

Information Quality, Disagreement and Political Polarisation

R. Emre Aytimur*

Richard M. H. Suen†

This Version: 5th May, 2026.

Abstract

How does the quality of information available to voters shape political polarisation? We address this question within an electoral competition model where voters infer an unknown state from a noisy and potentially biased signal. Voters' policy preferences are determined by their posterior belief, which is unobservable to the parties when they choose their platforms. Greater uncertainty about the election outcome increases parties' incentives to polarise. We show that improvements in information quality can either amplify or reduce polarisation, depending on the extent of disagreement between voters' and politicians' prior beliefs. We also analyse the welfare implications of both polarisation and information quality, and show that more precise information can reduce voters' welfare when there is significant disagreement.

Keywords: Polarisation, Voter Information, Bayesian Learning, Election

JEL Classification: D72, D80

**School of Economics, University of Leicester, Leicester LE1 7RH, United Kingdom. Email: rea22@le.ac.uk

†Corresponding Author: School of Economics, University of Leicester, Leicester LE1 7RH, United Kingdom.
Email: mhs15@le.ac.uk

1 Introduction

A well-informed electorate is essential for the functioning of representative democracy. Knowledge and information about the state of the world help voters form opinions and choose representatives that best align with their interests. In practice, however, voters and politicians often make decisions under imperfect information and uncertainty, particularly when policies have long-term consequences or concern unprecedented events such as a pandemic. The absence of definitive evidence leaves room for biased information and disagreement. Recent major elections have revealed the extent to which misinformation and political polarisation can shape democratic outcomes (Grinberg *et al.*, 2019; Chen *et al.*, 2021; Munger *et al.*, 2022). Understanding how political information affects polarisation is therefore of pressing importance. In this paper, we provide a systematic study on how the beliefs and information possessed by voters and politicians affect the extent of policy divergence and voters’ welfare.

The main novelty of our analysis is to allow for disagreement between voters’ and politicians’ subjective beliefs. This feature aligns with the empirical evidence showing that political elites and their staff often misperceive their constituent’s opinions (Broochman and Skovron, 2018; Hertel-Fernandez *et al.*, 2019; Pereira, 2021; Kärnä and Öhberg, 2023). Existing studies suggest several explanations, e.g., politicians’ views may reflect their own socioeconomic backgrounds; they may be disproportionately influenced by organised interests such as activists or lobbyists (Giger and Klüver, 2016); or they may simply disregard opinions that conflict with their own (Butler and Dynes, 2016). We do not take a stance on the origin of such disagreement, but instead focus on its implications for political polarisation. Our main contribution is twofold: First, we show that voters’ and politicians’ beliefs each have different effects on policy polarisation. Introducing disagreement between these two groups allows us to disentangle these effects and analyse them separately.¹ Second, we show that disagreement is a significant factor in determining both the positive and normative implications of information quality.

Our model builds on the canonical framework of electoral competition with policy-motivated parties (Calvert, 1985; Roemer, 1994; Bernhardt *et al.*, 2009), where the decisive voter’s preferred policy depends on an unknown state of the world. The hidden state can represent unforeseen events (e.g., an economic crisis or war) or the socially optimal policy response to complex issues (e.g., immigration, public health, economic reform). Unlike these earlier studies, we model voters

¹We abstract from belief heterogeneity among voters and between the political parties. Further discussion about this can be found at the end of Section 2.1.

as Bayesian learners who update their subjective prior belief upon the arrival of new information.

In our model, voters’ prior belief includes any information and knowledge that they possess *before* the political parties announce their platforms. After knowing the platforms but before the election, voters receive a public noisy signal about the hidden state.² This represents the dissemination of political news and information from both formal channels (such as mainstream news media, official government announcements, opinion polls, etc.) and informal ones (such as social media). We assume that the public signal is not only fraught with potential errors, it may also be biased.³ The potential bias is captured by an additive random term. Voters have prior belief about the hidden state and the unknown bias, but they cannot separately identify these factors from the observed signal.⁴ After receiving the new information, voters form posterior belief that determine their policy preferences, but these are unobserved by the parties. As shown in the literature cited in the paragraph above, political parties tend to polarise when they are unsure about voters’ preferences. This happens because moderation yields limited electoral advantage under uncertainty. In our setting, this uncertainty is endogenously determined by voters’ learning. Since the parties do not observe the signal when setting platforms, they must form expectations about voters’ preferences based on their own subjective belief, which may be different from those of voters. Specifically, we assume that voters and politicians agree to disagree on their prior beliefs about the hidden state and the bias term.⁵

The key insight of this setup is that voters’ and politicians’ prior beliefs affect electoral outcome through different channels. First, if the politicians’ prior belief is diffuse, or if they believe the precision of the signal is low, they will find it difficult to predict voters’ updated belief and policy preferences. This increases the uncertainty faced by the political parties which incentivise them to polarise. We refer to this as the *uncertainty effect*, which is determined by signal precision

²The timing of the signal makes clear that we are focusing on the voters’ learning process after the parties announced their platforms but before the election. This setup captures the following ideas: Before entering the election booth, voters are free to adjust their policy stance in response to any new information. But political parties are less likely to make significant platform changes before the election in fear of any potential damages on their credibility and reputation.

³Bias in news and information can arise due to various reasons. For instance, any ill-devised sampling and data collection method can result in systemic distortions (bias) in measurements (e.g., opinion poll based on a non-representative sample). This type of distortions is inherent in the data collection method and thus will affect all the measurements (signals) generated by the same method. Another possible reason is sensationalism in journalism, i.e., tactics used by editors to attract viewers. An obvious example is exaggeration: labelling minor achievements as “*groundbreaking*” and minor setback as “*catastrophic*”. A third possible reason is partisan bias of news media. There is extensive empirical evidence on the pervasiveness of biased reporting in mass media. See, for instance, Puglisi and Snyder (2015) for a comprehensive survey.

⁴A similar learning model with biased signal is considered in Little and Pepinsky (2021) and Little *et al.* (2022).

⁵This heterogeneous-prior assumption is intended to capture the idea that voters and politicians often differ intrinsically in their experiences or perspectives (e.g., due to differences in socioeconomic background), leading them to form different worldviews and interpret information differently.

and the belief held by the politicians. Second, if voters are uncertain about the hidden state in their initial belief, or if they think the signal is precise, they will rely heavily on the signal when updating their belief. This again makes voters' preferences harder to predict *ex ante* and increases the uncertainty faced by the parties. We refer to this as the *learning effect*, which depends on the voters' responsiveness to the realised signal. Within this framework, we present three sets of results.

Our first set of results considers how improvement in signal precision affects equilibrium polarisation. On the one hand, a more precise signal makes voters more confident on the learning process, thereby amplifying the learning effect and increasing polarisation. On the other hand, a less noisy signal allows the parties to better predict voters' beliefs, reducing the uncertainty effect and lowering polarisation. Which force dominates depends crucially on the assumption regarding disagreement. If voters and politicians share the same prior, then the learning effect *always* dominates. In other words, improvement in signal precision always promotes polarisation when there is no disagreement. This result continues to hold if voters are more confident in their prior than the politicians. However, if the politicians have a sufficiently more precise prior than the voters, then the uncertainty effect dominates, and greater signal precision reduces polarisation. This result suggests that for policy issues where voters hold firm prior views, better information leads to more polarisation; but for issues that they are unfamiliar with and parties have more precise views than voters, better information can reduce polarisation.

These findings are valid under both biased and unbiased signals. They are also robust in a more general model with multiple correlated signals.⁶ In this setting, reducing the correlation between signals is akin to improving their precision. Intuitively, voters will have more confidence on the learning process if they perceive the signals as independent opinions rather than "echo chambers". This will strengthen the learning effect and promote polarisation. Reducing the correlation also means that it is easier for the parties to predict the signals. This weakens the uncertainty effect and reduces polarisation. As before, the net effect depends on the two sides' prior beliefs.

Our second set of results concerns signal bias. We assume that both voters and politicians expect the signals to be symmetrically distributed around zero in their prior beliefs.⁷ A key factor here is the correlation between the bias term and the hidden state. If the two are independent,

⁶The details of this general model can be found in Sections C and E of the Online Supplementary Materials.

⁷Normalising the prior mean of the bias term to zero has no effect on the voters' learning process. This is true because voters will subtract the prior mean of the hidden state and the bias term from the observed signal when formulating their posterior estimate. See Lemma 1 for a formal statement of this result.

then the bias merely adds more noise to the signal. The more interesting case is when the two are correlated. The signal is called exaggerating [resp., contradicting] if the bias term is positively [resp., negatively] correlated with the hidden state.⁸ When voters expect the signal to be contradicting, they may engage in what we call *defiant learning*, i.e., they update their belief in the *opposite* direction as suggested by the realised signal. This arises not from the magnitude of the bias but from its uncertainty and negative correlation with the hidden state. This type of learning behaviour is not possible under the conventional Bayesian model with unbiased signal.

Our final set of results concerns the implications of policy polarisation on the voters' *ex ante* welfare (i.e., welfare before the signal is realised). Policy divergence between two ideologically differentiated parties can serve as an insurance against the uncertainty about the state of the world. Therefore, policy polarisation can be welfare-improving for risk-averse voters. We refer to this as the *insurance effect* of polarisation. In the welfare analysis, we seek to address the following questions: (i) How does the disagreement between voters and politicians affect the welfare implications of polarisation? (ii) Does improving signal precision always improve welfare?

Regarding the first question, we find that when there is little or no disagreement between voters and politicians, voters strictly prefer a society with partisan parties and a positive level of polarisation to an otherwise identical society but with more congruent parties and convergent platforms. Intuitively, equilibrium policies are very polarised only when the parties are highly uncertain about the election outcome. If voters and politicians share a common prior, then this high level of uncertainty is shared by the voters. As a result, voters will be more willing to accept a large extent of polarisation as insurance. However, when disagreement is substantial, particularly when voters hold much more precise priors than the politicians, then polarisation can become excessive and welfare-reducing. This happens because voters are relatively certain about the election outcome so the insurance effect is limited, but politicians with a diffuse prior have strong incentive to polarise. The result is over-polarisation from the voters' perspective.

As for the second question, improvement in signal precision can be welfare-reducing when there is significant disagreement between voters' and politicians' beliefs. This can happen in two ways: First, better information increases the parties' perceived uncertainty by intensifying the learning effect, which leads to excessive polarisation and welfare loss. Second, better information lowers the uncertainty faced by the parties and induces them to narrow the gap between their platforms, thereby reducing the insurance benefit of polarisation. These results show that the

⁸Examples of contradicting signal include misinformation that discredit the scientific evidence behind man-made climate changes or the efficacy of vaccination.

welfare effects of better information crucially depend on the degree of disagreement between voters and politicians.

Related Literature Our analysis builds upon Bernhardt *et al.* (2009) but extends their framework in two key dimensions. First, we introduce learning which provides a channel through which new information can affect voters' policy preferences. This enables a systematic study on how information precision and signal bias shape equilibrium polarisation and welfare. Second, we introduce disagreement in prior beliefs between voters and parties, allowing us to disentangle two distinct mechanisms: an uncertainty effect, arising from parties' difficulty in anticipating voters' belief; and a learning effect, driven by voters' responsiveness to new information. In addition, while Bernhardt *et al.* (2009) highlight the insurance role of polarisation when voters and politicians share the same belief, we show that disagreement can render polarisation excessive and that, in some cases, greater information precision may reduce welfare.

