Misspecified Politics and the Recurrence of Populism

Gilat Levy, Ronny Razin, Alwyn Young*

September 4, 2021

We develop a dynamic model of political competition between two groups that differ in their subjective model of the data generating process for a common outcome. One group has a simpler model than the other group as they ignore some relevant policy variables. We show that policy cycles must arise and that simple world views -which can be interpreted as populist world views- imply extreme policy choices. Periods in which those with a more complex model govern increase the specification error of the simpler world view, leading the latter to overestimate the positive impact of a few extreme policy actions.

"Democracy is complex, populism is simple" (Dahrendorf, 2007)

Voters differ not merely in their economic interests and preferences, but also in their fundamental understanding of the data generating process that underlies observed outcomes. Consequently, because they consider the same historical data through the prism of different models, even fully rational and otherwise similar voters can have persistent differences of opinion. In politics, such differences in model specification translate into differences in realized policy decisions when different groups are in power. The consequent interplay between world views, beliefs and policy can generate systematic correlations across observed data that sustain differing beliefs and biases.

Indeed, understanding the implications of differing world views can shed light on one aspect of populism. While the amorphous concept of "populism" has perhaps as many definitions as authors, the simplicity of populist world views are an important aspect of such movements. Motivated by the experience of populism in Latin America, Dornbusch and Edwards (1991)

*Levy: London School of Economics (e-mail: g.levy1@lse.ac.uk); Razin: London School of Economics (e-mail: r.razin@lse.ac.uk); Young: London School of Economics (e-mail: a.young@lse.ac.uk). This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 681579.

suggest that populism is "an approach to economics that emphasizes growth and income redistribution and de-emphasizes the risks of inflation and deficit finance, external constraints and the reaction of economic agents to aggressive nonmarket policies." Under this view, populist policies are motivated by world views that focus only on a subset of factors (for example, only short-run considerations) compared to a more complex macroeconomic model of growth and inflation suggested by experts and adopted by other political players.[1] The more recent incidences of populism in the western world seem to be centered on a simple ethos of "the people" versus the "elite".[2] This new rhetoric centers on the "will of the people" which, as some recent papers argued, has to be simplified to capture the common ground of many.[3] Similarly, many theories view the defining features of recent populism movements as anti-expert, anti-science and against the rule of law, all complex features of liberal well-functioning democracies.[4] Anti-pluralism, anti-immigration and nationalist views espoused by populists also necessitate a simple definition of group identities.

In practice, when in power or in opposition, populist politicians often offer narrow and extreme solutions, sometimes to detrimental affect. Dornbusch and Edwards (1991) analyse the different stages, as well as the grave consequences, of populist economic reforms in Latin America. Penal populism often overemphasizes the importance of tougher legislation and police funding, ignoring other issues such as the complex intersections of economic inequality, inequality in healthcare, opportunities, mental health issues and structural discrimination.[5]

To focus on the implications of simplistic world views on politics, we consider political competition between groups that share the same interests and preferences over common

[1]See also Guiso et al (2017) and Bernhardt et al (2019) who consider populist politicians who cater to voters' short-term interests and ignore the long-term consequences (e.g., when enacting protectionist trade policies).

[2]See Mudde (2004).

[3]See Sonin (2020) and the survey by Guriev and Papaioannou (2020). A recent literature analyses politicians speeches and party manifestos to compare the complexity of language used by populist politicians to others. Decasdri and Boussalis (2019) analyse a corpus of 78,855 utterances from the most recent Italian parliament and show that a change in allegiance from a populist to a mainstream parliamentary group increases a lawmaker's plenary spoken language complexity. Bischof and Senninger (2018) analyse a measure of complexity to assess the language of manifestos in Austria and Germany in the period 1945–2013. It shows that differences between parties exist and support is found for the conjecture about populist parties as they employ significantly less complex language in their manifestos. Chen, Yan and Hu (2019) compared Clinton's and Trumps' campaign speeches during the 2016 general election showing that Clinton used a more diverse vocabulary compared with Trump.

[4]See also in Sonin (2020) a discussion of how one feature of populism is to offer simple solutions to complicated problems, such as checks and balances.

[5]Enns (2014) documents how shifts in public opinion about the penal system has affected politicians to offer more simplistic policy prescriptions and has increased incarceration. Jennings et al (2017) show how similar penal populist trends affect policy in the UK.

outcomes but differ in their subjective model of the causes of these outcomes. Specifically, we consider the following dynamic model. The common outcome is a linear function of a set of relevant policy variables as well as a random shock. Everyone in the polity is interested in maximizing this outcome (subject to a resource constraint), but individuals differ in their subjective models of the relation between policy variables and the outcome. A subjective model is also linear and considers a set of policies to be relevant. We analyze a polity with a complex type, and a simple type: The simple type's subjective model consists a subset of the relevant policies that comprise the subjective model of the complex type. For example, while a complex subjective model may consider prevention of crime as best treated with a range of policies involving investment in policing but also in employment, education and welfare, a simple model may view crime as stemming from a single cause, lack of law and order due to inadequate police funding. In our dynamic model, both groups start with a prior and learn over time, via the prism of their subjective models, about the parameters determining the effect of each policy variable they deem relevant.

We assume that political competition takes a simple form so that the group that wins is the one that has a higher intensity of preferences (that is, the group that is more keen on winning the election rather than letting the other side win). This group chooses its ideal configuration of policies which are then implemented with small "bureaucratic" noise. At every period the outcome is observed and both groups use OLS to update their beliefs. Our model is then a social learning environment in which the group that takes an action is chosen endogenously, and proponents of both simple and complex solutions learn from the actual outcome delivered by themselves as well as by the rival group. Note that observations are not iid over time as learning and hence current policies depends on previous shocks.

Our key result is that the dynamic process converges and exhibits two important features: perpetual *political cycles*, as well as *extreme policies* advocated by the group holding the simple world view. We first show how the political process involves perpetual political cycles. The intuition for this result is as follows. When only one group is in power indefinitely, in the limit both groups are not "surprised" by the average output that the incumbent produces. Both groups' beliefs will allow them to explain how to produce this average output. These beliefs are different though, as they have different subjective models to explain the outcomes. In contrast to the incumbent group that actually gets to implement its optimal policies given their beliefs, the opposition group strongly believes that a different set of policies is optimal. They understand how to produce the current output, but think they can do more by switching to their preferred policy vector (that depends on their beliefs about the parameters). The opposition's view that they can produce more than the current average output -the optimal one in the eyes of the incumbent- makes them more politically engaged and forces a regime change. This contradicts the supposition that one group is in power indefinitely. The political cycles result relies then on all key features of the model: Social learning, different models of the world, and the mechanism of power transition that relies on intensity of preferences.

Specifically, assume for example that the simple type (S) are in power indefinitely. In this case their belief and those of the complex type (C) will converge to explain the average output produced by S. But C would also believe that by shifting some resources from $S's$ narrow set of policies to the whole vector of policies, they could use the resources more efficiently and generate higher output. This allows them to gain higher intensity and implies that S cannot stay in power indefinitely. Alternatively, if C is in power indefinitely, it is now S that becomes more intense. It understands in the limit how to produce the average output that C is producing. However, as they are focused on a narrow set of policies, they attribute the success of C to these policies being highly effective. They (wrongly) believe that C wastes resources on irrelevant policies and that by shifting resources to their smaller set of relevant policies they can substantially increase output. By becoming increasingly convinced of the effectiveness of a more decisive narrow policy agenda they are then mobilized to win the election. Thus, the economy suffers from inevitable political cycles and the recurrence of narrow and inefficiently extreme policies.

The group holding the simple world view observe their complex rivals invest in policies which they deem irrelevant. For example, if the outcome is crime prevention, S see that C invest in welfare schemes and social integration programs which they deem as irrelevant and wasteful. One possible interpretation for this is that S believe that their rivals target public money to narrow special interests rather than for the "common good" of the majority group. This interpretation bodes well with the anti-elite theory of populism, ascribing to populist supporters the frustration with policies of the elite which they see as unhelpful or not benefiting the "people".[6]

A second feature of the dynamic process is that simple world views imply extreme policy prescriptions. While extreme policies often involve simple rhetoric, we show that having a simple world view in a social learning environment implies extreme policies and beliefs on all the policy variables considered relevant by S. Specifically, the beliefs of S about the effectiveness of policy instruments converge to a multiple, larger than one, of the corresponding beliefs of C. This arises as S learn through the prism of their model both from their own policy choices but also from the policies implemented by C. As a result, when in power, S implement a narrow and exaggerated version of complex policies. Indeed an additional

[6] As Mudde (2004) writes, "In the populist mind, the elite are the henchmen of 'special interests'.... in contemporary populism a 'new class' has been identified, that of the 'progressives' and the 'politically correct'... In the following decades populists from all ideological persuasions would attack the dictatorship of the progressives."

frequent theme in the literature is that the policies of populist politicians are extreme, misguided and harmful to the very groups that support them (e.g., Dornbusch and Edwards 1991). Our framework provides an explanation for the recurrence of subpar outcomes that are supported by rational voters.

While our result is that regime change is inevitable, it is also hastened by negative shocks to the economy. When a negative shock arises, the intensity of the group in power falls by more than that of the opposition group. This arises as the group in power actually implements its ideal policy and hence learns more precisely that such policy is not effective. This accords with the conventional wisdom that large negative shocks trigger populism but might also end its term.

Our paper complements the growing literature about populism by uncovering two aspects of the dynamic political process. First, we highlight a novel mechanism for political cycles when misspecified simple and complex world views are held by different groups in the electorate. In a model with rational individuals we show how the dynamics of learning through misspecified models and endogenous power shifts renders political cycles to be natural and inevitable.

Second we provide a rationale for why simple world views imply extreme and suboptimal policy prescriptions. In this sense, our paper adds to the literature of political-economy models of sub-optimal populist policies. Acemoglu et al (2013) model left-wing populist policies that are both harmful to elites and not in the interests of the majority poor as arising from the need for politicians to signal that they are not influenced by rich right-wing interests. Di Tella and Rotemberg (2016) analyze populism in a behavioural model in which voters are betrayal averse and may prefer incompetent leaders so as to minimize the chance of suffering from betrayal. Guiso et al (2017) define a populist party as one that champions short-term redistributive policies while discounting claims regarding long-term costs as representing elite interests. Bernhardt et al (2019) show how office seeking-demagogues who cater to voters' short term desires compete successfully with far-sighted representatives who guard the long-run interests of voters. Morelli et al (2020) show how in a world with information costs incompetent politicians who simplistically commit to fixed policies can be successful.[7] Our framework expands this literature by linking the pursuit of suboptimal policies to the bias created by a misspecified interpretation of outcomes under optimal policies.

Our theoretical contribution is to establish convergence in a learning environment with a misspecified model. Convergence of beliefs in such environments is not guaranteed, and is especially problematic with multidimensional state spaces (Heidhues et al 2018, Bohren and

[7] For more examples see the recent survey by Guriev and Papaioannou (2020).

Hauser 2019, Esponda et al 2021, and Frick et al 2020). Our paper provides an example of how convergence can be proven in a model with multiple agents, a multidimensional state space and continuous actions. Specifically, we use noise in the implementation of policies to establish convergence in a Bayesian OLS framework.

Interest in learning with misspecified models dates back at least to Arrow and Green (1973), with examples including Bray (1982), Nyarko (1991), Esponda (2008) and, most recently, Esponda and Pouzo (2016) and Molavi (2019). Several recent papers feature interactions between competing subjective models that share features of our framework. Mailath and Samuelson (2019) consider individuals with heterogenous models who exchange beliefs sequentially once they receive a one-off (private) data and characterize conditions under which beliefs converge. Eliaz and Spiegler (2019) present a static model of political competition based upon competing narratives that draw voters' attention to different causal variables and mechanisms. They focus on a static equilibrium and on the possibility of "false positive" variables (which are not necessarily policy variables). Montiel Olea et al (2019), with auctions as a motivation, consider competition between agents that use simple or complex models to explain a given set of exogenous data and find that simpler agents have greater confidence in their estimates in smaller data sets and less confidence asymptotically. In our framework the endogenous data produced by actors with different specifications generates persistent biases and differences in beliefs that asymptotically keep both types politically competitive.

The paper proceeds as follows: Section I presents our basic framework, wherein voters differ in their beliefs regarding the possible determinants of common outcomes. Section II establishes convergence and explains the two key results of cycles and extremism. In Section III we discuss several extensions and modelling assumptions. An appendix contains proofs not in the text.

I. The Model

The Economic Environment: We consider a common outcome $y \in R$ whose realization at time t is governed by the data generating process:

$$y_t = (\mathbf{x}_t + \mathbf{n}_t)'\boldsymbol{\beta} + \varepsilon_t \tag{1.1}$$

where $\mathbf{x}_t$ and $\boldsymbol{\beta}$ are vectors of k policy actions in $\mathbb{R}^k$ and associated parameters, and $\varepsilon_t \in \mathbb{R}$, a mean zero iid normally distributed random shock.[8] We assume that all elements of $\boldsymbol{\beta}$ are non zero. The term $\mathbf{n}_t \in \mathbb{R}^k$ is a $k-$vector of policy noise which could be thought of as small

[8] If the equation for y includes a constant term, we assume it is known by all parties and subtracted from y. This assumption can be relaxed and the results of the model retained as discussed in III.C and III.D below.

policy implementation shocks, and this noise allows us to insure convergence as we discuss in Section II.E. The components of noise $\mathbf{n}_t$ are iid with zero mean and diagonal covariance matrix $\sigma_n^2\mathbf{I}_k$, and are independent of both the policy vector $\mathbf{x}_t$ and the shock to outcomes ε_t. We add noise to all relevant k policies, but alternatively we could add noise to only the set of policies that are implemented at each period and the results would be the same.

Although y is described as a single outcome, one can equally think of it as a weighted average of multiple outcomes that are influenced by $\mathbf{x}_t$.[9] Below, we use bold letters to denote vectors and matrices and when it does not lead to confusion often drop the subscript t, writing $\mathbf{x}$, $y, \mathbf{n}$ and ε.

Subjective Models: We assume that citizens are divided into two "types" based upon their subjective model about which of the unknown parameters in $\boldsymbol{\beta}$ can potentially be non-zero. We shall focus our analysis on the case where "complex" types (C) that believe all elements of $\boldsymbol{\beta}$ might be non-zero compete politically with "simple" types (S) whose model is misspecified, in that they exclude some relevant policies and are hence certain that for these policies, the elements of $\boldsymbol{\beta}$ are zero. We can easily extend our analysis to the case in which $C's$ model is also misspecified, and to the case in which also S considers some relevant policies which are not considered by C (see Section III.A). We assume that both groups know that ε is normally distributed.[10]

Example 1 (Tackling crime: *Social policies versus law and order*): Let $y = \beta_1 x_1 + \beta_2 x_2 + \varepsilon$ denote an aggregate measure of welfare which is negatively related to the rate of crime. Assume that x_1 is the level of investment in policing and law and order, and x_2 is the level of investment in youth services, education/employment opportunities, or integration programs. Suppose that S believes that $\beta_2 = 0$ so that only law and order is relevant. In this case S believes that investment in education, x_2, is wasteful. This world view might come from a belief that crime is affected by individual characteristics and can only be influenced by deterrence. C on the other hand, believes that a combination of both policies is effective. The limit case when $\beta_2 \to 0$ would be the case where group S also has the correct model.

In the general model group S has a set of $k_s < k$ policies that it deems relevant, while it

[9] If utility is a weighted average of j components each with $y_{jt} = (\mathbf{x}_t + \mathbf{n}_t)'\boldsymbol{\beta}_j + \varepsilon_{jt}$, then the outcome, parameters and error term in 1.1 are simply the weighted average of those components.

