American Economic Review 2024, 114(11): 3427–3468
https://doi.org/10.1257/aer.20201365

Institution Building without Commitment†

By Marco Bassetto, Zhen Huo, and José-Víctor Ríos-Rull*

We propose a theory of gradualism in the implementation of good policies, suitable for environments featuring time consistency. We downplay the role of the initial period by allowing agents both to wait for future agents to start equilibrium play and to restart the equilibrium by ignoring past history. The allocation gradually transits toward one that weighs both short- and long-term concerns, stopping short of the Ramsey outcome but greatly improving upon Markovian equilibria. We use the theory to account for the slow emergence of both climate policies and the reduction of global tariff rates. (*JEL* C73, E21, E61, F13, H30)

> 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 9, 1950

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 policymaker to delegate to future policymakers the implementation because of the initial costs involved in carrying out a fast adoption.[1] To make the case, we postulate environments where policymakers 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

* Bassetto: Federal Reserve Bank of Minneapolis (email: bassetto@nber.org); Huo: Department of Economics, Yale University (email: zhen.huo@yale.edu); Ríos-Rull: Department of Economics, University of Pennsylvania, CAERP, CEPR, NBER, and University College London, Ronald O. Perelman Center for Political Science and Economics (email: vr0j@upenn.edu). Mikhail Golosov was the coeditor for this article. We are grateful to three anonymous referees and the editor for extensive feedback. For useful comments, we thank Jess Benhabib, Gadi Barlevy, 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 many conferences and seminars. An older version of this paper circulated under the title "Organizational Equililbrium 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.

† Go to https://doi.org/10.1257/aer.20201365 to visit the article page for additional materials and author disclosure statement(s).

[1] As 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 the 2008 G8 Outreach Session (https://www.fmprc.gov.cn/mfa_eng/xw/zyjh/202405/t20240530_11339984.html, accessed September 11, 2024).

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 policymakers have to address an issue where short-term costs are at odds with long-term benefits and a time-inconsistency problem arises. 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 reversions 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 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 yields 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 tax,[3] 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 Golosov 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 gradual 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,

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

[3] Based on the World Bank Carbon Pricing Dashboard, we document that (i) the carbon taxes rise gradually over the last 30 years in Scandinavian countries that first initiated the carbon tax system and (ii) 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.

we see the implementation of free trade as a conflict between the import substitute producers and long-run growth (Autor, Dorn, and Hanson 2013; Acemoglu et al. 2016; Pierce and Schott 2016; Caliendo, Dvorkin, and Parro 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 gradual 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 who see themselves in a similar spot —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 and eliminates the intrinsic advantage that is often associated with the first agent. A second condition, a *no-delay condition*, requires that the first agent does not want to sit out and let the equilibrium process start in the following period, effectively inducing gradualism. If the sequence of equilibrium allocation favored the society's long-run interest with the current generation making sacrifices, the initial policymaker would prefer to delay the implementation to future agents. The no-delay condition excludes this type of outcome. 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 *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 0 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 de-emphasize the preeminent role of the time 0 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 0 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 sequential 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 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, Quadrini, and Ríos-Rull 2005; Bassetto and Sargent 2006; Klein, Krusell, and Ríos-Rull 2008; Bassetto 2008; Krusell, Kuruşçu, and Smith 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, Ray, and Yeltekin (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) and Acemoglu, Egorov, and Sonin (2012, 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 where the notion of no-delay

[4] We 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 forever, 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.

[5] On the technical side, while the approach in Abreu, Pearce, and Stacchetti (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.

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 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 "rescaled actions" 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 it 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, 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

[6] Kocherlakota (1996) 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.

[7] The same type of separability property is also explored in Halac and Yared (2014) and Halac and Yared (2017), but they focus on different subgame-perfect equilibrium refinements where agents' type is private information.

[8] Our approach encompasses the more specific cases introduced by Brendon and Ellison (2018) in their latest version to account for state variables.

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 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 vanishes 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 I. We define organizational equilibrium for separable economies in Section II, where we also describe the connections to game theory and the strategy to study nonseparable economies by using approximating separable economies. We also show how to tackle economies where a large decision-maker such as 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 III 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 multicountry model generates them as an organizational equilibrium in Section IV. Section V concludes.

I. 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 IIF. 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 $k_t = k_0^{\alpha^t} \prod_{j=0}^{t-1} s_j^{\alpha^{t-j-1}}$. The corresponding lifetime utility for the agent in period 0 is

$$\begin{aligned} U(k_0, s_0, s_1, \ldots) &= \log\left([1 - s_0] k_0^{\alpha}\right) + \delta \sum_{j=1}^{\infty} \beta^j \log\left([1 - s_j] k_j^{\alpha}\right) \\ &= \frac{\alpha(1 - \alpha\beta + \delta\alpha\beta)}{1 - \alpha\beta} \log k_0 + V(s_0, s_1, \ldots), \end{aligned}$$

[9] Jackson 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 online Appendix C.

where

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

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

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

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, $s_0^R < \alpha\beta$.

[10] While 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.

