Productive Policy Competition and Asymmetric Extremism*

Alexander V. Hirsch†

June 20, 2025

Abstract

Viable policies must be developed by individuals and groups with the expertise and willingness to do so. We study a model of costly policy development in which competing developers may differ in their intrinsic ideological extremism and ability at crafting high quality policies. The unique equilibrium exhibits unequal participation, inefficiently unpredictable and extreme policies and outcomes, wasted effort, and an apparent advantage for extreme policies. While asymmetries between the developers always reduce observable competition, they can nevertheless benefit the decisionmaker. This contrasts starkly with the classic all-pay contest used to study lobbying and electoral competition, and is rooted in the fact the developers care about *which* policy they may lose to, rather than simply winning or losing. The model provides a novel rationale for why extreme actors may come to dominate policymaking that is rooted in the nature of policy development, and highlights the difficulty in assessing the normative implications of such dominance.

*I thank Ken Shotts, Michael Gibilisco, Adam Meirowitz, John Kastellec, Dan Kovenock, Joanna Huey, Avi Acharya, Alessandra Casella, Ron Siegel, Leeat Yariv, Federico Echenique, Salvatore Nunnari, Betsy Sinclair, Jon Eguia, Craig Volden, Scott Ashworth, and seminar audiences at Yale, USC, the University of Utah (Eccles School), and the 2018 SAET conference for helpful comments and advice.

†California Institute of Technology. Division of the Humanities and Social Sciences, MC 228-77, Pasadena, CA 91125. Email: avhirsch@hss.caltech.edu

Ideological competition is a key driver of policymaking in democracies; citizens, parties, and interest groups compete via elections to elect representatives who share their ideological interests, who in turn compete via the rules and procedures of government to enact public policies that reflect those interests. Correspondingly, political scientists have devoted considerable attention to these two processes, in order to better understand why some public policies become law and not others.

Public policy scholars, however, have long recognized the importance of an intermediate step in the policy process – how viable policy alternatives are actually *developed*. In his sweeping work, Kingdon (1984) describes policy development as a necessary precondition for change – "before a subject can attain a solid position on a decision agenda, a viable alternative [must be] available for decisionmakers to consider" (p142) – and recounts a Presidential staffer's perspective as follows:

> Just attending to all the technical details of putting together a real proposal takes a lot of time.... It's one thing to lay out a statement of principles or a general proposal, but it's quite another thing to staff out all the technical work that is required to actually put a real detailed proposal together (p132).

Given the costs and uncertain rewards to policy development, who will "invest the resources – time, energy, reputation, and sometimes money" (Kingdon (1984), p. 122) to do it? When and why?

In answering these questions, it is important to consider that political competition is rarely *balanced*. For one, ideological extremists on one side of an issue typically do not face equally extreme opposition on the other. In the United States evidence suggests that the Republican party has moved away from the median voter faster than the Democratic party (e.g. Grossmann and Hopkins (2016); McCarty (2015)). Within U.S. institutions that craft policy like legislatures and bureaucracies, asymmetric extremism is the norm, since elected or appointed decisionmakers are typically better aligned with one internal faction over others (Lewis (2008)). In addition, asymmetries in ability or resources are common. In issue areas dominated by interest group politics, the primary axis of conflict is often between poorly funded public interest groups and well-resourced business interests (Kerwin and Furlong, 2018; Yackee and Yackee, 2006)). At a higher level, several observers have argued that there is an expertise deficit in today's U.S. Republican party (Bartlett, 2017).

A now sizable literature has studied costly policy development by using models in which developers strategically craft policies that consist of both an *ideology* and a level of *policy-specific* quality, which must be generated at an up-front cost.[1] But the effects of asymmetries in developer extremism and ability remain poorly understood, as no previous work has studied asymmetric policy developers who are engaged in fully open competition where neither is privileged by the agenda.[2] In this paper we analyze the effects of such asymmetries by extending the competitive policy development model of Hirsch and Shotts (2015), in which developers simultaneously craft policies for consideration by a decisionmaker.[3] In so doing, we uncover a number of novel results about asymmetric policy competition, with surprising implications for interpreting real-world patterns of policy development.

Asymmetric Extremism A common throughline of previous work on policy development is that ideological extremism is not an unalloyed bad. Specifically, when opposing sides of a policy conflict become more polarized, their disagreement over "the shape of public policy" motivates them to make more beneficial quality investments (e.g. Hirsch and Shotts (2015)). It is unknown, however, whether such beneficial effects extend to asymmetric extremism; and indeed, it is reasonable to

[1]See for example Hirsch and Shotts (2015); Hitt, Volden and Wiseman (2017); Lax and Cameron (2007); Londregan (2000); Ting (2011) and Turner (2017).

[2]Lax and Cameron (2007) and (briefly) Hitt, Volden and Wiseman (2017) both analyze competitive policy development models where one developer is assumed to be a privileged "first mover."

[3]Hirsch and Shotts (2015) analytically characterize equilibrium in the symmetric case but only implicitly characterize the asymmetric case; they show equilibria must be mixed, with one developer always crafting a policy strictly benefitting the decisionmaker. Section III.A also shows that when one developer exactly shares the decisionmaker's ideal there is no competition in equilibrium.

suppose that they do not. This is because of a well-known property of contest models known as the *discouragement effect*, which states that as a contest participant becomes weaker he reduces his effort (anticipating a likely loss), thereby causing the stronger participant to strategically reduce his effort as well (Chowdhury, Esteve-Gonzalez and Mukherjee (2022); Crisman-Cox and Gibilisco (2024)). For example, in the *all-pay contest* – a workhorse model that has been previously applied to campaign and policy competition – asymmetries *always harm* a decisionmaker due to this discouragement effect (e.g. Hillman and Riley (1989); Meirowitz (2008)).

In our model, asymmetries in the developers' ideological extremism do indeed create asymmetries in strength that generate patterns of behavior resembling the all-pay contest. With asymmetric extremism there is a "stronger" developer who always crafts a new policy, and a "weaker" one who sometimes declines to do so. Intuition would suggest that it is the more moderate developer who will be stronger, by virtue of his better alignment with the decisionmaker. In fact, however, the reverse is true; it is the *more extreme* developer who is stronger, by virtue of his greater motivation to "affect the shape of public policy" (Hillman and Riley (1989)). This in turn discourages the "weaker" player (the more moderate developer) from participating at all, as in the discouragement effect.

However, our model also generates predictions about the quality and extremism of the policies that are proposed and enacted. How will the extremist craft policy as compared to the moderate? He will naturally craft a more extreme policy (in a first order stochastic sense). However, because he is more motivated, he will *also* craft a higher quality policy (again in a first order stochastic sense). In fact, his policy will be *so much* higher quality despite its greater extremism that it will also be *better for the decisionmaker* (again in a first order stochastic sense). In equilibrium the decisionmaker's choices will therefore appear to be *biased toward the extremist.* Our model thus provides a novel account of how extremists may come to dominate policymaking in a particular issue domain, rooted in the nature of productive policy development rather than capture or some other systemic failure.

Finally, what happens as an extremist developer becomes *yet more* extreme? Similar to the all-pay contest, his more-moderate competitor will become less likely to craft a new policy, expecting to be outmatched. He will also moderate his policy when he crafts one, expecting to fail at using quality to move the ideological outcome in his direction. Conversely, the increasingly extreme developer will craft an increasingly extreme policy – both because his preferences are becoming more extreme, and because he is becoming increasingly likely to succeed at exploiting quality to achieve ideological gains. Surprisingly however – and in stark contrast to the all-pay contest – this increasingly extreme developer will *also* craft an increasingly appealing policy for the decisionmaker *despite* its greater extremism. This happens because, in our model, an increasingly outmatched developer cannot simply opt out of policy development without eliciting an even more extreme policy from his competitor; any attempt at "unilateral disarmament" invites an ideological reaction that pulls him back into the contest over policy. In equilibrium, an increasingly-dominant extremist must instead "force" out a moderate competitor by crafting a policy with enough quality to be increasingly appealing to the decisionmaker, and therefore increasingly difficult to beat despite its greater extremism.

The consequence of the preceding is stark. While an increasingly-extreme developer crafts an increasingly extreme policy, *and also* further discourages the moderate developer from competing, the decisionmaker nevertheless becomes *increasingly better off* in equilibrium. Put more intuitively, unilateral extremism results in less (and in the limit no) observable competition, but an increasingly better-off decisionmaker. In addition to contrasting with the all-pay contest (where asymmetries always harm the decisionmaker), this result is also much stronger than the analogous result in Hirsch and Shotts (2015) that symmetric polarization benefits the decisionmaker, because there is no equally-extreme opposition to counterbalance an increasingly extreme developer.

Asymmetric Ability We also use our model to examine the effect of asymmetries in ability at crafting high quality policies. Interestingly, our analysis uncovers that asymmetric ability is *observationally equivalent* to asymmetric extremism, in the sense of yielding nearly identical patterns of competition. Specifically, a developer who is no more intrinsically extreme than his competitor, but who is more skilled at producing high quality policies, will be more active, and develop a policy that is more extreme but also higher quality and better for the decisionmaker. Conversely, a developer who is no less ideologically extreme than his competitor, but who is less capable of producing high quality policies, will be less active, and develop a policy that is more moderate but also lower quality and worse for the decisionmaker. Finally, as a more-expert developer becomes increasingly expert, his policy becomes increasingly extreme but also higher quality and better for the decisionmaker, while his competitor increasingly moderates his policy and becomes increasingly unlikely to develop one. The competitor becomes increasingly worse off, and the decisionmaker becomes increasingly better off despite the growing imbalance in participation and extremism of the expert's policy.

Our results thus illustrate how asymmetric extremism may be not only a "cause" (when it describes the underlying preferences of political actors) but also a "consequence" (when it describes the observed behavior of political actors). This has far reaching implications for measures of elite policy preferences derived from their observable behaviors like votes and bill sponsorship (see Clinton (2012) for a review) – simply put, a political actor's underlying extremism cannot be straightforwardly extracted from his observable behavior. Our results also yield an important lesson for the design of effective political institutions; that there is really no way to allocate ostensibly "nonpartisan" resources to develop policy (like expert staff and budgets) without having "ideological" effects (Reynolds (2020)). Specifically, when a developer becomes endowed with greater ability in our model, he simply becomes better at generating the "public good" of policy quality. But because quality cannot be "transferred" between policies, this capacity nevertheless has ideological consequences, benefitting himself (with more appealing policy outcomes) and the decisionmaker (with higher quality policies) at the expense of an ideological competitor.

Robustness and Intuition Finally, to both explore robustness and develop the intuition for our main results, we conclude by considering a series of model variants. In the first, the developers cannot choose the extremism of their policy; they must instead each craft a policy at a fixed exogenous location. We show that with this change the benefits of asymmetric extremism and ability vanish; this illustrates the importance of the developers' ability to strategically choose the extremism of their policy for our results. In the second, the developers receive a fixed payoff when losing, so that they only care about policy when they win. We again show that the benefits of asymmetric extremism and ability vanish, illustrating the importance of full policy motivations for our results. After analyzing these two variants we return to the main model, and more carefully flesh out why it is only the *combination* of the preceding features that yields a potential benefit from asymmetries. Finally, we consider a third variant where players' preferences over policy are linear and the costs of generating quality are quadratic, and verify that our main results about asymmetries are robust to these alternative functional forms.

Related Literature

The classical approach to studying policy expertise supposes that a policy *outcome* results from the sum of a policy *choice* and an unknown *state of the world*. This approach has been widely used to study many institutional environments; its central tension is that privately-informed experts worry their that expertise will be exploited to implement a policy outcome contrary to their interests. Our model is part of a growing literature that captures an alternative conceptualization of expertise, in which experts make *policy-specific* investments that they use to achieve a particular ideological goal (e.g. Callander (2008); Hirsch and Shotts (2012); Hitt, Volden and Wiseman (2017); Lax and Cameron (2007); Londregan (2000); Ting (2011); Turner (2017)). Rather than fear that their expertise will be exploited, experts try to exploit costly policy-specific quality investments to persuade decision makers to accept policies that promote their own ideological interests.[4] Most closely related

[4]The framework pioneered by Callander (2008) (in which the mapping between policies and outcomes is the realized path of a Brownian motion) is a microfoundation that can capture both approaches, depending on the variance of the Brownian motion and corresponding "complexity" of the mapping. "Simple" mappings (with a zero-variance Brownian motion) correspond to the classi-

are works by Lax and Cameron (2007) and Hitt, Volden and Wiseman (2017), who study competitive policy development when the developers craft policies in a predetermined order. These models are better suited to studying institutions with structured agenda procedures like the US Supreme Court or House of Representatives, and yield very different patterns of competition that more closely resemble entry deterrence models of market competition.