[10] An important assumption in our analysis is that the types do not change their model when faced with data. As mentioned above, there is a growing literature that adopts this assumption. How and whether economic agents change their models is an open question. In our political application, changing models may be difficult as it is not obvious for example that the simple type has data about policy variables that she deems irrelevant. Gagnon-Bartsch et al (2018) provide a more general theory to justify why individuals will only pay attention to variables that are relevant to their model.

believes that the effect of all other policies on y is null. We will use the subscript $i \in \{S, C\}$ to denote the group, where $\mathbf{x}_i$ and $\bar{\boldsymbol{\beta}}_i$ denote the policy choice and the mean beliefs of group i. Unless otherwise specified, all vectors of policies and beliefs will be $k-$vectors, where zeroes will be used for elements of the vector which are null. Specifically, $\mathbf{x}_s$ and $\bar{\boldsymbol{\beta}}_s$ are $k-$vectors, with zeroes in all the elements pertaining to the $k - k_s$ policies that S deems irrelevant. Below we assume linear utility; together with the linear formulation of y, this implies that only mean beliefs matter, and we henceforth denote the vector of mean beliefs at period t by $\bar{\boldsymbol{\beta}}_{st}$ and $\bar{\boldsymbol{\beta}}_{ct}$.

Although the subjective model of $i \in \{S, C\}$ is fixed, the beliefs of type $i \in \{S, C\}$ about the magnitude of the elements and in particular the expectation of these beliefs, $\bar{\boldsymbol{\beta}}_{it}$, are updated over time according to OLS, and are consistent with Bayes rule. We use $\mathbf{H}_t = \mathbf{X}_t + \mathbf{N}_t$ to denote the t x k history of desired policy and iid noise and assume that prior beliefs are normally distributed. As our results are in any case asymptotic, normal beliefs of this sort can be justified by the observation of a long pre-history of policy, as under fairly general conditions the likelihood function determines the shape of the posterior (Zellner 1971).[11] As the normally distributed error ε is independent of contemporaneous policy, mean beliefs reduce to the standard OLS parameter estimates with each type focusing on their relevant policy regressors.

Preferences and optimal policies: We model utility with the minimal structure that allows for a tractable presentation. Specifically, we assume the utility citizens derive from the common outcome is linear:

$$U_t(y_t) = y_t, \tag{1.2}$$

and that the choice of policies is subject to the budget constraint $\mathbf{x}_t'\mathbf{x}_t \leq R$, where R is some bounded, exogenously-given, resource. The quadratic form of the constraint allows us not to worry about the signs of the elements of $\boldsymbol{\beta}$ or $\mathbf{x}$.

Given the above, it readily follows that at any period, given some mean beliefs $\bar{\boldsymbol{\beta}}_i$ for type $i \in \{S, C\}$, the optimal myopic policy solves

$$\max_{\mathbf{x}_i \in \mathbb{R}^{k_i}} \ \mathbf{x}_i'\bar{\boldsymbol{\beta}}_i + \lambda_i(R - \mathbf{x}_i'\mathbf{x}_i) \tag{1.3}$$

resulting in

$$\lambda_i = \frac{1}{2}\sqrt{(\bar{\boldsymbol{\beta}}_i'\bar{\boldsymbol{\beta}}_i)/R}, \ \mathbf{x}_i^* = \bar{\boldsymbol{\beta}}_i\sqrt{R}/\sqrt{\bar{\boldsymbol{\beta}}_i'\bar{\boldsymbol{\beta}}_i} \Rightarrow \tag{1.4}$$

[11] Specifically, consider prior beliefs across relevant policies that are normally distributed with mean $\bar{\boldsymbol{\beta}}_0$ and joint covariance matrix $\sigma_0^2\mathbf{V}_0^{-1}$, while the prior probability density function on σ_0^2 is inverted gamma. We then define the pre-history such that $\mathbf{V}_0 = \mathbf{H}_0'\mathbf{H}_0$ and $\bar{\boldsymbol{\beta}}_0 = (\mathbf{H}_0'\mathbf{H}_0)^{-1}\mathbf{H}_0'\mathbf{y}_0$.

$$\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*] \equiv \mathbf{x}_i^{*\prime}\bar{\boldsymbol{\beta}}_i = \sqrt{\bar{\boldsymbol{\beta}}_i'\bar{\boldsymbol{\beta}}_i}\sqrt{R}$$

While the solution to the Lagrangian problem is straightforward, we note here that given the constraint R, types which have more extreme parameter estimates, as measured by $\bar{\boldsymbol{\beta}}_i'\bar{\boldsymbol{\beta}}_i$, believe they know how to pursue more effective policies, as measured by their expected outcome when choosing their optimal policies, $\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*]$, and consequently feel more constrained by the resource limitation R, as measured by λ_i.

In each period political competition will determine which type chooses current period policies:

The political competition: We first define the notion of intensity of preferences. Let

$$I_i = \bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*] - \bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_j^*], \tag{1.5}$$

where $\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_j^*]$ is type $i's$ expected outcome when type j chooses their optimal policy. The intensity of preferences of type i is therefore the loss this type incurs from $j's$ ideal policy compared to her own ideal policy, given her subjective model. I_i does not necessarily equal $-I_j$ as beliefs differ across the two types. Specifically:

$$I_s = \mathbf{x}_s^{*\prime}\bar{\boldsymbol{\beta}}_s - \mathbf{x}_c^{*\prime}\bar{\boldsymbol{\beta}}_s, \tag{1.6}$$

$$I_c = \mathbf{x}_c^{*\prime}\bar{\boldsymbol{\beta}}_c - \mathbf{x}_s^{*\prime}\bar{\boldsymbol{\beta}}_c.$$

We assume that at any period t, the type that has the higher intensity of preferences wins the election, and then implements her ideal policy in that period (we focus then on myopic choices of policies and discuss strategic choices of policies in Section III.B).

Below we construct a political competition model which rationalizes why intensity of preferences is an engine for power shifts. Assume that the polity consists of two equally sized groups, simple and complex, each a continuum. Each group is represented by a "citizen-candidate" that runs in the election and if elected, implements the type's ideal policy.[12] Voting is costly, but citizens vote because they believe that with some (exogenous) probability p their vote will be pivotal.[13] Consequently, a voter l of type i will vote (for their own representative) if the expected gain from the implementation of type $i's$ optimal policies relative to those of type j exceeds voter $l's$ cost of voting, c_l, i.e.:

$$pI_i > c_l \tag{1.7}$$

Assume that c_l is iid drawn from a distribution of voting costs $G(c)$ and that the cost distribution is the same for both groups. Thus, the vote share that candidates of each type

[12] Given how we model voting decisions, it is easy to see that the presence of such candidates, offering voters of each type their ideal policy, will drive out all other policy platforms.

[13] For simplicity we are not modelling strategic voting, i.e., p is not determined endogenously in the model.

garner will be an increasing function of the intensity of their type. Consequently, the election is won by the candidate representing the type with the greatest preference intensity. The results below can be generalized to allow for unequal group sizes and different distributions. For example, the case of unequal groups implies the smaller group will require a certain margin of voting preference intensity to motivate its base enough to win an election.

Before defining our equilibrium notion, we now characterize voters' intensity of preferences:

Lemma 1: *The following are equivalent:*

(i) The intensity of group i is greater than that of group j;

(ii) The magnitude of group i's belief vector is greater than that of group j, i.e., $\bar{\boldsymbol{\beta}}_i' \bar{\boldsymbol{\beta}}_i > \bar{\boldsymbol{\beta}}_j' \bar{\boldsymbol{\beta}}_j$;

(iii) Group i expects to achieve a higher outcome when in power than group j do when they are in power, i.e., $\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*] > \bar{y}[\bar{\boldsymbol{\beta}}_j, \mathbf{x}_j^*]$.

Intuitively, individuals with more extreme parameter estimates feel the resource constraint more keenly (as exemplified by the Lagrange parameter λ_i in 1.4) and hence lose more from a sub-optimal movement away from their constrained choice. Hence the dynamic change of power in our model will be determined by the relative magnitude of beliefs of the two types.[14] As $\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*] = \sqrt{\bar{\boldsymbol{\beta}}_i' \bar{\boldsymbol{\beta}}_i} \sqrt{R}$, intensity is also higher for the group that believes it can produce a higher level of output.

Dynamics: We consider then the following dynamic process:

1. In any period t, the winning type $i \in \{S, C\}$, chooses her ideal policy $\mathbf{x}_{it}^*$ given her beliefs, $\bar{\boldsymbol{\beta}}_{it}$.

2. Given $\mathbf{x}_{it}^*$, $y_t = (\mathbf{x}_{it}^* + \mathbf{n}_t)' \boldsymbol{\beta} + \varepsilon_t$ is realized (and utility U_t gained). Both types update their beliefs using OLS. Mean beliefs evolve to $\bar{\boldsymbol{\beta}}_{j(t+1)}$, for all $j \in \{S, C\}$.

3. Type S (C) wins the election at period $t+1$ if its intensity is higher, that is, $\bar{y}[\bar{\boldsymbol{\beta}}_{s(t+1)}, \mathbf{x}_{s(t+1)}^*] > (<) \ \bar{y}[\bar{\boldsymbol{\beta}}_{c(t+1)}, \mathbf{x}_{c(t+1)}^*]$. In the case of equal intensities, some tie breaking rule determines the winner.[15]

While the model of group S is misspecified, they use OLS estimation to rationally update their beliefs. Crucially, while the shocks are iid, the policies and hence regressors are not,

[14] To see the proof of Lemma 1, note that the gain in expected utility from pursuing an optimal policy $\mathbf{x}_i^*$ versus an alternative policy in which a $k \times 1$ vector $\boldsymbol{\delta}$ is added to $\mathbf{x}_i^*$ is given by $-\boldsymbol{\delta}' \bar{\boldsymbol{\beta}}_i$. Substituting using optimal policies and the fact that $-\boldsymbol{\delta}' \mathbf{x}_i^* = \frac{1}{2} \boldsymbol{\delta}' \boldsymbol{\delta}$, as both $\mathbf{x}_i^{*\prime} \mathbf{x}_i^*$ and $(\mathbf{x}_i^* + \boldsymbol{\delta})'(\mathbf{x}_i^* + \boldsymbol{\delta})$ equal R, we get that:

$$\bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^*] - \bar{y}[\bar{\boldsymbol{\beta}}_i, \mathbf{x}_i^* + \boldsymbol{\delta}] = \sqrt{\frac{\bar{\boldsymbol{\beta}}_i' \bar{\boldsymbol{\beta}}_i}{R}} \frac{\boldsymbol{\delta}' \boldsymbol{\delta}}{2}.$$

[15] The exact tie breaking rule is inconsequential.

as each type learns from the observed actions which themselves depend on the endogenously evolving beliefs.

Remark 1 (*Anti-elite sentiment*): Note that S observes that group C invests in policies that group S feels are irrelevant. In our model it is not necessary to assume that S knows that C has a different model. It can be the case that S believes that C is corrupt and invests in policies that do not benefit the general public but only a select group. This fits well with the anti-elite interpretation of populism ascribing to populist supporters frustration with policies of the liberal elite which they see as unhelpful or not benefiting the "people". For example, in relation to Example 1, they might view spending on welfare benefits or integration programs as wasteful and corrupt.

II. Perpetual Cycles and Extremist Populists

In this section we present Theorem 1, our main result, characterizing the limit dynamics of the model. In particular, there are perpetual political cycles with both groups taking power and relatively extreme policies espoused and implemented by group S when they are in power. To formalize the notion of political cycles, let θ_{jt} denote the random variable equal to the share of time $j \in \{S, C\}$ has been in power up to period t. Let $\boldsymbol{\beta}_s$ be the $k-$vector that agrees with the true parameters of $\boldsymbol{\beta}$ on all policies that group S deem relevant and has zero entries on all other $k - k_s$ policies and let $\tau^* = \sqrt{\boldsymbol{\beta}'\boldsymbol{\beta}/\boldsymbol{\beta}_s'\boldsymbol{\beta}_s} > 1$. We then have (for the proof see the Appendix):

Theorem 1: *For sufficiently small σ_n^2, the polity almost surely converges to: (i)* ***Political cycles:*** $\theta_{st} \overset{a.s}{\to} \theta_s = (1 - \tau^*\sigma_n^2)/(1 + \tau^*)$, $0 < \theta_s < 1$, *(ii)* ***Correct beliefs for C:*** $\bar{\boldsymbol{\beta}}_{ct} \overset{a.s}{\to} \bar{\boldsymbol{\beta}}_c = \boldsymbol{\beta}$, *(iii)* ***Colinear and extreme beliefs for*** $S : \bar{\boldsymbol{\beta}}_{st} \overset{a.s}{\to} \bar{\boldsymbol{\beta}}_s = \tau^*\boldsymbol{\beta}_s$.

Theorem 1 implies that the polity will perpetually oscillate between group C and group S holding power. Although the beliefs of both groups converge, they continuously oscillate to be on both sides of the equal intensity plane. As a result there is a perpetual change in power. Moreover, the change in power will imply dramatic policy changes; when S takes power it will nullify the policies that its model ignores and amplify its policy positions on the policies it deems relevant, to a multiple of the positions that C takes when it is in power.

We first discuss the intuition for the main findings of political cycles and extremism assuming that beliefs and the share of time that S is in power, θ_{st}, converge. We then provide a more technical discussion of how we prove convergence.

Given that C has the correct model and given the policy implementation noise, it is intuitive that C will eventually learn the true parameters of the model, and so $\bar{\boldsymbol{\beta}}_{ct} \overset{a.s}{\to} \boldsymbol{\beta}$. We focus then on the limit beliefs of S, $\bar{\boldsymbol{\beta}}_s$, as well as on the limit values θ_s and θ_c (where

$\theta_s + \theta_c = 1$). Let $\mathbf{x}_i^*$ denote the optimal k -vector of policies of group $i \in \{S, C\}$ given their expected limit beliefs $\bar{\boldsymbol{\beta}}_i$. For expositional reasons, we will henceforth consider the case of no policy noise, so that $\sigma_n^2 = 0$ (we reinstate the policy noise in Section II.E where we discuss convergence).

Let $\mathbf{\Phi}$ be a $k \times k$ matrix with a submatrix $\mathbf{I}_{k_s}$ for the rows and columns relating to $S's$ model and zeroes everywhere else. Given convergence, the OLS coefficients converge to satisfy the following equations:[16]

$$\mathbf{\Phi}[\theta_s \mathbf{x}_s^* (\mathbf{x}_s^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_s^{*\prime} \boldsymbol{\beta}) + \theta_c \mathbf{x}_c^* (\mathbf{x}_c^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_c^{*\prime} \boldsymbol{\beta})] = \mathbf{0}, \tag{2.1}$$

where $(\mathbf{x}_s^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_s^{*\prime} \boldsymbol{\beta})$ and $(\mathbf{x}_c^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_c^{*\prime} \boldsymbol{\beta})$ are the average mistakes that group S makes under her beliefs, when S is in power and when C is in power respectively. That is:

$$\mathbf{x}_s^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_s^{*\prime} \boldsymbol{\beta} = \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] - \bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*], \tag{2.2}$$

$$\mathbf{x}_c^{*\prime} \bar{\boldsymbol{\beta}}_s - \mathbf{x}_c^{*\prime} \boldsymbol{\beta} = \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_c^*] - \bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*].$$

These expressions of average mistakes will play an important role in the intuition we give below for our key results. The mistakes will determine the direction of change for the beliefs of group S, depending on who is in power. As we will show this will imply a law of motion for the beliefs of S in the direction of changing power through the relative intensity of the group's beliefs.

A. The Cycles of Populism

We now show how cycles must arise. When only one group is in power, let's say S, so that $\theta_s = 1$, this implies, from (2.1) and (2.2), that in the limit $S's$ beliefs are such that they are not "surprised" anymore by the *average* output they produce, and so are not mistaken on average:

$$\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*] \tag{2.3}$$

And trivially, C also predicts correctly the average output $\bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*]$. But note that C can do better than $\bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*]$. Its limit beliefs explain what S does, but they are different as they have a different model. Hence if it switches to its optimal policies given $\bar{\boldsymbol{\beta}}_c = \boldsymbol{\beta}$, namely $\mathbf{x}_c^*$, C can generate a higher output. Specifically, by shifting some resources from a narrow set of policies to the whole vector of policies, C uses the resources more efficiently and generates higher output. In other words,

$$\bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*] > \bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*] = \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*]. \tag{2.4}$$

[16]This is the first order condition derived when minimizing expected squared mistakes.

