

False Narratives and Political Mobilization*

Kfir Eliaz, Simone Galperti, and Ran Spiegler[†]

June 25, 2022

Abstract

We present an equilibrium model of politics in which political platforms compete over public opinion. A platform consists of a policy, a coalition of social groups with diverse intrinsic attitudes to policies, and a narrative. We conceptualize narratives as subjective models that attribute a commonly valued outcome to (potentially spurious) postulated causes. When quantified against empirical observations, these models generate a shared belief among coalition members over the outcome as a function of its postulated causes. The intensity of this belief and the members' intrinsic attitudes to the policy determine the strength of the coalition's mobilization. Only platforms that generate maximal mobilization prevail in equilibrium. Our equilibrium characterization demonstrates how false narratives can be detrimental for the common good, and how political fragmentation leads to their proliferation. The false narratives that emerge in equilibrium attribute good outcomes to the exclusion of social groups from ruling coalitions.

*Spiegler acknowledges financial support from ERC Advanced Investigator grant no. 692995. We thank Tuval Danenberg, Danil Dmitriev, Federica Izzo, Gilat Levy and Nathan Hancart for helpful comments.

[†]Eliaz: Tel Aviv University and the University of Utah. Galperti: UC San Diego. Spiegler: Tel Aviv University and University College London.

1 Introduction

Success in democratic politics requires the mobilization of public opinion, which takes various forms: rallies, petitions, social media activism, and ultimately voter turnout. Fluctuations in public opinion can determine which policies get implemented and which coalitions of social groups form around them (Burstein (2003)). In turn, opinion makers (politicians, news outlets, pundits) use past performance of policies and coalitions as raw material for shaping public opinions. This paper is an attempt to shed light on this interplay.

Our starting point is the idea that *narratives* are a powerful tool for mobilizing public opinion. This is a familiar idea with numerous expressions in academic and popular discourse. After Senator John Kerry lost the 2004 presidential elections, his political strategist Stanley Greenberg said that “a narrative is the key to everything” and that Republicans had “a narrative that motivated their voters.”¹ Shanahan et al. (2011) write: “Policy narratives are the lifeblood of politics. These strategically constructed ‘stories’ contain predictable elements and strategies whose aim is to influence public opinion toward support for a particular policy preference.”

We build on the approach of Eliaz and Spiegler (2020), who proposed to view political narratives as *causal models* that attribute public outcomes (such as national security, economic growth, or the climate) to other factors (such as policies or external variables). A narrative generates a probabilistic belief by being “estimated” against the observed correlation between the outcome and its postulated causes. The stronger this correlation, the stronger the belief that the narrative generates, and hence the extent to which the narrative mobilizes agents behind a political platform. This means that competition among social groups for political power is to some extent a battle between conflicting narratives over what drives public outcomes. Under this

¹See William Safire’s New York Times article titled “Narrative” (<https://www.nytimes.com/2004/12/05/magazine/narrative.html>).

modeling approach, a *false narrative* corresponds to a *misspecified* causal model. It can produce wrong beliefs because it imposes an incorrect causal meaning on the correlation it highlights.

While Eliaz and Spiegler (2020) assumed a representative agent, this paper considers a heterogeneous society consisting of multiple social groups. We think of a social group as a collection of agents having shared ideological, socio-economic or ethnic/religious characteristics, as well as a distinct political representation. For example, society can be divided into a left and a right—each having its own representative political party. Another example is the distinction between Flemish and French parties in Belgium, or the various religious parties in Israel. The departure from a representative-agent environment means that narratives can now also attribute outcomes to which social groups are in power. This generates new insights into the role of various kinds of false narratives in determining the structure of ruling coalitions—namely, which groups are represented in government—and ultimately the implemented policies.

Our model makes the stark, simplifying assumption that policies are the *only true cause* of public outcomes. However, ideological differences between social groups will naturally give rise to correlations between the structure of ruling coalitions, the policies they implement, and these policies' resulting outcomes. A false narrative can exploit this correlation and causally attribute the outcome *solely* to social groups' power status (i.e., whether they belong to the ruling coalition), even though in reality this correlation is due to confounding by the implemented policies.

For illustration, suppose that a certain coalition C systematically refrains from taxing wealth. As a result, social inequality will tend to rise when C is in power. A rival coalition C' may exploit this correlation and spin a false narrative that in order to reduce inequality, we only need to keep the social groups in C out of power. Because this narrative does not attribute the outcome to its true cause (namely, tax policy), it enables C' to gain support:

On the one hand, C' acts exactly like C by not proposing an unpopular wealth tax; on the other hand, it can claim that by elbowing out C it is doing something to lower inequality, which *is* popular. Thus, in a sense, C' uses C as a “scapegoat.” Our main objective in this paper is to understand how such false narratives can gain ascendance, what form they may take, and how they shape public policies and ruling coalitions.

In our setting, a *political platform* consists of a policy, a coalition of social groups, and a narrative. Given a long-run joint empirical distribution over prevailing platforms and public outcomes, different narratives may induce conflicting beliefs regarding the consequences of policies and coalitions. Moreover, changes in long-run frequencies may change the beliefs induced by a narrative, and therefore the extent to which it can mobilize a social group. We define an *equilibrium* as a probability distribution over platforms, such that every platform in its support maximizes the total mobilization of the social groups belonging to the platform’s coalition.

This definition captures in reduced form the idea that a platform’s success depends on the strength of its popular support (in terms of the number and size of social groups that rally behind it, and the intensity with which they do so). It does so in the spirit of competitive equilibrium, as in Rothschild and Stiglitz (1976). The backstory is that there is “free entry” of office-motivated political entrepreneurs who propose policy-narrative combinations. If a particular combination attracts stronger support than the current combination, it will overthrow it, such that the platform that eventually prevails is the one that maximizes total support. We present a convergence result that provides a dynamic foundation for the equilibrium concept.

In line with the competitive-equilibrium-like approach, we do not offer an explicit model of the political process. In particular, there will be no explicit game between platforms. This is in keeping with our interest in competition over public opinion in a broad sense, rather than in specific electoral procedures or post-election coalitional bargaining. However, in the

case of two social groups, our equilibrium concept is consistent with a two-party voting model with endogenous voter turnout.

Using this formalism, we obtain several qualitative insights. First, in addition to the true narrative that attributes outcomes to policies, two types of false narratives emerge in equilibrium. The first type is a “denialist” narrative that attributes outcomes to external forces outside society’s control. The other type is a “tribal” narrative that attributes a good public outcome to the *exclusion* of some social groups from the ruling coalition. In a political speech or a social-media post, such a narrative could appear as “national security was in good shape when the Left was out of power.”

Recent public debates over rising inflation are suggestive of these kinds of narratives: Much of the action in these debates involves competing claims over the causes of high inflation. Some narratives attributes it to government actions (fiscal expansion), others to external factors (global supply-chain disruptions), and yet others assign credit or blame for the level of inflation solely to the party in power, without attempting to link this to the party’s policy decisions. A selection of press quotes demonstrates the form that these conflicting narratives take:

“As prices have increased...some Democrats have landed on a new culprit: price gouging...For Democrats, it is a convenient explanation as inflation turns voters against President Biden. It lets Democrats deflect blame from their pandemic relief bill, the American Rescue Plan, which experts say helped increase prices.”²

“Democrats have blamed supply chain deficiencies due to COVID-19, as well as large corporations and monopolies.”³

²<https://www.nytimes.com/2022/06/14/briefing/inflation-supply-chain-greedflation.html>

³<https://fivethirtyeight.com/features/what-democrats-and-republicans-get-wrong-about-inflation/>

“As the midterm elections draw nearer, a central conservative narrative is coming into sharp focus: President Biden and the Democratic-controlled Congress have made a mess of the American economy.”⁴

The distinction between a false narrative that attributes outcomes to whoever is in power and a more accurate narrative that attributes outcomes to policies is also alluded to by Paul Krugman in a recent article about the politics of inflation:

“... voters aren’t saying, “Trimmed mean P.C.E. inflation is too high because fiscal policy was too expansionary.” They’re saying, “Gas and food were cheap, and now they’re expensive ...” So when people say — as they do — that gas and food were cheaper when Donald Trump was president, what do they imagine he could or would be doing to keep them low if he were still in office?”⁵

We wish to emphasize that we do not argue that our specific model matches the inflation scenario. Nevertheless, we believe it offers an insight into the interplay between the popularity of various types of false narratives, and the objective statistical reality that both feeds the narratives and gets shaped by them through the policy choices that different narratives promote.

Our second qualitative insight is that the false narratives that are employed in equilibrium sustain a policy that would not be taken if the only prevailing narrative were a true one (which correctly attributes outcomes to policies). The function of false narratives is to resolve the cognitive dissonance between the objective ineffectiveness of a policy for the common good

⁴<https://www.nytimes.com/2022/06/11/opinion/fed-federal-reserve-inflation-democrats.html>

⁵<https://www.nytimes.com/2022/06/02/opinion/inflation-biden.html>. See also Weaver (2013) and Sanders et al. (2017).

and that same policy’s intrinsic appeal for a social group (e.g., based on costs or benefits which affect that group). This is achieved by deflecting responsibility for the public outcome from its true cause to spurious causes. Moreover, when society’s structure of political representation is more fragmented and “tribal” (in a sense we make precise in our model), false narratives proliferate and lead to further crowding out of the true narrative and the policy it justifies.

Finally, we characterize the structure of coalitions that form in equilibrium. False narratives can give rise to coalitions that would not form if only the true narrative prevailed. In particular, when a political platform employs a “tribal” narrative, it may “scapegoat” social groups that support that platform’s policy (indeed, they implement the same policy when they are in power), yet their exclusion from the platform’s coalition is necessary for the narrative’s effectiveness. We also perform comparative statics with respect to the polarization of attitudes toward policies in society. We show that greater polarization can be detrimental for the common good through the proliferation of false narratives.

2 A Model

We begin by introducing the building blocks of our model.

Policies and outcomes

Our model examines public-opinion battles over a single salient issue, defined by a *public outcome* variable $y \in \{0, 1\}$. Assume there is a social consensus that $y = 1$ is a desirable outcome, and refer to it accordingly as the “good” outcome. For example, the issue can be economic growth such that $y = 1$ represents high growth. Let $a \in A = \{\ell, h\}$ be a *policy*. Policies cause outcomes according to the following objective conditional probability

distribution:

$$\Pr(y = 1 \mid a) = \begin{cases} q & \text{if } a = h \\ 0 & \text{if } a = \ell \end{cases} \quad (1)$$

where $q \in (0, 1]$.

Social groups and coalitions

Let $N = \{1, \dots, n\}$ be a set of *social groups*. A *coalition* is a non-empty, strict subset of groups $C \subset N$. Note that this rules out only the possibility of a “grand coalition” $C = N$ including *all* groups (which can be achieved via other more primitive assumptions as we do in Section 4).

The following notation will be useful. Let $x = (x_0, \dots, x_n)$ be a profile of binary variables, where $x_0 \in \{\ell, h\}$ and $x_i \in \{0, 1\}$ for every $i > 0$. For every $S \subset \{0, \dots, n\}$, denote $x_S = (x_i)_{i \in S}$. Define the following function that assigns values of x to policy-coalition pairs: for every (a, c) , $x_0(a, C) = a$; and for every $i > 0$, $x_i(a, C) = 1$ if and only if $i \in C$. For $i > 0$, we will refer to x_i as group i ’s “*power status*.”

Narratives

We define a narrative as a subset $S \subseteq \{0, \dots, n\}$, which is the set of indices of the components of x . The set S defines the components to which the narrative attributes the public outcome y . For example, $S = \{0, 2\}$ means that the narrative attributes the outcome to the policy and the power status of social group 2. Given a long-run joint probability distribution p over (x, y) —which we endogenize below—a narrative S generates a probabilistic belief over the outcome conditional on its postulated causes, given by $(p(y \mid x_S))$.⁶ In other words, the narrative draws attention to the correlation between y and x_S and imposes a causal meaning on this correlation.

Two narratives have a special status in our model. First, $S = \{0\}$ is the “*true narrative*” because it attributes y to its sole true cause a . Every other narrative is false in the sense that it attributes y to wrong causes.

⁶We use the abbreviated notation $(p(y \mid x_S)) = (p(y \mid x_S))_{x_S, y}$.

Second, $S = \emptyset$ is a “*denialist*” narrative that does not attribute y to any of the endogenous variables. In particular, it claims that the policy has no consequences. We interpret the denialist narrative as the attribution of outcomes to external variables.

A third type of narrative will also play a key role in our model. These are narratives $S \subseteq N$, where $S \neq \emptyset$. We refer to such narratives as “*tribal*” because they attribute outcomes to the power status of social groups, without mentioning policies.