The present study contributes to the growing literature that examines the effect of voter information on policy outcomes. Each of the studies discussed below, however, focuses on a different mechanism from the one that we considered. Gul and Pesendorfer (2012) consider how the media industry structure and voter polarisation can affect the candidate endorsement strategies of profit-maximising media and consequently policy choices of parties. When the number of media firms approaches infinity, each voter is able to find a media firm that endorses her favourite party in each state of the world. This means that the electorate behaves as if they are perfectly informed and this leads to policy polarisation. In Levy and Razin (2015), voters receive private signals about the state of the world, and if voters have correlation neglect, they become more sensitive to the signals and therefore their beliefs become more dispersed. Levy and Razin (2015) show that this does not necessarily lead to policy polarisation. They also note that correlation neglect of voters makes the information aggregation more efficient, which increases voter welfare. In Yuksel (2022), voters differ in the policy dimensions they find important, and when they specialise in their learning accordingly, this leads to more dispersed voter beliefs and consequently to more policy polarisation. However, for a given level of specialised learning, better access to information of voters leads to reduced party polarisation. In Yuksel (2022)'s model, policy polarisation always reduces voter welfare. In a related paper to Yuksel (2022), Perego and Yuksel (2022) show how media competition leads to informational specialisation across voters and consequently to social disagreement. Personalised demand for information leads to inefficient policy outcomes in

Matejka and Tabellini (2021). Similarly, personalised news aggregators result in different types of voters (centrist and extreme voters) receiving different information and potentially lead to policy polarisation in Hu *et al.* (2023). Finally, Vaeth (2025) shows that when voters with rational inattention acquire information to determine which party to support, their behaviour endogenously leads to both voter and party polarisation in a multidimensional policy setting.

There is also a large literature which offers different reasons for policy divergence, including policy motivation (Wittman, 1983; Calvert, 1985; Besley and Coate, 1997); entry deterrence (Palfrey, 1984; Callander, 2005); agency problems (Van Weelden, 2013); incomplete information among voters or candidates (Castanheira, 2003; Bernhardt *et al.*, 2007; Callander, 2008); and differential candidate valence (Bernhardt and Ingberman, 1985; Groseclose, 2001; Krasa and Polborn, 2012).

Finally, in the empirical literature on political polarisation, there is a consensus that party polarisation in the U.S. is on the rise (McCarty *et al.*, 2006), but whether voters have become more polarised is less clear (Barber and McCarty, 2015). The explanation we provide for party polarisation does not rely on voter polarisation, but rather on the interaction between voter information and the disagreement between voters and politicians.

The rest of the paper is organised as follows. Section 2 presents the model environment and preliminary results. Section 3 analyses how information and beliefs affect equilibrium polarisation. Section 4 presents the welfare analysis. Section 5 concludes.

2 The Model

2.1 The Environment

Consider an election in which two political parties, L and R , compete on a one-dimensional policy issue. Prior to the election, the two parties simultaneously announce a policy platform from the policy space $X \equiv \mathbb{R}$. The electorate consists of a continuum of voters with heterogeneous policy preferences. The size of the electorate is normalised to one. Each voter v 's policy preferences are determined by two factors: (i) a deterministic parameter $\delta_v \in \mathbb{R}$ which captures the voter's pre-existing political attitudes, and (ii) a random variable $s \in \mathbb{R}$ which captures the exogenous state of the world. In any given state s , voter v 's utility from policy $x \in \mathbb{R}$ is given by

$$U(x; \delta_v, s) = -(\delta_v + s - x)^2.$$

The median of δ_v is normalised to zero. The exact distribution of δ_v across voters is irrelevant to our analysis.

Voters do not observe the realisation of s at the time of the election. Instead, they receive a noisy and potentially biased public signal m , defined as

$$m = s + b + \varepsilon,$$

where $b \in \mathbb{R}$ is an unknown parameter that captures the inherent bias of the information source, and $\varepsilon \in \mathbb{R}$ is the error term. Voters share the same subjective prior belief about the hidden state and the unknown bias term.⁹ This is assumed to take the form of a bivariate normal distribution with mean vector $\boldsymbol{\mu}_0$ and covariance matrix $\boldsymbol{\Sigma}_0$ specified as

$$\boldsymbol{\mu}_0 = \begin{bmatrix} \mu_s \\ \mu_b \end{bmatrix} \quad \text{and} \quad \boldsymbol{\Sigma}_0 = \begin{bmatrix} \sigma_s^2 & \rho_{s,b}\sigma_s\sigma_b \\ \rho_{s,b}\sigma_s\sigma_b & \sigma_b^2 \end{bmatrix}.$$

In the above expressions, $\mu_s \in \mathbb{R}$ and $\sigma_s^2 > 0$ denote, respectively, the mean and variance of the marginal distribution of s ; while $\mu_b \in \mathbb{R}$ and $\sigma_b^2 > 0$ are the counterparts for the bias term. The parameter $\rho_{s,b} \in (-1, 1)$ measures the perceived correlation between s and b . A positive $\rho_{s,b}$ means that the bias term is expected to exaggerate or complement the effect of the hidden state. A negative value, on the other hand, means that voters expect the bias term to contradict or subdue the effect of s . The covariances between (s, b) and m will play a crucial role in subsequent analysis and are given by

$$\lambda_s \equiv \text{Cov}(s, m) = \sigma_s^2 + \rho_{s,b}\sigma_s\sigma_b, \tag{1}$$

$$\lambda_b \equiv \text{Cov}(b, m) = \sigma_b^2 + \rho_{s,b}\sigma_s\sigma_b.$$

The error term ε is drawn from the normal distribution $N(0, \sigma_\varepsilon^2)$, with precision $\tau_\varepsilon \equiv \sigma_\varepsilon^{-2} > 0$. This random variable is independent of the distribution of political attitudes δ_v and the voters' prior belief. These statistical properties are known to both voters and political parties. Before the signal is realised, voters' uncertainty about the signal is measured by the variance

$$\text{var}(m) = \lambda_s + \lambda_b + \sigma_\varepsilon^2. \tag{2}$$

Political parties' beliefs on (s, b) may differ from voters' beliefs. Specifically, we assume that

⁹We will remark on this and other model assumptions at the end of this section.

the two political parties share a common belief about (s, b) , which is given by a normal distribution $\mathbf{N}(\widehat{\boldsymbol{\mu}}_0, \widehat{\boldsymbol{\Sigma}}_0)$ with

$$\widehat{\boldsymbol{\mu}}_0 = \begin{bmatrix} \widehat{\mu}_s \\ \widehat{\mu}_b \end{bmatrix} \quad \text{and} \quad \widehat{\boldsymbol{\Sigma}}_0 = \begin{bmatrix} \widehat{\sigma}_s^2 & \widehat{\rho}_{s,b} \widehat{\sigma}_s \widehat{\sigma}_b \\ \widehat{\rho}_{s,b} \widehat{\sigma}_s \widehat{\sigma}_b & \widehat{\sigma}_b^2 \end{bmatrix}.$$

The elements of $\widehat{\boldsymbol{\mu}}_0$ and $\widehat{\boldsymbol{\Sigma}}_0$ can be interpreted similarly as those of $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$. The implied covariances between (s, b) and m are denoted by $\widehat{\lambda}_s \equiv Cov_p(s, m)$ and $\widehat{\lambda}_b \equiv Cov_p(b, m)$, where the subscript ‘‘p’’ indicates that these moments are derived from the parties’ belief. From the parties’ perspective, the variance of the signal is given by

$$var_p(m) = \widehat{\lambda}_s + \widehat{\lambda}_b + \sigma_\varepsilon^2.$$

Political parties are both office-motivated and policy-motivated. This means they not only care about winning, but also the policy implemented by the winner of the election. The parties’ preferences on policy x are represented by

$$V(x; \phi_k) = -(x - \phi_k)^2,$$

where $\phi_k \in \mathbb{R}$ is the ideal policy of party $k \in \{L, R\}$ with $\phi_R = -\phi_L = \phi > 0$. If R wins, then x_R is implemented and it receives a payoff of $-(x_R - \phi)^2 + \gamma$, where $\gamma \geq 0$ represents the additional benefits of holding office. If R loses, then its payoff is $-(x_L - \phi)^2$. The payoffs for L are defined symmetrically.

Events in the model unfold in three stages: First, parties L and R simultaneously propose policies x_L and x_R , respectively. Both parties are fully aware of the median value of δ_v , the probability distribution of ε and the voters’ prior belief about (s, b) . But they do not observe the realised value of (m, s, b) when they choose their platforms. In the second stage, the signal m is realised and each voter votes for a party. Finally, the party that garners a majority of votes wins and implements its proposed policy. Our equilibrium concept is subgame perfect Nash equilibrium.

Remarks on Model Assumptions Before proceeding further, we remark on several key assumptions in our model. The first one is the assumption that all voters share the same prior belief about (s, b) and receive the same signal. This is mainly for simplifying the analysis. Note

that if we allow for heterogeneity in both δ_v and the learning process (e.g., heterogeneous priors among voters or privately observed signals), then there will not be a single, deterministic decisive voter. Instead, the election outcome will be decided by voters whose ideal policy after observing the signals (δ_v^*) is at the median position among all voters, i.e., any v that satisfies

$$\delta_{med}^* = \delta_v + E_v(s | \mathcal{I}_v),$$

where \mathcal{I}_v denote the information received by voter v , which may include both public and private signals; and $E_v(s | \mathcal{I}_v)$ is the posterior mean of s conditional on \mathcal{I}_v . This will greatly complicate the analysis and also make it harder to interpret the results. Hence, we do not follow this route.

The second major assumption is that the two political parties share a common belief about (s, b) . This is mainly for the purpose of studying *symmetric* equilibrium. It is standard in the literature (Persson and Tabellini, 2002) to assume that the two political parties have the same utility function, receive the same benefits from holding office and that their policy ideals are equally distanced on both sides of the centrist position, for the purpose of obtaining a symmetric equilibrium.¹⁰ Our common prior assumption between the two parties can be viewed as a natural extension of this tradition.

Finally, it is assumed that there is no disagreement between voters and politicians regarding signal precision. Our analysis can be easily extended to accommodate this type of disagreement.

2.2 Voting decision

After observing the signal, voters update their belief about (s, b) using Bayes' rule. The resulting posterior belief is again a bivariate normal distribution. For the purpose of our analysis, it suffices to focus on the marginal distribution of s in the posterior belief, which is characterised in Lemma 1. All proofs can be found in the Online Supplementary Materials.

Lemma 1. *The marginal distribution of s in the voters' posterior belief is a normal distribution with mean*

$$E(s | m) = \mu_s + \frac{Cov(s, m)}{var(m)}(m - \mu_s - \mu_b), \quad (3)$$

and variance

$$var(s | m) = \sigma_s^2 - \frac{[Cov(s, m)]^2}{var(m)}, \quad (4)$$

¹⁰Studies that follow this approach include Ossokina and Swank (2004), Bernhardt *et al.* (2009) and Xefteris and Zudenkova (2018), among many others.

where $Cov(s, m)$ and $var(m)$ are given in (1) and (2), respectively.

