

CREDIBLE REFORMS: A ROBUST IMPLEMENTATION APPROACH

JUAN PABLO XANDRI*
(JOB MARKET PAPER)

ABSTRACT. I study the problem of a government with low credibility, who decides to make a reform to remove ex-post time inconsistent incentives due to lack of commitment. The government has to take a policy action, but has the ability to commit to limiting its discretionary power. If the public believed the reform solved this time inconsistency problem, the policy maker could achieve complete discretion. However, if the public does not believe the reform to be successful, some discretion must be sacrificed in order to induce public trust. With repeated interactions, the policy maker can build reputation about her reformed incentives. However, equilibrium reputation dynamics are extremely sensitive to assumptions about the public's beliefs, particularly after unexpected events. To overcome this limitation, I study the *optimal robust policy* that implements public trust for *all rationalizable beliefs*. I focus on robustness to all extensive-form rationalizable beliefs and provide a characterization. I show that the robust policy exhibits both partial and permanent reputation building along its path, as well as endogenous transitory reputation losses. In addition, I demonstrate that almost surely the policy maker eventually convinces the public she does not face a time consistency problem and she is able to do this with an exponential arrival rate. This implies that as we consider more patient policy makers, the payoff of robust policies converge to the complete information benchmark. I finally explore how further restrictions on beliefs alter optimal policy and accelerate reputation building.

JEL CLASSIFICATION CODES: E58, D81, D82, D83

1. Introduction

In the mind of policy makers, a reputation for credibility is a delicate and hard-won situation. Policy shifts are discussed with great care and concerns regarding how the public will react. By contrast, formal models of reputation employing insights from repeated games typically assume a perfect degree of certainty and coordination. The purpose of this paper is to build on this literature to model reputation in a way that reflects the uncertainty faced by policy makers.

Date: November 2012.

I am grateful to Daron Acemoglu, George-Marios Angeletos, Abhijit Banerjee, Robert Townsend, Ivan Werning, and Muhamet Yildiz who provided support, guidance, and extremely helpful suggestions. I am also thankful for the essential feedback provided by Gabriel Carroll, Arun Chandrasekhar, Sebastian Di Tella, Juan Dubra, Glenn Ellison, Alvaro Forteza, Ben Golub, Horacio Larreguy, Anton Kolotilin, Pablo Kurlat, Plamen Nenov, Juan Passadore, Juuso Toikka, Xiao Yu Wang, Alex Wolitzky, Luis Zermeno and Jan Zilinsky

[†]Massachusetts Institute of Technology. Email: jxandri@mit.edu.

Following the seminal papers of Kydland and Prescott [1977] and Barro and Gordon [1983], macroeconomic theory has extensively studied the so-called time inconsistency problem of government policy. In essence, all time inconsistency problems consist of an authority who needs some agents in the economy (e.g., consumers, financial sector) to trust her with a decision that will be taken on their behalf. In the canonical example of monetary policy, a policy maker wants the public to trust her announcements of inflation policy. The fundamental problem resides in the fact that even if the decision maker is allowed to announce what policy she plans to take, the agent's decision to trust decision ultimately depends on their *perception* or beliefs about the (ex-post) incentives the policy maker will face *after* they trust her. This creates a wedge between the ideal policies the authority would want to implement and the ones that she can credibly promise. In the context of our inflation example, since agents know that the government will have ex-post incentives to boost employment by increasing inflation and reducing real wages, this will result in an inefficiently high equilibrium inflation.

The literature has dealt with this problem in two ways. First, some authors have argued that the government should *forfeit some flexibility* through a formal arrangement (e.g., inflation targeting/caps and tax restrictions). Second, others have argued that the government should *modify its incentives*, either by having reputational concerns through repeated interactions (Stokey [1989, 1991] and Chari and Kehoe [1990, 1993a,b]) or by delegating the decision to an agency with different incentives that will limit its time inconsistency bias. Examples of such delegation include appointing a conservative central banker (Rogoff [1985]) or making the monetary authorities subject to a formal or informal incentive contract (Lohmann [1992], Persson and Tabellini [1993], Walsh [1995]). If full commitment to contingent policies is not available and flexibility is socially desirable, these “*incentive reforms*” may become a desirable solution. The key difference between both approaches is that policies that are enforced by incentive reforms are very sensitive to the assumption that the public knows exactly what the reformed incentives are. If the public believed that with a high enough probability, the government still has a time inconsistency problem, then the situation would remain unsolved. I will model this uncertainty as the public having incomplete information about the policy makers incentives, as in Barro [1986], Phelan [2006], but also allowing the public to have uncertainty about the governments expectations for the continuation game. The main goal of our paper will be to investigate if, through repeated interactions, the government can convince the public about its reformed incentives.

Using equilibrium analysis to answer this question typically relies on rather strong common knowledge assumptions as to how agents play, their priors on the government's type, as well as how all parties revise their beliefs. In the particular case of repeated games, predictions of a particular equilibrium may be extremely sensitive to assumptions as to how agents update their expectations about the continuation game *on all potential histories* that might be observed. These are complicated, high dimensional objects, of which the policy maker

may have little information about. This may be due to the difficulty in eliciting both the higher order beliefs from the public as well as contingent beliefs on nodes that may never be reached.

The approach I use is conceptually related to the robust mechanism design literature (Bergemann and Morris [2005, 2009]). The policy is required to implement trust along its path for all feasible agent beliefs within a large class. The class of beliefs that are deemed feasible is crucial to our exercise, since a larger set makes the analysis more robust, a smaller set makes it trivial; the feasible set I consider is discussed below. I study the case where the only constraints on the sets of beliefs are that they are consistent with common certainty of rationality: every agent knows that the other agents are rational, they know everyone knows this, and so on. This is the rationale behind rationalizability, which consists on an iterative deletion of dominated strategies. The present analysis requires beliefs to be such that agents not to question the government's rationality, unless proven otherwise, which is given by the solution concept of *strong* or *extensive-form rationalizability* of Pearce [1984], Battigalli and Siniscalchi [2002].

I show that this policy exhibits endogenous transitory gains and losses of reputation. Moreover, the policy achieves permanent separation (i.e. public is convinced about the success of the reform from then on) almost surely and it does so with an exponential arrival rate. As the discount factor of the policy maker increases, the expected payoff of this robust policy approximates the full commitment first-best benchmark. This policy will also be the *max-min* strategy for the policy maker regardless of their particular beliefs and hence provides a lower bound both for payoffs as well as the speed of separation of any strategy that is consistent with extensive form rationalizability.

To understand the intuition behind the results, suppose the public hypothesizes that he faces a time inconsistent policy maker and has observed her taking an action that did not maximize her spot utility. To fix ideas, suppose she took an action such that, if she was the time inconsistent type, gave her 10 utils. Meanwhile, she could have reacted in a manner that gave her 25 in spot utility instead. The implied opportunity cost paid by her would only be consistent with her being rational if she expected a net present value of at least 15 utils, and therefore the opportunity cost paid would have been a profitable investment. This further implies that the government beliefs and planned course of action from tomorrow onward must deliver (from the government's point of view) more than 15 utils, which is a constraint that rationality imposes on the government's expected future behavior. If however, the maximum feasible net present value attainable by a time inconsistent government was 10 utils, the public should then infer that the only possible time inconsistent type that they are facing is an irrational one. However, if such a history was actually consistent with the policy maker being time consistent (e.g. she had an opportunity cost of 2) then the public should be fully convinced from then on that they are facing the reformed government. Therefore, the implied spot opportunity cost paid by the time inconsistent type will be a measure of

reputation that places restrictions on what the public believes the policy maker will do in the future. I emphasize that this will be independent of their particular beliefs and relies only on an assumption of rationality. I also show that this is in fact the *only* robust restriction that strong common certainty of rationality imposes, making the implied opportunity cost paid the *only* relevant reputation measure in the robust policy. Moreover, I show that the optimal robust policy can be solved as a dynamic contracting problem with a *single promise keeping constraint*, analogous to [Thomas and Worrall \[1988\]](#), [Kocherlakota \[1996\]](#) and [Alvarez and Jermann \[2000\]](#) in the context of optimal risk sharing with limited commitment, which makes the analysis of the optimal robust policy quite tractable.

The rest of the paper is organized as follows. Section 2 describes two macroeconomic applications of time inconsistency, monetary policy and capital taxation, which are informed by our theoretical results. Section 3 provides a brief literature review. Section 4 introduces the stage binary action repeated game, and introduces the concepts of weak and strong rationalizability. Section 5 defines the concept of robust implementation and studies robust implementation for all weak and strong rationalizable outcomes. In section 6 I study the basic properties of the optimal robust strategy and the reputation formation process as well as the limiting behavior as policy makers become more patient. I also study how further restrictions on the set of feasible beliefs can help accelerate the reputation formation process and in particular find a characterization of restrictions that generate monotone robust policies (i.e. policies that exhibit only permanent gains of reputation on its path). In Section 7 I study some extensions to our model and discuss avenues for future research. Finally, Section 8 concludes.

2. Examples

I start with some time-inconsistency examples from the literature and use them to motivate my model and analysis. I focus on two of the most commonly studied questions in the macroeconomic literature: *capital taxation* and *monetary policy*. I will show that even if the policy maker undertakes a reform that solves for her time inconsistent bias, when agents have imperfect knowledge about the government objectives, a time inconsistency problem of government policy arises.

2.1. Capital Taxation. I use a modified version of [Phelan \[2006\]](#) and [Lu \[2012\]](#), where the time inconsistent type is a benevolent government, instead of just opportunistic. Consider an economy with two type of households: workers (w) and capitalists (k). There is a continuum of measure one of identical households, for each type. Capitalist households have an investment possibility and can invest $q \in [0, \bar{q}]$ units in a productive technology with a constant marginal benefit of 1 and a constant marginal cost of I . Workers do not have access to this technology and can only consume their own, fixed endowment of $e > 0$.

There is also a public good that can be produced by a government that has a marginal value of z_k to capitalist households and z_w to workers, where $z = (z_k, z_w)$ is a joint random

variable. The government taxes a portion $\tau(z)$ of capital income after the shock is realized, in order to finance the production of $r(z)$ units of public good. Given the expected policy $\{\tau(z), r(z)\}_{z \in Z}$, workers and capitalists households utilities are given by

$$(2.1) \quad U_w = e + \mathbb{E}_z [r(z) z_w]$$

$$(2.2) \quad U_k = (1 - \tau^e) q - Iq + \mathbb{E}_z [r(z) z_k]$$

where $\tau^e = \mathbb{E}[\tau(z)]$. A leading example is the case where the “public good” is simply redistribution from the capitalists to the workers. In this case $z_w > 0$ and $z_k = 0$.

The optimal investment decision for a capitalist is to invest $q_i = \bar{q}$ if $1 - \tau^e < I$, and 0 otherwise, since they do not internalize their marginal effect on the production of public good. As a benchmark, we will first solve for the policy $\{\tau_k(z), r_k(z)\}_{z \in Z}$ that maximizes only the capitalist households expected utility, subject to the government’s budget constraint:

$$(2.3) \quad \max_{q \in [0, \bar{q}], \tau_k(\cdot), r_k(\cdot)} (1 - \mathbb{E}[\tau_k(z)]) q - Iq + \mathbb{E}[r_k(z) z_k] \text{ s.t. } r_k(z) \leq \tau_k(z) q \text{ for all } z$$

Given q , the optimal policy involves full expropriation ($\tau_k(z) = 1, r_k(z) = q$) when $z_k \geq 1$ and zero taxes otherwise, which induces an expected tax rate of $\tau^e = \Pr(z_k \geq 1)$. If

$$(2.4) \quad I < \Pr(z_k \leq 1)$$

then households expecting policy $\{\tau_k(z), r_k(z)\}_{z \in Z}$ will choose $q = \bar{q}$. However, this will not be the policy chosen by a benevolent government that also values workers. After the households investment decision, and the state of nature has been realized, the government chooses public good production \tilde{r} and tax rate $\tilde{\tau}$ to solve:

$$(2.5) \quad \max_{\tilde{r}, \tilde{\tau}} \tilde{r} (z_k + \alpha z_w) + (1 - \tilde{\tau}) q \text{ s.t. } \tilde{r} \leq \tilde{\tau} q$$

where $\alpha \geq 0$ is the relative weight that the government puts on workers welfare.

Defining $z_g := z_k + \alpha z_w$, the marginal value of the public good between capitalists and the government will typically be different, unless $\alpha = 0$. Solving 2.5 gives $\tau_g^e = \Pr(z_k + \alpha z_w > 1)$. I will assume that $I > \Pr(z_k + \alpha z_w \leq 1)$, so capitalist households will optimally decide not to invest ($q = 0$) and no public good production will be feasible. Finally, I assume that the parameters of the model are such that a benevolent government would want to commit to the capitalist’s most preferred policy $\{\tau_k(z), r_k(z)\}_{z \in Z}$ if she was given the possibility.¹

To solve the “time inconsistency” problem, I first explore the possibility of introducing a cost to raise taxes. This means that if taxes are increased, the government has to pay a cost of $c > 0$. The government would then optimally choose taxes $\tau = 0$ and increase them only when needed. She solves

$$\max_{\tilde{r}, \tilde{\tau}} \tilde{r} (z_k + \alpha z_w) - c \text{ s.t. } \tilde{r} \leq \tilde{\tau} q.$$

¹This happens if $\Pr(z_k > 1) \mathbb{E}(z_k + \alpha z_w \mid z_k > 1) + \Pr(z_k \leq 1) > 0$.

In this case, the expected tax rate is now $\tau^e(c) = \Pr(z_k + \alpha z_w > \frac{c}{q})$. By setting $c = \bar{c}$ to solve $1 - \tau^e(c) = I$, the time inconsistent government can now induce households to invest, by credibly distorting its tax policy.

Another way to deal with the problem is to make an institutional reform and delegate the public good provision to a different policy maker, who has incentives aligned with the capitalist households. The new policy maker type now solves

$$\max_{\tilde{r}, \tilde{\tau}} \tilde{r}(z_k + \alpha_{new} z_w) + (1 - \tilde{\tau})q \text{ s.t. } \tilde{r} \leq \tilde{\tau}q.$$

By introducing a “pro-capitalist government” with $\alpha_{new} = 0$, the capitalists most desirable policy $\{\tau_k(z), r_k(z)\}_{z \in Z}$ would be credibly implemented without the need of setting a cost to increase capital taxes. Under some parametric assumptions, it will be socially desirable for the benevolent government (without taken into account the commitment cost payed) to delegate policy making to the “pro-capitalist” type that does not need to impose tax increase costs to convince households to invest².

However, if households were not convinced that they are indeed facing a reformed, pro-capitalist government, they will need some assurance (i.e. some restrictions to ex post increase taxes) in order to *trust* that the government will not expropriate their investments too often. In Phelan [2006], the analog to a pro-capitalist type is a commitment type (as in Fudenberg and Levine [1989]) that always pick the same tax rate. In Lu [2012], the government can make announcements, and can be either a committed type (i.e. one that is bound by the announcement) or a purely opportunistic type that may choose to deviate from the promised policy, which is analog to the benevolent type in our setting.

Formally, Let $\pi \in (0, 1)$ be the probability that capitalist households assign to the new government to actually be a pro-capitalist type. Then, if there is complete flexibility to increase taxes, the expected tax rate would be

$$(2.6) \quad \tau^e(\pi) = \pi \Pr(z_k > 1) + (1 - \pi) \Pr(z_k + \alpha z_w > 1)$$

Condition 2.6 implies that for sufficiently low π , we would have $1 - \tau^e(\pi) < I$ and capitalists will decide not to invest. Thus, as long as capitalists perceive that the new government might still be time inconsistent (modeled by a low π), it will be necessary to set some cost to raise taxes in order to induce capitalists to invest, even though the government is now a pro-capitalist type.

2.2. Monetary Policy. I use the framework in Obstfeld and Rogoff [1996]³. I assume that total output (in logs) y_t depends negatively on the real wage and some supply side shock z_t , according to

$$(2.7) \quad y_t = \bar{y} - [w_t - p_t(z_t)] - z_t$$

²This happens if $\Pr(z > 1) \mathbb{E}(z_k + \alpha z_w | z_k > 1) + \Pr(z_k \leq 1) > \Pr(z_k + \alpha z_w > \frac{c}{q}) \mathbb{E}(z_k + \alpha z_w | z_k + \alpha z_w > \frac{c}{q}) + \Pr(z_k + \alpha z_w \leq 1)$

³Section 9.5, pp 634-657.

where \bar{y} is the flexible price equilibrium level, z_t is a supply shock with $\mathbb{E}(z_t) = 0$ and $p_t(z_t)$ is the nominal price level at time t set by the monetary authority. In equilibrium, nominal wages are set according to $w_t = \mathbb{E}_{t-1}[p_t(z_t)]$, to match expected output to its natural level \bar{y} . A benevolent monetary authority observes the shock z_t and decides the inflation level in order to minimize deviations of output with respect to a social optimal output y^* and deviations of inflation from a zero inflation target:

$$(2.8) \quad \mathcal{L} = \frac{1}{2} (y_t - y_t^*)^2 + \frac{\chi}{2} \pi_t^2.$$

I assume that $k := y_t^* - \bar{y} > 0$. This measures the wedge between the output level targeted by authorities and the natural level of output, which are different due to market inefficiencies, even under flexible prices.⁴ Defining inflation as $\pi_t(z_t) := p_t(z_t) - p_{t-1}$, and using equation 2.7, together with the wage setting rule, the loss function simplifies to

$$(2.9) \quad \mathcal{L}(\pi, \pi^e, z) = \frac{1}{2} [\pi - \pi^e - z - k]^2 + \frac{\chi}{2} \pi^2,$$

where $\pi^e = \mathbb{E}[\pi(z)]$ are the expectations formed by the private sector about inflation (which should be correct under rational expectations). The *full commitment benchmark*, in which the monetary authority can commit, ex-ante, to a state contingent inflation policy $\pi(z)$, to solve

$$\min_{\pi(\cdot), \pi^e} \mathcal{L}(\pi, \pi^e, z) \quad \text{s.t.: } \pi^e = \mathbb{E}[\pi(z)]$$

with solution

$$(2.10) \quad \pi_c(z) = \frac{z}{1 + \chi} \quad \text{and} \quad \pi^e = \mathbb{E}_z[\pi_c(z)] = 0.$$

In contrast, when the monetary authority cannot commit to a state contingent policy, conditional on π^e and z , she chooses π to solve:

$$(2.11) \quad \min_{\pi \in \mathbb{R}} \mathcal{L}(\pi, \pi^e, z) \iff \pi_{nc}(z) = \frac{\pi^e + z + k}{1 + \chi}$$

By taking expectations on both sides of 2.11 we get $\pi^e = \frac{k}{\chi}$, which I will refer to the *time inconsistency bias*. Equilibrium inflation is then

$$(2.12) \quad \pi_{nc}(z) = \frac{k}{\chi} + \pi_c(z)$$

Output $y(z)$ is identical in both cases, for all shocks. However, $\mathbb{E}[\pi_{nc}^2(z)] = \mathbb{E}[\pi_c^2(z)] + \frac{k^2}{\chi^2}$ so the outcome with no commitment is strictly worse than the full commitment benchmark.

How can the monetary authority solve this problem? A first approach is to formally limit the flexibility of monetary policy by restricting the set of inflation levels the monetary

⁴See Rogoff [1985] and Obstfeld and Rogoff [1996] for a discussion of such potential inefficiencies.

authority can choose from. [Athey et al. \[2005\]](#) show that this can be optimally done by choosing an *inflation cap* $\bar{\pi}$ ⁵, such that $\pi(z) \leq \bar{\pi}$ for all z . Inflation policy is now

$$\pi(z | \bar{\pi}) = \min \left\{ \frac{\pi^e(\bar{\pi}) + z + k}{1 + \chi}, \bar{\pi} \right\}$$

where $\pi^e(\bar{\pi})$ solves the fixed point equation $\pi^e(\bar{\pi}) = \mathbb{E}_z \left[\min \left\{ \frac{\pi^e(\bar{\pi}) + z + k}{1 + \chi}, \bar{\pi} \right\} \right]$.

An alternative approach, first suggested by [Rogoff \[1985\]](#), is to introduce *institutional reforms to the monetary authority*, with the purpose of alleviating the time inconsistency bias by inducing changes in their preferences. Imagine first that the government can delegate the monetary policy to a policy maker type $\theta = new$, that wants to minimize a modified loss function

$$(2.13) \quad \mathcal{L}_{new}(\pi, \pi^e, z) = \frac{1}{2} [\pi - \pi^e - z - k^{new}]^2 + \frac{\chi^{new}}{2} \pi^2$$

[Rogoff \[1985\]](#) suggests placing a “conservative central banker”, that has $k^{new} = k$ but $\chi^{new} > \chi$, so that it places a greater importance on inflation stabilization than society does. From 2.12 we see that increasing the weight χ makes the effective inflation bias smaller, and hence may alleviate the time inconsistency problem at the expense of a milder reaction to supply shocks, as evidenced by equation 2.11.

By setting $k^{new} = 0$ the optimal policy with no commitment for $\theta = new$ would implement the full commitment solution. This would correspond to having a monetary authority that believes there are no market inefficiencies and wants to stabilize output around its flexible price equilibrium level \bar{y} . The same outcome can be implemented if, instead of changing the preference parameters, we add a linear term to the loss function:

$$\mathcal{L}_{new} = \mathcal{L} + \alpha [\eta\pi]$$

[Walsh \[1995\]](#) and [Persson and Tabellini \[1993\]](#) argue that this can be done by offering a contract to the central bank governor. This can be either a formal monetary contract⁶ or an informal relational contract under which realized levels of inflation affect the continuation values for the monetary authority (e.g. the governor could be fired if inflation reaches sufficiently high levels, as in [Lohmann \[1992\]](#)). Here $\alpha > 0$ represents the relative weight of

⁵In their paper, [Athey et al. \[2005\]](#) solve for the optimal dynamic mechanism for a time inconsistent policy maker, that has private information about the state of the economy, which is i.i.d across periods, and show that any optimal mechanism exhibits a constant inflation cap in all periods. In a static setting, shocks can be thought as private information for the monetary authority, so an inflation cap would also be a characteristic of the more general mechanism design problem:

$$\begin{aligned} & \max_{\pi(\cdot), \pi^e} \mathbb{E}_z [\mathcal{L}(\pi(z), \pi^e, z)] \\ \text{s.t.} & \quad \mathcal{L}[\pi(z), \pi^e, z] \geq \mathcal{L}[\pi(z'), \pi^e, z] \text{ for all } z, z' \in Z \end{aligned}$$

⁶Suppose the monetary authority minimizes $\hat{\mathcal{L}}_{old} = \mathcal{L}_{old} - \alpha u[\phi(\pi)]$ where $\phi(\pi)$ is a monetary reward function depending on realized inflation, and $\alpha > 0$ of monetary incentives relative to the monetary authorities “benevolent” incentives. See that by picking $\phi(\pi) = u^{-1}(-\eta\pi)$, a decreasing function of inflation, the contract will induce the linear component in 2.13, which coincides with the optimal contract in this setting.

his self-interest payoffs relative to the social welfare. By picking $\eta = \frac{k}{\alpha}$ the full commitment inflation policy would be implemented.

While the institutional reform route may seem desirable, *these institutional reforms may not be perfectly observed by the private sector*. The public might not be convinced that the monetary authority is now more conservative or have a smaller time inconsistency bias than the previous one. Such a problem is likely to be particularly acute because these are changes in preferences, which involve either delegation or perhaps informal relational contracts that are imperfectly observed. If this was the case, restrictions such as inflation caps might still be necessary. For example, take the institutional reform with $k^{new} = 0$, and no inflation caps are set. If the public assigns probability $\mu \in (0, 1)$ to the incentive reform being successful, expected inflation would then be

$$\pi^e = (1 - \mu) \frac{k}{\chi} > 0$$

Thus, as long as the public perceives there might still be a time inconsistency bias, institutional reforms might not be enough, and inflation caps might be necessary to implement smaller inflation expectations. The literature has studied the case of

3. Literature Review

The literature on time inconsistency of government policy is extensive, beginning with the seminal papers by Kydland and Prescott [1977] and Barro and Gordon [1983], where the idea of the commitment solution (i.e. choosing policy first) was first introduced. The reputation channel was first explored by Backus and Driffil [1984] and Barro [1986], who studied policy games in which a rational (albeit time inconsistent) government living for finitely many periods may find it optimal to imitate a “commitment type”. This commitment type is an irrational type that plays a constant strategy at all histories. They show (following the arguments in Kreps et al. [1982], Kreps and Wilson [1982], Milgrom and Roberts [1982]) that for long enough horizon, the unique sequential equilibrium of the game would involve the government imitating the commitment type for the first periods, and then playing mixed strategies, which imply a gradual reputation gain if she keeps imitating.