Platforms

We refer to policy-coalition-narrative triples (a, C, S) as *platforms*. These are the objects that vie for popularity in our model. Given an objective long-run probability distribution σ over platforms (a, C, S) , we can define an induced joint distribution over (a, C, S, y) that incorporates expression (1):

$$p_\sigma(a, C, S, y) = \sigma(a, C, S) \cdot \Pr(y \mid a).$$

Hereafter, we denote the support of σ by $Supp(\sigma)$.

Mobilization

The extent to which a platform can mobilize a social group depends on two factors: the intensity of the group’s belief that the platform induces a good public outcome and the group’s intrinsic attitude to the platform’s policy. Specifically, let $f_i : A \rightarrow \mathbb{R}_+$ be a function that measures group i ’s intrinsic attitude towards each policy. Denote $f = (f_i)_{i \in N}$. We refer to $f_i(a)$ as group i ’s “*mobilization potential*” for a . Intuitively, $f_i(a)$ reflects the specific consequences (such as costs) that policy a has for a group i and so its raw support for a , besides its general interest in the public outcome.

Denote $N^a = \{i \in N \mid f_i(a) > 0\}$ —i.e., this is the set of groups having positive mobilization potential for a . To avoid trivial cases, we assume that for every group $i \in N$, $f_i(a) > 0$ for least one policy a , and that for every a , $f_i(a) > 0$ for some $i \in N$.

Definition 1 (Platform Support) Fix a probability distribution σ over platforms. Then, group i 's support for a platform (a, C, S) is defined by

$$u_{i,\sigma}(a, C, S) = p_\sigma(y = 1 \mid x_S(a, C)) \cdot f_i(a) \quad (2)$$

Our notion of a social group's support for a platform takes a broad view of political mobilization to include not only voting, but also other kinds of political participation: rallies, petitions, protests, or social media activism. We assume that the actual mobilization of group i is *proportional* to its mobilization potential for the platform's policy, as well as to the belief—shaped by the narrative—that the platform induces a good outcome. The stronger a group's belief, the stronger its support for the platform. In Section 6.1, we provide a “microfoundation” that substantiates the multiplicative form of $u_{i,\sigma}$.

The reason why $u_{i,\sigma}$ is indexed by σ is that the belief $p_\sigma(y = 1 \mid x_S)$ may be sensitive to changes in σ . To see why, recall that y is purely a (probabilistic) function of a , so y is independent of C conditional on a . This property can be represented by the directed acyclic graph (DAG) $C \leftarrow a \rightarrow y$ (the direction of the link between C and a is arbitrary; what matters is that they are correlated because they are jointly determined, as we will describe below). If the narrative S does not attribute y to a —i.e., $S \subseteq N$ —it amounts to interpreting the long-run correlation between C and y as if it is causal, namely as if the DAG were $C \rightarrow y$. In reality, the correlation is due to confounding, since a is correlated with both y and C . The latter correlation is determined by σ as the following expression makes evident:

$$p_\sigma(y = 1 \mid x_S) = \sum_{a,C} p_\sigma(a, C \mid x_S) p(y = 1 \mid a), \quad (3)$$

where the term $p_\sigma(a, C \mid x_S)$ is determined by σ .

To illustrate (3) and its role in how a false narrative can induce a wrong

belief, suppose $S = \{i\}$ and $i \notin C$ for some $i \in N$. Then,

$$p_\sigma(y = 1 \mid x_S(\ell, C)) = \frac{q \sum_{C, S \mid i \notin C} \sigma(h, C, S)}{\sum_{a, C, S \mid i \notin C} \sigma(a, C, S)}$$

If $\sigma(h, C, S) > 0$ for some platforms in which $i \notin C$, then $p_\sigma(y = 1 \mid x_S(\ell, C)) > 0$. This means that platform (ℓ, C, S) will receive positive support (as long as C includes some j such that $f_j(\ell) > 0$), even though it objectively leads to $y = 0$ with certainty.

Admissible coalitions

We impose a restriction on feasible coalitions, which is based on a qualitative distinction between $f_i(a) > 0$ and $f_i(a) = 0$. Suppose group i is in fact intrinsically *opposed* to policy a . Then, a coalition that includes group i and advocates a benefits from ousting that group because it would act as a “fifth column.” But of course group i would never join this coalition in the first place. Therefore, we interpret $f_i(a) = 0$ as an assumption that group i is opposed to a and, hence, will never be in a coalition that advocates it. By assumption, the group satisfies $f_i(a') > 0$ for $a' \neq a$, which means that it could join coalitions that advocate a' . In this sense, rallying in favor of a' is like rallying against a .

Formally, given policy a , a coalition C is *admissible* if $f_i(a) > 0$ for all $i \in C$. Hereafter, we consider only platforms involving admissible coalitions.

Equilibrium

The feedback between σ and political mobilization impels us to introduce an *equilibrium* notion of dominant platforms.

Definition 2 (Equilibrium) *A probability distribution σ with full support over platforms with admissible coalitions is an ε -equilibrium if whenever $\sigma(a, C, S) > \varepsilon$, (a, C, S) maximizes*

$$U_\sigma(a, C, S) = \sum_{i \in C} u_{i, \sigma}(a, C, S) \tag{4}$$

A probability distribution σ (not necessarily with full support) is an equilibrium if it is the limit of ε -equilibria as $\varepsilon \rightarrow 0$.

We start from the notion of ε -equilibrium to ensure that $p_\sigma(y = 1 \mid x_S)$ is always well-defined, given that groups rely on long-run data to form beliefs. This “trembling hand” aspect plays a very limited role in our analysis.

Our equilibrium notion captures the idea that a platform’s political strength depends on how strongly it mobilizes its constituent groups. We view narrative-fueled political competition as a battle over public opinion. A platform is *dominant* given σ if it generates the largest aggregate mobilization of its coalition members. Therefore, the only *dominant* platforms in equilibrium are those that maximize (4), to which we refer as the *platform’s payoff*. When (a, C, S) is dominant, we say that C is a *ruling coalition*.

Note that if only the true narrative $S = \{0\}$ were feasible, then a platform with $a = \ell$ would generate zero payoffs by (1) and the definition of $u_{i,\sigma}$. Instead, a platform with $a = h$ always generates positive payoffs. In this case, $a = h$ would occur with probability one in equilibrium. We therefore refer to $a = h$ as the “*rational*” policy.

3 Two-Group Societies

We begin our analysis with the simple case of $n = 2$, where the only feasible coalitions are $\{1\}$ and $\{2\}$. We assume that the mobilization potential function satisfies $f_1(h) > f_2(h)$ and $f_2(\ell) > f_1(\ell)$. That is, policy h (ℓ) receives stronger raw support from group 1 (2).

This specification of our model is akin to a two-party system, in which exactly one party can be in power at any point in time. In this case, our equilibrium concept can be interpreted in terms of a two-party voting model: Supporters of each party vote non-strategically for it if their net anticipatory utility from their party’s platform is positive—otherwise they abstain (some-

what as in Levy et al. (2022)). We elaborate on this connection in Section 6.1.

The following are a few real-life examples of outcomes and policies that we have in mind. First, the issue is climate change and $a = h$ represents carbon taxation, which produces a common environmental benefit but induces differential costs among social groups (captured by f). Second, the issue is economic growth, where $a = h$ represents policies such as deregulation or other structural reforms that foster growth but inflict differential adjustment costs across society. Finally, the issue is national security, where $a = h$ represents an aggressive military strategy that mitigates security threats but involves sacrifices and moral judgments that vary across groups.

This simple setting allows us to reduce the set of relevant narratives. Since $x_1 = 1$ if and only if $x_2 = 0$, all tribal narratives $S \subseteq N$ are equivalent: They effectively say that “*things are good when group i is in power / group j is not in power.*” In addition, all narratives that weakly contain $\{0\}$ are equivalent, because $\Pr(y = 1 \mid a, C) \equiv \Pr(y = 1 \mid a)$. Thus, every feasible narrative is equivalent to one of the following three narratives: the true narrative $\{0\}$, the denialist narrative \emptyset , and the tribal narrative $\{1\}$ (or, equivalently, $\{2\}$). Therefore, hereafter we assume that only these three narratives—which we denote by *true*, *denial*, and *tribal* for expositional clarity—are feasible.

Proposition 1 *There is a unique equilibrium distribution over (a, C) . The only platforms that can be in $\text{Supp}(\sigma)$ are $(h, \{1\}, \text{true})$, $(\ell, \{2\}, \text{denial})$, and $(\ell, \{1\}, \text{tribal})$. Furthermore:*

(i) $\sigma(h, \{1\}, \text{true}) = \min \{1, f_1(h)/f_2(\ell)\};$

(ii) $\sigma(\ell, \{1\}, \text{tribal}) > 0$ only if $\sigma(\ell, \{2\}, \text{denial}) > 0$.

We present the proof of this result here, as it illustrates the basic logic and forces at the core of our model. All other proofs are relegated to the Appendix.

Proof of Proposition 1 We begin by writing the payoff U_σ for the platforms carried by the three relevant narratives:

$$\begin{aligned} U_\sigma(a, \{i\}, true) &= p_\sigma(y = 1 \mid a) \cdot f_i(a) \\ U_\sigma(a, \{i\}, denial) &= p_\sigma(y = 1) \cdot f_i(a) = q \cdot p_\sigma(a = h) \cdot f_i(a) \\ U_\sigma(a, \{i\}, tribal) &= p_\sigma(y = 1 \mid x_i = 1) \cdot f_i(a) \end{aligned}$$

The proof proceeds in steps.

Step 1 (i) If $\sigma(a, \{i\}, true) > 0$, then $a = h$ and $i = 1$. (ii) If $\sigma(h, \{i\}, S) > 0$, then $S = true$.

Proof. Consider any ε -equilibrium σ . Note that $p_\sigma(y = 1 \mid a = \ell) = 0$ and $p_\sigma(y = 1 \mid a = h) = q$. It follows that if $\sigma(a, \{i\}, true) > \varepsilon$ and hence $(a, \{i\}, true)$ maximizes U_σ , then $a = h$ and $i = 1$ because $f_1(h) > f_2(h)$. Now suppose $\sigma(h, \{i\}, S) > \varepsilon$. Since σ has full-support, $p_\sigma(y = 1 \mid x_{S'}) < q$ whenever $0 \notin S'$. This means that $U_\sigma(h, \{i\}, true) > U_\sigma(h, \{i\}, S')$ for every such S' ; hence, $S = true$. We have thus established that the claimed properties hold for any ε -equilibrium and, hence, in any limit of ε -equilibria.

■

Step 1 implies that if $(a, \{i\}, denial)$ or $(a, \{i\}, tribal)$ are in $Supp(\sigma)$, then $a = \ell$.

Step 2 If $\sigma(\ell, \{i\}, denial) > 0$, then $i = 2$.

Proof. This follows immediately from $f_2(\ell) > f_1(\ell)$. ■

Step 3 If $\sigma(\ell, \{i\}, tribal) > 0$, then $i = 1$.

Proof. Step 1(i) and $Pr(y = 1 \mid a = \ell) = 0$ imply that $p_\sigma(y = 1 \mid x_i = 1) > 0$ only if $i = 1$. Therefore, if $(a, \{i\}, tribal)$ is in $Supp(\sigma)$, then $i = 1$. ■

The previous steps pin down the three platforms that can be in $Supp(\sigma)$ for any equilibrium σ , as well as their payoffs:

$$\begin{aligned}
U_\sigma(h, \{1\}, true) &= q \cdot f_1(h) \\
U_\sigma(\ell, \{2\}, denial) &= q \cdot \sigma(h, \{1\}, true) \cdot f_2(\ell) \\
U_\sigma(\ell, \{1\}, tribal) &= q \cdot \frac{\sigma(h, \{1\}, true)}{\sigma(h, \{1\}, true) + \sigma(\ell, \{1\}, tribal)} \cdot f_1(\ell)
\end{aligned} \tag{5}$$

Step 4 In any equilibrium, $\sigma(h, \{1\}, true) > 0$.

Proof. If $\sigma(h, \{1\}, true) = 0$, the previous expressions become

$$U_\sigma(\ell, \{2\}, denial) = U_\sigma(\ell, \{1\}, tribal) = 0 < U_\sigma(h, \{1\}, true).$$

Therefore, $(h, \{1\}, true)$ is the only possible member of $Supp(\sigma)$, a contradiction. ■

Step 5 In any equilibrium, $\sigma(\ell, \{1\}, tribal) > 0$ only if $\sigma(\ell, \{2\}, denial) > 0$.

