proofs and Details in Section 3.2

The players in the game are nature, plus an infinity of players $0, 1, \dots$ indexed by the time at which they act. Nature moves first, choosing a time \hat{t} , from which record keeping is possible. We assume that this distribution has full support on \mathbb{N} .³⁷

Players take two actions:

- Player t chooses $a_t \in A$.
- In addition, a player may choose a record-keeping action $\rho_t \in \{S, C, H\}$, where S stands for starting record keeping, C stands for continuing record keeping, and H for hiding past records. Whether these actions are available at time t depends on the past in a way that we will soon make explicit.

We are now ready to define histories and information. The first history is \emptyset , at which stage nature moves. In all periods $t < \hat{t}$, players only choose the action a_t . While the history of play is $(\hat{t}, a_0, \dots, a_{t-1})$, their only information is that $t < \hat{t}$, and the current level of the state k_t : they do not observe any of the past players' actions, and they only know that record keeping is not yet possible. In period $t = \hat{t}$, the history of play is also $(\hat{t}, a_0, \dots, a_{t-1})$. Player $t = \hat{t}$ observes k_t , does not observe any of the actions taken by the past players, but it knows that $t \geq \hat{t}$ and that either $t = \hat{t}$ or $\rho_{t-1} = H$: that is, it knows that it is either the first player with the opportunity to set up record keeping, or the opportunity was available in the past, but player $t - 1$ chose not to adopt it and to hide the previous history. Player $t = \hat{t}$ is called to choose an action $\rho_t \in \{S, H\}$ as well as a_t . In period $\hat{t} + 1$ and all subsequent periods, the history of play is $(\hat{t}, a_0, \dots, a_{\hat{t}-1}, \rho_{\hat{t}}, a_{\hat{t}}, \rho_{\hat{t}+1}, a_{\hat{t}+1}, \dots, \rho_{t-1}, a_{t-1})$. In each of these periods, if $\rho_{t-1} = H$, then player t only knows that record keeping is possible and the level of k_t ; it does not know whether $t = \hat{t}$ or $\rho_{t-1} = H$. In this case, player t has the same options as player \hat{t} . Otherwise, let \tilde{t} be the last time action S was taken; player t then knows that $\hat{t} \leq \tilde{t}$ and it knows $(\rho_{\tilde{t}}, a_{\tilde{t}}, \rho_{\tilde{t}+1}, a_{\tilde{t}+1}, \dots, \rho_{t-1}, a_{t-1})$ (in addition to k_t). Player t has 3 options for ρ_t : first, it can choose $\rho_t = H$, in which the next player will start again with no record of the past; second, it can choose $\rho_t = C$, that is, to continue record keeping: in this case, player $t + 1$ will know $(\rho_{\tilde{t}}, a_{\tilde{t}}, \rho_{\tilde{t}+1}, a_{\tilde{t}+1}, \dots, \rho_t, a_t)$. Finally, it can restart the history ($\rho_t = S$), disavowing the past, but recording its own actions, in which case player $t + 1$ will only observe (ρ_t, a_t) . In all cases player $t + 1$ will observe k_{t+1} .

A strategy σ_t for player t is a mapping from the set of time- t histories, H^t , to the set of actions A and (when available) record-keeping choices $\rho_t \in \{H, S, C\}$, that is measurable with respect to the information available at time t . As before, a strategy profile σ is a sequence of strategies, one for each player. It is useful to distinguish between the two choices made by agents: accordingly, let σ_{a_t} be the component of $\sigma_t(h^t)$ that contains the prescribed action $a \in A$ after history h^t , and $\sigma_{\rho_t}(h^t)$ be the prescribed choice of record keeping. Analogously, we define $\sigma_a := \{\sigma_{a_t}\}_{t=0}^\infty$ and $\sigma_\rho := \{\sigma_{\rho_t}\}_{t=0}^\infty$.

We restrict attention to equilibria that satisfy Requirement 1: that is, they involve strategies that are inde-

³⁷It would be equivalent to assume that nature moves in each period up to \hat{t} , as long as the conditional hazard rate of the start of record keeping is the same. This is because nature's choice is not fully observed by the agents anyway.

pendent of k_t .

We define a *full-disclosure* equilibrium to be an equilibrium in which $\sigma_{\rho_t}(h^t) = S$ for all histories for which $t \geq \hat{t}$ and no previous record of play is known, and $\sigma_{\rho_t}(h^t) = C$ for all histories for which player t observes a record of past actions. In a full-disclosure equilibrium, players introduce record keeping as soon as possible and never erase any of the record available, independent of the actions of past players.

The proof of Proposition 2 relies on a sequence of lemmata:

Lemma 1. *Let σ be a sequential equilibrium that satisfy Requirement 1 in the game defined above. Then:*

1. *There exists a full-disclosure sequential equilibrium $\tilde{\sigma}$ that also satisfies that satisfy Requirement 1 and such that the same actions $\{a_t\}_{t=0}^{\infty}$ are taken on the equilibrium path under σ and $\tilde{\sigma}$.*
2. *If σ satisfies Requirement 2 from period \hat{t} (whatever \hat{t} turns out to be), then $\tilde{\sigma}$ can be chosen to also satisfy the same requirement.*

Proof. Our proof only looks at pure-strategy equilibria. It could be extended to mixed-strategy equilibria, in which players randomize over their choice of record keeping, using the same logic presented here, as long as a public randomization device is present that allows coordination across players. We omit the case of mixed-strategy equilibria for brevity.

1. Assumption 1 implies that, if future players do not condition their choices on the state k (but potentially condition their choices on all their remaining information in any arbitrary way), the optimal choice for a current player is independent of the current state. In looking at equilibria that satisfy Requirement 1, we can therefore leave the state k in the background and focus only on the history of actions, disclosures, and the time at which record keeping becomes available.

Let $\sigma = \{\sigma_t\}_{t=0}^{\infty}$ be the strategy profile of the sequential equilibrium that contains the equilibrium action path $\{a_t\}_{t=0}^{\infty}$.

We need to construct an alternative strategy profile $\tilde{\sigma}$ that contains the same equilibrium action path, but involves full disclosure. We will do so by creating a suitable mapping from the set of histories to itself, and setting $\tilde{\sigma}_{a_t}(h^t) = \sigma_{a_t}(\eta(h^t))$. η is constructed recursively as follows:

- For $t \leq \hat{t}$, $\eta(h^t) = h^t$.
- For $t > \hat{t}$ and histories in which $\rho_{t-1} = H$, $\eta(h^t) = h^t$.
- For $t > \hat{t} + 1$ and histories in which $\rho_{t-1} = S$ and $\sigma_{\rho,t}(h^{t-1}) = S$ or $\sigma_{\rho,t}(h^{t-1}) = C$, $\eta(h^t) = h^t$.
- For $t > \hat{t} + 1$ and histories in which $\rho_{t-1} = S$ and $\sigma_{\rho,t}(h^{t-1}) = H$, $\eta(h^t) = (h^{t-1}, H, a_{t-1,h^t})$, where a_{t-1,h^t} is the action taken in period $t - 1$ according to the history h^t .
- For $t > \hat{t}$ and histories in which $\rho_{t-1} = C$, we define η recursively as $\eta(h^t) = (\eta(h^{t-1}), \sigma_{\rho,t}(h^{t-1}), a_{t-1,h^t})$.

Furthermore, whenever $t \geq \hat{t}$, $\tilde{\sigma}_{\rho t} = S$ if no record keeping is currently in place, and $\tilde{\sigma}_{\rho t} = C$ otherwise, in line with the definition of a full-disclosure equilibrium.

In words, $\tilde{\sigma}$ is constructed from σ by assuming that agents take the same actions under the two strategy profiles whenever they do not observe the past. When past actions are observed from \bar{s} on, the strategy profile $\tilde{\sigma}$ prescribes that the agents take the same actions they would have taken under σ when faced with a history that has same choices for (a_0, \dots, a_{t-1}) , but in which past players from \bar{s} on chose to hide, start, or continue record-keeping according to the equilibrium profile σ . At the same time, $\tilde{\sigma}$ always prescribes full disclosure. Next, we verify that $\tilde{\sigma}$ is a measurable strategy with respect to the information sets available to the players at each point t . The choice of ρ only depends on whether record keeping is possible and whether it is inherited from the past, which is observable to an agent at the time it makes its choice. Furthermore, by construction, the mapping η is such that the prescribed action $\tilde{\sigma}_{a_t}(h^t)$ is the same for all histories that share the same observable record.³⁸

Next, we verify that $\tilde{\sigma}$ represents a sequential equilibrium. A player's payoff only depends on the current and future actions $a_t \in A$, and only indirectly on record keeping choices.

- In any period $t < \hat{t}$, the current choice of a_t by player t is not known to future players and therefore it has no impact on any future action. Furthermore, the two strategies σ and $\tilde{\sigma}$ imply the same sequence of future actions $(a_{t+1}, a_{t+2}, \dots)$ along the equilibrium path.³⁹ The optimality of $\tilde{\sigma}_t$ then follows directly from that of σ_t .
- Consider next periods $t \geq \hat{t}$ and histories h^t such that no record is available to player t . For such histories, $\eta(h^t) = h^t$. There are two possibilities. First, suppose that $\sigma_{\rho}(h^t) = S$. Then, no matter what choice of (ρ_t, a_t) player t takes, the equilibrium implies that future players will take the same actions $\{a_s\}_{s=t+1}^{\infty}$ under profiles σ and $\tilde{\sigma}$. Hence, $\tilde{\sigma}_t(h^t) = \sigma_t(h^t)$ is an optimal choice. Suppose instead that $\sigma_{\rho t}(h^t) = H$, that is, according to the equilibrium profile σ , player t should hide its action. In this case, η is such that player t gets the same payoff whether it chooses $\rho_t = S$ or $\rho_t = H$, since $\eta(h^t, H, a_t) = \eta(h^t, S, a_t)$: player t is indifferent between starting record keeping or not, because in either case future players will ignore its play and behave as if no record had been taken in t . Starting record keeping is thus weakly optimal, and taking the same action that would have been taken under the profile σ is optimal as well.⁴⁰
- Consider histories h^t in which a record is present. The reasoning is similar. If $\sigma_{\rho, t}(\eta(h^t)) = C$, then, no matter what choice of (ρ_t, a_t) player t takes, the equilibrium implies that future players will take the same actions $\{a_s\}_{s=t+1}^{\infty}$ under profiles $\tilde{\sigma}$ and history h^t as they would under σ and history $\eta(h^t)$. Hence, if $\sigma_t(\eta(h^t))$ is optimal (taking as given that σ will be followed in the future), then $\tilde{\sigma}_t(h^t)$ is also optimal, if future players play according to $\tilde{\sigma}$. If $\sigma_{\rho t}(\eta(h^t)) = H$, then under $\tilde{\sigma}$ future players will ignore past actions whether player t chooses $\rho_t = H$ or $\rho_t = C$, and their future actions will follow the course dictated by $\sigma|_{(h^t, H, a_t)}$. By the measurability restriction, $\sigma|_{(h^t, H, a_t)} = \sigma|_{(\eta(h^t), H, a_t)}$. Hence, player t is indifferent between playing C or H . If player t

³⁸This assumes that the property is true for σ , which must be the case for σ to be a valid strategy profile and therefore a valid equilibrium, provided that σ does not condition on k_t , which is guaranteed by Requirement 1.

³⁹Notice that future actions are in general uncertain and depend on the realization of \hat{t} , but their stochastic process is identical in the two equilibria.