In an infinite horizon setting, Fudenberg and Levine [1989] show that a long lived agent facing a sequence short lived agents can create a reputation for playing as the commitment type. By consistently playing the commitment strategy, the long lived agent can eventually convince the short lived agents that she will play as a committed type for the rest of the game. Celentani and Pesendorfer [1996] generalized this idea to the case of a government playing against a continuum of long-lived small players, whose preferences depend only on aggregate state variables. The atomistic nature of the small players allows them to use Fudenberg and Levine [1989] results to get bounds on equilibrium payoffs. Phelan [2006] studies the problem of optimal linear capital taxation, in a model with impermanent types, which can accommodate occasional losses of reputation. Rather than obtaining bounds, he

characterizes the optimal Markovian equilibrium of the game, as a function of the posterior the public has about the government's type.

A second strand of the literature on reputation focuses on a complete information benchmark with the goal of characterizing sustainable policies. These are policies that are the outcome of a subgame perfect equilibrium of the policy game, starting with [Stokey \[1989, 1991\]](#) and [Chari and Kehoe \[1990, 1993a,b\]](#). In such environments, governments may have incentives to behave well under the threat of punishment by switching to a bad equilibrium afterwards. See [Sargent and Ljungqvist \[2004\]](#)⁷ for a tractable unified framework to study these issues.⁸

This paper studies reputation formation on both dimensions: in terms of payoff heterogeneity and in terms of equilibrium punishments. The main point of departure is that instead of designing the optimal policy for a time inconsistent policy maker (that wants to behave as if she was time consistent), I focus on the opposite case. I consider the problem of a trustworthy policy maker (with no time inconsistency bias) who nevertheless may be perceived as opportunistic by the agents. Therefore, its goal is essentially to separate itself from the time inconsistent, untrustworthy type, if possible. The most related papers in spirit to mine are [Debortoli and Nunes \[2010\]](#), [King et al. \[2012\]](#) and particularly [Lu \[2012\]](#). [Debortoli and Nunes \[2010\]](#) study the optimal policy problem of a benevolent government that has access to a "loose commitment" technology, under which not all announcements can be guaranteed to be fulfilled. [Lu \[2012\]](#) explores the optimal policy of a committed government that worries she might be perceived as a government that cannot credibly commit to her announced policies. This paper also focuses on characterizing the optimal policy for the time consistent type (the committed type in her setting) instead of just studying the optimal policy of a time inconsistent type imitating a consistent one. [King et al. \[2012\]](#) apply these ideas in the context of the standard New Keynesian model, similar to our setup in subsection 2.2. Ultimately these papers study equilibrium in an environment where all players involved know that the government is ex-ante either a type that can commit or cannot commit (which holds for all subsequent periods). They then study a particular equilibrium refinement that happens to select the best equilibrium for the able-to-commit type. They also show that other equilibrium refinements such as the intuitive criterion (e.g., [Cho and Kreps \[1987\]](#)) select a different equilibrium. In macroeconomics, the most related paper to mine that studies robustness to specific refinements is [Pavan and Angeletos \[2012\]](#). They study the robust predictions of any equilibria in a global game setting with incomplete information.

⁷Chapter 16, pp 485-526

⁸This principle is also exploited in the relational contract literature ([Levin \[2003\]](#), [Baker and Murphy \[2002\]](#), [Bull \[1987\]](#)) where a principal announces a payment scheme after income is realized (the state-contingent policy) but has no commitment to it other than the one enforced by the threat of retaliation by the agent (not making effort, strike, quit, etc). Similar themes are studied in the literature on risk sharing with limited commitment ([Thomas and Worrall \[1988\]](#), [Kocherlakota \[1996\]](#), [Ligon \[1998\]](#) and [Ligon et al. \[2000\]](#)) where a transfer scheme conditional on the realization of income (the contingent policy) is enforced by threatening agents who deviate of excluding them from the social contract.

The literature on robust mechanism design is fairly recent, starting with partial robust implementation in Bergemann and Morris [2005], and robust implementation in Bergemann and Morris [2009]. The latter focuses on finding conditions on environments and social choice functions such that they are implemented under implemented for all possible beliefs, if the only thing that the mechanism designer knows about the agent’s beliefs is that they share common knowledge (or certainty) of rationality. When the environment is dynamic, different concepts of rationalizability may be used, like normal form Interim Correlated Rationalizability (as in Weinstein and Yildiz [2012]) and Interim Sequential Rationalizability (Penta [2011, 2012]), among others. Our paper focuses on the stronger assumption of strong certainty of rationality (Battigalli and Siniscalchi [1999, 2002, 2003]) which is also equivalent to Pearce [1984] notion of “Extensive-form rationalizability”. In a similar vein, the paper most related in spirit to mine is Wolitzky [2012]. He studies reputational bargaining in a continuous time setting in which agents announce bargaining postures that they may become committed to with a given positive probability. He characterizes the minimum payoff consistent with mutual knowledge of rationality between players (i.e., one round of knowledge of rationality), and the bargaining posture that she must announce in order to guarantee herself a payoff of at least this lower bound. A crucial difference to my setting is the commitment technology, which ensures certain expected payoffs to the other party, regardless whether they think they are facing a rational opponent or not. I characterize optimal robust policy in a repeated setting in which one can guarantee themselves the best payoff that is consistent with (strong) common knowledge of rationality.

4. The Model

I now introduce the framework and model. Section 4.1 describes the stage game and shows the multiplicity of equilibria. I then setup the repeated game in Section 4.2 and develop the concept of system of beliefs in Section 4.3. Section 4.4 introduces weak and strong rationalizability and Section 4.5 argues why we must turn to robustness relative to equilibrium refinements.

4.1. Stage Game. There are two players: a policy maker d (she) and an agent p (he). The agent represents the public. In the no-commitment benchmark, p is asked to trust a state-contingent decision to d , who after a state of nature $z \in Z$ is realized, has to choose a policy that affects both parties payoffs. For simplicity, I will assume that d has only two options: a “normal” policy that is optimal most of the time and an “emergency” policy that needs to be taken in certain instances. As a mnemonic device, we will write g for the normal policy (pushing a “green button”) and r for the emergency policy (pushing a “red button”). The extensive form game is described in Figure 1.




FIGURE 1. Stage game with time inconsistent type

The random shock $z = (U_p, U_{old})$ is the profile of relative utilities of the emergency action r with respect to g , for both the public p and the decision maker d . I assume this random shock to be an absolutely continuous random variable over $Z := [\underline{U}, \bar{U}]^2 \subset \mathbb{R}^2$, with density function $f(z)$. The subscript “old” serves to remind the reader that these are the preferences of the policy maker before a reform is undertaken, which will be described below.

To make this concrete, consider the capital taxation example described in Section 2.1. The government decides whether to expropriate capital to finance public good provision (the emergency policy with $r = q$) or to keep taxes at zero (the normal policy with $r = 0$). The government gets to make this decision only when capitalist households decide to invest $q_i = \bar{q}$ which corresponds to the public placing trust in the government as without any investment the government cannot finance public good provision in the first place. As such, T corresponds to “invest \bar{q} ”. The shock z can be written as

$$(4.1) \quad z = (U_p, U_{old}) := \bar{q} (z_k - 1, z_b - 1)$$

Returning to the general model, I will additionally assume that

$$(4.2) \quad \mathbb{E}_z \{ \max(0, U_p) \} > \underline{u}_p > \int_{U_{old} > 0} U_p f(z) dz$$

This implies that if p was guaranteed his most desirable policy, he would place trust in d .⁹ I will further assume that

$$(4.3) \quad \underline{u}_p, \underline{u}_{old} < 0$$

so that both parties would benefit from ensuring that the decision maker plays the normal action g for all shocks (and hence losing all flexibility), if such a policy was enforceable. In the capital taxation problem, this would correspond to a ban on positive taxes while in our inflation example this would correspond to a commitment to zero inflation.¹⁰ Although this would induce the public trust, this would come at the cost of losing all flexibility to optimally react to shocks.

As I discussed in the introduction, I will explore two (potentially complementary) ways to solve this *time inconsistency problem*: by credibly losing *some* flexibility to react to the economic shock, or by reforming the incentives of the decision maker, in order to alleviate the time inconsistency problem.

First, I introduce a commitment cost technology, under which the decision maker can choose, before the game starts, a cost $c \geq 0$ of taking the emergency action r . This can be interpreted as a *partial commitment* to the normal policy g , that includes an escape clause to break the commitment, forcing the decision maker to suffer a cost of $c \geq 0$ utils (as in [Lohmann \[1992\]](#)). Although d cannot commit to a complete contingent rule, I assume that the commitment cost chosen is binding. In the capital taxation model, this corresponds to the cost of increasing taxes chosen by the time inconsistent government, while in the inflation setting model this would intuitively translate to the inflation cap $\bar{\pi}$, in that it is a partial commitment chosen by the monetary authority. The modified stage game is illustrated in [Figure 2](#).

⁹In the context of the capital taxation problem, by setting $\underline{u}_p = -(1 - I) \mathbb{E}(z_k | z_b \geq 1) (\bar{q} - 1)$, the left hand side inequality of [4.2](#) corresponds to the solution to [2.3](#) in the capital taxation problem, together with assumption [2.4](#). The right hand side inequality corresponds to condition [??](#). Notice also that $\underline{u}_p < 0$

¹⁰Although this seems to be to an extreme policy to be seen in practice, hyperinflation stabilization programs usually involve drastic measures, that resemble losing all flexibility to stabilize output. For example, Zimbabwe in 2009 decided to abandon its currency (and hence most of its monetary policy) within the context of a severe hyperinflation (which reached a peak of 79,600,000% per month in November of 2008).




FIGURE 2. Stage game with commitment cost choice

I will allow $c = \infty$, the case where d decides to shut down the emergency action r , as I previously discussed, so that the commitment cost set is $C = \mathbb{R}_+ \cup \{\infty\}$. In this game, the decision maker would then choose the commitment cost $c = \bar{c}$ that makes p indifferent between trusting and not trusting:

$$(4.4) \quad \bar{c} := \min \left\{ c \geq 0 : \int_{U_{old} > c} U_p f(z) dz \geq \underline{u}_p \right\} < \infty$$

While this technology is available, I will consider the case in which the policy maker makes a reform, by effectively changing the ex-post incentives that the decision maker faces, by either delegating the decision to a different agent (like the conservative central banker of Rogoff [1985]) or by designing a contract for the decision maker (as in Walsh [1995] and Persson and Tabellini [1993]). I model this reform by creating a new policy maker type, $\theta = new$, with ex-post payoffs given by

$$(4.5) \quad U_{new}(z) = U_p$$

i.e. the reformed decision maker has the same ex-post incentives as the public. Condition 4.2 implies that if p knew he was facing this type of decision maker, he would trust her even with $c = 0$, so that no commitment cost would be necessary. This corresponds to the “pro-capitalist” government in the capital taxation model, which removes the time inconsistency of government policy.

The key problem I will study in this paper is the government’s lack of credibility. Even though an incentive reform has been carried out, the public may remain unconvinced that

the reform has been effective. For instance, investors might still believe that the government is not pro-capitalist enough, and will expropriate them too often, or that the new appointed central banker may not be a conservative type. I model this situation by introducing *payoff uncertainty* from the public side: p believes he is facing a reformed, time consistent decision maker $\theta = new$ with probability $\pi \in (0, 1)$ and otherwise faces the old, time inconsistent type $\theta = old$. This results on an incomplete information game, described in Figure 3.




FIGURE 3. Incomplete information game

This model follows an *executive approach* to optimal policy, where it is the decision maker herself who decides the commitment cost and the policy rule. This contrasts with the *legislative approach* studied in Canzoneri [1985] and Athey et al. [2005] who instead solve for the optimal mechanism design problem from the point of view of p . In Section (7) I briefly explore this route and find that in our setting, it will be detrimental to welfare, conditional on the government being of type $\theta = new$.

Because the commitment cost choice is taken after the type has been realized, this is effectively a *signaling game*. The choice of commitment cost could in principle help p infer d 's type. As a prelude to what follows it is useful to characterize the set of Perfect Bayesian Equilibria (PBE) of this game.

Let $\underline{c}(\pi)$ be the minimum cost that induces p to trust d in any pooling equilibrium:

$$(4.6) \quad \underline{c}(\pi) := \min \left\{ c \geq 0 : \pi \int_{U_p > c} U_p f(z) dz + (1 - \pi) \int_{U_{old} > c} U_p f(z) dz \geq \underline{u}_p \right\}.$$

It is easy to see that $\underline{c}(\pi)$ is decreasing in π and that $\underline{c}(\pi) < \underline{c}(0) = \bar{c}$ according to 4.4. Proposition 4.1 characterizes the set of all PBE of the stage game.

Proposition 4.1. *All PBE of the static game are pooling equilibria. For any $\hat{c} \in [\underline{c}(\pi), \bar{c}]$ there exists a PBE in which both types choose \hat{c} as the commitment cost.*

Proof. Any equilibrium must induce p to trust since d can always choose $c > \bar{U}$ (so it will never be optimal to take the emergency action) and get a payoff of $0 > \underline{u}_d$. There cannot be any pooling equilibrium with $c < \underline{c}(\pi)$, since the definition of $\underline{c}(\pi)$ implies it would give p less than his reservation utility \underline{u}_p . It cannot happen either if $c > \bar{c}$, since either type would deviate and choose \bar{c} and induce p to trust, regardless of his updated beliefs $\pi_p(\bar{c})$. This follows from

$$(4.7) \quad \begin{aligned} & \pi_p(\bar{c}) \int_{U_p > \bar{c}} U_p f(z) dz + [1 - \pi_p(\bar{c})] \int_{U_{old} > \bar{c}} U_p f(z) dz \stackrel{(1)}{>} \\ & \pi_p(\bar{c}) \mathbb{E}[\max(U_p - \bar{c}, 0)] + [1 - \pi_p(\bar{c})] \underline{u}_p \stackrel{(2)}{>} \pi_p(\bar{c}) \underline{u}_p + [1 - \pi_p(\bar{c})] \underline{u}_p = \underline{u}_p \end{aligned}$$

where (1) follows from definition 4.4 and (2) from the fact that $0 > \underline{u}_p$. I will now show that for any $\hat{c} \in [\underline{c}(\pi), \bar{c}]$ there exist a pooling equilibrium in which both $\theta = new$ and $\theta = old$ find it optimal to choose $c = \hat{c}$. Conjecture the following belief updating rule:

$$(4.8) \quad \pi_p^{\hat{c}}(c) := \begin{cases} 0 & \text{if } c < \hat{c} \\ \pi & \text{if } c \geq \hat{c}. \end{cases}$$

Under a pooling equilibrium, since $\hat{c} \geq \underline{c}(\pi)$, p will trust d . Neither type will deviate from \hat{c} since the optimal deviation that would make p trust would be to choose $\hat{c} = \bar{c}$. The non-existence of other PBE is left to the appendix. \square

Suppose now that, in light of the results of Proposition 4.1, the reformed decision maker $\theta = new$ is considering what commitment cost to choose. If we are thinking about a policy prescription that is supported by some PBE of this game, the multiplicity of equilibria in Proposition 4.1 requires some equilibrium refinement. If the policy prescription was the commitment cost supported by the *best equilibrium* of the game, Proposition 4.1 shows that $c = \underline{c}(\pi)$ is selected for either $\theta = new$ or $\theta = old$. We will argue that this will not be a particularly “robust” policy in the sense that expecting p to trust after observing $c = \underline{c}(\pi)$ relies on strong and sensitive assumptions about p ’s beliefs, which may be imperfectly known by the decision maker.

First, observe that the best PBE policy is very sensitive to the prior. If p ’s true prior were $\tilde{\pi} = \pi - \epsilon$ for some $\epsilon > 0$, then p would not trust after observing $c = \underline{c}(\pi)$. Second, even if π were commonly known, after the commitment cost has been chosen, p updates

beliefs to $\pi_p(c) \in (0, 1)$. As illustrated by the proof of Proposition 4.1 (in particular the belief updating rule in 4.8), the indeterminacy of beliefs following zero probability events generates a large set of potential beliefs that can arise in equilibrium. As such, p 's behavior will depend on the complete specification of her updates beliefs for *all off-equilibrium costs*, not just the candidate equilibrium one. Therefore, small changes in the updating rule (for example, by changing \hat{c} in 4.8) generates potentially very different behavior for p .¹¹

Our main question becomes whether we can choose a policy that is robust to mis-specifications of both the prior π and the updating rule $\pi_p(c)$. It is clear that by choosing $c = \infty$ and effectively removing all flexibility, a rational p would trust d independently of his beliefs. However, we can do better. Inequality 4.7 implies that if $c = \bar{c}$ and p will find it optimal to trust irrespective of the updating rule $\pi_p(c)$ if he still believes he is facing a rational decision maker. I will show that in fact $c = \bar{c}$ is the *only* robust policy, when the only assumption we make about p 's beliefs is that they are consistent with strong common certainty of rationality; i.e. p believes he is facing a rational d if her observed past behavior is consistent with common knowledge of rationality.

Since the difference between the reformed and the time inconsistent type is about their ex-post incentives, Proposition 4.1 gives a negative result: types cannot separate in any equilibrium of the stage game, by their choice of c . Only by having repeated interactions can the reformed decision maker hope to convince p of the success of the reform, trying to signal her type through her reactions to the realized shocks. Throughout the remainder of the paper I will investigate whether robust policies, such as the one I found in the static game, can eventually convince p that $\theta = new$, regardless of his particular belief updating rule.

4.2. Repeated game: Setup and basic notation. I extend the stage game to an infinite horizon setting: $\tau \in \{0, 1, \dots\}$. I assume that d is infinitely lived and that types are permanent. That is, at $\tau = 0$ nature chooses $\theta = new$ with probability π_{new} . d has discounted expected utility with discount factor $\beta_\theta \in (0, 1)$. For notational ease, I will assume $\beta_{old} = \beta_{new} = \beta$. I will specify when the results are sensitive to this assumption. Shocks are iid across periods: $z_\tau := (U_{p,\tau}, U_{old,\tau}) \sim_{i.i.d} f(z_\tau)$. I assume that there is a sequence of myopic short run players p_τ (or equivalently $\beta_p = 0$) which is a standard assumption in the reputation literature (Fudenberg and Levine [1989], Phelan [2006]). This will be without loss of generality for most applications to macroeconomic models applications.¹² At every

¹¹More generally, we apply Aumann and Brandenburger [1995] results to the interim normal form of this game, finding tight sufficient conditions for any particular Bayesian equilibria (not necessarily perfect) to be the expected solution outcome: (a) There is common knowledge of rationality, (b) the strategies of both $\theta = new$ and $\theta = old$ prescribed by the Bayesian equilibrium are common knowledge, and (c) the inference rule $\pi_p(\cdot)$ is also common knowledge

¹²Celentani and Pesendorfer [1996] show that this assumption is without loss of generality when p is modeled as representative agent for a continuum of atomistic and anonymous patient agents. In particular, the capital taxation model of section 2.1 satisfies these assumptions when capitalist households have a common discount rate $\delta_k \in (0, 1)$.

period, d chooses $c_\tau \geq 0$ which is binding only for that period. The policy maker can change its choice freely in every period. I also assume that all past history of actions and shocks (except for d 's payoff type) is observed by all players at every node in the game tree. I will further assume that the structure of the game described so far is *common knowledge* for both players and that agents know their own payoff parameters.¹³

A stage τ *outcome* is a 4-tuple $h_\tau = (c_\tau, a_\tau, z_\tau, r_\tau)$ where c_τ is the commitment cost, $a_\tau \in \{0, 1\}$ is the trust decision, and $r_\tau \in \{0, 1\}$ is the contingent policy, where $r_\tau = 1$ if d chooses the emergency action, and $r_\tau = 0$ otherwise. A *history* up to time τ is defined as

$$h^\tau := (h_0, h_1, \dots, h_{\tau-1}).$$

I will refer to a “partial history” as a history plus part of the stage game. For example, when p is called to decide whether to trust or not, his information set is (h^τ, c_τ) which is the history of events so far, plus the commitment cost in the (not yet revealed) history $h^{\tau+1}$. The set of all partial histories will be denoted as \mathcal{H} , and $\mathcal{H}_i \subseteq \mathcal{H}$ is the set of histories in which agent $i \in \{p, d\}$ has to take an action.

A *strategy* for the policy maker is a function $\sigma_d : \mathcal{H}_d \rightarrow C \times \{0, 1\}^Z$ that specifies, at the start of every period τ , a commitment cost c_τ and the contingent choice provided p trusts. Then, we can always write a strategy as a pair:

$$(4.9) \quad \sigma_d(h^\tau) = (c^{\sigma_d}(h^\tau), r^{\sigma_d}(h^\tau, \cdot)) : Z \rightarrow \{0, 1\}$$

where the choice is a commitment cost $c^{\sigma_d}(h^\tau) \in C$ and a *policy rule function* $r^{\sigma_d}(h^\tau, z)$ of the shock. The superscript σ_d serves to remind the reader these objects are part of a single strategy σ_d . Likewise, a strategy σ_p for p is a function $\sigma_p : \mathcal{H}_p \rightarrow \{0, 1\}$ that assigns to every observed history, his trust decision

$$\sigma_p(h^\tau, c_\tau) = a^{\sigma_p}(h^\tau, c_\tau) = \begin{cases} 1 & \text{if } p \text{ trusts} \\ 0 & \text{if } p \text{ does not trust.} \end{cases}$$

The set of strategies of each agent as Σ_i for $i \in \{d, p\}$ are compact sets when endowed with the product topology.¹⁴ Also let $\Sigma = \Sigma_d \times \Sigma_p$ be the set of strategy profiles $\sigma = (\sigma_d, \sigma_p)$.

If player $i \in \{d, p\}$ plays strategy σ_i , the set of histories that will be consistent with σ_i is denoted $\mathcal{H}(\sigma_i) \subset \mathcal{H}$. For a history $h \in \mathcal{H}$ we say a strategy σ_i is *consistent with* h if $h \in \mathcal{H}(\sigma_i)$. Let $\Sigma_i(h) = \{\sigma_i \in \Sigma_i : h \in \mathcal{H}(\sigma_i)\}$ be the set of strategies consistent with h .

Given a strategy profile $\sigma = (\sigma_p, \sigma_d)$ let $W_\theta(\sigma | h)$ be the expected continuation utility for d 's type $\theta \in \{old, new\}$ given history h

$$(4.10) \quad W_\theta(\sigma | h) := (1 - \beta) \mathbb{E} \left\{ \sum_{s=\tau}^{\infty} \beta^{s-\tau} \left[a_s r_s (U_{\theta,s} - c_s) + (1 - a_s) \underline{u}_p \right] \mid h \right\}$$

¹³These are the basic assumptions in Penta [2012].

¹⁴A sequence $\{\sigma_{i,n}\}_{n \in \mathbb{N}}$ converges to σ_i in the product topology if and only if $\sigma_{i,n}(h) \rightarrow \sigma_i(h)$ for all $h \in \mathcal{H}_i$

where $c_s = c^{\sigma_d}(h^s)$, $a_s = a^{\sigma_p}(h^s, c_s)$ and $r_s = r^{\sigma_d}(h^s, z_s)$. Likewise, denote $V(\sigma | h)$ for the spot utility for agent p at history h

$$(4.11) \quad V(\sigma | h) := a_\tau \mathbb{E}_z(r_\tau U_{p,\tau}) + (1 - a_\tau) \underline{u}_p.$$

In Appendix (B) I show that when endowed with the product topology, $W_\theta(\cdot | h)$ and $V(\cdot | h)$ are continuous functions of the strategy profile σ for all $h \in \mathcal{H}$

4.3. Systems of Beliefs. Agents form beliefs both about the payoff types of the other player, as well as the strategies that they may be planning to play. In static games, such beliefs are characterized by some distribution $\pi \in \Delta(\Theta_{-i} \times S_{-i})$ where Θ_{-i} is the set of types of the other agent and S_{-i} their strategy set. In our particular game, $\Theta_d = \{new, old\}$ and $\Theta_p = \{p\}$. In dynamic settings however, agents may *revise their beliefs* after observing the history of play. This revision is described by a *conditional probability system*, that respects Bayes rule whenever possible.

Formally, let $X_i = \Theta_{-i} \times \Sigma_{-i}$, \mathcal{X}_i the Borel σ -algebra generated by the product topology on X_i and $\mathcal{K}_i = \{A \in \mathcal{X}_i : \text{proj}_{\Sigma_{-i}} A = \Sigma_{-i}(h) \text{ for some } h \in \mathcal{H}\}$ be the set of all events observed by i . A *system of beliefs* π_i on X_i is a mapping $\pi_i : \mathcal{K}_i \rightarrow \Delta(X_i)$ such that:

- (1) Given $A \in \mathcal{K}_i$, $\pi_i(\cdot | A)$ is a probability measure over $\Theta_{-i} \times \Sigma_{-i}$.¹⁵
- (2) If $A \subseteq B \subseteq C$ with $B, C \in \mathcal{K}_i$, then $\pi_i(A | B) \pi_i(B | C) = \pi_i(A | C)$.