Proof. Suppose $\sigma(\ell, \{1\}, tribal) > 0 = \sigma(\ell, \{2\}, denial)$. Then,

$$\sigma(h, \{1\}, true) + \sigma(\ell, \{1\}, tribal) = 1,$$

so that

$$U_\sigma(\ell, \{1\}, tribal) = q \cdot \sigma(h, \{1\}, true) \cdot f_1(\ell)$$

But $f_2(\ell) > f_1(\ell)$ then implies that $U_\sigma(\ell, \{1\}, tribal) < U_\sigma(\ell, \{2\}, denial)$, which contradicts $\sigma(\ell, \{1\}, tribal) > 0$. ■

We are now able to show that an equilibrium exists and it is unique. By Steps 1-3, there are three possible cases for $Supp(\sigma)$: all three platforms or $\{(h, \{1\}, true), (\ell, \{2\}, denial)\}$ or $\{(h, \{1\}, true)\}$. By the definition of equilibrium, each member of $Supp(\sigma)$ has to maximize U_σ .

Case I: $f_1(h) \geq f_2(\ell)$. In this case, $U_\sigma(h, \{1\}, true) > U_\sigma(\ell, \{2\}, denial)$ whenever $\sigma(h, \{1\}, true) < 1$. Step 5 then implies $Supp(\sigma) = \{(h, \{1\}, true)\}$. Indeed, when $\sigma(h, \{1\}, true) = 1$,

$$U_\sigma(h, \{1\}, true) \geq U_\sigma(\ell, \{2\}, denial), U_\sigma(\ell, \{1\}, tribal)$$

Thus, $\sigma(h, \{1\}, true) = 1$ is the unique equilibrium.

Case II: $f_2(\ell) > f_1(h)$. In this case, $U_\sigma(h, \{1\}, true) < U_\sigma(\ell, \{2\}, denial)$ if $\sigma(h, \{1\}, true) = 1$. Therefore, $\sigma(h, \{1\}, true) < 1$. Step 5 then implies that $(\ell, \{2\}, denial) \in Supp(\sigma)$. Now consider two sub-cases. First, let $f_2(\ell) > f_1(h) \geq f_1(\ell)$. Then, $U_\sigma(\ell, \{1\}, tribal) < U_\sigma(h, \{1\}, true)$ whenever $\sigma(\ell, \{1\}, tribal) > 0$. Therefore,

$$\sigma(h, \{1\}, true) = \frac{f_1(h)}{f_2(\ell)} \quad \sigma(\ell, \{2\}, denial) = \frac{f_2(\ell) - f_1(h)}{f_2(\ell)}$$

is the unique solution of

$$U_\sigma(\ell, \{2\}, denial) = U_\sigma(h, \{1\}, true) \geq U_\sigma(\ell, \{1\}, tribal).$$

Second, let $f_2(\ell) > f_1(\ell) > f_1(h)$. Then,

$$\begin{aligned} \sigma(h, \{1\}, true) &= \frac{f_1(h)}{f_2(\ell)} \\ \sigma(\ell, \{2\}, denial) &= \frac{f_2(\ell) - f_1(\ell)}{f_2(\ell)} \\ \sigma(\ell, \{1\}, tribal) &= \frac{f_1(\ell) - f_1(h)}{f_2(\ell)} \end{aligned}$$

is the unique solution of

$$U_\sigma(\ell, \{2\}, denial) = U_\sigma(h, \{1\}, true) = U_\sigma(\ell, \{1\}, tribal)$$

This completes the proof. \square

To interpret the equilibrium, consider the case in which

$$f_2(\ell) > f_1(\ell) > f_1(h) > f_2(h)$$

so that all three platforms described in the proposition are in $Supp(\sigma)$. The distribution σ describes the long-run frequencies with which each of these platforms prevail. When $(h, \{1\}, true)$ prevails, group 1 in power, implements the rational policy h , and employs the true narrative that attributes outcomes to policies. When $(\ell, \{2\}, denial)$ prevails, group 2 is in power, implements policy ℓ , and employs the denial narrative that implicitly attributes outcomes to external factors. Finally, when $(\ell, \{1\}, tribal)$ prevails, group 1 is in power, implements ℓ , and employs the tribal narrative that attributes outcomes to who is in power (without referring to policies).

A dynamic process

For an intuition for this equilibrium, it is useful to have a dynamic process in mind. Imagine that initially there are some random fluctuations over (a, C) and that only the true narrative is considered. Then, policy ℓ garners no support because it ensures $y = 0$. The true narrative can only justify policy h , which gets stronger support from group 1. Thus, the prevailing platform is $(h, \{1\}, true)$, and it generates a payoff of $q \cdot f_1(h)$.

Now suppose that at some point, the platform $(\ell, \{2\}, denial)$ arises. Since approximately only policy h that has ever been taken, the denialist narrative induces the belief $\Pr(y = 1) \approx q$. Because $f_2(\ell) > f_1(h)$, the new platform generates stronger support than the “incumbent” platform $(h, \{1\}, true)$. As a result, the new platform becomes dominant, displacing the old one. Since the new platform involves $a = \ell$, the historical frequency of $a = h$ gradually declines, lowering $\Pr(y = 1)$.

As this process continues, the denialist platform’s payoff will drop below $q \cdot f_1(\ell)$. When this happens, a third narrative can gain traction and give

rise to yet another platform $(\ell, \{1\}, \text{tribal})$. In the path described above, $a = h$ if and only if $x_1 = 1$ (approximately). This implies the historical conditional probability $\Pr(y = 1 \mid x_1 = 1) \approx q$. Consequently, a narrative arguing that things are good when group 1 is in power (or, equivalently, when group 2 is out of power) can mobilize group 1 behind policy ℓ . The payoff of $(\ell, \{1\}, \text{tribal})$ is approximately $q \cdot f_1(\ell)$. Since $f_1(\ell) > f_1(h)$, this payoff exceeds that of the two previous platforms and $(\ell, \{1\}, \text{tribal})$ becomes dominant. While it dominates, it weakens the correlation between x_1 and y and therefore lowers its own payoff. It also weakens the appeal of the denial narrative by lowering the frequency of $y = 1$. This brings the platform carried by the true narrative back in vogue.

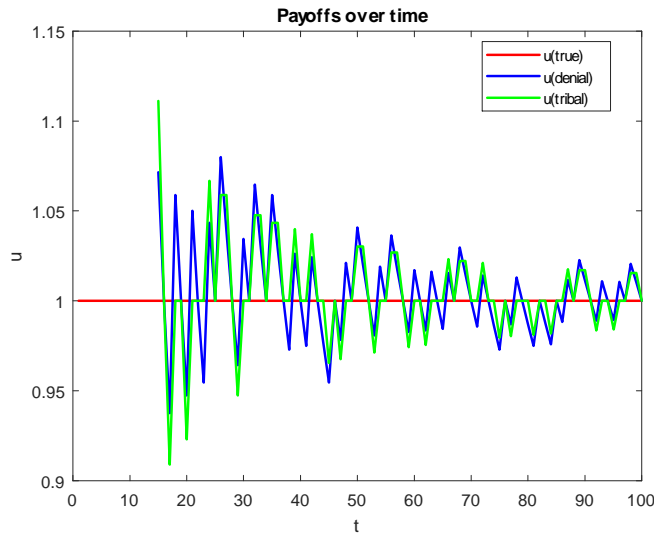


Figure 1

The subsequent dynamic repeats this cycle, albeit with smaller payoff swings. In the long run, the process reaches a configuration in which all three platforms generate the same payoff $q \cdot f_1(h)$. From this point on, any deviation that increases the long-run frequency of one of the three platforms

will trigger an offsetting dynamic response. That is, the equilibrium is dynamically stable. Figure 1 shows the evolution of the three platforms' payoffs, assuming $f_2(\ell) = 3$, $f_1(\ell) = 2$, and $f_1(h) = 1$.

We wish to highlight a few significant features of Proposition 1. The rational policy must be played with positive probability in equilibrium. The reason is that any platform carried by a false narrative free-rides on episodes with $a = h$. Also, note that a platform that advocates $a = h$ will generate its largest support if it employs the true narrative, which highlights the correlation between a and y (since this correlation is stronger than the correlation between y and any other variable).

However, when $f_2(\ell) > f_1(h)$, policy ℓ has stronger raw support than the rational policy h . In this case, false narratives allow ℓ to gain dominance at the expense of h . They enable supporters to “eat their cake and have it:” On the one hand, they are attracted to the policy; on the other hand, the narrative distracts them from the adverse objective consequences of this policy. The equilibrium probability of $a = h$ is determined by the ratio $f_1(h)/f_2(\ell)$. What makes policy ℓ not only popular but also “populist” is that it necessitates a false narrative to mobilize public opinion.

The distinction between the two false narratives—i.e., the denial and tribal narratives—is irrelevant for the equilibrium probability of $a = h$. However, it does play an important role for the identity of the group in power. When $f_1(\ell) > f_1(h)$, the tribal narrative enables group 1 to displace group 2, even though it adopts the same “populist” policy ℓ . The reason is that group 1 can milk its reputation for achieving a good outcome—thanks to its actual historical tendency to implement $a = h$, which is greater than its rival's. It does so by highlighting the historical correlation between $y = 1$ and being in power (or, equivalently, group 2 being out of power). In the next section, we will see that in societies with more than two groups, tribal narratives are relevant not only for the power structure, but also for the equilibrium frequency of the rational policy.

4 Fragmented Societies

In this section we characterize equilibria in the presence of more than two social groups (i.e., $n > 2$). This creates a distinction between varieties of tribal narratives absent in the case of $n = 2$: There will be a difference between “exclusionary” narratives of the form “things are good when these groups are *out* of power” and “inclusionary” narratives of the form “things are good when these groups are *in* power.” We will see how proliferation of exclusionary narratives can have a detrimental effect on the equilibrium probability of the rational policy. It also leads to new coalitional structures that would not arise if only the true narrative prevailed.

Toward this end, we impose some structure on the set of feasible narratives. Let \mathcal{S} be a family of subsets of N . A narrative S is feasible if and only if $S \setminus \{0\} \in \mathcal{S}$. We interpret each element in \mathcal{S} as a collection of social groups that can be *clearly identified* by a common label or defining attribute. For instance, \mathcal{S} can represent a division of society along political leanings or ethnic-religious affiliations, or it can represent salient collections of social groups (such as “unionized workers” or “the economic elite”).

The motivation for restricting the domain of feasible narratives is that, depending on the context, not every collection of social groups can be identified and held accountable for outcomes, because it may lack a clear representation in governing institutions. In some political systems (e.g., Israel) there are political parties that directly represent specific ethno-religious groups. Consequently, there is data about their power status and how it is correlated with outcomes, thus making a narrative that exploits this correlation quantifiable. In other systems (e.g., the US), the mapping between social groups and political representation is more blurred, thus restricting the supply of similar narratives.

Recall that $N^a = \{i \in N \mid f_i(a) > 0\}$. Henceforth, we assume that the sets $N \setminus N^h$ (groups that support only ℓ), $N \setminus N^\ell$ (groups that support only h), and $N^\ell \cap N^h$ (groups that support both policies) are non-empty

and that they are all in \mathcal{S} . We sometimes refer to these broad categories as “right,” “left,” and “center.” These categories are always feasible as tribal narratives. We also assume that any other $S \in \mathcal{S}$ is a *subset* of one of these three categories. Finally, we assume that the *true and denialist narratives are always feasible*.

An illustrative example

Let $n = 4$ and $\mathcal{S} = \{\{1\}, \{2\}, \{3\}, \{4\}, \{3, 4\}\}$. The left is $\{1\}$, the center is $\{2\}$, and the right is $\{3, 4\}$; $\{3\}$ and $\{4\}$ represent sub-divisions of the right (e.g., moderates and extremists). Let $f_3 \equiv f_4$. Assume further that $f_2(\ell) > f_1(h) + f_2(h)$, namely the center’s raw support for ℓ is stronger than the raw support for h among the center-left. The following distribution is an equilibrium (we will shortly see that it is the “essentially unique” equilibrium):

σ	<i>policy</i>	<i>coalition</i>	<i>narrative</i>
$\frac{f_1(h)+f_2(h)}{f_2(\ell)+f_3(\ell)+f_4(\ell)}$	h	$\{1, 2\}$	<i>true</i>
$\frac{f_2(\ell)-f_1(h)-f_2(h)}{f_2(\ell)+f_3(\ell)+f_4(\ell)}$	ℓ	$\{2\}$	$\{3, 4\}$
$\frac{f_3(\ell)+f_4(\ell)}{2[f_2(\ell)+f_3(\ell)+f_4(\ell)]}$	ℓ	$\{2, 3\}$	$\{4\}$
<i>ditto</i>	ℓ	$\{2, 4\}$	$\{3\}$

As in the previous section, policy ℓ is played with positive probability, sustained by false narratives. In contrast to that case, however, here all false narratives take the exclusionary tribal form. For example, in the platform $(\ell, \{2\}, \{3, 4\})$, the center attributes a good outcome to keeping the right out of power. Furthermore, the equilibrium exhibits *endogenous fragmentation*: The center and each faction of the right sometimes form a coalition, using a false narrative that attributes the good outcome to keeping the remaining right-wing group out of power.