The intended empirical domain of our model is policymaking settings that are "healthy" in two particular senses – (1) there exists *some* common ground between competing actors in the form of policy attributes that they all value, and (2) they are both able and willing to channel their ideological disagreements into productive investments in these attributes. A now-sizable literature applies such models to a range of institutional settings, thereby implicitly or explicitly supposing that they (at least sometimes) exhibit these features. An early example is Londregan (2000), who posited that competing branches of the Chilean government "weigh policy alternatives in terms of ideology, about which they disagree, and on the basis of shared public policy values, such as the desire for efficiency." Subsequent applications include intra and inter-court bargaining (Clark and Carrubba (2012); Lax and Cameron (2007)) (with opinion attributes like "persuasiveness, clarity, and craftsmanship" valued by all judges); Congressional delegation to the bureaucracy (Huber and McCarty (2004); Ting (2011)) (with "effective implementation" in the sense of "whether regulations are enforced, revenues are collected, benefits are distributed, and programs are completed" valued by legislators and bureaucrats); judicial oversight of the bureaucracy (Turner (2017)) (with "policy precision" valued by both risk-averse judges and bureaucrats); and legislatures (Hirsch and Shotts

cal approach. Maximally complex mappings (with an infinite-variance Brownian motion) resemble policy-specific quality in several ways. Specifically, policy-specific quality is effectively a reduced form of Callander's maximally complex variant if the quality investment is assumed to be binary, and a policy's ideology is unverifiable. An alternative microfoundation for maximally complex policy domains is offered in seminal work by Aghion and Tirole (1997) on "real authority."

(2012); Hitt, Volden and Wiseman (2017)) (with the "costs and benefits [of policies] across an array of societally valued criteria" being valued similarly by all legislators).

Our model is also related to a literature that studies contests in the context of lobbying and elections. Foundational work by Tullock (1980) modelled lobbying as a process by which competing groups exert wasteful effort to increase their chance of securing "politically-contestable rents." Important follow-on work by Hillman and Riley (1989) studied political contests in which these rents always fall to the group exerting the most effort, and groups could value control of policy differently.[5] This model is now known as the *all-pay contest* due to its close relationship to the all-pay auction format, in which a prize is awarded to the highest bidder but all bidders pay their bids (Baye, Kovenock and de Vries (1996); Siegel (2009)). Our model is closely related to the all-pay contest in two ways; the cost of generating quality is paid up-front by a developer regardless of whether his policy is ultimately implemented, and a policy strictly preferred by the decisionmaker is implemented with certainty rather than probabilistically. However, a key distinction is that our developers' payoffs from winning and losing are *endogenous* to the policies that they develop, as befits a setting where the participants are ideologically motivated rather than "rent-seeking."[6]

Finally, our model is related to several Downsian election models in which candidates, much like our policy developers, choose both ideological platforms and costly valence-generating spending (e.g. Ashworth and Bueno de Mesquita (2009); Balart, Casas and Troumpounis (2022); Hirsch (2023); Serra (2010); Wiseman (2006)). With the exception of Hirsch (2023), however, these works study candidates who sequentially choose platforms and then spending, spending and then platforms, or move in a pre-determined order. Such models in turn feature strategic forces that are related to an

[5] See Meirowitz (2008) for an application to campaign competition and the incumbency advantage.

[6] In the terminology of Baye, Kovenock and de Vries (2012), our contest has a "second order rank-order spillover" – the strategy of the "first ranked" player (the winner) directly affects the payoff of the "second ranked" player (the loser) because it is the policy he must actually live under.

earlier generation of Downsian models that study candidates with exogenous valence advantages (e.g. Groseclose (2001), Aragones and Palfrey (2002)). Indeed, our developers' attempts to parlay quality investments into support for more extreme policies is akin to the force underlying the influential "marginality hypothesis" in that literature (Groseclose (2001) p. 863), whereby an incumbent will "parlay the advantage into a policy position that is closer to her ideal point."

The Model

Two developers, labelled -1 (left) and 1 (right), craft competing policies for consideration by a decisionmaker (DM), labelled player 0. A policy (y, q) consists of an *ideology* $y \in \mathbb{R}$ and a level of *quality* $q \in [0, \infty) = \mathbb{R}^+$. All players are purely policy-motivated, in the sense that their final policy payoffs depend only on the ideology and quality of the final policy. The utility of player i for a policy (y, q) is $U_i(y, q) = \lambda q - (y - x_i)^2$. The parameter x_i is player i's ideological ideal point, the decisionmaker is located at 0, the left developer is distance $|x_{-1}|$ to her left, and the right developer is distance $|x_1|$ to her right. A developer's distance $|x_i|$ from the decisionmaker reflects his ideological extremism. A policy's quality q is a public good that all players value at weight λ; higher λ means that the players collectively place greater weight on policy quality vs. ideology.

The game proceeds in two stages. In the first, the developers simultaneously select the ideology and quality of their respective policies (y_i, q_i). Endowing a policy with quality q_i costs $c_i(q_i) = a_i q_i$ up front, which reflects the initial time and energy needed to improve the policy's quality. The parameter a_i is developer i's marginal cost of increasing quality, and reflects his *ability* at doing so. $\alpha_i = \frac{a_i}{\lambda}$ denotes the ratio of a developer's marginal cost of quality to its marginal benefit, and is the *weighted marginal cost* of quality once its intrinsic value is taken into account. We assume that this parameter is greater than 1 for both developers, implying that neither would invest in quality for its own sake. In the second stage the DM chooses a final policy to implement. This may be one of the two policies created by the developers, or any other policy from a set of outside options $\mathbb{O}$ that only includes policies weakly worse for the DM than her ideal point with 0-quality. (The set of outside options may also be empty, meaning that the DM *must* choose one of the developers' policies.[7])

Preliminary Analysis

The Monopolist's Problem It is helpful to first consider the model with only one developer, i.e. a "monopolist" (see also Hitt, Volden and Wiseman (2017) and Hirsch and Shotts (2018)). W.l.o.g. suppose he is the right developer ($i = 1$). The monopolist's problem is depicted in the left panel of Figure 1; ideology is on the x-axis and quality is on the y-axis. The shaded region depicts the set of policies that the decisionmaker would be willing to implement in lieu of (y_0, q_0), which denotes the best policy she can implement without the developer's help (i.e., her "outside option"). To clarify incentives we will temporarily allow this policy to be strictly better than $(0, 0)$ (the decisionmaker's ideal point with 0-quality), but recall that in the main competitive model we have assumed the decisionmaker's best outside option to be no better than $(0, 0)$.[8]

The developer must choose both *whether* to craft a new policy that the decisionmaker is willing to implement (in the shaded region), and if so exactly *which* policy (y, q) to develop. This problem

[7]The original Hirsch and Shotts (2015) treatment assumes that this set is exactly equal to the decisionmaker's ideal point with zero quality. In the Appendix we show that the exact set is irrelevant for the structure of equilibrium, and may even be empty.

[8]This treatment of the monopoly model further assumes that the decisionmaker's best outside option (y_0, q_0) is *no worse* than $(0, 0)$ (as in Hirsch and Shotts (2018)); this assumption proxies for an "open rule" in which the decisionmaker can choose any 0-quality policy in lieu of the monopolist's policy. The monopoly analysis herein is thus effectively a generalization of the "open rule" model studied in Hitt, Volden and Wiseman (2017) Prop. 3 that allows for positive-quality status quos. If the rule were instead "closed" (so that the decisionmaker could not unilaterally implement all 0-quality policies), then with a monopoly developer there would be additional cases to consider in which the monopolist "crafts" a 0-quality policy that the decisionmaker cannot access on her own in lieu of developing a new one; see Hitt, Volden and Wiseman (2017) Prop. 2.

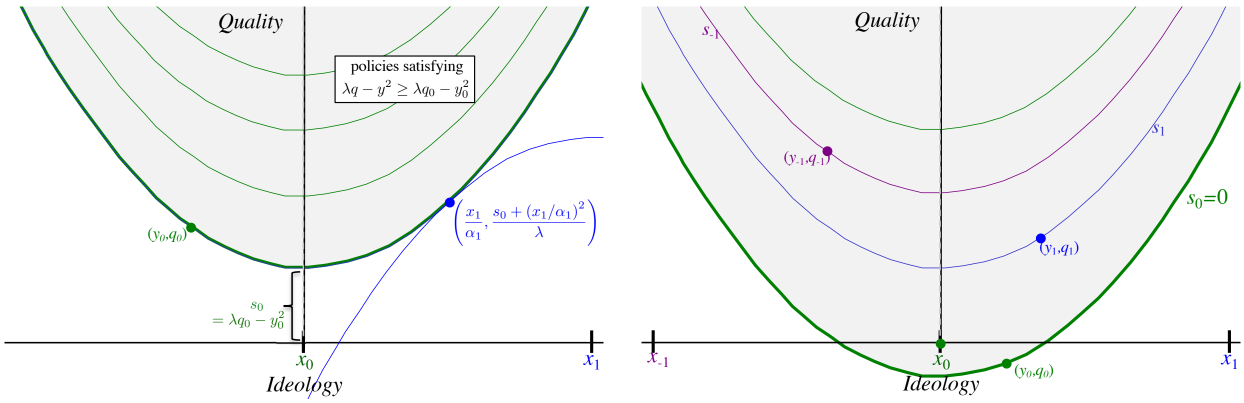

Figure 1: *The Developer's Problem.* The left panel depicts the problem faced by a "monopolist" when the decisionmaker's outside option is (y_0, q_0); the green curves are the decisionmaker's indifference curves, the shaded region depicts the policies the decisionmaker weakly prefers to her outside option, and the blue curve depicts the policies that the developer is indifferent over crafting conditional on acceptance. The right panel depicts the competitive problem, where the decisionmaker chooses her favorite between the policies crafted by the two developers or her best outside option.

can be understood using the inequality

$$\underbrace{\arg\max}_{\{(y,q):\lambda q - y^2 \geq \lambda q_0 - y_0^2\}} \left\{ \left(\lambda q - (y - x_1)^2 \right) - a_1 q \geq \lambda q_0 - (y_0 - x_1)^2 \right\} \tag{1}$$

The policy (y, q) that maximizes the left hand side is optimal *if* the developer chooses to be active (i.e., invest effort in developing a new policy), and depends on both the developer's ideology x_1 and ability a_1. *Whether* the developer will be active in turn depends on whether the left hand side (his utility from the developing the optimal policy) exceeds the right hand side (his utility from the decisionmaker's "outside option" (y_0, q_0)). Importantly, the outside option appears on both sides of the inequality because it functions as both a *constraint on* and a *motive for* policy development. It is a constraint because the developer must craft something at least as good as (y_0, q_0) for it to be adopted; the higher is the decisionmaker's indifference curve (in green) passing through (y_0, q_0), the harder it is to "beat." It is a motive because the developer must live with (y_0, q_0) if he doesn't develop something else; while all outside options on the same curve are equally difficult to beat, those further to the left are worse for the developer, and so more strongly incentivize policy development.

To solve this problem and aid in the subsequent analysis, we reparameterize policies (y, q) in terms of their ideology y and the *utility they give the decisionmaker* – we henceforth call this a policy's *score*, and denote it s. Now observe that $s = \lambda q - y^2$, implying that a score-s policy with ideology y must have quality $q = \frac{s+y^2}{\lambda}$. Next recall that $\alpha_i = \frac{a_i}{\lambda}$ is the marginal cost of generating quality weighted by the marginal benefit; then substituting into (1) yields the revised problem

$$\underset{\{(s,y):s \geq s_0\}}{\arg\max} \left\{ \underbrace{-(\alpha_1 - 1)\, s}_{\text{score effect}} + \underbrace{2x_1 y - x_1^2 - \alpha_1 y^2}_{\text{ideology effect}} \right\} \geq s_0 + 2y_0 x_1 - x_1^2, \tag{2}$$

where s and s_0 are the score of the developer's new policy and the decisionmaker's outside option, respectively. Now it is easy to see that *if* the developer crafts a new policy, it will be *just good enough* for the decisionmaker to be willing to choose it over her outside option, and thus leave the decisionmaker indifferent with her outside option ($s^* = s_0$). The developer will then optimally choose the ideology y^* of his policy by trading off the *linear ideological benefit* $2x_1 y$ of moving the policy outcome in his desired direction (along the score curve s_0) against the *quadratic up-front cost* $\alpha_1 y^2$ of compensating the decisionmaker for a more extreme policy with additional quality. Differentiating with respect to y and setting equal to 0 yields an optimal ideology of $y^* = \frac{x_1}{\alpha_1}$, which is a weighted average (by $\frac{1}{\alpha_1}$) of the developer's and decisionmaker's ideal ideologies.

The Competitive Problem When the developer faces a competitor, the decisionmaker's best outside option may no longer be an exogenous policy, but instead the policy (s_{-1}, y_{-1}) crafted by the competitor. The setup of the competitive problem is depicted in the right panel of Figure 1. A developer's policy choice is a two-dimensional "bid" (s_i, y_i) consisting of a policy's score s_i and ideology y_i. After seeing the two policies, the decisionmaker chooses the one with the highest score (i.e., on the highest indifference curve in Figure 1) or her best outside option. The developers thus compete in a contest over policy-development, where a policy's likelihood of winning is determined by its score, and its ideology affects both its up-front cost to craft and the benefit of winning.