⁴⁰Since future players will ignore the action a_t , player t will maximize its payoff assuming that its action does not affect the future, as if no record were taken, just as it would under the strategy σ , which prescribes hiding the record.

chooses to restart the record, then the future players' actions will evolve according to $\sigma|_{(h^t, S, a_t)}$. Measurability implies again that $\sigma|_{(h^t, S, a_t)} = \sigma|_{(\eta(h^t), S, a_t)}$. Since σ is an equilibrium profile, playing $\sigma_t(\eta(h^t))$ (which in this case involves hiding the record from future players) is weakly better than playing S along with any of the possible actions, under the assumption that future players will follow the same profile σ . It follows that the consequences of playing H vs. S and any action a_t in period t under history h^t when future players will follow $\tilde{\sigma}$ are the same as those of playing the corresponding actions under history $\eta(h^t)$ when future players will follow σ . Hence, if $\sigma(\eta(h^t)) = H$, playing S is a (weakly) dominated choice. In sum, in this case player t is indifferent between H and C , and it weakly prefers either to S , which ensures that it is optimal for its to play C . Furthermore, choosing $a_t = \tilde{\sigma}_{a_t}(h^t) = \sigma_{a_t}(\eta(h^t))$ is optimal because it involves a static optimization taking as given the future choices (that will be independent of the current a_t and will be the same under h^t and $\tilde{\sigma}$ as they are under $\eta(h^t)$ and σ). The last case to consider is one in which $\sigma_{\rho_t}(\eta(h^t)) = S$; this case is similar to the previous one. Specifically, the measurability restriction implies $\sigma|_{(h^t, S, a_t)} = \sigma|_{(\eta(h^t), S, a_t)}$. Furthermore, if player t chooses $\rho_t = C$, $\tilde{\sigma}$ is such that future players will choose the same sequence of actions whether player t chooses $\rho_t = S$ or $\rho_t = C$: these actions will only depend on a_t , which is the only element of the record that is passed to future players according to the strategy σ_t . If player t chooses $\rho_t = H$, the future equilibrium path unfolds according to $\tilde{\sigma}|_{(h^t, H, a_t)} = \sigma|_{h^t, H, a_t} = \sigma|_{(\eta(h^t), H, a_t)}$, where the last equality follows the usual measurability restriction. If $\sigma_t(\eta(h^t)) = S$, then playing $\rho_t = S$ is weakly better than playing $\rho_t = C$ at $\eta(h^t)$ if σ will be followed in the future; this then implies that S (and the best action a_t conditional on S) is weakly better than C (and the best a_t conditional on C) at history h^t if $\tilde{\sigma}$ will be played in the future. This establishes that, under $\tilde{\sigma}$, playing C yields the same payoff as playing S , and a weakly better payoff than playing H . So, playing C is optimal. Finally, the usual equivalence of future consequences implies that playing $a_t = \tilde{\sigma}_{a_t}(h^t) = \sigma_{a_t}(\eta(h^t))$ is optimal.

2. Note that $\tilde{\sigma}$ is constructed so that the actions on the equilibrium path starting from any history h^t (whether the history itself is on or off equilibrium) are the same as the actions on the equilibrium path starting from $\eta(h^t)$ when σ is played. The mapping η is such that histories with $t \geq \hat{t}$ are mapped into histories with $t \geq \hat{t}$. If V is symmetric, then it achieves the same action payoff V following any history that has $t \geq \hat{t}$; as a consequence, the same property is inherited by $\tilde{\sigma}$. This implies that the set of values attainable by sequential equilibria satisfying Requirement 1 from period \hat{t} is the same as the set of values attainable by full-disclosure sequential equilibria satisfying Requirement 1 from the same period; the maxima of the two sets will thus coincide, completing the proof.⁴¹

□

Lemma 2. *Let $\tilde{\sigma}$ be a full-disclosure state-independent sequential equilibrium for the game in which history can be hidden. Then:*

1. $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$ is a subgame-perfect equilibrium for the game where record-keeping starts at time 0, and it

⁴¹The inability to keep records for periods before \hat{t} will in general imply that the payoff in previous periods is lower.

also satisfies state independence (Requirement 1);⁴²

2. If $\tilde{\sigma}$ is symmetric from period \hat{t} on, then $\tilde{\sigma}_a|_{h^{\hat{t}}}$ is also symmetric.

Proof. 1. In the game in which history can be hidden, in period \hat{t} , player \hat{t} starts with no information about the past, just as in period 0 of the game where record-keeping starts at time 0. Furthermore, $\tilde{\sigma}$ is such that records will be kept from \hat{t} on. Take as given the choice of ρ_t dictated by $\tilde{\sigma}$, and focus on the choice of a_t . In order for $\tilde{\sigma}$ to represent a sequential equilibrium, at any time $t \geq \hat{t}$ and after any history h^t it must be the case that $\sigma_{a_t}(h^t)$ (along with starting record keeping if no record is present or continuing it otherwise) is optimal, conditional on the fact that future players will continue to play $\tilde{\sigma}$. Let $h_{a,\hat{t}}^t$ represent the subcomponent of history h^t that captures the history of actions $(a_{\hat{t}}, a_{\hat{t}+1}, \dots, a_t)$. Since $\tilde{\sigma}$ implies that future players will behave in such a way that the entire history of play from \hat{t} is known, it then follows that $\sigma_{a_t}(h^t)$ must be optimal in the game where record keeping starts in period 0 after history $h_{a,\hat{t}}^t$, assuming that future players will play according to the strategy profile $\tilde{\sigma}_a|_{h^{\hat{t}}}$.

2. Symmetry implies that the action payoff V on the equilibrium path conditional on attaining any history h^t with $t \geq \hat{t}$ is the same. This property is inherited by $\tilde{\sigma}_a|_{h^{\hat{t}}}$ in any subgame following a history $h_{a,\hat{t}}^t$, since the action paths coincide going forward.

□

Lemma 3. *Let σ be a symmetric state-independent subgame-perfect equilibrium of the game where record-keeping starts in period 0. Then, if and only if σ satisfies Requirement 3 as well, there exists a state-independent full-disclosure sequential equilibrium $\tilde{\sigma}$ of the game where history can be hidden, which is symmetric from period \hat{t} and is such that $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$.*

Proof. Assume first that σ satisfies Requirement 3. The condition $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$ fully characterizes $\tilde{\sigma}_a$ from period \hat{t} on. To see this, let h^t be an arbitrary history in which $t > \hat{t} + s$ and player t observes $(a_{t-s}, a_{t+1-s}, \dots, a_{t-1})$; this implies that either player $t - s - 1$ chose to hide records, or player $t - s$ chose to restart them, while all subsequent players up to t chose to continue record keeping. This history is in the same information set as a history with the same sequence of actions $(a_{t-s}, a_{t+1-s}, \dots, a_{t-1})$ in which $\hat{t} = t - s$ and players adopted full disclosure; actions for such history are determined by $\tilde{\sigma}_a|_{h^{\hat{t}}} \equiv \sigma$. This observation also implies that, following any history, the sequence of actions a that are predicted to happen along a continuation equilibrium according to $\tilde{\sigma}$ is the same as those in a corresponding history in the game where record-keeping starts in period 0 under σ . If all histories under σ are followed by the same equilibrium action payoff \bar{V} , then the same value carries over to $\tilde{\sigma}$. To verify that $\tilde{\sigma}$ is indeed optimal after any history h^t , $t \geq \hat{t}$, we denote $h_{a_s}^t = (a_s, \dots, a_t)$ to be the record available to player t after history h^t and proceed as follows:⁴³

⁴²Note that, without further assumptions, $\tilde{\sigma}_a|_{h^{\hat{t}}}$ may depend on the precise realization of \hat{t} . The property still holds: in this case, each possible continuation strategy $\tilde{\sigma}_a|_{h^{\hat{t}}}$ is a subgame-perfect equilibrium of the game where record-keeping starts at time 0.

⁴³Along the equilibrium path, the record available should start from period \hat{t} , but we need to verify optimality even for histories that are not on the equilibrium path.

- Player t does not have an incentive to choose $\rho_t = C$ and any action $a \neq \sigma_t(h_{a,\hat{t}}^t)$. Assuming that future players will follow $\tilde{\sigma}$, the consequences of such a choice would be the same as those of choosing $a \neq \sigma_t(h_{a,\hat{t}}^t)$ after history $h_{a,\hat{t}}^t$ in the game where record-keeping starts in period 0 when future players follow σ ; since σ represents an equilibrium, choosing $a \neq \sigma_t(h_{a,\hat{t}}^t)$ is weakly worse.
- Player t does not have an incentive to choose $\rho_t = S$ and any action $a \in A$. Following such a choice, player $t + 1$ will behave as if $\hat{t} = t$, and future actions will unfold according to the strategy profile σ . Requirement 3 implies that, whatever action player t chooses, it would be (weakly) better off playing $\rho_t = S$ and $a = \sigma(\emptyset)$, that is, choosing to restart record keeping and playing the first action of the strategy profile of the game where record-keeping starts in period 0. This latter choice gives an action payoff of \bar{V} , which is the same as that obtained by continuing record keeping and following $\tilde{\sigma}$.
- Player t does not have an incentive to choose $\rho_t = H$ and any action $a \in A$. Following such a choice, player $t + 1$ and subsequent players will follow the strategy σ as if the game in which record-keeping starts in period 0 took place from that point on. Requirement 3 implies that, faced with this prospect, player t does not have any action that can guarantee a payoff higher than \bar{V} for herself.

To finish establishing the “if” part of the Lemma, the last step is to construct the strategy profile $\tilde{\sigma}$ in periods $t < \hat{t}$. In these periods, the actions taken by player t will not be observed by future players; as long as $\tilde{\sigma}$ is independent of the state, the actions of the current player will thus have no consequences on the actions taken by future players. We thus need to prove existence of a sequence of actions $(\tilde{a}_0, \tilde{a}_1, \dots)$ that will be taken by players in period t if $t < \hat{t}$, and that are optimal given that the same sequence will be continued up to the unknown time \hat{t} and given that starting in period \hat{t} actions will unfold according to the equilibrium path dictated by σ . Given σ , consider a correspondence $M : A^\infty \rightrightarrows A^\infty$ that associates to a sequence (a_0, a_1, \dots) all the sequences such that player t is choosing optimally given that (a_0, a_1, \dots) will be followed up to period \hat{t} and σ will be followed from period \hat{t} on. By Assumptions 2 and 4 and the theorem of the maximum, M is nonempty, compact- and convex-valued, upper hemicontinuous, and independent of the state. By Kakutani’s fixed-point theorem, M has a fixed point, which can be used as our desired sequence $(\tilde{a}_0, \tilde{a}_1, \dots)$.

Conversely, suppose that σ does not satisfy Requirement 3. We know from the previous part of the proof that player $t \geq \hat{t}$ can attain the action payoff \bar{V} by continuing record keeping and following the strategy $\tilde{\sigma}$, but also by playing $\rho_t = S$ and $a = \sigma(\emptyset)$, effectively starting the sequence (a_0, a_1, \dots) of Requirement 3. However, if player t hides the record and chooses $\rho_t = H$, then the strategy profile $\tilde{\sigma}$ implies that record keeping will start in period $t + 1$ and the actions (a_0, a_1, \dots) will unfold from period $t + 1$ instead. If Requirement 3 fails, there exists an action \tilde{a} such that $V(\tilde{a}, a_0, a_1, \dots) > V(a_0, a_1, \dots) = \bar{V}$, which yields a higher payoff than following $\tilde{\sigma}$; this would imply that $\tilde{\sigma}$ is not an equilibrium strategy profile. \square

We are now ready to prove Proposition 2.