Finally, note that as in the previous section, the equilibrium probability of the rational policy h is equal to the ratio between the total mobilization potential for h among the groups that implement h in equilibrium (i.e.,

center-left) and the total mobilization potential for ℓ among the groups that implement ℓ in equilibrium (i.e., center-right). We will see in the next subsection that this is not a general feature. \square

We now establish existence of equilibrium and show that the joint equilibrium distribution over policies and coalitions is unique. The proof provides an algorithm for computing this distribution. Since multiple narratives can induce the same beliefs, we introduce a natural refinement that pins down the equilibrium distribution over platforms.

Definition 3 (Essential Equilibrium) *An equilibrium σ is essential if it satisfies two properties:*

- (i) *If $(a, C, S) \in \text{Supp}(\sigma)$ and $S \neq \{0\}$, then $U_\sigma(a, C, \{0\}) < U_\sigma(a, C, S)$.*
- (ii) *If $(a, C, S) \in \text{Supp}(\sigma)$ and $S \subseteq N$, then $U_\sigma(a, C, S') < U_\sigma(a, C, S)$ for all $S' \subset S$.*

This definition means that if several narratives that accompany (a, C) induce the same payoff, there is a lexicographically secondary criterion that favors *true over false* narratives (condition (i)) and *simple over complex* narratives (condition (ii)). We introduce refinement (i) in order to isolate the cases where false narratives are necessary for the platform's dominance. We introduce refinement (ii) in order to isolate the components of false narratives that are essential for the belief they generate.

Our main result uses the following piece of notation:

$$F(a, M) = \sum_{i \in M} f_i(a), \quad M \subseteq N.$$

That is, $F(a, M)$ is the aggregate mobilization potential for policy a in the collection of groups M . Recall that N^a is the set of groups with positive mobilization potential for a .

Proposition 2 *There exists a unique essential equilibrium σ^* . This equilibrium satisfies the following properties:*

- (i) *The rational policy $a = h$ is realized only as part of the platform $(h, N^h, \{0\})$.*
- (ii) *The probability of $a = h$ is 1 if $F(h, N) \geq F(\ell, N)$, and belongs to $(0, F(h, N)/F(\ell, N)]$ otherwise.*
- (iii) *Every platform $(\ell, C, S) \in \text{Supp}(\sigma^*)$ satisfies $S \subseteq N \setminus N^h$ and $C = N^\ell \setminus S$.*

Remark 1 *While Proposition 2 focuses on essential equilibrium, its proof effectively establishes uniqueness of the equilibrium distribution over (a, C) . The restriction to essential equilibria serves to pin down the essential components of the narratives that accompany each policy-coalition pair.*

The unique essential equilibrium has the following noteworthy features. First, the rational policy h is always taken with positive probability. Second, this is the only policy taken in equilibrium if and only if it has higher aggregate intrinsic support than ℓ (i.e., $F(h, N) \geq F(\ell, N)$). Third, when policy ℓ is taken with positive probability, it is accompanied by tribal false narratives that take the exclusionary form: They identify a collection $S \in \mathcal{S}$ of social groups that oppose h but are not part of the coalition supporting ℓ , and essentially argue that “things are good when S is out of power.” When $S = \emptyset$, this is reduced to the denialist narrative.

In the case of $n = 2$, exclusionary and inclusionary tribal narratives were equivalent. This is no longer the case when $n > 2$. What makes exclusionary tribal narratives effective in this context? When a social group opposes h , there is positive correlation between the good outcome and the group being out of power. The exclusionary tribal narrative allows a coalition to exploit this correlation to generate a false belief among its members that the very

exclusion of some groups from the coalition will lead to a good outcome, while advocating policy ℓ . This device enables the coalition to “*have its cake and eat it:*” The coalition reaps the mobilization benefits of the intrinsically more attractive policy ℓ , while using the tribal narrative to deflect responsibility for a bad outcome away from this policy and instead “scapegoat” the excluded groups for it.

Exclusionary tribal narratives balance a trade-off between breadth and intensity of the support they generate. Excluding groups from a coalition is costly because it forgoes their support and thus lowers its aggregate mobilization potential. However, if this exclusion is not too frequent, its correlation with $a = h$ (and hence $y = 1$) remains sufficiently strong, thus generating intense support from the coalition’s members. At one extreme, the denialist narrative garners the largest coalition because it does not exclude any group, but this comes at the cost of a weaker belief of $y = 1$ and therefore weaker mobilization of coalition members.

In fact, the denialist narrative plays an important role in the equilibrium characterization. The upper bound on $\sigma^*(h, N^h, \{0\})$ in Proposition 2 is the probability of $a = h$ in a model in which the only feasible narratives are the true and denialist ones. To see why, note that in such a model the payoff from platform $(h, N^h, \{0\})$ is $q \cdot F(h, N)$, whereas the payoff from $(\ell, N^\ell, \emptyset)$ is $q \cdot \sigma^*(h, N^h, \{0\}) \cdot F(\ell, N)$. Since $(h, N^h, \{0\})$ is always in the support of an essential equilibrium, the payoff from $(\ell, N^\ell, \emptyset)$ cannot exceed the payoff from $(h, N^h, \{0\})$. This implies the upper bound on $\sigma^*(h, N^h, \{0\})$.

Why do equilibrium platforms advocating policy ℓ refrain from employing “inclusionary” tribal narratives S that attribute the outcome to members of the platform’s coalition? The answer is that in order to generate a positive belief, S must also be contained in the coalition that supports policy h . But since the inclusion of S in a ruling coalition is associated with a good outcome, S cannot be part of any “exclusionary” tribal narrative that accompanies policy ℓ . This means that S will be part of every coalition that supports ℓ .

It follows that S will *always* be in power. As a result, its power status is uncorrelated with y and therefore a redundant aspect of any narrative that cannot add to its explanatory power. Part (ii) of the definition of essential equilibrium thus rules out the use of inclusionary tribal narratives.

5 Two Special Cases

In this section, we provide a richer characterization for two specifications of \mathcal{S} . Throughout this section, we assume $F(\ell, N^h) > F(h, N^h)$ —i.e., among the groups that support h , there is greater aggregate mobilization potential for ℓ . This means that while h is the better policy for y , ℓ has larger intrinsic appeal.

5.1 Narratives Based on a Social Taxonomy

In this sub-section, we assume that \mathcal{S} takes the form of a sequence of progressively finer *partitions* $\Pi = (\pi_1, \dots, \pi_K)$. That is, for every $k = 2, \dots, K$, π_k is a partition of N that is a refinement of the partition π_{k-1} . Note that by the general restrictions we imposed on the domain of feasible narratives, π_1 consists of the non-empty sets $N^\ell \cap N^h$, $N \setminus N^h$ and $N \setminus N^\ell$. We refer to Π as a *social taxonomy*. We sometimes abuse notation and write $S \in \Pi$ to refer to a cell in one of the partitions in Π .

We interpret each cell in a partition as a collection of social groups that has a clear label or a defining attribute. For example, the broadest partition can consist of the left, the right, and the center, with finer partitions representing finer distinctions (e.g., “progressives” and “moderates” within the Democratic wing in US politics). Another example is the division of Israeli society into increasingly fine ethnic-religious groups (e.g., “secular Jewish” vs. “religious Jewish” at a coarse level, then splitting the latter group into “nationalist-religious” and “ultra-orthodox” sub-categories). Finally, a typical socioeconomic division of Italian society is between employees, retirees,

and entrepreneurs. Employees are further divided into those who are unionized and highly protected and those who are not.

We impose some structure on Π : For every layer $k > 1$ there is an integer $r_k > 1$ such that, for every $S \in \pi_{k-1}$, there are *exactly* r_k sets $S' \in \pi_k$ that satisfy $S' \subset S$. In other words, at every layer $k - 1 < K$ in the social taxonomy, each category is split into the same number r_k of sub-categories in the next layer. For $k = 1$, let $r_1 = 1$. Denote

$$R = \sum_{k=1}^K (r_k - 1)$$

This number is a natural measure of how fragmented the social taxonomy is, because it increases with the number of levels K and with the number of splits-per-cell r_k at each layer k of the social taxonomy.

Proposition 3 *The unique essential equilibrium σ^* satisfies*

$$\sigma^*(h, N^h, \{0\}) = \frac{F(h, N)}{F(\ell, N^h) + \max(R, 1) \cdot F(\ell, N \setminus N^h)} \quad (6)$$

Moreover:

- (i) *Every $S \subseteq N \setminus N^h$ that belongs to one of the partitions in Π is used as a narrative in σ^* .*
- (ii) *The denialist narrative is used in the support of σ^* if and only if $K = 1$.*

Let us highlight a few features of the unique essential equilibrium in the social-taxonomy case. First, we obtain an explicit formula for the probability that the rational policy is taken, which decreases with R . Thus, political fragmentation can be detrimental to the implementation of socially beneficial policies because it creates more room for false tribal narratives.

Second, every social category that is weakly finer than $N \setminus N^h$ serves as an exclusionary tribal narrative in the equilibrium. To see the intuition, suppose

that some category in the social taxonomy is invoked by an exclusionary tribal narrative, but one of its direct sub-categories is never used. That is, the support of σ includes a platform that uses some narrative $S \in \pi_k$, but there is a narrative $S' \subset S$ that belongs to π_{k+1} and is not employed by any platform. Because of the nested partition structure of Π , the cells that weakly contain S and S' are the same. That is, realization of σ^* in which the members of S are not in power are the same as the realizations in which the members of S' are not in power. This means that the narratives S and S' generate the same beliefs. However, the larger coalition that attributes the outcome to S' would generate stronger mobilization. Therefore, a platform that employs the narrative S' would have a strictly higher payoff, which cannot happen in equilibrium.

Third, the denialist narrative is used only in the special case in which the social taxonomy lacks a finer distinction than the broad classification given by π_1 . In this case, $\sigma^*(h, N^h, \{0\})$ is equal to the upper bound $F(h, N)/F(\ell, N)$. The only other case in which the latter property holds is when the social taxonomy breaks $N \setminus N^h$ into precisely two sub-categories, and there is no finer categorization (in this case, $R = 1$).

Formula (6) enables us to subject $\sigma^*(h, N^h, \{0\})$ to some simple comparative statics with respect to the primitives of the model. First, a *richer taxonomy* (as captured by increasing R) lowers the equilibrium value of $\sigma^*(h, N^h, \{0\})$, due to the proliferation of false narratives in equilibrium. In this sense, a more fragmented political system intensifies the “tribal” character of public opinion, and this has a cost in terms of the common good.

Now consider changes in the mobilization potential function that reflect *more polarized attitudes* toward policy ℓ . Specifically, suppose that we lower $F(\ell, N^h)$ and raise $F(\ell, N \setminus N^h)$, while keeping $F(\ell, N)$ fixed (without upsetting the inequality $F(\ell, N^h) > F(h, N^h)$ that we assumed at the outset). This captures a shift of support for ℓ from the “center” to the “right”, resulting in a more polarized society. When the social taxonomy satisfies $R > 1$, this

shift lowers $\sigma^*(h, N^h, \{0\})$. In this sense, greater polarization is detrimental for the common good.

5.2 A Rich Domain of Narratives

In this sub-section, we consider the extreme case in which \mathcal{S} is the set of *all* subsets $S \subseteq N$ that are weakly contained in $N \setminus N^h$, $N \setminus N^\ell$, or $N^\ell \cap N^h$. We refer to this specification of \mathcal{S} as the “rich” narrative domain.

Proposition 4 *Under the rich narrative domain, the unique essential equilibrium σ^* satisfies the following properties:*

- (i) $\sigma^*(h, N^h, \{0\}) = \frac{F(h, N)}{F(\ell, N)}$
- (ii) *The rest of $\text{Supp}(\sigma^*)$ consists of all platforms $(\ell, N^\ell \setminus S, S)$ such that $S \subseteq N \setminus N^h$ and $|S| \geq |N \setminus N^h| - 1$.*

Thus, when the narrative domain is rich, the equilibrium structure is simple. The rational policy is played with the maximal possible probability, as given by our main result. As usual, the rational policy is carried by the true narrative and supported by the coalition N^h . All false narratives are tribal. In each of them, the collection of groups that is held accountable is either $N \setminus N^h$ (i.e., everyone that opposes h) or a subset that excludes exactly one group from this set. In these cases, this excluded group joins the coalition supporting ℓ .

We can see that even though the domain of possible tribal narratives is maximally rich, only a small subset of them is used. The result also demonstrates the non-monotonicity of the effect of tribalism and political fragmentation on the equilibrium probability of the rational policy. The rich domain represents a larger scope for tribal narratives than the taxonomies studied in the previous sub-section. Nevertheless, the probability of rational behavior is higher than for most social taxonomies. The reason is that if a

platform uses an exclusionary tribal narrative that is given by a small set S , the collection of narratives in \mathcal{S} that contain S is larger than in the taxonomy case. This means that in the rich domain case, using small narratives S significantly dilutes the belief that $y = 1$, because there are many instances where S is not in power and $a = \ell$. This effect deters the entry of platforms that rely on such small narratives, and therefore blocks their proliferation.