[11] More 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 nondifferentiable 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

$$u_c\big(f(k)-g(k)\big) \;=\; \beta u_c\Big(f\big(g(k)\big)-g\big(g(k)\big)\Big)\Big[\delta f_k\big(g(k)\big)+(1-\delta)\,g_k\big(g(k)\big)\Big],$$

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

$$s^{M} \;=\; \frac{\alpha\delta\beta}{1-\alpha\beta+\delta\alpha\beta}. \tag{3}$$

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

$$\max_{s}\;\; u\big(c_0\big)+\delta\sum_{t=1}^{\infty}\beta^{t}u\big(c_t\big)$$

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

The solution to this problem is given by

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

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

Toward 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 $V\big(s^{M},s^{R},s^{R},\ldots\big)$ is higher than $V\big(s^{R},s^{R},s^{R},\ldots\big)$ and it is feasible for future agents as well. If a time-t agent were 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;[12] in this equilibrium 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 $V\big(s^M, s^B, s^B, \cdots\big) > V\big(s^B, s^B, \cdots\big)$; by sitting out and playing Markov, 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 $\{s_0^*, s_1^*, s_2^*, \cdots\}$, such that

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

and also

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

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

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

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.[13] We will show in Section IID that the maximum constant value of $\bar{V}$ that can be attained is given by $V\big(s^B, s^B, s^B, \cdots\big)$, 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 that restarting from s_0 amounts to a sufficient (although often stronger than necessary) punishment to deter any deviation along the transition.

[12] The outcome in such best subgame-perfect equilibrium starts with $s_0 = s^M$, and the subsequent saving rates are larger than s^M but could be lower than s^R.

[13] Our 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).

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^*:

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

The solid red line in Figure 1 visualizes this recursive dynamic between s_t and s_{t+1}.[14] Following this relationship, there are a continuum of paths indexed by different starting points 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.

II. Organizational Equilibrium

This section contains the core of our equilibrium analysis. After we describe the economic environment (preferences and technology), we set up 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 IIB, 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 IIC, 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 who 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 $U(k_t, a_t, a_{t+1}, a_{t+2}, \ldots)$. The first assumption is that functions U and F are independent of calendar time, which allows meaningful welfare comparisons across decision-makers.

[14] The dashed blue lines represent condition (7) associated with smaller $\bar{V}$.

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:

(*i*) *At any point in time t, the set of feasible actions A is independent of the state* k_t;

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

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

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 rescale actions in such a way that it does.[15] 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.[16] 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

$$\tilde{a}_t = \gamma(a_t, k_t), \tag{9}$$

$$k_{t+1} = F(k_t, \tilde{a}_t).$$

A. *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 $h^0 := \varnothing$. 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

[15] Notice 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 I, any monotone transformation of the saving rate would be an equally valid action, yielding the same equilibrium outcomes.

[16] Note 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.

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 shorthand notation:

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

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*

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

  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 2 corresponds to Kocherlakota's (1996) definition of reconsideration-proof equilibrium. In particular, no restarting is the natural adaptation of Kocherlakota's (1996) 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 that 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 (2018) 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.[17] 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 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 A.1 in online Appendix A proves existence of an equilibrium satisfying Requirements 1 and 2 under the following technical assumption.

ASSUMPTION 2: *A is a convex compact subset of a locally convex topological linear space with topology ρ_x.*

- *V is quasi-concave over A^∞.*
- *V is continuous over A^∞ with respect to the product topology ρ_x^∞.*

There can be many such equilibria. In Section I, 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

[17] As an example, no reconsideration-proof equilibrium would exist in the example of Section I.

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 IIB, 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* $a \in A$,

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

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 that is assumed to have played their best one-shot response before time 0. 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, 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 $\tilde{V} : A \times \mathbb{R} \to \mathbb{R}$ and $\hat{V} : A^{\infty} \to \mathbb{R}$ such that, for all sequences $\big(a_t, a_{t+1}, a_{t+2}, \ldots\big) \in A^{\infty}$,*

$$V\big(a_t, a_{t+1}, a_{t+2}, \ldots\big) = \tilde{V}\Big(a_t, \hat{V}\big(a_{t+1}, a_{t+2}, \ldots\big)\Big), \tag{10}$$

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; online 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 that 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 online Appendix A.2.

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 $a_{s+1}, a_{s+2}, \ldots$. 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 IV, 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 toward cooperation from this new, lower level.

B. *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 0 agent. This game is another representation

of environments with time-consistency problems, which 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 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 0 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 0 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 0 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 IIA 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 $(a_0, \ldots, a_{t-1})$.[18] 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 $(a_{\tilde{t}}, \ldots, a_t)$, or it can start a new organization, in which case only a_t is known to player $t+1$, or it can

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

discontinue the current organization without replacing it, in which case player $t+1$ cannot condition its actions on any of the past actions $(a_0, \ldots, a_{t-1})$.[19]

With the limitations on record keeping described above, the game unfolds otherwise as in Section IIA, 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 IIA. To quickly distinguish between the two games, we will from now on refer to the game of Section IIA 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.[20] We relegate the notation defining histories, information sets, and strategies to online 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:

(*i*) *At any point in time t, the set of feasible actions A is independent of the state k_t;*

(*ii*) *There exist functions $\bar{v} : K \times \mathbb{R}_{++}$, $\bar{\bar{v}} : K \times \mathbb{R}$, and $V : A^{\infty} \to \mathbb{R}$, such that*

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

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 $(a_{\hat{t}}, a_{\hat{t}+1}, a_{\hat{t}+2}, \ldots)$ be the path implied by the equilibrium, conditional on $\hat{t}$. Then $(a_{\hat{t}}, a_{\hat{t}+1}, a_{\hat{t}+2}, \ldots)$ is an organizational equilibrium.*

PROOF:

See online 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 IIA also describes the equilibrium path of this new game from the point at which record keeping becomes possible. The no-delay

[19] Even 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.