I write $\pi_i(A | h) = \pi_i(A | \Sigma_i(h))$ for $A \subset \Theta_{-i} \times \Sigma_{-i}$ for the probability assessment of event A conditional on history h . Denote $\Delta^{\mathcal{H}}(\Theta_{-i} \times \Sigma_{-i})$ to be the set of all systems of beliefs. Given $\pi_d \in \Delta^{\mathcal{H}}(\Theta_p \times \Sigma_p) = \Delta^{\mathcal{H}}(\Sigma_p)$ and strategy $\sigma_d \in \Sigma_d$, define $W_\theta^{\pi_d}(\sigma_d | h)$ as the expected continuation payoff for type $\theta \in \{old, new\}$ conditional on history h , under beliefs π_d :

$$(4.12) \quad W_\theta^{\pi_d}(\sigma_d | h) := \mathbb{E}_{\sigma_p}^{\pi_d} [W_\theta(\sigma_d, \sigma_p | h)].$$

Analogously, given $\pi_p \in \Delta^{\mathcal{H}}(\Theta_d \times \Sigma_d)$ and $\sigma_p \in \Sigma_p$ the expected utility of strategy σ_p under π_p is

$$(4.13) \quad V^{\pi_p}(\sigma_d | h) := \mathbb{E}_{\sigma_p}^{\pi_p} [V(\sigma_d, \sigma_p | h)].$$

For a given system of beliefs π_i we write $\sigma_i \in SBR_{\theta_i}(\pi_i)$ as the set of sequential best responses of type θ_i to beliefs π_i .¹⁶ The definitions in 4.10 and 4.11 are the special cases where agents put probability 1 on a specific strategy.

4.4. Weak and Strong Rationalizability. In this subsection I introduce the notions of weak and strong rationalizability. As mentioned in Section 4.1, our goal is to find strategies that are robust to changes in p 's beliefs so that p is induced to trust as long as there is

¹⁵We endow X_i with the Borel σ -algebra with respect to the product topology on $\Theta_{-i} \times \Sigma_{-i}$.

¹⁶A strategy σ_p is a *sequential best response to π_p* for all $(h^\tau, c_\tau) \in \mathcal{H}_p$ and all other strategies $\hat{\sigma}_p \in \Sigma_p$, we have $V^{\pi_p}(\sigma_p | h^\tau, c_\tau) \geq V^{\pi_p}(\hat{\sigma}_p | h^\tau, c_\tau)$. Likewise, σ_d is a *sequential best response to belief system π_d* for type $\theta \in \{new, bad\}$ if for all histories $h \in \mathcal{H}_d$ and all strategies $\hat{\sigma}_d \in \Sigma_d$ we have $W_\theta^{\pi_d}(\sigma_d | h) \geq W_\theta^{\pi_d}(\hat{\sigma}_d | h)$.

common certainty of rationality, which means that all agents are rational, all agents are certain that all agents are rational, all agents are certain that all agents know this, ad infinitum. In static settings, beliefs that satisfy these common knowledge assumptions have their support over the set of *rationalizable strategies*. This set is characterized by an iterative deletion process described in Pearce [1984]. The set of strategies are refined by eliminating those which are not a best response to *some* beliefs about the other agents strategies, which are themselves best responses to some other beliefs, and so on.

However, the possibility of reaching zero probability events also creates different ways to extend the concept of rationalizability, which hinge upon on our notion of “certainty or rationality”. An agent is *certain* about some event E if she believes that this event happens with probability 1.¹⁷ The possibility of reaching zero probability events also creates different notions of certainty, depending on how beliefs are revised after such unexpected events arise. We say that a history $h \in \mathcal{H}$ is *consistent with event* $E \subseteq \Theta_{-i} \times \Sigma_{-i}$ if there exist a strategy $\sigma_{-i} \in \text{proj}_{\Sigma_{-i}} E$ such that $h \in \mathcal{H}(\sigma_{-i})$. Abusing the notation somewhat I will write $h \in \mathcal{H}(E)$ for such histories.

Definition 4.1 (*Weak Certainty of event E*). A system of beliefs $\pi_i \in \Delta^{\mathcal{H}}(\Theta_{-i} \times \Sigma_{-i})$ is *weakly certain of event* $E \subseteq \Theta_{-i} \times \Sigma_{-i}$ if $\pi_i(E | h^0) = 1$ (i.e. π_i assigns probability 1 to event E at the beginning of the game).

Definition 4.2 (*Strong Certainty in event E*). A system of beliefs $\pi_i \in \Delta^{\mathcal{H}}(\Theta_{-i} \times \Sigma_{-i})$ is *strongly certain of event* $E \subseteq \Theta_{-i} \times \Sigma_{-i}$ if $\pi_i(E | h) = 1$ for all $h \in \mathcal{H}(E)$ (i.e. π_i assigns probability 1 to event E at all histories consistent with it).

To illustrate the difference between both concepts, suppose p has a belief system π_p , that is certain of some event E , and is also certain about a smaller event $F = \{(new, \sigma_{new}), (old, \sigma_{old})\} \subset E$. That is, he is certain about what strategy each type of player d chooses (which is the required assumption in the construction of a Bayesian equilibrium). However, π_p may be an incorrect prediction of d 's behavior. Take a history h in which p realizes that the observed history is *not* consistent with the strategies in F but it is nevertheless consistent with event E : i.e. $\{\sigma_{new}, \sigma_{old}\} \cap \Sigma_d(h) = \emptyset$ but $h \in \mathcal{H}(E)$. If π_p is *weakly certain* of event E , then after the unexpected move by d , no restrictions are imposed on the updated beliefs from history h on. In particular, he is not required to remain certain about event E , even if the observed history is consistent with it. On the other hand, if π_p is *strongly certain* about event E , he would realize his beliefs about event F were wrong, but his updated beliefs would remain certain about event E . In a way, the concept of strong certainty is similar

¹⁷When the event E is also true we say that the type *knows* E . This admits the possibility that an agent believes with probability one an event that is indeed false. In static games, because the game ends right after the payoffs are realized, there is no substantive difference between *certainty* and *knowledge*. In dynamic games the situation is more subtle, since an agent's beliefs may be proven wrong (or refuted) by the observed path of play. Because of this feature, the literature has focused on the concept of certainty (Ben-Porath [1997], Battigalli and Siniscalchi [1999], Penta [2011, 2012]) instead of knowledge, for dynamic games.

to an agent that knows that event E is true, and her updated beliefs should respect it as a “working hypothesis” (Battigalli and Siniscalchi [2002])

These two different notions of certainty will give rise to two different notions of rationalizability. Define the set of sequentially rational outcomes $R_i \subset \Theta_i \times \Sigma_i$ as

$$(4.14) \quad R_i = \left\{ (\theta_i, \sigma_i) : \sigma_i \in SBR_{\theta_i}(\pi_i) \text{ for some } \pi_i \in \Delta^{\mathcal{H}}(\Theta_{-i} \times \Sigma_{-i}) \right\}.$$

The set R_i gives all the strategies and payoff types such that σ_i is the sequential best response to some system of beliefs.

I will now formally follow the iterative procedure of Battigalli and Siniscalchi [2003]. For a given set $E \subset \Theta_{-i} \times \Sigma_{-i}$ write $\mathbf{W}_i(E) \subset \Delta^{\mathcal{H}}(\Theta_{-i} \times \Sigma_{-i})$ to be the set of beliefs π_i that are weakly certain of E . Analogously, define $\mathbf{S}_i(E) \subset \mathbf{W}_i(E)$ for the set of beliefs that are strongly certain of it.

I now give an iterative definition of the sets $WCR_i^k \subset \Theta_i \times \Sigma_i$ and $SCR_i^k \subset WCR_i^k$, the sets of type-strategy pairs for agent i that are consistent with k rounds of mutual weak (strong) certainty of rationality. For $k = 0$, define

$$WCR_i^0 = SCR_i^0 = R_i.$$

For $k > 1$, define iteratively:

$$(4.15) \quad WCR_i^k := \left\{ (\theta_i, \sigma_i) : \begin{cases} (1) : (\theta_i, \sigma_i) \in WCR_i^{k-1} \\ (2) : \exists \pi_i \in \mathbf{W}_i(WCR_{-i}^{k-1}) : \sigma_i \in SBR_{\theta_i}(\pi_i) \end{cases} \right\}$$

$$(4.16) \quad SCR_i^k := \left\{ (\theta_i, \sigma_i) : \begin{cases} (1) : (\theta_i, \sigma_i) \in SCR_i^{k-1} \\ (2) : \exists \pi_i \in \mathbf{S}_i(SCR_{-i}^{k-1}) : \sigma_i \in SBR_{\theta_i}(\pi_i) \end{cases} \right\}.$$

So intuitively, we start with beliefs that are weakly (strongly) certain of event $E = R_{-i}$ and then we proceed with an iterative deletion procedure, in which the set agent i is weakly (strongly) certain about is the set $E = WCR_{-i}^{k-1}$ and similarly for strong certainty. Finally, the sets of *weak and strong rationalizable outcomes* is defined as

$$(4.17) \quad WCR_i^\infty = \bigcap_{k \in \mathbb{N}} WCR_i^k$$

$$(4.18) \quad SCR_i^\infty = \bigcap_{k \in \mathbb{N}} SCR_i^k$$

The sets $WCR_i^\infty, SCR_i^\infty \subset \Sigma_i$ are the sets of strategies for i that are consistent with him having weak (strong) common certainty of rationality. I will denote \mathcal{B}_i^w and \mathcal{B}_i^s as the sets of weak and strong rationalizable beliefs for p , respectively

$$(4.19) \quad \mathcal{B}_i^w := \Delta^{\mathcal{H}}(WCR_{-i}^\infty) \text{ and } \mathcal{B}_i^s := \Delta^{\mathcal{H}}(SCR_{-i}^\infty).$$

I will say that a strategy-belief pair (σ_i, π_i) is θ_i -weak rationalizable (or simply p -weak rationalizable for the case $i = p$) whenever $\pi_i \in \mathcal{B}_i^w$ and $\sigma_i \in SBR_{\theta_i}(\pi_i)$. A history h is θ_i -weak rationalizable whenever $h \in \mathcal{H}(\sigma_i)$ for some weak rationalizable pair (σ_i, π_i) . I will refer to such pairs as a θ_i -weak rationalizations of h . I also define the analogous notions for strong rationalizability.

4.5. Discussion. In the above we have characterized the multiplicity of equilibria in the static game and have established the setup of the repeated game including belief systems and notions of rationalizability. The next section will turn to robust implementation. We briefly connect the concepts now, and argue that equilibrium refinements are not robust to a variety of perturbations we might think of.

First, and most importantly for our applications is the one considered in this paper which is robustness to *strategic uncertainty*. Recall that the static game had multiple equilibria. The dynamic game only exacerbates this problem via classical folk theorem like arguments.¹⁸ This suggests that in order for either to form predictions or to make policy recommendations, some equilibrium refinement is needed. In macroeconomics, the most commonly used refinement is using optimal or efficient equilibria. We note that in some contexts this may be reasonable. For example, it would be a reasonable prediction if agents could meet and agree upon a desired outcome before the game started and are able to decide both the expected behavior by all agents, the punishments that should be sanctioned to deviators, all of these subject to the constraint that these should be self-enforceable. However, our environment is one in which the public has no reason to agree with the time inconsistent type and as such refinement via the selection of optimal equilibrium seems suspect. A second limitation of equilibrium refinements is that they are very sensitive to common knowledge assumptions about the payoff structure of the game, which in our case is the fact that there are only two types for d and one payoff type for p , as well as their discount rate. If we allow the set of feasible payoff structures to be “big enough”, then picking a Nash equilibrium of the game and the belief systems that support it, then arbitrarily small perturbations on the beliefs may pick *any* other weak rationalizable outcome as the unique equilibrium of the perturbed game (e.g., [Weinstein and Yildiz \[2012, 2007\]](#), [Penta \[2012\]](#)). Formally, consider a modification to the game structure, that allows for a larger set of payoff types for each agent but retaining the same discount rates and the assumption that every agent knows its type throughout the game, and never doubts it.

Formally, if the set of feasible payoff structures satisfies a richness condition¹⁹ and agents know their own payoff type throughout the game, then given any Bayesian equilibrium profile σ_N and any other weakly rationalizable strategy profile $\hat{\sigma}$ of the original unperturbed

¹⁸See [Mailath and Samuelson \[2006\]](#) for an exhaustive review on these topics.

¹⁹Formally, for every strategy σ_i there exist a type $\hat{\theta}_i(\sigma_i) \in \hat{\Theta}_i$ such that σ_i is conditionally dominant for type $\hat{\theta}_i(\sigma_i)$ at every history consistent with it: i.e. $W_{\hat{\theta}_i}(\sigma_i, \sigma_{-i} | h) > W_{\hat{\theta}_i}(\hat{\sigma}_i, \sigma_{-i} | h)$ for all $\hat{\sigma}_i \in \Sigma_i, \sigma_{-i} \in \Sigma_{-i}, h \in \mathcal{H}_i(\sigma_i)$.

game, we can find a hierarchy of beliefs on the larger game that is arbitrarily close (in the product topology) to the beliefs associated with σ_N but is such that $\hat{\sigma}$ is selected as the unique equilibrium of the game with the perturbed beliefs. This then implies that under these assumptions, the only concept that is robust to small perturbations of beliefs is weak rationalizability, and hence only predictions that hold for all weak rationalizable strategy profiles are robust to these perturbations.²⁰

The previous result also implies that strong rationalizability will be also subject to this lack of robustness. However, the richness assumption may be too a stringent condition for our robustness exercise, since we are ultimately interested in modeling robustness to strategic uncertainty. Strong rationalizability, being a stronger solution concept may not be robust to all of these perturbations in payoff structures, but we briefly study some in Section 7 how to create policies that are robust to certain perturbations over the set of payoff type spaces.

I now discuss one of the main implications of Strong Rationalizability: The fact that agents can be convinced at some histories that certain payoff types are not consistent with the history observed. Suppose that p reaches a history that is not consistent with both strong common certainty of rationality and $\theta = old$, but it is consistent with $\theta = new$. That is, no strongly rationalizable strategy of the time inconsistent type could have generated the observed history. Strong common certainty of rationality implies that at these histories p must be certain that $\theta = new$ for all strong rationalizable continuation histories; it becomes *common knowledge* that $\theta = new$, and the game transforms in practice to a game of complete information.²¹ When this happens, we will say that $\theta = new$ has achieved *full or strong separation* from $\theta = old$. This will be one of the key ingredients of the robust optimal strategy of the section (5): the reformed decision maker can gain reputation by taking actions that $\theta = old$ decision maker would never take, or that at least would be very costly for her.

Although this consequence is of great practical relevance for the policy design, we can see now the potential lack of robustness of strong rationalizability to perturbations over the set of payoff types: after a history that is not *old*-rationalizable there might exist *some* alternative payoff type $\hat{\theta} \notin \{new, old\}$ that rationalizes the observed history, and hence we do not reach full separation afterwards. Then, for sufficiently rich payoff type spaces, full separation becomes impossible, since we could always find some payoff type that

²⁰Weinstein and Yildiz [2012] show that when we relax the restriction that all players know their own type at the beginning of the game (and never abandon this belief), then the only robust solution concept is normal form interim correlated rationalizability (ICR), extending their previous result on static games (Weinstein and Yildiz [2007]).

²¹If at some continuation history p observes behavior that is inconsistent with $\theta = new$ playing a strongly rationalizable strategy, p abandons the assumption of strong common certainty of rationality, which then allows him to believe that $\theta = old$ after all. When this happens, we apply the “best-rationalization principle” as in Battigalli and Siniscalchi [2002]. It states that whenever p arrives at such a history, she will believe that there are at least k -rounds of strong common certainty of rationality, with k the highest integer for which the history is consistent with k rounds of strong rationalizability.

rationalizes the observed behavior. On the other hand, if we have some restrictions about the set of possible payoff types, then separation may still be feasible, both from all other types (complete reputation gain) or a subset of potential inconsistent types. In Section 7 we discuss how, by explicitly modeling the extensions over the set of payoff types, study strong rationalizable outcomes.

5. Robust Implementation

This section introduces the notion of robust implementation to a given set of restrictions on p 's beliefs (subsection 5.1) and solves for the robust implementing policies for two important benchmarks: weak and strong rationalizable beliefs. Focusing on strong rationalizable implementation, I characterize the optimal strong rationalizable implementation by solving a recursive dynamic contracting problem with a single promise-keeping constraint. Moreover, for histories where robust separation has not occurred, the relevant reputation measure for d is the *implied spot opportunity cost or sacrifice* for $\theta = old$ of playing $r_{\tau-1}$, so only the immediate previous period matters in terms of building partial reputation. I show that on the outcome path of the optimal robust policy, $\theta = new$ gets both partial gains and (endogenous) losses of reputation until robust separation is achieved. After this, the game essentially becomes one with complete information.

5.1. Definition. The decision maker has some information about p 's beliefs or may be willing to make some assumptions about them. She considers that p 's possible beliefs lie in some subset $\mathcal{B}_p \subset \Delta^{\mathcal{H}}(\Theta_d \times \Sigma_d)$. Write $SBR_p(\mathcal{B}_p) = \bigcup_{\pi_p \in \mathcal{B}_p} SBR_p(\pi_p) \subset \Sigma_p$ as the set of all sequential best responses to beliefs in \mathcal{B}_p . We will say that a strategy σ_d is a *robust implementation of trust in \mathcal{B}_p* when it induces p to trust d at all $\tau = 0, 1, 2, \dots$, provided d knows that **(1)** p 's beliefs are in \mathcal{B}_p and **(2)** p is sequentially rational.

Definition 5.1 (*Robust Implementation*). A strategy $\sigma_d \in \Sigma_d$ robustly implements trust in \mathcal{B}_p if, for all histories $(h^\tau, c_\tau) \in \mathcal{H}_p(\sigma_d)$ we have

$$a^{\sigma_p}(h^\tau, c_\tau) = 1 \text{ for all } \sigma_p \in SBR_p(\mathcal{B}_p)$$

Under the assumptions on the stage game, for a given belief system π_p , its sequential best response will be generically unique. Therefore, if d knows both that p is rational and that he has beliefs π_p , then she can predict the strategy that p will choose. Following this reasoning, we can think of the different beliefs as different *types* for agent p , and d 's uncertainty simply as having an extended incomplete information game, in which now there is incomplete information about p , with set of types $\hat{\Theta}_p = \mathcal{B}_p$. This is the idea behind the notion of *epistemic type spaces* introduced by Harsanyi [1967/68]: uncertainty about other agents beliefs (and higher order beliefs) can be modeled as enlarging the set of types for each agent. The notion of robustness in the robust mechanism design literature (Bergemann and Morris [2005, 2009, 2011]) uses this notion of uncertainty. They seek to implement

a desired outcome for any *type space* that is consistent with some common knowledge of rationality assumptions. In Appendix A I discuss an alternative but equivalent approach to robust implementation using type space implementation, extending the concepts described in Bergemann and Morris [2009] to our dynamic settings and using instead the interactive epistemic characterization of Battigalli and Siniscalchi [2003].

Although both approaches are equivalent, the epistemic type space approach might be more useful when we want to place restrictions not only on what p expects d to play, but also about p 's beliefs about d 's beliefs as well, and higher order beliefs. In subsection (6.3) I show how relevant restrictions on robust policies map into assumptions about p 's higher order beliefs. In what follows I will study the robust policies for both weak and strong rationalizable beliefs.

5.2. Weak Rationalizable Implementation. I begin with the most lax notion of rationalizability at our disposal – weak rationalizability; i.e. $\mathcal{B}_p = \mathcal{B}_p^w$. Here I show that this notion of rationalizability delivers a rather stark and, in some sense, negative result: only by eliminating the emergency action entirely can d robustly implement trust. The intuition is that because the game I study is of private values,²² the public cares only about the decision maker's strategy. Given the multiplicity of commitment costs that are consistent with common certainty of rationality, at a history (in fact, in the first stage) there will be at least two costs that the policy maker can choose. The following type of beliefs are rationalizable: believe d is rational only if she takes one of these specified decisions but is actually thought to be irrational if she takes the other (and vice versa). Then it becomes impossible to implement trust in both belief types, unless d gets rid of the emergency action altogether.

Proposition 5.1. *The unique robust implementing policy in $\mathcal{B}_p = \mathcal{B}_p^w$ involves $c_\tau = \infty$ in all periods.*

Thus, the result is that weak rationalizability is too weak a concept to be used for our purposes. After an unexpected commitment cost choice, p could believe d to be *irrational* and never trust d again unless r is removed. This sort of reasoning does not take into account a restriction that, say, if p could find some other beliefs under which d would be rational, then this now becomes p 's working hypothesis. This is precisely the notion of strong rationalizability, which I explore below.

5.3. Strong Rationalizable Implementation. The next three subsections present the main results of the paper in which I characterize the optimal strong rationalizable implementation. In this subsection, I will show that an optimal robust implementing strategy

²²A Bayesian game is of *private values* if utility for each agent depends only on their own payoff parameter, and not about the other agents payoffs. Formally, $\Gamma = \{A_i, \Theta_i, u_i : A \times \Theta \rightarrow \mathbb{R}\}_{i=1}^I$ is of private values if $u_i(a, \theta) = u_i(a, \theta_i)$ for all $a \in A \equiv \times_{i=1}^I A_i$ and all $\theta \in \Theta \equiv \times_{i=1}^I \Theta_i$

corresponds to the sequential best response to some strong rationalizable system of beliefs. In that sense, an optimal robust implementing strategy will be equivalent to finding the most pessimistic beliefs that d could have about p 's behavior, that is consistent with common strong certainty of rationality. Most importantly, I will also show that any optimal robust strategy will be in fact the *min-max* strategy for d , delivering the best possible utility that d can guarantee at any continuation history, regardless of her system of beliefs.

To simplify notation, denote $\Sigma_i^* = \text{proj}_{\Sigma_i} SCR_i^\infty$ for the set of extensive form rationalizable strategies for agent i . Abusing the notation somewhat I will also write $\Sigma_\theta^* = \text{proj}_{\Sigma_d} \left\{ (\hat{\theta}, \hat{\sigma}_d) \in SCR_d^\infty : \hat{\theta} = \theta \right\}$ to represent the set of extensive form rationalizable strategies for type. Unless we specify otherwise, I will refer to strong rationalizability as rationalizability for the remainder of this paper. I specialize definition 5.1 to the case of strong rationalizability below.

Definition 5.2. A strategy $\sigma_d \in \Sigma_d$ is a *strong rationalizable implementation of trust* if for all $(h^\tau, c_\tau) \in \mathcal{H}_p(\sigma_d)$ and all strong rationalizable strategies $\sigma_p \in \Sigma_p^*$ we have $a^{\sigma_p}(h^\tau, c_\tau) = 1$. The set of all such strategies is denoted $\mathbf{SR} \subset \Sigma_d$.

For the rest of the paper, we will refer to these as *robust strategies*. With some abuse of notation, when $\sigma_d \in \mathbf{SR}$ we will write $W_\theta(\sigma_d | h)$ as its expected continuation payoff given history h , which is belief independent,²³ because its robustness implies that d should not face any uncertainty about p 's strategy if σ_d is in fact a robust strategy. We know this set is non-empty, since $c = \infty$ for all histories robustly implements trust for $\mathcal{B}_p^w \supset \mathcal{B}_p^s$. We will restrict our attention to *optimal robust strategies*: *i.e.* strategies that maximize the expected (ex-ante) payoff for d , at $\tau = 0$

$$(5.1) \quad \sigma_\theta^* \in \underset{\sigma_d \in \mathbf{SR}}{\text{argmax}} \mathbb{E} \left\{ (1 - \beta) \sum_{\tau=0}^{\infty} \beta^\tau [U_{\theta,\tau} - c_\tau^{\sigma_d}(h^\tau)] r_\tau^{\sigma_d}(h^\tau, z_\tau) \right\}$$

It can be easily shown that 5.1 implies that it must also be optimal for all histories on its generated path $h \in \mathcal{H}(\sigma_\theta^*)$. The goal for the rest of the paper is to characterize the optimal robust strategy for the reformed payoff type $\theta = \text{new}$. As a warm up, I solve for the optimal robust strategy in the stage game.

