Political social-learning: Short-term memory and cycles of polarisation

Gilat Levy and Ronny Razin

Online Appendix

1 Proofs for Section I

Proof of Lemma 1: Assume that party L offers l. If party R offers l too it attains $\frac{1}{2}\alpha$, whereas if it switches to r it attains $(1-\Pr(L \text{ wins}|l,r))(1+\alpha)$, where

$$(1-\Pr(L \text{ wins}|l,r)) = \begin{cases} 1 \text{ if } 1/2+\zeta E[\beta_r-\beta_l|H_t] > 1 \\ 0 \text{ if } 1/2+\zeta E[\beta_r-\beta_l|H_t] < 0 \\ 1/2+\zeta E[\beta_r-\beta_l|H_t] \text{ otherwise} \end{cases}.$$

Note that if $1/2+\zeta E[\beta_r-\beta_l|H_t]<0$, then party R indeed offers l, and when $1/2+\zeta E[\beta_r-\beta_l|H_t]>1$, party R will best respond by offering r. When $1/2+\zeta E[\beta_r-\beta_l|H_t]\in(0,1)$, then party R will offer l when $(1/2+\zeta E[\beta_r-\beta_l|H_t])(1+\alpha)<\frac{1}{2}\alpha$, which amounts to $E[\beta_l-\beta_r|H_t]>1/2\zeta(1+\alpha)$. Given the above, whenever $E[\beta_l-\beta_r|H_t]>1/2\zeta(1+\alpha)$, party R offers l when party L offers l. Note that if this is the case, party L for sure offers l. An analogous condition, $E[\beta_r-\beta_l|H_t]>1/2\zeta(1+\alpha)$, guarantees that a consensus on r is the unique equilibrium. In all other cases, polarisation must arise as the unique equilibrium, that is, when

$$|E[\beta_l-\beta_r|H_t]| < 1/2\zeta(1+\alpha).$$

In the non-generic cases in which $|E[\beta_l-\beta_r|H_t]|=1/2\zeta(1+\alpha)$ both polarisation and consensus on one of the policies will be an equilibrium.■

We repeat here for convenience some of the notation defined in the text. Denote the expected outcome when policy p is implemented and given degenerate beliefs on some parameters β, as $E[y|p,\beta]$.

The random history (that arises given the randomness in the election and the randomness of the shock ε, through its effect on beliefs), induces a probability distribution P over the set of infinite paths of histories $\mathbb{H}$. Thus, when we write "almost surely" we mean P-almost surely on $\mathbb{H}$.

Remember that for full history up to time t we define the associated distribution over implemented actions at time t, $\hat{\eta}_t(p)$, as the share of time policy p was implemented, and we let $\eta_t(\textit{polarisation})$ and $\eta_t(\textit{consensus})$ be the fraction of time in the full history up to time t that the two parties offered different platforms and the same policy respectively.

1.1 Proofs for Section II

Proof of Proposition 1: Voters' posterior after observing the history H_t satisfies the conditions of the martingale convergence theorem. Therefore, for almost any infinite path, voters' beliefs at any period t, μ_t, converge almost surely to some limit probability distribution μ_∞.

We now consider the measure one of all paths for which the posteriors converge. Consider first all paths for which, in the limit, $\left|E_{\mu_\infty}[\beta_l-\beta_r|H_t]\right|>\rho$. By Lemma 1, for these paths, parties will offer the same policy p, in line with the statement of Proposition 1.

Consider next the paths for which $\left|E_{\mu_\infty}[\beta_l-\beta_r|H_t]\right|<\rho$. By Lemma 1 party polarisation is the unique equilibrium in the limit on these paths. Assume by contradiction that the probability mass of this set of paths is strictly positive. Each of the parties is elected in equilibrium with a strictly positive probability due to $\alpha>0$ and Lemma 1. We then have a strictly positive measure of paths for which $\lim_{t\to\infty}\inf\hat{\eta}_t(l)$ and $\lim_{t\to\infty}\inf\hat{\eta}_t(r)$ are bounded away from zero.

As the choice of policies in the model is endogenous and as they affect learning, this implies that the process of observed outcomes is not iid. We therefore cannot use standard laws of large numbers to pin down what are the limit beliefs. For this reason, we use below a result from Esponda, Pouzo and Yamamoto (2021), henceforth EPY.