Equation (3) shows that only the difference $(m - \mu_s - \mu_b)$ matters when forming the posterior expectation $E(s | m)$. In particular, voters will adjust the observed signal either upward or downward according to the prior means (μ_s, μ_b) . Equation (4) shows that new information always reduces voters' uncertainty about the hidden state, i.e., $var(s | m) < \sigma_s^2$. This is true even when the signal is perceived to be biased and when it is negatively correlated to s . The size of the reduction (i.e., the gain from learning) is negatively related to σ_ε^2 , which means more can be learned from a more precise signal.

Given the posterior belief about s , voter v 's expected utility from policy x is given by

$$E[U(x; \delta_v, s) | m] = E[-(\delta_v + s - x)^2 | m]. \quad (5)$$

The voter's ideal policy, δ_v^* , is one that maximises (5), i.e.,

$$\begin{aligned} \delta_v^* &\equiv \arg \max_{x \in \mathbb{R}} \left\{ E[-(\delta_v + s - x)^2 | m] \right\} \\ &= \delta_v + E(s | m). \end{aligned} \quad (6)$$

Equations (3) and (6) together show how voters use the observed signal to form their policy preferences. If $x_R = x_L$, then voters are indifferent between the two parties. If $x_R \neq x_L$, then after observing m , voter v will choose x_R over x_L if and only if

$$\begin{aligned} -(\delta_v^* - x_R)^2 &> -(\delta_v^* - x_L)^2 \\ \Leftrightarrow (x_R - x_L)(\delta_v^* - \bar{x}) &> 0, \end{aligned} \quad (7)$$

where $\bar{x} = (x_L + x_R)/2$. Hence, voter v will support R if either (i) $x_R > x_L$ and $\delta_v^* > \bar{x}$, or (ii) $x_R < x_L$ and $\delta_v^* < \bar{x}$. The voter is indifferent between any $x_R \neq x_L$ if $\bar{x} = \delta_v^*$.

This analysis implies that the median voter theorem holds and the decisive median voter is the one with $\delta_v = 0$.

2.3 Policy decision

We now focus on the first stage of events and characterise the parties' policy choices. We begin by formulating the parties' winning probability. If $x_L = x_R$, then the winner is decided by a fair coin toss. Suppose $x_L \neq x_R$. R wins if it gains the median voter's support. Since the median value

of δ_v is zero, equation (6) implies that the median voter's ideal policy is captured by $E(s | m)$ alone. Hence, R wins if either (i) $x_R > x_L$ and $E(s | m) > \bar{x}$, or (ii) $x_R < x_L$ and $E(s | m) < \bar{x}$. The value of $E(s | m)$, however, is unknown to the parties as m is not yet revealed at this stage.

From the parties' perspective, $E(s | m)$ is a normal random variable with mean

$$\tilde{\mu} \equiv E_p[E(s | m)] = \mu_s + \frac{\text{Cov}(s, m)}{\text{var}(m)} [(\hat{\mu}_s - \mu_s) + (\hat{\mu}_b - \mu_b)] \quad (8)$$

and variance

$$\tilde{\sigma}^2 \equiv \text{var}_p\{E(s | m)\} = \left[\frac{\text{Cov}(s, m)}{\text{var}(m)} \right]^2 \text{var}_p(m). \quad (9)$$

A higher value of $\tilde{\sigma}^2$ means that the parties are more uncertain about the median voter's ideal policy and hence the election outcome. This variance is in the centre stage of our analysis, and will be referred to as the parties' perceived uncertainty of the election outcome. Equation (9) shows that this variable not only depends on the parties' subjective belief, but also on the voters' prior belief and the signal precision.

Let $H(\cdot)$ be the cumulative distribution function of $N(\tilde{\mu}, \tilde{\sigma}^2)$, and $h(\cdot)$ be the corresponding probability density function. Then R 's winning probability is given by

$$\Pr(R \text{ wins}) = \begin{cases} 1/2 & \text{if } x_R = x_L, \\ 1 - H(\bar{x}) & \text{if } x_R > x_L, \\ H(\bar{x}) & \text{if } x_R < x_L. \end{cases} \quad (10)$$

Notice that apart from $x_R = x_L$, the two parties have equal opportunity of winning if \bar{x} coincides with the median of $N(\tilde{\mu}, \tilde{\sigma}^2)$, which is $\tilde{\mu}$. We will refer to this as the centrist policy position. Party R is deemed as the "right-wing" party if its ideal policy $\phi_R = \phi$ is on the right side of the centrist position, i.e., $\phi > \tilde{\mu}$. Similarly, party L is the left-wing party if $\phi_L = -\phi < \tilde{\mu}$. To preserve the symmetry between parties, $\tilde{\mu}$ is set to zero from this point onward. This is achieved by setting $\mu_s = 0$ in the voters' prior belief and having $\boldsymbol{\mu}_0 = \hat{\boldsymbol{\mu}}_0$ so that voters' and parties' beliefs differ only in the covariance matrices.

Taking $x_L \in \mathbb{R}$ as given, party R 's policy choice problem is to choose $x_R \in \mathbb{R}$ so as to maximise its expected utility

$$\mathcal{W}_R(x_R; x_L) = \left[-(x_R - \phi)^2 + \gamma \right] \Pr(R \text{ wins}) - (x_L - \phi)^2 [1 - \Pr(R \text{ wins})],$$

subject to (10). Party L 's expected utility $\mathcal{W}_L(x_L; x_R)$ is similarly defined.

We denote the equilibrium policies as $(x_L^*, x_R^*) \in \mathbb{R}^2$. In particular, we are interested in symmetric equilibrium, i.e., one in which x_L^* and x_R^* are equidistant on both sides of the centrist position so that $x_R^* = -x_L^* = x_{eq}^* \geq 0$. Policy convergence [resp., divergence] is said to occur if $x_{eq}^* = 0$ [resp., $x_{eq}^* > 0$].

The following result, which is due to Bernhardt *et al.* (2009, Corollary 2), provides a detailed characterisation of symmetric equilibrium. For the sake of completeness, we present the proof in Section B of the Online Supplementary Materials. The model in Bernhardt *et al.* (2009) differs from ours in two important regards: In their framework, voters observe perfectly the realised value of s before the election (hence no learning) and share the same prior beliefs with politicians (no disagreement). Consequently, $E_p[E(s | m)]$ and $var_p\{E(s | m)\} \equiv \tilde{\sigma}^2$ reduce, respectively, to the exogenously given prior mean and prior variance of s . In the present study, voters have imperfect information about the state of the world and $\tilde{\sigma}$ is endogenously (and jointly) determined by their learning process and politicians' belief. This allows us to examine how the quality of information possessed by the two groups and the disagreement between them will affect political polarisation.

Proposition 1.

- (a) *If $\phi \leq \gamma h(0)/2$, then there exists a unique symmetric equilibrium in which both parties choose the same policy which is the centrist position, i.e., $x_{eq}^* = 0$.*
- (b) *If $\phi > \gamma h(0)/2$, then there exists a unique symmetric equilibrium in which the two parties choose different policies, i.e., $x_R^* = -x_L^* = x_{eq}^* > 0$, and x_{eq}^* is given by*

$$x_{eq}^* = \frac{2\phi - \gamma h(0)}{4h(0)\phi + 2}. \tag{11}$$

Proposition 1 shows that the additional benefits of holding office γ (which captures the strength of the parties' office motivation) must be sufficiently large in order to induce the parties to sacrifice their own political ideals (policy motivation) and move towards their opponent's policy position. Since $h(0) \equiv 1/(\tilde{\sigma}\sqrt{2\pi})$ for the normal distribution $N(0, \tilde{\sigma}^2)$,

$$\phi \leq \gamma h(0)/2 \quad \text{if and only if} \quad \tilde{\sigma} \leq \sigma_{\min} \equiv \gamma/(2\sqrt{2\pi}\phi).$$

This states that σ_{\min} is a unique threshold value of $\tilde{\sigma}$ below [resp., above] which policy convergence

[resp., divergence] will emerge. In addition, as can be seen by differentiating (11) with respect to $\tilde{\sigma}$, the extent of policy divergence is increasing in $\tilde{\sigma}$. Intuitively, higher uncertainty on the median voter's ideal policy makes policy moderation less rewarding for politicians. Hence, policy polarisation is more likely to happen and more severe when the parties are more uncertain about the election outcome. This central observation forms the basis for the analysis in Section 3 and is summarised in part (a) of Corollary 1. The second part of the corollary will prove useful in the welfare analysis in Section 4. It states that policy polarisation becomes more severe when the parties' ideologies are further apart. But there is a limit: as ϕ increases indefinitely, x_{eq}^* approaches the upper bound $\sqrt{\pi/2} \cdot \tilde{\sigma}$.

Corollary 1. *Assume $\phi > \gamma h(0)/2$. Then the following results hold.*

- (a) *The extent of polarisation x_{eq}^* is strictly increasing in $\tilde{\sigma}$.*
- (b) *The extent of polarisation x_{eq}^* is strictly increasing in ϕ and bounded above by $\sqrt{\pi/2} \cdot \tilde{\sigma}$.*

3 Equilibrium Polarisation

In this section, we examine how the quality of information possessed by voters and political parties will affect the extent of policy polarisation in equilibrium. We focus on two aspects of information and beliefs, namely (i) the precision of the error term in the signal and (ii) the disagreement between voters' and parties' prior beliefs, as captured by the differences between $\{\lambda_s, \lambda_b\}$ and $\{\hat{\lambda}_s, \hat{\lambda}_b\}$.¹¹ In light of Corollary 1, it suffices to analyse how these parameters will affect the parties' perceived uncertainty ($\tilde{\sigma}^2$) in the decision stage. In particular, any factor that makes the parties more uncertain about the election outcome will make polarisation more likely to emerge and more severe, while the opposite is true when $\tilde{\sigma}^2$ is reduced.

The effects of beliefs and information on polarisation can be broadly classified into two groups. In order to explain these, we first rewrite (3) and (9) as

$$E(s | m) = \psi m \quad \text{and} \quad \tilde{\sigma}^2 = \psi^2 \text{var}_p(m),$$

where the first equality is obtained by setting $\mu_s = \mu_b = 0$ and ψ is defined as

$$\psi \equiv \frac{\text{Cov}(s, m)}{\text{var}(m)} = \frac{\lambda_s}{\lambda_s + \lambda_b + \tau_\varepsilon^{-1}}. \quad (12)$$

¹¹The rationale for focusing on these covariances will be explained later in the section.

Remember that we assume $\mu_s = 0$ to preserve the symmetry between the parties. On the other hand, $\mu_b = 0$ is assumed solely for notational simplicity and has no bearing on the results.¹²

The coefficient ψ captures the responsiveness of voters' posterior expectation to the signal. It is solely determined by the learning process, therefore it depends on the quality of information available to the voters and their subjective prior belief but not on the parties' belief. Any increase in the *magnitude* of ψ , i.e., $|\psi|$ or ψ^2 , will encourage polarisation by raising the parties' perceived uncertainty about the election outcome. We refer to this mechanism as the *learning effect*. The variance $var_p(m)$, on the other hand, represents the parties' subjective uncertainty about the signal. It thus depends solely on the information available to them and their subjective belief, but not on the voters'. Any increase in $var_p(m)$ will increase the uncertainty faced by the parties in the decision stage and promote polarisation. We refer to this as the *uncertainty effect*.