Now recall that a monopolist will choose to either develop *no* policy, or to develop a policy *no better* than the decisionmaker's outside option. Applying this insight to the competitive model straightforwardly yields that there is no pure strategy equilibrium. If each developer i expected his competitor to craft a *specific* policy (s_{-i}, y_{-i}), then each would treat the other's policy like the outside option in the monopoly model, and equilibrium would require that both developers be crafting policies with the exact same score. But if both developers were crafting policies with the exact same score then they would "tie," and one would have a strict incentive to break it; either by developing a slightly higher-score policy, or by dropping out of policy development.

Deriving Equilibrium Strategies We last sketch how to solve for the competitive equilibrium (see Appendix A for details).[9] This material is not required to read and interpret our results, but is necessary background for the later discussion on robustness.

It is first necessary to identify some basic structure on the form of equilibrium strategies. In the Appendix we show that each developer's mixed strategy can be written as pair of univariate functions: (1) a *cumulative distribution function* $F_i(s)$ describing the probability that developer i crafts a policy with score less than or equal to s, and (2) an *ideology function* $y_i(s)$ describing the exact ideology that developer i targets when crafting a policy with score s.

We next show that any pair of equilibrium score CDFs $(F_{-1}(s), F_1(s))$ must satisfy some familiar properties from the all-pay contest; that they be *continuous* and *strictly increasing* over a common interval $[0, \bar{s}]$, with $F_k(0) > 0$ for at most one k. That is, each developer must *randomize smoothly* over crafting policies with scores in a common interval, at most one developer k can be inactive with strictly positive probability, and the other developer $-k$ must always craft a policy strictly better for the decisionmaker than $(0, 0)$ with probability 1. Intuitively, these properties follow from two features that our model shares with the all-pay contest: that (i) choosing a higher-score strategy must yield a higher probability of victory (since it is costly), and (ii) no developer may choose a costly strategy that could result in a "tie" (since he could endow that policy with just a little more quality to break the tie). An additional property of our equilibrium is that it is "as if" the decisionmaker's outside

[9] A similar but less general treatment can be found in Hirsch and Shotts (2015) Sections I.A-I.C.

option is exactly $(0,0)$ (even if in reality it is strictly worse) because this is the most competitive "free" policy to craft, and the developers wish to move ideology in opposite directions.

With these properties established, equilibrium then requires that every policy (s, y) in the support of each developer i's strategy (that is, satisfying both $s \in [0, \bar{s}]$ and $y = y_i(s)$) maximize the expression

$$\underbrace{F_{-i}(s)}_{\text{Pr win}}\underbrace{\left(\left(s+y^2\right)-\left(y-x_i\right)^2\right)}_{\text{utility from winning}}-\underbrace{\alpha_i\left(s+y^2\right)}_{\text{effort cost}}+\underbrace{\int_{s_i}^{\bar{s}}\left(\left(s_{-i}+\left[y_{-i}(s)\right]^2\right)-\left(y_{-i}(s)-x_i\right)^2\right) f_{-i}\left(s_{-i}\right) d s_{-i}}_{\text{expected utility when losing}},$$

so that it is optimal to randomize among them. The preceding objective function simplifies to:

$$-\left(\alpha_i-F_{-i}(s)\right) s+F_{-i}(s) \cdot\left(2 x_i y-x_i^2\right)-\alpha_i y^2+\int_s^{\bar{s}}\left(s_{-i}+2 x_i y_{-i}(s)-x_i^2\right) f_{-i}\left(s_{-i}\right) d s_{-i}, \quad (3)$$

which exhibits several important properties.

First, consider a developer i's incentive to craft a more ideologically-extreme policy *y holding the policy's score s fixed.* This problem is similar to the monopoly problem, in that there is a quadratic up-front cost $\alpha_i y^2$ of compensating the decisionmaker for a more ideologically extreme policy with additional up-front quality. However, there is now a crucial difference: the linear ideological benefit $2x_i y$ of crafting a more extreme policy becomes *weighted by* the (endogenous) probability $F_{-i}(s)$ that i's opponent crafts a lower-score policy (since otherwise i's policy won't be chosen). This yields a revised optimal ideology $y_i(s) = F_{-i}(s)\frac{x_i}{\alpha_i}$, where $F_{-i}(s)$ is endogenously determined. Crucially, a developer in the competitive model must therefore *behave more ideologically aggressively* at a score s when he believes such a policy is *more likely to win* (higher $F_i(s)$). Correspondingly, should one developer participate less in the contest, the other will strategically respond by becoming more ideologically aggressive, thereby raising the ideological stakes for *both* developers.

Second, consider a developer's incentive to craft a marginally higher score policy (assuming both only target optimal ideologies). As in the monopoly model, there is a marginal cost of doing so equal to $\alpha_i - F_i(s)$ (since i will enjoy the intrinsic benefit of his policy's quality when it wins). In contrast to the monopoly model, however, increasing score is also *strategically productive* because it increases the chance that i's policy wins. Specifically, should i's opponent craft an exactly score-s policy (which he will do with "probability" $f_{-i}(s)$), a marginally higher score will change the ideological outcome from $y_{-i}(s)$ (since i would have lost in this event) to $y_i(s)$ (since i will now win in this event). From this we see a second crucial property – that the *stakes* of the contest at a given score s, and thus each developer's willingness to target higher scores, depends on the *endogenous* policies that they are expected to craft. And these in turn depend on the developers' endogenous score CDFs, as previously described. This complex entanglement between the optimal choice of ideology and score yields the following implicit characterization of equilibrium in the form of a system of differential equations and boundary conditions (a more complete statement is in Appendix A).

Proposition A.1 *In any equilibrium,* $y_i(s) = \frac{x_i}{\alpha_i} F_{-i}(s)$, $F_k(0) > 0$ *for at most one developer* $k \in \{-1, 1\}$, $F_L(\bar{s}) = F_R(\bar{s}) = 1$, *and* $f_{-i}(s) = \frac{\alpha_i - F_{-i}(s)}{2x_i(y_i(s) - y_{-i}(s))}$ $\forall s \in [0, \bar{s}]$ and $i \in \{-1, 1\}$.

The key analytic difficulty in solving the system in Proposition A.1 is that it is *coupled*, since each developer i's objective function contains both *his opponent's* score CDF $F_{-i}(s)$ (which determines the probability that a score-s policy will be chosen) and *his own* score CDF $F_i(s)$ (which determines the ideology of the marginal score-s policy $y_{-i}(s) = \frac{x_{-i}}{\alpha_{-i}} F_i(s)$ that he will defeat if he increases his score). This mutual dependence in the system of differential equations is not present in the asymmetric all-pay contest, but arises naturally from our key assumptions that the developers strategically choose ideology, and also care about policy even when they lose. Our first main contribution relative to Hirsch and Shotts (2015) is to derive a closed-form solution to this coupled system in the general case (see Appendix B for details). This permits constructive proofs of equilibrium existence and uniqueness, and a comparative statics analysis of equilibrium policies, outcomes, and utilities.

Equilibrium Characterization

In equilibrium each developer mixes smoothly over a continuum of policies with ideologies between their own ideal point and the decisionmaker's. A developer's equilibrium mixed strategy can be written as a pair of functions $(q_i(\delta), G_i(\delta))$ that describe: (1) the level of quality $q_i(\delta)$ that developer i produces when he crafts a policy whose ideology is *distance* δ from the decisionmaker, and (2) a smooth *cumulative distribution function* $G_i(\delta)$ that describes the probability developer i will craft a policy with ideology closer to the decisionmaker than δ. The equilibrium values of these functions are as follows (see Appendix B for a detailed derivation).

Proposition 1. *For each developer $i \in \{-1, 1\}$, define the strictly* ***decreasing*** *function*

$$\epsilon_i(p) = \int_p^1 \frac{|x_i|}{\alpha_i - q} dq = |x_i| \log\left(\frac{\alpha_i - p}{\alpha_i - 1}\right).$$

Let $p_i(\epsilon) = \alpha_i - (\alpha_i - 1)\, e^{\frac{\epsilon}{|x_i|}}$ denote the well-defined inverse of $\epsilon_i(p)$, and let k denote the developer with the smallest value of $\epsilon_i(0)$.

- *When developer i crafts a policy whose ideology is distance δ from the decisionmaker, he targets ideology $i\delta$ and attaches quality $q_i(\delta) = \frac{\delta^2 + s_i(\delta)}{\lambda}$, where*

$$s_i(\delta) = 2\int_{\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)}^{\epsilon_k(0)} \left(\sum_{j\in\{-1,1\}} \frac{|x_j|}{\alpha_j} p_j(\epsilon)\right) d\epsilon$$

- *The probability that developer i crafts a policy closer to the decisionmaker than δ is*

$$G_i(\delta) = p_{-i}\left(\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right) = \alpha_{-i} - (\alpha_{-i} - 1)\left(\frac{x_i - i\delta}{x_i - x_i/\alpha_i}\right)^{\left|\frac{x_i}{x_{-i}}\right|},$$

- *Developer $-k$ is always active, and developer k is inactive with probability $p_{-k}(\epsilon_k(0)) = \alpha_{-k} - (\alpha_{-k} - 1)\left(\frac{\alpha_k}{\alpha_k - 1}\right)^{\left|\frac{x_k}{x_{-k}}\right|}$*

Although the equilibrium strategies are straightforward to express, they are somewhat hard to interpret from the equations alone. We therefore describe the structure of equilibrium and key properties with the aid of an example.[10]

[10]Note that the following properties are also implied by the implicit characterization of asymmetric equilibria in Hirsch and Shotts (2015).

Figure 2 depicts equilibrium strategies and outcomes, and compares when the developers are symmetrically extreme and capable (the top panels) to when they are equally capable but the right developer is more extreme (the bottom panels). The left panels depict the policies (both ideology and quality) that the left (purple) and right (blue) developers randomize over. The decisionmaker's indifference curves are in gray. The right panels depict the probability distributions (PDFs) over the ideology of the left (purple) and right (blue) developer's policies. *When* a developer chooses to craft a policy, its ideology is distributed over an interval with the depicted density. In the asymmetric case the left developer also sometimes chooses to craft *no* new policy. This is depicted in the bottom-left panel by the purple dot at the origin (the DM's ideal point with zero quality), and the probability this occurs is illustrated in the bottom-right panel by the height of the thick purple segment. The density over the ideology of the final policy chosen by the DM is depicted by the gray dashed lines.

When the developers are asymmetric, the unique equilibrium generically exhibits asymmetric participation in the policy process. One of the two developers (in the example the right developer) is always *active*, in the sense of developing a new policy with strictly positive quality. Moreover, any such policy is strictly better for the decisionmaker than $(0, 0)$ (her ideal point with 0-quality) – in Figure 2 all positive-quality policies are strictly above the DM's indifference curve through the origin. Competition thus strictly benefits the DM with probability 1 regardless of the developers' characteristics (and even though any quality invested in the losing policy is lost). The other developer (in the example the left developer, and more generally developer k defined in Proposition 1) is only *sometimes* active; with strictly positive probability he develops nothing.

Whenever a developer chooses to craft a new policy, its ideology strictly diverges from the DM's ideal (in Figure 2, all positive-quality policies have divergent ideologies). Active participation is thus always accompanied by an attempt to extract "ideological rents" in the form of a policy closer to one's ideal than the DM's ideal. These efforts result in both developers being harmed by the presence of competition relative to acting as "monopolists" – both invest in enough quality to compensate the