Specifically, note that the Kullback-Leibler divergence at time t between a distribution of posteriors induced by some vector of parameters β and the posterior induced by the true parameters β^*, given the fractions $\hat{\eta}_t(l)$ and $\hat{\eta}_t(r)$, is defined as:

$$KL(\beta|\hat{\eta}_t, \beta^*) = \sum_{p \in \{l,r\}} \hat{\eta}_t(p) \int_{\mathbb{R}} f(\varepsilon) \ln \frac{f(\varepsilon)}{f(E[y|p,\beta^*] + \varepsilon - E[y|p,\beta])} d\varepsilon$$

where $f(\varepsilon)$ is the density over ε, Normal with mean zero. The KL divergence value is always non-negative and the true state, β^*, is a minimizer of the KL value attaining $KL(\beta^*|\hat{\eta}_t, \beta^*) = 0$ regardless of $\hat{\eta}_t$. But as $\hat{\eta}_t(l)$ and $\hat{\eta}_t(r)$ are bounded away from zero, this means that the true state, β^*, satisfies $\beta^* = \arg\min_{\beta'} KL(\beta'|\hat{\eta}_t, \beta^*)$, i.e., is the unique such minimiser of the KL divergence. The result in EPY, in the context of our model, implies then that

$$\lim_{t \to \infty} \int_{B_t} KL(\beta|\hat{\eta}_t, \beta^*) d\mu_{t+1}(\beta) = 0 \text{ almost surely.}$$

That is, for any observed frequency of actions, the posterior beliefs will concentrate on values of β for which $KL(\beta|\hat{\eta}_t, \beta^*)$ is closest to its minimized value, which is zero.[1] This result implies, by continuity, that if for $\beta \in B$ the KL value is strictly positive, then a ball around β must have zero measure in the limit beliefs μ_∞. Thus, beliefs can only concentrate on a ball around β^*. As $\beta_l^* - \beta_r^* > \rho$, this is in contradiction to the supposition that the beliefs converge to satisfy $\left|E_{\mu_\infty}[\beta_l - \beta_r|H_t]\right| < \rho$.

Finally, consider the infinite paths along which beliefs converge to satisfy $\left|E_{\mu_\infty}[\beta_l - \beta_r|H_t]\right| = \rho$. If either $\hat{\eta}_t(l)$ or $\hat{\eta}_t(r)$ converge to zero in any subsequence, implying consensus almost surely. If $\hat{\eta}_t(l)$ and $\hat{\eta}_t(r)$ are both bounded away from zero, by similar arguments to the argument above, beliefs must converge to satisfy $\left|E_{\mu_\infty}[\beta_l - \beta_r|H_t]\right| > \rho$, a contradiction. This concludes the proof of the proposition.■

The following Lemma will be helpful in the proof of Propositions 2 below and all results that pertain to the case of $\sigma \to 0$. In particular, it will imply that beliefs that arise in the equilibrium sequence as $\sigma \to 0$ converge to equilibrium beliefs when $\sigma = 0$. Thus the sequence of equilibria as $\sigma \to 0$ will also converge to the limit equilibrium when $\sigma = 0$.

Lemma A1: *Assume that* $K \geq 2$ and $\sigma \to 0$. *(i) Suppose that there is a strictly positive measure of histories H_t such that only one policy p was implemented throughout the history. Then almost surely beliefs will concentrate on $(\beta_l^*, E(\beta_r|\beta_l^*))$ in period t. (ii) If there is a strictly positive measure of histories H_t such that both policy l and policy r were implemented, then almost surely beliefs will concentrate on (β_l^*, β_r^*) in period t.*

Proof of Lemma A1: (i) Assume one policy p is implemented for K periods in a strictly positive measure of histories. Note that with the normal distribution over the shocks, for any $\gamma', \gamma'' > 0$ there is a $\bar{\sigma} > 0$ such that for all $\sigma < \bar{\sigma}$ with probability $1 - \gamma'$ all the shocks in the K periods are in $[-\gamma'', \gamma'']$. As $\sigma \to 0$ the distribution of shocks concentrates on its expectation. As a result, when $\sigma \to 0$, with probability arbitrarily close to one, the posterior belief after any path will be concentrated on $(\beta_p^*, E[\beta_{-p}|\beta_p^*])$. (ii) Assume that both l and r have been implemented in a strictly positive measure of histories of K periods. Again, for any $\gamma', \gamma'' > 0$ there is a $\bar{\sigma} > 0$ such that for all $\sigma < \bar{\sigma}$ with probability $1 - \gamma'$ all the shocks in the K periods are in $[-\gamma'', \gamma'']$. As a result, when $\sigma \to 0$, with probability arbitrarily close to one, the posterior belief after almost any path will be concentrated on β^*.■

Proof of Proposition 2: Consider the limit when $\sigma = 0$. After histories H_t that contain two different implemented policies, parties will both offer the optimal policy l. Once there is a K-period history in which only this optimal policy is implemented, parties will polarise, and will continue to do so

[1] Our model satisfies assumptions 1-3 in EPY, which are all technical and relate to the compactness of B and continuity of the outcome function y. In EPY the policy function determining the mapping from beliefs to action is deterministic. In our model the action that is implemented at every period is random when parties polarise, given the shock ϕ_t. But this has no bearing on the proof of the result in EPY.

until two different policies are implemented at which point parties will revert to a consensus on policy l. As the consensus phase is on the correct policy l, the polarisation phase ends once r is selected, which happens with an interior probability $1/2 - \zeta(\beta_l^* - E(\beta_r|\beta_l^*))$ at any period (this probability is interior as parties only polarise when each has a strictly positive probability of winning). This probability allows us to calculate the expected length of the polarisation phase and hence the share of time that the correct policy is implemented. As the limit equilibrium is unique, by Lemma A.1, the result also holds for $\sigma \to 0$.■

Analysis of equilibria when $\sigma \to 0$ and when β^* does not satisfy Assumption 1:

The result below characterises the equilibria when $\sigma = 0$. We continue to assume that $\beta_l^* - \beta_r^* > \rho$, but now consider violations of Assumption 1.