6 Model Foundations

6.1 A “Microfoundation” for Mobilization Potentials

In this section, we illustrate how our definition of group i ’s support for platforms, $u_{i,\sigma}$, can be derived from a more detailed model in which *anticipatory utility* drives political participation.

Each social group i consists of a measure m_i of individuals. Each individual has a payoff function $v(y, a) = y - c_{a,i}$, where the idiosyncratic cost $c_{a,i}$ is drawn from the uniform distribution over $[0, \bar{c}_{a,i}]$. Thus, the constant $\bar{c}_{a,i} > 0$ is an action-specific characteristic of the social group i .

Suppose group i belongs to a coalition C that is part of the platform (a, C, S) . The group’s conditional belief over y is given by $p_\sigma(y \mid x_S(a, C))$. The platform mobilizes an individual in group i with cost $c_{a,i}$ if and only if his subjective anticipatory utility from the platform is positive—that is, if $p_\sigma(y = 1 \mid x_S(a, C)) > c_{a,i}$. As a result, the total mass of mobilized individuals from group i is

$$m_i \cdot \frac{p_\sigma(y = 1 \mid x_S(a, C))}{\bar{c}_{a,i}}$$

This is consistent with our definition of $u_{i,\sigma}$ with respect to a mobilization potential $f_i(a) = m_i/\bar{c}_{a,i}$.

It should be emphasized that this “microfoundation” by no means rationalizes the individuals’ behavior, because they are not engaged in conse-

quentialist analysis of their participation decision. Nevertheless, it offers a deeper understanding of the psychology of motivated reasoning behind our mobilization-potential function.

6.2 A Dynamic Convergence Result

In this section, we consider a simple and natural dynamic process that determines which platforms garner maximal support over time. We show that the process converges to the unique equilibrium distribution over policies and coalitions in our main result. This global convergence result provides a dynamic foundation for our equilibrium concept.

Time is discrete and denoted by $t = 1, 2, \dots$. In each period t , there is a distribution σ_t over platforms (a, C, S) , where $a \in \{\ell, h\}$, $C \subseteq N$, and $S \in \mathcal{S}$. Let the initial σ_1 be any distribution with full support over the set of platforms using admissible coalitions. Since the set of platforms is finite, this distribution is well-defined. The distribution σ_t evolves according to the following adjustment. For every $t \geq 2$, let

$$\overline{(a, C, S)}_t \in \arg \max_{(a', C', S')} U_{\sigma_t}(a', C', S'),$$

where ties can be broken arbitrarily. Then, let

$$\sigma_{t+1}(a, C, S) = \begin{cases} \frac{1}{t+1} + \frac{t}{t+1} \sigma_t(a, C, S) & \text{if } (a, C, S) = \overline{(a, C, S)}_t \\ \frac{t}{t+1} \sigma_t(a, C, S) & \text{otherwise.} \end{cases}$$

Thus, for t large enough, we can essentially view $\sigma_t(a, S, C)$ as the empirical frequency with which platform (a, C, S) has been dominant in the available history of data.

Proposition 5 *Every limit point σ of the process σ_t induces the same distribution over policy-coalition pairs (a, C) as that induced by the unique essential equilibrium σ^* .*

This result formalizes and generalizes the dynamic convergence process we discussed in the context of the two-group specification in Section 3.

Comment: An alternative dynamic interpretation

Throughout the paper, we interpreted the platform support function $u_{i,\sigma}(a, C, S)$ given by (2) as if agents have a long memory, which enables them to measure the long-run conditional frequency $p_\sigma(y = 1 \mid x_S(a, C))$. This interpretation also underlies the dynamic process considered in this sub-section.

An alternative interpretation is that agents reason anecdotally. When a combination a, C is proposed and a narrative S is suggested at some time period, agents are reminded of the last period in which $x_S(a, C)$ took the same value. For example, if the narrative attributes the outcome to whether group i is in power, agents' attention is drawn to the last period in which this group was indeed in power, and they record the realization of y during that period. This means that $p_\sigma(y = 1 \mid x_S(a, C))$ does not represent a probabilistic belief, but rather the probability that a random anecdote will be favorable for the platform (a, C, S) .

This alternative interpretation is realistic in the sense that it reflects voters' short memory and tendency to overreact to single, recent anecdotes. Whether a dynamic process based on this interpretation induces long-run frequencies of dominant platforms that match our notion of equilibrium is left as an open problem.

7 Related Literature

Eliaz and Spiegler (2020) presented the basic idea of formalizing political narratives as causal models whose adoption by agents is driven by motivated

reasoning. The present paper borrows these ingredients. The key difference is that Eliaz and Spiegler (2020) considered a representative-agent model, whereas the present paper assumes a heterogeneous society and focuses on the role of false narratives as the “glue” of social coalitions. Indeed, the existence of multiple social groups in the present model enables the new class of “tribal” narratives, which are moot in the single-agent model.

More broadly, this paper is related to a strand in the political-economics literature that studies voters’ belief formation according to misspecified subjective models or wrong causal attribution rules (e.g., Spiegler (2013), Esponda and Pouzo (2017), Izzo et al. (2021), and Levy et al. (2022)). In particular, the latter paper studies dynamic electoral competition between two candidates, each associated with a different subjective model of how two policy variables map into outcomes. One model is complete and correct; the other is a “simplistic” model that omits one of the policy variables. Voter participation is costly; stronger beliefs lead to larger voter turnout. The long-run behavior of this system involves ebbs and flows in the relative popularity of the two models.

The general program of studying the behavioral implications of misspecified causal models is due to Spiegler (2016; 2020). In their general form, causal models are formalized as directed acyclic graphs, following the Statistics/AI literature on graphical probabilistic models (Cowell et al. (1999), Pearl (2009)). The causal models in this paper fit into the graphical formalism (as we saw in Section 2), but do not require its heavy use because they take a relatively simple form (this form is related to misspecified models in otherwise very different works, such as Jehiel (2005), Eyster and Piccione (2013) or Mailath and Samuelson (2020)). Therefore, in this paper, graphical representations of causal models remained mostly in the background.

Given the fluidity of the notion of narratives, it naturally invites diverse formalizations. Bénabou et al. (2018) focus on moral decision-making and formalize narratives as messages or signals that can affect decision-makers’

beliefs regarding the externality of their policies. Levy and Razin (2021) use the term to describe information structures in game-theoretic settings that people postulate in order to explain observed behavior. Schwartzstein and Sunderam (2021a; 2021b) propose an alternative approach to “persuasion by models,” where models are formalized as likelihood functions and the criterion for selecting models is their success in accounting for historical observations. Shiller (2017) focuses on the spread of economic narratives in society, using an epidemiological analogy.

The political science literature has long acknowledged the power of narratives in garnering public support for policies and in mobilizing people to protests or rallies (see Polletta (2008)). In particular, the so-called “narrative policy framework” was developed as a systematic empirical framework for studying the role of stories or narratives in public policy. Studies employing this framework have argued that narratives have a greater influence on the opinions of policymakers and citizens than does scientific information (see, e.g., Shanahan et al. (2011), Jones and McBeth (2010), and Jones et al. (2014)).

Finally, there are a few recent attempts to study political and economic narratives empirically, using textual analysis. Mobilizing public opinion often takes the form of texts (speeches, op-eds, tweets). What we observe in these texts are qualitative stories more than bare quantitative beliefs. Textual analysis can elicit the element of causal attribution in political narratives from these texts (Ash et al. (2021) and Andre et al. (2022) have performed manual and machine analysis of these texts in order to elicit prevailing narratives in various contexts). Our formalism, which explicitly focuses on narratives as simple causal models, will hopefully provide tools for interpreting such textual data and linking them to economic and social indicators.

8 Conclusion

This paper has explored the role of false narratives in the mobilization of public opinion in heterogeneous societies. Our main insight is that false narratives enable social groups to dissociate the link between the intrinsic appeal of certain policies and their adverse outcome. They achieve this by attributing outcomes to spurious causes, exploiting historical correlations and misrepresenting them as causal. This typically takes the form of exclusionary tribal narratives, which argue that keeping certain social groups out of power leads to good outcomes. Such narratives are reminiscent of “scapegoating,” a type of narrative that is often used in the political arena.

This insight suggests a novel perspective into the idea of *retrospective voting* (see Healy and Malhotra (2013) for a review article, and Plescia and Kritzinger (2017) for an example that extends the concept to multi-party systems). This is the notion that voters punish or reward parties according to their performance (measured by certain outcomes) when they were in office. This view puts less emphasis on processes (i.e., the policies that ruling parties take) and more emphasis on outcomes. The conventional view is that retrospective voting is a “healthy” feature of democratic politics because it improves government accountability and helps selecting competent candidates. By comparison, our perspective is that attributing public outcomes to who is (or is not) in power rather than to the implemented policies can be a demagogic false narrative that is detrimental for public outcomes. This offers a new and more critical view of retrospective voting.

Appendix: Proofs

Proof of Proposition 2

We organize the proof in a series of steps. We will posit the existence of an essential equilibrium, characterize its properties, and then confirm that we indeed have an equilibrium. Hereafter, let σ be any candidate essential equilibrium. Note that by definition, $F(a, N) = F(a, N^a)$. We use the two notations interchangeably for expositional purposes. For convenience, hereafter we denote

$$d = F(\ell, N) - F(h, N) \tag{7}$$

Step 1 *There exists $(a, C, S) \in \text{Supp}(\sigma)$ such that $a = h$.*

Proof. Assume the contrary—i.e., $a = \ell$ for every $(a, C, S) \in \text{Supp}(\sigma)$. Then, $p_\sigma(y = 1) = 0$. Therefore,

$$U_\sigma(a, C, S) = p_\sigma(y = 1 \mid x_S(a, C)) = 0$$

for every $(a, C, S) \in \text{Supp}(\sigma)$. By the definition of equilibrium, σ is the limit of a sequence of ε -equilibria for some $\varepsilon \rightarrow 0$. Since $\sigma(a, C, S) > 0$, $\sigma_\varepsilon(a, C, S)$ is bounded away from zero, and therefore $U_{\sigma_\varepsilon}(a, C, S) \approx p_{\sigma_\varepsilon}(y = 1 \mid x_S(a, C)) \approx 0$, for some point along the sequence $\varepsilon \rightarrow 0$. By contrast, $U_{\sigma_\varepsilon}(h, N^h, \{0\}) = q \cdot F(h, N)$, which is bounded away from zero and therefore higher than $U_{\sigma_\varepsilon}(a, C, S)$. Since by assumption $(h, N^h, \{0\}) \notin \text{Supp}(\sigma)$, we have a contradiction. ■

This step is the only place in the proof where we use the trembles of ε -equilibria, to ensure that the payoff from $(h, N^h, \{0\})$ is well-defined. From now on, we focus on the $\varepsilon \rightarrow 0$ limit itself.

Step 2 *The only $(a, C, S) \in \text{Supp}(\sigma)$ with $a = h$ is $(h, N^h, \{0\})$.*

Proof. For every $(h, C, S) \in \text{Supp}(\sigma)$,

$$p_\sigma(y = 1 \mid x_S(h, C)) = q \cdot p_\sigma(x_0 = h \mid x_S(h, C)) \leq q$$

whereas by definition, $p_\sigma(y = 1 \mid x_0 = h) = q$. Moreover, $F(h, C) \leq F(h, N^h)$ by definition. It follows that $U_\sigma(h, C, S) \leq U_\sigma(h, N^h, \{0\})$. Condition (i) in the definition of essential equilibrium then implies that $S = \{0\}$. Since $F(h, N^h) > F(h, C')$ for every $C' \subset N^h$, it follows that $C = N^h$. ■

Corollary 1 *Equilibrium payoffs are*

$$U^* = q \cdot F(h, N^h) \tag{8}$$

This follows immediately from Steps 1 and 2. Note that the expression for U^* is independent of σ . From now on, we use the notation

$$\alpha = \sigma(h, N^h, \{0\})$$

Step 3 *If $x_S(\ell, C) = x_S(h, N^h)$, then*

$$p_\sigma(y = 1 \mid x_S(\ell, C)) = \frac{q\alpha}{\alpha + \sum_{C', S' \mid x_S(\ell, C') = x_S(\ell, C)} \sigma(\ell, C', S')} \tag{9}$$

Otherwise, $p_\sigma(y = 1 \mid x_S(\ell, C)) = 0$.