The above description makes clear that voters' and political parties' prior beliefs each affect polarisation through a different pathway. In order to disentangle these two effects, it is necessary to decouple these two sets of beliefs.

3.1 Signal Precision and Polarisation

Any change in signal precision will induce both learning and uncertainty effects. On the one hand, higher signal precision will encourage voters to become more responsive to the signal when updating their belief. This will enhance polarisation by strengthening the learning effect. On the other hand, a more precise signal means that it is more predictable or less uncertain. This will alleviate the uncertainty effect and lower polarisation. These results are formally stated in Lemma 2.

Lemma 2. *An increase in signal precision will strengthen the learning effect [i.e. $d\psi^2/d\tau_\varepsilon > 0$] and weaken the uncertainty effect [i.e. $d[var_p(m)]/d\tau_\varepsilon < 0$].*

Proposition 2 below provides a necessary and sufficient condition under which one effect will dominate the other.

Proposition 2. *Holding $\{\lambda_s, \lambda_b, \hat{\lambda}_s, \hat{\lambda}_b\}$ constant,*

$$\frac{d\tilde{\sigma}^2}{d\tau_\varepsilon} \geq 0 \quad \text{if and only if} \quad 2var_p(m) \geq var(m). \quad (13)$$

¹²This is because voters will subtract μ_b from the realised signal when forming their posterior estimate, so that $E(s | m) = \psi(m - \mu_b)$. All of our subsequent results will remain valid if we replace m with $\hat{m} \equiv m - \mu_b$. In particular, the second moments $var(m)$ and $var_p(m)$ are the same as $var(\hat{m})$ and $var_p(\hat{m})$.

The condition in (13) is derived from the following decomposition of $\ln \tilde{\sigma}^2$,

$$\frac{d \ln \tilde{\sigma}^2}{d \ln \tau_\varepsilon} = \frac{d \ln \psi^2}{d \ln \tau_\varepsilon} + \frac{d \ln \text{var}_p(m)}{d \ln \tau_\varepsilon}. \quad (14)$$

The first term on the right captures the changes in $\tilde{\sigma}^2$ due to the learning effect, while the second term captures the contribution of the uncertainty effect. As shown in the proof of Proposition 2, the first term is inversely related to $\text{var}(m)$, i.e.,

$$\frac{d \ln \psi^2}{d \ln \tau_\varepsilon} = \frac{2}{\tau_\varepsilon} \frac{1}{\text{var}(m)} > 0. \quad (15)$$

If voters are highly uncertain about the signal to begin with (e.g., due to a large value of $\lambda_s + \lambda_b$), then a one-percentage increase in τ_ε will have a small impact on $\text{var}(m)$ and the outcome of the learning process. Similarly, the contribution of the uncertainty effect is determined by

$$\frac{d \ln \text{var}_p(m)}{d \ln \tau_\varepsilon} = -\frac{1}{\tau_\varepsilon} \frac{1}{\text{var}_p(m)} < 0. \quad (16)$$

Equation (16) states that the responsiveness of $\text{var}_p(m)$ to τ_ε (i.e., the magnitude of the derivative) tends to be small if the political parties are highly uncertain about the signal to begin with (e.g., due to a large value of $\hat{\lambda}_s + \hat{\lambda}_b$). The condition in (13) follows by combining (14)-(16).

If voters and political parties share a common prior belief so that $\lambda_s = \hat{\lambda}_s$ and $\lambda_b = \hat{\lambda}_b$, then they are equally uncertain about the signal, i.e., $\text{var}(m) = \text{var}_p(m)$. Proposition 2 then implies that a more precise signal will *always* lead to an increase in $\tilde{\sigma}^2$ and hence polarisation.¹³ In other words, the learning effect always dominates under the common prior assumption. The dominance of the learning effect remains valid if $2\text{var}_p(m) > \text{var}(m)$, but is overturned when $2\text{var}_p(m) < \text{var}(m)$. To give a concrete example, suppose both voters and political parties expect the bias term to be independent of the hidden state, so that $\rho_{s,b} = \hat{\rho}_{s,b} = 0$. In addition, they share the same degree of uncertainty regarding the bias, so that $\sigma_b^2 = \hat{\sigma}_b^2 > 0$. These two together imply $\lambda_b = \hat{\lambda}_b$, so that the two groups differ only in terms of σ_s^2 and $\hat{\sigma}_s^2$. Then an improvement in signal precision will reduce polarisation if the voters are sufficiently more uncertain (or less knowledgeable) on the hidden state, i.e. $\sigma_s^2 > 2\hat{\sigma}_s^2 + \sigma_b^2 + \sigma_\varepsilon^2$, so that $\text{var}(m) > 2\text{var}_p(m)$ is implied. This example shows that disagreement in σ_b^2 or in λ_b are not necessary for the result in Proposition 2.

¹³A similar result is reported in Gul and Pesendorfer (2012, Lemma 2). The main focus of Gul and Pesendorfer (2012), however, is on the relation between media competition and party polarisation.

We now examine how the relationship between signal precision and polarisation changes as signal precision increases. Expanding the condition $2\text{var}_p(m) \geq \text{var}(m)$ gives

$$2(\widehat{\lambda}_s + \widehat{\lambda}_b) + 2\tau_\varepsilon^{-1} \geq \lambda_s + \lambda_b + \tau_\varepsilon^{-1}$$

$$\Leftrightarrow \tau_\varepsilon^{-1} \geq \lambda_s + \lambda_b - 2(\widehat{\lambda}_s + \widehat{\lambda}_b).$$

If $2(\widehat{\lambda}_s + \widehat{\lambda}_b) \geq \lambda_s + \lambda_b$, which includes the case when there is no disagreement, then

$$\tau_\varepsilon^{-1} > \lambda_s + \lambda_b - 2(\widehat{\lambda}_s + \widehat{\lambda}_b) \quad \text{and} \quad 2\text{var}_p(m) > \text{var}(m)$$

are true for any $\tau_\varepsilon > 0$. Then by Proposition 2 and Corollary 1, both the parties' perceived uncertainty and equilibrium polarisation are monotonically increasing in τ_ε . This happens because the magnitude of the uncertainty effect is limited and insignificant when $(\widehat{\lambda}_s + \widehat{\lambda}_b)$ is large [equation (16)]. Consequently, the positive learning effect always dominates when signal precision improves. But if $\lambda_s + \lambda_b > 2(\widehat{\lambda}_s + \widehat{\lambda}_b)$, then there exists a unique threshold value, given by

$$\tau'_\varepsilon \equiv \frac{1}{\lambda_s + \lambda_b - 2(\widehat{\lambda}_s + \widehat{\lambda}_b)} > 0, \quad (17)$$

such that $\tilde{\sigma}^2$ is strictly increasing when $\tau_\varepsilon < \tau'_\varepsilon$ and strictly decreasing when $\tau_\varepsilon > \tau'_\varepsilon$. This hump-shaped pattern suggests that the positive learning effect first dominates when τ_ε is low, but it is eventually overpowered by the negative uncertainty effect as signal precision continues to increase. This hump-shaped pattern is possible only if disagreement between voters and politicians is present. We summarise these observations in Corollary 2.

Corollary 2.

- (a) *Suppose $2(\widehat{\lambda}_s + \widehat{\lambda}_b) \geq \lambda_s + \lambda_b$. Then the parties' perceived uncertainty $\tilde{\sigma}^2$ is strictly increasing in τ_ε .*
- (b) *Suppose $\lambda_s + \lambda_b > 2(\widehat{\lambda}_s + \widehat{\lambda}_b)$. Then there exists a unique threshold value $\tau'_\varepsilon > 0$ such that*

$$\frac{d\tilde{\sigma}^2}{d\tau_\varepsilon} \geq 0 \quad \text{if and only if} \quad \tau_\varepsilon \leq \tau'_\varepsilon.$$

Extension In Section C of the Online Supplementary Materials, we present an extended model featuring $n \geq 1$ potentially biased signals and derive a general version of (3) and (4). We then explore two special cases of the multiple-signal model in Sections D and E of the online materials. In there, we show that the result in Proposition 2 remains robust in an environment with multiple unbiased, correlated signals. We also show that reducing the correlation between signals has the same effect as increasing their precision, as mentioned in the Introduction.

3.2 More on Learning Effect

Apart from signal precision, the learning effect is also shaped by the voters' initial belief, as captured by λ_s and λ_b . The purpose of this section is to examine how changes in λ_s and λ_b will affect the learning effect. The results reported below can help us better understand the determinants of this effect.

The coefficient ψ , which captures the learning effect, is strongly influenced by the voters' perception of the bias term. To highlight the significance of this perception, we contrast our model to the "conventional" model in which the signal is truly unbiased (i.e., $b = 0$ with probability one) and it is common knowledge among the voters. This means there is no uncertainty about b in their initial belief so that $\sigma_b = \lambda_b = 0$. It follows that $\lambda_s = \sigma_s^2 > 0$ and

$$\psi = \frac{\lambda_s}{\lambda_s + \tau_\varepsilon^{-1}} > 0, \quad (18)$$

Therefore, in the conventional model, voters will always update their belief in the same direction as suggested by the signal.

This is not necessarily true when they are unsure about the bias term. In this case, the covariance λ_s can be either positive or negative, depending on the parameters in Σ_0 . If $\lambda_s > 0$, and hence $\psi > 0$, voters will respond to the signal as in the conventional model. We refer to this as *conventional learning*. But if $\lambda_s < 0$ so that $\psi < 0$, voters will revise their estimate for the hidden state in the *opposite* direction as suggested by the signal. We refer to this as *defiant learning*.¹⁴ Lemma 3 states the conditions under which defiant learning will emerge.¹⁵

¹⁴In the knife-edge case in which $\sigma_s + \rho_{s,b}\sigma_b = 0$, both λ_s and ψ become zero. This means voters do not respond to the signal at all, so that $E(s | m) = \mu_s = 0$. We do not consider this case in our analysis.

¹⁵It should be understood that conventional learning prevails under any $\sigma_s > 0$, $\sigma_b > 0$ and $\rho_{s,b} \in (-1, 1)$ that do not satisfy the conditions in (19).

Lemma 3. *Defiant learning happens if and only if*

$$0 < \sigma_s < \sigma_b \quad \text{and} \quad -1 < \rho_{s,b} < -\frac{\sigma_s}{\sigma_b}. \quad (19)$$

The above lemma establishes that defiant learning arises when (i) voters are more uncertain about the bias than about the hidden state, and (ii) the two unobservables, s and b , are perceived to be strongly negatively correlated. Greater uncertainty about the bias (i.e., a higher value of σ_b) expands the range of $\rho_{s,b}$ values for which defiant learning occurs.

As an illustration, suppose voters receive a news report claiming that a mass vaccination programme is more effective than previously expected. However, they are uncertain about the biasedness of the news channel (i.e., $\sigma_b > 0$). In particular, they suspect that the bias is negatively correlated with the true programme efficacy (i.e., $\rho_{s,b} < 0$). In this case, they may attribute the favourable report to a stronger-than-usual bias rather than genuine performance. When this suspicion is sufficiently strong (i.e., $\lambda_s < 0 < \lambda_b$), voters update their beliefs downward despite receiving seemingly positive information.