Proposition A1 *Let $\sigma = 0$ and $K \geq 2$. If Assumption 1 is violated, then the equilibrium is perpetual consensus or perpetual polarisation, unless $\{|\beta_l^* - E[\beta_r|\beta_l^*]| < \rho$ and $E[\beta_l|\beta_r^*] - \beta_r^* > \rho\}$ in which case the unique equilibrium is the same cycle as in Proposition 2.*

Proof of Proposition A1:

If Assumption 1 is violated then there are three cases to consider:

1. $|\beta_l^* - E[\beta_r|\beta_l^*]| > \rho$ and $|\beta_r^* - E[\beta_l|\beta_r^*]| < \rho$.
2. $|\beta_l^* - E[\beta_r|\beta_l^*]| > \rho$ and $|\beta_r^* - E[\beta_l|\beta_r^*]| > \rho$.
3. $|\beta_l^* - E[\beta_r|\beta_l^*]| < \rho$ and $|\beta_r^* - E[\beta_l|\beta_r^*]| > \rho$.

Case 1: Remember that the optimal policy is l. Assume that only policy l was implemented in the observed history, then in the next stage the parties will be in consensus on policy l if $\beta_l^* > E[\beta_r|\beta_l^*]$ and on policy r if $\beta_l^* < E[\beta_r|\beta_l^*]$. In the former case, this will be an absorbing state of consensus. In the latter case, after one period of implementing r the polity immediately learns that the optimal policy is l. Parties will then be in consensus on l for K periods, and then switching to a consensus on r for one period, and so on. Thus in this case, we have perpetual consensus.

If the history is composed of observations of the two policies, again we move to a consensus phase in which l is implemented, which takes us to the analysis above and hence we reach perpetual consensus.

Assume now that only r was implemented in the observed history. Then, given that $|\beta_r^* - E[\beta_l|\beta_r^*]| < \rho$, parties polarise, and once l is implemented for the first time, voters learn that the optimal policy is l, implying that l will be implemented for K periods and as we have shown above, this will lead to perpetual consensus.

Case 2: In this case, whatever the history in terms of implemented policies, parties will be in consensus on some policy p as either learning the whole state or learning just one parameter implies a consensus. So in equilibrium we will have consensus forever after some period.

Case 3: If $\beta_r^* - E[\beta_l|\beta_r^*] > \rho$ consensus on r is an absorbing state, as a history in which voters only observed r leads to the beliefs above and parties will converge on offering r by Lemma 1. Depending on the initial history, and specifically if it does not contain only observations of r, an additional equilibrium that is identical to the one in Proposition 2 may arise. This equilibrium is sustained as along its path we never reach a consensus on the policy r. When $E[\beta_l|\beta_r^*] - \beta_r^* > \rho$ the unique equilibrium is the cycle we have in Proposition 2. If we start with an initial history in which only r was implemented and we hence attain these beliefs, in the next period parties will both espouse l and so we revert to the equilibrium cycle of Proposition 2.■

Example for "scale-free" learning discussed in Section II.B:

To capture the limits of learning from β_p, let

$$\bar{\Delta}(\beta_p) = \sup_{p \in \{l,r\},\beta} |\beta_p - E[\beta_{-p}|\beta_p]|$$

We construct a sequence of priors $\{G_n(\beta_l, \beta_r)\}_{n=1}^{\infty}$ for which in the limit $\bar{\Delta}_n(\beta_p)$ converges to zero for any β_p. This implies that for almost any state of the world $\beta_l^* - \beta_r^* > \bar{\Delta}(\beta_p) \approx 0$, and so Assumption 1 is satisfied whenever $\beta_l^* - \beta_r^* > \rho$. To see this, re-parameterise the distribution $G_n(\beta_l, \beta_r)$ into the parameter space (v, δ) where $v = \frac{\beta_l + \beta_r}{2}$ represents the scale, the mid-point between the utilities, and $\delta = \beta_l - \beta_r$ represents the utility difference. Let $\hat{G}_n(v, \delta) \equiv G_n(v + \frac{\delta}{2}, v - \frac{\delta}{2})$ be the transformed distribution function that satisfies independence between v and δ. Let $\hat{G}_2(\delta|v) = \hat{G}_2(\delta)$ be the marginal over δ, which we assume is independent of v and n and symmetric around zero so that $\int_{-\infty}^{\infty} \delta \hat{g}_2(\delta) d\delta = 0$. Let $\hat{G}_{1,n}(v)$ denote the marginal of this distribution over scale, and we assume that it is uniform on $[-D_n, D_n]$. In addition, we assume that for any v and x, $E[\delta|\delta > x]$ is finite. We now show that as $D_n \to \infty$, $\lim_{n\to\infty} \bar{\Delta}_n(\beta_p) = 0$ for any β_p.