[20] Note, 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 independent 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.

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 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.

C. *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. Online Appendix C discusses the role of these conditions in greater detail and shows how to apply similar ideas to characterize organizational equilibria in examples where these specific assumptions fail.

The following proposition 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.*

(*i*) *No restarting:*

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

(*ii*) *Optimality: No other sequence satisfying no restarting achieves a higher constant value;*

(*iii*) *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.

PROOF:

See online Appendix D.1.

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 three 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 I are true more generally.

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

$$\hat{V}(a_t, a_{t+1}, a_{t+2}, \ldots) \equiv W\big(a_t, \hat{V}(a_{t+1}, a_{t+2}, \ldots)\big). \tag{12}$$

This assumption of course holds in the example of Section I and in many other applications of economic interest.[21] 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

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

[21] The key difference between Assumption 3 and Assumption 5 is that the same function $\hat{V}$ 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 $\hat{V}$ 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 Assumption 5 fails, computing an organizational equilibrium may be more challenging, but often there are alternative ways of making the problem recursive. As an example, online 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.

PROOF:
See online Appendix D.2.

Proposition 4 uses values as a state variable in ways similar to Abreu, Pearce, and Stacchetti (1986, 1990)—henceforth, 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 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:*

(*i*) *The action payoff of an organizational equilibrium coincides with the maximum value that* $V(a,a,a,\dots)$ *can attain under a constant action.*

(*ii*) *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.*

(*iii*) $\hat{V}(a_t,a_{t+1},a_{t+2},\dots)$ *is increasing over time for any organizational equilibrium, and it converges to the value associated with the constant action profile that maximizes* $V(a,a,a,\dots)$.

(*iv*) *If* $\hat{V}$ *is a strictly quasi-concave 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* $\hat{V}(a_0,a_1,a_2,\dots)$ *is strictly below* $\hat{V}(a,a,a,\dots)$.

PROOF:
See online Appendix D.3.

COROLLARY 1: *Under Assumptions 2, 3, and 5, there exists a sequence* $\{\bar{a}_t\}_{t=0}^{\infty}$ *that satisfies Properties* (*i*), (*ii*), *and* (*iii*) *of Proposition 3.*

PROOF:
See online Appendix D.4.

Part (i) 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 (ii) 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. Parts (iii) and (iv) in turn show that convergence takes time. Whether this is interpreted as a gradual buildup 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 I 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).

D. *Application to the Growth Model*

We now apply the results that we derived for a generic case to the specific application of Section I. 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 $1 - \epsilon$.[22] 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

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

[22] We 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.

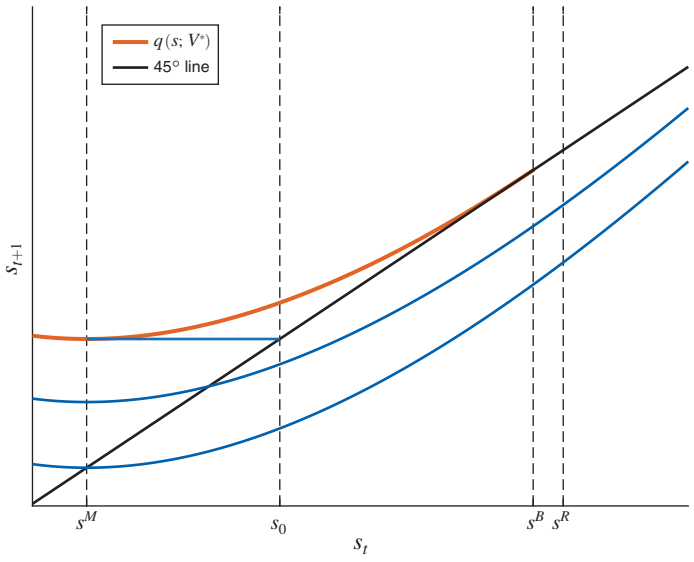

FIGURE 1. EVOLUTION OF THE SAVING RATE FOR SEQUENCES THAT ATTAIN A CONSTANT VALUE

so that equation (7) can be rewritten as $s_{t+1} = q\left(s_t; \bar{V}\right)$. 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; by Proposition 5, this is the payoff attained by the action s that maximizes $V(s, s, s, \ldots)$, 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^*.[23]

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 $V(s_0, s_1, \ldots) = V\left(s^B, s^B, \ldots\right) = V^*$ 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 $V\left(s^M, s_0, s_1, \ldots\right) \leq V^*$, so that the initial agent would not gain from waiting for the sequence to start in the next period.[24] To guarantee the no-delay condition is satisfied, the sequence of saving rates has to start with a low enough point to prevent the initial agent from sitting

[23] For 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.

[24] In general, the no-delay condition requires $\max_s V(s, s_0, s_1, \ldots) \leq V(s_0, s_1, \ldots)$. 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.

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.[25]

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 that 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 $\delta \to 1$, 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 zero when $s = s^M$, and equals to one when $s = s^B$.*

PROOF:

See online Appendix E.

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 and the Markov equilibrium.[26]

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.,

[25] For 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 that 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).