Proof. Suppose $0 \notin S$. By definition,

$$p_\sigma(y = 1 \mid x_S(\ell, C)) = \frac{q \cdot \sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma(h, C', S')}{\sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma(a', C', S')}$$

By Step 2, the numerator can be rewritten as

$$q \cdot \sigma(h, N^h, \{0\}) \cdot \mathbf{1}[x_S(\ell, C) = x_S(h, N^h)]$$

which delivers (9). (Note that when $0 \notin S$, $x_S(\ell, C) = x_S(h, C')$ if and only if $S \cap C = S \cap C'$.) Now suppose $0 \in S$. Then,

$$p_\sigma(y = 1 \mid x_S(\ell, C)) = p_\sigma(y = 1 \mid x_0 = \ell) = 0 \quad (10)$$

■

Corollary 2 *For every $(\ell, C, S) \in \text{Supp}(\sigma)$, $0 \notin S$.*

Proof. Suppose $0 \in S$. By (10), $U_\sigma(\ell, C, S) = 0 < U^*$, hence $(\ell, C, S) \notin \text{Supp}(\sigma)$. ■

Step 4 *If $F(\ell, N) \leq F(h, N)$, then $\alpha = 1$. If $F(\ell, N) > F(h, N)$, then*

$$\alpha \leq \frac{F(h, N)}{F(\ell, N)}$$

Proof. Suppose $F(\ell, N) \leq F(h, N)$, but $\alpha < 1$. Then, there exists $(\ell, C, S) \in \text{Supp}(\sigma)$, which implies by definition that the denominator of (9) is greater than α and hence $p_\sigma(y = 1 \mid x_S(\ell, C)) < q$. It follows that

$$U_\sigma(\ell, C, S) = p_\sigma(y = 1 \mid x_S(\ell, C)) \cdot F(\ell, C) < q \cdot F(\ell, N) \leq q \cdot F(h, N) = U^*$$

which is a contradiction. Thus, in this case, $\alpha = 1$. Suppose $F(\ell, N) > F(h, N)$. If $\alpha = 1$, then

$$U_\sigma(\ell, N^\ell, \emptyset) = p_\sigma(y = 1)F(\ell, N) = qF(\ell, N) > U^*$$

which is again a contradiction. Thus, in this case, $\alpha < 1$. Since we must have $U_\sigma(\ell, N^\ell, \emptyset) \leq U^*$ in any equilibrium, and since $p_\sigma(y = 1) = q\alpha$, it follows that $q\alpha \cdot F(\ell, N) \leq q \cdot F(h, N)$. This implies that upper bound on α in this case. Note that this upper bound relies on the assumption that the denialist narrative is feasible (i.e., $\emptyset \in \mathcal{S}$). ■

The next step establishes that in equilibrium, false narratives take the “exclusionary” tribal form (including the denialist narrative as a special case).

Step 5 *If $(\ell, C, S) \in \text{Supp}(\sigma)$, then $S \subseteq N \setminus N^h$ and $C = N^\ell \setminus S$.*

Proof. The proof consists of two steps. First, we show that $N^h \cap N^\ell \subseteq C$ for every $(a, C, S) \in \text{Supp}(\sigma)$. Then, we show that this implies the claim. By Step 3, $0 \notin S$ for every $(\ell, C, S) \in \text{Supp}(\sigma)$. By the restriction we imposed on the domain of feasible narratives, S is a weak subset of $N \setminus N^\ell$, $N \setminus N^h$ or $N^h \cap N^\ell$. Let us consider each of these cases.

First, suppose $S \subseteq N \setminus N^\ell$. By our definition of admissible coalitions, any platform (ℓ, C, S) satisfies $C \cap S = \emptyset$ —i.e., $x_i(\ell, C) = 0$ for every $i \in S$. By Step 3, $p_\sigma(y = 1 \mid x_S(\ell, C)) = 0$, hence $U_\sigma(\ell, C, S) = 0$, a contradiction.

Second, suppose $S \subseteq N \setminus N^h$. If $C \subset N^h \cap N^\ell$, then $U_\sigma(\ell, C, S) < U_\sigma(\ell, N^h \cap N^\ell, S)$. The reason is that $p_\sigma(y = 1 \mid x_S(\ell, C)) = p_\sigma(y = 1 \mid x_S(\ell, N^h \cap N^\ell))$ (since the narrative S ignores the power status of groups in $N^h \cap N^\ell$), and $f_i(\ell) > 0$ for every $i \in N^h \cap N^\ell$. Since (ℓ, C, S) fails to maximize U_σ , we obtain a contradiction. It follows that $N^h \cap N^\ell \subseteq C$.

The only remaining case is thus $S \subseteq N^h \cap N^\ell$. By steps 1 and 2, $x_i(\ell, N^h) = 1$ for every $i \in S$. By Step 3, $p_\sigma(y = 1 \mid x_S(\ell, C)) > 0$ only if $x_i(\ell, C) = 1$ for every $i \in S$ —i.e., $S \subseteq C$. Since $f_i(\ell) > 0$ for every $i \in N^h \cap N^\ell$, C is an admissible coalition given policy ℓ . If $C \subset N^h \cap N^\ell$, then $U_\sigma(\ell, C, S) < U_\sigma(\ell, N^h \cap N^\ell, S)$, by the same argument as in the previous case. Therefore, $N^h \cap N^\ell \subseteq C$.

We have thus established that $S \subseteq N^\ell$ for every $(\ell, C, S) \in \text{Supp}(\sigma)$, and that $N^h \cap N^\ell \subseteq C$ for every platform in $\text{Supp}(\sigma)$. Suppose $S \subseteq N^h \cap N^\ell$ and $S \neq \emptyset$ for some $(\ell, C, S) \in \text{Supp}(\sigma)$. Then,

$$p_\sigma(y = 1 \mid x_S(\ell, C)) = p_\sigma(y = 1 \mid x_i(\ell, C) = 1 \text{ for every } i \in S) = p_\sigma(y = 1)$$

Therefore,

$$U_\sigma(\ell, C, \emptyset) = U_\sigma(\ell, C, S)$$

By condition (ii) in the definition of essential equilibrium, $(\ell, C, S) \notin \text{Supp}(\sigma)$, a contradiction.

The only remaining possibility is that $S \subseteq N \setminus N^h$. As we saw, this implies $N^h \cap N^\ell \subseteq C$. However, admissibility of C given ℓ requires $C \subseteq N^\ell$. It follows that $C = N^\ell \setminus S$. ■

The last step implies that platforms $(\ell, C, S) \in \text{Supp}(\sigma)$ are entirely pinned down by S . Therefore, in the rest of the proof, we use the notation $\bar{\sigma}(S) = \sigma(\ell, C, S)$.

Step 6 $(\alpha, \bar{\sigma})$ is an equilibrium if and only if, for all $S \in \mathcal{S}$ that satisfy $S \subseteq N \setminus N^h$,

$$\alpha \cdot \frac{d - F(\ell, S)}{F(h, N)} \leq \sum_{S' \in \mathcal{S} | S' \supseteq S} \bar{\sigma}(S') \quad (11)$$

with equality if $\bar{\sigma}(S) > 0$. (Recall that d is defined by (7).)

Proof. Using Definition 2, σ is an equilibrium if and only if $U_\sigma(\ell, C, S) \leq U^*$ for all (ℓ, C, S) , with equality if $\sigma(\ell, C, S) > 0$. By Step 5, we can restrict attention to platforms (ℓ, C, S) for which $S \subseteq N \setminus N^h$ and $C = N^\ell \setminus S$. Plugging expressions (8) and (9), the inequality $U_\sigma(\ell, C, S) \leq U^*$ can be rewritten as a linear inequality in σ :

$$\alpha \cdot \frac{F(\ell, C) - F(h, N)}{F(h, N)} \leq \sum_{C', S' | x_S(\ell, C') = x_S(\ell, C)} \sigma(\ell, C', S')$$

Note that $F(\ell, C) - F(h, N) = d - F(\ell, S)$, such that the L.H.S of the last inequality becomes the L.H.S of (11). Finally, by Step 5, $(\ell, C', S') \in \text{Supp}(\sigma)$ satisfies $x_S(\ell, C') = x_S(\ell, C)$ if and only if $S' \supseteq S$ (since $C = N^\ell \setminus S$ and $C' = N^\ell \setminus S'$). This means that we can replace the R.H.S of the last inequality with the R.H.S of (11). ■

The last step immediately implies the following observation, which means that exclusionary tribal narratives will “scapegoat” a collection of groups only when its aggregate mobilization potential is not too large.

Corollary 3 *For every $S \subseteq N \setminus N^h$, if $F(\ell, S) \geq d$, then $\bar{\sigma}(S) = 0$.*

Inequalities (11) enable us to construct an algorithm that pins down the essential-equilibrium distribution.

An algorithm for computing $\bar{\sigma}(S)$ for all $S \in \mathcal{S}$ such that $S \subseteq N \setminus N^h$

Let

$$\bar{\mathcal{S}} = \{S \in \mathcal{S} \mid S \subseteq N \setminus N^h \text{ and } F(\ell, S) < d\}.$$

Define

$$\bar{\mathcal{S}}_1 = \{S \in \bar{\mathcal{S}} \mid \text{there is no } S' \in \bar{\mathcal{S}} \text{ such that } S \subset S'\}$$

Now, for every $k > 1$, define $\bar{\mathcal{S}}_k$ recursively as follows:

$$\bar{\mathcal{S}}_k = \{S \in \bar{\mathcal{S}} \mid \text{there is no } S' \in \bar{\mathcal{S}} \setminus \cup_{j < k} \bar{\mathcal{S}}_j \text{ such that } S \subset S'\}$$

Since $\bar{\mathcal{S}}$ is finite, in this way we obtain a finite sequence $\{\bar{\mathcal{S}}_k\}_{k=1}^K$. This sequence exhausts all the candidate narratives, as it identifies all exclusionary tribal narratives allowed by the primitive \mathcal{S} . The algorithm starts from the “top layer” of $\bar{\mathcal{S}}$ (i.e., $\bar{\mathcal{S}}_1$) and then proceeds to the other layers in order. For every $S \in \bar{\mathcal{S}}_1$, (11) can be written as

$$\bar{\sigma}(S) \geq \alpha \cdot \frac{d - F(\ell, S)}{F(h, N)}$$

By the definition of $\bar{\mathcal{S}}$, the R.H.S is strictly positive for every $S \in \bar{\mathcal{S}}_1$, which implies that S is in the equilibrium support and therefore the inequality must hold with equality. This pins down $\bar{\sigma}(S)$. For every $S \in \bar{\mathcal{S}}$, denote $\mathcal{H}(S) = \{S' \in \bar{\mathcal{S}} \mid S \subset S'\}$. By definition, if $S \in \bar{\mathcal{S}}_k$, then $\mathcal{H}(S) \subseteq \cup_{j < k} \bar{\mathcal{S}}_j$. Now, by induction, suppose that for all $j < k$ and every $S \in \bar{\mathcal{S}}_j$, there exists $w(S) \geq 0$ such that $\bar{\sigma}(S) = \alpha w(S)$. For $S \in \bar{\mathcal{S}}_1$, we have already established

that, where $w(S) = (d - F(\ell, S))/F(h, N)$. For every $S \in \overline{\mathcal{S}}_k$, (11) becomes

$$\bar{\sigma}(S) = \max \left\{ 0, \alpha \cdot \frac{d - F(\ell, S)}{F(h, N)} - \alpha \sum_{S' \in \mathcal{H}(S)} w(S') \right\}$$

where $w(S')$ is well-defined for all $S' \in \mathcal{H}(S)$, by the inductive argument. This confirms that $\bar{\sigma}(S) = \alpha w(S)$, where

$$w(S) = \max \left\{ 0, \frac{d - F(\ell, S)}{F(h, N)} - \sum_{S' \in \mathcal{H}(S)} w(S') \right\}$$

completing the inductive argument, and thus the definition of the algorithm for computing $\bar{\sigma}(S)$.