For any β_l^* we have that,

$$
\begin{aligned}
&\beta_l^* - E_{G_n(\beta_l,\beta_r)}[\beta_r|\beta_l^*] \\
&= \beta_l^* - \int_{v,\delta \text{ so that } v+\frac{\delta}{2}=\beta_l^*} (v - \frac{\delta}{2}) \frac{\hat{g}_{1,n}(v)\hat{g}_2(\delta)}{\int_{v',\delta' \text{ so that } v'+\frac{\delta'}{2}=\beta_l^*} \hat{g}_{1,n}(v')\hat{g}_2(\delta')dv'd\delta'} dv d\delta \\
&= \beta_l^* - \int_{2(\beta_l^*-D_n)}^{2(\beta_l^*+D_n)} (\beta_l^* - \delta)\hat{g}_2(\delta) d\delta \\
&\to_{D_n\to\infty} \int_{-\infty}^{\infty} \delta \hat{g}_2(\delta) d\delta = 0 \blacksquare
\end{aligned}
$$

Proof of Proposition 3:

Step A1: *For a large enough K, there is no positive measure of paths along which there is a subsequence $\{t_n\}_{n=1}^{\infty}$ such that $\hat{\eta}_{t_n}(p) \to 1$ for some $p \in \{l, r\}$.*

Proof of Step A1: To see this, let us assume to the contrary that there exists such a subsequence t_n which on a strictly positive measure of paths satisfies that $\hat{\eta}_{t_n}(p) \to 1$ for some $p \in \{l, r\}$. For any t, denote the preceding K periods of history as the $K - window$ at t.

Claim A1: *For a large enough K, along the subsequence $\{t_n\}_{n=1}^{\infty}$, after a $K - window$ in which only one policy $p \in \{l, r\}$ was implemented, almost surely the next period will involve parties polarising with a strictly positive probability.*

Proof of Claim A1: Consider $t_n \to \infty$ and then a large enough K. Then for each $K - window$ with a fixed p, beliefs will concentrate on β_p^* and $E[\beta_{-p}|\beta_p^*]$ with a strictly positive probability. However, as $|\beta_p^* - E[\beta_{-p}|\beta_p^*]| < \rho$ and by Lemma 1, parties will polarise in the next period with strictly positive probability.$\square_{claim\ A.2}$

We can now use Claim A1 to prove Step A1. As $\hat{\eta}_{t_n}(p) \to 1$, the fraction of these $K - windows$ with only p implemented within the window must be going to one. By Claim A.2, each of these will lead to polarisation with a strictly positive probability almost surely, and so $\eta_{t_n}(polarisation)$ is in the order of $1/K$. But as each party wins with strictly positive probability when there is polarisation, this contradicts the supposition that $\hat{\eta}_{t_n}(p) \to 1$.$\square$

Step A2: *For a large enough K, $\liminf_{t\to\infty} \eta_t(polarisation) > 0$ almost surely.*

Proof of Step A2: Suppose not, and so there is a positive measure of paths along which there is a subsequence t_n such that $\eta_{t_n}(polarisation) \to 0$. This implies that if we look at the $K - windows$ along these paths almost all of them include no polarisation. Following from step A1, it cannot be that there is strictly positive measure of $K - windows$ with only one policy implemented as then we would have $\eta_{t_n}(polarisation)$ bounded from zero as we showed above.

Thus the only possibility that remains is that in almost all $K - windows$, at least two policies p and p' are implemented, and that parties will shift from a consensus on one policy p to a consensus on another policy p' (a "consensus-switch").

So assume that in almost all $K - windows$, at least two policies p and p' are implemented. Assume first that in all these K-*windows* the ratio of the share of time that p was implemented compared to the share of time that p' was implemented, converges to some finite $c > 0$. But then, as beliefs in almost all such K-*window* must converge to β^* when K grows large, after almost all such K-*window* both parties will choose the optimal policy l and so $\hat{\eta}_{t_n}(l) \to 1$, a contradiction to a finite c.

Thus we must have a strictly positive measure of $K - windows$ for which this ratio of implemented policies converges to zero or infinity. Let p' denote the policy implemented most times in the $K - windows$. Note that this ratio has to converge to infinity slow enough so that overall beliefs do not converge to $(\beta_{p'}^*, E[\beta_p|\beta_{p'}^*])$, as then we would have polarisation after such histories implying a contradiction to $\eta_{t_n}(polarisation) \to 0$.