This, by Lemma 1, implies then that C becomes more intense than S when S is assumed to hold power indefinitely. We then have a contradiction to this assumption, and so S must be replaced and cannot be in power for ever.

The exact same argument implies that if C were in power indefinitely, group S will become more intense. When $\theta_c = 1$, again, the beliefs of S converge to explain the average output produced by C; in the limit S is not surprised by what C is producing, with $\bar{\boldsymbol{\beta}}_s$ solving $\mathbf{x}_c^{*\prime}\bar{\boldsymbol{\beta}}_s = \mathbf{x}_c^{*\prime}\boldsymbol{\beta}$. But given these beliefs, S realises that it can produce more by shifting resources to its own narrow set of policies. Namely:

$$\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] > \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_c^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*] = \bar{y}[\bar{\boldsymbol{\beta}}_c, \mathbf{x}_c^*] \tag{2.5}$$

where the first inequality follows from the fact that $\mathbf{x}_s^*$ maximizes output given $\bar{\boldsymbol{\beta}}_s$, and the other equalities follows from the observation that learning in the limit implies that both groups have beliefs that explain expected output.

In other words, when one group is in power long enough, both groups learn to explain the average output it produces. But while the group in power also gets to implement its ideal policies given these beliefs, the group in opposition believes it can do better; this implies that it becomes more intense and a power shift is inevitable. Thus long term dynamics must include political cycles. This can also be interpreted as a form of *incumbency disadvantage.* While the incumbent party implements its ideal policies given its beliefs, the opposition party finds the incumbent's policies wasteful, either as the incumbent invests in what it finds to be irrelevant policies (as is the case when S is in opposition), or as the incumbent invests too much in some policies (as in the case when C is in the opposition). In either case this leads the opposition party to believe it can induce a better outcome and makes it more motivated to replace the incumbent.

B. Fighting Crime: Cycles of Investment in Law and Order

To get some intuition about the cycles and long term dynamics, we now return to Example 1, where the true model is $y = \beta_1 x_1 + \beta_2 x_2 + \varepsilon$. We focus on crime prevention, y, with the policies being either investment in law and order, x_1, or in social policies, x_2 (these could be integration/employment opportunities or social welfare policies). In this example S believes that $\beta_2 = 0$ so that only law and order, x_1, is relevant. Recall that $C's$ beliefs converge to the true values of β_1 and β_2. For simplicity of exposition assume now that C already hold the true beliefs and so we can focus on the evolution of the beliefs of S, which we denote by $\bar{\beta}_1$. Note that in this simple model $x_{1,s}^* = \sqrt{R}$ for all beliefs of S, and so let us simplify further and assume $R = 1$.

In the criminology literature the populist tendency towards policies that are "tough on crime" has been termed penal populism. Enns (2014) documents the historical patterns of both voter attitudes towards law and order and policy outcome measures in the US. As a proxy for actual law and order policies, Enns (2014) uses changes in incarceration rates. The data shows that both attitudes towards "being tough on crime" and changes in incarceration have been decreasing from 1950-1970, increasing from 1970-mid 1990s and decreasing until 2010. While these trends are somewhat correlated with the level of crime, this relation exists even when controlling for crime rates and other economic variables. Similar patterns are shown in Jennings et al (2017) for the UK.

We now illustrate how the long term dynamics imply cycles of investment in law and order in our example. Whenever S is in power, they invest all resources in law and order, but over time become disillusioned with the benefits of higher police funding. In contrast, when C is in power, investment in law and order is lower (as it is accompanied with other policies), but over time in opposition S becomes convinced that tougher measures and more investment in policing are crucial.

Consider the case when S is in power for a sufficiently long time. When in power they have the true model to assess the impact of their investment in x_1; since they set $x_2 = 0$ their learning about x_1 is not biased. As their policy on x_1 is fixed, their beliefs converge to the true impact of law and order, β_1. It is then easy to see how at some point, once the beliefs of S are sufficiently close to β_1, C has greater intensity as it can produce a more effective crime prevention outcome by spreading resources efficiently on both policies. This implies that S cannot stay in power for too long.

Alternatively, if C is in power for a sufficiently long time, $S's$ beliefs will suffer from an omitted variable bias. Again, as $C's$ action is fixed, it is straightforward to show that $\bar{\beta}_1$ will converge to solve:

$$\bar{\beta}_1 x^*_{1,c} = \beta_1 x^*_{1,c} + \beta_2 x^*_{2,c} \Rightarrow \tag{2.6}$$

$$\bar{\beta}_1 = \beta_1 + \beta_2(\beta_2/\beta_1)$$

where we had substituted for the optimal policies of C. As derived in the previous section, this implies that at some point, S develops greater intensity as S believes that substituting $x^*_{1,s}$ for $x^*_{1,c}$ will produce greater output on average.[17] This greater intensity of S implies that C cannot stay in power for too long, and cycles must arise.

The implication of the political cycles result is that the beliefs of S must converge to satisfy *equal intensity,* as in any other case one group will be in power indefinitely. This pins down the (excessive) beliefs of S as follows:

$$y[\bar{\beta}_1, x_{1,s}] = y[\boldsymbol{\beta}, x^*_c] \Rightarrow \bar{\beta}_1 = \sqrt{(\beta_1)^2 + (\beta_2)^2} > \beta_1. \tag{2.7}$$

[17]This arises as using Lemma 1, the intensity of S, $(\bar{\beta}_1)^2 = (\beta_1 + \beta_2 \frac{\beta_2}{\beta_1})^2$, is greater than $\beta_1^2 + \beta_2^2$, which is the intensity of C.

Figure 1 below describes the asymptotic beliefs of S, close to the equal intensity beliefs defined above (note that these beliefs must be "sandwiched" between the limit beliefs that arises when each group is in power indefinitely). Close to the equal intensity beliefs, whenever the intensity of preferences of S is larger than that of C, it gains power and implements its ideal policy. But then, on average, S becomes disappointed in the outcomes it generates and moderates its beliefs towards the true β_1. Simple voters are then systematically disappointed by the outcomes of their extreme investment in law and order. This leads to a gradual diminution of beliefs, until those with more complex views once again take power. But whenever $S's$ intensity falls below that of C, and C gains power, S starts to inflate again the effectiveness of law and order. The surprising success of C's policies (which includes an array of other policies such as investment in education, integration and employment) gradually convinces simple voters of the value of law and order policies as they believe the success of C stems from these policies only. This omitted variable bias that affects their beliefs increases their probability of voting in favour of populist politicians who advocate narrow and extreme solutions to complex problems.

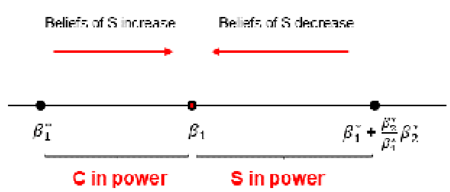

Figure 1: The dynamics of $S's$ beliefs

The equal intensity beliefs is, as can be seen in the Figure, a basin of attraction for the dynamics described above. Such a basin of attraction also implies convergence of the beliefs of S as well as of the share of time that each group is in power. Compared to our general model, in this example there are two simplifying features that make convergence easier to prove. First, the discrete set of actions for each group allows us to get a clear handle on the limit dynamics of the beliefs. Second, the unidimensionality of the beliefs of S implies only two directions in which the beliefs of S can move (up or down). In contrast, in our general model we have continuous actions and multidimensional policy space. We expand on this in Section II.E.

Note that the equal intensity beliefs of S is such that S does not correctly conjecture the average output at each regime. In fact, the beliefs of S evolves to balance the mistakes in these conjectures *across* the two regimes. This feature is behind the extremism result that we show more generally below. These mistakes are also the engine of the dynamics of power shifts in our model. While the above example illustrates systemic regime changes over time, random shocks to the economy also play an important role in regime change, and we explore this in Section II.D.

C. Simplicity implies Extremism

Many historical and contemporary examples of populistic politicians are considered to offer extreme rhetoric and policy prescriptions. Examples include left-wing economic populism in Latin America, or more recently anti-immigration and anti rule-of-law rhetoric (and policies) of right-wing populists in Europe and the US. Among the latter, examples abound, such as the policies of Victor Orban in Hungary, Poland's Law and Justice Party, the far-right Alternative for Germany (AfD), the National Front in France or Lega Nord in Italy, and the Republican party in the US under Trump's leadership.

In Example 1 above, we saw that the simple converge to believe that law and order is more effective compared with the beliefs of the complex type on these type of polices. This implies in general that their policy prescriptions for law and order are also more extreme. We now show that the relation between a simple model and extreme policy prescriptions holds in the general model when S considers more than one relevant policy. We show that $S's$ beliefs and policies will be more extreme on each of the policies which both groups find relevant. In particular, $S's$ beliefs must be colinear with those of C on the relevant shared policies.

Extremism in our model is a result of the fact that S learns through the prism of her model both from her own policy choices but also from the policies implemented by C. If S is the only group in society, then as we saw above it would have a correctly specified model of how output is generated. Along with our assumption of small implementation noise, S will then learn the true parameter values on the policies it considers. However, as power shifts are inevitable, the learning of S is substantially different as C is also in power, implying an omitted variable bias which as we show below, takes a specific form.

To understand the result of colinearity and extremism, suppose first that upon convergence, the actions of S are *not* colinear with those of C. This implies, from (2.1), that the beliefs of S will evolve to fully explain the average output at each regime. Specifically, if optimal actions are not colinear, a solution to (2.1) can be achieved by solving the following two equations:

$$\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_c^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*], \ \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*] \tag{2.8}$$

These two equalities are linearly independent due to the fact that the policies are not colinear, and thus a solution exists. On an intuitive level, without colinearity there is enough variation in the data so that S will be able to correctly conjecture the average output delivered by each regime. But C, having the correct model, will also learn how to do this. This leads to a contradiction to the equal intensity condition. Specifically, given that the actions above are also optimal, S will now have greater intensity than C:

$$\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] > \bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_c^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*] = \bar{y}[\bar{\boldsymbol{\beta}}_c, \mathbf{x}_c^*]. \tag{2.9}$$

But this is in contradiction to our cycles result, which demands equal intensity. In other words, S cannot learn too much: Policies must be colinear to limit the learning of S and specifically its ability to predict expected output at each regime. Instead, the beliefs of S will evolve to balance its prediction mistakes *across* the two regimes.

We can now fully derive the beliefs of S and show that S must hold more extreme beliefs than those of C. To see this, remember that in the long run the two types have equal intensity, i.e.,

$$\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*] = \bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*] \tag{2.10}$$

which by Lemma 1 implies that:

$$\bar{\boldsymbol{\beta}}_s' \bar{\boldsymbol{\beta}}_s = \boldsymbol{\beta}' \boldsymbol{\beta} \tag{2.11}$$

The colinearity result implies that $\bar{\boldsymbol{\beta}}_s = \tau \boldsymbol{\beta}_s$ for some τ. Plugging this into (2.11) pins down the degree of colinearity τ^* :

$$(\tau^*)^2 (\boldsymbol{\beta}_s' \boldsymbol{\beta}_s) = \boldsymbol{\beta}' \boldsymbol{\beta} \Rightarrow \tau^* = \sqrt{(\boldsymbol{\beta}' \boldsymbol{\beta}) / (\boldsymbol{\beta}_s' \boldsymbol{\beta}_s)} > 1 \Rightarrow \tag{2.12}$$

$$\bar{\boldsymbol{\beta}}_s = \sqrt{(\boldsymbol{\beta}' \boldsymbol{\beta}) / (\boldsymbol{\beta}_s' \boldsymbol{\beta}_s)} \boldsymbol{\beta}_s$$

Thus, S is more bold in its policy prescriptions and so our model implies that simplicity implies extremism.

D. Dynamics of Power Shifts

We now explore the comparative statics of the political cycles and how the true data generating process affects the cycle dynamics. First, we solve for the limit share of time that each group is in power. To solve for θ_s, we plug the expression for $\bar{\boldsymbol{\beta}}_s$ from (2.12) in the OLS condition (2.1), where $\bar{\boldsymbol{\beta}}_s$ is also required to explain mistakes across the two regimes. Noting that $\bar{\boldsymbol{\beta}}_c = \boldsymbol{\beta}$, we then get:

$$\theta_s = 1/(1 + \tau^*), \tag{2.13}$$

where it is easy to see that θ_s is lower when τ^* is higher. The colinearity parameter measures the relative importance of the parameters *not* considered by S. Therefore we have:

Observation 1: *The more important are the policy variables that S ignores, the more extreme are S's beliefs, and the less time it spends in power.*

Intuitively, to generate more extreme beliefs in equilibrium, S needs to suffer from a higher omitted variable bias, which arises when C is in power more often. Thus, political cycles must result in just enough omitted variable bias to equate intensity. This implies that if S is extremely wrong to ignore some policies, so that τ^* is very large, then it spends very little time in power, but when it does, its policies are very biased. Alternatively when S is almost correct and τ^* is close to one, then θ_s is close to a half.

The deterministic average power sharing, θ_s, captures the systematic changes of power. This results from the fact that the beliefs of S must balance the "mistakes" in its predictions across the two regimes. Thus, S is continuously "surprised" by its prediction for the average output for *each* regime. As in Figure 1, when they are in power, S is surprised that its policies have underperformed, as the true average outcome, $\bar{y}[\boldsymbol{\beta}, \mathbf{x}_s^*]$, is strictly lower than their expected outcome, $\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_s^*]$. When C is in power, group S is surprised by the (wrongly attributed) success of her narrow set of relevant policies, $\bar{y}[\boldsymbol{\beta}, \mathbf{x}_c^*]$, as compared to what she expected, $\bar{y}[\bar{\boldsymbol{\beta}}_s, \mathbf{x}_c^*]$. Thus the change in the beliefs of S contains a systematic component: A gradual increase in bias and intensity when C is in power and a gradual reduction of the bias when they are in power.

But beyond these two systemic changes in power, another source for power shifts arises from the random shocks ε. Indeed a common thread in the literature on populism is how economic shocks are more likely to lead frustrated voters to support populist movements. We now analyse how the random shocks ε affect power switches. We focus on outcomes far enough into the process so that beliefs are close to their limiting values:

Proposition 1: *A negative (positive) ε shock to y affects the intensity of the incumbent party relatively more than that of the opposition. Thus a negative (positive) ε shock to y hastens (delays) a regime change.*[18]

Random shocks change estimates of the effectiveness of policy, but these effects are stronger for the incumbent party which is implementing its desired policy combination. Specifically, when the simple group is in power, a negative shock reduces their intensity, as their belief in the effectiveness of the policies they deem relevant falls. Complex beliefs in these same policies also fall, but the complex beliefs in the efficacy of policies the simple deem irrelevant,

[18]The proof of Proposition 1 is in the Online Appendix, Section II.

and hence do not implement, rises, as the poor outcome under simple rule convinces the complex that these neglected policies are more effective than previously thought. These two effects offset each other, and complex intensity remains constant.

When the complex are in power, a negative shock reduces the belief in the effectiveness of policies of both types, but the effects on intensity are greater for the complex, for whom intensity depends upon a wider range of policies, all of which are seen to be failing. In sum, a negative shock hastens the transfer of power, with positive shocks having the opposite effect.

Simulation: We now simulate the dynamic process outlined in Example 1, where we describe the mean results for the following parameter values: $\beta_1 = \beta_2 = 1$, $R = 1$, and the policy shock n is normally distributed with $\sigma_n^2 = 0.01$. As a result, $\tau^* = \sqrt{2}$ and thus $\theta_s \approx 0.408$.[19] As for the relative share of time in power, this satisfies $\frac{\theta_c}{\theta_s} \approx 1.45 \approx \tau^*$.

Simulations are run for 10 million periods, and the reported results are for the last 1 million periods. We vary the variance of the outcome shock ε, denoted below by σ^2, increasing it from close to zero to be sufficiently large.

We report in Table 1 on the following long term statistics: First, μ_s denotes the mean number of periods of an episode of simple rule, and μ_c denote the mean number of periods of an episode of complex rule. Second, we denote by π_i the fraction of transitions from i to j which involve a negative outcome shock.