Step 7 *The algorithm establishes existence and uniqueness of an essential equilibrium.*

Proof. Since $(\alpha, \bar{\sigma})$ must define a probability distribution, we must have

$$\alpha + \sum_{S \in \overline{\mathcal{S}}} \bar{\sigma}(S) = 1$$

Moreover, we have obtained unique expressions for each $\bar{\sigma}(S)$ that depend multiplicatively on α . This pins down the value of α :

$$\alpha = \frac{1}{1 + \sum_{S \in \overline{\mathcal{S}}} w(S)}$$

Thus, we have pinned down $(\alpha, \bar{\sigma})$. Since this pair satisfies all the inequalities (11), it is therefore an equilibrium by construction. ■

Proof of Proposition 3

We use the algorithm in the proof of Proposition 2 to characterize the equilibrium. We begin by calculating the equilibrium probabilities of the cells

that comprise Π . Note that $N \setminus N^h \in \pi_1$ and $F(\ell, N^h) > F(h, N^h)$ imply that $F(\ell, S) < d$ for every $S \subseteq N \setminus N^h$ that belongs to one of the partitions in Π . Given this, for $N \setminus N^h \in \pi_1$, we must have

$$\bar{\sigma}(N \setminus N^h) = \alpha \cdot \frac{d - F(\ell, N \setminus N^h)}{F(h, N)} > 0$$

Let $k > 1$. Given $S_k \in \pi_k$ such that $S_k \subset N \setminus N^h$, the collection of sets $\mathcal{H}(S_k) = \{S' \in \bar{\mathcal{S}} \mid S_k \subset S'\}$ in the proof of Proposition 2 takes the form of a chain $\{S_j\}_{j=1}^{k-1}$ that satisfies $S_j \in \pi_j$ and $S_{j+1} \subset S_j \subset N \setminus N^h$ for all $j < k$. For $S_2 \in \pi_2$, we must have

$$\begin{aligned} \bar{\sigma}(S_2) &= \frac{1}{F(h, N)} \max\{0, \alpha(d - F(\ell, S_2)) - \alpha(d - F(\ell, S_1))\} \\ &= \frac{1}{F(h, N)} \max\{0, \alpha F(\ell, S_1 \setminus S_2)\} > 0 \end{aligned}$$

Thus, the coefficient $w(S_2)$ in the proof of Proposition 2 takes the form $F(\ell, S_1 \setminus S_2)/F(h, N)$. By induction,

$$\bar{\sigma}(S_k) = \alpha \frac{F(\ell, S_{k-1} \setminus S_k)}{F(h, N)} > 0 \quad (12)$$

for every $S_k \in \pi_k$ such that $S_k \subset N \setminus N^h$ and $k = 2, \dots, K$. This completes the characterization of the $\bar{\sigma}(S)$ for every cell S in one of the partitions in Π . Now, consider $S = \emptyset$. In this case, we need

$$\alpha \frac{d}{F(h, N)} - \sum_{S \in \Pi \mid S \subseteq N \setminus N^h} \bar{\sigma}(S) \leq \bar{\sigma}(\emptyset) \quad (13)$$

with equality if $\bar{\sigma}(\emptyset) > 0$. Suppose that $K = 1$ (such that $R = 0$), and hence only $N \setminus N^h$ itself belongs to Π . Then,

$$\alpha \frac{d}{F(h, N)} - \bar{\sigma}(N \setminus N^h) = \alpha \frac{F(\ell, N \setminus N^h)}{F(h, N)} > 0$$

which implies that

$$\bar{\sigma}(\emptyset) = \alpha \frac{F(\ell, N \setminus N^h)}{F(h, N)} > 0$$

If instead $R \geq 1$, given our assumptions on Π it means that at least two subsets $S \subset N \setminus N^h$ belong to π_2 . In this case,

$$\begin{aligned} \sum_{S \in \Pi | S \subseteq N \setminus N^h} \bar{\sigma}(S) &\geq \frac{\alpha}{F(h, N)} \{d - F(\ell, N \setminus N^h)\} + \sum_{S \in \pi_2 | S \subseteq N \setminus N^h} \bar{\sigma}(S) \\ &= \frac{\alpha}{F(h, N)} (d - F(\ell, N \setminus N^h)) \\ &\quad + \frac{\alpha}{F(h, N)} \sum_{S \in \pi_2 | S \subseteq N \setminus N^h} F(\ell, N \setminus \{N^h \cup S\}) \\ &\geq \frac{\alpha}{F(h, N)} (d - F(\ell, N \setminus N^h)) + \frac{\alpha}{F(h, N)} F(\ell, N \setminus N^h) \\ &= \alpha \frac{d}{F(h, N)} \end{aligned}$$

The last inequality follows from the fact that since π_2 is a collection of disjoint sets whose union is $N \setminus N^h$, the summation over $S \in \pi_2 \mid S \subseteq N \setminus N^h$ counts $F(\ell, S)$ at least once, and potentially more than once. Plugging this inequality in (13), we obtain $\bar{\sigma}(\emptyset) = 0$. It remains to calculate α . Let's start from the case of $R = 0$. Plugging the expressions we obtained for $\bar{\sigma}(N \setminus N^h)$ and $\bar{\sigma}(\emptyset)$, we have

$$1 = \alpha + \bar{\sigma}(N \setminus N^h) + \bar{\sigma}(\emptyset) = \alpha \left\{ 1 + \frac{d - F(\ell, N \setminus N^h)}{F(h, N)} + \frac{F(\ell, N \setminus N^h)}{F(h, N)} \right\}$$

such that

$$\alpha = \frac{F(h, N)}{F(\ell, N)}$$

For all other cases (where $R > 0$), we saw that $\bar{\sigma}(\emptyset) = 0$. For every $S_k \in \pi_k$ such that $S_k \subset N \setminus N^h$, let S_{k-1} be again the antecedent of S_k in the chain $\{S_j\}_{j=1}^{k-1}$ that we used in the construction above. For every $S \in \pi_k$, let $P(S)$ be the unique cell $S' \in \pi_{k-1}$ such that $S \subseteq S'$. Given this, and plugging (12),

we have

$$\begin{aligned}
1 &= \alpha + \sum_{S \in \Pi | S \subseteq N \setminus N^h} \bar{\sigma}(S) \\
&= \frac{\alpha}{F(h, N)} \left\{ F(h, N) + d - F(\ell, N \setminus N^h) + \sum_{k=2}^K \sum_{S \in \pi_k | S \subseteq N \setminus N^h} F(\ell, P(S) \setminus S) \right\} \\
&= \frac{\alpha}{F(h, N)} \left\{ F(\ell, N^h) + \sum_{k=2}^K \sum_{S \in \pi_k | S \subseteq N \setminus N^h} F(\ell, P(S) \setminus S) \right\}
\end{aligned}$$

To further simplify this expression, we now use the assumption that each cell in π_{k-1} has exactly r_k subsets in π_k . Using this, we can rewrite the last condition as

$$\begin{aligned}
1 &= \frac{\alpha}{F(h, N)} \left\{ F(\ell, N^h) + F(\ell, N \setminus N^h) \sum_{k=2}^K (r_k - 1) \right\} \\
&= \frac{\alpha}{F(h, N)} \{ F(\ell, N^h) + F(\ell, N \setminus N^h) \cdot R \}
\end{aligned}$$

such that

$$\alpha = \frac{F(h, N)}{F(\ell, N^h) + F(\ell, N \setminus N^h) \cdot R}$$

This completes the proof.

Proof of Proposition 4

Let σ be the unique essential equilibrium. Since $F(\ell, N) \geq F(\ell, N^h) > F(h, N)$, Proposition 2 implies that $\sigma(h, N^h, \{0\}) = \alpha \in (0, 1)$. Let us now activate the algorithm described in the proof of Proposition 2. Recall that we can identify all the platforms in which $a = \ell$ with the narrative $S \subseteq N \setminus N^h$. Therefore, from now on we will use the abbreviated notation $\bar{\sigma}(S) = \sigma(\ell, C, S)$ from the proof of Proposition 2. Since $F(\ell, N^h) > F(h, N)$,

we have

$$\bar{\sigma}(N \setminus N^h) = \alpha \cdot \frac{F(\ell, N^h) - F(h, N)}{F(h, N)} > 0$$

Now consider a narrative $S = N \setminus (N^h \cup \{i\})$ for some $i \in N \setminus N^h$. The rich domain assumption means that $S \in \mathcal{S}$. By definition, $S \not\subseteq S'$ for any $S' \neq S$ such that $S' \subset N \setminus N^h$. Therefore, if $\bar{\sigma}(S) = 0$, then the following inequality must hold:

$$\alpha \cdot \frac{d - F(\ell, S)}{F(h, N)} = \alpha \cdot \frac{F(\ell, N^h \cup \{i\}) - F(h, N)}{F(h, N)} \leq \bar{\sigma}(N \setminus N^h)$$

which is a contradiction since

$$F(\ell, N^h \cup \{i\}) = F(\ell, N^h) + f_i(\ell) > F(\ell, N^h).$$

It follows that $\bar{\sigma}(N \setminus (N^h \cup \{i\})) > 0$ for every $i \in N \setminus N^h$, and given by

$$\bar{\sigma}(N \setminus (N^h \cup \{i\})) = \alpha \cdot \frac{f_i(\ell)}{F(h, N)}$$

We will now show that the support of $\bar{\sigma}$ contains no other narratives. Assume the contrary. Denote $|N \setminus N^h| = s^*$. Let $S \subset N \setminus N^h$ such that $|S| = s < s^* - 1$. In particular, select S to be a maximal set in this class—i.e., there exist no other such S that contains it. Then, if $\bar{\sigma}(S) > 0$, it must satisfy the equation

$$\alpha \cdot \frac{d - F(\ell, S)}{F(h, N)} - \bar{\sigma}(N \setminus N^h) - \sum_{i \in N \setminus (N^h \cup S)} \bar{\sigma}(N \setminus (N^h \cup \{i\})) = \bar{\sigma}(S)$$

Plugging our expressions for $\bar{\sigma}(N \setminus N^h)$ and $\bar{\sigma}(N \setminus (N^h \cup \{i\}))$, the L.H.S becomes

$$\alpha \cdot \left(\frac{d - F(\ell, S)}{F(h, N)} - \frac{F(\ell, N^h) - F(h, N)}{F(h, N)} - \frac{F(\ell, N \setminus (N^h \cup S))}{F(h, N)} \right) = 0$$

Thus, we obtain a contradiction. It remains to obtain the exact value of α .
By construction,

$$\alpha + \bar{\sigma}(N \setminus N^h) + \sum_{i \in N \setminus N^h} \bar{\sigma}(N \setminus (N^h \cup \{i\})) = 1$$

Plugging our expressions for $\bar{\sigma}(N \setminus N^h)$ and $\bar{\sigma}(N \setminus (N^h \cup \{i\}))$, we obtain $\alpha = F(h, N)/F(\ell, N)$. This completes the proof.

Proof of Proposition 5

In this proof, we denote platforms by z whenever convenient to simplify notation. For every t , let $\bar{z}_t = (\bar{a}_t, \bar{C}_t, \bar{S}_t) \in \arg \max_z U_{\sigma_t}(z)$ be the dominant platform at period t and let $\bar{U}_{\sigma_t} = U_{\sigma_t}(\bar{z}_t)$ be the payoff it generates. Note that if there exists T such that $\bar{z}_t \neq (a, C, S)$ for all $t \geq T$, then $\sigma_t(a, C, S) \rightarrow 0$ as $t \rightarrow \infty$. Recall that $U^* = q \cdot F(h, N^h) > 0$. The proof proceeds stepwise.

Step 8 $\bar{U}_{\sigma_t} \geq U^*$ for every t .

Proof. Since σ_1 has full support, $\sigma_t(h, N^h, \{0\}) > 0$ for every finite t ; therefore, $\bar{U}_{\sigma_t} \geq U_{\sigma_t}(h, N^h, \{0\}) = U^*$ for every t . ■

Step 9 If $\bar{z}_t = (h, C, S)$, then $C = N^h$ and $U_{\sigma_t}(h, C, S) = U^*$.

Proof. For every platform (h, C, S) such that $C \subset N^h$, $U_{\sigma_t}(h, C, S) < U_{\sigma_t}(h, N^h, \{0\})$ because $Pr_{\sigma_t}(y = 1 \mid x_S(h, C)) \leq q$ and $F(h, C) < F(h, N^h)$. This also implies that $U_{\sigma_t}(h, N^h, S) \leq U^*$ for all S and hence the last equality.

■

Step 10 For all t , there exists $t' > t$ such that $\bar{z}_{t'} = (h, N^h, S)$ for some S .

Proof. Step 8 implies that

$$\liminf_{t \rightarrow \infty} \bar{U}_{\sigma_t} \geq U^*.$$

Suppose there exists t such that $\bar{z}_{t'} = (\ell, \bar{C}_{t'}, \bar{S}_{t'})$ for all $t' \geq t$. This implies that $Pr_{\sigma_t}(y = 1 \mid x_{\bar{S}_t}(\bar{a}_t, \bar{C}_t)) \rightarrow 0$, which is inconsistent with $\liminf_{t \rightarrow \infty} \bar{U}_{\sigma_t} > 0$. ■

Step 11 $\liminf \bar{U}_{\sigma_t} = U^*$.