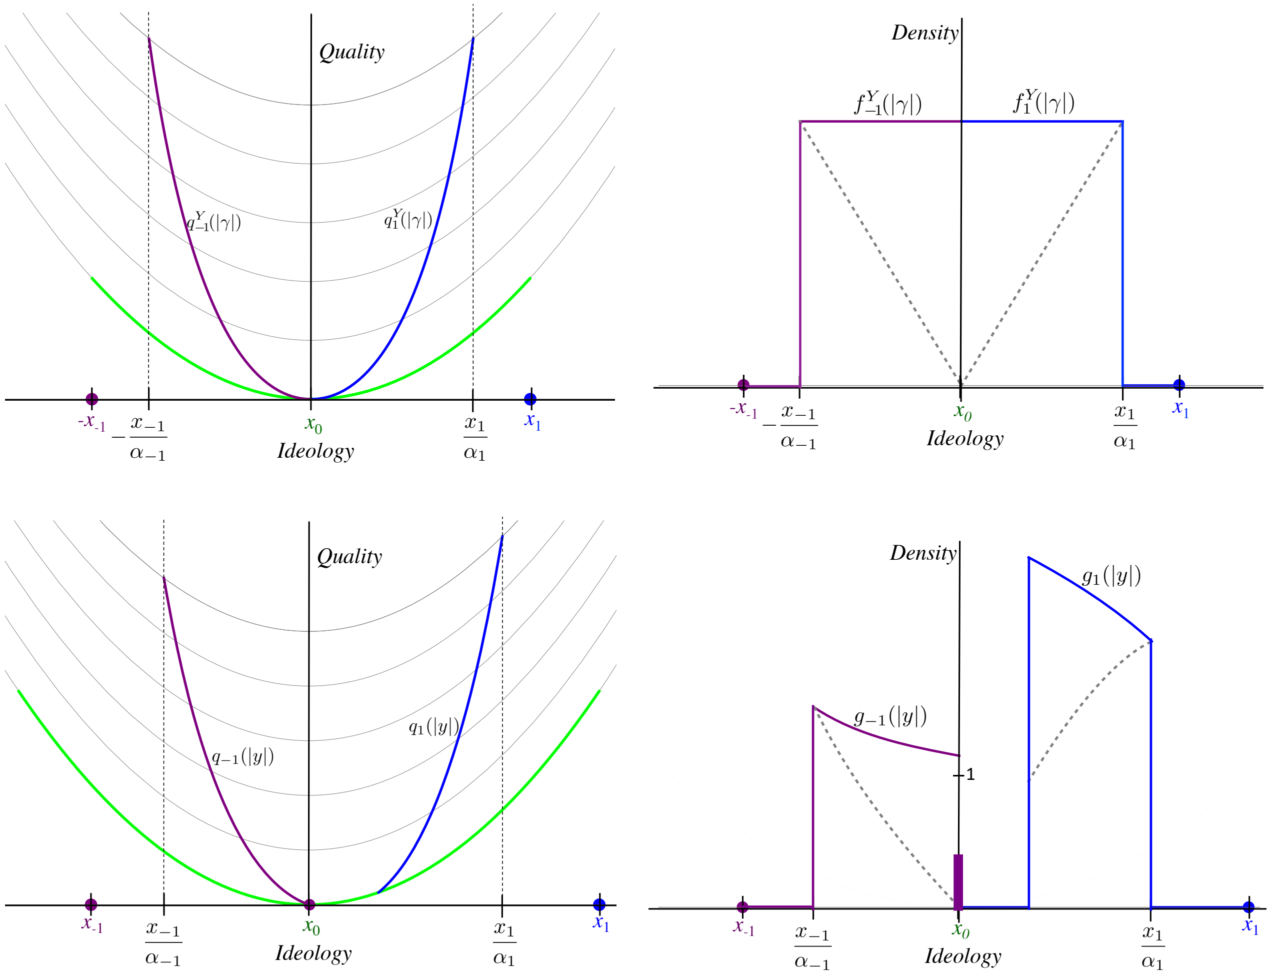

Figure 2: *Equilibrium with Symmetric vs. Asymmetric Extremism*

DM for her ideological losses from selecting their policies, but not enough to compensate each other.

The equilibrium also exhibits a variety of inefficiencies. Except in the special case of symmetric developers, the expected ideology of the final policy generically differs from both the DM's ideal, as well as the ideology that maximizes aggregate utility. The ideology of the final policy is also uncertain ex-ante, which harms all participants in the policy process due to risk aversion (its distribution is depicted by the dashed grey lines in the right panels of Figure 2). Finally, because the developers must make their quality investments before they know which policy will be chosen, all of the benefits of the effort invested in the losing policy are wasted.

The Politics of Asymmetric Extremism

We next turn to the politics of *asymmetric extremism*, by studying equilibrium when the developers are equally capable ($\alpha_1 = \alpha_{-1}$) but one is more ideologically extreme ($|x_i| \neq |x_{-i}|$). For convenience we call the more extreme developer "the extremist" and the other "the moderate."

Proposition 2. *If the developers are equally skilled but i is more extreme ($|x_i| > |x_{-i}|, \alpha_i = \alpha_{-i}$),*

- *the extremist always develops a new policy, while the moderate only sometimes does*
- *the extremist's policy is first-order stochastically more extreme than the moderate's policy, but also first-order stochastically higher quality and better for the decisionmaker*
- *the extremist's policy is strictly more likely to be chosen*

Recall that equilibrium in an example with asymmetric extremism is depicted in Figure 2.

Proposition 2 first characterizes the form of asymmetric participation that arises with asymmetric extremism – it is the *extremist* who always develops a new policy, while the moderate only sometimes does so (despite being better aligned with the decisionmaker). The extremist also develops a first-order stochastically more extreme policy than the moderate. Surprisingly, however, his policy actually performs better than the moderate's, because it is so much higher quality so as to be first-order stochastically better for the decisionmaker *despite* its greater extremism. What explains the extremist's dominance of the policy process despite his more ideologically-extreme policy? For one, the extremist is advantaged because his extremism makes him *more motivated* to invest in enough quality to compensate the decisionmaker for a more extreme policy. In addition, the extremist is not disadvantaged by his extremism because he is still free to strategically moderate his policy if doing so is necessary to ensure victory over his opponent's undesirable policy.

We next examine what happens to the developers' policies when one *becomes* more extreme.

Proposition 3. *If developer i becomes intrinsically more extreme (higher $|x_i|$), his* ***own strategy*** *and his* ***opponent's strategy*** *are affected in the following ways:*

(Own strategy)

- *if he previously did not always develop a policy, he becomes strictly more likely to do so*
- *his policy becomes first-order stochastically more extreme*
- *his policy becomes first-order stochastically higher quality, better for the decisionmaker, and strictly more likely to be chosen*

(Opponent's strategy)

- *if he did not always develop a policy, he becomes strictly less likely to do so*
- *his policy become first-order stochastically more moderate*
- *there is no unambiguous first-order stochastic change to his policy's quality or appeal to the decisionmaker, but his policy becomes strictly less likely to be chosen*

Although the above comparative statics apply to any configuration of preferences and costs, they are easiest to discuss in the special case of developers who are equally capable ($\alpha_1 = \alpha_{-1}$).

When a developer i's underlying preferences become unilaterally more extreme, the effects on his own participation and his opponent's participation are quite natural; if he is initially the moderate he becomes strictly more likely to be active, and if he is initially the extremist his competitor becomes strictly less likely to be active. Both results follow from a property that our model shares with the all-pay contest – that *balance* in the participants' motivation and abilities maximizes the likelihood that both players participate.

When developer i becomes more extreme, his policy also naturally becomes first-order stochastically more extreme; this is both because he has more extreme intrinsic preferences, and because his greater likelihood of winning the contest makes him more ideologically aggressive. Interestingly, his *competitor* also *moderates* his policy (first-order stochastically) despite no change in his underlying preferences or abilities. This moderation is driven *not* by the competitor's desire to make his policy more competitive when facing a more extreme competitor, but rather his acceptance of the fact that he is less likely to win and move the ideological outcome in his direction.

Finally, when developer i becomes more extreme, his policy becomes *first order stochastically better for the decisionmaker* despite being first-order stochastically more extreme.[11] This effect contrasts starkly with the asymmetric all-pay contest (where the strategy of the stronger player does not change as he becomes even stronger), and derives from a combination of the developers' strategic policy choice and their fear of losing to an ideologically-distant policy. Specifically, in both the classic all-pay contest and in our model, a weaker player (here the moderate) must become increasingly discouraged from participating when facing an increasingly strong competitor (here the extremist); the simple reason is that he is increasingly outmatched. In the all pay contest, this "discouragement" of the weaker player has no impact on the behavior of the stronger one.[12] In our model, in contrast, it causes the stronger player to behave more ideologically aggressively; this raises the ideological "stakes" of the contest, reinvigorating the weaker player's desire to participate. To prevent this reinvigoration, an increasingly extreme developer must therefore *also* crafting an increasingly appealing policy for the decisionmaker that is more difficult to beat.

With this unusual effect in mind, we last examine the consequences of a developer i becoming more extreme for the other players' welfare. We begin with his competitor.

Proposition 4. *If developer i becomes intrinsically more ideologically extreme (higher $|x_i|$), the equilibrium utility of his competitor $-i$ decreases*

[11] There is no unambiguous first-order stochastic change in the appeal of his opponent's policy to the decisionmaker, meaning that it never becomes unambiguously better or worse overall.

[12] This is because the stronger player's behavior is entirely driven by the need to encourage the weaker player's participation, and neither his stakes nor abilities have changed.

The effect of unilateral extremism on a competitor's welfare is thus unambiguous. Despite the greater quality of the extremist's policy, the moderate is harmed; this greater quality is insufficient to compensate the moderate for the policy's greater extremism. While intuitive, this effect also differs from the all-pay contest, where the payoff of the weaker player is unaffected by the characteristics of the stronger one (Hillman and Riley (1989); Siegel (2009)).

Finally, the effect of unilateral extremism on the decisionmaker's welfare is surprising.

Proposition 5. *Unilateral changes in extremism have the following effects on the decisionmaker.*

- *If the developers are symmetrically capable and extreme ($|x_i| = |x_{-i}|, \alpha_i = \alpha_{-i}$) and developer i becomes intrinsically more extreme, the decisionmaker's utility locally increases*
- *As a developer becomes intrinsically more extreme ($|x_i| \to \infty$) the competitor's probability of developing a policy approaches* 0*, but the decisionmaker's utility approaches infinity*

While characterizing the precise local effect of a more-extreme developer is difficult, the broader relationship is simple and striking; the decisionmaker strongly benefits from unilateral extremism. If the developers begin symmetrically capable and extreme and one becomes more extreme, the decisionmaker benefits despite the resulting imbalance in participation. And as a developer becomes increasingly extreme, the decisionmaker becomes increasingly better off, even though his competitor also becomes vanishingly likely to participate. Unilateral extremism thus strongly benefits the decisionmaker, even though it decreases (and in the limit eliminates) observable competition.

This effect contrasts strikingly with the all-pay contest, where asymmetries always harm the decisionmaker due to the discouragement effect (Hillman and Riley (1989)). In our model, part of the discouragement effect is still present – the weaker player (the moderate) must be increasingly discouraged by the growing motivation of the stronger player (the extremist). The extremist, however, does *not* strategically respond by also weakly reducing his participation, as in the all pay contest. Instead, he crafts an increasingly extreme but also increasingly appealing policy, whose beneficial effect is sufficient enough to outweigh the cost to the decisionmaker of the moderate's discouragement.[13]

The Politics of Asymmetric Ability

We last turn to the politics of *asymmetric ability*, by studying equilibrium when the developers are equally extreme ($|x_i| = |x_{-i}|$) but one is more skilled at producing quality ($\alpha_i < \alpha_{-i}$). For convenience we call the more capable developer "the expert" and the other "the amateur." Our main finding is that asymmetric ability and extremism are effectively *observationally equivalent*.

Proposition 6. *If the developers are equally extreme but i is more skilled ($|x_i| = |x_{-i}|, \alpha_i < \alpha_{-i}$), then the equilibrium pattern of competition is identical to the pattern described in Proposition 2, in which the developers are equally capable but i is more extreme ($|x_i| > |x_{-i}|, \alpha_i = \alpha_{-i}$).*

The observational equivalence between asymmetric ability and extremism can be seen in Figure 3, which compares equilibrium with symmetric developers to equilibrium with asymmetric ability. The expert exploits his greater ability at crafting high quality policies to craft a more competitive but also more extreme policy, consistent with the finding in Hitt, Volden and Wiseman (2017) that "more effective lawmakers" (i.e., with lower α_i) "are more likely to offer successful proposals." The amateur reacts by both disengaging from policy development, and by moderating his policy when

[13] A broader contest theory literature studies implications of the discouragement effect (Chowdhury, Esteve-Gonzalez and Mukherjee (2022)). In many perfectly discriminating contests (where the outcome follows deterministically from the players' strategies, like the all pay contest) decisionmakers are harmed by asymmetries because of the discouragement effect – ours is a notable exception. However, if there is enough "noise" in the outcome, then the decisionmaker can also benefit from large asymmetries in an otherwise standard model. For example, in a Tullock (1980) contest, total effort approaches zero as one player's prize value approaches ∞ (meaning the decisionmaker is harmed from large asymmetries) unless there are weakly decreasing returns to scale ($r \leq 1$); this generates sufficient uncertainty in the outcome to tamp down the strength of the discouragement effect.