	$\sigma^2 = .01$	$\sigma^2 = .1$	$\sigma^2 = 1$	$\sigma^2 = 10$	$\sigma^2 = 100$
θ_s	.409	.408	.409	.408	.408
μ_s	1.09	1.41	2.89	7.90	25.5
μ_c	1.58	2.04	4.17	11.48	37.0
π_s	.526	.662	.908	.987	.999
π_c	.549	.660	.887	.984	.997
μ_c/μ_s	1.45	1.45	1.44	1.45	1.45

Table 1: Simulation of long term transition of power for varying variance of ε.

As can be seen in the table, μ_s and μ_c increase with the variance of ε, where $\mu_c \approx \tau^* \mu_s$. Also, for both types, the fraction of transitions of power π that involve a negative shock increases, from 0.5 to 1. This confirms our analytical results reported in Proposition 1, showing that a negative shock hastens transition of power.

[19] Equation (2.13) reports θ_s for the case where $\sigma_n^2 = 0$. The general formulation as we show in the Appendix satisfies:

$$\theta_s = \frac{1 - \tau^* \frac{\sigma_n^2}{R}}{1 + \tau^*}.$$

The simulation results illustrate the interplay between the systemic components of power dynamics, which are derived from the equal intensity and colinearity conditions, and those determined directly by the noise ε. The larger is the variance of ε, the more likely are paths in which beliefs wander further away from equal intensity. This lengthens the stay in power of incumbents, as one good shock allows them to be in power for a longer time (noting that future shocks have mean zero). Moreover, the shock is dominating the variation of intensity. Being far away from equal intensity implies that the systematic component cannot easily shift beliefs across the equal intensity point, but a big negative shock will do so.

E. Convergence

In general, establishing convergence with misspecified models is problematic even with exogenous iid data (see Berk 1966). Having endogenous data, as we have in our model, introduces more challenges as observations are non iid. As we mentioned in the introduction, substantial progress has been made in the literature analyzing the convergence properties of misspecified models with non iid data.[20] But with respect to this literature, our model is further complicated by having multiple players, continuous actions, and a multidimensional state space, assumptions which are all important for the derivation of a meaningful colinearity and extremism result.

A first step to prove convergence in this environment of non iid data is to establish a law of large numbers for our framework. This is done in Section (a) in the Appendix. To do so, we rely on the fact that at period t, the regressors $\mathbf{x}_t$ and the shock ε_t are independent of each other. While the regressors depend on past realizations of the shock, they are not correlated with the current one.[21]

A second issue, and one that arises with multiple dimensions of policies, is misidentification of parameter values. This arises even for rational players, as a continuum of parameter values can explain long run average output.[22] In our model, absent the policy noise $\mathbf{n}$, C's limit beliefs are undetermined. Whatever $C's$ long-term beliefs are, the equilibrium characteristics will always involve cycles and extremism as identified above. However, such multiplicity of possible beliefs introduces additional challenges for establishing convergence as it is hard to rule out that groups' beliefs perpetually "travel" along this continuum of limit beliefs. The first role of the policy noise, $\mathbf{n}$, is then to allow us to establish that the beliefs of C, who has the correct model, converges to the true parameter values. Together with the law of large numbers, it is straightforward to show this when allowing for policy noise, which we do in

[20] See for example Esponda, Pouzo and Yamamaoto (2021) and Frick, Iijima and Ishii (2020).

[21] These types of law of large numbers for non iid data are established in other papers in this literature, using similar methods albeit in different environment as we focus on OLS. See for example Heidhues et al (2018).

[22] Specifically for C, these are all the $\bar{\boldsymbol{\beta}}_c$ vectors that satisfy $\boldsymbol{\beta}'\mathbf{x}_c^* = \bar{\boldsymbol{\beta}}_c' x_c^*$ and $\boldsymbol{\beta}'\mathbf{x}_s^* = \bar{\boldsymbol{\beta}}_c' x_s^*$.

Section (b) in the Appendix.

The third issue involves the combination of our multiple dimensions of policies and continuous actions. Esponda et al (2021) use stochastic approximation methods to characterise conditions for stability of steady state. These conditions rely on a discrete set of actions and are moreover easier to check and apply when a unidimensional state space is involved. Heidhues et al (2021) analyze convergence using stochastic approximation techniques in a model with continuous actions but a unidimensional state space, using a normal prior and noise. To see why continuous actions and multidimensionality create difficulties in proving convergence, consider the situation where beliefs move slightly away from a steady state. If actions are discrete, then by continuity of utilities, there is a range of beliefs for which the same action is taken and so one can check whether beliefs return to steady state level or move further away to establish stability. With continuous actions, when beliefs change slightly, actions change too, changing beliefs further and so on. On top of this, the multidimensionality of the belief space poses an additional difficulty in pinning down the direction of change of these beliefs.

We now sketch how our method of proof deals with those issues. After establishing the convergence of $C's$ beliefs to the true parameter values, we take these as given (so, we are looking far enough into the learning process). We then prove the cycles result, namely, that the ratio of the shares of time that each group is in power must be bounded strictly between 0 and 1 (Section b(i) in the Appendix). Together with the OLS formulation, this allows us to derive a deterministic law of motion for the asymptotic beliefs of S.

We can then establish the following monotonicity properties of this law of motion. When C is in power, and similar to the dynamics in Example 1, the beliefs of S on all parameter values include an omitted variable bias and are then pushed to be more extreme. But in addition, the policy noise $\mathbf{n}$ pushes these values to be colinear with the true parameters: This noise, similar to small experimentation, puts some discipline on how beliefs move and this discipline implies movement towards the true relative values of the parameters that S considers. Alternatively, when S is in power, and similar to the dynamics in Example 1, there is no omitted variable bias. As long as policies are not colinear, there is also sufficient variation in policies. This, along with the policy noise $\mathbf{n}$, pushes the beliefs of S towards the true values (and hence also colinear beliefs), as S has the correct model when they are in power.

The discipline imposed by the policy noise (and policy variation) on the changes in beliefs allows us to establish that the beliefs of S must converge to the colinear multiplication of the true value on the dimensions they consider. Our cycle dynamics imply that the beliefs must also converge to satisfy equal intensity, and so the fraction of time C are in power (implying an omitted variable bias for S) fluctuates up and down until asymptotic power sharing results in just enough bias to reach equal intensity. Equal intensity and colinearity pin then a unique belief to which S's beliefs must converge. The law of motion of S beliefs and the convergence to a unique value are established in Sections b(ii)-(iv) of the Appendix.

III. Extensions

In this Section we present some additional results and discuss alternative modelling assumptions.

A. Overlapping versus non-overlapping world views

Wrong complex world view. Above we considered an environment in which the beliefs of the complex are correctly specified, in that they include all relevant policies, whereas the simple type erroneously exclude a subset of these. The fact that the complex consider all relevant policies does not matter at all, and it is sufficient for our results that the simple consider a subset of the relevant policies that the complex consider.

Overlapping world views. One can consider the model in which the two groups have overlap in the policies that both consider relevant, but that each group considers in addition an exclusive set of additional relevant policies. We can show that if beliefs converge, they can converge to a case in which one of the groups "becomes" a simple group in the sense that it abandons its exclusive relevant dimensions. In this case our results about political cycles as well as the extremism of the simple group's policies still hold. However, we do not have convergence results for this type of environment.

Non-overlapping world views. Another possibility is that there is no overlap between the world views of the groups, so that each group considers a mutually exclusive set of relevant policies. In these cases we can have one group perpetually in power. For example, if one group has a low prior on how its policies affect output, then the other group can potentially be in power in the long term.

Irrelevant policies. In the Appendix we consider an additional extension in which the relation between the two groups' world views is the same (simple versus complex) when it comes to relevant policies but both types can also consider policies which are irrelevant. We show that all the results reported in the paper hold for this more general model. Our assumption of noise implies that both groups abandon the non-relevant variables in the long term and hence the asymptotic equilibrium looks exactly like the one in our basic model.

B. Strategic Politicians

We use a simple political model in which intensity of preferences is the key to electoral success. We adopt a citizen candidate model so that politicians choose policies myopically, and offer voters exactly their ideal polices. There are many ways to justify the assumption of myopia in politics as politicians have limited terms, policy choices have high stakes and it is very complicated to predict the influence of current policies on future behaviors. We now discuss two alternative assumptions.

Strategically affecting the beliefs of the other group: One element of our model is that we assume that the winning politician implements her myopic ideal policy. That is: (i) she does not engage in experimentation in order to enhance her learning; (ii) she does not use today's policy choice to manipulate future learning and actions of others.

With regards to (i), our assumption of policy implementation noise captures some form of experimentation. Indeed, this feature of the model is the reason why the complex end up converging on the true parameters of their model.

More sophisticated forward-looking strategic behavior along the lines of (ii) might alter some of our positive results but not the qualitative effect of the simple group's influence on policy outcomes. While this is beyond the scope of our analysis, we conjecture that even in such a model, in the long term, the simple group's misspecified model will affect policies. Specifically, it is not possible for the complex group to be perpetually in power implementing their ideal policy, as in such a case the simple's estimates must converge to induce them to have higher intensity. As a result, even if the complex converge to be in power perpetually, they must implement long-term policies that prevent the simple from obtaining higher intensity; such policies have therefore to be biased.

Office-motivated politicians: One may imagine other models of political competition, e.g., a probabilistic voting model with office motivated politicians, which essentially implies that politicians choose policies to maximize average welfare. While this would yield different policies as well as learning patterns, a key feature of our analysis will remain: In equilibria, policies will cater to group S to some degree. That is, the omitted variable bias in $S's$ beliefs would mean that they would prefer stronger policies on the policies they deem effective. Any policy that maximizes welfare will then exhibit such a bias.

C. Endogenous resources and other utility functions

In our model we have assumed a fixed resource constraint R. We can extend the model to allow the different types to endogenously choose their desired level of resources. In particular, we can assume that the utility function of citizens is given by:

$$U_t = y_t + V(R_t),$$

where as before $R_t = \mathbf{x}_t'\mathbf{x}_t$ represents the resources used in implementing policy $\mathbf{x}_t$ for y_t, while V represents the utility derived from policy outcomes over which there is no disagreement regarding causal mechanisms. V is a reduced form, representing the utility that can be achieved in other policy areas given the allocation of resources to y_t, and the assumptions

$V' < 0$ and $V'' < 0$ are natural. To derive analytical results, we work with a second-order approximation of V as a quadratic function of R_t. We can then show that intensity of preferences is also an increasing function of the magnitude of beliefs. Assuming that R_t is bounded from above, we can then extend all our convergence results accordingly.

We also assume a simple utility function that is linear in y, which implies that utility is a function of mean beliefs only. For more general utilities the whole distribution of beliefs would matter. Montiel Olea et al (2019) show that in a model with exogenous data, complex models (which abide with the truth) would induce lower variance of their beliefs when data is sufficiently large. This would imply an advantage to the complex group. Thus, our results hold as long as individuals are not too risk averse.

D. Other data generating processes

We assumed that the true data generating process is a simple linear function of the policy variables. This is not important for our key results of cycles and extremism. Whatever is the true data generating process, when one group is in power for a sufficiently long time, then the beliefs of both groups evolve to explain the average output produced by the incumbent, which then implies that the group in opposition becomes more intense. Similarly, colinearity of beliefs and policies must arise given the same logic above, as otherwise with sufficient variation both groups will be able to agree on how to produce the average output of each regime, contradicting equal intensity.

Another possible change in the data generating process is to add an unknown constant term. In the baseline model, the simple bias due to the collinearity of asymptotically constant complex policies can be completely absorbed into an unknown constant term and policy cycles do not arise. However, if there is any variation in the resources available due to varying outside opportunity costs, the asymptotic policies implemented by the complex will fluctuate in tandem with the resources R_t, and all of results concerning cycles and evolving simple bias described above are re-established. This is easily modelled through time varying fluctuations in $V_t(R_t)$ as in section III.C above, but adds little to our story, while greatly complicating the algebra.

E. Relation to Berk-Nash equilibrium

To conclude the discussion, we examine the relation between our results above and a static notion of equilibrium in the spirit of Berk-Nash equilibrium (Esponda and Pouzo 2016). A Berk-Nash equilibrium is a static solution concept for a dynamic game of players with misspecified models where actions are optimal given beliefs and beliefs rationalize the observed output which arises given the actions played. Berk (1966) shows for the case of iid data that beliefs stemming from a misspecified model will concentrate on those that minimize the Kullback–Leibler (KL) distance to the true beliefs.[23] Using this notion of minimizing the KL distance, Esponda and Pouzo (2016) define a Berk-Nash equilibrium.

In the Online Appendix (Section III), we analyse the set of Berk-Nash equilibria in our model. We then show, using similar intuition to the one provided in Section II.A, that the unique Berk-Nash equilibrium is analogous to the dynamic one identified in Theorem 1 (See Proposition III.1 in the Online Appendix). Specifically, our dynamic cycles result is translated in the static solution to a power sharing probability between the two groups.[24]

IV. Conclusion

Our analysis has shown how simplistic beliefs can persist in political competition against a more accurate and complex view of the world, delivering sub-par outcomes on each outing in power and yet returning to dominate the political landscape over and over again. In the framework presented above simplistic beliefs arise as a consequence of a primitive assumption of misspecification, but we recognize that there are deeper questions to explore. A recent examination of European Social Survey data by Guiso et al (2017) finds that the responsiveness of the electorate to populist ideas and the supply of populist politicians increases in periods of economic insecurity. Social and economic transformation, and the insecurity and inequality it can engender, may create environments in which opportunistic politicians are able to plant erroneously simplistic world views into the electorate. Linking belief formation, at its most fundamental level, to ongoing economic and political events allows a richer characterization of political cycles, and may be a good avenue for future work.

[23] Intuitively, minimising the Kullback–Leibler distance is similar to maximising the likelihood of previous observations.

[24] For a related result in a static model, see Eliaz and Spiegler (2019).

APPENDIX: PROOF OF CONVERGENCE IN A GENERALIZED MODEL (GENERALIZATION OF THEOREM 1)

In this appendix we prove almost sure convergence to the beliefs and share of time each type is in power given in the paper in a generalized framework. Specifically, while in the paper all k potential policies were relevant (i.e., had non-zero effects), in this appendix we allow that some may be irrelevant and have zero effects. While the beliefs of "complex" types are correctly specified, in that they include all relevant policies, "simple" types erroneously exclude a subset of these. The prior beliefs of both types may include some irrelevant policies that have zero effects, and we impose no a priori restriction on the relative number of policies, k_s and k_c, each type believes may be relevant, other than that their union covers the set of k policies that are systematically implemented. The monikers "complex" and "simple" derive from the fact that the endogenous asymptotic equilibrium looks much like that assumed in the paper, where the non-zero beliefs of the complex are broader than those of the simple.

We begin by establishing notation. Let $\mathbf{y} = \mathbf{H}\boldsymbol{\beta} + \boldsymbol{\varepsilon}$, where $\mathbf{y}$ is the t x 1 history of outcomes, $\mathbf{H} = \mathbf{X} + \mathbf{N}$ the t x k history of observed policy and noise, $\boldsymbol{\beta}$ the unknown k x 1 vector of parameters, and $\boldsymbol{\varepsilon}$ the unobserved t x 1 vector of outcome shocks. $\mathbf{H}_i$ and $\mathbf{H}_{\sim i}$ denote the k_i and $k_{\sim i}$ columns of $\mathbf{H}$ deemed relevant and irrelevant by type i and $\boldsymbol{\beta}_i$ and $\boldsymbol{\beta}_{\sim i}$ the corresponding parameters.[23] $\mathbf{H}_{ij}$ and $\mathbf{H}_{\sim ij}$ are the rows of $\mathbf{H}_i$ and $\mathbf{H}_{\sim i}$ associated with the t_j periods when type j is in power, with $t_i + t_j = t$. We use the notation $\mathbf{H}_{\bullet j}$ to denote the t_j x k history of all policies during the periods type j is in power. $\mathbf{I}_k$ and $\mathbf{0}_{n\,x\,m}$ denote the identity matrix and matrix of zeros of the indicated dimensions. We assume that the rows and columns of $\mathbf{N}$ are independently and identically distributed with mean $\mathbf{0}_{k\,x\,1}$, covariance matrix $\sigma_n^2\mathbf{I}_k$ and bounded fourth moments, while the rows of $\boldsymbol{\varepsilon}$ are independently and identically normally distributed with mean zero and variance σ^2.