Example 5.1 (Strong Rationalizable Implementation in the static game). Take the stage game of Section 4.1. Then, there exist only two robust commitment cost choices: $c = \bar{c}$ and $c = \infty$, which implies that the optimal robust implementation of trust is $c = \bar{c}$, for either $\theta \in \{\text{new}, \text{old}\}$. This follows from the argument in the proof of Proposition 4.1: Irrespective of the system of beliefs $\pi_p(c)$, if p is certain that he is facing a rational d , he will find it optimal to trust. For $c < \bar{c}$ there always exist strong rationalizable beliefs that induces p not to trust (see the construction of belief system in 4.8), and for $c > \bar{c}$ p cannot

²³This is because any robust strategy has payoffs $W_\theta(\sigma_d | h) = \mathbb{E} \left\{ (1 - \beta) \sum_{\tau=0}^{\infty} \beta^\tau r^{\sigma_d}(h^\tau, z_\tau) [U_{\theta,\tau} - c^{\sigma_d}(h^\tau)] | h \right\}$

expect to be facing a rational decision maker, since it would have been a dominant strategy just to play $c = \bar{c}$, that induces p to trust and gives d a strictly higher payoff, regardless of her type.

Theorem B.3 in Appendix B.2 shows that Σ_i^* is a compact set, and so are the subsets $\Sigma_i^*(h) \subset \Sigma_i^*$ of history consistent strategies. This implies that existence of payoff functions $\underline{W}_\theta, \bar{W}_\theta : \mathcal{H}_d \rightarrow \mathbb{R}$ for $\theta \in \{new, old\}$ such that, for all $h \in \mathcal{H}$

$$(5.2) \quad \underline{W}_\theta(h) \leq W_\theta^{\pi_d}(\sigma_d | h) \leq \bar{W}_\theta(h) \text{ for all } \pi_d \in \mathcal{B}_d^s, \sigma_d \in SBR_\theta(\pi_d).$$

I will refer to $\underline{W}_\theta(\cdot)$ and $\bar{W}_\theta(\cdot)$ as the *best* and *worst strong rationalizable payoffs* for type θ . I will also write $\underline{\mathbb{W}}_\theta := \underline{W}_\theta(h^0)$ and $\bar{\mathbb{W}}_\theta = \bar{W}_\theta(h^0)$ for the ex-ante worst (and best) rationalizable payoff, from $\tau = 0$ perspective. The first result relates these bounds to robust implementation. Any optimal robust policy is extensive form rationalizable (i.e. it corresponds to the best response of some rationalizable beliefs) and delivers the worst rationalizable payoff $\underline{W}_\theta(h)$ at all histories and types $\theta \in \{new, old\}$.

Lemma 5.1. *Let σ_θ^* be the optimal robust strategy for type θ . Then $\sigma_\theta^* \in \Sigma_\theta^*$ – that is, there exist a strongly rationalizable system of beliefs $\underline{\pi}_\theta$ such that $\sigma_\theta^* \in SBR_\theta(\underline{\pi}_\theta)$. Moreover, for all $h \in \mathcal{H}$ and any strategy $\sigma_d \in \mathbf{SR} \cap \Sigma_\theta^*(h)$ we have*

$$W_\theta(\sigma_d | h) = \underline{W}_\theta(h),$$

i.e., all robust, strong rationalizable strategies are payoff equivalent, and deliver the worst strong rationalizable payoff at all histories.

Proof. The first part is a consequence of Lemma (C.2) in Appendix C. For the second result, take a robust and strong rationalizable strategy σ_d and suppose there exist a history h and a strong rationalizable pair $(\hat{\sigma}_d, \hat{\pi}_d)$ that deliver an expected payoff that is less than the payoff of the robust policy:

$$W_\theta^{\hat{\pi}_d}(\hat{\sigma}_d | h) < W_\theta(\sigma_d | h).$$

However, if $\hat{\pi}_d$ has common strong certainty of rationality, then she is also certain that p plays strong rationalizable strategies (Proposition 3.10 in Battigalli and Bonanno [1999]), and hence she should be also certain that by following the robust strategy σ_d from history h on she will get a higher expected payoff. Since this is true for any rationalizable belief, $\hat{\sigma}_d$ cannot be the sequential best response for beliefs $\hat{\pi}_d$ (since it is conditionally dominated by σ_d at h), reaching a contradiction \square

The proof of Lemma 5.1 implies a very important corollary. The optimal robust strategy is the min-max strategies for type $\theta \in \{new, old\}$ (as in Mailath and Samuelson [2006]) across all beliefs that are consistent with common strong certainty of rationality. The beliefs $\underline{\pi}_\theta$ corresponds to the min-max beliefs for type θ , the most pessimistic beliefs that type θ can have about the strategy that p may be playing. Therefore, at history h^τ , the optimal robust

policy gives the *best payoff that type θ can guarantee herself*, regardless of her beliefs, as long as p plays some strong rationalizable strategy.

Note here that the worst rationalizable payoff does not coincide with the payoff of the worst Bayesian equilibrium of the extensive form game. Common strong certainty of rationality, strictly speaking, is neither a stronger nor weaker solution concept than Bayesian equilibrium. Applying [Aumann and Brandenburger \[1995\]](#) to the interim normal form game, for a particular Bayesian equilibria to be the predicted outcome of the game, we need the common prior assumption (i.e. both players know $\pi = \Pr(\theta = \text{new})$) together with weak common *knowledge* of rationality and beliefs (i.e. weak common certainty, plus the requirement that the beliefs are correct). While the common certainty of rationality is weaker than strong certainty, this characterization implies a much stronger condition. Agents have common knowledge about the *strategies* that each other will play and these beliefs must be correct.

5.4. Observed Sacrifice and Strong Rationalizability. In this subsection, I derive the central results of our paper which will govern the shape of the optimal robust policy. I will give a characterization of strong rationalizable strategies in terms of the observed behavior and conclude that when studying the feasible rationalizable continuation strategies, we must only take into account the previous period implied opportunity cost payed by d . In this sense, the set of rationalizable continuation strategies is Markovian, with a state variable that is observable by all agents in the game.

Consider a history $(h^\tau, c_\tau) \in \mathcal{H}_p$ observed by agent p . He wants to know whether the observed history of past actions of d can be rationalized as coming from a rational type θ . He also wants to know whether he can derive restrictions about the strategies that type θ may plan to follow in the continuation game. In particular, since p is myopic, he only needs to know what restrictions rationalizability would impose on the state contingent policy function $r(z) = r^{\sigma_d}(h^\tau, z)$, if he decided to trust. Formally, given a history (h^τ, c_τ) we say that policy function $r(z)$ is θ -rationalizable at (h^τ, c_τ) if there exist some θ -rationalizable pair (σ_d, π_d) that is consistent with history (h^τ, c_τ) and that delivers the policy function $r^{\sigma_d}(h^\tau, z) = r(z)$. Our goal for this section will be to characterize the set of all θ -rationalizable policy functions.

Consider histories (h^τ, c_τ) such that $r_{\tau-1} = 0$ and $U_{\theta, \tau-1} - c_{\tau-1} > 0$. Let $h^\tau(r=1)$ be the continuation history had d chosen $r_{\tau-1} = 1$ instead. Then, a θ -rationalizable pair (σ_d, π_d) is consistent with the observed h^τ if and only if

$$(5.3) \quad \beta W_\theta^{\pi_d}(\sigma_d | h^\tau) \geq (1 - \beta)(U_{\theta, \tau-1} - c_{\tau-1}) + \beta W_\theta^{\pi_d}(\sigma_d | h^\tau(r=1)).$$

To interpret condition 5.7, define $S_{\theta, \tau-1} := U_{\theta, \tau-1} - c_{\tau-1} > 0$ as the sacrificed spot utility for type θ (the implied opportunity cost payed by playing $r_{\tau-1} = 0$) and let

$$(5.4) \quad \text{NPV}_\theta^{\pi_d}(\sigma_d | h^\tau) := \frac{\beta}{1 - \beta} [W_\theta^{\pi_d}(\sigma_d | h^\tau) - W_\theta^{\pi_d}(\sigma_d | h^\tau(r=1))]$$

be the *net present continuation value* under pair (σ_d, π_d) of having played $r_{\tau-1} = 0$, with respect to the alternative continuation had she played $r_{\tau-1} = 1$ instead. This gives a very intuitive characterization of condition 5.3: we can rewrite it as

$$(5.5) \quad S_{\theta, \tau-1} \leq \mathbf{NPV}_{\theta}^{\pi_d}(\sigma_d \mid h^{\tau})$$

so it is optimal for type θ to “invest” an opportunity cost of $S_{\theta, \tau-1}$ utils today (by not following the spot optimal strategy) only if she expects a net present flow (relative to the alternative) that will compensate for the investment. We can further refine condition 5.3 by first showing that

$$(5.6) \quad W_{\theta}^{\pi_d}(\sigma_d \mid h^{\tau}(r=1)) \geq \underline{W}_{\theta}(\sigma_d \mid h^{\tau}(r=1)) \geq \underline{\mathbb{W}}_{\theta}$$

using equation 5.2 and the fact that the ex-ante worst rationalizable payoff is always a lower bound. This together with 5.3 implies a simple *necessary* condition for θ -rationalizability. If a pair (σ_d, π_d) θ -rationalizes (h^{τ}, c_{τ}) , then

$$(5.7) \quad W_{\theta}^{\pi_d}(\sigma_d \mid h^{\tau}) \geq \frac{1-\beta}{\beta} S_{\theta, \tau-1} + \underline{\mathbb{W}}_{\theta}$$

Condition 5.7 also holds for any other history (h^{τ}, c_{τ}) , generalizing the definition of sacrificed utility as

$$(5.8) \quad S_{\theta, \tau-1} := \max_{\tilde{r} \in \{0,1\}} (U_{\theta, \tau-1} - c_{\tau-1}) \tilde{r} - (U_{\theta, \tau-1} - c_{\tau-1}) r_{\tau-1}$$

Condition 5.7 places restrictions on θ -rationalizing pairs (σ_d, π_d) (and hence over policy functions $r(z)$) based only on the previous period outcome, disregarding the information in the observed history up to $\tau - 1$. In principle, there could be further restrictions that strong rationalizability could impose over the set of rationalizable policies p may expect after history (h^{τ}, c_{τ}) . A striking feature of strong rationalizability is that in fact, 5.7 is also *sufficient*: whether a policy function $r(\cdot)$ pair is strong rationalizable or not depends only on the observed past sacrificed utility. Proposition 5.2 states the core result of this paper.

Proposition 5.2. *Take a θ -rationalizable history $(h^{\tau}, c_{\tau}) \in \mathcal{H}_p$ and a policy function $r : Z \rightarrow \{0, 1\}$. Then $r(\cdot)$ is θ -rationalizable at (h^{τ}, c_{τ}) if and only if there exists a measurable function $w : Z \rightarrow [\underline{\mathbb{W}}_{\theta}, \overline{\mathbb{W}}_{\theta}]$ such that*

$$(5.9) \quad (1-\beta)(U_{\theta} - c_{\tau})r(z) + \beta w(z) \geq (1-\beta)(U_{\theta} - c_{\tau})\hat{r} + \beta \underline{\mathbb{W}}_{\theta} \text{ for all } \hat{r} \in \{0, 1\}, z \in Z$$

and

$$(5.10) \quad \mathbb{E}_z [(1-\beta)(U_{\theta} - c_{\tau})r(z) + \beta w(z)] \geq \frac{1-\beta}{\beta} S_{\theta, \tau-1} + \underline{\mathbb{W}}_{\theta}$$

Condition 5.9 is analogous to the Abreu et al. [1990] notion of *enforceability*. A policy $r(z)$ will be enforceable (or implementable) at some history only if we can find a continuation payoff function that enforces it on the set of strong rationalizable payoffs $[\underline{\mathbb{W}}_{\theta}, \overline{\mathbb{W}}_{\theta}]$. This argument employs the same tools and insights as in Abreu et al. [1990]. Condition 5.10

is the translation of condition 5.7 into this notation. Observe that it resembles a promise keeping constraint in a dynamic contracting problem: the expected value of following a rationalizable strategy σ_d at this history (given by the right hand side of 5.10) must be greater than the value implied by the implied opportunity cost paid in the previous period, which can be thought of as the utility “promised” by some rationalizable pair (σ_d, π_d) . Its proof resembles closely the well known “optimal punishment” argument in Abreu [1988]: any strong rationalizable outcome can be enforced by switching to the worst rationalizable payoff upon observing a deviation from the prescribed path of play. This means that without loss of generality, we can check whether a policy $r(\cdot)$ is θ -rationalizable if it is implementable whenever type θ thinks that if she deviated, she will have to play the optimal robust policy from then on.

A natural question to ask at this stage is whether the set of θ -rationalizing policies at history h^τ be empty. If this was the case, then p becomes certain, from τ on, that he cannot be facing a decision maker with type θ and strong rationalizable beliefs, since no rationalizable strategy could have generated the observed history h^τ . Proposition 5.2 allows us to answer this as well. First, let

$$S_\theta^{\max} := \frac{\beta}{1-\beta} (\overline{\mathbb{W}}_\theta - \underline{\mathbb{W}}_\theta)$$

be the maximum sacrifice level for type θ , that is consistent with common strong certainty of rationality.

Corollary 5.1. *History $h^\tau \in \mathcal{H}_d$ is θ -rationalizable if and only if $S_{\theta,k} \leq S_\theta^{\max}$ for all $k \leq \tau - 1$.*

To understand the necessary part of Corollary 5.1 note that for any rationalizable strategy σ_d , any belief π_d and any history h^τ we must have $W_\theta^{\pi_d}(\sigma_d | h^\tau) \leq \overline{\mathbb{W}}_\theta$. Then, a necessary condition for the set $\Sigma_\theta^*(h^\tau)$ to be non-empty is that

$$(5.11) \quad \overline{\mathbb{W}}_\theta \geq \frac{1-\beta}{\beta} S_{\theta,\tau-1} + \underline{\mathbb{W}}_\theta \iff S_{\theta,\tau-1} \leq S_\theta^{\max}.$$

Therefore, if $S_{\theta,\tau-1} > S_\theta^{\max}$ then $\Sigma_\theta^*(h^\tau) = \emptyset$; i.e., the observed history is not consistent with d having both strong common certainty of rationality and have payoff type θ .

To illustrate the powerful implications of Corollary 5.1, take a history h^τ such that (1) : $h^{\tau-1}$ is both *new* and *old*-rationalizable (2) : $r_{\tau-1} = 0$, (3) : $U_{old,\tau-1} - c_{\tau-1} > S_{old}^{\max}$, and (4) : $U_{new,\tau-1} - c_{\tau-1} \leq S_{new}^{\max}$. At such history, we must conclude that *for all (rationalizable) continuation stories, p should be strongly certain that he is facing $\theta = new$* . We will refer to these events as *strong separation*, or simply “separation” when it is obvious from the context.

5.5. Robust Commitment Cost Choice. In this subsection I will use the results of subsection 5.4 to characterize $c^*(h^\tau) := c^{\sigma_d^*}(h^\tau)$, the optimal robust commitment cost chosen at any *new*-rationalizable history h^τ . Based on the Markovian structure of the set of rationalizable policy functions, I will show that the robust commitment cost choice will be a

function only of the past sacrifice of the time inconsistent type: i.e. $c^*(h^\tau) = f(S_{old, \tau-1})$ for some (non increasing) function f . This results illustrates the reputation formation process in a robust optimal strategy: histories with higher implied sacrifices decrease the commitment cost necessary for d to induce p to trust for any rationalizable belief she might have. Therefore, $S_{\tau-1}$ is an endogenously derived reputation measure, since higher implied sacrifices at time $\tau - 1$ will result in smaller commitment costs in the next period.

Moreover, I will show that the relationship between past sacrifices and robust commitment cost is determined by the solution to a recursive program that resembles a dynamic contracting problem with promise keeping constraints, such as [Thomas and Worrall \[1988\]](#), [Kocherlakota \[1996\]](#) and [Alvarez and Jermann \[2000\]](#).

Given history $(h^\tau, c_\tau) \in \mathcal{H}_p$ and a system of beliefs π_p , p would decide to trust d at (h^τ, c_τ) if and only if $\mathbb{E}_{z, \hat{\sigma}_d}^{\pi_p} \{ \mathbb{E}_z [r^{\hat{\sigma}_p}(h^\tau, z) U_p] \mid h^\tau, c_\tau \} \geq \underline{u}_p$. Therefore, p would trust for d for all strongly rationalizable beliefs if and only if²⁴

$$(5.12) \quad \underline{V}(h^\tau, c_\tau) := \min_{\sigma_d \in \Sigma_d^*(h^\tau, c_\tau)} \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) U_p] \geq \underline{u}_p$$

which corresponds to p 's *worst rationalizable payoff* of trusting. It can be easily shown that optimal choice of c_τ would then be the minimum one that satisfies condition 5.12

$$(5.13) \quad c^*(h^\tau) := \min \left\{ c \in C : \underline{V}(h^\tau, c) \geq \underline{u}_p \right\}$$

Then, in order to characterize the optimal robust commitment cost choice, we need first to solve the problem in 5.12. The rest of this subsection will then be devoted to solve this programming problem, in order to solve for $c^*(h^\tau)$.

We can rewrite 5.12 as

$$(5.14) \quad \underline{V}(h^\tau, c) = \min_{\theta \in \{new, old\}} \left\{ \min_{\sigma_d \in \Sigma_\theta^*(h^\tau, c)} \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) U_p] \right\}$$

so we solve for the worst rationalizable payoff for p of facing each potential type θ , and then focus on the minimum of the two. While it could potentially happen that the minimum is achieved at a strategy in Σ_{new}^* , I will show that this is not the case: when solving 5.14 we will only need to consider the rationalizable strategies for $\theta = old$. Moreover, I show that when p is actually certain he is facing $\theta = new$, then the *unique* strong rationalizable outcome coincides with the static first best, with zero commitment cost and full flexibility.

Lemma 5.2. *For all new and old-rationalizable histories $h^\tau \in \mathcal{H}$ and $c \in C$, we have*

$$(5.15) \quad \underline{V}(h^\tau, c) = \min_{\sigma_d \in \Sigma_{old}^*(h^\tau, c)} \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) U_p]$$

²⁴Because of Fubini's theorem, we can write $\mathbb{E}_{z, \hat{\sigma}_d}^{\pi_p} [r^{\hat{\sigma}_p}(h^\tau, z) U_p \mid h^\tau, c_\tau] = \mathbb{E}_z \{ \mathbb{E}_{\hat{\sigma}_d}^{\pi_p} [r^{\hat{\sigma}_d}(h^\tau, z_\tau) \mid h^\tau, c_\tau] U_p \}$ which corresponds to the expected value over a mixed strategy $\hat{\sigma}_d$ with expected policy $\mathbb{E} [r^{\hat{\sigma}_d}(h^\tau, c_\tau)] = \mathbb{E}_{\hat{\sigma}_d} [r^{\hat{\sigma}_d}(h^\tau, z_\tau) \mid h^\tau, c_\tau]$. Then, the minimum rationalizable payoff of trusting is the one that assigns probability 1 to the worst rationalizable strategy for p on that history.

Moreover, for histories where there's common certainty that $\theta = \text{new}$, there is a unique rationalizable continuation strategy σ_{FB} , which corresponds to the spot First Best; i.e.

$$(5.16) \quad c^{\sigma_{FB}}(h^\tau) = 0 \text{ and } r^{\sigma_{FB}}(h^\tau, z_\tau) = \begin{cases} 0 & \text{if } U_p \leq 0 \\ 1 & \text{if } U_p > 0 \end{cases}$$

Then, to find $c^*(h^\tau)$ we need to solve the sequence problem 5.15. But then, Proposition 5.2 readily implies that the solution to 5.15 depends only on the past sacrifice for $\theta = \text{old}$, i.e.

$$\underline{V}(h^\tau, c_\tau) = \mathcal{V}(S_{old, \tau-1}, c_\tau)$$

where the function \mathcal{V} is defined by the following programming problem:

$$(5.17) \quad \mathcal{V}(s, c) := \min_{r(\cdot), w(\cdot)} \mathbb{E}[r(z) U_p]$$

subject to

$$(5.18) \quad (1 - \delta)(U_{old} - c_\tau) r(z) + \delta w(z) \geq (1 - \delta)(U_{old} - c_\tau) \hat{r} + \delta \underline{W}_{old} \text{ for all } \hat{r} \in \{0, 1\}, z \in Z$$

$$(5.19) \quad \mathbb{E}_z [(1 - \delta)(U_{old} - c) r(z) + \delta w(z)] \geq \frac{1 - \delta}{\delta} s + \underline{W}_{old}$$

$$(5.20) \quad w(z) \in [\underline{W}_{old}, \overline{W}_{old}] \text{ for all } z \in Z$$

For the rest of the paper, I will refer to $S_{old, \tau-1}$ simply as “observed sacrifice” and write it as $S_{\tau-1}$, since it will be the only relevant measure for the optimal robust policy. See that past history enters the problem only through the sacrifice measure s , which in turn only appears in the “promise keeping” constraint 5.19. Notice that the higher the observed sacrifice in the previous period, the smaller the feasible set for program 5.17, which then implies that $\mathcal{V}(s, c)$ is a weakly increasing function of s , which is intuitive: the more implied sacrifice exerted by $\theta = \text{old}$, the tighter the restriction that strong rationalizability imposes on continuation strategies, which in turns makes the worst rationalizable strategy be, perhaps, not as bad for p . When $s > S_{old}^{\max}$ the feasible set is empty, there's no pair $(r(\cdot), w(\cdot))$ that satisfy conditions 5.19 and 5.20 simultaneously.

In Appendix A I study in detail the solution $(\underline{r}(\cdot), \underline{w}(\cdot))$ to 5.17. In Proposition A.1 I show that there exist a threshold $\hat{S} \in (0, S_{old}^{\max})$ such that if $s \leq \hat{S}$, the promise keeping constraint 5.19 does not bind, and hence it is identical to the solution of 5.17 when $s = 0$, and $\mathcal{V}(s, c) = \mathcal{V}(0, c)$ for all $s \in [0, \hat{S}]$. When $s \in (\hat{S}, S_{old}^{\max})$ the promise keeping constraint starts binding, making $\mathcal{V}(s, c)$ strictly decreasing in this interval.

Once we solve for $\mathcal{V}(s, c)$, we can solve for the optimal robust commitment cost, which solves

$$(5.21) \quad c^*(h^\tau) = \min \left\{ c \in C : \mathcal{V}(S_{\tau-1}, c) \geq \underline{u}_p \right\} \iff$$

$$c^*(h^\tau) = \mathbf{c}(S_{\tau-1})$$

where the function $\mathbf{c}(s)$ is the one that solves the right hand side of equation 5.21. Given these findings, I summarize the implications for the optimal commitment cost choice in the next Proposition.

Proposition 5.3 (Robust Commitment Choice). *Let c_0^* be the optimal robust commitment cost at the initial history h^0 . Then, there exist a function $\mathbf{c} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ such that, for all old-rationalizable histories $h^\tau \in \mathcal{H}$:*

$$(5.22) \quad c^*(h^\tau) = \mathbf{c}(S_{\tau-1})$$

Moreover, the function $\mathbf{c}(\cdot)$ has the following properties:

- (1) $\mathbf{c}(s) = c_0^*$ for all $s \in [0, \hat{S}]$
- (2) If $\hat{S} < s' < s < S_{old}^{\max} \implies \mathbf{c}(s) < \mathbf{c}(s')$
- (3) $\mathbf{c}(s) = 0$ for all $s > S_{old}^{\max}$
- (4) $\mathcal{V}[s, \mathbf{c}(s)] = \underline{u}_p$ for all $s \in [0, S_{old}^{\max}]$




FIGURE 4. Robust commitment cost function

Proposition 5.3 finally makes the connection between the sacrifice measure $S_{\tau-1}$ and the concept of reputation formation. I provide here some intuition. For small observed sacrifices, p cannot discard the possibility that if $\theta = old$, she is expecting to behave according to the strategy that solves $\mathcal{V}(0, c)$ (i.e. with no sacrifice) from tomorrow on. Therefore, a robust choice for the commitment cost should prescribe exactly the same solution as in $\tau = 0$: the game basically “resets” and all reputation is lost at these histories.

For intermediate sacrifices, p still cannot rule out that $\theta = old$, but can nevertheless impose some restrictions on the set of rationalizable strategies, which are stronger the bigger the

sacrifice observed. When sacrifice is bigger than the maximum possible rationalizable net present value gain of any continuation value for $\theta = old$, the decision maker achieves strong separation, and hence she knows p is certain $\theta = new$ for all rationalizable continuation strategies, and therefore play the first best strategy with no commitment costs,

5.6. Optimal Robust Policy. In the previous section, I characterized a lower bound on the commitment cost that d must choose in all *old*-rationalizable histories he might be in, in order to guarantee p 's trust. Because d knows higher sacrifices lower future commitment costs, a reputational trade-off arises: acting in a way that generates big implied sacrifice would reduce commitment in the future, but it may be at the expense of current opportunity costs. In this subsection we will characterize the optimal robust policy $r^*(z)$ for the time consistent type $\theta = new$.