Proof of Proposition 2. In the game in which record-keeping starts in period 0, let σ be a strategy profile whose equilibrium path is an organizational equilibrium. By Lemma 3 we can find a state-independent strategy profile $\tilde{\sigma}$ for the game in which history can be hidden that attains the same equilibrium path from \hat{t} on, whatever

the realization of \hat{t} ; this equilibrium is also symmetric. To complete the proof, we need to show that there is no other state-independent equilibrium which is symmetric from period \hat{t} on and attains a higher payoff from that point onwards. By contradiction, suppose that such an equilibrium existed, let it be $\bar{\sigma}$. From Lemma 1, we can assume without loss of generality that $\bar{\sigma}$ involves full revelation. Lemma 2 implies that $\bar{\sigma}|_{h,\hat{t}}$ is a symmetric state-independent equilibrium of the game in which history can be hidden, which would then achieve a higher payoff than σ ; however, this would imply that σ does not satisfy Requirement 2 and therefore that its equilibrium path is not an organizational equilibrium, establishing a contradiction. \square

C Further Discussion of Assumptions 1, 3, 4, and 5.

Assumption 1 is central to our definition. By ensuring that the preference ordering over sequences of actions is independent of the state, it provides a way of achieving a meaningful comparison across different periods of time (or different histories) for which the state variable is different. Section 3.5 provides an example where this assumption fails and illustrates a way we construct an approximating economy that satisfies it. Without uncertainty, utility functions are only identified up to monotone transformations. In this case, it can be shown that Assumptions 1 and 4 are equivalent. However, in the game of Section 3.2, uncertainty is present, and we need the separability property to apply to *lotteries* about future outcomes. In this case, utility functions are identified up to affine transformations, and Assumption 4 is stronger than Assumption 1. Nonetheless, all of the separable preferences that we use in practice satisfy it. A (contrived) example of preferences that satisfies Assumption 1 but not Assumption 4 is one in which we amend the preferences of Section 2 to be

$$E_t \left[u(c_t) + \delta \sum_{\tau=1}^{\infty} \beta^\tau u(c_{t+\tau}) \right]^\xi,$$

with $\xi < 1$: in addition to the standard risk aversion period by period (embedded in u), these preferences exhibit risk aversion over the entire infinite sequence. When Assumption 1 holds but Assumption 4 fails, an organizational equilibrium still exists, but the interpretation based on the alternative game of Section 3.2 does not necessarily apply. An avenue to generalize the results to this case would be to study the limiting behavior of the game of Section 3.2 to the probability of record-keeping being available in each period converging to 1.

The other Assumptions that we introduce are sufficient conditions that allow us to derive our results in a clean way, but there is often an alternative way to derive similar results in economies that do not satisfy them, in particular by relying on Proposition 3.

Specifically, we use Assumption 3 to prove that the game that includes only the action payoff $V(a_t, a_{t+1}, \dots)$ has a reconsideration-proof equilibrium that satisfies the no-restarting condition. Weak separability between the initial action and the following sequence of actions allows us to find a worst continuation sequence that is a sufficient deterrent for all possible deviations. When Assumption 3 fails, the worst continuation may depend on the action taken, so for example the threat of restarting might work for the action to be taken in period $t+1$, but not in period $t+2$. Nonetheless, checking whether this is the case in an application is not a difficult exercise. As an example, consider the following modification of preferences and technology of Section 2, that

induce a violation of Assumption 3. At time t , preferences in terms of consumption sequences are given by

$$\frac{1}{1-\sigma} \left[c_t^{(1-\iota)(1-\sigma)} + \beta\delta \sum_{v=0}^{\infty} \beta^v \left(\frac{c_{t+v+1}}{c_{t+v}} \right)^{(1-\sigma)} \right],$$

with $\sigma \neq 1$ and $\iota \in (0, 1]$,⁴⁴

$$k_{t+1} = Ak_t - c_t.$$

Compared to the standard case, continuation preferences embed habit formation.⁴⁵ When we express this problem as preferences over a sequence of saving rates, so as to isolate the role of capital, we obtain

$$\frac{(Ak_t)^{(1-\iota)(1-\sigma)}}{1-\sigma} \left[(1-s_t)^{(1-\iota)(1-\sigma)} + \beta\delta \sum_{v=0}^{\infty} \beta^v \left(A^{v(1-\iota)+1} s_{t+v} (1-s_{t+v+1}) (1-s_{t+v})^{-\iota} \prod_{n=0}^{v-1} s_{t+n}^{1-\iota} \right)^{1-\sigma} \right].$$

For these preferences, the marginal rate of substitution between s_{t+1} and s_{t+2} depends on s_t , so that separability of s_t from the remaining sequence fails. Nonetheless, we can establish whether an organizational equilibrium exists by computing it from a recursive structure. Even when Assumption 3 fails, the proof of Theorem 1 implies that a reconsideration-proof equilibrium of the game where preferences are given by the action component only exists. On the path of play implied by such an equilibrium, the value from the sequence of actions (excluding the separable state) is constant:

$$\frac{1}{1-\sigma} \left[(1-s_t)^{(1-\iota)(1-\sigma)} + \beta\delta \sum_{v=0}^{\infty} \beta^v \left(A^{v(1-\iota)+1} s_{t+v} (1-s_{t+v+1}) (1-s_{t+v})^{-\iota} \prod_{n=0}^{v-1} s_{t+n}^{1-\iota} \right)^{1-\sigma} \right] = \bar{V}.$$

Using the fact that the players at time t and $t+1$ attain the same value \bar{V} , we derive a recursive expression similar to the one we derive in the applications of the main text:

$$\begin{aligned} \bar{V}(1-\beta(As_t)^{(1-\iota)(1-\sigma)}) &= (1-s_t)^{(1-\iota)(1-\sigma)} \\ &+ \beta\delta(As_t(1-s_{t+1})(1-s_t)^{-\iota})^{1-\sigma} - \beta(As_t(1-s_{t+1}))^{(1-\iota)(1-\sigma)} \end{aligned} \quad (22)$$

For any given \bar{V} , equation (22) is a difference equation in the saving rates that can be solved numerically. For any given parameter combination, we can check whether this difference equation implies monotonic convergence to a steady state. As long as δ and ι are such that there is an incentive to undersave in the first period, that was the case in the numerical examples we tried. When this is the case, we can proceed as in the main text:

- Find the steady state that maximizes the value \bar{V} ;
- From equation (22), derive the function that maps s_t into s_{t+1} ;
- For any potential initial starting point s_0 , we can compute the payoff that a player at time t receives if she expects restarting from s_0 to happen in period $t+1$ and plays the best response to it, and compare it

⁴⁴When $\sigma = 1$ we obtain the logarithmic case, that preserves Assumption 3 even with the habit-formation specification here. $\iota = 0$ is the standard case in which Assumption 3 also applies.

⁴⁵Introducing habit formation over the initial time- t consumption would break separability between the state and the actions.

to \bar{V} . If a value s_0 can be found such that the threat of reversion to s_0 in the future is enough to (weakly) deter any action, we have found an organizational equilibrium. Such a value for s_0 is guaranteed to exist under Assumption 3, and not here. Nonetheless, in the numerical examples we tried, there is an interval of values of s_0 where the condition is satisfied, just as in our applications in the main text, so an organizational equilibrium exists; as in the main text, we pick the highest saving in the interval where no-restarting applies based on Pareto optimality (though another choice would also be valid and converge to the same constant saving rate in the long run).

Finally, we used Assumption 5 to guarantee the existence of an organizational equilibrium which is recursive in the continuation value. As always in infinite-horizon models, a recursive structure is of great help for computations. Assumption 5 implies that the preference disagreement between the players moving at t and $t + 1$ only concerns the action taken at $t + 1$: conditional on the action taken at $t + 1$, they agree on their preference ordering over sequences of actions from $t + 2$ on. This allows us to use the continuation value $\widehat{V}(a_{t+2}, a_{t+3}, \dots)$ as a state in computing the equilibrium path recursively. Even when Assumption 5 fails, there may be other ways of obtaining a recursive representation. As an example, we consider here a variant of the consumption-saving problem of Section 2. We now assume that the planner is seeking to maximize the utility of a two-person household where both members have standard time-consistent preferences and share consumption, but they differ in their discount factor, generating time-inconsistency for the planner as in Jackson and Yariv (2014, 2015). Preferences at time t are thus given by

$$\sum_{v=0}^{\infty} (\beta_h^v + \lambda \beta_\ell^v) \log(c_{t+v}),$$

with $0 < \beta_\ell < \beta_h < 1$, and $\lambda > 0$ being a measure of the relative Pareto weight of the impatient member. Section 2 is a limiting case of these preferences as $\beta_\ell = 0$, $\beta_h = \beta$, and $\delta = 1/(1 + \lambda)$. When $\beta_\ell > 0$, Assumption 5 fails, as we can see considering the relative discount factor between periods $t + 2$ and $t + 3$. From the perspective of period t , the relative discount factor is $(\beta_h^3 + \beta_\ell^3)/(\beta_h^2 + \beta_\ell^2)$, while from the perspective of period $t + 1$ it is $(\beta_h^2 + \beta_\ell^2)/(\beta_h + \beta_\ell)$. As a consequence, the players at t and $t + 1$ differ not only in the relative valuation of saving in period $t + 1$, but also on saving in any future period. We can nonetheless retrieve a recursive structure for this game as well. Specifically, let $V_{\ell,t}(s_t, s_{t+1}, \dots)$ and $V_{h,t}(s_t, s_{t+1}, \dots)$ be the values accruing to the impatient and the patient member of the household respectively, when the planner chooses a sequence of saving rates (s_t, s_{t+1}, \dots) , excluding the additive utility from initial capital $\alpha/(1 - \alpha\beta_i)k_t$ for $i = h, \ell$. Since each member has standard time-consistent preferences, we can express these values recursively:


$$V_{i,t} = \log(1 - s_t) + \frac{\alpha\beta_i}{1 - \alpha\beta_i} \log s_t + \beta_i V_{i,t+1}, \quad i = h, \ell. \quad (23)$$

Since this economy satisfies the conditions of Proposition 1, an organizational equilibrium exists. The utility attained by the planner in such an equilibrium is a constant value $\bar{V} = V_{h,t} + \lambda V_{\ell,t}$. From the Pareto weighted sum of the two equations we then obtain

$$(1 + \lambda) \log(1 - s_t) + \left(\frac{\alpha\beta_h}{1 - \alpha\beta_h} + \frac{\alpha\beta_\ell}{1 - \alpha\beta_\ell} \right) \log s_t - \lambda(\beta_h - \beta_\ell) V_{\ell,t+1} = \bar{V}(1 - \beta_h). \quad (24)$$

For any given value \bar{V} , equation (24) admits a unique solution for s_t as a function of $V_{\ell,t+1}$;⁴⁶ we can substitute this solution into (23) for $i = \ell$ and obtain a difference equation in $V_{\ell,t}$. We have thus obtained a recursive representation in terms of the value $V_{\ell,t}$. This difference equation can alternatively be expressed in terms of s_t , since (24) implies a one-to-one correspondence. In our numerical evaluations, this difference equation behaves in the same way as it does in the baseline case of Section 2, so that the same procedure described above for the habit-formation economy can be used again here to compute the organizational equilibrium. Figure 7 plots such an example for the transition function between s_t and s_{t+1} . When $\beta_L = 0$, this economy becomes the standard quasi-hyperbolic discounting example in our baseline analysis. When $\beta_L > 0$, the dynamics needs to be computed based on equations (23) and (24). In the end, the transitional dynamics are similar qualitatively.

FIGURE 7: Evolution of the Saving Rate in Jackson and Yariv (2014)



D Proofs of Section 3.3.