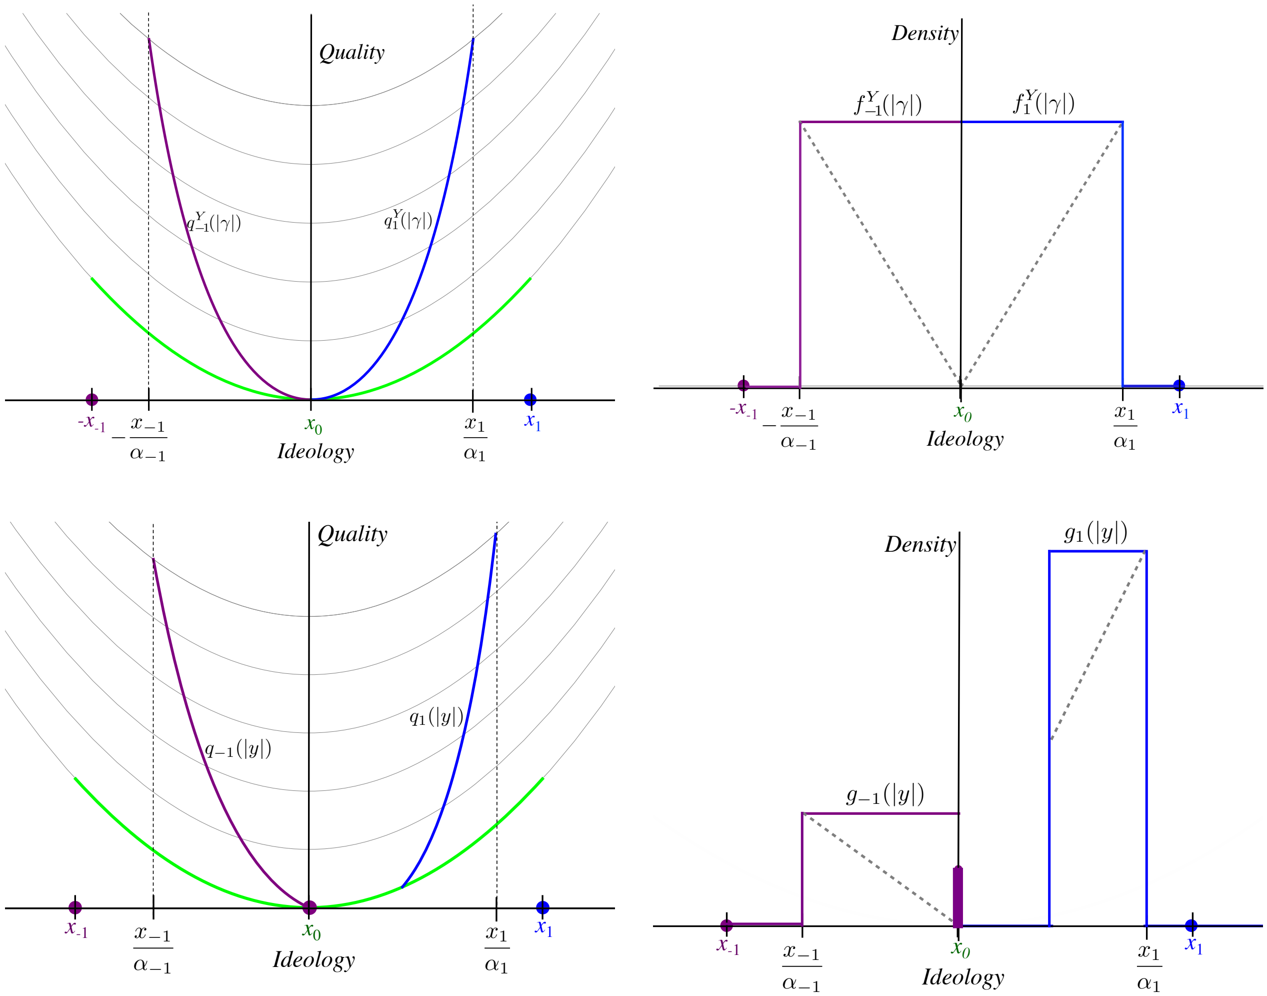

Figure 3: *Equilibrium with Symmetric vs. Asymmetric Ability*

he crafts one. The key empirical implication is that observably-extreme *behavior* by one political faction may not actually reflect greater underlying extremism, but rather greater ability at crafting "good policy" that is appealing on non-ideological grounds.[14]

The observational similarity between asymmetric extremism and ability extends to several consequences of one developer becoming more skilled.

[14]There is not an exact isomorphism between extremism and ability, unlike the isomorphism in the all-pay contest between the benefit from winning and the cost of effort. The reason is that our developers intrinsically value quality, so an up-front investment in quality is only partially "all-pay."

Proposition 7. *If developer i becomes more skilled (lower α_i)*

- *his own strategy and his opponent's strategy are affected in the same ways as when he becomes more ideologically extreme (higher $|x_i|$)*
- *the equilibrium utility of his competitor decreases*

A developer becoming more skilled thus increases his own activity (if he was the amateur) and makes his policy more extreme, higher quality, and better for the decisionmaker. It further decreases his competitor's activity (if he was the amateur) and makes his policy more moderate. Finally, it unambiguously harms his competitor.

The fact that unilaterally greater skill harms a competitor is somewhat surprising, given that the skill in question is at making *common value quality investments that benefit everyone.* Indeed, it is a striking demonstration of how "good policy" considerations cannot really be considered separately from "ideological" ones even if are *theoretically* distinct, because of how strategic actors will exploit their skill at crafting good policy to achieve their ideological goals. It also clarifies that strong disagreement between opposing sides of a policy conflict is not itself prima facie evidence that the policy domain lacks areas of agreement; indeed, such disagreements *should* emerge because the developers strategically use quality to curry favor with the decisionmaker, rather than the opposition.

We last examine the decisionmaker's welfare.

Proposition 8. *Unilateral changes in ability have the following effects on the decisionmaker.*

- *If the developers begin symmetric ($|x_i| = |x_{-i}|, \alpha_i = \alpha_{-i}$) and developer i becomes more skilled, then the decisionmaker's utility locally increases.*
- *As a developer becomes increasingly skilled ($\alpha_i \to 1$), the competitor's probability of developing a policy approaches 0, but the decisionmaker's utility approaches a strictly positive limit that is is strictly increasing in the competitor's extremism and ability.*

The decisionmaker thus benefits when a developer becomes unilaterally more skilled, even though he crafts a more extreme policy and his competitor participates less. And as one developer becomes increasingly skilled ($\alpha_i \to 1$), the effect again resembles that of a developer becoming increasingly extreme ($|x_i| \to \infty$), but with some notable differences. As before, the "weaker" developer (here the amateur, previously the moderate) is eventually driven out, but the decisionmaker still benefits from his *potential* participation (in the sense that her utility is bounded away from her utility under monopoly). The decisionmaker's utility doesn't increase unboundedly, however; rather, it approaches a strictly positive value that is generally higher than her utility under symmetry.[15]

Notably, the decisionmaker's limiting utility as the expert becomes more capable still depends on the traits of the amateur, even though he is effectively driven out of the contest. In other words, an *apparent* expert monopolist will still appreciably react to the traits of a largely inactive amateur; the empirical implication is that a seemingly irrelevant participant in the policy process can still critically influence behavior and outcomes. Interestingly, this result does not require a sequential agenda for policy development (as in Hitt, Volden and Wiseman (2017) and Lax and Cameron (2007)), where the first mover can explicitly attempt to craft a policy that deters the second mover from participating.

Strategic Extremism, Policy Spillovers, and Robustness

To sharpen the intuition for why asymmetries can benefit the decisionmaker in our model, we last consider several variants and compare them to the main model (see Appendix E for details). For our subsequent results we will say that *asymmetries always harm the decisionmaker* if she becomes strictly worse off whenever one of two symmetric developers becomes differentially extreme or capable – this is a core property of the two player all-pay contest.

[15] Specifically, if the developers begin symmetric and developer i becomes arbitrarily skilled, the decisionmaker will be better off than under symmetry as long as $\alpha_{-i} \geq \underline{\alpha} \approx 1.0435$.

Strategic Extremism A crucial property of our model is that the developers strategically *choose* the extremism of their policy. To see this, consider a variant where which each developer $i \in \{-1, 1\}$ may only target a fixed ideology $y_i = iy \neq 0$ (further assume that the decisionmaker has no outside option).[16] Then the objective function in equation 3 need only be maximized with respect to score s. As in the main model the developers' equilibrium scores CDFs must be continuous and strictly increasing over a common interval $[\underline{s}, \bar{s}]$ with $F_k(\underline{s}) > 0$ for at most one k.[17] Differentiating w.r.t. s and setting equal to 0 then yields that they must satisfy the modified differential equations

$$f_k(s) = \frac{\alpha_{-k} - F_{-k}(s)}{2x_{-k}(y_{-k} - y_k)} \quad \text{and} \quad f_{-k}(s) = \frac{\alpha_k - F_{-k}(s)}{2x_k(y_k - y_{-k})} \quad \forall s \in [\underline{s}, \bar{s}]. \tag{4}$$

Thus, in this variant each developer i's ideological "stakes" in the contest are fixed at $2x_i(y_i - y_{-i})$ rather than endogenous, making it more closely resembles the all-pay contest.[18]

Now consider what happens when the stronger developer $-k$ becomes yet stronger, either by becoming more extreme (higher x_{-k}) or more capable (lower α_{-k}). From equation 4, the CDF of *weaker* developer k must become *flatter* ($f_k(s)$ decreases $\forall s$) to preserve the stronger developer's willingness to mix over a common score interval $[\underline{s}, \bar{s}]$. But since both developers' CDFs must equal 1 at the same top score ($F_k(\bar{s}) = F_{-k}(\bar{s}) = 1$), this can only be accomplished by strictly increasing the weaker developer's probability of inactivity $F_k(\underline{s})$, so that $F_k(s)$ *first-order stochastically decreases*

[16] If the developers were constrained to ideologies with different degrees of extremism ($|y_i| < |y_{-i}|$) then developer i could achieve a higher score proposal at no cost, so it would be as if he has a "head start" in the contest (Kirkegaard (2012); see also Meirowitz (2008)).

[17] Unlike the main model, the low score $\underline{s} = -y^2$ is that of the developers' ideals y_i with 0 quality.

[18] With the assumption that the decisionmaker has no outside option, the only difference with the all-pay contest is that the developers intrinsically value quality, so the marginal cost $\alpha_i - F_{-i}(s)$ of increasing score s also depends on the probability $F_{-i}(s)$ a score-s policy wins. If the decisionmaker had an outside option (s, q) with $q \in (-y, y)$, then it would be *as if* the stakes are not fixed, because a developer could "compromise" by allowing the decisionmaker to enact $(0, q)$ instead of (s_i, y_i).

(i.e., increase $\forall s \in [\underline{s}, \bar{s}]$). This is exactly the first part of the discouragement effect; as a stronger developer becomes yet stronger, his opponent becomes less active.

What then happens to the stronger developer's score CDF $F_{-k}(s)$? As in the classic all-pay contest (but unlike in our model), the answer is *nothing at all.* The reason is simple: the stronger developer's score CDF is determined *only* by need to keep the weaker developer willing to target scores in $[\underline{s}, \bar{s}]$, and neither the weaker developer's preferences and abilities (x_k, α_k) nor his ideological stakes in the contest $2x_k(y_k - y_{-k})$ have changed. We thus recover the classic discouragement effect.

Proposition 9. *If the developers are constrained to craft policies with fixed ideologies $y_i = iy \neq 0$ and the decisionmaker has no outside options, then asymmetries always harm the decisionmaker.*

Policy Spillovers Another crucial property of our model is that the developers care about policy when they lose. To see this, consider a variant where the developers can choose the ideology of their policies, but receive a fixed payoff from losing.[19] Then the objective function of a developer i is:

$$-(\alpha_i - F_{-i}(s))\, s + F_{-i}(s) \cdot \left(2x_i y - x_i^2\right) - \alpha_i y^2 + \int_s^{\bar{s}} \left(-x_i^2\right) f_{-i}(s_{-i})\, ds_{-i},$$

Maximizing first with respect to y, the optimal ideology to target is $y_i(s) = F_{-i}(s)\frac{x_i}{\alpha_i}$ (as in the main model). Also as before, the developers' score CDFs must be continuous and strictly increasing over a common interval $[0, \bar{s}]$ with $F_k(0) > 0$ for at most one k. Differentiating w.r.t. s and setting equal to 0 yields that they must satisfy the modified differential equations

$$f_k(s) = \frac{\alpha_{-k} - F_k(s)}{2x_{-k}y_{-k}(s)} \text{ and } f_{-k}(s) = \frac{\alpha_k - F_{-k}(s)}{2x_k y_k(s)} \;\; \forall s \in [0, \bar{s}]. \tag{5}$$

This variant more closely resembles our main model in that each developer i's ideological stakes $2x_i y_i(s)$ are endogenous; but this is insufficient to yield a benefits from asymmetries. The reason is that the magnitude of these stakes depends only on a developer's *own* endogenous extremism $y_i(s)$, which in turn depends *only* on his opponent's score CDF $F_{-i}(s)$. Thus, as in the previous

[19] For simplicity we assume this to be the payoff from policy $(0, 0)$.

variant, when a stronger developer $-k$ becomes yet stronger (higher x_{-k} or lower α_{-k}) the score CDF $F_k(s)$ of the weaker developer first-order stochastically decreases, while that of the stronger developer $F_{-k}(s)$ remains unchanged. We then again recover the classic discouragement effect.