Formally, let s be the current observed sacrifice in a history h^τ that is both *old* and *new*-rationalizable. Proposition 5.3 implies that s is the only relevant state variable. Also, denote $r^*(z)$ to be the optimal robust policy when past observed sacrifice is s and the realized shock is z , and let $s'(z)$ be the sacrifice that will be observed in the next period under the optimal robust strategy σ_d^* . Using the fact that $c^*(h^\tau) = \mathbf{c}(s)$, the law of motion for the observed sacrifice for $\theta = old$ at *old*-rationalizable histories is

$$(5.23) \quad s'(z) = \max_{\tilde{r} \in \{0,1\}} [U_{old} - \mathbf{c}(s)] \tilde{r} - [U_{old} - \mathbf{c}(s)] r(z)$$

When $s'(z) > S_{old}^{\max}$, because of Corollary 5.1 we know that in the next period $\theta = new$ would have achieved strong separation, and hence the expected payoff from next period on is the first best payoff, $\mathbb{E}_z \{\max(0, U_p)\}$. If $s'(z) \leq S_{old}^{\max}$, the Markovian nature of the commitment cost process implies that at $\tau + 1$ the optimal policy problem will be identical to the problem at τ , with the difference that now $s = s'(z)$.

Finally, let $\mathcal{W}(s)$ be the value function of the optimal robust implementing strategy for *old*-rationalizable histories h^τ with past period sacrifice $S_{\tau-1} = s$. We can solve for it recursively, by solving the following Bellman equation:

$$(5.24) \quad \mathcal{W}(s) = \max_{r(\cdot), s'(\cdot)} (1 - \beta) \mathbb{E}_z \{ [U_p - \mathbf{c}(s)] r(z) \} + \beta \mathbb{E}_z \{ \mathcal{W}[s'(z)] \} \quad \text{for all } s \leq S_{old}^{\max}$$

subject to 5.23, and $\mathcal{W}(s) = \mathbb{E}_z \{\max(0, U_p)\}$ for all $s > S_{old}^{\max}$. See that this is a very standard Bellman equation with a one-dimensional state variable, and the standard fixed point arguments of Stokey et al. [1989] would work to characterize the value function $\mathcal{W}(s)$. Also note that because of Lemma 5.1, the function $\mathcal{W}(s)$ gives the worst strong rationalizable payoff; i.e. for all *old*-rationalizable histories h^τ we must have $\underline{W}_{new}(h^\tau) = \mathcal{W}(S_{\tau-1})$. Therefore, the ex-ante payoff of the optimal strong rationalizable strategy is

$$(5.25) \quad \underline{W}_{new} := \mathcal{W}(0)$$

In the remainder of this section, I solve for the optimal robust policy $r^*(z)$ and the law of motion for the sacrifice process $s'(z)$. Figure 6 previews the shape of the optimal policy

$r^*(z) = r^*(U_p, U_{old})$ over the set of states $Z = [\underline{U}, \bar{U}]^2 \subset \mathbb{R}^2$. Regions where $r^*(U_p, U_{old}) = 1$ (i.e. d takes the emergency action) are depicted in red, and $r^*(U_p, U_{old}) = 0$ in green. In the bottom we include the spot optimum strategy for $\theta = new$ and for agent p .²⁵ In the right margin, we draw the analogous scale for $\theta = old$.




FIGURE 5. Static Best policy $r^{spot}(z)$ for $\theta = new$

²⁵The spot optimal policy for p is defined as $r_p^{spot}(z) := 1 \iff U_p \geq 0$. For and for type θ we have $r_\theta^{spot}(z) := 1 \iff U_\theta \geq c$




FIGURE 6. Optimal robust policy $r^*(z)$ for $\theta = new$

By comparing the optimal robust policy of Figure 6 to the spot optimal strategy for $\theta = new$ in Figure 5, we see how the optimal robust policy is distorted from the spot optimum towards actions that are spot inefficient for $\theta = old$. For example, consider the region where $U_{old} > c$ so that $r_{old}^{spot}(z) = 1$. Figure 6 shows then that the optimal policy over this region prescribes $r^*(U_p, U_{old}) = 0$ on a strictly larger set of states than $r_{new}^{spot}(z)$. The intuition for this phenomenon is simple: if d plays $r_\tau = 0$, then p will observe an opportunity cost paid by $\theta = old$ of $S_\tau = U_{old} - c > 0$ utils. This will result in a smaller commitment cost at $\tau + 1$ than the one implied by playing $r = 1$ and resetting to the time $\tau = 0$ robust policy from tomorrow on. When the relative reduction in future commitment costs is big enough (i.e. when observed implied sacrifice for $\theta = old$) then the optimal strategy will be to choose $r = 0$. In the rest of this section I will formally characterize both the robust policy $r^*(z)$ and the next period sacrifice $s'(z)$, which governs the reputation formation process.

First, take the region $R_1 = \{z \in Z : U_{old} > c + S_{old}^{max}\}$, which corresponds to the upper-most horizontal strip of Figure 6. For any z in this region, the unique rationalizable policy for $\theta = old$ is to play $r(z) = 1$. This is because if she played $r = 0$, the implied sacrifice $S = U_{old} - c$ would be strictly greater than any potential net present value gain from switching to the best rationalizable payoff (given by S_{old}^{max}). Therefore, if $\theta = new$ chooses $r^*(z) = 0$ she would strongly separate from tomorrow on, achieving the first best payoff $\mathbb{E}_z[\max(0, U_p)]$. If, however, she chose $r^*(z) = 1$ then $s'(z) = 0$ and next period the

commitment cost gets reset to c_0^* , getting a continuation value of \underline{W}_{new} . Therefore, $r^*(z) = 0$ over region R_1 if and only if

$$\beta \mathbb{E}_z [\max(0, U_p)] \geq (1 - \beta)(U_p - c) + \beta \underline{W}_{new} \iff$$

$$(5.26) \quad U_p \leq c + S_{new}^{\max}$$

If $U_p \leq c$ then by playing $r^*(z) = 0$ type $\theta = new$ would maximize both her spot and her continuation values, achieving strong separation from $\tau + 1$ on. Even when $U_p > c$, $\theta = new$ could still find it optimal to sacrifice spot gains for the strong separation that would be achieved in the next period. Therefore, when the time inconsistent type has a unique rationalizable strategy, the good type would optimally *invest* in reputation, sacrificing present utility to achieve strong separation in the next period.

Second, take region $R_2 = \{z \in Z : U_{old} \in (c + \hat{S}, c + S_{old}^{\max})\}$. In this region, $\theta = old$ preferred strategy is still $r = 1$, but now $r = 0$ is also *old*-rationalizable. By playing $r = 0$, $\theta = new$ cannot achieve separation in the next period, but she still can decrease the commitment cost in the next period to $c_{\tau+1} = c(U_{old} - c)$. Therefore, the same analysis from region R_1 applies here, with the only difference that now the continuation value will be $\mathcal{W}(U_{old} - c) < \mathbb{E}_z [\max(0, U_p)]$. Then, $r^*(z) = 0$ over region R_2 if and only if

$$U_p - c \leq \frac{\beta}{1 - \beta} [\mathcal{W}(U_{old} - c) - \underline{W}_{new}] := \phi(U_{old} - c)$$

where ϕ is an increasing function of the implied sacrifice $S = U_{old} - c$ of playing $r = 0$ for the time inconsistent type. As before, whenever $U_p < c$ then $r = 0$ will be optimal. When $U_p > c$ her decision will depend on two variables: the spot disutility by not choosing $r = 1$ ($U_p - c$) and the reputation value gained by choosing $r = 0$, $\phi(U_{old} - c)$. States with very high disutility for $r = 0$ would only prescribe it as an optimal policy for states with high potential sacrifice.

Finally, study region $R_3 = \{z \in Z : U_{old} \in (c, c + \hat{S})\}$. See that regardless of the the action, the sacrifice potential $U_{old} - c$ is too small to make the commitment cost in the next period to be smaller than its maximum possible level, c_0^* . Therefore, regardless of the policy chosen, in the next period reputation will be lost. Therefore the optimal robust policy is just the spot optimal policy: $r^*(z) = 1 \iff U_p > c$.

Is easy to see that when $U_{old} < c$, then the optimal robust policy analysis will be identical, but with the role of each policy reverted, since it will be now when $r = 1$ that sacrifice may be signaled. Moreover, because of the symmetry of the sacrifice measure around c , the optimal policy will be “anti-symmetric”: identical to the policy in regions R_1 to R_3 , but switching the policies.

In Figure 7 we illustrate the map $s'(z)$. The dark blue areas correspond to $s'(z) = 0$ (i.e. all reputation is lost in the next period), and white regions are those in which $\theta = new$ achieves full separation from tomorrow on. In the shaded areas, lighter color illustrate higher




FIGURE 7. Optimal Robust Sacrifice

sacrifice levels, and hence smaller commitment costs in the next period (but not zero, as in full separation).

5.7. Discussion. The reason why the robust policy problem ends up being quite tractable is exactly because of the robustness condition: when having to make sure that p trusts in all histories and for all rationalizable beliefs, the *worst types* (in the sense of beliefs) that p might be facing when deciding whether to trust or not, may correspond to very different beliefs about d 's behavior. Because of this disconnect, we are able to separate the problems of commitment cost choice and of the optimal robust policy rule $r^*(z)$. When more restrictions are imposed (for example, a belief set $\mathcal{B} = \{\pi_{new}, \pi_{old}\}$, as in any Bayesian equilibrium), this separation will be broken.

In terms of the optimal robust policy, note that there exist regions where p and both types of decision maker would unanimously prefer certain strategy to be played, but because of reputation building motives $\theta = new$ would still want to do exactly the opposite of the Pareto optimal decision. For example, when $U_{old} > c + S_{old}^{max}$ and $U_p \in (c, c + S_{new}^{max})$ all agents prefer $r = 1$, but the optimal policy prescribes the normal action $r = 0$.

In the context of the capital taxation model, this would correspond with states where the marginal utility of the public good is both sufficiently high for both workers and capitalists, so that both households would agree that the ex-post optimal strategy would be to expropriate. By the lens of our model, we can summarize the policy maker's decision by the following

argument: “Even as a pro-capitalist government, I am tempted to expropriate capitalists. However, the incentives for a benevolent, time inconsistent government to expropriate out would be much higher than mine. Therefore, by not expropriating, I can show that I am in fact, not the time inconsistent type”. Notice also that regardless of the beliefs that d may have, any strong rationalizable strategy of d should also achieve separation at $\tau + 1$, if she decides to play $r = 0$. This then gives a *robust prediction* about d 's behavior, as long it is consistent with common strong certainty of rationality.

A perhaps troubling feature of the robust policy is the impermanence of reputation gain: only the sacrifice of the previous period matters, but past sacrifices do not provide relevant information for reputation building. In the next section I find conditions on the set of beliefs \mathcal{B} so that the optimal robust implementation exhibits permanent reputation gains, and hence all past sacrifices give some information about the continuation strategies that the decision maker may be planning to follow.

6. Basic Properties of Strong Rationalizable Implementation

In this section I will study some features of the optimal robust policy. I will first show how present potential sacrifices may affect the distribution of future sacrifices, creating a momentum for reputation formation. I will also show that the observed sacrifice process achieves almost surely the bound S_{old}^{\max} . Hence, by playing the robust policy d will eventually convince p that $\theta = new$, with probability one. Moreover, the speed of convergence to the absorbing complete information stage (where p is certain that $\theta = new$) is exponential, which is also the convergence rate of the best equilibrium of the game, so that by using the robust policy we still get the same rate of convergence (although at a strictly lower speed).

I also study the asymptotic behavior of the robust policy as both the time consistent and the time inconsistent type become more patient, and show that as the discount rate approaches unity, the worst rationalizable payoff \underline{W}_{new} converges to the first best payoff, and hence the value of the robust policy converges uniformly to the first best payoff (e.g. for all histories). This further implies that the expected value of any strong rationalizable strategy that $\theta = new$ may follow converges to the first best payoff as well, an analog result to [Fudenberg and Levine \[1989\]](#)

6.1. Dynamics of the optimal robust policy. We saw in the previous section that only immediate past behavior builds reputation, and past histories are irrelevant. However, it seems intuitive that there should be some momentum in reputation gaining. The basic idea is that gaining reputation at time τ will lower the commitment cost in the next period. The lowering of commitment cost will allow the reformed type to exploit the difference in ex-post payoffs between both types, which is the source of the difference between her and the time inconsistent decision maker, and therefore making that the commitment cost in $\tau + 1$ should also go down even more, in a probabilistic way. However, the degree of generality I have been using so far does not allow for an easy characterization of the stochastic process followed

by the commitment cost $\mathbf{c}(S_{\tau-1})$. Therefore, in this section I will show a somewhat weak momentum result, using a plausible assumption on the primitives of the model

Assumption 1 (Status quo). *In the static version of the game, we have*

$$(6.1) \quad \Pr \left(\underset{\tilde{r} \in \{0,1\}}{\operatorname{argmax}} (U_{old} - \bar{c}) \tilde{r} = 0 \right) > \Pr \left(\underset{\tilde{r} \in \{0,1\}}{\operatorname{argmax}} (U_{old} - \bar{c}) \tilde{r} = 1 \right)$$

i.e. the optimal static decision rule for the bad type induces the green button to be played more often than the red button.

This assumption further reinforces our interpretation of the green button strategy ($r = 0$) as the status quo: it is the strategy that both the good guy, and a trustworthy bad type would play most often.

As we saw in the previous section, the main driver of reputation building is the sacrifice potential $|U_{old,\tau} - c_\tau|$, a exogenous variable for d given the commitment cost chosen. When the sacrifice potential is high is when d may decide to invest in reputation building, and moreover, conditional on observing a sacrifice, higher sacrifice potential imply lower commitment cost in the next period. While I cannot provide a characterization of the commitment cost process, I can show that the expected value of the sacrifice potential goes up when the commitment cost decreases.

Proposition 6.1. *If Assumption 1 holds, then*

$$(6.2) \quad s \geq s' \text{ implies } \mathbb{E}_z [|U_{old} - \mathbf{c}(s)|] \geq \mathbb{E}_z [|U_{old} - \mathbf{c}(s')|]$$

so that after higher observed sacrifices, we expect higher potential sacrifices. If $s > s' > \hat{S}$, then the inequality in 6.2 is strict.

To show our second result on the dynamics of the optimal robust policy, I will need this very important lemma.

Lemma 6.1. *For all old- rationalizable histories $h^\tau \in \mathcal{H}(\sigma_d^*)$, we have that:*

$$(6.3) \quad \Pr(U_p > c^*(h^\tau) - S_{new}^{\max}, U_{old} < c^*(h^\tau) - S_{old}^{\max}) > 0$$

$$(6.4) \quad \Pr(U_p < c^*(h^\tau) + S_{new}^{\max}, U_{old} > c^*(h^\tau) + S_{old}^{\max}) > 0$$

$$(6.5) \quad \Pr(S_\tau \geq S_{old}^{\max}) > \underline{q} > 0 \text{ for all } h^{\tau+1} \in \mathcal{H}(\sigma_d^*)$$

where

$$(6.6) \quad \underline{q} := \Pr(U_p > c_0^* - S_{new}^{\max}, U_{old} < \bar{c} - S_{new}^{\max}) + \Pr(U_p < \bar{c} + S_{new}^{\max}, U_{old} > c_0^* + S_{old}^{\max})$$

Proof. See Appendix C □

Lemma 6.1 is important on it's own, and states first that all the regions considered in the optimal robust policy of Section (5.6) have positive probability, and hence separation will

surely occur (something we didn't have to check before). Moreover, I get a uniform non-zero lower bound on the probability of separating at any history, that can be easily calculated. With it, I can show its speed of convergence to strong separation

Proposition 6.2. *For all $\tau \in \mathbb{N}$*

$$(6.7) \quad \Pr(\text{Separating before } \tau \text{ periods}) > 1 - (1 - \underline{q})^\tau$$

Proof. In every trial history (*new* and *old*-rationalizable) there is at least probability \underline{q} of separating. Since shocks are i.i.d this implies that

$$\Pr(S_{old,k} < S_{old}^{\max} \text{ for all } k \leq \tau) < (1 - \underline{q})^\tau$$

and hence $\Pr(\text{Separating before } \tau \text{ periods}) = 1 - \Pr(S_{old,k} < S_{old}^{\max} \text{ for all } k \leq \tau) = 1 - (1 - \underline{q})^\tau$ \square

This proposition states one of the most important results: the probability of reaching separation from the time inconsistent type is exponentially decreasing in τ . A perhaps even more important corollary is that in fact, for any belief restriction \mathcal{B}_p that is consistent with strong common certainty of rationality, (i.e. $\mathcal{B}_p \subseteq \mathcal{B}_p^s$) we will also achieve separation in exponential time, the probability of separation can only be higher for any smaller belief sets. In the next section we will explore some “reasonable” restrictions we could impose, and see how the solution would be improved.

The second important corollary is that eventually $S_{old,\tau} > S_{old}^{\max}$ almost surely (and states there after separation), so that d will surely separate eventually from the time inconsistent type.

6.2. First Best Approximation by patient players. In this subsection the assumption $\beta_{old} = \beta_{new} = \beta$ will be significant, since we will be increasing both discount rates. I will show that as both types become more patient, the payoff of the robust policy for $\theta = new$ converges to the payoff of the stage game after separation. Now, the probability of separation for history h^τ will be denoted as $q(h^\tau, \beta)$. I first show that these probabilities are uniformly bound away from zero for all $\delta \in (0, 1)$ and all rationalizable histories

Lemma 6.2. *Let $q(h^\tau, \beta)$ be the ex-ante probability of separation under the optimal robust strategy σ_d^* . Then, there exist $\hat{q} > 0$ such that $q(h^\tau, \beta) > \hat{q}$ for all $h^\tau \in \mathcal{H}(\sigma_d^*), \beta \in (0, 1)$*

Proof. See Appendix C \square

The previous lemma shows the existence of a number $\hat{q} > 0$ such that no matter the discount rate δ , the probability of reaching separation in any *new* and *old*-rationalizable histories is greater than \hat{q} . Since shocks are i.i.d, even if history may exhibit time dependence, we can bound the expected time of separation by a geometric random variable with success probability \hat{q} . Since once we reach separation, the unique rationalizable outcome is the First Best (i.e. no commitment, spot optimum policy for $\theta = new$) and the speed of

convergence is exponential for this random variable, then for a sufficiently patient decision maker d , the expected payoff of the robust policy will be very close to the first best (i.e. the expected time for separation is very small in utility terms). This is what I show in the following proposition

Proposition 6.3. *Let $\mathbb{E}_z \{\max(0, U_p)\}$ be the first best payoff, corresponding to the case where p is certain that $\theta = \text{new}$, and let $\underline{\mathbb{W}}_{\text{new}}(\beta)$ be the ex-ante expected payoff for the optimal robust policy. Then*

$$(6.8) \quad \underline{\mathbb{W}}_{\text{new}}(\beta) \rightarrow \mathbb{E}_z \{\max(0, U_p)\} \text{ as } \beta \rightarrow 1$$

Proof. For the robust policy, we always can bound it as

$$\underline{\mathbb{W}}_{\text{new}}(\beta) \geq \mathbb{E}_\tau \{\beta^\tau \mathbb{E}_z [\max(0, U_p)]\}$$

where $\tau \sim \text{Geom}(\hat{q})$. This is true since $\mathbb{E}_z \{r(h^\tau, z)(U_p(z) - c(h^\tau))\} \geq 0$ by Lemma C.2. Since we always have that the contemporaneous utility greater than zero and separation is achieved with a probability greater than \hat{q} in any period, we have that this is a lower bound for the robust policy payoff. See also that

$$\mathbb{E}_\tau(\beta^\tau) = \sum_{\tau=1}^{\infty} \beta^\tau (1 - \hat{q})^{\tau-1} \hat{q} = \frac{\hat{q}}{1 - \hat{q}} \frac{\beta(1 - \hat{q})}{1 - \beta(1 - \hat{q})} = \frac{\beta \hat{q}}{1 - \beta(1 - \hat{q})}$$

Therefore

$$\begin{aligned} \mathbb{E}[\max(0, U_p)] - \underline{\mathbb{W}}_{\text{new}}(\beta) &\leq \mathbb{E}_z[\max(0, U_p)](1 - \mathbb{E}_\tau(\beta^\tau)) = \\ &= \mathbb{E}_z[\max(0, U_p)] \left(1 - \frac{\beta \hat{q}}{1 - \beta(1 - \hat{q})}\right) = \mathbb{E}_z[\max(0, U_p)] \left(\frac{1 - \beta}{1 - \beta(1 - \hat{q})}\right) \rightarrow 0 \end{aligned}$$

as $\beta \rightarrow 1$. □

If d is patient enough, because of discounting and the exponential speed of convergence to separation, the payoff of the robust policy will be very close to the first best payoff. Therefore, if events of distrust are sufficiently bad (as in our infinite cost interpretation of p 's distrust), the risk of using a weaker solution concept may be substantial, if we are not quite sure about the restrictions implied by it, while the potential increase in payoffs would be almost irrelevant if d is patient enough.

6.3. Restrictions on Beliefs. In certain situations, the policy maker may have more information about the public p 's beliefs. I describe how this may be incorporated into the problem. Notice that without any restrictions, there is an aspect of *paranoia* in the problem. Paranoia means that a past sacrifice of the government says nothing about the future sacrifices. Moreover, gains of reputation are not permanent. A natural question to ask is what restrictions on beliefs make reputation gains permanent? That is, when is $c_\tau \geq c_{\tau+k}$ for every k ?

Formally, say a strategy $\sigma_d \in \Sigma_d$ exhibits *permanent reputation gains* if and only if $c^{\sigma_d}(h^{\tau+1}) \leq c^{\sigma_d}(h^\tau)$ for all histories $h^\tau, h^{\tau+1} \in \mathcal{H}(\sigma_d)$. We already know that the optimal robust strategy does not satisfy this property. The goal then is to find what type of restrictions on beliefs should we impose to get a robust implementing strategy that exhibits permanent reputation gains. Say a belief system π_d is $\hat{\sigma}_d$ -*nondecreasing* if and only if, for all histories $h^\tau, h^{\tau+1} \in \mathcal{H}(\hat{\sigma}_d)$,

$$\mathbf{NPV}_{old}^{\pi_d}(\hat{\sigma}_d | h^{\tau+1}) \geq \mathbf{NPV}_{old}^{\pi_d}(\hat{\sigma}_d | h^\tau).$$

This means that under belief π , $\theta = old$ cannot get less than what she expected in the previous period, by playing strategy $\hat{\sigma}_d$. Denote also $\mathbf{ND}(\hat{\sigma}_d) \subset \Sigma_d$ be the set of best responses to $\hat{\sigma}_d$ -nondecreasing beliefs.

Proposition 6.4. *Take a belief restriction set $\mathcal{B}_p \subseteq \mathcal{B}_p^s$ and σ_d^* the optimal robust policy in \mathcal{B}_p . Then*

$$\sigma_d^* \text{ exhibits permanent reputation gains} \iff \mathcal{B}_p \subseteq \mathbf{S}_p(\{\theta = old\} \times \mathbf{ND}(\sigma_d^*)).$$

That is, p is strongly certain that $\theta = old$ has σ_d^ -nondecreasing beliefs.*

Proof. I show necessity. Take a *old*-rationalizable pair (σ_d, π_d) and a history $h^\tau \in \mathcal{H}(\sigma_d^*)$ such that $c_k = c_0^*$ for all $k < \tau - 1$ and $c_\tau < c_0^*$. This is a history where there has been only one gain in reputation so far, and which has been realized only in the last period. The fact that the commitment cost decreased has, as a necessary condition, that the observed sacrifice should be higher than certain threshold level \hat{S} ; i.e.

$$(6.9) \quad \mathbf{NPV}_{old}^{\pi_d}(\sigma_d | h^\tau) \geq S_\tau$$

Also, because $h^\tau \in \mathcal{H}(\sigma_d^*)$ condition (6.9) also holds for σ_d^* . Then, the only way the commitment cost could go up in some other history $h^{\tau+s} \in \mathcal{H}(\sigma_d^*)$, is that $\mathbf{NPV}_{old}^{\pi_d}(\sigma_d | h^{\tau+s}) < S_\tau$. But since π_d is σ_d^* -nondecreasing, we have

$$\mathbf{NPV}_{old}^{\pi_d}(\sigma_d^* | h^{\tau+s}) \geq \mathbf{NPV}_{old}^{\pi_d}(\sigma_d^* | h^\tau) \geq S_\tau > \mathbf{NPV}_{old}^{\pi_d}(\sigma_d | h^{\tau+s})$$

implying that σ_d is dominated by σ_d^* at $h^{\tau+s}$. Then, the fact that the net present value is always increasing will imply that the resulting commitment cost will be always non-increasing. \square

Intuitively, to get a monotonic robust policy, the assumption we need to make on p 's beliefs about the time inconsistent type are the following: if $\theta = old$ and we have observed that $S_{\tau-1} = S$, then fact that she was willing to sacrifice S utils will “stick”, and p will always think that $\theta = old$ will not settle for any smaller net present value. This feature of beliefs are actually pretty common in dynamic adverse selection and signaling problems.