Let us examine then what happens to $\beta^*_{p'} - E[\beta_p|\beta^*_{p'}, H_{t_n}]$, the beliefs attained for a large K, at a path where mostly p' is implemented. Note that for large K, $E[\beta_{p'}|H_{t_n}] - E[\beta_p|H_{t_n}]$ is arbitrarily close to $\beta^*_{p'} - E[\beta_p|\beta^*_{p'}, H_{t_n}]$. But as β^* satisfies Assumption 1, we have that $\left|\beta^*_{p'} - E[\beta_p|\beta^*_{p'}]\right| < \rho$. As we look at a strictly positive measure of paths, we can use iterated expectation to conclude that $E[\beta^*_{p'} - E[\beta_p|\beta^*_{p'}, H_{t_n}]] = \beta^*_{p'} - E[\beta_p|\beta^*_{p'}] \leq \left|\beta^*_{p'} - E[\beta_p|\beta^*_{p'}]\right| < \rho$. This implies that $\beta^*_{p'} - E[\beta_p|\beta^*_{p'}, H_{t_n}] < \rho$ with a strictly positive probability. As a result, for a strictly positive measure of paths we should have polarisation and hence a contradiction to $\eta_{t_n}(polarisation) \to 0$.□

Step A3: *For a large enough K, $\lim_{t\to\infty} \sup_t \eta_t(polarisation) < 1$ almost surely.*

Suppose not, and so there is a positive measure of paths along which there is a subsequence t_n such that $\eta_{t_n}(polarisation) \to 1$. This implies that if we look at all the $K-windows$ almost all of them include polarisation at every period, implying that for all windows there exist at least two different policies p and p' implemented with a strictly positive probability. As a result, for a large enough K and $t_n \to \infty$, after almost all the $K-windows$ we have that, as in Proposition 1, beliefs almost surely concentrate on a ball around β^*. This implies that both parties must choose the optimal policy after almost all these $K-windows$, a contradiction to $\eta_{t_n}(polarisation) \to 1$.□

Step A4: *For a large enough K, for any $\sigma > 0$, there exists $\eta_K > 0$ such that*

$$\min\{\lim_{t\to\infty} \inf \eta_t(polarisation), \lim_{t\to\infty} \inf \eta_t(consensus)\} > \eta_K.$$

Suppose the statement is not true. Steps A2 and A3 imply that there exists a large enough K such that for any $\sigma > 0$, $\min\{\lim_{t\to\infty} \inf \eta_t(polarisation), \lim_{t\to\infty} \inf \eta_t(consensus)\} > 0$ almost surely. So for the statement to be wrong we must have that

$\lim_{\sigma\to 0} \min\{\lim_{t\to\infty} \inf \eta_t(polarisation), \lim_{t\to\infty} \inf \eta_t(consensus)\} = 0$ with strictly positive probability. But in Proposition 2 we have shown that at $\sigma = 0$,

$\min\{\lim_{t\to\infty} \inf \eta_t(polarisation), \lim_{t\to\infty} \inf \eta_t(consensus)\} > 1/K$. Therefore, by continuity there must be an $\eta_K > 0$ that satisfies the statement of Step A4.□

The above concludes part (i). To consider part (ii) of the Proposition, note that consensus on policy p arises when:

$$\left|E[(\beta_p - \beta_{-p})|H_t]\right| > \rho.$$

As $K < \infty$ there is always a strictly positive probability that the above inequality arises for the wrong policy.■

Proof of Proposition 4: Assume that $\lim_{K\to\infty} \limsup_{t\to\infty} \eta_t(polarisation)/(\eta_t(consensus)) > \psi$ for some $\psi > 0$, and so there is a strictly positive measure of paths for which for any convergent sequence $\{t_n\}$, $\lim_{K\to\infty} \lim_{t_n\to\infty} \hat{\eta}_{t_n}(p)/\hat{\eta}_{t_n}(p') \geq c$, for any p, p', where $c \in (0, \infty)$. As $K \to \infty$, by continuity, and using similar arguments as in Proposition 1, this implies that $\lim_{K\to\infty} \lim_{t_n\to\infty} E[\beta_l - \beta_r|H_{t_n}] \to \beta^*_l - \beta^*_r > \rho$, and so parties must converge in the long run on l almost surely, and so we must have $\lim_{K\to\infty} \liminf_{t_n\to\infty} \eta_t(polarisation)/(\eta_t(consensus)) = 0$ almost surely, a contradiction.■

1.2 Proofs for Section III

Proof of Proposition 5:

We prove three results that together suffice for the proof of Proposition 5.