This mechanism depends on the negative correlation $\rho_{s,b} < 0$ but not on the mean bias μ_b . Hence, defiant learning can arise regardless of whether the source is expected to be unbiased on average ($\mu_b = 0$) or systematically biased ($\mu_b \neq 0$).

Given the opposite responses under conventional and defiant learning, it is natural to ask whether and how the learning effect may differ between the two. Regarding improvement in signal precision, Lemma 2 shows that an increase in τ_ε *always* strengthens the learning effect, i.e. under both conventional and defiant learning. At first glance, this seems to be no different from the prediction of the conventional model [equation (18)], but there is an important difference. Under defiant learning, more responsive voters means that they are more defiant (i.e., ψ becomes more negative).

Our next result examines how changes in λ_s or λ_b will affect the strength of the learning effect. As is evident from (12), the covariances λ_s and λ_b summarise the effects of $\{\sigma_s, \sigma_b, \rho_{s,b}\}$ on ψ . Focusing on changes in λ_s and λ_b proves to be a more fruitful approach in organising our results and generating intuition. It is also straightforward to compare to the conventional model, which corresponds to the case when $\lambda_s = \sigma_s^2$ and $\lambda_b = 0$.

We note that the same change in λ_s or λ_b can be caused by more than one type of changes in $\{\sigma_s, \sigma_b, \rho_{s,b}\}$. Likewise, a change in one of $\{\sigma_s, \sigma_b, \rho_{s,b}\}$ generally affects both λ_s and λ_b . For

instance, when $\rho_{s,b} \neq 0$ an increase in σ_s will affect λ_s and λ_b simultaneously. In the proof of Proposition 3, we will explain in detail how one can relate changes in $\{\lambda_s, \lambda_b\}$ to those in $\{\sigma_s, \sigma_b, \rho_{s,b}\}$.

Proposition 3.

- (a) *Under both conventional and defiant learning, the learning effect is strengthened when there is a decrease in λ_b , i.e., $d\psi^2/d\lambda_b < 0$.*
- (b) *Under conventional learning (i.e., when $\lambda_s > 0$),*

$$\frac{d\psi^2}{d\lambda_s} \geq 0 \quad \text{if and only if} \quad \lambda_b + \tau_\varepsilon^{-1} \geq 0.$$

- (c) *Under defiant learning (i.e., when $\lambda_s < 0$), the learning effect is strengthened when there is a decrease in λ_s , i.e., $d\psi^2/d\lambda_s < 0$.*

Two main lessons can be learned from Proposition 3. First, holding λ_s and τ_ε constant, lowering the value of λ_b will induce a one-to-one decrease in $var(m)$ and have the same effect as an improvement in signal precision. As an illustration, suppose now we compare the current model with $\{\lambda_s, \lambda_b, \tau_\varepsilon\}$ and $\lambda_b < 0$ to the conventional model with the same $\{\lambda_s, \tau_\varepsilon\}$ but no bias, i.e. $\lambda_b = 0$.¹⁶ Then part (a) of the proposition states that the learning effect is stronger in the former than in the latter. This happens because when λ_b is strictly negative, the bias term tends to mitigate or counteract the effect of the hidden state [i.e., $\lambda_b < 0$ implies $\rho_{s,b} < -\sigma_b/\sigma_s$], which then reduce the variance of the signal.

Second, in the presence of an unknown bias, the learning effect may become weaker when the signal is more positively associated to the hidden state. This is in stark contrast with the prediction of the conventional model. Intuitively, a higher positive covariance between m and s suggests that there is more to learn from the signal. It is, however, important to note that any increase in λ_s will increase both the numerator and denominator in (18) [and likewise in (12)]. The learning effect is strengthened if and only if the increase in λ_s induces a less than proportionate increase in $var(m)$.¹⁷ In particular, when $\lambda_s > 0$ and $\lambda_b + \tau_\varepsilon^{-1} < 0$, any increase in λ_s will induce a greater than proportionate increase in $var(m)$ and weaken the learning effect. The prediction

¹⁶More precisely, we are comparing the current model with parameters $\{\sigma_s, \sigma_b, \rho_{s,b}\}$, where $\lambda_b = \sigma_b^2 + \rho_{s,b}\sigma_s\sigma_b < 0$, to a conventional model with parameters $\{\bar{\sigma}_s, \bar{\sigma}_b, \bar{\rho}_{s,b}\}$, where $\bar{\sigma}_b = \bar{\rho}_{s,b} = 0$ and $\bar{\sigma}_s^2 = \lambda_s = \sigma_s^2 + \rho_{s,b}\sigma_s\sigma_b$.

¹⁷Recall that the real-valued function $f(x) = x/(x+c)$ is strictly increasing if $c > 0$ and strictly decreasing if $c < 0$.

of the conventional model is also invalid under defiant learning. When λ_s is negative, any increase means that it moves closer to zero, which makes the voters less responsive to the signal. As a result, the learning effect is diminished.

It is also possible to derive the effects of $\{\sigma_s, \sigma_b, \rho_{s,b}\}$ on ψ^2 , but the conditions are in general more complicated and harder to interpret. Here we only present the results related to changes in $\rho_{s,b}$. Further results can be found in Section F of the Online Supplementary Materials.

Lemma 4.

(a) *Suppose $\sigma_b \leq \sigma_s$. Then*

$$\frac{d\psi^2}{d\rho_{s,b}} \geq 0 \quad \text{if and only if} \quad \sigma_b^2 + \tau_\varepsilon^{-1} \geq \sigma_s^2.$$

(b) *Suppose $\sigma_b > \sigma_s$. Then*

$$\frac{d\psi^2}{d\rho_{s,b}} \geq 0 \quad \text{if and only if} \quad \rho_{s,b} \geq -\frac{\sigma_s}{\sigma_b}.$$

Using the definition of $Cov(s, m)$ and $var(m)$, we can rewrite (12) as

$$\psi = \frac{Cov(s, m)}{var(m)} = \frac{\sigma_s^2 + \rho_{s,b}\sigma_s\sigma_b}{\sigma_s^2 + \sigma_b^2 + \sigma_\varepsilon^2 + 2\rho_{s,b}\sigma_s\sigma_b}. \quad (20)$$

From this equation, it is evident that any changes in $\rho_{s,b}$ will affect both $Cov(s, m)$ and $var(m)$. As shown in Lemma 3, if $\sigma_b \leq \sigma_s$ then defiant learning is not possible. In this case, an increase in $\rho_{s,b}$ will intensify the learning effect if and only if it induces a larger percentage increase in $Cov(s, m)$ than in $var(m)$. This happens when $\sigma_b^2 + \tau_\varepsilon^{-1} > \sigma_s^2$. When $\sigma_b > \sigma_s$, defiant learning happens if $-1 < \rho_{s,b} < -\sigma_s/\sigma_b$ is true. When $\rho_{s,b}$ increases toward $-\sigma_s/\sigma_b$ within this range, both $Cov(s, m)$ and ψ^2 will approach zero which means there is not much to learn from the signal, hence the learning effect diminishes. But once $\rho_{s,b}$ increases beyond $-\sigma_s/\sigma_b$, conventional learning emerges. Any further increase in $\rho_{s,b}$ will increase $Cov(s, m)$ more (in percentage terms) than $var(m)$, and strengthen the learning effect. This gives rise to a non-monotonic relationship between $\rho_{s,b}$ and ψ^2 when $\sigma_b > \sigma_s$.

We conclude this section by pointing out the implications of Proposition 3 on equilibrium polarisation. When voters and political parties have different prior beliefs so that $(\lambda_s, \lambda_b) \neq (\widehat{\lambda}_s, \widehat{\lambda}_b)$, λ_s and λ_b will only affect ψ^2 but not $var_p(m)$. It follows that their effects on ψ^2 can

be directly translated into their effects on $\tilde{\sigma}^2$ and polarisation. The results are summarised in Corollary 3 which follows immediately from Proposition 3, hence the proof is omitted.

Corollary 3 *Suppose $(\lambda_s, \lambda_b) \neq (\hat{\lambda}_s, \hat{\lambda}_b)$.*

- (a) *Holding λ_s and τ_ε constant, a decrease in λ_b will promote polarisation.*
- (b) *Suppose $\lambda_s > 0$ and $\lambda_b + \tau_\varepsilon^{-1} > 0$. Then, holding λ_b and τ_ε constant, an increase in λ_s will promote polarisation.*
- (c) *Suppose either (i) $\lambda_s > 0$ and $\lambda_b + \tau_\varepsilon^{-1} < 0$, or (ii) $\lambda_s < 0$. Then holding λ_b and τ_ε constant, an increase in λ_s will suppress polarisation.*

3.3 More on Uncertainty Effect

Apart from signal precision, the uncertainty effect is also determined by the parties' initial belief, as captured by $\hat{\lambda}_s$ and $\hat{\lambda}_b$. Since $var_p(m) = \hat{\lambda}_s + \hat{\lambda}_b + \tau_\varepsilon^{-1}$, any increase in either $\hat{\lambda}_s$ or $\hat{\lambda}_b$ will induce a one-to-one increase in $var_p(m)$ and strengthen the uncertainty effect. When disagreement is present so that $(\lambda_s, \lambda_b) \neq (\hat{\lambda}_s, \hat{\lambda}_b)$, neither $\hat{\lambda}_s$ nor $\hat{\lambda}_b$ will affect ψ . Therefore, any increase in $\hat{\lambda}_s$ or $\hat{\lambda}_b$ will directly contribute to an increase in the parties' perceived uncertainty and promote polarisation.

3.4 Disagreement and Polarisation

In this section, we are interested in the question: How does disagreement itself affect polarisation? To address this, we first rewrite the political parties' perceived uncertainty as

$$\tilde{\sigma}^2 = \underbrace{\frac{[Cov(s, m)]^2}{var(m)}}_{(\sigma^\dagger)^2} \cdot \frac{var_p(m)}{var(m)}. \quad (21)$$

Notice that

$$(\sigma^\dagger)^2 = \frac{[Cov(s, m)]^2}{var(m)} = \frac{[Cov(s, m)]^2}{var^2(m)} var(m) = \psi^2 var(m),$$

therefore $(\sigma^\dagger)^2$ represents voters' perceived uncertainty about the election outcome, which is the counterpart of parties' perceived uncertainty $\tilde{\sigma}^2$. If voters and parties share a common prior so that $var(m) = var_p(m)$, then (21) will be reduced to $\tilde{\sigma}^2 = (\sigma^\dagger)^2$. Equation (21) has two other immediate implications. First, disagreement does not always have an impact on

polarisation. More specifically, voters and political parties may have different prior beliefs so that $(\lambda_s, \lambda_b) \neq (\widehat{\lambda}_s, \widehat{\lambda}_b)$. But if they are equally uncertain about the signal so that $\text{var}(m) = \text{var}_p(m)$, which happens when $\lambda_s + \lambda_b = \widehat{\lambda}_s + \widehat{\lambda}_b$, then the differences in beliefs will have no impact on $\widetilde{\sigma}^2$ and hence polarisation.