(a) *Preliminaries: Standard Matrix Algebra Results & Important Lemmas*

A symmetric positive definite matrix $\mathbf{V}$ allows the spectral decomposition $\mathbf{E\Lambda E}'$, where $\boldsymbol{\Lambda}$ is the diagonal matrix of strictly positive eigenvalues and $\mathbf{E}$ is a matrix whose columns are the corresponding mutually orthogonal eigenvectors, with $\mathbf{EE}' = \mathbf{E}'\mathbf{E} = \mathbf{I}$. $\mathbf{V}^{-1} = \mathbf{E\Lambda}^{-1}\mathbf{E}'$, i.e., the inverse of $\mathbf{V}$ has the same eigenvectors as $\mathbf{V}$ and eigenvalues equal to the inverse of those of $\mathbf{V}$. We can also define $\mathbf{V}^{-½} = \mathbf{E\Lambda}^{-½}\mathbf{E}'$ as $\mathbf{V}^{-½}\mathbf{V}^{-½} = \mathbf{E\Lambda}^{-½}\mathbf{E}'\mathbf{E\Lambda}^{-½}\mathbf{E}' = \mathbf{E\Lambda}^{-1}\mathbf{E}'$. In a similar

[23] In order to simplify discussion of matrix inverses and eigenvalues, we depart from the practice used in the text above of inserting 0s for policies that a type believes are irrelevant, and instead modify the dimensions of matrices and vectors to cover only the policies that a type deems relevant. Thus, for example, $\boldsymbol{\beta}_i$ is k_i x 1 not k x 1 with $k_{\sim i}$ 0s inserted.

spirit, $\mathbf{V}^{-2} = \mathbf{V}^{-1}\mathbf{V}^{-1}$ has eigenvalues equal to the square of those of $\mathbf{V}^{-1}$ and the same eigenvectors. For a rank one update of $\mathbf{V}$ using the vector $\mathbf{x}$, the Sherman-Morrison formula tells us that $(\mathbf{V}+\mathbf{xx}')^{-1} = \mathbf{V}^{-1} - \mathbf{V}^{-1}\mathbf{xx}'\mathbf{V}^{-1}/(1+\mathbf{x}'\mathbf{V}^{-1}\mathbf{x})$, while the eigenvalues of the matrix $(\mathbf{V} + c\mathbf{I})$, with c a constant, are given by $\mathbf{\Lambda} + c\mathbf{I}$, and the eigenvectors are the same as those of $\mathbf{V}$. The eigenvalues of $\mathbf{V}$ are all weakly increasing following a rank-one update (Golub 1973), so if $\mathbf{V}$ is initially positive definite (with strictly positive eigenvalues) it remains so following a sequence of rank-one updates. The maximum across all possible vectors $\mathbf{x}$ of the Rayleigh quotient $\mathbf{x}'\mathbf{V}\mathbf{x}/\mathbf{x}'\mathbf{x}$ is the maximum eigenvalue of $\mathbf{V}$, which we denote with $\lambda_{\max}(\mathbf{V})$, with $\lambda_{\min}(\mathbf{V})$ denoting the minimum eigenvalue.

The following two lemmas are used repeatedly in our proofs:

(A1) $$\text{(a) } \frac{\mathbf{N}'\boldsymbol{\varepsilon}}{t} \xrightarrow{a.s.} \mathbf{0}_{kx1};\ \text{(b) } \frac{\mathbf{N}'\mathbf{N}}{t} \xrightarrow{a.s.} \sigma_n^2\mathbf{I}_k;\ \text{(c) } \frac{\mathbf{X}'\boldsymbol{\varepsilon}}{t} \xrightarrow{a.s.} \mathbf{0}_{kx1};\ \text{(d) } \frac{\mathbf{X}'\mathbf{N}}{t} \xrightarrow{a.s.} \mathbf{0}_{kxk}.$$

(A2) $$\boldsymbol{\varepsilon}'\mathbf{H}_i'(\mathbf{H}_i'\mathbf{H}_i)^{-1}(\mathbf{H}_i'\mathbf{H}_i)^{-1}\mathbf{H}_i\boldsymbol{\varepsilon} \xrightarrow{a.s.} 0.$$

(A1a) and (A1b) follow immediately from the strong law of large numbers, as the product of two iid random variables is an iid random variable in its own right and hence its average converges almost surely to the expectation.

The i^{th} term of (A1c) or i x j^{th} term of (A1d) can be expressed as

(A3) $$v_t = \sum_{s=1}^{t} \frac{x_{is}\eta_s}{t},$$

where η_s is either ε_s or n_{js}. We use μ_i to denote the i^{th} moment of η_s, with $\mu_1 = 0$ and μ_2 and μ_4 bounded following from the assumptions articulated earlier above. As η_s is independent of contemporaneous policy and past shocks and policy, applying the law of iterated expectations (i.e., taking the expectation at time 0 of the expectation at time 1 of the expectation at time 2 ...) one sees that

(A4) $$E\left[\sum_{s=1}^{t} x_{is}\eta_s\right] = 0,\quad E\left[\left(\sum_{s=1}^{t} x_{is}\eta_s\right)^2\right] = \sum_{s=1}^{t} E(x_{is}^2)\mu_2 \le tR\mu_2,$$

$$E\left[\left(\sum_{s=1}^{t} x_{is}\eta_s\right)^4\right] = 2\sum_{r=1}^{t}\sum_{s=r+1}^{t} E(x_{ir}x_{is})\mu_2^2 + \sum_{s=1}^{t} E(x_{is}^4)\mu_4 \le R^2[t(t-1)\mu_2^2 + t\mu_4]$$

where we use the fact that x_{is}^2 is bounded by the total resources devoted to policy (R) and that bounded 4$^{\text{th}}$ moments, by Jensen's Inequality, imply bounded 3$^{\text{rd}}$ moments, so that we can comfortably say that terms involving $\mu_3\mu_1$ equal 0.

From the Borel-Cantelli lemma, we know that if for all $a > 0$

(A5) $$\sum_{t=1}^{\infty} P(|v_t| > a) < \infty$$

then v_t converges almost surely to 0, as the probability v_t deviates by more than a from 0 an infinite number of times is zero. Recalling Markov's Inequality, that for a non-negative random variable x and $a > 0$, $P(x \geq a) \leq E(x)/a$, we see that

(A6) $$P(v_t^4 \geq a^4) \leq \frac{E(v_t^4)}{a^4} \Rightarrow P(|v_t| > a) \leq \frac{E(v_t^4)}{a^4} \leq \frac{R^2}{a^4}\left[\frac{(t-1)\mu_2^2 + \mu_4}{t^3}\right]$$

Since the sum from t equals 1 to infinity of the last expression is finite, we see that v_t converges almost surely to 0, establishing (A1c) and (A1d).

Turning to (A2), we begin by noting that

(A7) $$\frac{\boldsymbol{\varepsilon}'\mathbf{H}_i}{t}\left[\frac{\mathbf{H}_i'\mathbf{H}_i}{t}\right]^{-1}\left[\frac{\mathbf{H}_i'\mathbf{H}_i}{t}\right]^{-1}\frac{\mathbf{H}_i'\boldsymbol{\varepsilon}}{t} \leq \frac{\boldsymbol{\varepsilon}'\mathbf{H}_i}{t}\frac{\mathbf{H}_i'\boldsymbol{\varepsilon}}{t}\lambda_{\max}\left(\left[\frac{\mathbf{H}_i'\mathbf{H}_i}{t}\right]^{-1}\left[\frac{\mathbf{H}_i'\mathbf{H}_i}{t}\right]^{-1}\right)$$

$$= \frac{\boldsymbol{\varepsilon}'\mathbf{H}_i}{t}\frac{\mathbf{H}_i'\boldsymbol{\varepsilon}}{t}\lambda_{\min}\left(\frac{\mathbf{H}_i'\mathbf{H}_i}{t}\right)^{-2} \leq \frac{\boldsymbol{\varepsilon}'\mathbf{H}_i}{t}\frac{\mathbf{H}_i'\boldsymbol{\varepsilon}}{t}\lambda_{\min}\left(\frac{\mathbf{H}_i'\mathbf{H}_i - \mathbf{X}_i'\mathbf{X}_i}{t}\right)^{-2}$$

where in the first inequality we use the properties of the Rayleigh quotient, in the following equality the relation between the eigenvalues of matrix products and inverses, and in the final inequality the fact that in the t rank one updates of matrix $\mathbf{H}_i'\mathbf{H}_i - \mathbf{X}_i'\mathbf{X}_i$ to $\mathbf{H}_i'\mathbf{H}_i$ the eigenvalues are always weakly increasing. Noting that $\mathbf{H}_i'\mathbf{H}_i - \mathbf{X}_i'\mathbf{X}_i = \mathbf{X}_i'\mathbf{N}_i + \mathbf{N}_i'\mathbf{X}_i + \mathbf{N}_i'\mathbf{N}_i$ and applying (A1), we see that

(A8) $$\frac{\boldsymbol{\varepsilon}'\mathbf{H}_i}{t}\frac{\mathbf{H}_i'\boldsymbol{\varepsilon}}{t}\lambda_{\min}\left(\frac{\mathbf{H}_i'\mathbf{H}_i - \mathbf{X}_i'\mathbf{X}_i}{t}\right)^{-2} \overset{a.s.}{\rightarrow} \mathbf{0}_{kx1}'\mathbf{0}_{kx1}\lambda_{\min}(\sigma_n^2\mathbf{I}_k)^{-2} = \frac{0}{(\sigma_n^2)^2} = 0.$$

Since $\boldsymbol{\varepsilon}'\mathbf{H}_i(\mathbf{H}_i'\mathbf{H}_i)^{-1}(\mathbf{H}_i'\mathbf{H}_i)^{-1}\mathbf{H}_i'\boldsymbol{\varepsilon}$ is a non-negative random variable bounded from above by a random variable which almost surely converges to zero, it follows that (A2) is true.

Standard econometric proofs start off by assuming that $(\mathbf{H}_i'\mathbf{H}_i/t)^{-1}$ converges almost surely or in probability to a positive definite matrix, arguing that $\mathbf{H}_i'\boldsymbol{\varepsilon}/t$ similarly converges to a vector of 0s, and then drawing conclusions about the convergence of $(\mathbf{H}_i'\mathbf{H}_i)^{-1}\mathbf{H}_i'\boldsymbol{\varepsilon}$. In our case, since the regressors are endogenous, we cannot make a priori assumptions about whether $(\mathbf{H}_i'\mathbf{H}_i/t)^{-1}$ converges. However, as (A2) shows, a quadratic form based upon $(\mathbf{H}_i'\mathbf{H}_i/t)^{-1}$ is easily shown to be bounded and to converge almost surely to zero provided there is minimal noise. In the proofs below we make use of such quadratic forms to prove that beliefs and other objects of interest converge.

(b) *Convergence in the Generalized Model*

We assume that prior beliefs for each type across the policies they believe are relevant are normally distributed with mean $\bar{\boldsymbol{\beta}}_{i0}$ and joint covariance matrix $\sigma_{i0}^2\mathbf{V}_{i0}^{-1}$, while the prior probability density function on σ_{i0} is inverted gamma. Following the observation of the t x 1 history of outcomes $\mathbf{y}$ and t x k_i history of implemented policy deemed relevant by i, $\mathbf{H}_i$, such beliefs give rise to mean posterior beliefs[24]

(A9) $$\bar{\boldsymbol{\beta}}_i = (\mathbf{V}_{i0} + \mathbf{H}_i'\mathbf{H}_i)^{-1}(\mathbf{V}_{i0}\bar{\boldsymbol{\beta}}_{i0} + \mathbf{H}_i'\mathbf{y}).$$

However, since one can easily define a finite "pre-history" of policy $\mathbf{H}_{i0}$ and outcomes $\mathbf{y}_0$ such that $\mathbf{V}_{i0} = \mathbf{H}_{i0}'\mathbf{H}_{i0}$ and $\bar{\boldsymbol{\beta}}_{i0} = (\mathbf{H}_{i0}'\mathbf{H}_{i0})^{-1}\mathbf{H}_{i0}'\mathbf{y}_0$, and our results will be asymptotic, we simplify algebra by including these pre-histories in $\mathbf{H}$ and $\mathbf{y}$ and simply writing beliefs as

(A10) $$\bar{\boldsymbol{\beta}}_i = (\mathbf{H}_i'\mathbf{H}_i)^{-1}\mathbf{H}_i'\mathbf{y}.$$

The complex's model incorporates the effects of all policies whose effects are non-zero and their mean beliefs are given by

(A11) $$\bar{\boldsymbol{\beta}}_c = (\mathbf{H}_c'\mathbf{H}_c)^{-1}\mathbf{H}_c'\mathbf{y} = (\mathbf{H}_c'\mathbf{H}_c)^{-1}\mathbf{H}_c'(\mathbf{H}_c\boldsymbol{\beta}_c + \mathbf{H}_{\sim c}\boldsymbol{\beta}_{\sim c} + \boldsymbol{\varepsilon}) = \boldsymbol{\beta}_c + (\mathbf{H}_c'\mathbf{H}_c)^{-1}\mathbf{H}_c'\boldsymbol{\varepsilon},$$

$$\text{so } (\bar{\boldsymbol{\beta}}_c - \boldsymbol{\beta}_c)'(\bar{\boldsymbol{\beta}}_c - \boldsymbol{\beta}_c) = \boldsymbol{\varepsilon}'\mathbf{H}_c(\mathbf{H}_c'\mathbf{H}_c)^{-1}(\mathbf{H}_c'\mathbf{H}_c)^{-1}\mathbf{H}_c'\boldsymbol{\varepsilon} \overset{a.s.}{\to} 0,$$

where the first line uses the fact that all elements of $\boldsymbol{\beta}_{\sim c}$ are zero and the second line follows from Lemma (A2) above. Consequently, we know that the beliefs of the complex converge almost surely to the true parameter values

(A12) $$\bar{\boldsymbol{\beta}}_c \overset{a.s.}{\to} \boldsymbol{\beta}_c,$$

and in the limit the complex almost surely implement policies

(A13) $$\boldsymbol{\beta}\sqrt{R/\boldsymbol{\beta}'\boldsymbol{\beta}}$$

where we use the fact that since the elements of $\boldsymbol{\beta}_{\sim c}$ are all zero we can express complex policies in the areas they believe are irrelevant in terms of these parameters as well. The remainder of this appendix is devoted to proving that simple beliefs $\bar{\boldsymbol{\beta}}_s$ converge to the steady state values $\tau^*\boldsymbol{\beta}_s$, where $\tau^* = \sqrt{\boldsymbol{\beta}'\boldsymbol{\beta}/\boldsymbol{\beta}_s'\boldsymbol{\beta}_s}$. We note that τ^* is strictly greater than 1, as we assume that simple beliefs are misspecified, so $\boldsymbol{\beta}_{\sim s} \neq \mathbf{0}_{k_{\sim s}x1}$.