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

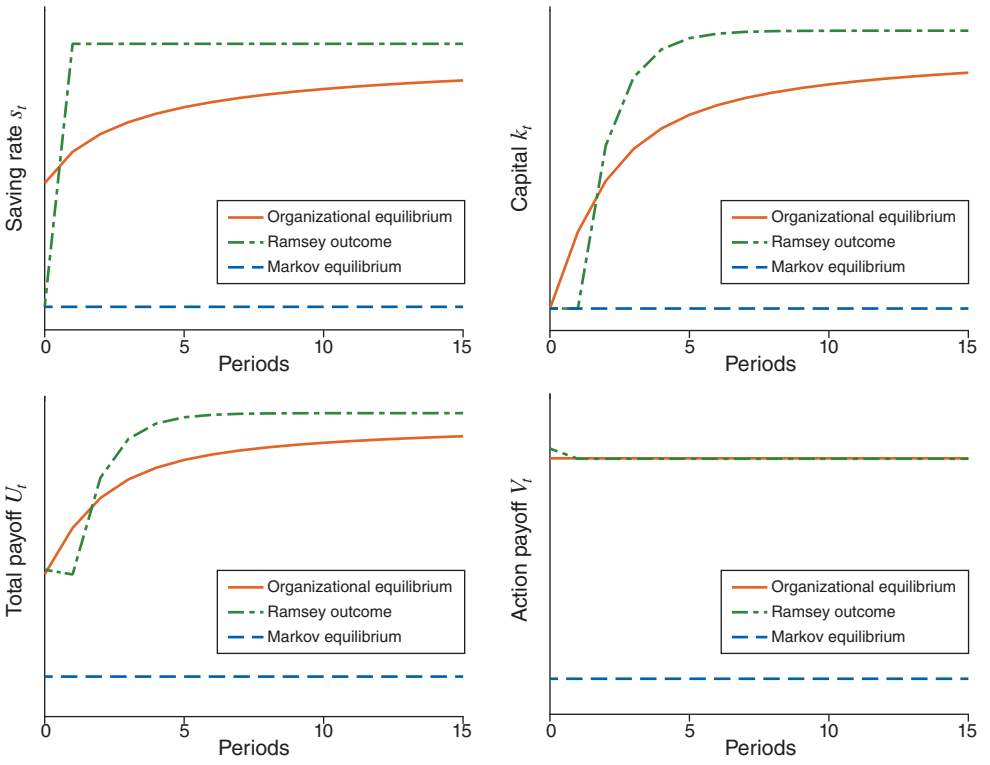

FIGURE 2. TRANSITION PATHS IN VARIOUS EQUILIBRIA

$k_0 = k^M$. The top 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 onward, 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

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

The total payoff U_t and the action payoff V_t are depicted in the bottom 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 payoff, while eventually capital falls below the Ramsey outcome, and late generations lose from this.

E. *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 work 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:

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

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 $k_{t+1} \;=\; F\big(k_t, a_t\big)$. 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 lifetime 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 $\bar{k} = F(\bar{k}, \bar{a})$.[27] We propose a relatively flexible first-order approximation of our original economy around any such stationary point:

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

That is, we allow flexible function choices when approximating the actions a_t in the preference and the technology, respectively.[28] 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

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

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

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

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

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

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

[28] The choice of the approximating function for the state k_t is irrelevant, as eventually the equilibrium is about the action a_t.

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 satisfies[29]

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

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 online Appendix F, we further illustrate this strategy in a nonseparable 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 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.

F. *Agents as Policymakers*

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 III and IV. 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.[30]

[29] Such 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.

[30] The setup here uses ideas from Stokey (1991).

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 $\Psi(k_t, a_t, s_t, a_{t+1}, s_{t+1}, a_{t+2}, s_{t+2}, \ldots)$.[31]

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

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

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

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

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 online Appendix G, we adapt the analysis of Section II to this new environment, and we provide an explicit game that describes the strategic interaction between the government and the households over time.[32]

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, which we adopt in what follows.

In Sections IIA and IIB, 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

[31] As in Section II, sometimes it may be necessary to transform the original government action so that it is feasible independent 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 rescaling may be needed for the household choices.

[32] In 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).

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.

III. Climate Change Mitigation

The discussion of climate change policies 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 have led to increasingly ambitious goals. These include the UN Framework Convention on Climate Change in 1992, the Kyoto Protocol in 2005, the Paris Agreement in 2015, and the recent agreement to end fossil fuels at 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 Golosov et al. (2014) when specifying the preferences, the production technology, and the carbon evolution dynamics. However, we depart from their work in that we allow quasi-geometric discounting, which captures the idea that short-sighted politicians may have a bias in favor of current generations. Assume that the objective of the policymaker at period t is given by

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

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

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

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

The use of the energy input generates carbon in the atmosphere, which negatively affects production efficiency. We denote the amount of carbon accumulated in the atmosphere as q_t. The parameters α and ν control the relative importance of various production inputs, and γ determines the extent to which carbon damages 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 Golosov et al. (2014), coal is the most important type of fossil fuel energy that calls for government intervention, and its supply is practically unlimited.

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 of 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 $\{s_0, s_1, \ldots\}$, a sequence of labor allocation $\{n_0, n_1, \ldots\}$, and the initial state $(k_0, q_{1,-1}, q_{2,-1})$, the payoff of the policymaker at time 0 can be expressed as

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

where the part involving the allocation of labor on energy production can be expressed as[33]

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

[33] The expression for W is given by

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

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, characterizing the equilibrium allocation of labor is sufficient to understand the dynamics of the carbon tax Λ_t.