Note that the important restriction is about p 's higher order beliefs: they are not about what p thinks d will do, but rather what p believes d believes about the continuation game. While working directly over system of beliefs can always be implemented, assumptions about

higher order beliefs are not very transparent in this framework. In Appendix (A) I explore a different approach, by modeling restrictions on beliefs as type spaces, which allow the modeling of assumptions on higher order beliefs more tractable.

7. Extensions And Further Research

I now address several extensions of the model and strategies for future research. First, a natural alternative is to take a *legislative approach*, as in [Athey et al. \[2005\]](#) and [Persson and Tabellini \[1993\]](#). The policy maker may have delegated the commitment choice to the public. The idea here is that if one delegates the commitment cost to the public then certainly one will have robust implementation. The relevant source of uncertainty in the problem is that the public mistrusts the government. The intuition comes from contract theory: we should give control rights precisely to the party who has the first-order inability to trust. However, this will come at a cost in terms of efficiency. Specifically, the public would always put a higher commitment cost, to make the optimal policy for $\theta = old$ not drive him to indifference between trusting or not. As such, the public would increase commitment costs relative to the levels chosen by the new regime government. Therefore, it is easy to show that, if the government has the same robustness concerns, then the executive approach is superior for her in terms of welfare, given their information.

Second, we may consider robustness to not just a single time inconsistent “old type” but a multitude of time inconsistent types. Is straightforward to see that Proposition (5.2) would still be true for any type space Θ_d and hence the characterization of the robust commitment cost would be identical to condition (5.14), only that now $\underline{V}(h^\tau, c_\tau)$ is defined as

$$(7.1) \quad \underline{V}(h^\tau, c_\tau) = \min_{\theta \in \Theta_d} \mathcal{V}_\theta(S_{\theta, \tau-1}, c_\tau)$$

where the function $\mathcal{V}_\theta(c, s)$ is the minimum problem in (5.17) for a given payoff type θ . Therefore, this will be equivalent to our dynamic contracting characterization of the problem above but with multiple types. In the case of a finite type set $\Theta_d = \{\theta_1, \theta_2, \dots, \theta_k\}$ where we now have the vector of observed sacrifices $S_{\tau-1} = (S_{\theta_1, \tau-1}, S_{\theta_2, \tau-1}, \dots, S_{\theta_k, \tau-1})$ as the state variables for the implied promise keeping constraints. The solution would exhibit separation from certain types across time, and if the other types satisfy the same assumptions made about $\theta = old$, then it will also eventually convince p about her being the time consistent type.

A third extension is to an environment in which d has an imperfect signal about p 's perceived incentives of the time inconsistent type. If signals are bounded and its support may be affected by some signal that d observes, then robust policy would be qualitatively identical.

Finally, looking forward, I would like to extend our analysis to situations in which there are a continuum of strategies and policies available. This will allow researchers to apply this robust modeling approach to various macroeconomic applications of interest, as the inflation

setting model of subsection (2.2). It is easy to see how Proposition (5.2) would remain valid on more general models, so that the Markovian nature of reputation formation would be a very general characteristic of this type of robustness.

8. Conclusions

I have studied the problem of a government with low credibility. A government faces ex-post time inconsistent incentives due to lack of commitment, such as an incentive to tax capital or an incentive to allow for undesirably high levels of inflation. The government undergoes a reform in order to remove these incentives; however, the reform is successful only if the public actually believes that the government has truly reformed its ways. As such, the crux of the problem relies on the government building reputation in the eyes of the public.

After arguing that the typical approach to this problem relies on equilibrium concepts, which are highly sensitive to small perturbations about the public's beliefs, I turned to studying the problem through the lens of optimal robust policy that will implement the public's trust over any rationalizable belief that any party can hold. Focusing on robustness to all extensive-form rationalizable beliefs, I characterize the solution as well as the speed of reputation acquisition.

This is a particularly desirable property from the point of view of macroeconomic mechanism design. Equilibrium type solution concepts rely on every party knowing every higher order belief of every other party involved in the interaction. This is an extremely high dimensional object and in all likelihood it may be very difficult to believe that such an assumption really holds in settings in which one agent is trying to convince the other agent that he is not adversarial. Furthermore, equilibrium concepts rely on high dimensional belief functions off the path of play – that is, nodes or histories that may never be reached. This sort of sensitivity is problematic when advising a policy maker as small deviations in how a party truly conjectures some off the path of play belief may severely affect the policy maker's ability to obtain trust. This sort of analysis, studying optimal robust policy, can be a very powerful tool within macroeconomic policy making.

APPENDIX A. TYPE SPACES

As mentioned before, the decision maker d acting as a “policy maker”, may have some information about people's beliefs about what strategy may be played by d , as well as as beliefs d may hold, which may involve assumptions about higher order beliefs (what agent i believes about j , what i believes about j 's beliefs about her beliefs, and so on. As reviewed in Bergemann and Morris [2005] the literature on epistemic game theory distinguishes between two approaches: an explicit and an implicit approach. In the *explicit approach*, beliefs are modeled as subjective probability measures over (1) the other players strategy, and (2) probability measures over the beliefs of the other agents (which are themselves probability

measures over the player’s own strategy), etc. In the *implicit approach*, beliefs are formed over other player strategies and “types”, where a type is directly mapped to a belief over strategies and types played by the other agent. We will argue that for most applications, the implicit approach will be a more tractable modeling assumption.

Formally, the explicit approach consists on modeling people’s beliefs as a hierarchy of beliefs over the types and strategies that d could play, the beliefs that d may have about p ’s beliefs, and so on, ad infinitum. This means that a *belief hierarchy* is a sequence of measures $(\pi_0, \pi_1, \pi_2, \dots)$ where π_0 is a CPS over $\{old, new\} \times \Sigma_d$ (the payoff types of d and with the strategies that she can choose from), π_1 is a measure over the space of such probability systems (i.e. $\pi_1 \in \Delta^{\mathcal{H}}(\Delta^{\mathcal{H}}(\{old, new\} \times \Sigma_d))$), $\pi_2 \in \Delta^{\mathcal{H}}(\Delta^{\mathcal{H}}(\Delta^{\mathcal{H}}(\{old, new\} \times \Sigma_d)))$ and so on. The set of all possible coherent hierarchies of beliefs²⁶ is denoted by H_p^* , and can be shown to have desirable topological properties²⁷ Therefore, information about people’s beliefs can be then represented as restrictions over the set of all coherent hierarchies; i.e. we can represent our information on beliefs as a subset $I \subset H_p^*$

While this approach has the advantage of being explicit about higher order beliefs, it is cumbersome to work with. Note that we have to consider beliefs over exponentially larger spaces. As an alternative, Harsanyi proposed an implicit approach (Harsanyi [1967/68]). He suggested that one could bundle all relevant information about beliefs and payoff parameters into different “(epistemic) types” of agents, in the same way that we think about payoff types. As such, we can model the system as a Bayesian game, with a larger type space. We formalize this idea for our context:

Definition A.1 (Type Space). A type space \mathcal{T} is a 5-tuple $(T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$ where T_i are sets of types for each agent, $\hat{\theta} : T_d \rightarrow \{old, new\}$ is a function that assigns to each type $t_d \in T_d$ to a payoff type $\hat{\theta}(t_d)$, and $\hat{\pi}_i : T_i \rightarrow \Delta^{\mathcal{H}}(T_{-i} \times \Sigma_{-i})$ assigns a CPS $\hat{\pi}_i(t_i)$ over strategy and type pairs (t_{-i}, σ_{-i})

I intersect the approaches of Bergemann and Morris [2005] and Battigalli and Siniscalchi [1999], a type encodes it’s payoff type (since d knows its type) together with his beliefs about the other agent. The set of “states of the world” Ω_{-i} that agent i form beliefs over is then the set of pairs $\omega_{-i} = (t_{-i}, \sigma_{-i})$ of strategies and types of the other agent. Unlike Bergemann and Morris [2005] however, agents hold beliefs over other player’s types and strategies as well, since in a dynamic environment these are not perfectly observed.²⁸ This method of representing restrictions on beliefs has the advantage of being more compact (since only first order beliefs have to be specified) and also being a natural generalization of Bayesian games. For example, suppose the information we have about p is that p thinks that $\theta = old$ could have two possible beliefs about the future play of the game: optimistic (expecting her

²⁶See Definition B.1 in Appendix B

²⁷See Appendix B for a formal treatment of the topological properties of the set of all coherent hierarchies.

²⁸This is the *interactive epistemic characterization* of Ben-Porath [1997] and Battigalli and Siniscalchi [1999]. This definition also corresponds to Penta [2012] notion of “conjectures”

best equilibrium to be played) or pessimistic (expecting her worst equilibrium). If this was the case, we could model this situation by simply augmenting the type space by creating two copies of the old type: an “optimistic old type” t_O and a “pessimistic old type” t_W with the same payoff parameter ($\hat{\theta}(t_O) = \hat{\theta}(t_W) = old$) but with different beliefs $\hat{\pi}_d(t_O) \neq \hat{\pi}_d(t_W)$. The type sets assigned to d would then be $T_d = \{new, t_O, t_P\}$.

In compact static games, Mertens and Zamir [1985] and Heifetz [1993] showed that these two approaches are equivalent: for any subset of possible hierarchies of beliefs $I \subset H^*$ there exist a type space \mathcal{T} that generates the exact same belief hierarchies²⁹ and vice versa.

In particular, if no restrictions are imposed on hierarchies (i.e. $I = H^*$) there exist a *universal type space* \mathcal{T}^* which is capable of generating all possible hierarchies of beliefs. In another paper, (Xandri [2012]) I extend Battigalli and Siniscalchi [1999] to (non-metrizable)³⁰ topological spaces, to show that this is also true in a relevant class of extensive form games. In Appendix B I provide an application to our particular setting, and also a formal description of how we can make the mapping between these two approaches.

Because of this equivalence result, I will use the implicit approach throughout this paper, when modeling d 's information and assumptions about p 's beliefs. I will further consider only *compact type spaces*, where the sets T_i are compact and Hausdorff topological spaces (with some topology) and the belief functions $\pi_i(t_i)$ are continuous in the weak convergence sense: if a sequence $t_{i,n} \rightarrow t_i$ then $\pi_i(t_{i,n})(\cdot | h)$ converges in distribution to $\pi_i(t_i)(\cdot | h)$ for all $h \in \mathcal{H}$. For most applications this will not be restrictive, since any type space that is “closed” is homeomorph to a subset of the universal type space \mathcal{T}^* ³¹ which is itself compact (see Theorem B.1 in Appendix B), making \mathcal{T} itself compact.

For an epistemic type $t_d \in T_d$ and a strategy $\sigma_d \in \Sigma_d$ define the expected continuation value for type t_d as

$$(A.1) \quad W^{t_d}(\sigma_d | h) = W_{\hat{\theta}}^{\hat{\pi}}(\sigma_d | h) \text{ with } \hat{\theta} = \hat{\theta}(t_d), \hat{\pi} = \hat{\pi}_d(t_d)$$

Likewise, given a type $t_p \in T_p$ and a strategy $\sigma_p \in \Sigma_p$ define p 's expected value as

$$V^{t_p}(\sigma_p | h) = V^{\hat{\pi}}(\sigma_p | h) \text{ where } \hat{\pi} = \hat{\pi}(t_p)$$

Also, we write $SBR(t_i) = SBR_{\hat{\theta}(t_i)}[\hat{\pi}(t_i)]$. An agent is then *sequentially rational* if the strategy she chooses is a sequential best response to her beliefs: i.e. $\sigma_i \in SBR(t_i)$. The interactive epistemic representation of types allows us to easily write this assumption; as the subset of *sequentially rational states* $R_i \subset T_i \times \Sigma_i$ defined as:

²⁹Formally, there is always a *belief morphism* between both types of spaces, as studied in Appendix B

³⁰This is a relevant extension, since a very large class of relevant games in economics cannot be modeled with metrizable type spaces. For example, any infinitely repeated dynamic game with a continuous strategy space (such as Cournot duopoly, or most macro applications) are not metrizable.

³¹This is the main lesson from the Universal Type Space Theorem of Mertens and Zamir [1985], Heifetz [1993] and Brandenburger and Dekel [1993]. See Theorem B.1 in Appendix B

$$(A.2) \quad R_i := \{(t_i, \sigma_i) \in T_i \times \Sigma_i : \sigma_i \in SBR_i(t_i)\}$$

for $i \in \{d, p\}$. We write $\Sigma_i^*(\mathcal{T}) \subseteq \Sigma_i$ as the set of all sequentially rational strategies.

Definition A.2 (Robust Implementation). Given a type space $\mathcal{T} = (T_d, T_p, \hat{\theta}, \hat{\pi}_p, \hat{\pi}_d)$ we say that a strategy σ_d is a *robust implementation of trust* if and only if for all histories $(h^\tau, c_\tau) \in \mathcal{H}_p(\sigma_d)$, all $t_p \in T_p$ and all $\sigma_p \in SBR(t_p)$ we have $a^{\sigma_p}(h^\tau, c_\tau) = 1$

Besides the information we have about beliefs (modeled by a type space \mathcal{T}) we might also know (or be willing to assume) some common certainty restrictions on agents beliefs. Following the construction of subsection 4.4 we can extend the definitions of weak and strong rationalizable sets to type spaces, where the sets $WCR_i^k(\mathcal{T})$ and $SCR_i^k(\mathcal{T}) \subset T_i \times \Sigma_i$ correspond to all the weak and strong rationalizable pairs. In Proposition B.3 we adapt the result of Battigalli [2003] to show that for compact type spaces, these sets are compact, which also implies that the set of weak and strong rationalizable sets $WCR_i^\infty(\mathcal{T})$ and $SCR_i^\infty(\mathcal{T})$ are non-empty, compact subsets of T_i . Together with the upper hemicontinuity of the sequential best response correspondence, this implies that the sets of weakly and strongly rationalizable strategies

$$(A.3) \quad \Sigma_i^w := SBR_i\{\mathbf{W}[WCR_i^\infty(\mathcal{T})]\}, \Sigma_i^s := SBR_i\{\mathbf{S}[SCR_i^\infty(\mathcal{T})]\}$$

are compact with respect to the product topology. With this formulation, we can work with common certainty assumptions (of rationality or other assumptions about beliefs) and still retain the type space representation we have been considering. When besides the restrictions on beliefs modeled by the type space \mathcal{T} , d is also willing to make assumptions about common certainty of rationality, this can be thought of as refining the type space by getting a subspace $\hat{\mathcal{T}} \subset \mathcal{T}$, formed only by types that have survived the iterative deletion procedure just described. That is $t_i \in \hat{T}_i = WCR_i^\infty(\mathcal{T})$ if we use common weak certainty, and analogously with $SCR_i^\infty(\mathcal{T})$ for strong certainty.

With some abuse of notation, we will denote the type space resulting of this refinement as $\hat{\mathcal{T}} = \mathbf{WCR}^\infty(\mathcal{T})$ and $\hat{\mathcal{T}} = \mathbf{SCR}^\infty(\mathcal{T})$ for the case of strong certainty. When a type space \mathcal{T} can be written as it's own weak rationalizable refinement (i.e. $\mathcal{T} = \mathbf{WCR}^\infty(\mathcal{T})$) we will say that \mathcal{T} is consistent with weak common certainty of rationality, and analogously for strong certainty

Armed with this concepts, we can give a definition of robust implementation that relates to Bergemann and Morris [2009].

Definition A.3 (Weak Robust Implementation). A strategy σ_d is a weak robust implementation of trust if it implements it for all type spaces \mathcal{T} such that $\mathcal{T} = \mathbf{WCR}^\infty(\mathcal{T})$

In Appendix B I show how this is actually equivalent to doing robustness with the belief space \mathcal{B}_p^w , and likewise for \mathcal{B}_p^s

APPENDIX B. UNIVERSAL TYPE SPACE AND STRONG RATIONALIZABLE STRATEGIES

In this section I adapt the results in Xandri [2012] on the characterization of the Universal Type Space theorem and the study of the topological properties of the sets of weak and strong rationalizable strategies to my setup, for any compact type space we might consider. It generalizes Battigalli and Siniscalchi [1999, 2003], Battigalli [2003] to general topological spaces, which is necessary because their results do not apply to my paper. Their results require either finite strategies or finite periods, because they are obtained by extending Brandenburger and Dekel [1993] to dynamic settings, which works with complete metrizable strategy spaces. However, infinitely repeated games typically involve using the weak convergence topology on the set of strategies, which is non-metrizable if, for example, agents have a continuum of actions in each period. This section is organized as follows: in subsection B.1 I introduce and give some results on the topology of strategy spaces. In subsection B.2 I introduce the notion of hierarchies of beliefs, and study their topological properties. In subsection B.3 I provide a version of the Universal Type Space Theorem (as in Mertens and Zamir [1985], Battigalli and Siniscalchi [1999]) and finally, in subsection B.4 we apply the results we found in the previous sections to characterize the compactness of the set of weak and strong rationalizable strategies, a crucial result for the model studied in this paper.

B.1. Topological Properties of Strategy Spaces. We have that the set from which p chooses is clearly Hausdorff, regular and compact $S_p := \{0, 1\}$.

We will now show that the set from which d chooses is also Hausdorff, compact and regular since it is the product of two compact, regular and Hausdorff spaces:

$$(B.1) \quad S_d := C \times G$$

where $G = \mathcal{M}(Z, \{0, 1\})$, the set of measurable functions $g : Z \rightarrow \{0, 1\}$. $C \subset \mathbb{R}$ is compact by assumption, (and Hausdorff and regular because \mathbb{R} is). We will show that G is also a Hausdorff, compact and regular space with the product topology:i.e. point-wise convergence

$$(B.2) \quad g_n(\cdot) \rightarrow g(\cdot) \iff g_n(z) \rightarrow g(z) \text{ for all } z \in Z$$

The compactness follows from 3 reasons:

- (1) $G \subset \{0, 1\}^Z = \prod_{z \in Z} \{0, 1\} = \hat{G}$ which is a compact space with the product topology described B.2, because of Tychonoff's Theorem. It is also Hausdorff and regular (Theorem 31.2 in Munkres [2000]).
- (2) G is a closed subset of \hat{G} , because of the Dominated convergence theorem (Theorem 2.24 in Folland [1999])
- (3) G is therefore compact (since it is a closed subset of a Hausdorff compact space, Theorem 26.2 in Munkres [2000])

The fact that G is regular and Hausdorff is simply because it is a subspace of \hat{G} , which is a regular and Hausdorff space itself (Theorem 31.2 in Munkres [2000]). Because C and G and

both Hausdorff, compact and regular, S_d is also Hausdorff, compact and regular. Strategies for d are functions

$$\sigma_d : \mathcal{H}_d \rightarrow C \times G$$

which can be written as

$$\Sigma_d \equiv (C \times G)^{\mathcal{H}_d} = S_d^{\mathcal{H}_d}$$

which by Tychonoff's Theorem and Theorem 31.2 in Munkres [2000], is also Hausdorff, compact and regular, with the product topology; i.e.

$$(B.3) \quad \sigma_d^{(n)} \rightarrow \sigma_d \text{ if and only if } \begin{cases} c_n(h^\tau) \rightarrow c(h^\tau) & \text{for all histories } h^\tau \in \mathcal{H}_d \\ g_n(h^\tau, z) \rightarrow g(h^\tau, z_\tau) & \text{for all } (h^\tau, c_\tau, 1, z_\tau) \in \mathcal{H}_d \end{cases}$$

and

$$(B.4) \quad \sigma_p^{(n)} \rightarrow \sigma_p \text{ if and only if } a_n(h^\tau, c_\tau) \rightarrow a(h^\tau, c_\tau)$$

Note that, because of the boundedness of both c and $g(\cdot)$ we can apply the Dominated convergence theorem (Theorem 2.24 in Folland [1999]) to show that the function $V(\cdot | h)$ defined in 4.11 is a continuous function of $\sigma \in \Sigma$, and using Theorem Theorem 2.25 in Folland [1999], we can also show that the continuation value function $W_\theta(\cdot | h)$ defined in 4.10 is also continuous function of $\sigma \in \Sigma$ with the product topology. We summarize the results of this subsection in the following Lemma.

Lemma B.1 (Topology of Σ_i). *The strategy spaces Σ_i for $i \in \{p, d\}$ are Hausdorff, compact and regular topological spaces, with the topology of point-wise convergence (as in B.3 and B.4). Moreover, for all histories $h \in \mathcal{H}_i$, the conditional expected utility functions $V(\sigma | h)$ and $W_\theta(\sigma | h)$ as defined in 4.11 and 4.10 are continuous*

B.2. Hierarchies of Beliefs. Given a topological space (X, τ) , define $\Delta(X)$ as the set of all Borel probability measures on X . If X is a compact, Hausdorff space, then $\Delta(X)$ is also a Hausdorff and compact topological space (Theorem 3 in Heifetz [1993]) with the *weak-** topology. This is the topology of the convergence in distribution: a sequence $\{\lambda_n\}_{n \in \mathbb{N}}$ converges in distribution to λ (written as $\lambda_n \rightsquigarrow \lambda$) if and only if

$$(B.5) \quad \int f(x) d\lambda_n(x) \rightarrow \int f(x) d\lambda(x) \text{ for all } f \in \mathcal{M}(X, \mathbb{R})$$

where $\mathcal{M}(X, \mathbb{R})$ is the set of all measurable functions with respect to the Borel σ -algebra. Moreover, all Borel probability measures on X are also **regular** (Theorem 5, Heifetz [1993]). Therefore, using Tychonoff's theorem, the set $[\Delta(X)]^{\mathcal{H}}$ is also a compact, Hausdorff space with the product topology (having point-wise the weak topology). The set of conditional probability systems on X , which we write $\Delta^{\mathcal{H}}(X)$ is a closed subset of $[\Delta(X)]^{\mathcal{H}}$ (Lemma 1 in Battigalli and Siniscalchi [1999]), and therefore inherits compactness and Hausdorff property. We will say a conditional probability system π is **regular** if and only if $\pi(\cdot | h)$

is a regular measure over X for all $h \in \mathcal{H}$. These results are summarized in the following Lemma

Lemma B.2 (Topology of $\Delta^{\mathcal{H}}(X)$). *Given a Hausdorff and compact topological space X and a family \mathcal{H} of histories, the space $\Delta^{\mathcal{H}}(X)$ of conditional probability systems on X is also a Hausdorff, compact space with the product topology of convergence in distribution: i.e. given a sequence $\{\pi_n\} \in \Delta^{\mathcal{H}}(X)$ and $\pi \in \Delta^{\mathcal{H}}(X)$, we say*

$$(B.6) \quad \pi_n \rightarrow \pi \text{ in } \Delta^{\mathcal{H}}(X) \iff \pi_n(\cdot | h) \rightsquigarrow \pi(\cdot | h) \text{ for all } h \in \mathcal{H}$$

Moreover, every $\pi \in \Delta^{\mathcal{H}}(X)$ is regular.