The simple's mean beliefs are the coefficient estimates in the misspecified regression

(A14) $$\bar{\boldsymbol{\beta}}_s = (\mathbf{H}_s'\mathbf{H}_s)^{-1}\mathbf{H}_s'\mathbf{y} = (\mathbf{H}_s'\mathbf{H}_s)^{-1}\mathbf{H}_s'\mathbf{H}\boldsymbol{\beta} + (\mathbf{H}_s'\mathbf{H}_s)^{-1}\mathbf{H}_s'\boldsymbol{\varepsilon},$$

[24] This is a standard OLS Bayesian result (Zellner 1971). Our model is somewhat different than the standard framework in that the regressors are determined by past realizations of the error term. However, since each period's disturbance ε is independent of the row of regressors $\mathbf{h}_i$, provided some initial prior exists the period by period recursive application of the updating formula (B1) aggregates across t periods to the result given above.

so with a similar use of Lemma (A2) we have

(A15) $$\overline{\boldsymbol{\beta}}_s - (\mathbf{H}_s'\mathbf{H}_s)^{-1}\mathbf{H}_s'\mathbf{H}\boldsymbol{\beta} \overset{a.s.}{\rightarrow} 0.$$

Recalling that we use the notation $\mathbf{H}_{\bullet j}$ and $\mathbf{H}_{ij}$ to denote the history of all policies and only the policies type i deems relevant, respectively, during the t_j periods type j is in power, with $\mathbf{H}_s'\mathbf{H}_s = \mathbf{H}_{ss}'\mathbf{H}_{ss} + \mathbf{H}_{sc}'\mathbf{H}_{sc}$ and $\mathbf{H}_s'\mathbf{H}\boldsymbol{\beta} = \mathbf{H}_{ss}'\mathbf{H}_{ss}\boldsymbol{\beta}_s + \mathbf{H}_{ss}'\mathbf{H}_{\sim ss}\boldsymbol{\beta}_{\sim s} + \mathbf{H}_{sc}'\mathbf{H}_{\bullet c}\boldsymbol{\beta}$ we see that we have

(A16) $$\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\rightarrow} \boldsymbol{\beta}_s + (\mathbf{H}_{ss}'\mathbf{H}_{ss} + \mathbf{H}_{sc}'\mathbf{H}_{sc})^{-1}[-\mathbf{H}_{sc}'\mathbf{H}_{sc}\boldsymbol{\beta}_s + \mathbf{H}_{ss}'\mathbf{H}_{\sim ss}\boldsymbol{\beta}_{\sim s} + \mathbf{H}_{sc}'\mathbf{H}_{\bullet c}\boldsymbol{\beta}].$$

The remainder of this appendix proves that this equation implies that simple beliefs converge to $\tau^*\boldsymbol{\beta}_s$, as indicated in Theorem 1.

(i) *Boundedness of t_s/t_c*

We begin by considering the possibility that the complex are in power only a finite number of times. In this case, as the simple will be in power an infinite number of times, we can use Lemma (A1) to calculate the following limits

(A17) $$\frac{\mathbf{H}_{ss}'\mathbf{H}_{\sim ss}}{t_sR} = \frac{\mathbf{X}_{ss}'\mathbf{N}_{\sim ss}}{t_sR} + \frac{\mathbf{N}_{ss}'\mathbf{N}_{\sim ss}}{t_sR} \overset{a.s.}{\rightarrow} \mathbf{0}_{k_s x k_{\sim s}},$$

$$\frac{\mathbf{H}_{ss}'\mathbf{H}_{ss}}{t_sR} - \frac{\mathbf{X}_{ss}'\mathbf{X}_{ss}}{t_sR} = \frac{\mathbf{N}_{ss}'\mathbf{X}_{ss}}{t_sR} + \frac{\mathbf{X}_{ss}'\mathbf{N}_{ss}}{t_sR} + \frac{\mathbf{N}_{ss}'\mathbf{N}_{ss}}{t_sR} \overset{a.s.}{\rightarrow} \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\& \frac{\mathbf{H}_{sc}'\mathbf{H}_{sc} - \mathbf{X}_{sc}'\mathbf{X}_{sc}}{t_sR} \rightarrow \mathbf{0}_{k_s x k_s}, \frac{\mathbf{H}_{sc}'\mathbf{H}_{sc}}{t_sR} \rightarrow \mathbf{0}_{k_s x k_s}, \& \frac{\mathbf{H}_{sc}'\mathbf{H}_{\bullet c}}{t_sR} \rightarrow \mathbf{0}_{k_s x k},$$

where in the first line we make use of the fact that $\mathbf{H}_{\sim ss} = \mathbf{N}_{\sim ss}$, as the simple set all policies they believe are irrelevant to zero, and in the last line that we are dividing the sum of a finite number of random variables by a number (t_s) that goes to infinity. Following the approach of the proof of Lemma (A2) earlier, we can then argue that:

(A18) $$\mathbf{v}'(\mathbf{H}_s'\mathbf{H}_s)^{-2}\mathbf{v} \le \frac{\mathbf{v}'\mathbf{v}}{\lambda_{\min}(\mathbf{H}_s'\mathbf{H}_s)^2} \le \frac{\mathbf{v}'\mathbf{v}/(t_sR)^2}{\lambda_{\min}\left((\mathbf{H}_s'\mathbf{H}_s - \mathbf{X}_s'\mathbf{X}_s)/t_sR\right)^2} \overset{a.s.}{\rightarrow} \frac{\mathbf{0}_{k_s x1}'\mathbf{0}_{k_s x1}}{(\sigma_n^2/R)^2} = 0,$$

$$\Rightarrow \mathbf{v}'(\mathbf{H}_s'\mathbf{H}_s)^{-2}\mathbf{v} \overset{a.s.}{\rightarrow} 0, \text{ where } \frac{\mathbf{v}}{t_sR} = -\frac{\mathbf{H}_{sc}'\mathbf{H}_{sc}}{t_sR}\boldsymbol{\beta}_s + \frac{\mathbf{H}_{ss}'\mathbf{H}_{\sim ss}}{t_sR}\boldsymbol{\beta}_{\sim s} + \frac{\mathbf{H}_{sc}'\mathbf{H}_{\bullet c}}{t_sR}\boldsymbol{\beta} \overset{a.s.}{\rightarrow} \mathbf{0}_{k_s x1}.$$

Combined with (A16), this implies that simple beliefs almost surely converge on the true parameter values, i.e., $\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\rightarrow} \boldsymbol{\beta}_s$. In this case, however, the simple have strictly lower intensity than the complex and hence must lose power to the complex. In sum, if the complex are in power only a finite number of times, then with the exception of paths whose total probability measure is zero, on each and every other path the limit of simple intensity is strictly less than the limit of complex intensity, ensuring that the complex are in the limit always in power, thereby establishing a contradiction.

We now consider the possibility that the simple are in power only a finite number of times. In this case, as the complex will be in power an infinite number of times, we use Lemma (A1) and (A12) to calculate

(A19) $$\frac{\mathbf{H}'_{sc}\mathbf{H}_{sc}}{t_c R} = \frac{\mathbf{X}'_{sc}\mathbf{X}_{sc}}{t_c R} + \frac{\mathbf{N}'_{sc}\mathbf{X}_{sc}}{t_c R} + \frac{\mathbf{X}'_{sc}\mathbf{N}_{sc}}{t_c R} + \frac{\mathbf{N}'_{sc}\mathbf{N}_{sc}}{t_c R} \overset{a.s.}{\to} \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'_s}{\boldsymbol{\beta}'\boldsymbol{\beta}} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\frac{\mathbf{H}'_{sc}\mathbf{H}_{\bullet c}}{t_c R} = \frac{\mathbf{X}'_{sc}\mathbf{X}_{\bullet c}}{t_c R} + \frac{\mathbf{N}'_{sc}\mathbf{X}_{\bullet c}}{t_c R} + \frac{\mathbf{X}'_{sc}\mathbf{N}_{\bullet c}}{t_c R} + \frac{\mathbf{N}'_{sc}\mathbf{N}_{\bullet c}}{t_c R} \overset{a.s.}{\to} \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'}{\boldsymbol{\beta}'\boldsymbol{\beta}} + \frac{\sigma_n^2}{R}[\mathbf{I}_{k_s}, \mathbf{0}_{k_s x k_{\sim s}}],$$

$$\& \ \frac{\mathbf{H}'_{ss}\mathbf{H}_{ss}}{t_c R} \to \mathbf{0}_{k_s x k_s} \ \& \ \frac{\mathbf{H}'_{ss}\mathbf{H}_{\sim ss}}{t_c R} \to \mathbf{0}_{k_s x k_{\sim s}},$$

where once again in the last line we make use of the fact that we are dividing the sum of a finite number of random variables by a number (t_c) that goes to infinity. Applying these limits to (A16), we see that if the simple are only in power a finite number of times

(A20) $$\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\to} \boldsymbol{\beta}_s + \left[\frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'_s}{\boldsymbol{\beta}'\boldsymbol{\beta}} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s}\right]^{-1}\left(-\left[\frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'_s}{\boldsymbol{\beta}'\boldsymbol{\beta}} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s}\right]\boldsymbol{\beta}_s + \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'}{\boldsymbol{\beta}'\boldsymbol{\beta}}\boldsymbol{\beta} + \frac{\sigma_n^2}{R}[\mathbf{I}_{k_s}, \mathbf{0}_{k_s x k_{\sim s}}]\boldsymbol{\beta}\right)$$

$$= \boldsymbol{\beta}_s + \left[(R/\sigma_n^2)\mathbf{I}_{k_s} - \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'_s(R/\sigma_n^2)^2/\boldsymbol{\beta}'\boldsymbol{\beta}}{1+\boldsymbol{\beta}'_s\boldsymbol{\beta}_s(R/\sigma_n^2)/\boldsymbol{\beta}'\boldsymbol{\beta}}\right]\left(1 - \frac{\boldsymbol{\beta}'_s\boldsymbol{\beta}_s}{\boldsymbol{\beta}'\boldsymbol{\beta}}\right)\boldsymbol{\beta}_s = \boldsymbol{\beta}_s + \frac{(1-\boldsymbol{\beta}'_s\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta})}{\sigma_n^2/R + \boldsymbol{\beta}'_s\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta}}\boldsymbol{\beta}_s.$$

From (A20), we see that as the ratio of noise to the information revealed by policy (σ_n^2/R) goes to infinity, $\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\to} \boldsymbol{\beta}_s$. This implies that asymptotically the simple have strictly lower voting intensity than the complex, which is consistent with their being in power only a finite number of times. In contrast, as σ_n^2/R goes to 0, (A20) reduces to $\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\to} \boldsymbol{\beta}_s(\boldsymbol{\beta}'\boldsymbol{\beta}/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s)$, which implies that, with the exception of paths of probability measure zero, on each and every other path in the limit the simple's voting intensity is strictly greater than that of the complex (as $\overline{\boldsymbol{\beta}}'_s\overline{\boldsymbol{\beta}}_s \overset{a.s.}{\to} (\boldsymbol{\beta}'\boldsymbol{\beta})^2/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s > \boldsymbol{\beta}'\boldsymbol{\beta} \overset{a.s.}{\leftarrow} \overline{\boldsymbol{\beta}}'_c\overline{\boldsymbol{\beta}}_c$), thereby contradicting the assumption that the simple are only in power a finite number of times.

Having established that both types must be in power an infinite number of times, we can use Lemma (A1) to recalculate limits for terms that appear on the right hand side of (A16) as both t_s and t_c go to infinity:

(A21) $$\frac{\mathbf{H}'_{ss}\mathbf{H}_{\sim ss}}{t_sR} = \frac{\mathbf{X}'_{ss}\mathbf{N}_{\sim ss}}{t_sR} + \frac{\mathbf{N}'_{ss}\mathbf{N}_{\sim ss}}{t_sR} \overset{a.s.}{\to} \mathbf{0}_{k_s x k_{\sim s}},$$

$$\frac{\mathbf{H}'_{ss}\mathbf{H}_{ss}}{t_sR} - \frac{\mathbf{X}'_{ss}\mathbf{X}_{ss}}{t_sR} = \frac{\mathbf{N}'_{ss}\mathbf{X}_{ss}}{t_sR} + \frac{\mathbf{X}'_{ss}\mathbf{N}_{ss}}{t_sR} + \frac{\mathbf{N}'_{ss}\mathbf{N}_{ss}}{t_sR} \overset{a.s.}{\to} \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\frac{\mathbf{H}'_{sc}\mathbf{H}_{sc} - \mathbf{X}'_{sc}\mathbf{X}_{sc}}{t_cR} = \frac{\mathbf{N}'_{sc}\mathbf{X}_{sc}}{t_cR} + \frac{\mathbf{X}'_{sc}\mathbf{N}_{sc}}{t_cR} + \frac{\mathbf{N}'_{sc}\mathbf{N}_{sc}}{t_cR} \overset{a.s.}{\to} \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\frac{\mathbf{H}'_{sc}\mathbf{H}_{sc}}{t_cR} = \frac{\mathbf{X}'_{sc}\mathbf{X}_{sc}}{t_cR} + \frac{\mathbf{N}'_{sc}\mathbf{X}_{sc}}{t_cR} + \frac{\mathbf{X}'_{sc}\mathbf{N}_{sc}}{t_cR} + \frac{\mathbf{N}'_{sc}\mathbf{N}_{sc}}{t_cR} \overset{a.s.}{\to} \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'_s}{\boldsymbol{\beta}'\boldsymbol{\beta}} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\& \frac{\mathbf{H}'_{sc}\mathbf{H}_{\bullet c}}{t_cR} = \frac{\mathbf{X}'_{sc}\mathbf{X}_{\bullet c}}{t_cR} + \frac{\mathbf{N}'_{sc}\mathbf{X}_{\bullet c}}{t_cR} + \frac{\mathbf{X}'_{sc}\mathbf{N}_{\bullet c}}{t_cR} + \frac{\mathbf{N}'_{sc}\mathbf{N}_{\bullet c}}{t_cR} \overset{a.s.}{\to} \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}'}{\boldsymbol{\beta}'\boldsymbol{\beta}} + [\frac{\sigma_n^2}{R}\mathbf{I}_{k_s}, \mathbf{0}_{k_s x k_{\sim s}}].$$

With regards to (A16) as a whole, using arguments as in Lemma (A2), we see that

(A22) $$(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s)'(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s) \overset{a.s.}{\to} \mathbf{v}'(\mathbf{H}'_s\mathbf{H}_s)^{-2}\mathbf{v} \le \frac{\mathbf{v}'\mathbf{v}/(t_sR)^2}{\lambda_{\min}\left((\mathbf{H}'_s\mathbf{H}_s - \mathbf{X}'_s\mathbf{X}_s)/t_sR\right)^2},$$

$$\text{with } \frac{\mathbf{v}}{t_sR} = \frac{t_c}{t_s}\left[-\frac{\mathbf{H}'_{sc}\mathbf{H}_{sc}}{t_cR}\boldsymbol{\beta}_s + \frac{\mathbf{H}'_{sc}\mathbf{H}_{\bullet c}}{t_cR}\boldsymbol{\beta}\right] + \frac{\mathbf{H}'_{ss}\mathbf{H}_{\sim ss}}{t_sR}\boldsymbol{\beta}_{\sim s} = \frac{t_c}{t_s}\mathbf{a} + \mathbf{b},$$

$$\frac{\mathbf{H}'_s\mathbf{H}_s - \mathbf{X}'_s\mathbf{X}_s}{t_sR} = \frac{\mathbf{H}'_{ss}\mathbf{H}_{ss}}{t_sR} - \frac{\mathbf{X}'_{ss}\mathbf{X}_{ss}}{t_sR} + \frac{t_c}{t_s}\left[\frac{\mathbf{H}'_{sc}\mathbf{H}_{sc} - \mathbf{X}'_{sc}\mathbf{X}_{sc}}{t_cR}\right] = \mathbf{c} + \frac{t_c}{t_s}\mathbf{d}$$

$$\text{where } \mathbf{a} \overset{a.s.}{\to} \left(1 - \frac{\boldsymbol{\beta}'_s\boldsymbol{\beta}_s}{\boldsymbol{\beta}'\boldsymbol{\beta}}\right)\boldsymbol{\beta}_s,\ \mathbf{b} \overset{a.s.}{\to} \mathbf{0}_{k_s x 1},\ \mathbf{c} \overset{a.s.}{\to} \frac{\sigma_n^2}{R}\mathbf{I}_{k_s},\ \&\ \mathbf{d} \overset{a.s.}{\to} \frac{\sigma_n^2}{R}\mathbf{I}_{k_s}.$$

The almost sure limits of **a**, **b**, **c** and **d**, tell us that, outside of paths of probability measure zero, for sufficiently large t_s and t_c, or equivalently sufficiently large t, the last term on the first line of (A22) can be made arbitrarily close to $\mathbf{a}'\mathbf{a}(1 + t_s/t_c)^{-2}(\sigma_n^2/R)^{-2}$. Together with the left-hand side of the first line of (A22), which shows that for sufficiently large t $(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s)'(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s)$ can be made arbitrarily close to $\mathbf{v}'(\mathbf{H}'_s\mathbf{H}_s)^{-2}\mathbf{v}$, this implies that outside of paths of probability measure zero in the limit $\bar{\boldsymbol{\beta}}_s$ lies within a sphere around $\boldsymbol{\beta}_s$ whose radius is a decreasing function of t_s/t_c.. By the triangle inequality there exists a κ^* such that the limiting value of simple intensity, for all possible beliefs residing in the sphere determined by $t_s/t_c = \kappa^*$, must be strictly less than the limiting value of complex intensity $(\boldsymbol{\beta}'\boldsymbol{\beta} > \boldsymbol{\beta}'_s\boldsymbol{\beta}_s)$.