D.1 Proof of Proposition 3

Proof. Suppose first that $\{\bar{a}_t\}_{t=0}^{\infty}$ is a sequence satisfying the three properties in the proposition. We construct a subgame-perfect equilibrium strategy profile as follows.⁴⁷ We start with $\sigma_0(\emptyset) = \bar{a}_0$. Let \tilde{h}^t , $t \geq 1$ be an arbitrary history whose predecessors are $(\emptyset, \tilde{h}^0, \tilde{h}^1, \dots, \tilde{h}^{t-1})$. If $a_s = \sigma_s(\tilde{h}^{s-1})$, $s = 0, \dots, t-1$, set $\sigma_t(\tilde{h}^t) = \bar{a}_t$. Otherwise, let $\tilde{t} := \max\{s : a_s \neq \sigma_s(\tilde{h}^{s-1})\}$ and set $\sigma_t(\tilde{h}^t) = \bar{a}_{t-1-\tilde{t}}$. In words, this strategy punishes any deviation by restarting the continuation equilibrium from the same equilibrium path that is supposed to prevail

⁴⁶More precisely, the equation admits at most one solution, and may have none. However, since we know that an organizational equilibrium exists, a solution has to exist for the appropriate range of values of \bar{V} and V_{ℓ} .

⁴⁷We defined an organizational equilibrium within the context of the game of Section 3.1, so the proposition is proven in the context of this game, although of course the results apply to the game of Section 3.2 when Assumption 4 is satisfied.

in period 0. Properties 1 and 3 ensure that such a punishment is sufficient to deter deviations, both in the initial period and in any subsequent period and history. This equilibrium is state independent (Requirement 1) and symmetric, since the equilibrium path of play attains an action value \bar{V} independent of the past history. No equilibrium can attain a higher constant value. Suppose such an equilibrium existed, and let $\{a_t^B\}_{t=0}^\infty$ be its equilibrium path, which attains a constant $V^B > \bar{V}$. Then we would have $V(a_t^B, a_{t+1}^B, \dots) = V^B > \bar{V}, \forall t \geq 0$, which would contradict property 2 of our initial sequence. Therefore, the newly constructed subgame-perfect equilibrium satisfies Requirement 2. Finally, Requirement 3 is a direct analog of the third property that we imposed on the sequence.

Suppose now that a sequence satisfying the 3 properties of the proposition exists and its value is \bar{V} . Requirement 2 implies that all organizational equilibria feature a path of constant value \bar{V} as well, which implies that they satisfy the first two properties; the third property follows directly from Requirement 3. \square

D.2 Proof of Proposition 4.

To prove this we rely on a useful lemma, which introduces a convenient way of representing equilibria through their values, similarly to Abreu, Pierce, and Stacchetti's (1986; 1990) method.⁴⁸

Lemma 4. *Let $V^* \in \mathbb{R}$ and $\hat{\mathcal{V}} \subset \mathbb{R}$ be a value and a set of continuation values that satisfy the following properties:*

1.

$$\forall a \in A \quad \exists \hat{v} \in \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) \leq V^*;$$

2.

$$\forall v \in \hat{\mathcal{V}} \quad \exists (a, \hat{v}) \in A \times \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) = V^* \wedge W(a, \hat{v}) = v.$$

3. *There exists no value $V^{**} > V^*$ and set $\hat{\mathcal{V}}$ that satisfies properties 1 and 2; furthermore, there is no set $\hat{\mathcal{V}}_a \supset \hat{\mathcal{V}}$ that satisfies properties 1 and 2 together with V^* .*

Then:

- *Construct an arbitrary sequence of actions $\{a_t^*\}_{t=0}^\infty$ recursively as follows. In period 0, pick $\hat{v}_0^* \in \hat{\mathcal{V}}$ and $(a_0^*, \hat{v}_1^*) \in A \times \hat{\mathcal{V}}$ such that $\tilde{V}(a_0^*, \hat{v}_1^*) = V^*$ and $W(a_0^*, \hat{v}_1^*) = \hat{v}_0^*$. In each subsequent period, pick $(a_t^*, \hat{v}_{t+1}^*) \in A \times \hat{\mathcal{V}}$ such that $\tilde{V}(a_t^*, \hat{v}_{t+1}^*) = V^*$ and $W(a_t^*, \hat{v}_{t+1}^*) = \hat{v}_t^*$. Constructing such a sequence is possible by the definition of V^* and $\hat{\mathcal{V}}$. The sequence so constructed is the outcome of a reconsideration-proof equilibrium;*

⁴⁸Note, however, that we cannot adopt their method to recursively compute the desired sets. Given V^* , $\hat{\mathcal{V}}$ can be computed recursively as in Abreu, Pierce, and Stacchetti. However, without further assumptions the set of values of V^* for which $\hat{\mathcal{V}}$ is defined need not be convex, which makes finding its maximum difficult.

- If $\{a_t^*\}_{t=0}^\infty$ is the equilibrium path of a reconsideration-proof equilibrium, $\tilde{V}(a_0^*, a_1^*, \dots) = V^*$ and $\hat{V}(a_t^*, a_{t+1}^*, \dots) \in \hat{\mathcal{V}}$ for any $t > 0$.

Proof.

First, we prove that the recursively-constructed sequence $\{a_t^*\}_{t=0}^\infty$ satisfies

$$\tilde{V}(a_t^*, \hat{V}(a_{t+1}^*, a_{t+2}^*, \dots)) = V^* \quad \forall t \geq 0 \quad (25)$$

and

$$\hat{V}(a_t^*, a_{t+1}^*, a_{t+2}^*, \dots) \in \hat{\mathcal{V}} \quad \forall t \geq 0. \quad (26)$$

Note that, if $\hat{v}_T^* = \hat{V}(a_T^*, a_{T+1}^*, a_{T+2}^*, \dots)$ for some period T , iterating backwards we find that $\hat{v}_t^* = \hat{V}(a_t^*, a_{t+1}^*, a_{t+2}^*, \dots)$ for all $t < T$, so that equations (25) and (26) hold.

Define

$$\{\underline{a}_t\}_{t=0}^\infty \in \arg \min_{\{a_t\}_{t=0}^\infty} \hat{V}(a_0, a_1, \dots)$$

and similarly let $\{\bar{a}_t\}_{t=0}^\infty$ be a sequence that attains the maximum. Both exist by the compactness of A and the continuity of \hat{V} (in the product topology).

Next, truncate the sequence $\{a_t^*\}_{t=0}^\infty$ at time $S > T$ and replace the continuation with $\{\underline{a}_t\}_{t=0}^\infty$ or $\{\bar{a}_t\}_{t=0}^\infty$. By Assumption 5 and the monotonicity of W , we have

$$\hat{V}(a_T^*, a_{T+1}^*, \dots, a_S^*, \underline{a}_0, \underline{a}_1, \dots) \leq \hat{V}(a_T^*, a_{T+1}^*, \dots, a_S^*, a_{S+1}^*, a_{S+2}^*, \dots) \leq \hat{V}(a_T^*, a_{T+1}^*, \dots, a_S^*, \bar{a}_0, \bar{a}_1, \dots) \quad (27)$$

and

$$\begin{aligned} \hat{V}(a_T^*, a_{T+1}^*, \dots, a_S^*, \underline{a}_0, \underline{a}_1, \dots) &= W(a_T^*, W(a_{T+1}^*, \dots, W(a_S^*, W(\underline{a}_0, W(\underline{a}_1, \dots)))) \dots) \leq \\ W(a_T^*, W(a_{T+1}^*, \dots, W(a_S^*, \hat{v}_S^*))) &= \hat{v}_T^* \leq \\ W(a_T^*, W(a_{T+1}^*, \dots, W(a_S^*, W(\bar{a}_0, W(\bar{a}_1, \dots)))) \dots) &= \hat{V}(a_T^*, a_{T+1}^*, \dots, a_S^*, \bar{a}_0, \bar{a}_1, \dots). \end{aligned} \quad (28)$$

Taking limits as $S \rightarrow \infty$ in equations (27) and (28) and exploiting the continuity of \hat{V} according to the product topology, the left-most and right-most expressions in the inequalities converge to the same value, which then implies that indeed $\hat{v}_T^* = \hat{V}(a_{T+1}^*, a_{T+2}^*, a_{T+3}^*, \dots)$ and (25) and (26) hold.

To complete the proof of the first point, we need to show that there exists no symmetric subgame-perfect equilibrium whose payoff is strictly greater than V^* . By contradiction, suppose that there is such an equilibrium with value $V^{**} > V^*$. Let σ^{**} be the strategy profile representing one such equilibrium. Define

$$\hat{\mathcal{V}}_b := \{v : v = \hat{V}(a_{t+1}^{**}|h^t, a_{t+2}^{**}|h^t, a_{t+3}^{**}|h^t, \dots), h^t \in A^t\},$$

where $\{a_s^{**}|h^t\}_{s=t+1}^\infty$ is the equilibrium path implied by the strategy profile σ^{**} following a history h^t . The

pair $(V^{**}, \hat{\mathcal{V}}_b)$ satisfies property 1 in the lemma, since otherwise σ_0^{**} would not be optimal at time 0. It also satisfies property 2 since σ^{**} is symmetric and by the definition of $\hat{\mathcal{V}}_b$. But then this implies that property 3 in the lemma does not hold for V^* , establishing a contradiction.

In the previous point we proved that, given V^* and $\hat{\mathcal{V}}$, we can construct a reconsideration-proof equilibrium of value V^* . Since all reconsideration-proof equilibria must have the same value, it must be the case that $\tilde{V}(a_0^*, a_1^*, \dots) = V^*$. Furthermore, repeating the steps of the previous point, we can prove that the value V^* and the set

$$\hat{\mathcal{V}}_a := \{v : v = \hat{V}(a_{t+1}^*|_{h^t}, a_{t+2}^*|_{h^t}, a_{t+3}^*|_{h^t}, \dots), h^t \in A^t\}$$

satisfy properties 1 and 2. By the definition of $\hat{\mathcal{V}}$, it follows that $\hat{\mathcal{V}}_a \subseteq \hat{\mathcal{V}}$. \square

While not essential for the proof of Proposition 4, the following lemma is useful for computations:

Lemma 5. *The set $\hat{\mathcal{V}}$ defined in Lemma 4 is convex.*⁴⁹

Proof. We first define the set $\hat{\mathcal{V}}_c$ by relaxing property 2 in Lemma 4 to be the following:

$$\forall v \in \hat{\mathcal{V}}_c \quad \exists (a, \hat{v}) \in A \times \hat{\mathcal{V}} : \tilde{V}(a, \hat{v}) \geq V^* \wedge W(a, \hat{v}) = v. \quad (29)$$

We will later prove that $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$.

Simple case. First, if $\hat{\mathcal{V}}_c$ is a singleton, then it is necessarily convex and $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$: by property 3 of Lemma 4, V^* should be raised until $\tilde{V}(a, \hat{v}) = V^*$ at the single element $\hat{v} \in \hat{\mathcal{V}}_c$, with no effect on property 2 and relaxing the constraint in property 1.

From now on, we study the case in which $\hat{\mathcal{V}}_c$ contains at least two values.

Step 1. To prove that $\hat{\mathcal{V}}_c$ is convex, we prove that its convex hull, $\text{Co}(\hat{\mathcal{V}}_c)$, satisfies properties 1 and 2 as well (and of course $\text{Co}(\hat{\mathcal{V}}_c) \supset \hat{\mathcal{V}}_c$ unless $\hat{\mathcal{V}}_c$ is convex as well). Property 1 is immediate from the monotonicity of \tilde{V} . Let $v_1, v_2 \in \hat{\mathcal{V}}_c$, and let $(a_1, \hat{v}_1), (a_2, \hat{v}_2)$ elements of $A \times \hat{\mathcal{V}}_c$ be two pairs of actions and continuation values that satisfy property 2 of Lemma 4. Consider their convex combination $(\alpha v_1 + (1 - \alpha)v_2, \alpha \hat{v}_1 + (1 - \alpha)\hat{v}_2)$, $\alpha \in [0, 1]$. Since \tilde{V} is continuous and quasiconcave and W is continuous, $\tilde{V}(\alpha v_1 + (1 - \alpha)v_2, \alpha \hat{v}_1 + (1 - \alpha)\hat{v}_2) \geq V^*$, and $W(\alpha v_1 + (1 - \alpha)v_2, \alpha \hat{v}_1 + (1 - \alpha)\hat{v}_2)$ takes all values in $[v_1, v_2]$ as α varies between 0 and 1. Hence, all intermediate values satisfy property 2 as well, which completes the proof that $\text{Co}(\hat{\mathcal{V}}_c)$ satisfies property 2.