Proposition 10. *If the developers only care about the policy outcome when they win, then asymmetries always harm the decisionmaker.*

Main Model We now return to the main model, to reexamine how the *combination* of strategic extremism and policy spillovers breaks the classic discouragement effect. Recall the developers' score CDFs are continuous and strictly increasing over $[0, \bar{s}]$ with $F_k(0) > 0$ for at most one k, and satisfy

$$f_k(s) = \frac{\alpha_{-k} - F_k(s)}{2x_{-k}(y_{-k}(s) - y_k(s))} \quad \text{and} \quad f_{-k}(s) = \frac{\alpha_k - F_{-k}(s)}{2x_k(y_k(s) - y_{-k}(s))} \quad \forall s \in [0, \bar{s}]$$

A developer's ideological stakes $2x_i(y_i(s) - y_{-i}(s))$ at a score s are now endogenous to both his own policy $y_i(s)$ and his opponent's policy $y_{-i}(s)$. Again consider what happens as the stronger developer $-k$ becomes yet stronger (higher x_{-k} or lower α_{-k}). As before, the score CDF of the weaker developer k must become flatter ($f_k(s)$ decreases $\forall s$) to incentivize the stronger developer to target scores in $[0, \bar{s}]$; if both developers' score CDFs still equal 1 at the same top score ($F_k(\bar{s}) = F_{-k}(\bar{s}) = 1$) then the weaker developer's score CDF must again first-order stochastically decrease.

But this is not the end of the story; now the effect of these changes "spills over" onto the incentives of the weaker developer. Why? Because the weaker developer ideological stakes depend on the stronger developer's endogenous extremism $y_{-k}(s) = \frac{x_{-k}}{\alpha_{-k}} F_k(s)$, which in turn depends on the weaker developer's participation $F_k(s)$. Intuitively, should the weaker developer k start dropping out of the contest, the stronger developer $-k$ will react with more ideological aggression, raising the stakes to k and *pulling him back in.* To ensure the weaker developer remains willing to target scores in $[0, \bar{s}]$ (i.e., to keep him from rushing all the way back in), the score CDF of the stronger developer k must then *also* become flatter ($f_{-k}(s)$ decreases $\forall s$). Finally, since the stronger developer $-k$ must always be active ($F_{-k}(0) = 0$), his score CDF must *first order stochastically increase* ($F_{-k}(s)$

decreases $\forall s$). In other words, the stronger developer *must craft better policies for the decisionmaker* to incentivize his weaker opponent to participate less despite the fact that doing so will trigger more ideological aggression. The combined effect for the decisionmaker is beneficial, despite increasingly extreme policies being crafted by an increasingly dominant developer.

Quadratic Costs and Linear Benefits Lastly, suppose the players' utilities are *linear* in ideology ($U_D(y,q) = \lambda q - |y|$ and $U_i(y,q) = \lambda q + x_i y$ for $i \in \{-1,1\}$ where $|x_i| > 1$), while the costs of generating quality are *quadratic* ($c_i(q) = \frac{a_i}{2}q^2$). With these functional forms a score-s policy with ideology y has quality $q = \frac{s+|y|}{\lambda}$, and the objective function of a developer i becomes

$$-\frac{\alpha_i}{2}(s+|y|)^2 + F_{-i}(s)(s+|y|+x_i y) + \int_s^{\bar{s}} (s+|y_{-i}(s)|+x_i y_{-i}(s)) f_{-i}(s_{-i}ds_{-i}),$$

where $\alpha_i = \frac{a_i}{\lambda^2}$. This has both similarities and differences with the baseline model (see equation 3).

With respect to similarities, the marginal benefit of moving the ideological outcome in i's direction *holding score fixed* remains linear. With respect to differences, the marginal cost of compensating the decisionmaker for ideological movements holding score fixed becomes $\alpha_i(s+y)$ (instead of $2\alpha_i$), which depends directly on the score s being targetted. Intuitively, with quadratic quality costs it becomes more expensive to compensate the decisionmaker for ideological concessions on higher score policies because a developer must first generate the baseline quality needed to achieve those scores. Consequently, a developer's optimal ideology becomes $y_i(s) = i \cdot \max\left\{\frac{1+|x_i|}{\alpha_i}F_{-i}(s) - s, 0\right\}$. This optimum shares a crucial property with the baseline model; that a developer becomes more ideologically aggressive ($|y_i(s)|$ increases) when his opponent participates less (higher $F_{-i}(s)$). However, it also differs in an interesting way; higher score policies *within the support of a developer's strategy* are no longer necessarily more extreme due to the higher marginal cost of quality on higher scores.

As in the baseline model, the developers' score CDFs must be continuous and strictly increasing over a common interval $[0,\bar{s}]$ with $F_k(0) > 0$ for at most one $k \in \{-1,1\}$. Differentiating w.r.t. s

and setting equal to 0 yields the modified differential equations

$$f_k(s) = \frac{\alpha_{-k}(s - ky_{-k}(s)) - F_k(s)}{(x_{-k} - k)\, y_{-k}(s) - (x_{-k} + k)\, y_k(s)} \text{ and } f_{-k}(s) = \frac{\alpha_k(s + ky_k(s)) - F_{-k}(s)}{(x_k + k)\, y_k(s) - (x_k - k)\, y_{-k}(s)} \quad (6)$$

for all $s \in [0, \bar{s}]$. This system lacks the simplicity of baseline model; both because the optimal ideology might be the "corner" solution of the decisionmaker's ideal, and because the score s also enters the differential equations directly (rather than only indirectly via the score CDFs $F_i(s)$). It thus does not appear to admit a closed form solution in the asymmetric case (we compute these numerically – see Appendix for details). Nevertheless, it retains the key properties that generate a benefit from asymmetries; as the stronger developer $-k$ becomes yet stronger the weaker developer k becomes discouraged, which then elicits more ideological aggression from the stronger developer. As a result, asymmetries that result from one developer becoming more extreme or capable lead him to craft a first-order stochastically better policy for the decisionmaker despite the discouragement of his competitor, potentially benefitting the decisionmaker.

Figure 4 depicts an example in which asymmetric extremism benefits the decisionmaker; as before the top panels depict symmetric developers while the bottom panels depict a more extreme right developer, and the left panels depict the policies the developers randomize over (note the decisionmaker's indifference curves now linear). Positive-quality policies no longer necessarily diverge from the decisionmaker's ideal, nor are higher score policies within a mixed strategy necessarily more extreme. The right panels depict the developers' equilibrium score CDFs; the blue curve depicts the score CDF of the right developer, the purple curve depicts the score CDF of the left developer, and the black curve depicts the CDF of the maximum score policy (i.e., the decisionmaker's utility). A more-extreme right developer crafts a more ideologically extreme policy that is also better for the decisionmaker (both first-order stochastically), and the left developer becomes both strictly less likely to participate and crafts policies over a wider range of scores. The net effect is a first-order stochastic increase in the decisionmaker's equilibrium utility (i.e., a rightward shift in the black curve).

Numerical results broadly illustrate a benefit from asymmetries in the quadratic-linear variant.

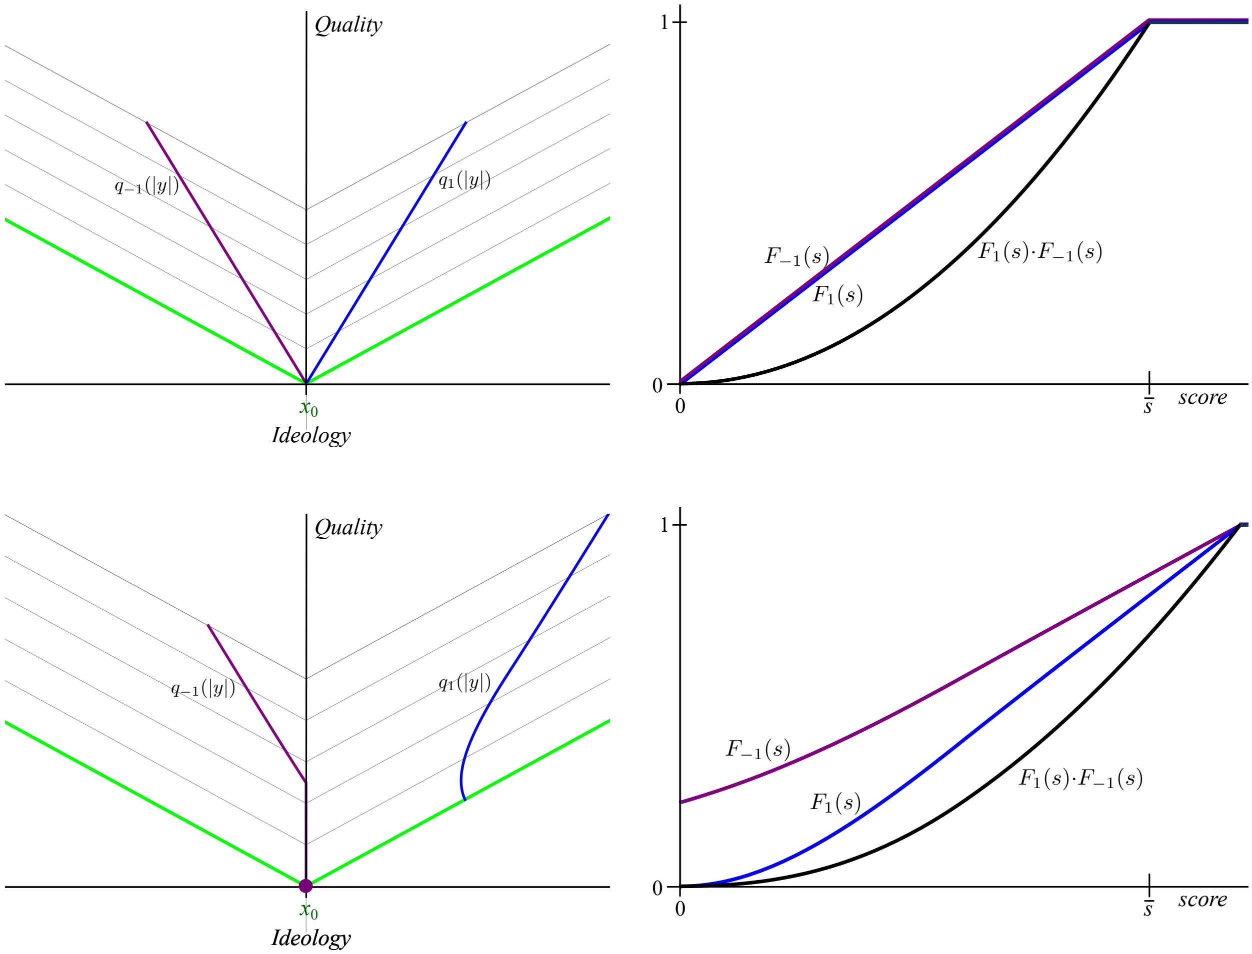

Figure 4: *Equilibrium with Symmetric vs. Asymmetric Extremism (Linear Ideological Preferences and Quadratic Costs of Quality)*

Figure 5 is a contour plot of the *net benefit* to the decisionmaker of a more-extreme vs. equally extreme right developer x_1 for a range of extremism of the left developer x_{-1} (holding a common α fixed). In the green regions the decisionmaker benefits from the right developer's greater extremism, whereas in the red regions she is harmed. When the left developer is relatively moderate the decisionmaker is locally harmed by a more extreme right developer x_1 due to the discouragement effect (in contrast to the main model) – but she nevertheless benefits once the right developer becomes sufficiently extreme. When the left developer is more extreme the decisionmaker always benefits from a more-extreme right developer x_1. Numerical results for the effects of asymmetric ability are

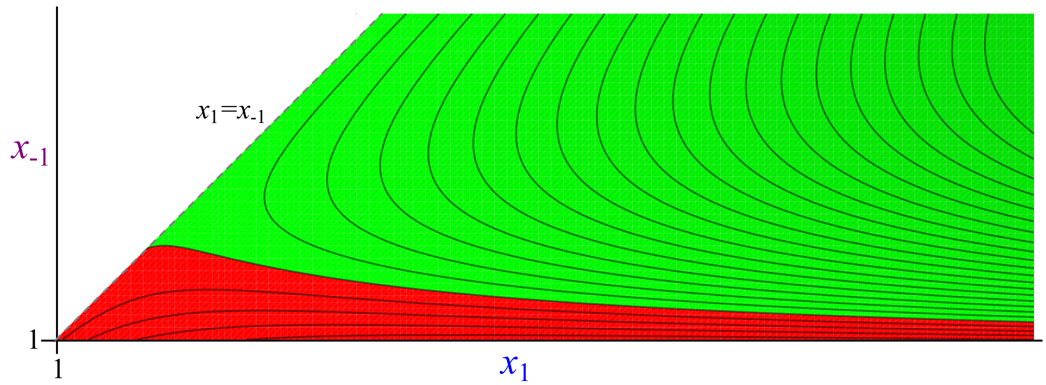

Figure 5: *Net benefit to decisionmaker of right developer being more extreme than left developer*

broadly similar – the decisionmaker is only harmed by having a more capable right developer (as compared to symmetry) when the left developer is already very capable to begin with, so that the impact of his discouragement from asymmetries on the decisionmaker's utility is large.