A useful corollary of Lemma B.2 will be needed for characterizing the best reply correspondence. Given a type space $\mathcal{T} = (T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$, define now the functions $\mathcal{V}(\cdot | h) : \Sigma_p \times \Delta^{\mathcal{H}}(T_d \times \Sigma_d) \rightarrow \mathbb{R}$ and $\mathcal{W}_\theta(\cdot | p) : \Sigma_d \times \Delta^{\mathcal{H}}(T_p \times \Sigma_p)$ as

$$(B.7) \quad \mathcal{V}_p(\sigma_p, \pi_p | h) := V^{\pi_p}(\sigma_p | h)$$

and

$$(B.8) \quad \mathcal{W}_\theta(\sigma_d, \pi_d | h) := W_\theta^{\pi_d}(\sigma_d | h)$$

Corollary B.1 (Continuity of Expected Utility over types). *If T_i are compact topological spaces for $i = 1, 2$, then the functions $\mathcal{V}(\cdot | h)$ and $\mathcal{W}_\theta(\cdot | h)$ are continuous functions (in the weak topology).*

Proof. Since both functions are linear functionals in the space $\Delta(T_{-i} \times \Sigma_{-i})$, for continuity I only need to show boundedness of both functions. This follows directly from Lemma B.2 and the Dominated convergence theorem (which makes the convergence the weak convergence). The continuity of $\mathcal{V}(\cdot | h)$ with respect to σ_d has already been established in Lemma B.1 \square

I will now replicate here the inductive construction of the set of hierarchies of beliefs, as in Battigalli and Siniscalchi [1999]: Define first

$$(B.9) \quad \begin{aligned} X_p^0 &:= \Theta_d \times \Sigma_d \text{ and } X_d^0 = \Sigma_p \\ X_i^1 &:= X_i^0 \times \Delta^{\mathcal{H}_i}(X_j^0) \text{ for } i \in \{d, p\} \end{aligned}$$

and in general

$$(B.10) \quad X_i^k := X_i^{k-1} \times \Delta^{\mathcal{H}_i}(X_j^{k-1})$$

Proposition B.1. *X_i^k is a Hausdorff and compact topological space for all $k = 0, 1, 2, \dots$ and $i \in \{d, p\}$. Moreover, $x \in X_i^k \iff x = (x_{k-1}, \pi_1, \pi_2, \dots, \pi_{k-1})$ where $x_{k-1} \in X_i^{k-1}$ and π_s is a regular CPS on X_j^s for all $s = 1, 2, \dots, k-1$.*

Proof. By induction, I will show that X_i^k is compact, Hausdorff, and it consists of regular measures on its previous ‘‘level’’. Clearly is true for $k = 0$, since from B.9 and Lemma B.1, we

know that X_i^0 is a Hausdorff and compact topological space. Now, assuming $\{X_i^{k-1}\}_{i \in \{p,d\}}$ is Hausdorff and compact, I need to show that $\{X_i^k\}_{i \in \{p,d\}}$ is also Hausdorff and compact. Using Lemma B.2 we then know that $\Delta^{\mathcal{H}}(X_j^{k-1})$ is compact and Hausdorff, and consists of regular measures. This together with definition B.10 gives the desired result. The second result follows from Battigalli and Siniscalchi [1999] which show that we can write X_i^k simply as

$$(B.11) \quad X_i^k = \Sigma_j \times \prod_{s=0}^{k-1} \Delta^{\mathcal{H}_i}(X_j^s)$$

□

Define the set of hierarchies of beliefs for agent $i \in \{p,d\}$ to be the set $H_i = \lim_{k \rightarrow \infty} X_i^k$, which can be written (according to B.11) as

$$(B.12) \quad H_i = \prod_{k=1}^{\infty} \Delta^{\mathcal{H}_i}(X_j^k)$$

So, an element $\mathbf{h} = (\pi_0, \pi_1, \dots) \in H_i$ consists on a CPS π_0 on Σ_j (the strategies of the other agent), a CPS π_1 on $\Delta^{\mathcal{H}_j}(\Sigma_i)$ (the CPS's of j about i 's strategies), a CPS π_2 on $\Delta^{\mathcal{H}_j}(\Delta^{\mathcal{H}_i}(\Delta^{\mathcal{H}_j}(\Sigma_d)))$, and so on. Clearly the space H_i is compact and Hausdorff, because of Proposition B.1 and Tychonoff's theorem. We summarize these results below

Proposition B.2 (Topology of H_i). *The set of hierarchies of beliefs H_i for $i \in \{p,d\}$ as defined in B.12 are Hausdorff and compact topological spaces, with the point-wise convergence in each level:*

$$(B.13) \quad \mathbf{h}_n = (\pi_n^k)_{k \in \mathbb{N}} \rightarrow \mathbf{h} = (\pi^k)_{k \in \mathbb{N}} \iff \pi_n^k(\cdot | h) \rightsquigarrow \pi^k(\cdot | h) \text{ for all } k \in \mathbb{N}, h \in \mathcal{H}_i$$

Moreover, for all hierarchies $\mathbf{h} = (\pi^k)_{k \in \mathbb{N}}$ and all $k \in \mathbb{N}$, we have that $\pi^k \in \Delta^{\mathcal{H}_i}(X_j^k)$ is a regular CPS

B.3. Construction of the Universal Type Space. Not all hierarchies of beliefs will be “rational”, in the sense that upper level beliefs (say, k -order beliefs) may not be consistent with lower level beliefs. We say that a hierarchy $h \in H_i$ is *coherent* when different levels of beliefs are consistent with each other. The formal definition is given by Battigalli and Siniscalchi [1999] (Definition 1) :

Definition B.1 (Coherency). A hierarchy of beliefs $\mathbf{h} \in H_i$ is *coherent* if and only if

$$(B.14) \quad \text{mrg}_{X_j^{k-1}} \pi^{k+1}(\cdot | h) = \pi^k(\cdot | h) \text{ for all } h \in \mathcal{H}_i, k \in \mathbb{N}$$

where $\text{mrg}_{X_j^{k-1}} \pi^{k+1}$ is the marginal of measure π^{k+1} on the projection X^{k-1} .

Definition B.1 is also identical to the definition of projective sequence of regular Borel probability measures, as in Heifetz [1993] (Definition 7), since by B.2 we know that all

measures (and their projections) involved in the hierarchies are regular probability measures. We write H_i^* mean the set of coherent hierarchies of beliefs for agent i . Is easy to see that coherency is a closed restriction of the space H_i , which readily implies that H_i^* is itself a Hausdorff and compact subspace of H_i .

Define

$$(B.15) \quad T_p^* = H_p^*, T_d^* = \Theta_d \times H_d^*$$

So, the universal type sets are simply the sets of all coherent hierarchies of beliefs for each agent. See also that the identity mappings make sense, in that an element in H_i^* is precisely a coherent CPS on the elements of H_j^* . In order to understand the maximality property of the type space we want to construct, in the sense that all other type spaces are in some way embedded into it, I need to define the concept of *type-morphisms*. Given two type spaces $\mathcal{T} = (T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$ and $\mathcal{T}' = (T'_p, T'_d, \hat{\theta}'(\cdot), \hat{\pi}'_p(\cdot), \hat{\pi}'_d(\cdot))$ and a function $\varphi_i : \Sigma_j \times T_j \rightarrow \Sigma_j \times T'_j$ with $i \in \{p, d\}$ and $j \neq i$, define $\tilde{\varphi}_i : \Delta^{\mathcal{H}_i}(\Sigma_j \times T_j) \rightarrow \Delta^{\mathcal{H}_i}(\Sigma_j \times T'_j)$ as the function associating to each CPS π_i on $\Sigma_j \times T_j$ the induced CPS $\tilde{\varphi}_{ii}(\mu_i)$ over $\Sigma_j \times T'_j$, as defined in Battigalli and Siniscalchi [1999] (Subsection 3.1). Formally, given $\pi_i \in \Delta^{\mathcal{H}_i}(\Sigma_j \times T_j)$ we have

$$(B.16) \quad \hat{\varphi}_i(\mu_i)(A | h) = \pi_i(\varphi_i^{-1}(A) | h) \text{ for all measurable } A \subset \Sigma_j \times T'_j, h \in \mathcal{H}_i$$

i.e. it gives events in $\Sigma_j \times T'_j$ the probability according to μ_i in the pre-image of that event in $\Sigma_j \times T_j$.

Definition B.2 (Type-morphisms (Battigalli and Siniscalchi [1999])). Given two type spaces $\mathcal{T} = (T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$ and $\mathcal{T}' = (T'_p, T'_d, \hat{\theta}'(\cdot), \hat{\pi}'_p(\cdot), \hat{\pi}'_d(\cdot))$ we say that a pair of functions $\varphi = (\varphi_p, \varphi_d)$ where $\varphi_i : T_i \rightarrow T'_i$ is a **type morphism** from \mathcal{T} to \mathcal{T}' if and only if the functions φ_i are continuous, and satisfy

$$(B.17) \quad \hat{\pi}'_i[\varphi_i(t_i)] = \tilde{\varphi}_i[\hat{\pi}_i(t_i)] \text{ for all } t_i \in T_i, i \in \{p, d\}$$

and

$$(B.18) \quad \hat{\theta}'[\varphi_d(t_d)] = \hat{\theta}(t_d) \text{ for all } t_d \in T_d$$

When φ is a homeomorphism we say that \mathcal{T} and \mathcal{T}' are *type-isomorphic*.

Conditions B.17 and B.18 state that the beliefs and utility parameters (respectively) of all types in \mathcal{T} can be mapped (in a continuous way) into beliefs and parameters of \mathcal{T}' . The intuitive idea of this definition is that \mathcal{T} is “smaller” than \mathcal{T}' , since every type in T_i can be mapped to a subset of types in T'_i (i.e. the image of φ) that have essentially the same epistemic properties: same beliefs and same utility parameters.

The following Theorem is a simple consequence of Theorem 8 in Heifetz [1993] and Proposition 3 in Battigalli and Siniscalchi [1999], adapted to the modified topological assumptions of this model.

Theorem B.1 (Universal Type Space Theorem). *The sets T_p^* and T_d^* defined in B.15 satisfy:*

$$(B.19) \quad T_p^* \text{ is homeomorphic to } \Delta^{\mathcal{H}_d}(T_d^* \times \Sigma_d)$$

and

$$(B.20) \quad T_d^* \text{ is homeomorphic to } \Theta_d \times \Delta^{\mathcal{H}_d}(T_p^* \times \Sigma_p)$$

with homeomorphisms $Q_p : T_p^* \rightarrow \Delta^{\mathcal{H}_p}(T_d^* \times \Sigma_d)$ and $Q_d : T_d^* \rightarrow \Theta_d \times \Delta^{\mathcal{H}_d}(T_p^* \times \Sigma_p)$. The type space $\mathcal{T}^* = (T_p^*, T_d^*, \hat{\theta}^*(\cdot), \hat{\pi}_p^*(\cdot), \hat{\pi}_d^*(\cdot))$ with $(\hat{\theta}^*(\cdot), \hat{\pi}_d^*(\cdot)) = Q_d(\cdot)$ and $\hat{\pi}_p^*(\cdot) = Q_p(\cdot)$ is called the **Universal Type Space**, and has the following property: for any other type space \mathcal{T} there exists a (unique) type morphism φ between \mathcal{T} and \mathcal{T}^*

Proof. Proposition B.2 tells us that all measures in a hierarchy are regular measures. This, together with Theorems 8 and 9 in Heifetz [1993] proves conditions B.19 and B.20, by applying the Generalized consistency theorem to each individual history $h \in \mathcal{H}$ and constructing the homeomorphism by defining it history by history. The universality condition is an almost direct application of Proposition 3 in Battigalli and Siniscalchi [1999] since we can easily replicate the proof step by step with our topological assumptions. \square

B.4. Topology of Rationalizable sets. In this section I show that the set of rationalizable strategies for any compact type space is in fact, a compact subset of the set of strategies, which I characterized in subsection B.1. Moreover, the set of strongly rationalizable strategies will be in fact, a subset of the weak rationalizable strategy set, implying that strong rationalizability is a closed, stronger solution concept. This will be useful when using the structure theorems in Weinstein and Yildiz [2012]. The main tool I will be using to prove stated in Ausubel and Deneckere [1993]

Theorem B.2 (Berge's Theorem of the Maximum (Berge [1963])). *Let X and Y be topological spaces, with Y regular, a continuous function $f : X \times Y \rightarrow \mathbb{R}$ and a continuous, non-empty and compact valued correspondence $\Gamma : X \rightrightarrows Y$. Then the function*

$$M(x) := \max_{y \in \Gamma(x)} f(x, y)$$

is well defined and continuous, and moreover, the correspondence

$$g(x) := \arg \max_{y \in \Gamma(x)} f(x, y)$$

is non-empty, compact valued and upper hemi-continuous.

The most important consequence of the theorem of the maximum is the continuity and upper hemicontinuity of the value and best response functions, respectively

Proposition B.3 (Continuity of Sequential Best Responses). *For any type space $\mathcal{T} = (T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$ and any $t_i \in T_i$ define the sequential best response correspondence*

$SBR_i : T_i \rightarrow \Sigma_i$ that gives the sequential best responses for type i . Then, if $\hat{\pi}_i(\cdot)$ for $i \in \{p, d\}$ are continuous functions, $SBR_i(t_i)$ is a non-empty, compact valued and upper-hemi continuous correspondence.

Proof. I will only show the continuity of SBR_d , since SBR_p follows a similar (and easier) argument. Corollary B.1 tells us that $\mathcal{W}_\theta(\sigma_d, \pi_d | h)$ is a continuous function of the CPS π_d , and Proposition B.1 tells us that it is also a continuous function of σ_d (taking π_d as given), which makes $\mathcal{W}_\theta(\cdot | h)$ a continuous function over $\Sigma_d \times \Delta^{\mathcal{H}_d}(T_p \times \Sigma_p)$ (with the product topology). Proposition B.1 also implies that the set Σ_d is regular, Hausdorff and compact. The domain of the program is Σ_d , which is a constant correspondence, hence continuous, non-empty and compact valued (since Σ_d is compact). Therefore, we can apply the theorem of the maximum B.2 to show that the correspondence

$$\phi(\theta, \pi_d | h) := \arg \max_{\sigma_d \in \Sigma_d} \mathcal{W}_\theta(\sigma_d, \pi_d | h)$$

is a continuous, non-empty and compact valued u.h.c correspondence of (θ, π_d) for all $h \in \mathcal{H}$, (continuity on θ comes for free with it's finiteness) which therefore implies that the correspondence $\hat{\phi}(\theta, \pi) = (\phi(\theta, \pi_d | h))_{h \in \mathcal{H}}$ is also a continuous, non-empty and compact valued u.h.c correspondence. The desired result then follows from the continuity of $\hat{\pi}_d$, since

$$SBR_d(t_d) = \hat{\phi}[\hat{\theta}(t_d), \hat{\mu}_d(t_d)]$$

a composition of an u.h.c correspondence with a continuous function, which is also a u.h.c correspondence, as we wanted to show. Compactness also follows from continuity (Weierstrass). \square

Now I present the main result of this section

Theorem B.3 (Topological Properties of Weak and Strong rationalizability). *Take a compact type space $\mathcal{T} = (T_p, T_d, \hat{\theta}(\cdot), \hat{\pi}_p(\cdot), \hat{\pi}_d(\cdot))$ and recall the definitions of $WCR_i^k(\mathcal{T}) \subset \Sigma_i$ and $SCR_i^k(\mathcal{T}) \subset \Sigma_i$ as the set of weak and strong rationalizable strategies for type space \mathcal{T} . Then:*

- (1) *The sets $WCR_i^k(\mathcal{T})$ and $SCR_i^k(\mathcal{T})$ are non-empty, compact, Hausdorff and regular spaces, and satisfy $WCR_i^k(\mathcal{T}) \subseteq SCR_i^k(\mathcal{T})$ for all $k \in \mathbb{N}, i \in \{p, d\}$*
- (2) *The rationalizable sets $WCR_i^\infty(\mathcal{T}) \subseteq SCR_i^\infty(\mathcal{T})$ are also non-empty, compact, Hausdorff and regular spaces for $i \in \{p, d\}$*
- (3) *The set of all weak rationalizable strategies $WCR_i^\infty \subset \Sigma_i$ and strong rationalizable strategies $SCR_i^\infty \subset \Sigma_i$ are non-empty, compact, Hausdorff and regular spaces, and satisfy $WCR_i^\infty \subset SCR_i^\infty$ for $i \in \{p, d\}$.*

Proof. (1) follows directly from Propositions 3.5 and 3.6 in Battigalli [2003], since the strategy space Σ_i is compact (Proposition B.1) and the best response correspondences are u.h.c, non-empty compact valued (Proposition B.3). The restrictions on the rationalizable sets

can also be mapped as restrictions on the type space, as shown by Battigalli and Siniscalchi [2003]. The fact that $WCR_i^\infty(\mathcal{T})$ and $SCR_i^\infty(\mathcal{T})$ are non-empty follows from the compactness and non-emptiness proved in follows from (2) and the generalization of Cantor's Theorem, which states that the intersection of a decreasing sequence of non-empty compact sets is non-empty (Theorem 26.9 in Munkres [2000]). Since $WCR_i^k(\mathcal{T})$ and $SCR_i^k(\mathcal{T})$ are compact, they are also closed sets, which make $WCR_i^\infty(\mathcal{T})$ and $SCR_i^\infty(\mathcal{T})$ closed. Because Σ_i is a Hausdorff space, this also implies that $WCR_i^\infty(\mathcal{T})$ and $SCR_i^\infty(\mathcal{T})$ are also compact spaces (Theorem 26.2 Munkres [2000]). Regularity follows from regularity of Σ_i , and therefore we have shown (2). For (3) we use the universal type space theorem B.1 to be able to write

$$(B.21) \quad WCR_i^* \equiv \bigcup_{\mathcal{T}: \mathcal{T} \text{ is a type space}} WCR_i^\infty(\mathcal{T}) = WCR_i^\infty(\mathcal{T}^*)$$

and

$$(B.22) \quad SCR_i^* \equiv \bigcup_{\mathcal{T}: \mathcal{T} \text{ is a type space}} SCR_i^\infty(\mathcal{T}) = SCR_i^\infty(\mathcal{T}^*)$$

and we use again this theorem to recall that the type space \mathcal{T}^* consists of compact type spaces T_i^* with continuous belief functions $\hat{\pi}_i^*$. Therefore we can apply the result in (3) for the particular case of $\mathcal{T} = \mathcal{T}^*$. \square

APPENDIX C. PROOFS AND SUPPLEMENTARY RESULTS

I will need some extra notation for the proofs in this section. Given an appended history $h^s = (h^\tau, h^k)$, I write $h^s \sim h^\tau = h^k$ for the tail of the history. Also, whenever we can decompose h^s in this manner, I will say that h^τ precedes h^s and write $h^\tau \prec h^s$.

Lemma C.1. *Take a history h^τ and θ -rationalization (σ_d, π_d) . Also, let $\nu = (\hat{\sigma}_d, \hat{\pi}_d)$ be another θ -rationalizable pair that satisfies:*

$$(C.1) \quad W_{\theta}^{\hat{\pi}_d}(\hat{\sigma}_d | h^0) \geq \frac{1-\delta}{\delta} S_{\theta, \tau-1} + \mathbb{W}_{\theta}$$

Then, there exists a pair $(\sigma_d^\nu, \pi_d^\nu)$ that also θ -rationalizes h^τ and is such that

$$(C.2) \quad \sigma_d^\nu(h^s) = \hat{\sigma}_d(h^s \sim h^\tau), \pi_d^\nu(\cdot | h^s) = \hat{\pi}_d(\cdot | h^s \sim h^\tau)$$

for all histories $h^s \succ h^\tau$

Proof. Define the pair $(\sigma_d^\nu, \pi_d^\nu)$ for any history \tilde{h}^s as

$$(C.3) \quad \sigma_d^\nu(\tilde{h}^s) := \begin{cases} \sigma_d(\tilde{h}^s) & \text{if } s < \tau \text{ or } \tilde{h}^s = h^\tau \\ \sigma_\theta^*(\tilde{h}^s \sim \tilde{h}^\tau) & \text{if } s \geq \tau \text{ and } h^\tau \not\prec \tilde{h}^s \\ \hat{\sigma}_d(\tilde{h}^s \sim h^\tau) & \text{if } s \geq \tau \text{ and } h^\tau \prec \tilde{h}^s \end{cases}$$

and for any measurable set $A \subset \Sigma_p$

$$(C.4) \quad \pi_d^k(A | \tilde{h}^s) := \begin{cases} \pi_d(A | \tilde{h}^s) & \text{if } s < \tau \\ \underline{\pi}_\theta(A | \tilde{h}^s \sim \tilde{h}^\tau) & \text{if } s \geq \tau \text{ and } h^\tau \not\prec \tilde{h}^s \\ \hat{\pi}_d(A | \tilde{h}^s \sim h^\tau) & \text{if } s \geq \tau \text{ and } h^\tau \preceq \tilde{h}^s \end{cases}$$

so the pair $(\sigma_d^\nu, \pi_d^\nu)$ coincides with (σ_d, π_d) for any histories of length less than $\tau - 1$, and strategies also do it up to time τ . If at history $(h^{\tau-1}, c_{\tau-1}, a_{\tau-1}, z_{\tau-1})$ d deviates from $r = r^{\sigma_d}(h^{\tau-1}, z_{\tau-1})$ going to h^τ , then type θ believes that she will switch to the optimal strong rationalizable strategy from then on, to which the best response is σ_θ^* and the expected payoff is

$$W_\theta^{\pi_d^\nu}(\sigma_d^\nu | h^\tau) = W_\theta^{\underline{\pi}_\theta}(\sigma_d^* | h^0) = \underline{\mathbb{W}}_\theta$$

which is a rationalizable continuation pair. Same is true for the continuations at all histories after h^τ , and so the pair $(\sigma_d^\nu, \pi_d^\nu)$ is rationalizable. Then, to finish our proof, we need to show that it is consistent with h^τ only at $r_{\tau-1}$. Consider first the case where $r_{\tau-1} = 0$ and $S_{\theta, \tau-1} = U_{\theta, \tau-1} - c_{\tau-1} > 0$. Then, the optimal choice under (σ_d^k, π_d^k) is

$$\begin{aligned} \delta W_\theta^{\pi_d^\nu}(\sigma_d^\nu | h^\tau) &\geq (1 - \delta)(U_{\theta, \tau-1} - c_\tau) + \delta \underline{\mathbb{W}}_\theta \iff \\ W_\theta^{\hat{\pi}_d}(\hat{\sigma}_d | h^0) &\geq \frac{1 - \delta}{\delta} S_{\theta, \tau-1} + \underline{\mathbb{W}}_\theta \end{aligned}$$

which is the assumption made in C.1. The other cases are shown in a similar fashion. \square

Proof of Proposition 5.2. Given the functions $(r(\cdot), w(\cdot))$ that satisfy conditions 5.9 and 5.10, I need to construct a θ -rationalizable pair (σ_d, π_d) such that $r^{\sigma_d}(h^\tau, z_\tau) = r(z_\tau)$ for all $z \in Z$. Because the set of rationalizable payoffs is convex, we know that for any $w \in [\underline{\mathbb{W}}_\theta, \overline{\mathbb{W}}_\theta]$ there exist some rationalizable pair (σ_w, π_w) such that

$$W_\theta^{\pi_w}(\sigma_w | h^0) = w$$

then, for all $z \in Z$ we can find a rationalizable pair $(\hat{\sigma}_z, \hat{\pi}_z)$ such that

$$(C.5) \quad W_\theta^{\hat{\pi}_z}(\hat{\sigma}_z | h^0) = w(z)$$

which are rationalizable continuations from time 0 perspective. Moreover, see that that $r(z)$ solves the IC constraint 5.9 for this continuations, which means that it would be the best response at $\tau = 0$ if θ expected the continuation values $w(z)$ starting from $\tau = 1$. Formally, let $h^1(z) = (c_0, a_0, z_0 = z, r_0 = r(z))$ and define the strategy σ_0 as

$$\sigma_0(h^\tau) = \begin{cases} (c_\tau, r(\cdot)) & \text{if } h = h^0 \\ \sigma_z(h^s \sim h^1(z)) & \text{if } h^1(z) \prec h^s \\ \sigma_\theta^*(h^s \sim h^1) & \text{otherwise} \end{cases}$$

i.e. upon deviations in the first period, goes to the optimal robust strategy, and by following the proposed policy $r(z)$ it continues prescribing strategy σ_z after that history, which gives an expected payoff of $w(z)$. This then implies that the policy function is θ -rationalizable at h^0 , and that its expected payoff is

$$W_\theta^{\pi_0}(\sigma_0 | h^0) = \mathbb{E}_z[(1 - \delta)r(z)(U_\theta - c_\tau) + \delta w(z)] \geq \frac{1 - \delta}{\delta} S_{\theta, \tau-1} + \underline{W}_\theta$$

But then we can use Lemma C.1 for the pair $(\hat{\sigma}_d, \hat{\pi}_d) = (\sigma_0, \pi_0)$, finishing the proof. \square

Lemma C.2. *For any strong rationalizable strategy $\sigma_d \in \Sigma_{new}^*$, and any new-rationalizable history, we have*

$$(C.6) \quad \mathbb{E}_z \{r^{\sigma_d}(h^\tau, z) [U_p - c^{\sigma_d}(h^\tau)]\} \geq 0$$

Proof. The proof will follow from 2 steps:

Step 1: Let $\underline{W}_{new} \geq S_{new}^{\max}$. This is equivalent to showing

$$\begin{aligned} \underline{W}_{new} &\geq \frac{\delta}{1 - \delta} \{ \mathbb{E}_z [\max(0, U_p)] - \underline{W}_{new} \} \iff \\ &\underline{W}_{new} \geq \delta \mathbb{E}_z [\max(0, U_p)] \end{aligned}$$

Suppose $\underline{W}_{new} < \delta \mathbb{E}_z [\max(0, U_p)]$. Then the following strategy would be strongly rationalizable: prohibit $r = 1$ at h^τ and in $\tau + 1$ d separates completely. See that since type $\theta = old$ never prohibits r in any rationalizable strategy, then strong certainty of rationality would imply that $\theta = new$ from then on. Therefore, this strategy would then be a robust one, and therefore $\underline{W}_{new} \geq \delta \mathbb{E}_z [\max(0, U_{new})]$ from the fact that \underline{W}_{new} is the maximum utility over robust strategies, and thus reaching a contradiction.