Step 2. To prove that $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$, proceed as follows. Define $\underline{v}_c := \min\{\hat{\mathcal{V}}_c\}$ and $\bar{v}_c := \max\{\hat{\mathcal{V}}_c\}$.⁵⁰ By definition, we can find $(\underline{a}, \underline{\hat{v}})$ and $(\bar{a}, \bar{\hat{v}})$ such that

$$\tilde{V}(\underline{a}, \underline{\hat{v}}) \geq V^* \wedge W(\underline{a}, \underline{\hat{v}}) = \underline{v}$$

⁴⁹Lemma 4 defines a unique set, since the union of all sets satisfying properties 1 and 2 satisfies properties 1 and 2 as well.

⁵⁰It is straightforward to prove that $\hat{\mathcal{V}}_c$ is closed, by the continuity of the functions defining it.

and

$$\tilde{V}(\bar{a}, \bar{v}) \geq V^* \wedge W(\bar{a}, \bar{v}) = \bar{v}.$$

Since A is convex, we can construct within it a line from \underline{a} to \bar{a} by defining $a(\alpha) := \alpha \underline{a} + (1 - \alpha) \bar{a}$, $\alpha \in [0, 1]$. By the quasiconcavity of \tilde{V} , we know

$$\tilde{V}(a(\alpha), \alpha \underline{v} + (1 - \alpha) \bar{v}) \geq V^*.$$

By property 1 of Lemma 4, for each action $a(\alpha)$ and the monotonicity and continuity of \tilde{V} we have

$$\tilde{V}(a(\alpha), \underline{v}) \leq V^*.$$

Since $\hat{\mathcal{V}}_c$ is convex, we can find a (unique) value $\hat{v}(\alpha)$ such that

$$\tilde{V}(a(\alpha), \hat{v}(\alpha)) = V^*.$$

Monotonicity and continuity of \tilde{V} imply that $\hat{v}(\alpha)$ is a continuous function. It then follows that $\hat{V}(a(\alpha), \hat{v}(\alpha))$ is a continuous function of α . As $\alpha \in [0, 1]$, this function must take all values between \underline{v} and \bar{v} , proving that the property 2 of Lemma 4 is satisfied by $\hat{\mathcal{V}}_c$ and thus $\hat{\mathcal{V}}_c = \hat{\mathcal{V}}$. \square

We are now ready to prove Proposition 4.

Proof. The second property of the value V^* and the set $\hat{\mathcal{V}}$ in Lemma 4 implies that we can construct a function $g : \hat{\mathcal{V}} \rightarrow \mathbb{R} \times \hat{\mathcal{V}}$ with the property that $\tilde{V}(g(v)) = V^*$ and $W(g(v)) = v$. Starting from any value $v_0 \in \hat{\mathcal{V}}$, we can construct recursively a path $(a_t, v_{t+1}) = g(v_t)$. By Lemma 4, this is the equilibrium path of a reconsideration-proof equilibrium. It will thus be an organizational equilibrium provided that

$$V(a_t, v_{t+1}) \geq \max_a \tilde{V}(a, v_0) \quad \forall t.$$

By the definition of \mathcal{V} , this property is satisfied by its least element, \underline{v} ⁵¹ hence, it will be satisfied provided that the initial value v_0 is sufficiently low. \square

D.3 Proof of Proposition 5

Proof. Define a correspondence $\zeta : \mathbb{R} \times \mathbb{R} \rightrightarrows \mathbb{R}$ as follows:

$$v \in \zeta(v', v^*) \iff \exists a \in A : \begin{cases} \tilde{V}(a, v') = v^* \\ W(a, v') = v. \end{cases} \quad (30)$$

⁵¹By the monotonicity of \tilde{V} in its second argument and the property 1 of \mathcal{V} , $\tilde{V}(a, \underline{v}) \leq V^*$ for all $a \in A$.

In words, given (v^*, v') , v belongs to the correspondence if there is an action a which, together with a continuation value v' , yields utility v^* when evaluated according to the decision maker's preferences (\tilde{V}) and utility v when evaluated with its continuation utility function W .

We prove that there exists a value v^* for which ζ is nonempty and admits a fixed point in continuation utilities ($v = v'$). We do so by proving that a Markov equilibrium (a^M, v^M) exists, such that

$$v^* = \tilde{V}(a^M, v^M) = \max_a \tilde{V}(a, v^M) \quad (31)$$

and

$$v^M = W(a^M, v^M). \quad (32)$$

To prove the existence of a Markov equilibrium, we construct a correspondence $\hat{a}(\cdot)$ from A into itself by setting

$$\hat{a}(a) = \max_{a_0 \in A} \hat{V}(a_0, a, a, \dots).$$

By the usual compactness and continuity properties, this correspondence is nonempty, compact-valued, and upper hemicontinuous. Quasiconcavity of \hat{V} ensures that it is also convex-valued. Hence, the correspondence has a fixed point by Kakutani's theorem; let a^M be one such fixed point. Given Assumption 5, letting $v^M := \hat{V}(a^M, a^M, a^M, \dots)$, equations (31) and (32) are satisfied.

We thus know $v^M \in \zeta(v^M, \tilde{V}(a^M, v^M))$. Once again, our assumptions about compactness and continuity imply that the correspondence ζ is upper hemicontinuous. Let V^* be the maximal value for which ζ admits a fixed point in continuation utilities. In the proofs below, it is useful to establish that

$$v \in \zeta(v', V^*) \implies v \leq v'. \quad (33)$$

Suppose (33) is not satisfied. Let (a, v') be such that $V(a, v') = V^*$ and $W(a, v') > v'$. Holding the action a fixed, continuity and monotonicity imply that higher values of v' lead to higher values of $V(a, v')$ and $W(a, v')$. As long as $W(a, v') > v'$, we know that $v' < \max_{\{a_t\}_{t=0}^\infty} \hat{V}(a_0, a_1, \dots)$ and can thus be raised further. Eventually, we will attain a value $v^h > v'$ for which $W(a, v^h) = v^h$ (this has to happen, since $W(a, v')$ is bounded by the maximum above). Let $V^h := V(a, v^h) > V^*$. We just established that a fixed point of $\zeta(\cdot, V^h)$ exists, which contradicts the assumption that V^* is the highest value for which a fixed point can be found.

In our next step, we prove that there are no symmetric equilibria with value $V^{**} > V^*$. By the definition of V^* , given any combination of an action and a continuation utility (a, v') , if $\tilde{V}(a, v') = V^{**}$ then $W(a, v') < v'$. This implies that any equilibrium path with value V^{**} would feature a strictly increasing sequence of continuation values; convergence is ruled out, because continuity and compactness would imply that the limiting point would be a fixed point of ζ , which is inconsistent with $V^{**} > V^*$. Since the set of possible continuation values is bounded by

$$\max_{\{a_t\}_{t=0}^\infty} \hat{V}(a_0, a_1, \dots),$$

no such equilibrium path can exist.

We now prove that there exist symmetric equilibria with value V^* , which then implies that any such equilibrium is reconsideration proof. Let v^{SS} be the maximal fixed point of $\zeta(\cdot, V^*)$. For any continuation value $v > v^{SS}$, a repetition of the arguments described above for V^{**} imply that no equilibrium path would be possible.⁵² We prove instead that there exists a convex set $\mathcal{V} = [v_\ell, v^{SS}]$ which, together with V^* , satisfies the properties of Lemma 4, where

$$v_\ell := \min_{v' \leq v^{SS}} \min \zeta(v', V^*). \quad (34)$$

To do so, prove first that, for any action $a \in A$, $\tilde{V}(a, \min_{\{a_t\}_{t=0}^\infty} \widehat{V}(a_0, a_1, \dots)) \leq V^*$. By contradiction, suppose that an action a_L such that $\tilde{V}(a_L, \min_{\{a_t\}_{t=0}^\infty} \widehat{V}(a_0, a_1, \dots)) > V^*$ existed. We could then repeat the same steps used to prove (33) and construct a steady state with value higher than V^* .

Since $\tilde{V}(a, \min_{\{a_t\}_{t=0}^\infty} \widehat{V}(a_0, a_1, \dots)) \leq V^* \quad \forall a \in A$, we can define

$$v'_{\min} := \min_{(a, v')} v' := \tilde{V}(a, v') = V^*.$$

Since there exists an action a^{SS} such that $V(a^{SS}, v^{SS}) = V^*$, $v'_{\min} \leq v^{SS}$. Also, by equations (33) and (34), $v_\ell \leq v'_{\min}$. Hence, $\tilde{V}(a, v_\ell) \leq V^* \quad \forall a \in A$: Property 1 of Lemma 4 is satisfied by the value V^* and the continuation set $[v_\ell, v^{SS}]$. To prove Property 2, let a_ℓ and v'_ℓ be such that $W(a_\ell, v'_\ell) = v_\ell$ and $\tilde{V}(a_\ell, v'_\ell) = V^*$, and $\lambda \in [0, 1]$.⁵³ As we just established, $\tilde{V}(\lambda a_\ell + (1 - \lambda)a^{SS}, v_\ell) \leq V^*$. By quasiconcavity, $\tilde{V}(\lambda a_\ell + (1 - \lambda)a^{SS}, \lambda v'_\ell + (1 - \lambda)v^{SS}) \geq V^*$. Strict monotonicity implies that there exists a unique value v_λ such that $\tilde{V}(\lambda a^{SS} + (1 - \lambda)a_\ell, v_\lambda) = V^*$, which must vary continuously with λ by the continuity of \tilde{V} . It follows that $W(\lambda a^{SS} + (1 - \lambda)a_\ell, v_\lambda)$ is a continuous function of λ and it takes all values between v_ℓ and v^{SS} , proving that Property 2 of Lemma 4 holds. Finally, from equations (33) and (34), we know that any value $v \notin [v_\ell, v^{SS}]$ could only be attained by some action a with a continuation value $v' > v^{SS}$, which would lead to nonexistence in subsequent periods. Hence, $[v_\ell, v^{SS}]$ is the largest set that satisfies Properties 1 and 2 of Lemma 4 together with the value V^* , completing the proof that a reconsideration-proof equilibrium has value V^* , and thus that in turn the organizational equilibrium with the state variable is also associated with an action value V^* . Our construction also proved that V^* is the maximal action payoff that can be attained by a constant action.

Finally, suppose that \widehat{V} is strictly quasiconcave. Let a^{SS} be the unique action that attains $\max_a V(a, a, \dots)$. If this steady state is not a Markov equilibrium, then $a^{SS} < \max_a \tilde{V}(a, v^{SS})$. In this case, a sequence that starts at a^{SS} and stays constant violates the no-delay condition.

Finally, we prove part 2 of the proposition. The Ramsey outcome is the allocation that attains the highest payoff, and so by definition an organizational equilibrium cannot do better. If there is no constant allocation that attains the Ramsey outcome, then it means that the best constant allocation attains a payoff strictly smaller than Ramsey; Proposition 5 proves that the payoff of an organizational equilibrium coincides with that of the best constant allocation, and is thus strictly worse than Ramsey as well. When a constant allocation

⁵²If along the equilibrium path, for some $T \geq 0$, $v_T > v^{SS}$, then $v_t > v^{SS}$ for all $t > T$. Since $\{v_t\}$ is bounded and monotonically increasing, the limiting point will be a fixed point of ζ , which is incompatible with v^{SS} being the largest fixed point.