Second, suppose $\text{var}(m) \neq \text{var}_p(m)$. Then disagreement can either enhance or suppress polarisation. In particular, disagreement is polarisation-enhancing (i.e., disagreement generates a larger extent of polarisation than that in the common prior case) if $\widehat{\lambda}_s + \widehat{\lambda}_b > \lambda_s + \lambda_b$ so that $\text{var}_p(m) > \text{var}(m)$; and polarisation-reducing if $\widehat{\lambda}_s + \widehat{\lambda}_b < \lambda_s + \lambda_b$. Going back to our earlier example, suppose $\rho_{s,b} = \widehat{\rho}_{s,b} = 0$ and $\sigma_b^2 = \widehat{\sigma}_b^2 > 0$ so that $\lambda_b = \widehat{\lambda}_b$. Then disagreement is polarisation-reducing if the political parties are more certain about the hidden state than the voters, i.e., $\widehat{\sigma}_s^2 < \sigma_s^2$.

4 Welfare Analysis

In this section we focus on voters' *ex ante* welfare (i.e., welfare before the signal is realised) in a polarised equilibrium. We present two sets of results. The first one concerns the welfare implications of parties' ideological differences in such an equilibrium. The second one concerns the welfare implications of an improvement in signal precision. In both instances, disagreement between voters' and politicians' beliefs plays a crucial role in shaping the results.

4.1 Ideological Polarisation

We start by deriving a measure of *ex ante* welfare. Consider a symmetric equilibrium in which $x_R^* = -x_L^* = x_{eq}^* \geq 0$. The existence and uniqueness of such an equilibrium is guaranteed by Proposition 1. Note that the magnitude of x_{eq}^* is known under given value of $\{\psi, \gamma, \widetilde{\sigma}\}$ and is independent of the signal. The realised value of m , however, decides which party wins. If the median voter's ideal policy, $E(s | m) = \psi m$, is strictly positive, then R wins and the implemented policy is x_{eq}^* . But if $\psi m < 0$, then L is the winner and $-x_{eq}^*$ will be implemented.¹⁸ Therefore, before m is realised, voter v 's expected utility is given by

$$W(x_{eq}^*; \delta_v) = \int_0^\infty E[U(x_{eq}^*; \delta_v, s) | m] dG(\psi m) + \int_{-\infty}^0 E[U(-x_{eq}^*; \delta_v, s) | m] dG(\psi m),$$

¹⁸Since m is drawn from a continuous distribution, the random event $\{\psi m = 0\}$ happens with zero probability. Hence, it is omitted.

for any $\delta_v \in \mathbb{R}$, where $G(\psi m)$ is the cumulative distribution function of ψm and $E[U(x; \delta_v, s) | m]$ is the expected utility from any implemented policy $x \in \mathbb{R}$ based on the posterior belief of s [see equation (5)]. The following lemma provides an explicit formula for $W(x_{eq}^*; \delta_v)$.

Lemma 5. *In any symmetric equilibrium with $x_R^* = -x_L^* = x_{eq}^* \geq 0$,*

$$W(x_{eq}^*; \delta_v) = \left[2\sqrt{\frac{2}{\pi}}\sigma^\dagger - x_{eq}^* \right] x_{eq}^* - (\delta_v^2 + \tau_s^{-1}), \quad (22)$$

where $\sigma^\dagger \equiv \sqrt{\text{var}(\psi m)} = |\psi| \cdot \sqrt{\text{var}(m)}$.

Equation (22) implies that all voters will agree on the same optimal level of polarisation, because δ_v induces only a constant downward shift in the expected utility.

In the convergent equilibrium, i.e., when $x_{eq}^* = 0$, voter v 's expected utility is simply

$$W(0; \delta_v) = -(\delta_v^2 + \tau_s^{-1}).$$

Using this, we can rewrite (22) as

$$W(x_{eq}^*; \delta_v) - W(0; \delta_v) = \left[2\sqrt{\frac{2}{\pi}}\sigma^\dagger - x_{eq}^* \right] x_{eq}^*, \quad (23)$$

which measures the welfare gain or loss due to policy polarisation in equilibrium. We will use this as our measure of *ex ante* welfare in any symmetric polarised equilibrium. Two observations are immediate from (23). First, it is independent of δ_v . Therefore, the analysis below applies unanimously to *all* voters. Second, this welfare measure is a hump-shaped function in x_{eq}^* . In particular, polarisation is welfare-improving if and only if

$$0 \leq x_{eq}^* = \frac{2\phi - \gamma h(0)}{4h(0)\phi + 2} \leq 2\sqrt{\frac{2}{\pi}}\sigma^\dagger, \quad (24)$$

where the equality in the middle comes from (11), and the gain in welfare is highest at the mid-point (x_{mid}) of this range. A graphical illustration of (23) is shown in Figure 1.

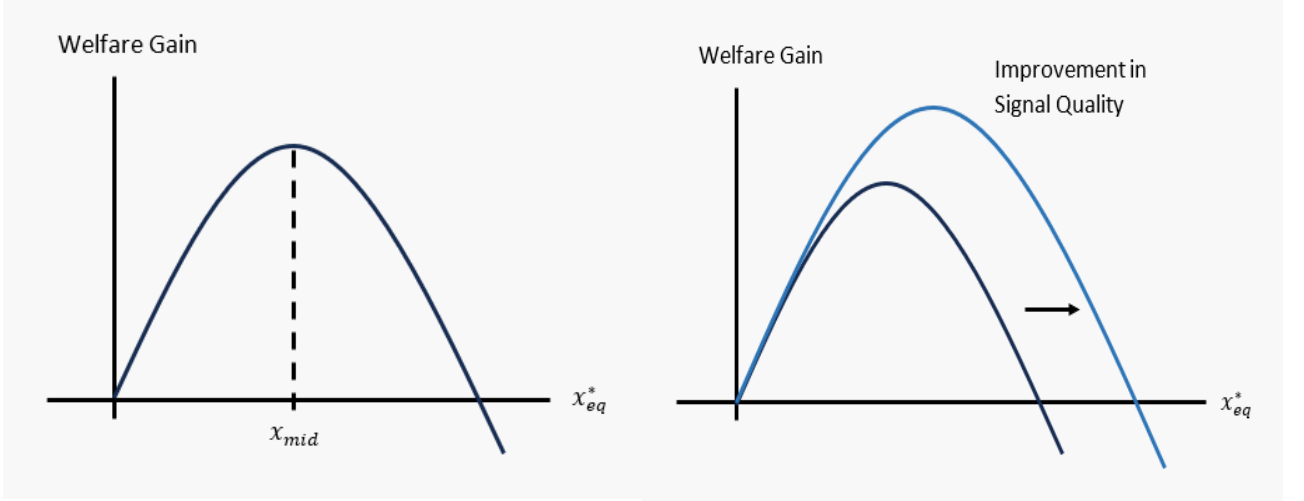


Figure 1

Figure 2

Policy polarisation can be welfare-improving as it provides voters a partial insurance against uncertainty on the state of the world. This is because polarisation gives voters some choice once the signal is realised and they know more about the state of the world. We refer to this as the *insurance effect* of policy polarisation.¹⁹ This insurance is more valuable to voters when they face greater uncertainty (i.e., when σ^\dagger increases). Notice that their perceived uncertainty increases when they expect to receive a more precise signal since they will be more responsive to it. In other words, σ^\dagger is strictly increasing in τ_ε . Holding x_{eq}^* constant, any increase in σ^\dagger will generate a greater welfare gain from polarisation. In terms of the diagram, this will take the form of an upward and outward shift of the graph (see Figure 2). The benefit of policy polarisation, however, is not without bound. When x_{eq}^* continues to increase, the welfare gain from divergent policies will diminish and eventually turn into a welfare loss.

We next analyse the welfare implications of parties' ideological divergence. The first inequality in (24), which is derived in Proposition 1, states that polarised equilibrium exists only if the two parties' ideological differences are sufficiently large. In other words,

$$x_{eq}^* \geq 0 \quad \text{if and only if} \quad \phi \geq \frac{\gamma h(0)}{2} = \frac{\gamma}{2\sqrt{2\pi}\tilde{\sigma}} \equiv \phi_{\min}.$$

Our next result examines the conditions under which the second inequality in (24) is also satisfied.

¹⁹See Bernhardt *et al.* (2009) and Aytimur *et al.* (2025) for a more in-depth discussion about this insurance (or hedging) effect.

Proposition 4.

(a) Suppose the following condition is satisfied,

$$\frac{\tilde{\sigma}}{\sigma^\dagger} \equiv \sqrt{\frac{\text{var}_p(m)}{\text{var}(m)}} \leq \frac{4}{\pi} \simeq 1.273. \quad (25)$$

Then polarisation is welfare-improving, i.e., $W(x_{eq}^*; \delta_v) \geq W(0; \delta_v)$, for any $x_{eq}^* \geq 0$, or equivalently, $\phi \geq \phi_{\min}$.

(b) Suppose the following condition is satisfied,

$$\frac{\tilde{\sigma}}{\sigma^\dagger} \equiv \sqrt{\frac{\text{var}_p(m)}{\text{var}(m)}} > \frac{4}{\pi} \simeq 1.273.$$

Then polarisation is welfare-improving if and only if

$$\phi_{\min} \leq \phi \leq \frac{\sqrt{\pi} (8\sigma^\dagger \cdot \tilde{\sigma} + \gamma)}{2\sqrt{2} (\pi\tilde{\sigma} - 4\sigma^\dagger)}. \quad (26)$$

The above results can be explained using the insurance effect of policy polarisation and the uncertainty effect mentioned in Section 3. The former states that the benefits of policy polarisation as an insurance mechanism increases when voters become more uncertain about the election outcome (i.e., when σ^\dagger increases). The latter states that when the parties become more uncertain about the signal and hence the election outcome (i.e., when $\tilde{\sigma}$ increases), they will have a greater incentive to polarise, leading to a higher value of x_{eq}^* .

Whether polarisation is welfare-improving or welfare-reducing ultimately depends on (i) the relative magnitude or ratio between $\tilde{\sigma}$ and σ^\dagger , and (ii) the extent of ideological polarisation ϕ . When condition (25) holds, *all* equilibrium levels of polarisation are welfare-improving. As shown in part (b) of Corollary 1, x_{eq}^* is strictly increasing in ϕ but bounded above by

$$\lim_{\phi \rightarrow \infty} x_{eq}^* = \sqrt{\frac{\pi}{2}} \tilde{\sigma}.$$

Polarisation is always welfare-improving if this limit lies within the range specified in (24), which happens under condition (25). This corresponds to an environment in which the insurance benefits from policy divergence are substantial (e.g., when voters hold imprecise prior beliefs or receive highly informative signals, leading to a large value of σ^\dagger), or in which the uncertainty effect remains weak (e.g., when politicians possess more precise prior belief about the signal), limiting

the value of $\tilde{\sigma}$ and the extent of policy polarisation. Note that condition (25) includes the case when voters and politicians share a common prior (i.e., $\tilde{\sigma} = \sigma^\dagger$).