Discussion and Conclusion

We have studied policy development by competing developers who differ in their ideological extremism and/or ability at crafting high quality policies. We have shown that a decisionmaker can strongly benefit from such asymmetries (at the expense of one developer) despite increasingly imbalanced policies and outcomes, and even seemingly absent observable competition. The model thus provides a novel rationale for how ideological extremism may come to dominate policymaking (as in Osborne, Rosenthal and Turner (2000)) that is rooted in the nature of *productive* policy competition rather than dysfunction, bias, capture, or some other systemic failure. It further illustrates how asymmetrically extreme behavior may result from asymmetric ability rather than asymmetrically-extreme preferences, which has implications for how political scientists measure the preferences of political actors from their observed behavior. Finally, it shows how ostensibly nonpartisan "good policy" considerations and partisan ideological ones are inextricably linked, because of how strategic actors exploit ability at crafting good policy to gain ideological influence.

Given the surprising nature of our findings, it is worth briefly remarking on the boundaries of our model's empirical domain. Clearly, in some issue areas disagreement may become so pathological that participants value the "quality" of ideologically-distant policies *negatively.* Consider for example the politics of reproductive rights; pro-life voters likely place an intrinsic negative value on many policy attributes that pro-choice voters would associate with quality, such as population coverage and cost-effectiveness. Hirsch and Shotts (2015) consider a variant of the symmetric policy development model in which the developers can value the quality of each others' policies negatively, and show that this actually strengthens the benefits of polarization by raising the intensity of competition. We conjecture that this effect would similarly extend to the benefits of asymmetries; indeed, the driving force of our model is not that the developers have a shared notion of quality between each other, but rather that each has a shared notion of quality with the decisionmaker.

Second, there are issue domains where shared notions of quality are meaningfully present, but are superseded by other (possibly strategic) considerations. For example, when policymaking is dynamic, implementing a high-quality policy "today" might improve one actor's control over policymaking "tomorrow"; this can give a competing actor the incentive to sabotage the policy (in the sense of damaging quality that they intrinsically value) to improve their prospects for future control (Gieczewski and Li (2022); Hirsch and Kastellec (2022)). The applicability of our model requires that such destructive means of policy influence be absent, relatively costly to employ (as compared to the productive means we study), or prohibited by either formal rules or shared norms of governance.

Finally, our analysis suggests several avenues for follow-on work. One is to directly consider the politics of destructive influence, by studying developers who can also sabotage each others' policies or pursue other unproductive activities (like bribing the decisionmaker or engaging in advertising, lobbying, or grassroots mobilization). Hirsch and Shotts (2015) study an extension with the option of sabotage (i.e., costly up-front effort to reduce the quality of an opponent's policy), and show that it will not be used in equilibrium unless it is significantly cheaper than investing in quality. However, they do not solve for equilibrium when this condition fails. It would be interesting to study the effects of one actor choosing to specialize in productive policy development when the other chooses to specialize in sabotage; might a decisionmaker actually benefit from the presence of the saboteur because he better motivates a productive competitor?

A second avenue (following the classical literature on policy expertise) is to study how political institutions can be *designed* to encourage effective policymaking in competitive policy environments.[20] What if the developers can be chosen by the decisionmaker, as in the literature on legislative committee composition (Krehbiel (1992); see also Hirsch and Shotts (2012))? What if there are existing developers – when would the decisionmaker want to subsidize their activities and how? What if it is not the identities of the developers under consideration but that of the decisionmaker, as in a President appointing an agency head to consider proposals from career staff and outside groups (Lewis (2008))? How would the decisionmaker bias the preferences of the appointee to shape productive competition between the developers? And what if the decisionmaker is not a unitary actor but a collective choice body, as in a legislature; what sort of collective choice rules will best encourage the development of high quality legislation? We hope to explore these and other avenues in future work.

References

Aghion, Philippe, and Jean Tirole. 1997. "Formal and Real Authority in Organizations." *Journal of Political Economy*, 105(1): 1–29.

Aragones, Enriqueta, and Thomas R Palfrey. 2002. "Mixed Equilibrium in a Downsian Model with a Favored Candidate." *Journal of Economic Theory*, 103(1): 131–161.

Ashworth, Scott, and Ethan Bueno de Mesquita. 2009. "Elections with Platform and Valence Competition." *Games and Economic Behavior*, 67(1): 191–216.

Balart, Pau, Agustin Casas, and Orestis Troumpounis. 2022. "Technological change, campaign spending and polarization." *Journal of Public Economics*, 211: 104666.

[20] Hirsch and Shotts (2018) conduct a similar exercise when there is a monopoly developer.

Bartlett, Bruce. 2017. "Trump Is What Happens When a Political Party Abandons Ideas." *POLITICO Magazine*. (www.politico.com/magazine/story/2017/06/24/intellectual-conservatives-lost-republican-trump-215259).

Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 1996. "The All-Pay Auction with Complete Information." *Economic Theory*, 8(2): 291–305.

Baye, Michael R., Dan Kovenock, and Casper G. de Vries. 2012. "Contests with rank-order spillovers." *Economic Theory*, 51(2): 315–350.

Callander, Steven. 2008. "A Theory of Policy Expertise." *Quarterly Journal of Political Science*, 3(2): 123–140.

Chowdhury, Subhasish M., Patricia Esteve-Gonzalez, and Anwesha Mukherjee. 2022. "Heterogeneity, Leveling the Playing Field, and Affirmative Action in Contests." *Southern Economic Journal*, 89: 924–74.

Clark, Tom S., and Clifford J. Carrubba. 2012. "A Theory of Opinion Writing in a Political Hierarchy." *Journal of Politics*, 74(2): 584–603.

Clinton, Joshua D. 2012. "Using Roll Call Estimates to Test Models of Politics." *Annual Review of Political Science*, 15(1): 79–99.

Crisman-Cox, Casey, and Michael Gibilisco. 2024. "Tug of War: The Heterogeneous Effects of Outbidding between Terrorist Groups." Working paper.

Gieczewski, German, and Christopher Li. 2022. "Dynamic Policy Sabotage." *American Journal of Political Science*, 66(3): 617–629.

Groseclose, Tim. 2001. "A Model of Candidate Location When One Candidate Has a Valence Advantage." *American Journal of Political Science*, 45(4): 862–886.

Grossmann, Matt, and David A. Hopkins. 2016. *Asymmetric Politics: Ideological Republicans and Group Interest Democrats.* New York:Oxford University Press.

Hillman, Arye L., and John G. Riley. 1989. "Politically Contestable Rents and Transfers." *Economics & Politics*, 1(1): 17–39.

Hirsch, Alexander V. 2023. "Polarization and Campaign Spending in Elections." *The Journal of Politics*, 85(1): 240–249.

Hirsch, Alexander V., and Jonathan P. Kastellec. 2022. "A Theory of Policy Sabotage." *Journal of Theoretical Politics*, 34(2): 191–218.

Hirsch, Alexander V., and Kenneth W. Shotts. 2012. "Policy-Specific Information and Informal Agenda Power." *American Journal of Political Science*, 56(1): 67–83.

Hirsch, Alexander V., and Kenneth W. Shotts. 2015. "Competitive Policy Development." *American Economic Review*, 105(4): 1646–64.

Hirsch, Alexander V., and Kenneth W. Shotts. 2018. "Policy-Development Monopolies: Adverse Consequences and Institutional Responses." *The Journal of Politics*, 80(4): 1339–1354.

Hitt, Matthew P., Craig Volden, and Alan E. Wiseman. 2017. "Spatial Models of Legislative Effectiveness." *American Journal of Political Science*, 61(3): 575–590.

Huber, John D., and Nolan McCarty. 2004. "Bureaucratic Capacity, Delegation, and Political Reform." *American Political Science Review*, 98(3): 481–494.

Kerwin, Cornelius Martin, and Scott R. Furlong. 2018. *Rulemaking: How Government Agencies Write Law and Make Policy.* . 5th edition ed., Thousand Oaks, CA:CQ Press.

Kingdon, John W. 1984. *Agendas, Alternatives, and Public Policies.* Boston:Little, Brown.

Kirkegaard, Ren. 2012. "Favoritism in asymmetric contests: Head starts and handicaps." *Games and Economic Behavior*, 76(1): 226–248.

Krehbiel, Keith. 1992. *Information and Legislative Organization.* Ann Arbor:University of Michigan Press.

Lax, Jeffrey R., and Charles M. Cameron. 2007. "Bargaining and Opinion Assignment on the US Supreme Court." *Journal of Law, Economics, and Organization*, 23(2): 276–302.

Lewis, David E. 2008. *The Politics of Presidential Appointments.* Princeton:Princeton University Press.

Londregan, John B. 2000. *Legislative Institutions and Ideology in Chile.* New York:Cambridge University Press.

McCarty, Nolan. 2015. "What We Know and Do Not Know about Our Polarized Politics." In *Political Polarization in American Politics.* , ed. Daniel J. Hopkins and John Sides. London:Bloomsbury.

Meirowitz, Adam. 2008. "Electoral Contests, Incumbency Advantages, and Campaign Finance." *Journal of Politics*, 70(03): 681–699.

Osborne, Martin J., Jeffrey S. Rosenthal, and Matthew A. Turner. 2000. "Meetings with Costly Participation." *The American Economic Review*, 90(4): 927–943.

Reynolds, Molly E. 2020. "The Decline in Congressional Capacity." In *Congress Overwhelmed: The Decline in Congressional Capacity and Prospects for Reform.* , ed. Timothy M. LaPira, Lee Drutman and Kevin R. Kosar. Chicago, IL:University of Chicago Press.

Serra, Gilles. 2010. "Polarization of What? A Model of Elections with Endogenous Valence." *The Journal of Politics*, 72(2): 426–437.

Siegel, Ron. 2009. "All-Pay Contests." *Econometrica*, 77(1): 71–92.

Ting, Michael M. 2011. "Organizational Capacity." *Journal of Law, Economics, and Organization*, 27(2): 245–71.

Tullock, Gordon. 1980. "Efficient Rent Seeking." In *Toward a Theory of the Rent Seeking Society.* , ed. James M. Buchanan, Robert D. Tollison and Gordon Tullock. College Station, TX:Texas A&M University Press.

Turner, Ian R. 2017. "Working smart and hard? Agency effort, judicial review, and policy precision." *Journal of Theoretical Politics*, 29(1): 69–96.

Wiseman, Alan E. 2006. "A Theory of Partisan Support and Entry Deterrence in Electoral Competition." *Journal of Theoretical Politics*, 18(2): 123–158.

Yackee, Jason Webb, and Susan Webb Yackee. 2006. "A Bias Towards Business? Assessing Interest Group Influence on the U.S. Bureaucracy." *Journal of Politics*, 68(1): 128–139.

Productive Policy Competition and Asymmetric Extremism

Online Appendix

This Appendix is divided into five parts. Appendix A is a general analysis of the model concluding with a statement of necessary and sufficient conditions for equilibrium. Appendix B derives the closed-form characterization of the equilibrium strategies given in main text Proposition 1. Appendix C analyzes properties of equilibrium using this characterization. Appendix D describes where to locate results in the main text propositions in the general model analysis in Appendices B-C. Appendix E analyzes several variants of the model (including one with alternative functional forms) to isolate which properties of the main model are crucial for our results.

A General Equilibrium Characterization

We begin with a (slightly) more general formulation of the model than stated in the main text. Two developers labelled -1 (left) and 1 (right) craft competing policies for consideration by a decisionmaker (DM), labelled player 0. A policy (γ, q) consists of an *ideology* $\gamma \in \mathbb{R}$ and a level of *quality* $q \in [0, \infty) = \mathbb{R}^+$. Utility over policies takes the form

$$U_i(\gamma, q) = \lambda q - (\gamma - X_i)^2,$$

where X_i is player $i's$ ideological ideal point, and λ is the weight all players place on quality. The developers' ideal points are on either side of the decisionmaker ($X_{-1} < X_D < X_1$).

The game is as follows. First, the developers simultaneously craft policies (γ_i, q_i); crafting a policy with quality q_i costs $c_i(q_i) = a_i q_i$, where $a_i > \lambda$. Second, the DM chooses one of the two policies or something else from an exogenous set of outside options $\mathbb{O}$, where $\mathbb{O}$ may contain the DM's ideal point with no quality $(0, 0)$ and/or policies that are strictly worse (and can be empty).

A.1 Preliminary Analysis

The game is a multidimensional contest in which the scoring rule applied to "bids" (γ, q) is just the DM's utility $U_D(\gamma, q) = \lambda q - (X_D - \gamma)^2$. To facilitate the analysis we thus reparameterize policies (γ, q) to be expressed in terms of (s, y), where $y = \gamma - X_D$ is the (signed) distance of a policy's ideology from the DM's ideal, and $s = \lambda q - y^2$ is the DM's utility for a policy or its *score*. The implied quality of a policy (s, y) is then $q = \frac{s+y^2}{\lambda}$. Using this we re-express the developers' utility and cost functions in terms of (s, y). Note that the decisionmaker's ideal point with 0-quality has exactly 0 score, and is the most competitive "free" policy to craft.

Definition A.1.

1. *Player i's utility for policy (s, y) is*

$$V_i(s, y) = U_i\left(y + X_D, \frac{s + y^2}{\lambda}\right) = -x_i^2 + s + 2x_i y$$

where $x_i = X_i - X_D$ is the (signed) distance of i's ideal from the DM.