⁵³We have $v_\ell \leq v'_\ell \leq v^{SS}$ by (33) and (34).

a^{SS} attains the Ramsey outcome, it must be the case that

$$a_{SS} \in \arg \max_a V(a, a^{SS}, a^{SS}, \dots);$$

this implies that $(a^{SS}, a^{SS}, a^{SS}, \dots)$ is also a Markov equilibrium, and that a^{SS} achieves the highest payoff among constant allocations, which (by Proposition 5) is also the payoff of an organizational equilibrium. In particular, (a^{SS}, a^{SS}, \dots) is an organizational equilibrium.

A state-independent Markov equilibrium cannot depend on the past nor on calendar time, and so it is a constant sequence (a, a, \dots) . An organizational equilibrium attains the same payoff as the best constant allocation; hence, it can be no worse than the best Markov equilibrium, and is strictly better whenever the best constant allocations do not correspond to a Markov equilibrium.

□

D.4 Proof of Corollary 1

Proof. This proof follows closely that of Proposition 5. Let ζ, V^*, v_ℓ , and v^{SS} be defined as in that proof. The proof of Proposition 5 rules out symmetric equilibria with values higher than V^* by showing that there does not exist a sequence of actions that has a constant value higher than V^* . It also shows how to construct a sequence such that $\tilde{V}(a_0, \hat{V}(a_1, a_2, \dots)) = V^*$ and $\hat{V}(a_0, a_1, \dots) = v$ for any value in $v \in [v_\ell, v^{SS}]$; any such sequence satisfies properties 1 and 2 of Proposition 3. Let $\{\bar{a}_t\}_{t=0}^\infty$ be such that $\tilde{V}(\bar{a}_0, \hat{V}(\bar{a}_1, \bar{a}_2, \dots)) = V^*$ and $\hat{V}(\bar{a}_0, \bar{a}_1, \dots) = v_\ell$. The proof of Proposition 5 establishes that $\tilde{V}(a, v_\ell) \leq V^* \quad \forall a \in A$. Hence, $\{\bar{a}_t\}_{t=0}^\infty$ satisfies Property 3 of Proposition 3 as well. □

E Proofs for Section 3.4

In this appendix, we establish that the slope of the transition function in the organizational equilibrium equals to 1 when approaching to the steady state and equals to 0 when starting at the saving rate in the Markov equilibrium. Furthermore, the slope is positive between the steady state and the Markov saving rate.

Given the transition function (13), the slope of it can be expressed as

$$\frac{\partial s_{t+1}}{\partial s_t} = -\exp \left\{ \frac{-(1-\beta)V^* + \frac{\delta\alpha\beta}{1-\alpha\beta} \log s_t + \log(1-s_t)}{\beta(1-\delta)} \right\} \left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s_t} - \frac{1}{1-s_t}}{\beta(1-\delta)} \right).$$

With a constant saving rate s , the lifetime action payoff is

$$(1-\beta)\bar{V} = \frac{\delta\alpha\beta}{1-\alpha\beta} \log s + \log(1-s) - \beta(1-\delta) \log(1-s).$$

The optimal constant saving rate s^* satisfies

$$\frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s^*} - \frac{1}{1-s^*} - \frac{\beta(1-\delta)}{1-s^*} = 0,$$

and the action payoff V^* satisfies

$$(1-\beta)V^* = \frac{\delta\alpha\beta}{1-\alpha\beta} \log s^* + \log(1-s^*) - \beta(1-\delta) \log(1-s^*).$$

Therefore, we have

$$\begin{aligned} \left. \frac{\partial s_{t+1}}{\partial s_t} \right|_{s_t=s^*} &= -\exp \left\{ \frac{-(1-\beta)V^* + \frac{\delta\alpha\beta}{1-\alpha\beta} \log s^* + \log(1-s^*)}{\beta(1-\delta)} \right\} \left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s^*} - \frac{1}{1-s^*}}{\beta(1-\delta)} \right) \\ &= -(1-s^*) \left(\frac{\frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s^*} - \frac{1}{1-s^*}}{\beta(1-\delta)} \right) \\ &= 1 \end{aligned}$$

In the Markov equilibrium, the saving rate s^M maximizes the part involving only the current saving rate:

$$\frac{\delta\alpha\beta}{1-\alpha\beta} \log s + \log(1-s),$$

which implies that

$$\frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s^M} - \frac{1}{1-s^M} = 0.$$

As a result, $\left. \frac{\partial s_{t+1}}{\partial s_t} \right|_{s_t=s^M} = 0$.

Denote $\chi(s_t) \equiv \frac{\delta\alpha\beta}{1-\alpha\beta} \frac{1}{s_t} - \frac{1}{1-s_t}$. Notice that: (1) $\chi(s_t)$ is decreasing in s_t when $s_t \in (0, 1)$; (2) $\chi(s_t) = 0$ when $s_t = s^M$. It follows that, $\frac{\partial s_{t+1}}{\partial s_t} > 0$ when $s_t > s^M$.

F Example of Approximating Strategy

As an example to illustrate the approximating strategy, we revisit the quasi-geometric discounting economy with partial depreciation and CRRA utility function and apply our approximation strategy. Compared with the environment in Section 2, we modify the period utility function and the law of motion to be

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}, \quad k_{t+1} = f(k_t) - c_t + (1-d)k_t,$$

where $d \in (0, 1)$.

Let s_t denote the saving rate. Mapping to the general setup, we have

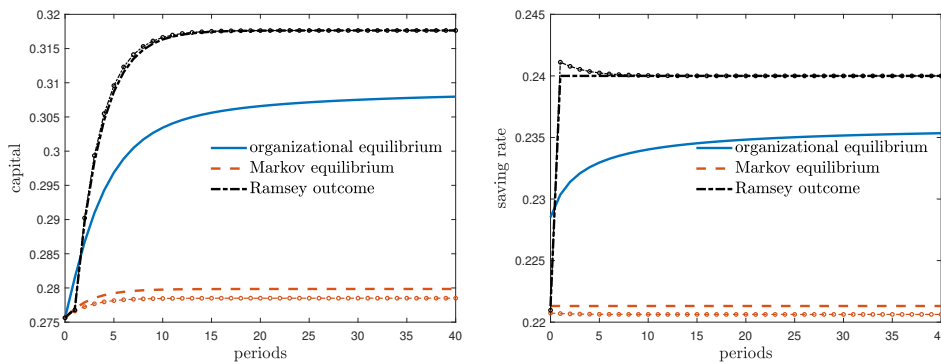
$$\begin{aligned} P(k, s) &= u((1 - s)f(k)), \\ Q(k, s) &= \beta(\delta - 1)u((1 - s)f(k)), \\ F(k, s) &= sf(k) + (1 - d)k. \end{aligned}$$

In this economy, the action s_t and the states k_t are not separable. To proceed, we choose $m(s) = \frac{s^{1-\sigma}}{1-\sigma}$ to approximate the utility function and $g(s) = \log(1 - s)$ to approximate the technology. In this approximating economy, the organizational equilibrium can be constructed.

The blue solid line in Figure displays the transition paths of the capital shock and the saving rate of the organizational equilibrium in this approximating economy. The red dashed line and the black broken line correspond to the Markov equilibrium and the Ramsey outcome in the approximating economy. Similar to our baseline analysis, the organizational equilibrium gradually transits from being close to the Markov equilibrium towards being close to the Ramsey outcome.⁵⁴

To evaluate this approximation, we also compute the the Markov equilibrium and the Ramsey outcome in the original economy via global solutions, which are shown by the lines with circle markers. The outcomes in the approximating economy and the original economy are close to each other not only in the steady states but also along the entire transition paths.

FIGURE 8: Transition Paths in Approximating Economy



Note: The dotted lines in red and black represent the true solutions to the Markov equilibrium and the Ramsey outcome, respectively. The dashed lines in red and black represent the solutions in the approximating economy according to our strategy.

⁵⁴We set $\beta = 0.8$, $\delta = 0.9$, $d = 0.5$, $\sigma = 2$, $\alpha = 0.36$.

G Organizational Equilibrium in Policy Problems

In Section 3, there is one player for each period. Here, the policymaker is still represented by one player for each period, but we also include a continuum of identical households that face a dynamic problem.⁵⁵ In this appendix, we describe explicitly the strategic interaction between the government and the households at different points in time. The game unfolds as follows. In each period, the government in power takes an action $a \in A$ first. Then, the households move simultaneously. Each household takes an action $s \in S$. The aggregate state for next period evolves according to $k' = F(k, a, s)$. A full description would require us to specify what happens when households take different actions, so that, while they are identical ex ante, they may end up being different ex post. However, in most of the applications that are of interest, the household optimization problem has a unique solution. Hence, there can be no equilibrium in which identical households take different actions. Moreover, a deviation by a single household has no effect on aggregates. We exploit these properties and specify the evolution of the economy and preferences only after histories in which (almost) all households have taken the same action. Starting from an arbitrary period t and state k_t , household preferences are given by a function

$$Z(k_t, \{a_v, s_v, s_v^-\}_{v=t}^\infty), \quad (35)$$

where s_v represents the action taken by the individual household, and s_v^- is the action taken by (almost) all other households. We assume that S is a convex compact subset of a locally convex topological linear space and that Z is jointly continuous in all of its arguments (in the product topology), strictly quasiconcave in the own action sequence $\{s_v\}_{v=t}^\infty$, and weakly separable between the state and the remaining arguments. We also assume that household preferences are time consistent. More precisely, we assume that, given an initial level of the state k_t and a sequence of other households' actions $\{a_v, s_v\}_{v=t}^\infty$,

$$\begin{aligned} Z(k_t, \{a_v, s_v, s_v\}_{v=t}^\infty) &= \max_{\{\tilde{s}_v\}_{v=t}^\infty} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^\infty) \implies Z(F(k_t, a_t, s_t), \{a_v, s_v, s_v\}_{v=t+1}^\infty) = \\ &= \max_{\{\tilde{s}_v\}_{v=t+1}^\infty} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t+1}^\infty). \end{aligned} \quad (36)$$

Equation (36) states that, if it is optimal from period t to follow the same sequence of actions that all other households are taking, then it is also optimal to follow that sequence in subsequent periods, as long as other households also continue to do the same. Notice that we exploit the fact that each household has no effect on the aggregates to leave the continuation preferences over several histories unspecified; this is convenient, because it prevents us from having to explicitly introduce individual state variables. To be concrete, consider the taxation game to which we apply this general definition; in that game, s_t is the individual saving rate. Equation (36) is written from the perspective of a household that starts with the same level of k_t as the aggregate, which allows us not to draw a distinction between the two. If that household finds it optimal to follow the same saving rate as all other households, then it will optimally choose to have the same level of k_{t+1} , and equation (36) ensures that the continuation plan will remain optimal from period $t + 1$ onwards. If instead the household chooses a different saving rate from others, then it would potentially enter period $t + 1$ with a different level of the state from the aggregate; however, whenever this choice does not maximize (35),

⁵⁵The notion of an equilibrium can be readily extended to environments with finite types of households or to economies with overlapping generations. Extending organizational equilibrium to economies with a continuum of types could be done by interacting the analysis here with distributional notions of equilibrium as in Jovanovic and Rosenthal (1988).

we know this would not be an optimal individual choice without need to specify the entire continuation path; moreover, the individual deviation does not affect aggregate incentives; hence, we do not need to keep track of it for the purpose of computing other households' best response either.