Claim A2 *Let $\sigma \to 0$. For small enough λ:*

(i) The equilibrium in the model with long-term memory follows the following phases:

(1) A phase of consensus on l that continues $\lfloor \bar{t}(\lambda) \rfloor$ periods, where $\bar{t}(\lambda)$ satisfies $\phi(\bar{t}(\lambda)) = \varphi \equiv ((\beta^{a*}_l - \beta^{a*}_r) - \rho)/(1-\pi)(\beta^{b*}_r - \beta^{a*}_r) \in (0, 1)$. At period $\lfloor \bar{t}(\lambda) \rfloor + 1$ we move to phase (2).

(2) A phase of polarisation which lasts until r is implemented. After r is implemented, voters learn the current state. If voters learn that the state is a we move to phase (1). If voters learn that the state is b we move to phase (3).

(3) A phase of consensus on r until the state changes to state a at which point we revert to (1).

(ii) The long run proportion of time implementing the wrong policy can be arbitrarily close to $(1-\pi)$ as λ goes to zero by setting $\beta^{a*}_l - (\pi\beta^{a*}_r + (1-\pi)\beta^{b*}_r)$ arbitrarily close to ρ.

Proof of Claim A2: We begin by proving (i). Consider first the case in which voters are sure that the state is a at some point in time which we name time 0, and that there is a consensus on l for

t periods after. Note that even though the voters cannot see a change of state in these periods, voters are aware that the state may have changed. They will then compute the probability that a change had occurred in this time frame of t periods. Specifically, the probability of a change at period t (and not before) is $(1-\lambda)^{t-1}\lambda$, and the probability that a state had changed before time t is $\phi(t) = 1-(1-\lambda)^t$. Note that the expected time in which a change actually occurs is $1/\lambda$.

Along this path, voters cannot be certain whether a change has occurred. To vote, the voters will compute the expected difference in outcomes next period, between implementing policy l rather than r:

$$\phi(t)(\beta_l^{a*} - (\pi\beta_r^{a*} + (1-\pi)\beta_r^{b*})) + (1-\phi(t))(\beta_l^{a*} - \beta_r^{a*}).$$

As long as this expression is higher than ρ then parties will remain in consensus on l but once this expression is smaller than ρ parties will polarise and subsequently voters will learn the true state once r is chosen.

As $\phi(t)$ is increasing, the expression above is decreasing, and we can calculate $\bar{t}$ (possibly a non-integer), which satisfies:

$$\phi(\bar{t})(\beta_l^{a*} - (\pi\beta_r^{a*} + (1-\pi)\beta_r^{b*})) + (1-\phi(\bar{t}))(\beta_l^{a*} - \beta_r^{a*}) = \rho$$

$$\Leftrightarrow \phi(\bar{t}) = ((\beta_l^{a*} - \beta_r^{a*}) - \rho)/(1-\pi)(\beta_r^{b*} - \beta_r^{a*}) \equiv \varphi \in (0,1).$$

The last statement that $\varphi \in (0,1)$ follows from our assumptions on the parameters. Note that φ could be made arbitrarily close to one if $(\beta_l^{a*} - (\pi\beta_r^{a*} + (1-\pi)\beta_r^{b*}))$ is close to ρ (from below).

Thus, the equilibrium will have consensus up to the time $\lfloor\bar{t}\rfloor$. Computing the period this happens, $\lfloor\bar{t}\rfloor$, we get:

$$\phi(\bar{t}) = 1-(1-\lambda)^{\bar{t}} = \varphi$$
$$\rightarrow \bar{t}(\lambda) = \ln(1-\varphi)/\ln(1-\lambda).$$

Consider now that we are in a history in which there was consensus on r for some periods. At any period voters can learn the state. If the state is b the consensus on r continues, and if the state is a then parties switch to a consensus on l and the continuation follows as above. When the polity is in the polarisation phase, once policy r is implemented voters will know the state and we move to one of the consensus periods above. This concludes the proof of (i).

We now prove (ii): We focus on the relevant phase in which mistakes can happen, which is phase 1, on which there is a consensus on l. The expected time in which a change actually occurs is $1/\lambda$. To assess how much time passes until polarisation follows the expected change in the state, we look at $\lim_{\lambda\to 0} \bar{t}(\lambda)/(1/\lambda) = \lim_{\lambda\to 0} \lambda \ln(1-\varphi)/\ln(1-\lambda) = -\ln(1-\varphi)$ where the last equality follows, as when $\lambda \to 0$ the expected time of change in state, $1/\lambda$, grows large at the same rate as $1/(-\ln(1-\lambda))$. As we can take φ to be as close to one as we want (by choosing $(\beta_l^{a*} - (\pi\beta_r^{a*} + (1-\pi)\beta_r^{b*}))$ close to ρ) this implies that voters, even though they might be quite sure that a change has happened, may delay indefinitely their switching to a polarisation phase and continue to choose the wrong policy.