The main message of part (a) can be summarised as follows: When there is no disagreement between voters and politicians, or when voters are more uncertain about the electoral outcome than the parties, all voters will be strictly better off in a society with partisan differentiation (i.e., $\phi > \phi_{\min}$ and $x_{eq}^* > 0$) than in an otherwise identical society but with more congruent parties and convergent platforms (i.e., $\phi < \phi_{\min}$ and $x_{eq}^* = 0$).

Part (b) of the proposition concerns the case when $var_p(m)$ is sufficiently larger than $var(m)$ so that (25) no longer holds. Intuitively, this means the insurance benefits of polarisation are limited but the political parties (driven by a significant uncertainty effect) have strong incentive to polarise. In this case, whether or not polarisation remains welfare-improving hinges on ϕ . If the two parties' ideological differences are “moderate”, i.e., fall into the range in (26), then polarisation continues to be welfare-improving.²⁰ But if ϕ is “large” so that

$$\phi > \frac{\sqrt{\pi} (8\sigma^\dagger \cdot \tilde{\sigma} + \gamma)}{2\sqrt{2} [\pi\tilde{\sigma} - 4\sigma^\dagger]},$$

then the resulting policy polarisation will fall outside of the range in (24), i.e.,

$$x_{eq}^* > 2\sqrt{\frac{2}{\pi}}\sigma^\dagger.$$

In this case, polarisation is suboptimal and all voters will be strictly better off in an otherwise identical society but with more congruent parties and convergent policy platforms.²¹

4.2 Improvement in Signal Precision

Since any changes in τ_ε will have no impact on the voters' expected utility in the convergent equilibrium, i.e., $W(0; \delta_v)$, it suffices to focus on how these changes will affect $W(x_{eq}^*; \delta_v)$. Our next result shows that, when voters and politicians share a common prior, a more precise signal will improve all voters' welfare in any symmetric polarised equilibrium.

²⁰In the proof of Proposition 4, it is shown that for any $\gamma \geq 0$ and $(\sigma^\dagger, \tilde{\sigma})$ that satisfy $\tilde{\sigma} > 4\sigma^\dagger/\pi > 0$, the range in (26) is nonempty.

²¹Our Proposition 4 is similar in spirit to Proposition 8 in Bernhardt *et al.* (2009, p.578). However, in their model, the parties' perceived uncertainty about the median voter's policy preference is an exogenous parameter and there is no disagreement between voters' and parties' beliefs. Our result shows that the extent and the direction of disagreement can reverse the conclusion that polarisation is welfare-improving.

Proposition 5. *Suppose there is no disagreement between voters' and politicians' prior beliefs, i.e., $(\lambda_s, \lambda_b) = (\widehat{\lambda}_s, \widehat{\lambda}_b)$, so that $\text{var}(m) = \text{var}_p(m)$. Then any improvement in signal precision will unanimously improve voters' welfare in any symmetric polarised equilibrium, i.e.,*

$$\frac{dW(x_{eq}^*; \delta_v)}{d\tau_\varepsilon} > 0, \quad \text{for any } x_{eq}^* > 0 \text{ and for all } \delta_v.$$

The derivative in Proposition 5 can be expressed as

$$\frac{dW(x_{eq}^*; \delta_v)}{d\tau_\varepsilon} = 2 \underbrace{\left(\sqrt{\frac{2}{\pi}} \sigma^\dagger - x_{eq}^* \right)}_A \frac{d\tilde{\sigma}}{d\tau_\varepsilon} \cdot \underbrace{\frac{dx_{eq}^*}{d\tilde{\sigma}}}_{(+)} + 2 \underbrace{\sqrt{\frac{2}{\pi}} x_{eq}^*}_{(+)} \cdot \underbrace{\frac{d\sigma^\dagger}{d\tau_\varepsilon}}_{(+)} \quad (27)$$

Equation (27) shows that improvement in signal precision affects voters' welfare through two channels. The first one is through affecting the parties' perceived uncertainty about the election outcome, which in turn affect their platform choices.²² This channel is represented by the first term on the right side of (27). The second channel is by affecting the voters' perceived uncertainty and is captured by the second term on the right.

When voters and parties share a common prior (i.e., $\tilde{\sigma} = \sigma^\dagger$), an increase in τ_ε will increase both as shown in Proposition 2. A larger σ^\dagger implies more volatile expected outcomes, which enhances the insurance value of policy divergence and voters' welfare. This positive effect on welfare is captured by the second term in (27). Graphically, this is represented by the same outward shift as shown in Figure 2. At the same time, higher $\tilde{\sigma}$ induces greater polarisation (i.e., x_{eq}^* increases). Depending on whether x_{eq}^* lies below or above the midpoint x_{mid} , the expression A (and hence the entire first term) in (27) may be positive or negative. In the proof of Proposition 5, we show that in the case when $A < 0$, the positive insurance effect will dominate, ensuring that welfare rises with signal precision when there is no disagreement.

This result, however, may not hold when disagreement is present. Our next result describes two scenarios in which improvement in signal precision can be welfare-reducing. In the first scenario, large ideological differences between the parties coexist with significant polarisation-enhancing disagreement between voters and politicians (i.e., when $\widehat{\lambda}_s + \widehat{\lambda}_b$ is sufficiently greater than $\lambda_s + \lambda_b$). The intuition is shown in Figure 3. Suppose x_{eq}^* is large initially due to a

²²It can be seen from (11) that any changes in τ_ε will affect x_{eq}^* through $\tilde{\sigma}$ [hidden in the term $h(0)$] alone. This justifies the decomposition

$$\frac{dx_{eq}^*}{d\tau_\varepsilon} = \frac{d\tilde{\sigma}}{d\tau_\varepsilon} \cdot \frac{dx_{eq}^*}{d\tilde{\sigma}}.$$

high value of ϕ so that it is located on the downward-sloping side of the hump-shaped graph (point P). As mentioned above, an increase in τ_ε raises the voters' uncertainty about the election outcome (i.e., σ^\dagger), which in turn increases the insurance benefits and welfare gain from policy divergence. This is again represented by a strictly positive second term in (27). But now the first term is unambiguously strictly negative since any increase in τ_ε will promote polarisation when $\widehat{\lambda}_s + \widehat{\lambda}_b > \lambda_s + \lambda_b$ [which implies $var_p(m) > var(m)$], and x_{eq}^* is greater than the midpoint x_{mid} initially. As a result, any further increase in polarisation will reduce welfare. Figure 3 depicts a situation in which this negative effect is dominating so that voters' welfare decreases. A rigorous characterisation of this scenario is provided in part (a) of Proposition 6.




Figure 3

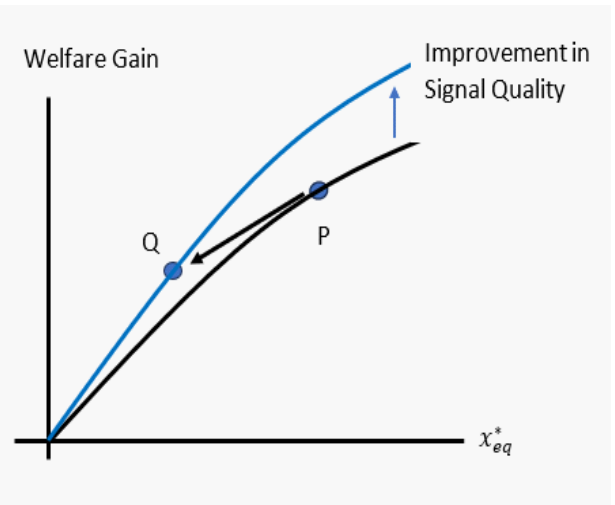


Figure 4

Better signal precision can also be welfare-reducing when disagreement is polarisation-reducing (i.e., when $\lambda_s + \lambda_b$ is sufficiently greater than $\widehat{\lambda}_s + \widehat{\lambda}_b$). This scenario is described in part (b) of Proposition 6 which focuses on the case when $\lambda_s > 0$ and $\lambda_b \geq 0$.²³ As before, an increase in τ_ε will increase σ^\dagger and the insurance benefits of divergent policies. But these benefits are subject to diminishing returns. Specifically, when τ_ε is already large, any further increase will have a negligible impact on σ^\dagger . As a result, the second term in (27) will diminish to zero as τ_ε continues to increase. But on the other hand, Proposition 2 states that a higher value of τ_ε will suppress polarisation when $var(m) > 2var_p(m)$ [which happens when $\lambda_s + \lambda_b \gg \widehat{\lambda}_s + \widehat{\lambda}_b$]. In addition, when λ_s is large, x_{eq}^* will be lower than the mid-point x_{mid} (i.e., on the upward-sloping side of the graph). Therefore, any reduction in x_{eq}^* will lower voters' welfare. This scenario is depicted in Figure 4.

²³If $\lambda_b = \widehat{\lambda}_b = 0$, so that $\lambda_s = \sigma_s^2$ and $\widehat{\lambda}_s = \widehat{\sigma}_s^2$, then the threshold value λ_c in part (b) of Proposition 6 is zero. The set of sufficient conditions in part (b) will then be reduced to $\widehat{\sigma}_s > \sigma_{\min}$ and $\sigma_s > (\pi/2)\widehat{\sigma}_s$.

Proposition 6.

- (a) Suppose $\widehat{\lambda}_s + \widehat{\lambda}_b > 1.32(\lambda_s + \lambda_b)$. Then there exists a threshold value of signal precision, denoted by $\tau_c > 0$, such that

$$\lim_{\phi \rightarrow \infty} \left[\frac{dW(x_{eq}^*; \delta_v)}{d\tau_\varepsilon} \right] < 0 \quad \text{for all } \tau_\varepsilon > \tau_c \text{ and for all } \delta_v \in \mathbb{R}.$$

- (b) Suppose $\lambda_s > 0$, $\lambda_b \geq 0$, $\widehat{\lambda}_s + \widehat{\lambda}_b > \sigma_{\min}^2$ and $\lambda_s + \lambda_b > (\pi/2)^2 (\widehat{\lambda}_s + \widehat{\lambda}_b)$. Then there exists a threshold value of λ_s , denoted by $\lambda_c > 0$, such that

$$\lim_{\tau_\varepsilon \rightarrow \infty} \left[\frac{dW(x_{eq}^*; \delta_v)}{d\tau_\varepsilon} \right] < 0 \quad \text{for all } \lambda_s > \lambda_c \text{ and for all } \delta_v \in \mathbb{R}.$$

We now present two groups of numerical examples that can help illustrate the results in Proposition 6.²⁴ The first one is for the scenario described in part (a). Set $\gamma = 1$, $\lambda_s = \lambda_b = 0.2$ and $\widehat{\lambda}_s = \widehat{\lambda}_b = 0.5$. As shown in the proof of part(a), the larger the differences between $(\widehat{\lambda}_s + \widehat{\lambda}_b)$ and $(\lambda_s + \lambda_b)$, the lower the threshold value τ_c will be. Our choice of $\{\lambda_s, \lambda_b, \widehat{\lambda}_s, \widehat{\lambda}_b\}$ allows us to demonstrate the changes in $dW(x_{eq}^*; \delta_v)/d\tau_\varepsilon$ within a compact range of τ_ε . To capture the effect of ϕ , we consider five different values of this parameter, $\{3, 5, 10, 25, 35\}$. In each case, we compute the welfare gains from polarisation, $[W(x_{eq}^*; \delta_v) - W(0; \delta_v)]$, over a range of τ_ε . The results are shown in Figure 5. According to part (b) of Corollary 1, polarisation increases as ϕ increases. Figure 5 shows that this leads to a decrease in the welfare gain from polarisation at every value of τ_ε . In terms of Figure 3, this represents a movement along the downward-sloping side of the same graph as x_{eq}^* increases (i.e., when σ^\dagger is kept constant). In all the cases considered, increasing τ_ε from an initially low value will create a large welfare gain from polarisation. The concavity of these graphs, however, suggest a diminishing marginal gain in welfare from signal precision improvement. In particular, for sufficiently large values of ϕ , any further increase in τ_ε will reduce the welfare gain, i.e.,

$$\frac{dW(x_{eq}^*; \delta_v)}{d\tau_\varepsilon} < 0$$

as predicted by Proposition 6, and eventually turn the gain into a loss, i.e., $W(x_{eq}^*; \delta_v) < W(0; \delta_v)$.²⁵

²⁴The results reported below are robust to a wide range of values of $\{\gamma, \phi, \lambda_s, \lambda_b, \widehat{\lambda}_s, \widehat{\lambda}_b\}$, hence it is easy to construct other examples that can deliver the same message. Several other cases can be found in Section G of the Online Supplementary Materials.