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

$$\text{Ramsey outcome:} \quad \Lambda^R = \gamma\left(\frac{\varphi_L}{1-\beta} + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta}\right),$$

$$\text{Markov equilibrium:} \quad \Lambda^M = \gamma\left[\frac{\varphi_L}{1-\beta}\big(1-\beta+\delta\beta\big) + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta} \times \big(1-\varphi\beta+\delta\varphi\beta\big)\right],$$

$$\text{Organizational equilibrium:} \quad \Lambda^O = \gamma\left[\frac{\varphi_L}{1-\beta}\left(1-\beta+\frac{\delta\beta}{1-\beta(1-\delta)}\right) + \frac{(1-\varphi_L)\varphi_0}{1-\varphi\beta}\left(1-\varphi\beta+\frac{\delta\varphi\beta}{1-\beta(1-\delta)}\right)\right],$$

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

Let us unpack these expressions. In 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 Golosov et al. (2014). In the Markov economy, the additional discounting resulting from $\delta < 1$ induces the policymaker to behave in a more myopic way, as captured by the terms $1-\beta+\delta\beta < 1$ and $1-\varphi\beta+\delta\varphi\beta < 1$. Finally, in the organizational equilibrium, each policymaker directly takes into account the welfare of all subsequent policymakers 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 Golosov et al. (2014) and set $\delta = 0.95$ at annual frequency to allow the time-inconsistency problem 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

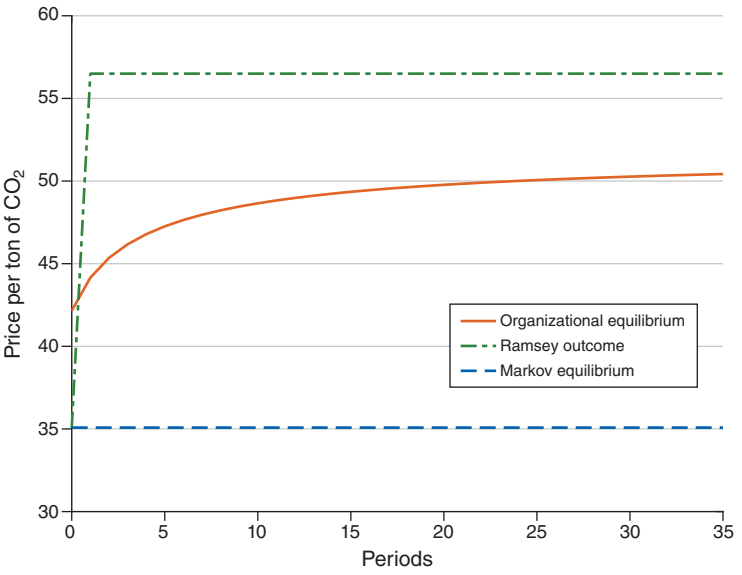

FIGURE 3. TRANSITION PATH OF CARBON TAX

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

an organizational equilibrium, policymakers overcome the temptation to ignore the carbon externality only slowly and partially. In contrast, such implicit coordination is absent in the Markov equilibrium, and the carbon tax remains constant at a lower level. In the Ramsey outcome, the initial policymaker enjoys the privilege and sets a low tax level at the beginning, but a sudden jump to a high level follows immediately after that.

The gradualism resulting from the trade-off between current and future policymakers in the organizational equilibrium speaks to the observed evolution in the data.[34] The left panel of Figure 4 presents the average carbon tax rates over time in Scandinavian countries, which were the first group of countries that implemented the carbon tax. Instead of a steady level or a disruptive jump, the carbon tax rates in these countries display a gradually 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 tax. Ceteris paribus, the carbon tax rate is increasing in the years since implementation in an organizational equilibrium. After controlling for the quantity of carbon emissions covered by the tax and the GDP per capita, we find that the carbon tax rates are positively correlated with the years since implementation, which lends additional credibility to our theory. The right panel of Figure 4 displays the corresponding binscatter plot.

[34] The main data source for the carbon tax is the Carbon Pricing Dashboard (World Bank 2023).

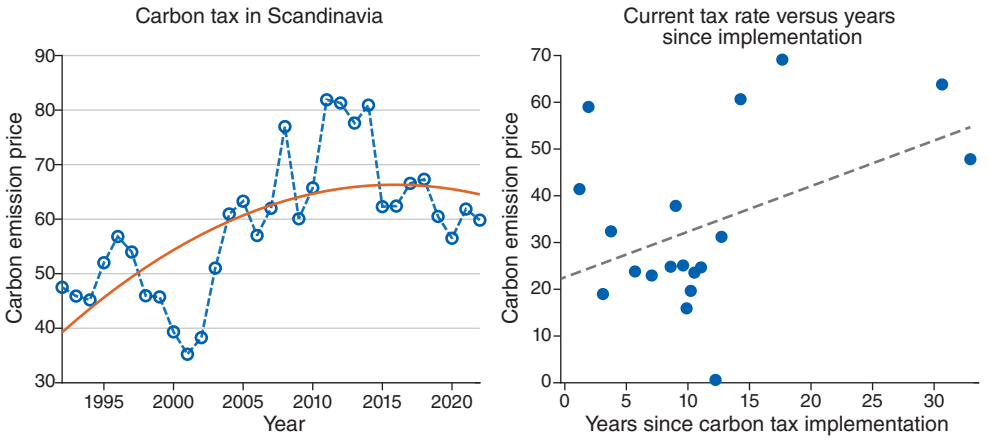

FIGURE 4. EVOLUTION OF CARBON TAXES

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 a country-specific GDP deflator (World Bank 1990–2021), with 2010 as the base year for prices. The red solid line represents a quadratic fit of the data over time. In the right panel, we consider the following regression: $\tau_j = \beta_0 + \beta_1 t_j +$ controls $+$ 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 CO_2 emission and the GDP per capita (World Bank 2022) as controls. The estimated coefficient β_1 is 1.11, with p-value 0.017.