We define a competitive equilibrium from period t and a state k_t as a sequence $\{a_v, s_v\}_{v=t}^\infty$, such that

$$Z(k_t, \{a_v, s_v\}_{v=t}^\infty) = \max_{\{\tilde{s}_v\}_{v=t}^\infty} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^\infty).$$

Proposition 7. *Given any sequence of policy actions $\{a_v\}_{v=t}^\infty$, a competitive equilibrium exists.*

Proof. Fix k_t and $\{a_v\}_{v=t}^\infty$. Given our assumptions on S and Z , the best-response function

$$br(\{s_v\}_{v=0}^\infty) := \arg \max_{\{\tilde{s}_v\}_{v=t}^\infty} Z(k_t, \{a_v, \tilde{s}_v, s_v\}_{v=t}^\infty)$$

is well defined and continuous. By Brouwer's theorem, it admits a fixed point, which is a competitive equilibrium. \square

Equation (36) ensures that the continuation of a competitive equilibrium is a competitive equilibrium itself. Also, the separability assumption about Z implies that, if $\{a_v, s_v\}_{v=t}^\infty$ is a competitive equilibrium from a state k_t , then it is also a competitive equilibrium from any other state k'_t .

In what follows, we proceed by assuming that the competitive equilibrium is unique given a sequence of policy actions, which can be verified in each specific application.⁵⁶

At time t , government preferences are given by a function $\Psi^g(k_t, a_t, s_t, a_{t+1}, s_{t+1}, a_{t+2}, s_{t+2}, \dots)$. We assume that this function is also weakly separable in k_t and its other arguments. For each given sequence of government actions $\{a_s\}_{s=t}^\infty$, a unique competitive equilibrium exists. The resulting sequence of private sector actions is given by a sequence $\{s_s\}_{s=t}^\infty$, which is independent of k_t , since household preferences are also separable in k_t . We thus specify the government utility from its sequence of actions as that experienced in the competitive equilibrium associated with those actions. With this specification, government preferences can be represented as in equation (8), and an organizational equilibrium can be defined in the same way as in Section 2. Existence of an organizational equilibrium is guaranteed by Proposition 1 when Assumptions 2 and 3 hold. However, these assumptions are significantly more restrictive in tax applications. As is well known, optimal tax problems frequently feature nonconvexities, in which case existence may have to be established in the specific application, as we do in our examples. Moreover, anticipation effects from the competitive equilibrium imply that Assumption 3 often does not hold either. It is worth noting that this assumption can be weakened. Its central role in our proof of Proposition 1 is to establish that the continuation sequence $(a_{t+1}^E, a_{t+2}^E, \dots)$ in equation (20) can be made independent of the current deviation a . In our policy applications, we prove this result

⁵⁶Non-uniqueness can be accommodated by assuming a selection rule on how households coordinate when multiple equilibria are possible, as long as this rule has the properties that the continuation of a selected competitive equilibrium is selected itself as a continuation competitive equilibrium and that the selection is continuous.

by showing instead that the static best-response $\arg \max_a V(a, a_0, a_1, a_2, \dots)$ is independent of the sequence $\{a_t\}_{t=0}^\infty$: hence, any continuation which deters deviation to this action will also be sufficient to deter deviation to any other choice.

As we did for the simpler case of Section 3, we relate an organizational equilibrium to a strategic notion of equilibrium. To do so, we need to keep track of histories of play. A symmetric history of play is a record of all actions taken in the past; we distinguish between histories at which the government is called to play, which are given by $h^0 := \emptyset$ and

$$h^t := (a_0, s_0, a_1, s_1, \dots, a_{t-1}, s_{t-1}), \quad t > 0,$$

and histories at which households are called to play, that take the form of $h^{p,0} := a_0$ and

$$h^{p,t} := (a_0, s_0, a_1, s_1, \dots, a_{t-1}, s_{t-1}, a_t), \quad t > 0.$$

Let H be the set of histories at which the government is called to play, and H^p the set of histories at which households are called to play. For the reasons discussed above, we only keep track of histories in which almost all households have taken the same action.

A strategy for the households is a mapping $\sigma^p : H^p \rightarrow S$; likewise, a government strategy is a mapping $\sigma : H \rightarrow A$. A *symmetric strategy profile* is a pair (σ^p, σ) , representing how all households and the government will act following any symmetric history; it recursively induces a path of play $\{a_t, s_t\}_{t=0}^\infty$.

A symmetric strategy profile (σ^p, σ) is a sequential equilibrium if the following is true:

- Given that the government will follow σ and other households will follow σ^p , the actions dictated by σ^p are optimal for each household. After any history $h^{p,t}$, each household takes as given the government policy action a_t and the initial state k_t , which is recursively determined by the history of past play. Moreover, the strategy σ^p followed by other households and the government strategy σ determine the *future* path of aggregate play, $\{s_v, a_{v+1}\}_{v=t}^\infty$. Household optimality requires that the sequence of actions prescribed by σ^p is optimal along this path: equivalently stated, it requires the actions prescribed by σ^p to be a competitive equilibrium from period t on, following any arbitrary (symmetric) history.
- Given that households will follow the strategy σ^p and that future governments will follow the strategy σ , and given any past history h^t , the current government choice $\sigma(h^t)$ is optimal.

Proposition 8. *Given any organizational equilibrium, there exists a sequential equilibrium whose outcome coincides with the organizational equilibrium.*

Proof. Let $(a_0^*, a_1^*, a_2^*, \dots)$ be an organizational equilibrium, and let (s_0^*, s_1^*, \dots) be the competitive-equilibrium associated with it. We construct a strategy profile recursively as follows:

- $\sigma(\emptyset) = a_0^*$;
- For any $t > 0$ and any history $h^t = (a_0, \dots, a_{t-1})$ such that $a_s = a_s^* \quad \forall s = 0, \dots, t-1$, $\sigma(h^t) = a_t^*$;

- For any $t > 0$ and any history $h^t = (a_0, \dots, a_{t-1})$ such that $\exists s : a_s \neq a_s^*$, define $T := \max\{s < t : a_s \neq \sigma(a_0, \dots, a_{s-1})\}$ and set $\sigma(h^t) = a_{t-1-T}^*$.
- For any history $h_t^p = (a_0, s_0, a_1, s_1, \dots, a_{t-1}, s_{t-1}, a_t)$ at which households are called to play, let $\{a_s^e\}_{s=t+1}^\infty$ be the sequence of government actions that follow from period $t+1$ if the government plays the continuation of the strategy σ defined above following (a_0, \dots, a_t) . Set $\sigma^p(h_t^p)$ to be the competitive equilibrium that is associated with $(a_t, a_{t+1}^e, a_{t+2}^e, \dots)$, which exists and is unique by assumption.

By construction, the household strategy satisfies the second condition for a sequential equilibrium for any history of play. For the government, following any history, the strategy prescribes to play the organizational equilibrium sequence, either from its beginning or from some element a_t^* , $t > 0$. Should the government deviate from its strategy, the continuation strategy restarts the organizational equilibrium sequence from a_0^* . By the definition of an organizational equilibrium, continuing along the sequence is always weakly preferred to playing the best one-shot deviation followed by a restart; hence, the government optimality condition is satisfied and the strategy above describes a sequential equilibrium. \square

H Details for Section 4

We first provide the details on the separable property of the model environment. Recall that the law of motion of the stock of carbon is given by

$$\begin{aligned} q_{1t} &= q_{1t-1} + \varphi_L e_t, \\ q_{2t} &= \varphi q_{2t-1} + (1 - \varphi_L) \varphi_0 e_t. \end{aligned}$$

It follows that

$$\begin{aligned} q_{1t} &= q_{1,-1} + \varphi_L A \sum_{j=0}^t (1 - n_j), \\ q_{2t} &= \varphi^t q_{2,-1} + (1 - \varphi_L) \varphi_0 A \sum_{j=0}^t \varphi^{t-j} (1 - n_j), \\ q_t &= q_{1t} + q_{2t} = q_{1,-1} + \varphi^t q_{2,-1} + A \sum_{j=0}^t (\varphi_L + (1 - \varphi_L) \varphi_0 \varphi^{t-j}) (1 - n_j). \end{aligned}$$

Given a sequence of $\{s_t\}$ and $\{n_t\}$, it implies that

$$\begin{aligned} \log c_t &= \log(1 - s_t) - \gamma q_t + \alpha \log k_t + (1 - \alpha - \nu) \log n_t + \nu \log(A(1 - n_t)) \\ &= \log(1 - s_t) - \gamma \left(q_{1,-1} + \varphi^t q_{2,-1} + A \sum_{j=0}^t (\varphi_L + (1 - \varphi_L) \varphi_0 \varphi^{t-j}) (1 - n_j) \right) \\ &\quad + \alpha \log k_t + (1 - \alpha - \nu) \log n_t + \nu \log(1 - n_t) + \nu \log A \end{aligned}$$

Turn to the part involving the saving rates. The sequence of capital is

$$\log k_t = \alpha^t \log k_0 + \sum_{j=0}^{t-1} \alpha^{t-j-1} \log s_j$$

The lifetime utility is therefore separable between initial states $(q_{1,-1}, q_{2,-1}, k_0)$ and the sequence of labor and saving rates

$$\begin{aligned} U_0 &= \log c_0 + \delta \sum_{j=1}^{\infty} \beta^j \log c_{t+j} \\ &= G(k_0, q_{1,-1}, q_{2,-1}) + W(s_0, s_1, \dots) + V(n_0, n_1, \dots). \end{aligned}$$

Here, $W(s_0, s_1, \dots)$ captures the impact of saving rates. Using the expression for $\log c_t$ and $\log k_t$, we have

$$W(s_0, s_1, \dots) = \log(1 - s_0) + \frac{\delta \alpha \beta}{1 - \alpha \beta} \log(s_0) + \delta \sum_{j=1}^{\infty} \beta^j \left(\log(1 - s_j) + \frac{\alpha \beta}{1 - \alpha \beta} \log(s_j) \right),$$

which is similar to the baseline quasi-geometric discounting model.

Next, consider the part involving the labor choice.

$$\begin{aligned} &V(n_0, n_1, \dots) \\ &= (1 - \alpha - \nu) \log n_0 + \nu \log(1 - n_0) - \gamma A \left(\varphi_L \frac{1 - \beta + \delta \beta}{1 - \beta} + (1 - \varphi_L) \varphi_0 \frac{1 - \varphi \beta + \delta \varphi \beta}{1 - \varphi \beta} \right) (1 - n_0) \\ &\quad + \delta \beta \left(-\gamma A \left(\frac{\varphi_L}{1 - \beta} + \frac{\varphi_0(1 - \varphi_L)}{1 - \beta \varphi} \right) \sum_{j=0}^{\infty} \beta^j (1 - n_{j+1}) + (1 - \alpha - \nu) \sum_{j=0}^{\infty} \beta^j \log n_{j+1} + \nu \sum_{j=0}^{\infty} \beta^j \log(1 - n_{j+1}) \right) \\ &= (1 - \alpha - \nu) \log n_0 + \nu \log(1 - n_0) - \gamma A \left(\varphi_L \frac{1 - \beta + \delta \beta}{1 - \beta} + (1 - \varphi_L) \varphi_0 \frac{1 - \varphi \beta + \delta \varphi \beta}{1 - \varphi \beta} \right) (1 - n_0) \\ &\quad - \beta(1 - \delta) \left((1 - \alpha - \nu) \log n_1 + \nu \log(1 - n_1) - \gamma A (\varphi_L + (1 - \varphi_L) \varphi_0) (1 - n_1) \right) \\ &\quad + \beta V(n_1, n_2, \dots). \end{aligned}$$

Characterization of steady state The organizational equilibrium requires that $V(n_0, n_1, \dots) = V(n_1, n_2, \dots) = \bar{V}$, which leads to

$$\begin{aligned} &-\gamma A \left(\varphi_L \frac{1 - \beta + \delta \beta}{1 - \beta} + (1 - \varphi_L) \varphi_0 \frac{1 - \varphi \beta + \delta \varphi \beta}{1 - \varphi \beta} \right) (1 - n_0) + (1 - \alpha - \nu) \log n_0 + \nu \log(1 - n_0) \\ &= \beta(1 - \delta) \left((1 - \alpha - \nu) \log n_1 + \nu \log(1 - n_1) - \gamma A (\varphi_L + (1 - \varphi_L) \varphi_0) (1 - n_1) \right) + (1 - \beta) \bar{V}. \end{aligned}$$