As the average number of periods between changes of the state is $1/\lambda$, it is instructive to count in terms of λ-periods which are blocks of $\lfloor 1/\lambda \rfloor$ periods. Note that whenever policy r is implemented, then, upon a change, with probability π state a is drawn and then detected, implying that the polity moves to implement l. Alternatively, when l is implemented, the polity will remain stuck on it for many λ-periods as $\lim_{\lambda\to 0} \bar{t}(\lambda)/(1/\lambda) = -\ln(1-\varphi) \to_{\varphi\to 1} \infty$. This implies that eventually the polity is stuck on l and so will implement the wrong decision whenever the state is b, which arises on average with probability $1-\pi$.■

We now consider voters with short-term memory. Since we assume that these voters are not aware that the state may change, we need to characterize how they update their beliefs. We can do this for strictly positive σ and characterise the beliefs as $\sigma \to 0$.

As $\lambda \to 0$ we only consider cases in which there was one change in the K period that the voters observe; the probability of two or more changes is negligible. Note further that by Lemma A1, when $\sigma \to 0$, observed outcomes will concentrate on β_r^{b*} or on β_r^{a*} when implementing policy r. We now consider what happens to beliefs when voters observe both such outcomes in their history. Note that beliefs can change only when r is implemented as $\beta_l^{b*} = \beta_l^{a*}$.

Claim A3: *Assume a K period history of play in which voters observe outcomes generated by policy r for $K_r \leq K$ periods in this history. For any $0.5|\beta_r^{b*} - \beta_r^{a*}| > \delta > 0$, as $\sigma \to 0$, voters' posteriors are given by:*

(i) If a majority (minority) of outcomes out of the K_r are in a $\delta-$ball around β_r^{b*} then the voters believe that the state is b (a) with probability converging to 1.

(ii) If the number of periods out of the K_r in which the outcome was in a $\delta-$ball around β_r^{b*} is equal to the number of periods out of the K_r in which the outcome was in a $\delta-$ball around β_r^{a*} then the voters beliefs converge to the prior.

Proof of Claim A3: If there are different states in these K_r periods, let us rename and order all observations (of $\{y_t\}_{t=1}^{t=K_r}$) under state b to be implemented at periods $t = 1$ to $t = T$, and all observations under state a to be implemented under renamed periods $t = T + 1$ to $t = K_r$. This is without loss of generality. Let

$$\delta(K_r) \equiv \frac{\Pi_{t=1}^{T} f_\sigma(\beta_r^{b*} + \varepsilon_t - \beta_r^{a*}) \Pi_{t=T+1}^{K_r} f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{a*})}{\Pi_{t=1}^{T} f_\sigma(\beta_r^{b*} + \varepsilon_t - \beta_r^{b*}) \Pi_{T+1}^{K_r} f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{b*})}$$

Then:

$$\Pr(\beta_r = \beta_r^{b*}) = (1 - \pi)/((1 - \pi) + \pi\delta(K_r))$$

As $\sigma \to 0$, by the same arguments as in Lemma A2, $\lim_{\sigma\to 0} f_\sigma(\beta_r^{b*} + \varepsilon_t - \beta_r^{a*})/f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{b*}) = 1$ and $\lim_{\sigma\to 0} f_\sigma(\beta_r^{b*} + \varepsilon_t - \beta_r^{b*})/f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{a*}) = 1$ and so $\delta(K_r) \to_{\sigma\to 0} 1$ if $T = K/2$, which implies that voters beliefs converge to the prior. If $T < K/2$,

$$\delta(K_r) \underset{\sigma\to 0}{\to}$$

$$\frac{\Pi_{t=2T+1}^{t=K_r} f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{a*})}{\Pi_{t=2T+1}^{K_r} f_\sigma(\beta_r^{a*} + \varepsilon_t - \beta_r^{b*})} \underset{\sigma\to 0}{\to} \infty$$

and so $\Pr(\beta_r = \beta_r^{b*}) \underset{\sigma\to 0}{\to} 0$ and similarly if $T > K_r/2$,

$$\delta(K_r) \underset{\sigma\to 0}{\to} 0$$

and so $\Pr(\beta_r = \beta_r^{b*}) \underset{\sigma\to 0}{\to} 1$.■

Claim A4:

When voters have short term memory, in the limit when $\sigma \to 0$, the long run fraction of time implementing the non-optimal policy is of the order of $1/K$, and the equilibrium has the following phases:

(i) A consensus phase of K periods with consensus on l, followed by phase (ii).

(ii) A polarisation phase until r is implemented. Once r is implemented, voters revert to phase (i) if the state is a, and move to phase (iii) if the state is b.

(iii) A phase of consensus on r, that ends with a switch to phase (i) $\lfloor K/2 \rfloor + 1$ periods after the state changed to a.