²⁵In Section G of the Online Supplementary Materials, we provide further examples to show that a large




Figure 5




Figure 6

In the second set of numerical examples, we set $\gamma = 1$, $\widehat{\lambda}_s = \widehat{\lambda}_b = \lambda_b = 0.1$ and $\phi = 1$. The last value is used to show that a large ideological difference between the two parties is not required for the result in part (b) of Proposition 6. We then compute $[W(x_{eq}^*; \delta_v) - W(0; \delta_v)]$ over a range polarisation-enhancing disagreement is needed for the results in Figure 5.

of τ_ε under five different values of λ_s , namely $\{0.1, 0.5, 1, 2, 5\}$. The results are shown in Figure 6. Under this parameterisation, policy polarisation does not exist when $\lambda_s = 0.1$, but gradually increases at each τ_ε when λ_s increases. Since x_{eq}^* is on the upward-sloping side of the graph, this induces an increase in the welfare gain from polarisation at each τ_ε , as shown in Figure 6. When $\lambda_s = 5$, the welfare gain from polarisation displays a clear hump-shaped pattern, i.e., it first increases and then decreases as the signal becomes increasingly precise. This is consistent with the prediction of part (b) of Proposition 6.

5 Conclusion

This paper develops a theoretical framework linking information quality, disagreement between voters and politicians, and political polarisation. By endogenising parties' perceived uncertainty about electoral outcomes, the model highlights how belief formation and information structure shape both policy divergence and welfare. The main novelty of this study is to allow for belief heterogeneity between voters and politicians, which gives rise to two mechanisms: a learning effect, reflecting voters' responsiveness to new signals; and an uncertainty effect, reflecting politicians' difficulty in anticipating voter beliefs. These forces jointly determine how changes in information precision affect equilibrium polarisation. The framework thus explains why improvements in media transparency or access to information can either intensify or mitigate polarization, depending on the extent and direction of disagreement. Therefore, similar societies can exhibit very different levels of polarisation depending on how information is interpreted across groups.

In terms of voters' welfare, polarisation can be beneficial when it provides insurance against uncertainty but becomes inefficient when perceived uncertainty among politicians is excessive. In addition, when disagreement is large, better information can reduce welfare. The value of information therefore depends not only on its accuracy but also on how beliefs and trust in information are distributed across political actors.

Future research could extend the framework to heterogeneous voter groups with different information sources or to dynamic environments in which parties and voters learn across elections. Such extensions would help explain how polarisation persists or evolves over time in response to changing informational and political conditions.²⁶

²⁶ Another potential extension is to incorporate disagreement among voters or politicians. In a preliminary analysis, we allow each political party to hold optimistic beliefs that the state of the world aligns with its own ideological position. The results indicate an additional learning effect: as voter learning intensifies, parties become more optimistic and further polarised, anticipating that voters will also learn that the state favours their respective positions. We are grateful to an anonymous referee for suggesting this extension.

References

- [1] Aytimur, R. E., A. Boukouras and R. M. H. Suen (2025) “How does Political Uncertainty affect the Optimal Degree of Policy Divergence?” *Social Choice and Welfare*, doi.org/10.1007/s00355-025-01595-w.
- [2] Barber, M. J., and N. McCarty (2015) “Causes and Consequences of Polarization,” In Persily, Nathaniel, ed., *Solutions to Political Polarization in America*. Cambridge University Press, pp. 19-53.
- [3] Bernhardt, D., J. Duggan, and F. Squintani (2007) “Electoral Competition with Privately Informed Candidates,” *Games and Economic Behavior*, vol. 58 (1), pp. 1-29.
- [4] Bernhardt, D., J. Duggan and F. Squintani (2009) “The Case for Responsible Parties,” *American Political Science Review*, vol. 103 (4), pp. 570- 587.
- [5] Bernhardt, M. D., and D. E. Ingberman (1985) “Candidate Reputations and the Incumbency Effect,” *Journal of Public Economics*, vol. 27 (1), pp. 47-67.
- [6] Besley, T., and S. Coate (1997) “An Economic Model of Representative Democracy,” *Quarterly Journal of Economics*, vol. 112 (1), pp. 85-114.
- [7] Broockman, D. E., and C. Skovron (2018) “Bias in Perceptions of Public Opinion among Political Elites,” *American Political Science Review*, vol. 112 (3), pp. 542-563.
- [8] Butler, D. M., and A. M. Dynes (2016) “How Politicians Discount the Opinions of Constituents with Whom They Disagree,” *American Journal of Political Science*, vol. 60 (4), pp. 975-989.
- [9] Callander, S. (2005) “Electoral Competition in Heterogeneous Districts,” *Journal of Political Economy*, vol. 113 (5), pp. 1116-1145.
- [10] Callander, S. (2008) “Political Motivations,” *Review of Economic Studies*, vol. 75 (3), pp. 671-697.
- [11] Calvert, R.L. (1985) “Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence,” *American Journal of Political Science* , vol. 29 (1), pp. 69-95.

- [12] Castanheira, M. (2003) “Why Vote for Losers?” *Journal of the European Economic Association*, vol. 1 (5), pp. 1207-1238.
- [13] Chen, E., H. Chang, A. Rao, K. Lerman, G. Cowan and E. Ferrara (2021) “COVID-19 Misinformation and the 2020 US Presidential Election,” *The Harvard Kennedy School Misinformation Review*, vol. 1 (7).
- [14] Giger, N., and H. Klüver (2016) “Voting Against Your Constituents? How Lobbying Affects Representation,” *American Journal of Political Science*, vol. 60 (1), pp. 190-205.
- [15] Grinberg, N., K. Joseph, L. Friedland, B. Swire-Thompson and D. Lazer (2019) “Fake News on Twitter during the 2016 U.S. Presidential Election,” *Science*, vol. 363, pp. 374-378.
- [16] Groseclose, T. (2001) “A Model of Candidate Location When One Candidate Has a Valence Advantage,” *American Journal of Political Science*, vol. 45 (4), pp. 862-886.
- [17] Gul, F., and W. Pesendorfer (2012) “Media and Policy,” unpublished manuscript.
- [18] Hertel-Fernandez, A., M. Mildenerger and L. C. Stokes (2019) “Legislative Staff and Representation in Congress,” *American Political Science Review*, vol. 113 (1), pp. 1-18.
- [19] Hu, L., A. Li, and I. Segal (2023) “The Politics of Personalized News Aggregation,” *Journal of Political Economy Microeconomics*, vol. 1 (3).
- [20] Kärnä, A., and P. Öhberg (2023) “Misrepresentation and Migration,” *Kyklos*, p.1-23.
- [21] Krasa, S., and M. K. Polborn (2012) “Political Competition Between Differentiated Candidates,” *Games and Economic Behavior*, vol. 76 (1), pp. 249-271.
- [22] Levy, G., and R. Razin (2015) “Does Polarisation of Opinions Lead to Polarisation of Platforms? The Case of Correlation Neglect,” *Quarterly Journal of Political Science*, vol. 10, pp. 321-355.
- [23] Little, A. T., and T. B. Pepinsky (2021) “Learning from Biased Research Designs,” *Journal of Politics*, vol. 83 (2).
- [24] Little, A. T., A. Mack and T. B. Pepinsky (2022) “Biased Learning from Elections,” unpublished manuscript.
- [25] Matejka, F., and G. Tabellini (2021) “Electoral Competition with Rationally Inattentive Voters,” *Journal of the European Economic Association*, vol. 19 (3), pp. 1899-1935.

- [26] McCarty, N., K. T. Poole, and H. Rosenthal (2006) *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge, Mass.: MIT Press.
- [27] Munger, K., P. J. Egan, J. Nagler, J. Ronen and J. Tucker (2022) “Political Knowledge and Misinformation in the Era of Social Media: Evidence From the 2015 UK Election,” *British Journal of Political Science*, vol. 52, pp. 107-127.
- [28] Ossokina, I. V., and O. H. Swank (2004) “The Optimal Degree of Polarisation,” *European Journal of Political Economy*, vol. 20(1), pp. 255-262.
- [29] Palfrey, T. R. (1984) “Spatial Equilibrium with Entry,” *Review of Economic Studies*, vol. 51 (1), pp. 139-156.
- [30] Perego, J., and S. Yuksel (2022) “Media Competition and Social Disagreement,” *Econometrica*, vol. 90 (1), pp. 223-265.
- [31] Pereira, M. (2021) “Understanding and Reducing Biases in Elite Beliefs About the Electorate,” *American Political Science Review*, vol. 115 (4), pp. 1308-1324.
- [32] Persson, T. and G. Tabellini (2002) *Political Economics: Explaining Economic Policy*, MIT Press.
- [33] Puglisi, R., and J. M. Snyder Jr. (2015) “Empirical Studies of Media Bias,” in *Handbook of Media Economics*, vol. 1, edited by S. P. Anderson, J. Waldfogel and D. Strömberg, pp. 647-667. Elviesier.
- [34] Roemer, J. E. (1994) “A Theory of Policy Differentiation in Single Issue Electoral Politics,” *Social Choice and Welfare*, vol. 11 (4), pp. 355-380.
- [35] van Weelden, R. (2013) “Candidates, Credibility, and Re-election Incentives,” *Review of Economic Studies*, vol. 80 (4), pp. 1622-1651.
- [36] Vaeth, M. (2025) “Rational Voter Learning, Issue Alignment, and Polarization,” *SSRN Working Paper*, available at: <https://ssrn.com/abstract=4640231>.
- [37] Wittman, D. (1983) “Candidate Motivation: A Synthesis of Alternative Theories,” *American Political Science Review*, vol. 77 (1), pp. 142-157.
- [38] Xefteris, D. and G. Zudenkova (2018) “Electoral Competition under Costly Policy Implementation,” *Social Choice and Welfare*, vol. 50, pp. 721-739.

- [39] Yuksel, S. (2022) “Specialized Learning and Political Polarization,” *International Economic Review*, vol. 63(1), pp. 457-474.