Consequently, for all values of $t_s/t_c > \kappa^*$ in the limit, on each and every path outside of a possible set of paths of probability measure zero, simple intensity must be strictly less than complex intensity and t_s/t_c must be falling in the next period. This tells us that outside of a set of paths of probability measure zero in the limit $t_s/t_c < \kappa^*$, as it cannot rise above κ^*. As a result, for a large enough t, it must be that t_s/t_c is bounded from above by κ^*.

(ii) *A deterministic equation of monotonic motion*

Taking the (bounded) value of t_s/t_c as given, and secure in the knowledge that both t_c and t_s go to infinity, we substitute into the right-hand side of (A16) using the limits in (A21)

(A23)
$$\bar{\boldsymbol\beta}_s \overset{a.s.}{\rightarrow} \boldsymbol\beta_s + \left[\frac{\mathbf{X}'_{ss}\mathbf{X}_{ss}}{t_c R} + \frac{t_s}{t_c}\frac{\sigma_n^2}{R}\mathbf{I}_{k_s} + \frac{\boldsymbol\beta_s\boldsymbol\beta'_s}{\boldsymbol\beta'\boldsymbol\beta} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s}\right]^{-1} *$$
$$\left(-\left[\frac{\boldsymbol\beta_s\boldsymbol\beta'_s}{\boldsymbol\beta'\boldsymbol\beta} + \frac{\sigma_n^2}{R}\mathbf{I}_{k_s}\right]\boldsymbol\beta_s + \left[\frac{\boldsymbol\beta_s\boldsymbol\beta'}{\boldsymbol\beta'\boldsymbol\beta} + \frac{\sigma_n^2}{R}[\mathbf{I}_{k_s}, \mathbf{0}_{k_s x k_{\sim s}}]\right]\boldsymbol\beta\right)$$

$$\Rightarrow \quad \bar{\boldsymbol\beta}_s \overset{a.s.}{\rightarrow} \boldsymbol\beta_s + c\mathbf{M}^{-1}\boldsymbol\beta_s, \text{ where } \mathbf{M} = \frac{\mathbf{X}'_{ss}\mathbf{X}_{ss}}{t_c R} + \frac{t_c + t_s}{t_c}\frac{\sigma_n^2}{R}\mathbf{I}_{k_s} + \frac{\boldsymbol\beta_s\boldsymbol\beta'_s}{\boldsymbol\beta'\boldsymbol\beta} \quad \& \quad c = 1 - \frac{\boldsymbol\beta'_s\boldsymbol\beta_s}{\boldsymbol\beta'\boldsymbol\beta},$$

where we note that

(A24)
$$\lambda_{\max}(\mathbf{M}^{-1}) = \lambda_{\min}(\mathbf{M})^{-1} \leq \lambda_{\min}\left(\mathbf{M} - \frac{\mathbf{X}'_{ss}\mathbf{X}_{ss}}{t_c R}\right)^{-1}$$
$$= \lambda_{\min}\left(\frac{t_c + t_s}{t_c}\frac{\sigma_n^2}{R}\mathbf{I}_s + \frac{\boldsymbol\beta_s\boldsymbol\beta'_s}{\boldsymbol\beta'\boldsymbol\beta}\right)^{-1} = \left(\frac{t_c + t_s}{t_c}\frac{\sigma_n^2}{R}\right)^{-1} \leq \frac{R}{\sigma_n^2}$$

so the product of $\mathbf{M}^{-1}$ times the almost sure limits of $\mathbf{H}'_{sc}\mathbf{H}_{sc}/t_c R$ and $\mathbf{H}'_{sc}\mathbf{H}_{\bullet c}/t_c R$, and $\mathbf{M}^{-1} * t_s/t_c$ times the almost sure limit of $\mathbf{H}'_{ss}\mathbf{H}_{\sim ss}/t_s R$, all as given in (A21), is bounded, thereby validating the transition from (A16) to (A23). (A23) states that there is an asymptotic relationship between simple beliefs (on the left hand side) and parameter values, the relative share of time each type is in power, and the past history of simple policy (on the right hand side). In principle, it is possible that (accounting for the influence of current beliefs on the evolution of the history of simple policy) both left and right hand sides never converge to constants, perpetually changing in a manner that maintains their asymptotic equality. We now show that this cannot arise, and that (A23) in fact defines a deterministic and monotonic motion which leads simple beliefs to converge to their steady state values.

From (A23), we see that the limiting intensity of the simple almost surely equals

(A25)
$$\bar{\boldsymbol\beta}'_s\bar{\boldsymbol\beta}_s \overset{a.s.}{\rightarrow} \boldsymbol\beta'_s\boldsymbol\beta_s + 2f(c\mathbf{M}^{-1}) + f(c^2\mathbf{M}^{-2}),$$

$$\text{where } f(c\mathbf{M}^{-1}) = c\boldsymbol\beta'_s\mathbf{M}^{-1}\boldsymbol\beta_s \quad \text{and} \quad f(c^2\mathbf{M}^{-2}) = c^2\boldsymbol\beta'_s\mathbf{M}^{-1}\mathbf{M}^{-1}\boldsymbol\beta_s$$

Figure A1: $f(c\mathbf{M}^{-1})$ and $f(c^2\mathbf{M}^{-2})$ for Two-Dimensional $\bar{\boldsymbol{\beta}}_s$

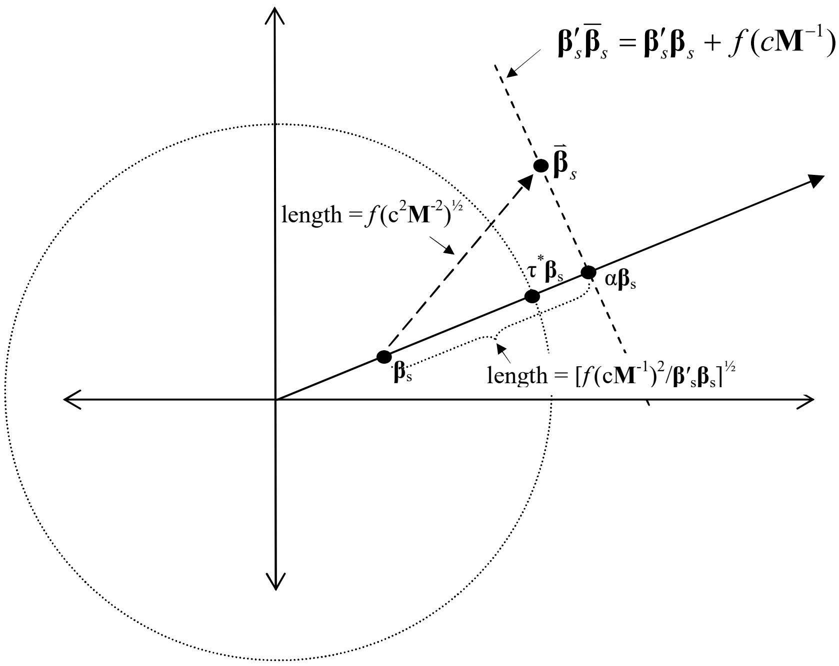

are quadratic forms involving $\boldsymbol{\beta}_s$. Moreover, from (A23) we also see that

(A26) $$\text{(a) } \boldsymbol{\beta}_s'(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s) \xrightarrow{a.s.} f(c\mathbf{M}^{-1}) \;\;\&\;\; \text{(b) } (\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s)'(\bar{\boldsymbol{\beta}}_s - \boldsymbol{\beta}_s) \xrightarrow{a.s.} f(c^2\mathbf{M}^{-2}),$$

so $f(c\mathbf{M}^{-1})$ is the equation of a plane perpendicular to the ray from the origin defined by the true parameter values, while $f(c^2\mathbf{M}^{-2})^{1/2}$ is the distance of the ray from the true parameter values to mean beliefs. These are illustrated graphically, for the case where $\boldsymbol{\beta}_s$ involves two policies, in Figure A1 above.

If $\alpha\boldsymbol{\beta}_s$ denotes the coordinates of the intersection of the plane defined by (A26a) with the ray from the origin defined by $\boldsymbol{\beta}_s$ (see Figure A1), we can substitute $\alpha\boldsymbol{\beta}_s$ for $\bar{\boldsymbol{\beta}}_s$ in (A26a)

(A27) $$\boldsymbol{\beta}_s'(\alpha\boldsymbol{\beta}_s - \boldsymbol{\beta}_s) \xrightarrow{a.s.} f(c\mathbf{M}^{-1}) \;\Rightarrow\; (\alpha-1)^2\boldsymbol{\beta}_s'\boldsymbol{\beta}_s \xrightarrow{a.s.} f(c\mathbf{M}^{-1})^2/\boldsymbol{\beta}_s'\boldsymbol{\beta}_s$$

However, the square of the length of the line segment from $\boldsymbol{\beta}_s$ to $\alpha\boldsymbol{\beta}_s$ is also $(\alpha-1)^2\boldsymbol{\beta}_s'\boldsymbol{\beta}_s$. By the Pythagorean theorem this must be less than or equal to the square of the length of the line segment from $\boldsymbol{\beta}_s$ to $\bar{\boldsymbol{\beta}}_s$, which equals $f(c^2\mathbf{M}^{-2})$. Consequently, $f(c^2\mathbf{M}^{-2}) \geq f(c\mathbf{M}^{-1})^2/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s$, with equality only when $\bar{\boldsymbol{\beta}}_s$ actually equals $\alpha\boldsymbol{\beta}_s$.[25] In sum, another interpretation of $f(c\mathbf{M}^{-1})$ is that it is proportional to the projection of the deviation of the simple's beliefs from the truth ($\boldsymbol{\beta}_s$) on the

[25]This result is also an implication of the generalized Cauchy-Schwarz inequality, which states that for a positive definite matrix **S**, and vectors **x** and **y**, $(\mathbf{x}'\mathbf{y})^2 \leq \mathbf{x}'\mathbf{S}\mathbf{y}*\mathbf{x}'\mathbf{S}^{-1}\mathbf{y}$ (Anderson 2003). Letting $\mathbf{x} = \mathbf{y} = c^{1/2}\boldsymbol{\beta}_s\mathbf{M}^{-1/2}$ and $\mathbf{S} = \mathbf{M}$, we have $(c\boldsymbol{\beta}'_s\mathbf{M}^{-1}\boldsymbol{\beta}_s)^2 \leq c\boldsymbol{\beta}'_s\boldsymbol{\beta}_s c\boldsymbol{\beta}'_s\mathbf{M}^{-1}\mathbf{M}^{-1}\boldsymbol{\beta}_s \rightarrow f(c^2\mathbf{M}^{-2}) \geq f(c\mathbf{M}^{-1})^2/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s$.

Figure A2: Asymptotic Phase Diagram for $f(c^2\mathbf{M}^{-2})$ and $f(c\mathbf{M}^{-1})$

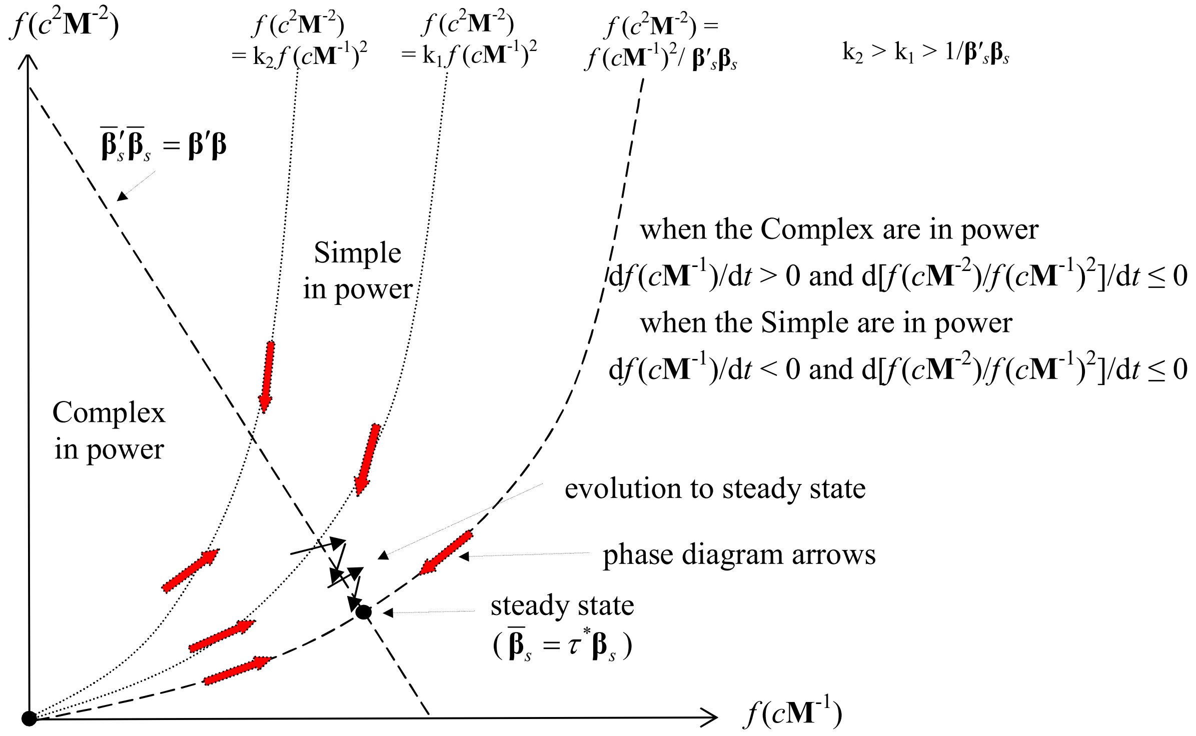

direction $\boldsymbol{\beta}_s$, a measure of bias, while the ratio $f(c^2\mathbf{M}^{-2})/[f(c\mathbf{M}^{-1})^2/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s]$ is the secant^2 of the angle of deviation from the direction $\boldsymbol{\beta}_s$.

Figure A2 draws the asymptotic phase diagram for $f(c\mathbf{M}^{-1})$ and $f(c^2\mathbf{M}^{-2})$, conditional on the knowledge that complex beliefs converge to the truth, neither type can be in power only a finite number of times and t_s/t_c is asymptotically bounded (results already proven above). We use this diagram to prove that simple beliefs converge to $\tau^*\boldsymbol{\beta}_s$.[26] The downward sloping dashed line, with slope -2, denotes the combinations that are consistent with $\bar{\boldsymbol{\beta}}'_s\bar{\boldsymbol{\beta}}_s = \boldsymbol{\beta}'\boldsymbol{\beta}$, i.e., the simple having the same voting intensity as the complex, based on (A25) above. Above the line the simple are in power, while below the line the complex are in power. Also drawn in the figure are "level curves" of the form $f(c^2\mathbf{M}^{-2}) = \mathrm{k}{*}f(c\mathbf{M}^{-1})^2$, with each curve defined by a different value of the constant k. The lowest curve, with $f(c^2\mathbf{M}^{-2}) = f(c^2\mathbf{M}^{-1})^2/\boldsymbol{\beta}'_s\boldsymbol{\beta}_s$ is, as noted in the previous paragraph, attained when $\bar{\boldsymbol{\beta}}_s$ is proportional to $\boldsymbol{\beta}_s$, and hence passes through the steady state. We prove the following results in the on-line appendix:

[26] Thus, to be clear, this diagram (which takes it as given that complex beliefs have converged to the truth) does not describe the actual convergence of simple and complex beliefs to their steady state values, but rather describes the dynamics of simple beliefs if complex beliefs have converged and simple beliefs have not.