Another important case is the gradual change of climate policy in China, which has the most GHG emissions and just started its emissions trading system in 2021. Speeches by successive presidents reveal the trade-off between current and future generations. In 2008, 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."[35] These statements reflect the conflict of interest between current and future generations and signal the incentive to delay the process of reducing carbon emissions. Our theory provides a natural solution that balances the interests across generations via a gradual reform.

IV. Tariff Policies

Like measures to address 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 came to cover an increasingly broad set of protectionist policies. We describe here a simple model that captures

[35] These remarks are from Hu Jintao's remarks at the G8 Outreach Session on July 9, 2008, and from Xi Jinping's speech at the United Nations General Assembly on September 22, 2020.

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 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

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

where ℓ_{it} and k_{it} are capital and labor used by firm i and k_t is aggregate capital. In the home country, 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.[36] As a consequence, wage rates in the two sectors may differ. Capital is free to flow across sectors.[37] 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

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

We normalize the price of good 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 output of good 2. When the government sets the tariff rate to be τ_t on imported goods, the price of good 1 in the home country becomes $p_{1t} = 1/(1+\tau_t)$. A higher tariff lowers the relative price of good 1 in country 1 since it prevents from fully exploiting its comparative advantage.

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

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

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

[37] This assumption is also not essential.

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

With a reduction of the tariff rate (both in the home country and in the foreign country), the exporting sector expands, which results in 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 the 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 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

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

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 policymaker 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 $\chi(\tau_t)$ depends only on the tariff and the formula is specified in online 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 the result that, given an initial capital stock k_0, the payoff to the policymaker from a sequence of tariff rates $\{\tau_t\}$ is

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

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 investment 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 policymaker 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 policymaker fails to internalize the benefit of lowering tariffs and inducing more capital accumulation in the future and leans toward protectionism. In contrast, in the organizational equilibrium, the action payoffs of policymakers in different periods are equalized. The forward-looking behavior allows the policymaker 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 has been a decline in tariffs over the course of globalization. Figure 5 displays the world average tariff from 1950 to 2017, which gradually decreases from around 16 percent to 4 percent. 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 increases 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 created strategic incentives for gradual liberalization.

Our theory provides an alternative rationale for the observed gradual decline. It is based on the intertemporal trade-off between the short-run welfare of workers, which calls for protectionism, and the long-run welfare of workers, which 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 seven years and choose $\alpha = 0.33$, $\beta = 0.7$, and $\delta = 0.65$. The elasticity of substitution is set to be 1.1, which lies in the middle of various empirical estimates.[38] Finally, we set

[38] See Boehm, Levchenko, and Pandalai-Nayar (2023) for a more detailed discussion.

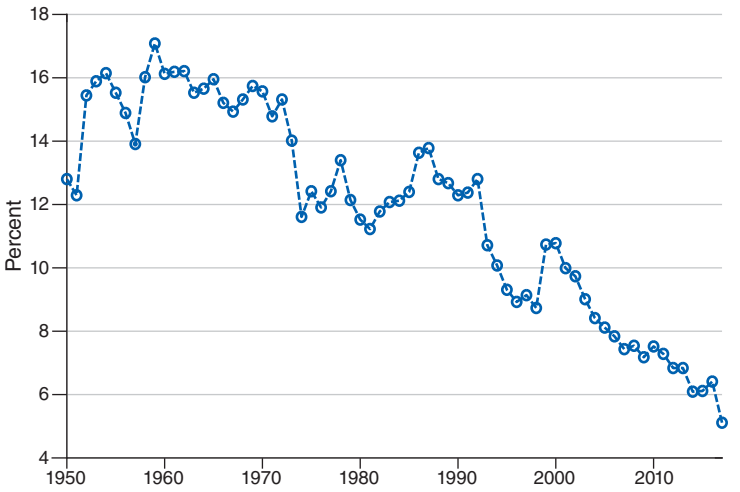

FIGURE 5. WORLD AVERAGE TARIFF

Note: This tariff series is the unweighted world average tariff from Figure 8 in Antràs (2020).

$A = 3.5$ 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 tariffs. The Markov equilibrium features a constant high tariff rate as policymakers behave in the most myopic way. In the Ramsey outcome, the policymaker sets a high level of tariff initially but quickly eliminates it afterward. With organizational equilibrium, the tariff starts from a level in the middle of the previous two cases and gradually converges to a low rate; this pattern is broadly consistent with the one 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 Markov equilibrium workers initially enjoy more consumption than they do in the organizational equilibrium, thanks to the higher 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 analyzing 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 $\tilde{\tau}$ instead of 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.