Proof of Claim A4:

Note that when voters know that the state is a, there is a consensus on l for K periods as no new information is generated when l is implemented. After K periods the voters have no knowledge about the state from the history and their belief about the state accords with the prior belief. A polarisation phase arises, and lasts until the first time r is implemented. When r is implemented the voters will learn the state immediately. If the state is a we go back to the consensus on l and such a consensus will last for K periods. If the state is b there will be consensus on r until the state changes to a and then, in accordance with Claim A3, the polity will revert to the consensus on l after $\lfloor K/2 \rfloor + 1$ periods (as $\lambda \to 0$ we consider only one change of policy per $\lambda-$*period*).

Note that the instances in which the polity makes a mistake are

1. There was a change of state to b within phase (i).
2. There was a change of state to b in phase (ii) when policy r hasn't been chosen yet.

3. The first time r is implemented in phase (ii) and the state is still a.

4. In phase (iii) when the state switches to a in the periods before the polity switches to phase (i).

When $\lambda \to 0$ the mistakes in (1), (2) and (4) are negligible in size as they happen only when a change of state occurs, on average once in any $\lambda - period$. However, mistakes in (3) happen in any phase (ii) which is recurring multiple times in any $\lambda - period$. Still, the probability of mistakes in (3) is smaller than $1/K$. This concludes the proof of the claim.■

This concludes the proof of the proposition.■

1.3 Proofs for Section IV

Proof of Lemma 2: The proof follows that of Lemma 1 by substituting the expression derived in the text in Section 2 for $\Pr(L \text{ wins}|l, r)$.■

Proof of Proposition 6:

Let $\sum_{j \leq \hat{\jmath}} w^j = \mu(\hat{\jmath})$. Without loss of generality assume that $\beta_l^* - E[\beta_r|\beta_l^*]$ and $\beta_l^* - \beta_r^*$ are such that for any $\hat{\jmath} \in \{0, 1, ..., m\}$ we have $L(\hat{\jmath}) \equiv |\mu(\hat{\jmath})(\beta_l^* - E[\beta_r|\beta_l^*]) + (1 - \mu(\hat{\jmath}))(\beta_l^* - \beta_r^*)| \neq \rho$. Note that $L(0) > \rho$ and $L(m) < \rho$. In addition, $L(\hat{\jmath})$ is either decreasing in $\hat{\jmath}$ or first decreasing and then increasing. This follows from the fact that $\mu(\hat{\jmath})(\beta_l^* - E[\beta_r|\beta_l^*]) + (1 - \mu(\hat{\jmath}))(\beta_l^* - \beta_r^*)$ is always decreasing in $\hat{\jmath}$, and it is either positive at $\hat{\jmath} = m$ or negative. In the latter case its absolute value $L(\hat{\jmath})$ is then decreasing and then increasing. This means that we can find a unique $j^* \in \{1, ..., m\}$ which is the solution to the following inequalities:

$$L(j^* - 1) > \rho > L(j^*)$$

Now consider the following limit equilibrium cycle (when $\sigma \to 0$):

(i) A consensus phase on l lasting exactly K^{j^*} periods after which we move to phase (ii) below.

(ii) A polarisation phase which continues until the party espousing r wins an election, in this case we go to phase (i).

To see that this is a unique equilibrium, note that along this cycle, in a consensus phase that lasted for $K < K^{j^*}$ periods, we will have by assumption 1, the equilibrium conjecture, and the definition of j^*:

$$\begin{aligned} &|\sum_{j \text{ such that } K^j < K^{j^*}} w^j(\beta_l^* - E[\beta_r|\beta_l^*]) + \sum_{j \text{ such that } K^{j^*} \geq K} w^j(\beta_l^* - \beta_r^*)| \\ &= L(j|j \leq j^* - 1) \geq L(j^* - 1) > \rho \end{aligned}$$

and so by Lemma 2 there will be consensus on l in the next election. After period K^{j^*} of the consensus phase we will have polarisation. As long as l is elected polarisation continues. The expected length of polarisation will depend on the probability that party L wins the election, as:

$$\Pr(L|p^L, p^R) = \begin{cases} 1 \text{ if } 1/2 + \zeta \sum_j w^j E[(\beta_l - \beta_r)|H_t^j] > 1 \\ 0 \text{ if } 1/2 + \zeta \sum_j w^j E[(\beta_l - \beta_r)|H_t^j] < 0 \\ 1/2 + \zeta \sum_j w^j E[(\beta_l - \beta_r)|H_t^j] \text{ otherwise} \end{cases}$$

Note that in this expression all groups' information affects the probability that L wins. As there is randomness in the election's outcome, this probability might change over time as $\sum_j w^j E[(\beta_l - \beta_r)|H_t^j]$ changes. In particular, as long as the polarisation phase lasts, more groups' histories will move from $E[(\beta_l - \beta_r)|H_t^j] = \beta_l^* - \beta_r^*$ to $E[(\beta_l - \beta_r)|H_t^j] = \beta_l^* - E[\beta_r|\beta_l^*]$ as their history progresses to contain only l being implemented.■