In the steady state, the best constant action maximizes the following object

$$n^O = \operatorname{argmax}_n -\gamma A \left(\varphi_L \left(\frac{1-\beta+\delta\beta}{1-\beta} - \beta(1-\delta) \right) + (1-\varphi_L)\varphi_0 \left(\frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} - \beta(1-\delta) \right) \right) (1-n) \\ + (1-\alpha-\nu)(1-\beta(1-\delta)) \log n + \nu(1-\beta(1-\delta)) \log(1-n).$$

The first-order condition implies

$$\Lambda^O + (1-\alpha-\nu) \frac{1}{A} \frac{1}{n^O} = \nu \frac{1}{A} \frac{1}{1-n^O},$$

where Λ^O in the organizational equilibrium is given by

$$\Lambda^O \equiv \gamma \left(\varphi_L \left(1 + \frac{\delta\beta}{(1-\beta)(1-\beta(1-\delta))} \right) + (1-\varphi_L)\varphi_0 \left(1 + \frac{\delta\varphi\beta}{(1-\varphi\beta)(1-\beta(1-\delta))} \right) \right).$$

When $\delta = 1$, the steady-state policy reconciles with the outcome characterized in [Golosov et al. \(2014\)](#), which corresponds to the Ramsey outcome in the long-run

$$\Lambda^R \equiv \gamma \left(\varphi_L \frac{1}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1}{1-\varphi\beta} \right).$$

The allocation of labor in the Markov equilibrium solves the action payoff taken future labor choice as given

$$n^M = \operatorname{argmax} -\gamma A \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right) (1-n) + (1-\alpha-\nu) \log n + \nu \log(1-n).$$

The implied tax is

$$\Lambda^M = \gamma \left(\varphi_L \frac{1-\beta+\delta\beta}{1-\beta} + (1-\varphi_L)\varphi_0 \frac{1-\varphi\beta+\delta\varphi\beta}{1-\varphi\beta} \right).$$

I Details for Section 5

Static equilibrium Recall that the final goods is an aggregator of the two intermediate goods

$$y_t = [0.5^{1-\rho} m_{1t}^\rho + 0.5^{1-\rho} m_{2t}^\rho]^{\frac{\rho-1}{\rho}}.$$

The aggregate price index \mathcal{P}_t is given by

$$\mathcal{P}_t = [0.5 p_{1t}^{\frac{\rho}{\rho-1}} + 0.5]^{\frac{\rho-1}{\rho}}$$

. The demand schedules for goods 1 and 2 satisfy

$$m_{1t} = \frac{1}{2} \left(\frac{p_{1t}}{\mathcal{P}_t} \right)^{\frac{1}{\rho-1}} y_t, \quad m_{2t} = \frac{1}{2} \left(\frac{1}{\mathcal{P}_t} \right)^{\frac{1}{\rho-1}} y_t,$$

The production functions in the two sectors are given by

$$y_{1t} = AL_{1t}^{1-\alpha} k_{1t}^\alpha k_t^{1-\alpha}, \quad \text{and} \quad y_{2t} = L_{2t}^{1-\alpha} k_{2t}^\alpha k_t^{1-\alpha}.$$

We impose fixed labor input in the two sectors, i.e., $L_{1t} = L_{2t} = 1$.

Denote the tariff rate as τ_t . In the home country, the price of goods 2 is 1 by normalization. Therefore, the price of goods 2 at the foreign country is $\frac{1}{1+\tau_t}$ due to the law of one price and the symmetric tariff rate assumption. Again, by symmetry, the price of goods 1 in country 1 is also $p_{1t} = \frac{1}{1+\tau_t}$. That is, a higher tariff rate lowers the more productive sector's relative price. It also implies that there is a one-to-one mapping between the tariff rate τ_t and the goods 1 price p_{1t} .

Since capital is free to flow across sectors, the return to capital is equalized across the two sectors

$$p_{1t} A k_{1t}^{\alpha-1} k_t^{1-\alpha} = k_{2t}^{\alpha-1} k_t^{1-\alpha}$$

which leads to

$$k_{2t} = \phi_t k_t \quad \text{where} \quad \phi_t = \frac{1}{1 + (p_{1t} A)^{\frac{1}{1-\alpha}}}.$$

It also follows that the output in the two sectors are given by

$$y_{1t} = A(1 - \phi_t)^\alpha k_t, \quad y_{2t} = \phi_t^\alpha k_t,$$

and the nominal GDP in country 1 can be expressed as

$$v_t = p_{1t} A k_{1t}^\alpha k_t^{1-\alpha} + k_{2t}^\alpha k_t^{1-\alpha} = k_t (p_{1t} A (1 - \phi_t)^\alpha + \phi_t^\alpha) = k_t \phi_t^{\alpha-1}.$$

The total expenditure and the nominal GDP is identical, which implies that

$$v_t = \mathcal{P}_t y_t = p_{1t} m_{1t} + m_{2t},$$

where the total usage of goods 1 and 2 follows from the demand schedule

$$m_{1t} = \frac{1}{2} \left(\frac{p_{1t}}{\mathcal{P}_t} \right)^{\frac{1}{\rho-1}} \frac{v_t}{\mathcal{P}_t}, \quad m_{2t} = \frac{1}{2} \left(\frac{1}{\mathcal{P}_t} \right)^{\frac{1}{\rho-1}} \frac{v_t}{\mathcal{P}_t}.$$

By symmetry, the terms of the trade is 1 in equilibrium. Denote the tariff revenue as \mathcal{R} , which can be derived as

$$\mathcal{R}_t = (m_{2t} - y_{2t}) \left(\frac{1}{p_{1t}} - 1 \right) = k_t \phi_t^\alpha \left(\frac{1}{\phi_t} \frac{1}{2} \mathcal{P}_t^{-\frac{\rho}{\rho-1}} - 1 \right) \left(\frac{1}{p_{1t}} - 1 \right).$$

Under the assumption that the tariff revenue is equally split between the two groups of workers, the consump-

tion of workers in industry 1 and 2 are

$$c_{1t} = \frac{\frac{1}{2}\mathcal{R}_t + w_{1t}}{\mathcal{P}_t} = \frac{\frac{1}{2}\mathcal{R}_t + (1 - \alpha)p_{1t}A(1 - \phi_t)^\alpha k_t}{\mathcal{P}_t} = \chi_1(\tau_t)k_t, \quad (37)$$

$$c_{2t} = \frac{\frac{1}{2}\mathcal{R}_t + w_{2t}}{\mathcal{P}_t} = \frac{\frac{1}{2}\mathcal{R}_t + (1 - \alpha)\phi_t^\alpha k_t}{\mathcal{P}_t} = \chi_2(\tau_t)k_t, \quad (38)$$

where p_{1t} , \mathcal{P}_t , ϕ_t are all functions of τ_t .

Capitalists The real return to capital is

$$r(\tau_t) = \frac{\alpha k_{2t}^{\alpha-1} k_t^{1-\alpha}}{\mathcal{P}_t} = \frac{\alpha \phi_t^{\alpha-1}}{\mathcal{P}_t}.$$

The problem of capitalists can be written as

$$\max \sum_{t=0}^{\infty} \beta^t \frac{c_t^{1-\sigma}}{1-\sigma}$$

subject to

$$c_t + k_{t+1} = (r(\tau_t) + 1 - \delta)k_t.$$

The Euler equation is

$$c_t^{-\sigma} = \beta(r(\tau_{t+1}) + 1 - \delta)c_{t+1}^{-\sigma}.$$

Denote s_t as the saving rate, the Euler equation can be expressed as

$$\left(\frac{1-s_t}{s_t}\right)^{-\sigma} = \beta(r(\tau_{t+1}) + 1 - \delta)^{1-\sigma} (1-s_{t+1})^{-\sigma}.$$

Meanwhile, given a sequence of saving rates and tariff rates, the capital evolution satisfies

$$k_t = \prod_{j=0}^{t-1} s_j (r(\tau_j) - \delta) k_0.$$

Welfare Suppose the policy maker's preference is to maximize

$$U = \sum_{t=0}^{\infty} \beta^t \left(\lambda \log c_{1t} + (1 - \lambda) \log c_{2t} \right)$$

Define

$$\chi(\tau) = \lambda \log \chi_1(\tau) + (1 - \lambda) \log \chi_2(\tau),$$

where $\chi_1(\cdot)$ and $\chi_2(\cdot)$ are defined in equation (37) and (38).

The total welfare is separable between capital and the trade policy

$$U = \frac{1}{1-\beta} \log k_0 + \sum_{t=0}^{\infty} \beta^t \chi(\tau_t) + \frac{\beta}{1-\beta} \sum_{t=0}^{\infty} \beta^t (\log s_t + \log(r(\tau_t) + 1 - \delta)).$$

Markov Equilibrium The policy maker in the Markov equilibrium takes future tariff rates as given and simply maximizes the following object

$$\chi(\tau) + \frac{\beta}{1-\beta} \log(r(\tau) + 1 - \delta)$$

which yields a constant tariff. Note that the impact of tariff on the saving rate is not taken into account.

Organization Equilibrium The steady-state allocation in the organizational equilibrium satisfies

$$\max_{s, \tau} \chi(\tau) + \frac{\beta}{1-\beta} \log(r(\tau) + 1 - \delta) + \frac{\beta}{1-\beta} \log s$$

subject to

$$\left(\frac{1-s}{s} \right)^{-\sigma} = \beta (r(\tau) + 1 - \delta)^{1-\sigma} (1-s)^{-\sigma}$$

With $\sigma < 1$, the saving rate is decreasing in τ , which gives the policy maker a larger incentive to lower the tariff rate.

Ramsey Outcome With commitment, the policy maker's problem is

$$\max_{\{\tau_0, \tau_1, \dots\}} \sum_{t=0}^{\infty} \beta^t \chi(\tau_t) + \frac{\beta}{1-\beta} \sum_{t=0}^{\infty} \beta^t (\log s_t + \log(r(\tau_t) + 1 - \delta))$$

subject to

$$\left(\frac{1-s_t}{s_t} \right)^{-\sigma} = \beta (r(\tau_{t+1}) + 1 - \delta)^{1-\sigma} (1-s_{t+1})^{-\sigma}$$

Let $\beta^t \mu_t$ denote the multiplier associated with the constraint involving s_t and s_{t+1} . For τ_0 , the choice is to maximize

$$\chi(\tau_0) + \frac{\beta}{1-\beta} \log(r(\tau_0) + 1 - \delta)$$

For s_0 , the first-order condition is

$$\frac{\beta}{1-\beta} \frac{1}{s_0} = -\mu_0 \sigma \left(\frac{1-s_0}{s_0} \right)^{-\sigma-1} \frac{1}{s_0^2}$$

For $t \geq 1$, the first-order condition with respect to τ_t is

$$\beta^t \left(\chi_{\tau}(\tau_t) + \frac{\beta}{1-\beta} \frac{r_{\tau}(\tau_t)}{r(\tau_t) + 1 - \delta} \right) = \beta^{t-1} \mu_{t-1} (1-\sigma) \beta (r(\tau_t) + 1 - \delta)^{-\sigma} (1-s_t)^{-\sigma} r_{\tau}(\tau_t).$$

The first order condition with respect to s_t is

$$\beta^t \left(\frac{\beta}{1-\beta} \frac{1}{s_t} \right) = \beta^{t-1} \mu_{t-1} \sigma \beta (r(\tau_t) + 1 - \delta)^{1-\sigma} (1 - s_t)^{-\sigma-1} - \beta^t \mu_t \sigma \left(\frac{1 - s_t}{s_t} \right)^{-\sigma-1} \frac{1}{s_t^2}.$$

These conditions characterize the transition dynamics.