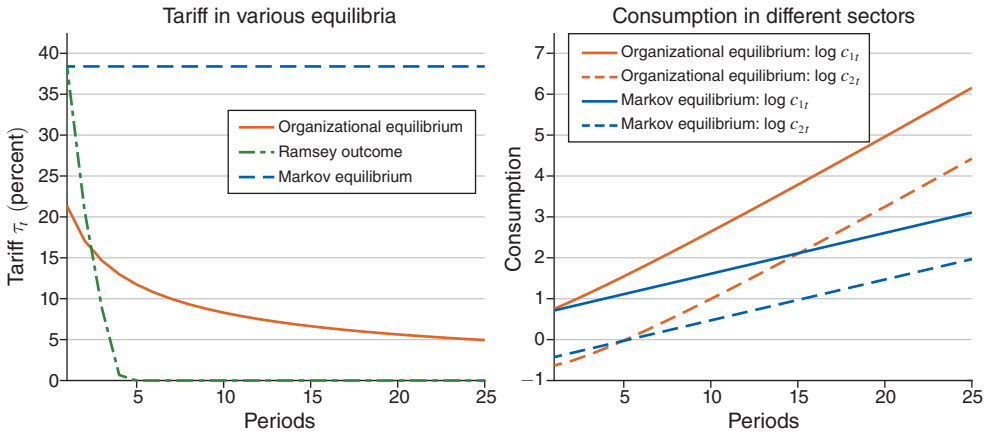

FIGURE 6. TRANSITION PATHS OF TARIFFS AND CONSUMPTION

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 the importing sector (c_{2t}) under the organizational equilibrium and the Markov equilibrium.

V. 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 1970s, 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 United States 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. 1986. "Optimal Cartel Equilibria with Imperfect Monitoring." *Journal of Economic Theory* 39 (1): 251–69.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.

Acemoglu, Daron, and James A. Robinson. 2000. "Why Did the West Extend the Franchise? Democracy, Inequality and Growth in Historical Perspective." *Quarterly Journal of Economics* 115 (4): 1167–99.

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics* 34 (1): S141–98.

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin. 2012. "Dynamics and Stability of Constitutions, Coalitions, and Clubs." *American Economic Review* 102 (4): 1446–76.

Acemoglu, Daron, Georgy Egorov, and Konstantin Sonin. 2015. "Political Economy in a Changing World." *Journal of Political Economy* 123 (5): 1038–86.

Ales, Laurence, and Christopher Sleet. 2014. "Revision Proofness." *Journal of Economic Theory* 152 (C): 324–55.

Antràs, Pol. 2020. "De-Globalisation? Global Value Chains in the post-COVID-19 Age." NBER Working Paper 28115.

Asheim, Geir B. 1997. "Individual and Collective Time-Consistency." *Review of Economic Studies* 64 (3): 427–43.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–68.

Azzimonti, Marina. 2011. "Barriers to Investment in Polarized Societies." *American Economic Review* 101 (5): 2182–2204.

Bassetto, Marco. 2005. "Equilibrium and Government Commitment." *Journal of Economic Theory* 124 (1): 79–105.

Bassetto, Marco. 2008. "Political Economy of Taxation in an Overlapping-Generations Economy." *Review of Economic Dynamics* 11 (1): 18–43.

Bassetto, Marco, Zhen Huo, and José-Víctor Ríos-Rull. 2024. *Data and Code for: "Institution Building without Commitment."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E201421V1.

Bassetto, Marco, and Thomas J. Sargent. 2006. "Politics and Efficiency of Separating Capital and Ordinary Government Budgets." *Quarterly Journal of Economics* 121 (4): 1167–1210.

Bernheim, B. Douglas, Debraj Ray, and Şevin Yeltekin. 2015. "Poverty and Self-Control." *Econometrica* 83 (5): 1877–1911.

Boehm, Christoph E., Andrei A. Levchenko, and Nitya Pandalai-Nayar. 2023. "The Long and Short (Run) of Trade Elasticities." *American Economic Review* 113 (4): 861–905.

Bond, Eric W., and Jee-Hyeong Park. 2002. "Gradualism in Trade Agreements with Asymmetric Countries." *Review of Economic Studies* 69 (2): 379–406.

Brendon, Charles, and Martin Ellison. 2018. "Time-Consistently Undominated Policies." Unpublished.

Caliendo, Lorenzo, Maximiliano Dvorkin, and Fernando Parro. 2019. "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock." *Econometrica* 87 (3): 741–835.

Cao, Dan, and Iván Werning. 2018. "Saving and Dissaving with Hyperbolic Discounting." *Econometrica* 86 (3): 805–57.

Chari, V. V., and Patrick J. Kehoe. 1990. "Sustainable Plans." *Journal of Political Economy* 98 (4): 784–802.

Chatterjee, Satyajit, and Burcu Eyigungor. 2016. "Continuous Markov Equilibria with Quasi-Geometric Discounting." *Journal of Economic Theory* 163 (C): 467–94.

Chisik, Richard. 2003. "Gradualism in Free Trade Agreements: A Theoretical Justification." *Journal of International Economics* 59 (2): 367–97.

Cohen, Daniel, and Philippe Michel. 1988. "How Should Control Theory Be Used to Calculate Time-Consistent Government Policy?" *Review of Economic Studies* 55 (2): 263–74.

Currie, David A., and Paul L. Levine. 1993. *Rules, Reputation and Macroeconomic Contracts.* Cambridge: Cambridge University Press.

Farrell, Joseph, and Eric Maskin. 1989. "Renegotiation in Repeated Games." *Games and Economic Behavior* 1 (4): 327–60.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski. 2014. "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica* 82 (1): 41–88.

Halac, Marina, and Pierre Yared. 2014. "Fiscal Rules and Discretion under Persistent Shocks." *Econometrica* 82 (5): 1557–1614.