Step 2: $\mathbb{E}_z [r^{\sigma_d}(h^\tau, z) (U_p - c^{\sigma_d}(h^\tau))] \geq 0$ for all $\sigma_d \in \Sigma_{new}^*$ and all rationalizable histories h^τ .

For any rationalizable strategy σ_d we have

$$(1 - \delta) \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) (U_p - c^{\sigma_d}(h^\tau))] + \delta \mathbb{E}_z [\max(0, U_p)] \geq W_{new}^{\sigma_d}(h^\tau) \geq \underline{W}_{new}$$

This also implies then that

$$\begin{aligned} (1 - \delta) \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) (U_p - c^{\sigma_d}(h^\tau))] &\geq \delta \underline{W}_{new} - \delta \mathbb{E}_z [\max(0, U_p)] + (1 - \delta) \underline{W}_{new} \iff \\ \mathbb{E}_z [r^{\sigma_d}(h^\tau, z) (U_p(z) - c^{\sigma_d}(h^\tau))] &\geq \underline{W}_{new} - S_{new}^{\max} \geq 0 \end{aligned}$$

using Step 1 in the last inequality. \square

Proof of Proposition 6.1 . Define $P(c) = \mathbb{E}_z [|U_{old} - c|]$. It can be expressed as

$$P(c) = \int_{z \in Z} |U_{old} - c| f(z) dz = \int_{\underline{U}}^c (c - u) f_o(u) du + \int_c^{\bar{U}} (u - c) f_o(u) du$$

where $f_o(u) := \int_{\underline{U}}^{\bar{U}} f(U_p, u) dU_p$ denotes the partial of U_{old} . Using Leibniz rule

$$P'(c) := \frac{\partial P(c)}{\partial c} = \int_{\underline{U}}^c f_o(u) du - \int_c^{\bar{U}} f_o(u) du = \Pr(U_{old} < c) - \Pr(U_{old} > c)$$

so $\frac{\partial P(c)}{\partial c} > 0 \iff \Pr(U_{old} < c) \geq \Pr(U_{old} > c)$ or equivalently $\Pr(U_{old} < c) \leq \frac{1}{2}$. Then, is easy to see that if condition 1 holds, then for all $c \geq \bar{c}$ we get $P'(c) > 0$ and hence P is increasing in c . Because $c(\cdot) \in [\bar{c}, c_0^*]$ for all $s \in [0, S_{old}^{\max}]$ and is weakly decreasing in s , the result holds. \square

Proof of Lemma 6.1 . I present the proof for the case with $s = 0$, which corresponds to the greatest commitment cost $c_0^* \geq \mathbf{c}(s)$ for all s . For smaller commitment costs the proof will be analogous. It follows from various steps:

Step 1: $\max |U_{old} - c_0^*| > S_{old}^{\max}$.

If this was not the case, then for all z , $c_0^* - S_{old}^{\max} \leq U_{old} \leq c_0^* + S_{old}^{\max}$. If this was the case, using Proposition A.1 we have that

$$\underline{V}(h^\tau, c_0^*) = \int_{z \in Z} \min(0, U_p) f(z) dz \leq \int_{z: U_{old} > 0} U_p dF(z) < \underline{u}_P$$

which violates the definition of c_0^* in 5.13

Step 2: $\min(U_{old}) = \underline{U} < c_0^* - S_{old}^{\max} < c_0^* < \bar{U} = \max(\bar{U})$

The right hand side inequality follows from the fact that if $\bar{U} \leq c_0^* - S_{old}^{\max}$ then

$$\underline{V}(h^\tau, c_0^*) = 0 > \underline{u}_P$$

which will never hold for c_0^* (since $\theta = old$ can drive them to indifference by decreasing the commitment cost enough). From step 1, we either must have that $c_0^* - S_{old}^{\max} > \underline{U}$ or $\bar{U} > c_0^* + S_{old}^{\max}$ (or both). Suppose that the result is not true, so that $\underline{U} \geq c_0^* - S_{old}^{\max}$. Suppose first that $c_0^* - S_{old}^{\max} < \bar{c}$. Then

$$\begin{aligned} \text{(C.7)} \quad \underline{V}(h^\tau, c_0^*) &= \int_{U_{old} > c_0^* - S_{old}^{\max}} \min(0, U_p) f(z) dz \\ &= \int_{U_{old} \in (c_0^* - S_{old}^{\max}, \bar{c})} \min(0, U_p) f(z) dz + \int_{U_p > \bar{c}} \min(0, U_p) f(z) dz \leq \\ &\quad \int_{U_p > \bar{c}} \min(0, U_p) f(z) dz < \int_{U_p > \bar{c}} U_p f(z) dz = \underline{u}_P \end{aligned}$$

violating the definition of c_0^* . If $\bar{c} \leq c_0^* - S_{old}^{\max}$ then

$$\begin{aligned} \underline{V}(h^\tau, c_0^*) &= \int_{U_p > c_0^* - S_{old}^{\max}} \min(0, U_p) f(z) dz < \int_{U_p > c_0^* - S_{old}^{\max}} U_p f(z) dz < \\ &< \int_{U_p > \bar{c}} U_p f(z) dz = \underline{u}_P \end{aligned}$$

from the definition of \bar{c} (since it's the minimum cost that achieves \underline{u}_P in the spot game). Therefore, we have shown that if $\underline{U} \leq c_0^* - S_{old}^{\max}$ then we have $\underline{V}(h^\tau, c_0^*) < \underline{u}_P$, violating the definition of c_0^* . Finally, to show $c_0^* > \bar{U}$, suppose that $\bar{U} \leq c_0^*$. Then any strategy consistent with this choice would give the $\theta = old$ an utility of 0, while we know we will make the reservation utility to be binding (i.e. choose the commitment cost a little smaller so that the contrarian behavior is enough to reach the reservation utility).

Step 3 : $\Pr(U_p > c_0^* - S_{new}^{\max}, U_{old} < c_0^* - S_{old}^{\max}) > 0$

Follows from the fact that $\bar{U} > c_0^* > c_0^* - S_{new}^{\max}$, Step 2 and the full support assumption.

Step 4: $\bar{U} > c_0^* + S_{old}^{\max}$

Suppose that this is not the case: then

$$\underline{V}(h^\tau, c_0^*) = \int_{U_{old} \in (c_0^* - S_{old}^{\max}, c_0^* + S_{old}^{\max})} \min(0, U_p) f(z) dz$$

but see that this is identical to expression C.7. Therefore, replicating the same proof as in Step 2, we conclude the result.

Step 5: $\Pr(U_p < c_0^* + S_{new}^{\max}, U_{old} > c_0^* + S_{old}^{\max}) > 0$

Since $\underline{U} < 0$ we clearly have that $\underline{U} < c_0^* + S_{new}^{\max}$. This, together with the Step 5 and the full support assumption proves the result. \square

Proof of Lemma 6.2 . I first show that for any *old*-rationalizable history h^τ we have $\inf_{\delta \in (0,1)} p(h^\tau, \delta) > 0$. I present the proof for when $c^*(h^\tau) = c_0^*$. Suppose not: then there exists an increasing sequence $\delta_n \in (0, 1)$ such that $p(h^\tau, \delta_n) > 0 \forall n \in \mathbb{N}$ and $p(h^\tau, \delta_n) \searrow 0$. For all δ define the expected utility for the people $\underline{v}(\delta_n) := \underline{V}(h^\tau, c_0^*) = \underline{u}_P$. For all n we have:

$$\underline{v}(\delta_n) < \int_{U_{old} \in (c_0^*(\delta_n) - S_{old}^{\max}(\delta_n), c_0^*(\delta_n) + S_{old}^{\max}(\delta_n))} \min(0, U_p) f(z) dz + p(h^\tau, \delta_n) \left[\max_{U_P \in [\underline{U}, \bar{U}]} (0, U_p) \right]$$

where the first term is the utility in the middle region, and the second term is the natural bound on all regions (particularly in separation regions). Taking limits as $n \rightarrow \infty$:

$$\underline{u}_P = \lim_{n \rightarrow \infty} \underline{v}(\delta_n) \leq \mathbb{E}[\min(0, U_p)] < \underline{u}_P$$

reaching a contradiction. \square

APPENDIX A. CHARACTERIZATION OF $\mathcal{V}(s, c)$

In this section I solve and analyze the solution to the programming problem in subsection (5.17)

Proposition A.1 (Rationalizable Contrarian Strategy). *Consider the programming problem 5.17. Then*

(1) We can rewrite it as

$$(A.1) \quad \mathcal{V}(s, c) = \max_{r(\cdot), n(\cdot)} \mathbb{E}_z [U_p r(z)]$$

$$(A.2) \quad s.t. : \begin{cases} \mathbb{E}_z [(U_{old} - c) r(z) + n(z)] \geq \frac{1}{\beta} s + \mathbb{W}_{old} & (PK \text{ for sacrifice}) \\ r(z) [U_{old} - c + n(z)] \geq 0 \text{ for all } z \in Z & (IC \text{ for } r = 1) \\ [1 - r(z)] [n(z) - U_{old} + c] \geq 0 \text{ for all } z \in Z & (IC \text{ for } r = 0) \\ n(z) \in [0, S_{old}^{\max}] \text{ for all } z \in Z & (Feasibility) \end{cases}$$

(2) There exist $\hat{S} \in (0, S_{old}^{\max})$ such that if for $s < \hat{S}$ then the solution policy $\underline{r}(z)$ is

$$(A.3) \quad \underline{r}(z) = \begin{cases} 1 & \text{if } U_{old} - c > S_{old}^{\max} \\ 1 & \text{if } U_{old} - c \in (-S_{old}^{\max}, S_{old}^{\max}) \text{ and } U_{old} < 0 \\ 0 & \text{if } U_{old} - c \in (-S_{old}^{\max}, S_{old}^{\max}) \text{ and } U_{old} > 0 \\ 0 & \text{if } U_{old} - c < -S_{old}^{\max} \end{cases}$$

(3) If $s \in [\hat{S}, S_{old}^{\max}]$, there exist a positive constant $\alpha(s) \in (0, 1)$ such that

$$(A.4) \quad \hat{r}(z) = \begin{cases} 1 & \text{if } U_{old} - c > S_{old}^{\max} \\ 1 & \text{if } U_{old} - c \in (-S_{old}^{\max}, S_{old}^{\max}) \text{ and } U_p < \gamma(s) (U_{old} - c) \\ 0 & \text{if } U_{old} - c \in (-S_{old}^{\max}, S_{old}^{\max}) \text{ and } U_p > \gamma(s) (U_{old} - c) \\ 0 & \text{if } U_{old} - c < S_{old}^{\max} \end{cases}$$

(4) For all $s \in (0, S^{\max})$ we have $\mathbf{c}(s) \in (\bar{c}, S^{\max})$

Proof. Define $n(z) = \frac{\beta}{1-\beta} [w(z) - \mathbb{W}_{old}]$. If $r(z) = 1$ then we can rewrite the enforceability constraint in 5.9 as $(1 - \beta)(U_{old} - c) + \beta w(z) \geq \beta \mathbb{W}_{old} \iff U_{old} - c + n(z) \geq 0$. Likewise, if $r(z) = 0$ the IC constraint is $\beta w(z) \geq (1 - \beta)(U_{old} - c) + \beta \mathbb{W}_{old} \iff n(z) - U_{old} + c \geq 0$. Finally, rewrite 5.19 as

$$\begin{aligned} \mathbb{E}_z [(1 - \beta)(U_{old} - c) r(z) + \beta (w(z) - \mathbb{W}_{old})] &\geq \left(\frac{1 - \beta}{\beta} \right) s + (1 - \beta) \mathbb{W}_{old} \iff \\ \mathbb{E}_z [(U_{old} - c) r(z) + n(z)] &\geq \frac{1}{\beta} s + \mathbb{W}_{old} \end{aligned}$$

See that for any $z : |U_{old} - c| < S_{old}^{\max}$ then any $r \in \{0, 1\}$ is implementable. However, if $U_{old} > c + S_{old}^{\max}$ then only $r = 1$ is implementable, and if $c - U_{old} < -S_{old}^{\max}$ then only $r = 0$ is implementable. Then, without the promise keeping constraint 5.19 the solution to A.1 is simple:

$$\underline{r}(z) := \begin{cases} 1 & \text{if } U_{old} > c + S_{old}^{\max} \\ 1 & \text{if } |U_{old} - c| < S_{old}^{\max}, U_p < 0 \\ 0 & \text{otherwise} \end{cases}$$

i.e. whenever both policies $r \in \{0, 1\}$ are rationalizable, $\theta = old$ picks the worst policy for p . We will refer to this policy as the *rationalizable contrarian policy*. It will be also the solution when $s = 0$ when the policy \underline{r} satisfies 5.19 with strict inequality. Define $\underline{n}(z)$ as the implementing continuation for $\underline{r}(z)$ that maximizes $\mathbb{E}\{((U_{old} - c)\underline{r}(z) + \underline{n}(z))\}$. Then, it will be also the solution of A.1 if and only if

$$s \leq \beta \mathbb{E}_z [(U_{old} - c)\underline{r}(z) + \underline{n}(z)] - \mathbb{W}_{old} \equiv \hat{s}$$

showing (2). For (3), ignoring for now the IC constraints, use the Lagrangian method (Luenberger [1997])

$$\mathcal{L} = \int U_p r(z) f(z) dz - \gamma \left\{ \int [(U_{old} - c)r(z) + n(z)] - \frac{1}{\beta}s - \mathbb{W}_{old} \right\}$$

where $\gamma \geq 0$ is the Lagrange multipliers of the problem.

$$\frac{\partial \mathcal{L}}{\partial r(z)} = U_p - \gamma(U_{old} - c)$$

then, if $r(z) = 1$ is implementable, the optimum will be $r(z) = 1 \iff U_p \leq \gamma(U_{old} - c)$. If we want to implement $r = 1$ we then set $n(z) = \min\{0, c - U_p\}$. Then, given γ we solve for $r(z | \gamma)$ and $n(z | \gamma)$, and we solve for γ using the promise keeping constraint

$$\int [r(z | \gamma)(U_{old} - c) + n(z | \gamma)] f(z) = \frac{1}{\beta}s + \mathbb{W}_{old}$$

which determines γ as a function of s , showing (3). □

Results are better explained using Figures (8) , (9) and (1) below.




FIGURE 8. Rationalizable Evil Agent strategy, with $s \leq \hat{s}$

See that for U_{old} above $c + S_{old}^{max}$ and below $c - S_{old}^{max}$ the unique rationalizable actions for b are $\hat{r}(z) = 1$ (red) and $\hat{r}(z) = 0$ (green) respectively, as we have seen before. When $U_{old} \in (c - S_{old}^{max}, c + S_{old}^{max})$, both $r = 0$ and $r = 1$ are rationalizable for any z in this region, by appropriately choosing the expected continuation payoffs. Therefore, the worst strategy that p could expect would be one of a *contrarian*: whenever p wants the green button to be played ($U_p < 0$), then the old type would play the opposite action. We can draw an obvious parallelism to the “evil agent” in the robustness literature of Hansen and Sargent [2011], with the restriction that instead of a pure evil agent, the *rationalizable evil agent*, that is only contrarian at states in which the utility of doing her most desirable action is not too high.

See that being a rationalizable contrarian is costly for $\theta = old$, since there are regions in which both p 's and $\theta = old$ most desired action coincide, as we see in the next figure (regions stressed in darker colors)




FIGURE 9. Self-contrarian regions for rationalizable evil agent

Then, when sacrifice is high enough, the disutility generated by the dark regions would not be consistent with the observed behavior. Therefore, to satisfy the “promise keeping” constraint, we must allow the “rationalizable evil agent” not to be fully contrarian, and play her desired action in some states, as we see in the figure below.




TABLE 1. Rationalizable evil agent strategy, with $s > \hat{s}$

Finally, it is easy to see that as the promise keeping constraint becomes more and more binding, the worst type's policy $\hat{r}(\cdot)$ resembles more and more the spot optimum policy $r_{old}^{spot}(\cdot)$. Then, $\mathbf{c}(s) > \bar{c}$ and it approximates it as the promise keeping constraint becomes more binding.

REFERENCES

Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2):pp. 383–396, 1988.

Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):pp. 1041–1063, 1990.

Fernando Alvarez and Urban J. Jermann. Efficiency, equilibrium, and asset pricing with risk of default. *Econometrica*, 68(4):775–797, 2000.

Susan Athey, Andrew Atkeson, and Patrick J. Kehoe. The optimal degree of discretion in monetary policy. *Econometrica*, 73(5):pp. 1431–1475, 2005.

Robert Aumann and Adam Brandenburger. Epistemic conditions for nash equilibrium. *Econometrica*, 63(5):pp. 1161–1180, 1995.

Lawrence M. Ausubel and Raymond J. Deneckere. A generalized theorem of the maximum. *Economic Theory*, 3:99–107, 1993.

- David Backus and John Driffil. Rational expectations and policy credibility following a change in regime. *Review of Economic Studies*, 52:211–221, 1984.
- Gibbons Baker and Murphy. Relational contracts and the theory of the firm. *The Quarterly Journal of Economics*, 117:39–84, 2002.
- Robert J. Barro. Reputation in a model of monetary policy with incomplete information. *Journal of Monetary Economics*, 17(1):3 – 20, 1986.
- Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101 – 121, 1983.
- Pierpaolo Battigalli. Rationalizability in infinite, dynamic games with incomplete information. *Research in Economics*, 57(1):1 – 38, 2003.
- Pierpaolo Battigalli and Giacomo Bonanno. Recent results on belief, knowledge and the epistemic foundations of game theory. *Research in Economics*, 53(2):149 – 225, 1999.
- Pierpaolo Battigalli and Marciano Siniscalchi. Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *Journal of Economic Theory*, 88(1):188 – 230, 1999.
- Pierpaolo Battigalli and Marciano Siniscalchi. Strong belief and forward induction reasoning. *Journal of Economic Theory*, 106(2):356 – 391, 2002.
- Pierpaolo Battigalli and Marciano Siniscalchi. Rationalization and incomplete information. *Advances in Theoretical Economics*, 3, 2003.
- Elchanan Ben-Porath. Rationality, nash equilibrium and backwards induction in perfect-information games. *The Review of Economic Studies*, 64(1):pp. 23–46, 1997.
- C. Berge. *Espaces topologiques (Translated by Patterson, E.M.): Topological Spaces*. Oliver and Boyd, 1963.
- Dirk Bergemann and Stephen Morris. Robust mechanism design. *Econometrica*, 73(6): 1771–1813, 2005.
- Dirk Bergemann and Stephen Morris. Robust implementation in direct mechanisms. *The Review of Economic Studies*, 76(4):pp. 1175–1204, 2009.
- Dirk Bergemann and Stephen Morris. Robust implementation in general mechanisms. *Games and Economic Behavior*, 71(2):261 – 281, 2011.
- Adam Brandenburger and Eddie Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189 – 198, 1993.
- Clive Bull. The existence of self-enforcing implicit contracts. *The Quarterly Journal of Economics*, 102:147–159, 1987.
- Matthew B. Canzoneri. Monetary policy games and the role of private information. *The American Economic Review*, 75(5):1056–1070, 1985.
- Marco Celentani and Wolfgang Pesendorfer. Reputation in dynamic games. *Journal of Economic Theory*, 70(1):109 – 132, 1996.
- V. V. Chari and Patrick J. Kehoe. Sustainable plans. *Journal of Political Economy*, 98(4): pp. 783–802, 1990.

- V. V. Chari and Patrick J. Kehoe. Sustainable plans and debt. *Journal of Economic Theory*, 61(2):230–261, 1993a.
- V. V. Chari and Patrick J. Kehoe. Sustainable plans and mutual default. *The Review of Economic Studies*, 60(1):175–195, 1993b.
- In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221, 1987.
- Daide Debortoli and Ricardo Nunes. Fiscal policy under loose commitment. *Journal of Economic Theory*, 145(3):1005 – 1032, 2010.
- Gerald B. Folland. *Real Analysis : modern techniques and their applications - 2nd Edition*. John Wiley & Sons, Inc., 1999.
- Drew Fudenberg and David K. Levine. Reputation and equilibrium selection in games with a patient player. *Econometrica*, 57(4):pp. 759–778, 1989.
- L.P. Hansen and T.J. Sargent. *Robustness*. Princeton university press, 2011.
- J Harsanyi. Games with incomplete information played by bayesian agents. *Management Science*, 14:159–182, 320–334, 486–502, 1967/68.
- Aviad Heifetz. The bayesian formulation of incomplete information - the non-compact case. *International Journal of Game Theory*, 21:329–338, 1993.
- Robert G. King, Y. K Lu, and Ernesto S. Pasten. Policy design with private sector skepticism in the textbook new keynesian model. *Working P*, 2012.
- Narayana R Kocherlakota. Implications of efficient risk sharing without commitment. *Review of Economic Studies*, 63(4):595–609, October 1996.
- David M Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27(2):253 – 279, 1982.
- David M Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners' dilemma. *Journal of Economic Theory*, 27(2):245 – 252, 1982.
- Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: The inconsistency of optimal plans. *Journal of Political Economy*, 85(3):pp. 473–492, 1977.
- Levin. Relational incentives contracts. *The American Economic Review*, 93:835–857, 2003.
- Ethan Ligon. Risk sharing and information in village economics. *Review of Economic Studies*, 65(4):847–64, 1998.
- Ethan Ligon, Jonathan P. Thomas, and Tim Worrall. Mutual insurance, individual savings, and limited commitment. *Review of Economic Dynamics*, 3(2):216 – 246, 2000.
- Susanne Lohmann. Optimal commitment in monetary policy: Credibility versus flexibility. *The American Economic Review*, 82(1):pp. 273–286, 1992.
- Yang K. Lu. Optimal policy with credibility concerns. *HK UST Working Paper*, 2012.
- D.G. Luenberger. *Optimization by vector space methods*. Wiley-Interscience, 1997.
- George J. Mailath and Larry Samuelson. *Repeated Games and Reputations: Long Run Relationships*. Oxford, 2006.

- Jean François Mertens and Shmuel Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- Paul Milgrom and John Roberts. Predation, reputation, and entry deterrence. *Journal of Economic Theory*, 27(2):280 – 312, 1982.
- James R. Munkres. *Topology - Second Edition*. Prentice Hall Inc., 2000.
- Maurice Obstfeld and Kenneth Rogoff. *Foundations of International Macroeconomics*. MIT Press, 1996.
- Alessandro Pavan and George-Marios Angeletos. Selection-free predictions in global games with endogenous information and multiple equilibria. *Theoretical Economics, Forthcoming*, 2012. URL <http://ideas.repec.org/a/the/publsh/1156.html>.
- David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):pp. 1029–1050, 1984.
- Antonio Penta. Robust dynamic mechanism design. *Working Paper*, 2011.
- Antonio Penta. Higher order uncertainty and information: Static and dynamic games. *Econometrica*, 80(2):631–660, March 2012.
- Torsten Persson and Guido Tabellini. Designing institutions for monetary stability. *Carnegie-Rochester Conference Series on Public Policy*, 39(0):53 – 84, 1993.
- Christopher Phelan. Public trust and government betrayal. *Journal of Economic Theory*, 130(1):27 – 43, 2006.
- Kenneth Rogoff. The optimal degree of commitment to an intermediate monetary target. *The Quarterly Journal of Economics*, 100(4):pp. 1169–1189, 1985.
- Thomas J. Sargent and Lars Ljungqvist. *Recursive Macroeconomic Theory*. MIT Press, 2004.
- Nancy L. Stokey. Reputation and time consistency. *The American Economic Review*, 79(2): pp. 134–139, 1989.
- Nancy L. Stokey. Credible public policy. *Journal of Economic Dynamics and Control*, 15(4):627 – 656, 1991.
- Nancy L. Stokey, Robert E. Lucas, and Edward C. Prescott. *Recursive methods in economics dynamics*. Presidents and Fellows of Harvard College, 1989.
- Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *Review of Economic Studies (1988)*, 55(4):541–554, 1988.
- Carl Walsh. Optimal contracts for central bankers. *The American Economic Review*, 85: 150–167, 1995.
- Jonathan Weinstein and Muhamet Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75(2):365–400, 2007.
- Jonathan Weinstein and Muhamet Yildiz. Robust predictions in infinite-horizon games - an unrefinable folk theorem. *Review of Economic Studies, forthcoming*, 2012.
- Alexander Wolitzky. Reputational bargaining with minimal knowledge of rationality. *Econometrica*, 80(5):2047–2087, 2012.

Juan Pablo Xandri. A note on interactive epistemology in dynamic games with discounting.
MIT Working Paper, 2012.