(A28a) If the complex are in power $\mathrm{d}f(c\mathbf{M}^{-1})/\mathrm{d}t > 0$ and $\mathrm{d}[f(c^2\mathbf{M}^{-2})/f(c\mathbf{M}^{-1})^2]/\mathrm{d}t \leq 0$, with equality only along the steady state level curve;

(A28b) If the simple are in power $\mathrm{d}f(c\mathbf{M}^{-1})/\mathrm{d}t < 0$ and $\mathrm{d}[f(c^2\mathbf{M}^{-2})/f(c\mathbf{M}^{-1})^2]/\mathrm{d}t \leq 0$, with equality only along the steady state level curve;

(A28c) No matter which type is in power, $\lim_{t\to\infty} \mathrm{d}f(c\mathbf{M}^{-1})/\mathrm{d}t = 0$.

Asymptotically, when the complex are in power, bias as measured by the projection onto the directional vector given by the truth monotonically increases ($\mathrm{d}f(c\mathbf{M}^{-1})/\mathrm{d}t > 0$), while when the simple are in power it monotonically declines ($\mathrm{d}f(c\mathbf{M}^{-1})/\mathrm{d}t < 0$). Regardless of which type is in power, the angle of the deviation of beliefs from the direction implied by true parameter values monotonically falls, $\mathrm{d}[f(c^2\mathbf{M}^{-2})/f(c\mathbf{M}^{-1})^2]/\mathrm{d}t \leq 0$. As shown formally below, this effect comes from two factors: (i) noise, which regardless of which type is in power lowers the directional deviation of beliefs from $\boldsymbol{\beta}_s$, and (ii) the policy actions of the simple which, insofar as they are not proportional to $\boldsymbol{\beta}_s$, when contrasted with the actions of the complex reveal information about the relative effects of the k_s policies the simple consider relevant. The asymptotic collinearity of complex actions means that the effects of policies the simple believe are irrelevant can be loaded upon on any of the policies they believe are relevant. The effects of this bias are expressed in the form of movements of the line defined by $\boldsymbol{\beta}_s'\overline{\boldsymbol{\beta}}_s = \boldsymbol{\beta}_s'\boldsymbol{\beta}_s + f(c\mathbf{M}^{-1})$, but simple beliefs in principle could lie anywhere on this line. It is noise, plus the contrast between the effects of simple and complex actions when simple policies are not collinear in the area of overlap, that gradually reduces the deviation along this line from the ray $\alpha\boldsymbol{\beta}_s$.

(A28a) and (A28b) together establish that in the limit simple beliefs almost surely evolve toward the steady state following zig-zag paths such as the one drawn in the figure. (A28c), along with the monotonicity of $f(c^2\mathbf{M}^{-2})/f(c\mathbf{M}^{-1})^2$, ensures that these movements eventually stop.

(iii) *Movement continues until the unique steady state is reached*

As a final step, we need to show that when simple beliefs stop moving they must be at the steady state given in the figure, i.e., they cannot converge on some earlier point in the phase diagram path. We will first show that if simple beliefs converge they must converge to a point on the lowest level curve of the phase diagram, where simple beliefs are proportional to $\boldsymbol{\beta}_s$, and then show that this implies convergence to the steady state.

We return to the equation $\overline{\boldsymbol{\beta}}_s = \boldsymbol{\beta}_s + c\mathbf{M}^{-1}\boldsymbol{\beta}_s$, as defined in (A23), plugging in the almost sure limit of $\mathbf{X}_{ss}'\mathbf{X}_{ss}/t_s$ given knowledge that simple beliefs almost surely converge

(A29) $$\mathbf{M} = \mathbf{V} + \frac{\boldsymbol{\beta}_s\boldsymbol{\beta}_s'}{\boldsymbol{\beta}'\boldsymbol{\beta}} \text{ with } \mathbf{V} = \frac{t_s}{t_c}\frac{\overline{\boldsymbol{\beta}}_s\overline{\boldsymbol{\beta}}_s'}{\overline{\boldsymbol{\beta}}_s'\overline{\boldsymbol{\beta}}_s} + \frac{t_c + t_s}{t_c}\frac{\sigma_n^2}{R}\mathbf{I}_{k_s},$$

$$\text{so } \mathbf{M}^{-1}\boldsymbol{\beta}_s = \left[\mathbf{V}^{-1} - \frac{\mathbf{V}^{-1}\boldsymbol{\beta}_s\boldsymbol{\beta}_s'\mathbf{V}^{-1}/\boldsymbol{\beta}'\boldsymbol{\beta}}{1+\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta}}\right]\boldsymbol{\beta}_s = \frac{\mathbf{V}^{-1}\boldsymbol{\beta}_s}{1+\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta}} \;\;\&\;\; \mathbf{V}^{-1} = t_c\left[\frac{R}{t\sigma_n^2}\mathbf{I}_{k_s} - \frac{t_s\overline{\boldsymbol{\beta}}_s\overline{\boldsymbol{\beta}}_s'(R/t\sigma_n^2)^2/\overline{\boldsymbol{\beta}}_s'\overline{\boldsymbol{\beta}}_s}{1+t_s(R/t\sigma_n^2)}\right]$$

$$\Rightarrow \;\; \overline{\boldsymbol{\beta}}_s \overset{a.s.}{\rightarrow} \boldsymbol{\beta}_s + \frac{ct_c}{1+\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta}}\left[\frac{R}{t\sigma_n^2}\boldsymbol{\beta}_s - \frac{t_s\overline{\boldsymbol{\beta}}_s'\boldsymbol{\beta}_s R^2/t^2\sigma_n^4\overline{\boldsymbol{\beta}}_s'\overline{\boldsymbol{\beta}}_s}{1+t_sR/t\sigma_n^2}\overline{\boldsymbol{\beta}}_s\right],$$

$$\Rightarrow \;\; \overline{\boldsymbol{\beta}}_s \overset{a.s.}{\rightarrow} \tau\boldsymbol{\beta}_s, \text{ where } \tau = \frac{1+\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta} + c(R/\sigma_n^2)(t_c/t)}{1+\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta} + \dfrac{ct_ct_s\overline{\boldsymbol{\beta}}_s'\boldsymbol{\beta}_s R^2/t^2\sigma_n^4\overline{\boldsymbol{\beta}}_s'\overline{\boldsymbol{\beta}}_s}{1+t_sR/t\sigma_n^2}},$$

so we see that almost surely, i.e., outside of a measure zero of paths, the limiting value of simple beliefs on each and every path must be proportional to $\boldsymbol{\beta}_s$. We use this fact to calculate $\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s$ and substitute in the expression for τ (using as well the fact that $\overline{\boldsymbol{\beta}}_s'\boldsymbol{\beta}_s/\overline{\boldsymbol{\beta}}_s'\overline{\boldsymbol{\beta}}_s = 1/\tau$)

(A30) $$\boldsymbol{\beta}_s'\mathbf{V}^{-1}\boldsymbol{\beta}_s = t_c\left[\frac{R\boldsymbol{\beta}_s'\boldsymbol{\beta}_s}{t\sigma_n^2} - \frac{t_s\boldsymbol{\beta}_s'\boldsymbol{\beta}_sR^2/t^2\sigma_n^4}{1+t_sR/t\sigma_n^2}\right] = \frac{t_cR\boldsymbol{\beta}_s'\boldsymbol{\beta}_s/t\sigma_n^2}{1+t_sR/t\sigma_n^2}$$

$$\Rightarrow \;\tau = \frac{1+\dfrac{\boldsymbol{\beta}_s'\boldsymbol{\beta}_s}{\boldsymbol{\beta}'\boldsymbol{\beta}}\dfrac{t_cR/t\sigma_n^2}{1+t_sR/t\sigma_n^2} + c(t_cR/t\sigma_n^2)}{1+\dfrac{\boldsymbol{\beta}_s'\boldsymbol{\beta}_s}{\boldsymbol{\beta}'\boldsymbol{\beta}}\dfrac{t_cR/t\sigma_n^2}{1+t_sR/t\sigma_n^2} + \dfrac{1}{\tau}\dfrac{ct_ct_sR^2/t^2\sigma_n^4}{1+t_sR/t\sigma_n^2}} \Rightarrow \; \tau = 1 + \frac{c(t_cR/t\sigma_n^2)}{1+(t_sR/t\sigma_n^2)+\dfrac{\boldsymbol{\beta}_s'\boldsymbol{\beta}_s}{\boldsymbol{\beta}'\boldsymbol{\beta}}(t_cR/t\sigma_n^2)}.$$

The right hand side of the last line is decreasing in t_s/t and increasing in t_c/t, so we have

(A31) $$\frac{d\tau}{d(t_s/t)} < 0, \;\; \frac{d\tau}{d(t_c/t)} > 0, \;\; \lim_{\frac{t_c}{t}\to 0, \frac{t_s}{t}\to 1} \tau = 1 \;\;\&\;\; \lim_{\frac{t_c}{t}\to 1, \frac{t_s}{t}\to 0} \tau = 1 + \frac{(1-\boldsymbol{\beta}_s'\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta})}{\sigma_n^2/R + \boldsymbol{\beta}_s'\boldsymbol{\beta}_s/\boldsymbol{\beta}'\boldsymbol{\beta}}$$

The last expression was encountered earlier in (A20) and as σ_n^2/R goes to zero leads to a bias level $\tau = \tau^{*2} > \tau^*$.

(A29) - (A31) together ensure that movement in Figure A2 continues until simple beliefs converge on the steady state with bias τ equal to τ^*. In the limit beliefs almost surely must be proportional to $\boldsymbol{\beta}_s$. When $t_s/t = 1 - t_c/t$ is such that in the limit bias is greater than τ^*, on all but a measure zero of paths the simple must certainly be in power and t_s/t will rise while t_c/t falls, ensuring that τ falls, with opposite effects when τ is less than τ^* and the complex are in power. For small enough σ_n^2/R the limiting values of τ as t_s/t goes to zero and one encompass τ^*, ensuring that the limit of t_s/t is, almost surely, the one consistent with bias equal to the steady state value τ^*, as given in the text.

REFERENCES

Acemoglu, Daron, Georgy Egorov and Konstantin Sonin. 2013. "A Political Theory of Populism". *Quarterly Journal of Economics* 128 (2): 771-805.

Anderson, T.W. 2003. *An Introduction to Multivariate Statistical Analysis*. Third Edition. New Jersey: John Wiley & Sons.

Arrow, Kenneth J. and Jerry R. Green. 1973. "Notes on Expectations Equilibria in Bayesian Settings." Working paper #33, Institute for Mathematical Studies in the Social Sciences, Stanford University.

Berk, Robert. 1966. "Limiting behavior of posterior distributions when the model is incorrect." *Annals of Mathematical Statistics* 37: 51–58.

Bernhardt, Dan, Stefan Krasa, and Mehdi Shadmehr. 2019. "Demagogues and the Fragility of Democracy." Manuscript, University of Illinois, Urbana.

Bischof, Daniel and Roman Senninger. 2018. "Simple politics for the people? Complexity in campaign messages and political knowledge." *European Journal of Political Research* 57 (2): 473-495.

Bohren, J. Aislinn and Daniel N. Hauser. 2019. "Social Learning with Model Misspecification: A Framework and a Characterization." Manuscript, 2019.

Bray, Margaret. 1982. "Learning, Estimation, and the Stability of Rational Expectations." *Journal of Economic Theory* 26: 318-339.

Chen, Xueliang, Yuanle Yan, and Jie Hu. 2019. "A Corpus-Based Study of HIllary Clinton's and Donald Trump's Linguistic Styles." *International Journal of English Linguistics* 9 (3): 13-22.

Dahrendorf, Ralf. 2007. "Acht Anmerkungen zum Populismus". *Eurozine* 2007 (first published Transit 25, 2003).

Decasdri, Silvia and Constantine Boussalis. 2019. "Truth in Simplicity: Populism and Language Complexity in the Italian Chamber of Deputies." Centro Einaudi working paper no. 8, 2019.

Di Tella, Rafael and Julio J. Rotemberg. 2016. "Populism and the Return of the `Paranoid Style': Some Evidence and a Simple Model of Demand for Incompetence as Insurance against Elite Betrayal". Manuscript, Harvard Business School.

Dornbusch, Rudiger and Sebastian Edwards. 1991. "The Macroeconomics of Populism." In Dornbusch & Edwards, eds, *The Macroeconomics of Populism in Latin America*. Chicago, University of Chicago Press.

Eliaz, Kfir and Ran Spiegler. 2019. "A Model of Competing Narratives." Manuscript.

Enns, Peter. 2014. "The Public's Increasing Punitiveness and its Influence on Mass Incarceration in the United States." *American Journal of Political Science* 58 (4): 857-872.

Esponda, Ignacio. 2008. "Behavioral Equilibrium in Economies with Adverse Selection." *American Economic Review* 98 (4): 1269-1291.

Esponda, Ignacio and Demian Pouzo. 2016. "Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica* 84 (3): 1093-1130.

Esponda, Ignacio, Demian Pouzo and Yuichi Yamamoto. 2021. "Asymptotic Behavior of Bayesian Learners with Misspecified Models." *Journal of Economic Theory* 195: 1-52.

Frick, Mira, Ryota Iijima and Yuhta Ishii. 2020. "Stability and Robustness in Misspecified Learning Models. Manuscript, 2020.

Gagnon-Bartsch, Tristan, Matthew Rabin, and Joshua Schwartzstein. 2019. "Channelled Attention and Stable Errors." Working Paper: HBS WP 18-108.

Golub, Gene M. 1973. "Some Modified Matrix Eigenvalue Problems." *SIAM Review* 15 (2): 318-334.

Guiso, Luigi, Helios Herrera, Massimo Morelli, and Tommaso Sonno. 2017. "Demand and Supply of Populism." Manuscript, Einaudi Institute for Economics and Finance.

Guriev, Sergei and Elias Papaioannou. 2020. "The Political Economy of Populism". Centre for Economic Policy Research discussion paper no. DP14433, 2020.

Heidhues, Paul, Botond Kőszegi, Philipp Strack. 2018. "Unrealistic Expectations and Misguided Learning". *Econometrica* 86 (4): 1159-1214

Heidhues, Paul, Botond Kőszegi, Philipp Strack. 2021. "Convergence in Models of Misspecified Learning." *Theoretical Economics* 16: 73-99.

Jennings, Will, Emily Gray, Stephen Farrall, and Colin Hay. 2017. "Penal Populism and the Public Thermostat: Crime, Public Punitiveness and Public Policy." *Governance* 30 (3): 463-481.

Mailath, George J. and Larry Samuelson. 2019. "Learning under Diverse World Views: Model-Based Inference." PIER working paper 19-018, University of Pennsylvania.

Molavi, Pooya. 2019. "Macroeconomics with Learning and Misspecification: A General Theory and Applications." Manuscript, 2019.

Montiel Olea, José Luis, Pietro Ortoleva, Mallesh Pai and Andrea Prat. 2019. "Competing Models". Manuscript, 2019.

Morelli, Massimo, Antonio Nicoló and Paolo Roberti. 2020. "Populism as Simplistic Commitment." Manuscript.

Mudde, Cas. 2004. "The Populist Zeitgeist". *Government and Opposition* 39 (4): 541-563.

Nyarko, Yaw. 1991. "Learning in Mis-specified Models and the Possibility of Cycles." *Journal of Economic Theory* 55 (2): 416-427.

Sonin, Konstantin. 2020. "The Historical Perspective on the Trump Puzzle: A Review of Barry Eichengreen's `The Populist Temptation: Economic Grievance and Political Reaction in the Modern Era'", Becker-Friedman Institute working paper, no. 2020-129, 2020.

Zellner, Arnold. 1971. *An Introduction to Bayesian Inference in Econometrics*. New Jersey: John Wiley & Sons.