Halac, Marina, and Pierre Yared. 2017. "Fiscal Rules and Discretion under Self-Enforcement." NBER Working Paper 23919.

Hall, George J., and Thomas J. Sargent. 2014. "Fiscal Discrimination in Three Wars." *Journal of Monetary Economics* 61: 148–66.

Hall, George J., and Thomas J. Sargent. 2015. "A History of U.S. Debt Limits." NBER Working Paper 21799.

Hall, George J., and Thomas J. Sargent. 2018. "Brief History of US Debt Limits Before 1939." *Proceedings of the National Academy of Sciences* 115 (12): 2942–45.

Jackson, Matthew O., and Leeat Yariv. 2014. "Present Bias and Collective Dynamic Choice in the Lab." *American Economic Review* 104 (12): 4184–4204.

Jackson, Matthew O., and Leeat Yariv. 2015. "Collective Dynamic Choice: The Necessity of Time Inconsistency." *American Economic Journal: Microeconomics* 7 (4): 150–78.

Klein, Paul, Per Krusell, and José-Víctor Ríos-Rull. 2008. "Time-Consistent Public Policy." *Review of Economic Studies* 75 (3): 789–808.

Klein, Paul, Vincenzo Quadrini, and José-Víctor Ríos-Rull. 2005. "Optimal Time–Consistent Taxation with International Mobility of Capital." *B.E. Journal of Macroeconomics: Advances in Macroeconomics* 5 (1): 1–34.

Klein, Paul, and José-Víctor Ríos-Rull. 2003. "Time-Consistent Optimal Fiscal Policy." *International Economic Review* 44 (4): 1217–45.

Kocherlakota, Narayana R. 1996. "Reconsideration-Proofness: A Refinement for Infinite Horizon Time Inconsistency." *Games and Economic Behavior* 15 (1): 33–54.

Krusell, Per, Burhanettin Kuruşçu, and Anthony A. Smith. 2010. "Temptation and Taxation." *Econometrica* 78 (6): 2063–94.

Krusell, Per, and José-Víctor Ríos-Rull. 1996. "Vested Interests in a Positive Theory of Stagnation and Growth." *Review of Economic Studies* 63 (2): 301–31.

Krusell, Per, and Anthony A. Smith. 2003. "Consumption-Savings Decisions with Quasi-Geometric Discounting." *Econometrica* 71 (1): 365–75.

Kubler, Felix. 2007. "Approximate Generalizations and Computational Experiments." *Econometrica* 75 (4): 967–92.

Kubler, Felix, and Karl Schmedders. 2005. "Approximate versus Exact Equilibria in Dynamic Economies." *Econometrica* 73 (4): 1205–35.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.

Maggi, Giovanni, and Andrés Rodríguez-Clare. 2007. "A Political-Economy Theory of Trade Agreements." *American Economic Review* 97 (4): 1374–1406.

Martin, Fernando M. 2011. "On the Joint Determination of Fiscal and Monetary Policy." *Journal of Monetary Economics* 58 (2): 132–45.

Matsuyama, Kiminori. 1990. "Perfect Equilibria in a Trade Liberalization Game." *American Economic Review* 80 (3): 480–92.

Nozawa, Wataru. 2018. "On Nonexistence of Reconsideration-proof Equilibrium with State Variables." *Theory and Decision* 85 (2): 253–73.

Ortigueira, Salvador. 2006. "Markov-Perfect Optimal Taxation." *Review of Economic Dynamics* 9 (1): 153–78.

Phelan, Christopher, and Ennio Stacchetti. 2001. "Sequential Equilibria in a Ramsey Taxes Model." *Econometrica* 69 (6): 1491–1518.

Phelps, Edmund S., and Robert A. Pollak. 1968. "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies* 35: 185–99.

Pierce, Justin R., and Peter K. Schott. 2016. "The Surprisingly Swift Decline of US Manufacturing Employment." *American Economic Review* 106 (7): 1632–62.

Piguillem, Facundo, and Alessandro Riboni. 2015. "Spending-Biased Legislators: Discipline Through Disagreement." *Quarterly Journal of Economics* 130 (2): 901–49.

Piguillem, Facundo, and Alessandro Riboni. 2020. "Fiscal Rules as Bargaining Chips." *Review of Economic Studies* 88 (5): 2439–78.

Prescott, Edward C., and Jose-Victor Rios-Rull. 2005. "On Equilibrium for Overlapping Generations Organizations." *International Economic Review* 46 (4): 1065–80.

Sargent, Thomas J. 2017. "Honoring Public Debts." Unpublished.
Stokey, Nancy L. 1991. "Credible Public Policy." *Journal of Economic Dynamics and Control* 15 (4): 627–56.
Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23 (3): 165–80.
World Bank. 1990–2021. "World Development Indicators, GDP deflator." World Bank. https://data.worldbank.org/indicator/NY.GDP.PCAP.CD (accessed December 22, 2023).
World Bank. 2022. "World Development Indicators, GDP per capita (current US dollar)." World Bank. https://data.worldbank.org/indicator/NY.GDP.DEFL.KD.ZG (accessed December 22, 2023).
World Bank. 2023. "Carbon Pricing Dashboard." World Bank. https://carbonpricingdashboard.worldbank.org/ (accessed October 7, 2023).
Zissimos, Ben. 2007. "The GATT and gradualism." *Journal of International Economics* 71 (2): 410–33.