Proof. We have already established that $\liminf_{t \rightarrow \infty} \bar{U}_{\sigma_t} \geq U^*$. Note that, if $\bar{U}_{\sigma_t} > U^*$, then $\bar{z}_t = (\ell, C, S)$ for some C and S , because $U_{\sigma_t}(h, C', S') \leq U^*$ for all C' and S' . Now suppose $\liminf_{t \rightarrow \infty} \bar{U}_{\sigma_t} > U^*$. Then, there exists T such that for all $t \geq T$, \bar{z}_t involves policy $a = \ell$. This contradicts Step 10. ■

Recall that

$$Pr_{\sigma_t}(y = 1 \mid x_S(a, C)) = q \cdot \frac{\sum_{C', S' \mid x_S(h, C') = x_S(a, C)} \sigma_t(h, C', S')}{\sum_{a', C', S' \mid x_S(a', C') = x_S(a, C)} \sigma_t(a', C', S')}$$

Step 12 If $\bar{z}_t = (h, N^h, \hat{S})$ and $x_S(h, N^h) = x_S(\ell, C)$, then

$$Pr_{\sigma_{t+1}}(y = 1 \mid x_S(\ell, C)) > Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C))$$

Proof. Given $\bar{z}_t = (h, N^h, \hat{S})$, for every (ℓ, C, S) such that $x_S(h, N^h) = x_S(\ell, C)$,

$$\begin{aligned} Pr_{\sigma_{t+1}}(y = 1 \mid x_S(\ell, C)) &= q \frac{\frac{1}{t+1} + \frac{t}{t+1} \sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\frac{1}{t+1} + \frac{t}{t+1} \sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &= q \frac{\frac{1}{t} + \sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\frac{1}{t} + \sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &> q \frac{\sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &= Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C)) \end{aligned}$$

■

Step 13 If $\bar{z}_t = (\ell, \hat{C}, \hat{S})$, then for every (ℓ, C, S) ,

$$Pr_{\sigma_{t+1}}(y = 1 \mid x_S(\ell, C)) \leq Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C))$$

with strict inequality if and only if $x_S(\ell, \hat{C}) = x_S(\ell, C)$.

Proof. If $\bar{z}_t = (\ell, \hat{C}, \hat{S})$ and $x_S(\ell, \hat{C}) \neq x_S(\ell, C)$, then by definition, $Pr_{\sigma_{t+1}}(y = 1 \mid x_S(\ell, C)) = Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C))$. Now suppose that $\bar{z}_t = (\ell, \hat{C}, \hat{S})$ and $x_S(\ell, \hat{C}) = x_S(\ell, C)$. Then,

$$\begin{aligned} Pr_{\sigma_{t+1}}(y = 1 \mid x_S(\ell, C)) &= q \frac{\frac{t}{t+1} \sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\frac{1}{t+1} + \frac{t}{t+1} \sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &= q \frac{\sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\frac{1}{t} + \sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &< q \frac{\sum_{C', S' \mid x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{a', C', S' \mid x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \\ &= Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C)) \end{aligned}$$

■

Step 14 If (ℓ, C, S) is such that $x_S(\ell, C) \neq x_S(h, N^h)$, then $\sigma_t(\ell, C, S) \rightarrow 0$ as $t \rightarrow \infty$.

Proof. Suppose $\sigma_t(\ell, C, S) \not\rightarrow 0$. Then, there exists a subsequence such that $\sigma_t(\ell, C, S) \rightarrow \hat{\sigma} > 0$, which implies that the denominator of $Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C))$ converges to a strictly positive number along the subsequence. However, the numerator of $Pr_{\sigma_t}(y = 1 \mid x_S(\ell, C))$ converges to zero by Step 9, because $\sigma_t(h, C', S') \rightarrow 0$ if $x_S(h, C') = x_S(\ell, C)$ and hence C'^h . Therefore, $U_{\sigma_t}(\ell, C, S) \rightarrow 0$ along the subsequence, which contradicts $\sigma_t(\ell, C, S) \rightarrow \hat{\sigma} > 0$. ■

Step 15 If (ℓ, C, S) is such that $x_S(\ell, C) = x_S(h, N^h)$, then

$$\liminf_{t \rightarrow \infty} \sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S') = \liminf_{t \rightarrow \infty} \sum_{S'} \sigma_t(h, N^h, S') \equiv \underline{\sigma} > 0$$

Proof. The first equality follows because $\sigma_t(h, C', S') \rightarrow 0$ if C'^h by Step 9 and because $x_S(\ell, C) = x_S(h, N^h)$. The last inequality is strict because, if $\underline{\sigma} = 0$, there exists a subsequence such that $\sum_{C', S'} \sigma_t(h, C', S') \rightarrow 0$ and hence $\sigma_t(\ell, C, S) \rightarrow \hat{\sigma} > 0$ for some (ℓ, C, S) such that $x_S(\ell, C) = x_S(h, N^h)$. However, in this case there exists T such that for all $t \geq T$ in this subsequence the numerator of $Pr_{\sigma_t}(y = 1 | x_S(\ell, C))$ becomes arbitrarily small and hence $U_{\sigma_t}(\ell, C, S) < U^*$, which is inconsistent with $\hat{\sigma} > 0$. ■

Step 16 $\limsup_{t \rightarrow \infty} \bar{U}_{\sigma_t} \leq U^*$.

Proof. Suppose $\limsup_{t \rightarrow \infty} \bar{U}_{\sigma_t} = \bar{U} > U^*$. Let

$$\bar{P} = \left\{ (\ell, C, S) \mid \limsup_{t \rightarrow \infty} U_{\sigma_t}(\ell, C, S) = \bar{U} \right\},$$

which must be non-empty because the set of platforms is finite. Note that $(\ell, C, S) \in \bar{P}$ only if $x_S(\ell, C) = x_S(h, N^h)$. By finiteness of \bar{P} , there exists a common subsequence, T , and $\varepsilon > 0$ such that for all $t' \geq T$ in this subsequence $U_{\sigma_{t'}}(\ell, C, S) \geq U^* + \varepsilon$ for all $(\ell, C, S) \in \bar{P}$. By Step 10, there must exist a $t > T$ (not necessarily in the subsequence) such that $\bar{z}_t = (h, N^h, S)$ and hence $\bar{U}_{\sigma_t} = U^*$. Therefore, $U_{\sigma_t}(\ell, C, S) \leq U^*$ for all $(\ell, C, S) \in \bar{P}$. By Step 12, for all $(\ell, C, S) \in \bar{P}$,

$$\begin{aligned}
\frac{U_{\sigma_{t+1}}(\ell, C, S)}{U_{\sigma_t}(\ell, C, S)} &= \frac{\left(\frac{\frac{1}{t} + \sum_{C', S' | x_S(h, C') = x_S(\ell, C)}{\frac{1}{t} + \sum_{a', C', S' | x_S(a', C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')} \right)}{\left(\frac{\sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{a', C', S' | x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \right)} \\
&< \frac{\left(\frac{\frac{1}{t} + \sum_{C', S' | x_S(h, C') = x_S(\ell, C)}{\sum_{a', C', S' | x_S(a', C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')} \right)}{\left(\frac{\sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')}{\sum_{a', C', S' | x_S(a', C') = x_S(\ell, C)} \sigma_t(a', C', S')} \right)} \\
&= \frac{\frac{1}{t}}{\sum_{C', S' | x_S(h, C') = x_S(\ell, C)} \sigma_t(h, C', S')} + 1
\end{aligned}$$

which converges to 1 as $t \rightarrow \infty$ by Step 15. Therefore, for every $\delta > 0$, we can pick T large enough such that, for all $t \geq T$ such that $\bar{z}_t = (h, C, S)$,

$$\frac{U_{\sigma_{t+1}}(\ell, C, S)}{U_{\sigma_t}(\ell, C, S)} \leq 1 + \delta$$

for all $(\ell, C, S) \in \bar{P}$. Finally, this means that we can also pick T and $t \geq T$ so that $\bar{z}_t = (h, C, S)$ and $U_{\sigma_{t+1}}(\ell, C, S) < U^* + \varepsilon$ for all $(\ell, C, S) \in \bar{P}$. Therefore, $U_{\sigma_{t+k}}(\ell, C, S) < U^* + \varepsilon$ for all $(\ell, C, S) \in \bar{P}$ and all $k \geq 1$, because by Step 13 the payoff of (ℓ, C, S) is weakly decreasing when $U_{\sigma_t}(\ell, C, S) > U^*$. We, thus, reach a contradiction. ■

Steps 11 and 16 imply that $\lim_{t \rightarrow \infty} \bar{U}_{\sigma_t} = U^*$. Now, denote by Σ the set of limit points of σ_t .

Step 17 *All $\sigma \in \Sigma$ must induce the same joint distribution over (a, C) , and this distribution must coincide with the unique equilibrium distribution.*

Proof. Note that $U_\sigma(z)$ is continuous in σ for all z . The previous conclusion implies that, for every $\sigma \in \Sigma$ and every z , $U_\sigma(z) \leq U^*$, with equality for $z \in \text{Supp}(\sigma)$. Propositions 1 and 2 established that every σ that satisfies this property induces the same distribution over (a, C) . This completes the proof. ■

References

- ANDRE, P., I. HAALAND, C. ROTH, AND J. WOHLFART (2022): “Narratives about the Macroeconomy,” *Working Paper*.
- ASH, E., G. GAUTHIER, AND P. WIDMER (2021): “Text semantics capture political and economic narratives,” *arXiv preprint arXiv:2108.01720*.
- BÉNABOU, R., A. FALK, AND J. TIROLE (2018): “Narratives, imperatives, and moral reasoning,” Tech. rep., National Bureau of Economic Research.
- BURSTEIN, P. (2003): “The Impact of Public Opinion on Public Policy: A Review and an Agenda,” *Political Research Quarterly*, 56, 29–40.
- COWELL, R. G., S. L. LAURITZEN, A. P. DAVID, D. J. SPIEGELHALTER, V. NAIR, J. LAWLESS, AND M. JORDAN (1999): *Probabilistic Networks and Expert Systems*, Berlin, Heidelberg: Springer-Verlag, 1st ed.
- ELIAZ, K. AND R. SPIEGLER (2020): “A Model of Competing Narratives,” *American Economic Review*, 110, 3786–3816.
- ESPONDA, I. AND D. POUZO (2017): “Conditional Retrospective Voting in Large Elections,” *American Economic Journal: Microeconomics*, 9, 54–75.
- EYSTER, E. AND M. PICCIONE (2013): “An Approach to Asset Pricing under Incomplete and Diverse Perceptions,” *Econometrica*, 81, 1483–1506.
- HEALY, A. AND N. MALHOTRA (2013): “Retrospective Voting Reconsidered,” *Annual Review of Political Science*, 16, 285–306.
- IZZO, F., G. J. MARTIN, AND S. CALLANDER (2021): “Ideological Competition,” *SocArXiv. February*, 19.
- JEHIEL, P. (2005): “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory*, 123, 81–104.

- JONES, M. D. AND M. K. MCBETH (2010): “A Narrative Policy Framework: Clear Enough to Be Wrong?” *Policy Studies Journal*, 38, 329–353.
- JONES, M. D., M. K. MCBETH, AND E. A. SHANAHAN (2014): “Introducing the Narrative Policy Framework,” *Palgrave Macmillan US*, 1–25.
- LEVY, G. AND R. RAZIN (2021): “A maximum likelihood approach to combining forecasts,” *Theoretical Economics*, 16, 49–71.
- LEVY, G., R. RAZIN, AND A. YOUNG (2022): “Misspecified politics and the recurrence of populism,” *American Economic Review*, 112, 928–62.
- MAILATH, G. J. AND L. SAMUELSON (2020): “Learning under Diverse World Views: Model-Based Inference,” *American Economic Review*, 110, 1464–1501.
- PEARL, J. (2009): *Causality: Models, Reasoning and Inference*, USA: Cambridge University Press, 2nd ed.
- PLESCIA, C. AND S. KRITZINGER (2017): “Retrospective Voting and Party Support at Elections: Credit and Blame for Government and Opposition,” *Journal of Elections, Public Opinion and Parties*, 27, 156–171, PMID: 28515772.
- POLLETTA, F. (2008): “Storytelling in politics,” *Contexts*, 7, 26–31.
- ROTHSCHILD, M. AND J. STIGLITZ (1976): “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *Quarterly Journal of Economics*, 90, 629–649.
- SANDERS, K., M. J. M. HURTADO, AND J. ZORAGASTUA (2017): “Populism and Exclusionary Narratives: The ‘Other’ in Podemos’s 2014 European Union Election Campaign,” *European Journal of Communication*, 32, 552–567.

- SCHWARTZSTEIN, J. AND A. SUNDERAM (2021a): “Shared Models in Networks, Organizations, and Groups,” *mimeo*, *Harvard University*.
- (2021b): “Using Models to Persuade,” *American Economic Review*, 111, 276–323.
- SHANAHAN, E. A., M. K. MCBETH, AND P. L. HATHAWAY (2011): “Narrative policy framework: The influence of media policy narratives on public opinion,” *Politics & Policy*, 39, 373–400.
- SHILLER, R. J. (2017): “Narrative Economics,” *American Economic Review*, 107, 967–1004.
- SPIEGLER, R. (2013): “Placebo Reforms,” *American Economic Review*, 103, 1490–1506.
- (2016): “Bayesian Networks and Boundedly Rational Expectations,” *Quarterly Journal of Economics*, 131, 1243–1290.
- (2020): “Behavioral Implications of Causal Misperceptions,” *Annual Review of Economics*, 12, 81–106.
- WEAVER, R. K. (2013): “Policy leadership and the blame trap: Seven strategies for avoiding policy stalemate,” *Governance Studies*, *Brookings Institution*.