2. *Developer i's cost to craft policy (s, y) is*

$$c_i\left(\frac{s + y^2}{\lambda}\right) = \frac{a_i}{\lambda}\left(s + y^2\right) = \alpha_i\left(s + y^2\right)$$

where $\alpha_i = \frac{a_i}{\lambda}$ is i's weighted marginal cost of generating quality.

Definition 1 reparameterizes policies into score and ideological distance (henceforth just ideology) (s, y), and the five primitives (X_i, a_i, λ) into four parameters (x_i, α_i) describing the developers' (signed) ideal ideological distance from the DM $x_i = X_i - X_D$ (henceforth just ideal ideology) and weighted marginal costs of generating quality $\alpha_i = \frac{a_i}{\lambda}$ (henceforth just costs).

A.1.1 Necessary and Sufficient Equilibrium Conditions

In the reparameterized game, a developer's pure strategy (s_i, y_i) is a two-dimensional element of $\mathbb{B} \equiv \{(s, y) \in R^2 \mid s + y^2 \geq 0\}$. A mixed strategy σ_i is a probability measure over the Borel subsets of $\mathbb{B}$, and let $F_i(s)$ denote the CDF over scores induced by i's mixed strategy σ_i.[21]

We now derive necessary and sufficient equilibrium conditions in a series of four lemmas. Let $\bar{\Pi}_i(s_i, y_i; \sigma_{-i})$ denote i's expected utility for crafting a policy (s_i, y_i) with $s_i \geq 0$ *if* a tie would be broken in his favor. Clearly this is i's expected utility from crafting a policy with any $s_i > 0$ where $-i$ has no atom, and i can always achieve utility arbitrarily close to $\bar{\Pi}_i(s_i, y_i; \sigma_{-i})$ by crafting $\varepsilon-$higher score policies. Now $\bar{\Pi}_i(s_i, y_i; \sigma_{-i}) =$

$$-\alpha_i\left(s_i + y_i^2\right) + F_{-i}(s_i) \cdot V_i(s_i, y_i) + \int\limits_{s_{-i} > s_i} V_i(s_{-i}, y_{-i})\, d\sigma_{-i}. \tag{A.1}$$

The first term is the up-front cost of generating the policy's quality. The second term is the probability i's policy is selected, times his utility for it. The third term is i's utility should he lose, which requires integrating over all the policies in the support of his opponent's mixed strategy with score higher than s_i. Taking the derivative with respect to y_i and setting equal to 0 yields the first Lemma.

Lemma A.1. *At any score $s_i > 0$ where $F_{-i}(\cdot)$ has no atom, the policy $(s_i, y_i^*(s_i))$, where $y_i^*(s_i) = F_{-i}(s_i) \cdot \frac{x_i}{\alpha_i}$, is the strictly best score-s_i policy.*

Proof: Straightforward. QED

Lemma A.1 states that at almost every score $s_i > 0$, developer i's unique best combination of ideology and quality to generate that score is just a weighted average of the developer's and DM's ideal ideologies $\frac{x_i}{\alpha_i}$, multiplied by the probability $F_{-i}(s_i)$ that i's opponent crafts a lower-score policy. Note that i's optimal ideology does not depend directly on his opponent $-i$'s ideologies, since a policy's ideology (holding score fixed) only matters conditional on winning. The optimal ideology also depends on the exact score s_i only *indirectly* through probability $F_{-i}(s_i)$ the policy wins the contest, since i's utility conditional on winning is additively separable in score and ideology.

[21] For technical convenience we restrict attention to strategies generating score CDFs that can be written as the sum of an absolutely continuous and a discrete distribution.

The second lemma establishes that at least one of the developers is always *active*, in the sense of crafting a policy with strictly positive score (all positive-score policies are positive-quality, but the reverse is not necessarily true).

Lemma A.2. *In equilibrium $F_k(0) > 0$ for at most one k.*

Proof: Suppose not, so $F_i(0) > 0\,\forall i$ in some equilibrium. Let U_i^* denote developer i's equilibrium utility, which can be achieved by mixing according to his strategy conditional on crafting a score-$s \leq 0$ policy. Let $\bar{y}^0$ denote the expected ideological outcome and $\bar{s}^0$ the expected score outcome conditional on both sides crafting score ≤ 0 policies. Since $x_L < 0 < x_R$, we have $V_k\left(\bar{s}^0, \bar{y}^0\right) \leq V_k(0,0)$ for at least one k, which implies k has a profitable deviation since $U_k^* \leq \bar{\Pi}_k(0,0;\sigma_{-k}) < \bar{\Pi}_k(0, y_k^*(0);\sigma_{-k})$ (since $F_{-k}(0) > 0$). QED

The third Lemma establishes that in equilibrium there is 0 probability of a tie at a positive score.

Lemma A.3. *In equilibrium there is 0-probability of a tie at scores $s > 0$.*

Proof: Suppose not, so each developer's strategy generates an atom of size $p_i^s > 0$ at some $s > 0$. Developer i achieves his equilibrium utility U_i^* by mixing according to his strategy conditional on a score-s policy. Let $\bar{y}^s$ denote the expected ideological outcome conditional on both sides crafting score-s policies; then $V_k(s, \bar{y}^s) \leq V_k(s, 0)$ for at least one k, who has a profitable deviation. If k's policy at score s is $(s, 0)$, then $U_k^* \leq \bar{\Pi}_k(s, 0;\sigma_{-k}) < \bar{\Pi}_k(s, y_k^*(s);\sigma_{-k})$ (since $F_{-k}(s) > 0$). If k sometimes crafts something else, then $U_k^* < \left(1 - \frac{p_{-k}}{F_{-k}(s)}\right)\bar{\Pi}_k(s, E[y_k|s];\sigma_{-k}) + \left(\frac{p_{-k}}{F_{-k}(s)}\right)\bar{\Pi}_k(s, 0;\sigma_{-k})$, which is k's utility if he were to instead craft $(s, 0)$ with probability $\frac{p_{-k}}{F_{-k}(s)}$, and the expected ideology $E[y_k|s]$ of his strategy at score s with the remaining probability (and always win ties). QED

Lemmas A.1 – A.3 jointly imply that in equilibrium, developer i can compute his expected utility *as if* his opponent only crafts policies of the form $\left(s_{-i}, y_{-i}^*(s_{-i})\right)$. The utility from crafting *any* policy

(s_i, y_i) with $s_i > 0$ where $-i$ has no atom (or a tie would be broken in i's favor) is therefore

$$\bar{\Pi}_i^* \left(s_i, y_i; F\right) = -\alpha_i \left(s_i + y_i^2\right) + F_{-i}\left(s_i\right) \cdot V_i \left(s_i, y_i\right) + \int_{s_i}^{\infty} V_i \left(s_{-i}, y_{-i}^* \left(s_{-i}\right)\right) dF_{-i}. \tag{A.2}$$

Developer i's utility from crafting the *best* policy with score s_i is $\bar{\Pi}_i^* \left(s_i, y_i^* \left(s_i\right); F\right)$, which we henceforth denote $\bar{\Pi}_i^* \left(s_i; F\right)$.

Fourth and finally, we establish that equilibrium score CDFs must satisfy the following natural properties arising from the all pay component of the contest.

Lemma A.4. *Support of the equilibrium score CDFs over* $\mathbb{R}^+$ *is common, convex, and includes* 0.

Proof: We first argue $\hat{s} > 0$ in support of $F_i \rightarrow F_{-i}(s) < F_{-i}(\hat{s})\ \forall s < \hat{s}$. Suppose not; so $\exists s < \hat{s}$ where $-i$ has no atom and $F_{-i}(s) = F_{-i}(\hat{s})$. Then $\bar{\Pi}_i(\hat{s}, y_i; F) - \bar{\Pi}_i(s, y_i; F) = -(\alpha_i - F_{-i}(\hat{s})) \cdot (\hat{s} - s) < 0$, implying i's best score-s policy is strictly better than his best score-$\hat{s}$ policy, a contradiction. We now argue this yields the desired properties. First, an $\hat{s} > 0$ in i's support but not $-i$ implies $\exists\ \delta > 0$ s.t. $F_{-i}(s - \delta) = F_{-i}(s)$. Next, if the common support were not convex or did not include 0, then there would $\exists\, \hat{s} > 0$ in the common support s.t. neither developer has support just below, so $F_i(s) < F_i(\hat{s})\ \forall i,\ s < \hat{s}$ would imply both developers have atoms at $\hat{s}$, a contradiction. QED

We conclude by combining the preceding lemmas to provide a preliminary characterization of all equilibria in the form of necessary and sufficient conditions.

Proposition A.1. *Necessary conditions for SPNE are as follows:*

1. ***(Ideological Optimality)*** *With probability* 1*, policies are either*

   (a) *negative score* $s_i \leq 0$ *and* 0*-quality* $\left(s_i + y_i^2 = 0\right)$

   (b) *positive score* $s_i > 0$ *with ideology* $y_i = y_i^*\left(s_i\right) = \left(\frac{x_i}{\alpha_i}\right) F_{-i}\left(s_i\right)$.

2. ***(Score Optimality)*** *The profile of score CDFs* (F_i, F_{-i}) *satisfy the following boundary conditions and differential equations.*

- ***(Boundary Conditions)*** *$F_k(0) > 0$ for at most one developer k, and there $\exists \bar{s} > 0$ such that $\lim_{s \to \bar{s}} \{F_i(s)\} = 1 \; \forall i$.*

- ***(Differential Equations)*** *For all i and $s \in [0, \bar{s}]$,*

$$\alpha_i - F_{-i}(s) = f_{-i}(s) \cdot 2x_i \left(y_i^*(s) - y_{-i}^*(s)\right)$$

The above and $F_i(s) = 0 \; \forall i, s < 0$ are sufficient for equilibrium.

Proof: (*Score Optimality*) A score $\hat{s} > 0$ in the common support implies $[0, \hat{s}]$ in the common support (by Lemma A.4) implying $\lim_{s \to \hat{s}^-} \left\{\bar{\Pi}_i(s; F)\right\} \geq U_i^*$. Equilibrium also requires $\bar{\Pi}_i(s; F) \leq U_i^* \; \forall s$ so $\bar{\Pi}_i(s; F) = U_i^* \; \forall s \in [0, \bar{s}]$, further implying the F's are absolutely continuous over $(0, \infty)$ (given our initial assumptions), and therefore $\frac{\partial}{\partial s}\left(\Pi_i^*(s; F)\right) = 0$ for almost all $s \in [0, \bar{s}]$. This straightforwardly yields the differential equations for score optimality, with the boundary conditions implied by Lemma A.4. (*Ideological Optimality*) At most one developer k crafts ≤ 0-score policies with positive probability, so $F_{-k}(0) = 0$. Such policies lose for sure and never influence a tie, and therefore must be 0-quality with probability 1, yielding property (a). Atomless score CDFs $\forall s > 0$ implies $(s, y_i^*(s))$ is the strictly best score-s policy (by Lemma A.1), yielding property (b). (*Sufficiency*) Necessary conditions imply all $(s, y_i^*(s))$ with $s \in (0, \bar{s}]$ yield a constant U_i^*. $F_{-k}(0) = 0$ implies k's strictly best score-0 policy is $(0, y_k^*(0)) = (0, 0)$ and yields $\bar{\Pi}_k(0; F)$, and $F_k(s) = 0$ for $s < 0$ implies k has a size $F_k(0)$ atom here. Thus both developers' mixed strategies yield U_i^*, and neither can profitably deviate to $s \in (0, \bar{s}]$. To see neither can profitably deviate to $s > \bar{s}$, observe $\Pi_i^*(s; F) - \Pi_i^*(\bar{s}; F) = -(\alpha_i - 1)(s - \bar{s}) < 0$. To see k cannot profitably deviate to $s_k \leq 0$, $F_{-k}(0) = 0$ implies such policies lose and never influence a tie, and so yield utility $\leq U_k^*$. To see $-k$ cannot profitably deviate to $s_{-k} \leq 0$, observe all such policies result in either $(0, y_{-k})$ or $(0, 0)$ when $s_k \leq 0$ (since the DM's other choices are $(0, 0)$ and $\mathbb{O}$), and thus yield utility $\leq \max\left\{\bar{\Pi}_{-k}(0, 0; F), \bar{\Pi}_{-k}(0, y_{-k}; F)\right\}$ which is $\leq U_{-k}^*$. QED

A.1.2 Preliminary Observations about Equilibria

Proposition A.1 implies that all equilibria have a simple form. At least one developer (henceforth labelled $-k$) is always active – thus, competition not only strictly benefits the DM in expectation, but with probability 1. The other developer (henceforth labelled k) may also always be active ($F_k(0) = 0$), or be inactive with strictly positive probability ($F_k(0) > 0$). Inactivity may manifest as crafting the DM's ideal point with no quality $(0, 0)$, or as "position-taking" with more distant 0-quality policies that lose for sure ($s_k < 0$ and $s_k + y_k^2 = 0$). However, any equilibrium exhibiting the latter is payoff-equivalent to one exhibiting the former; we thus focus on the former for comparative statics.[22] When either developer i is active, he mixes smoothly over the ideologically-optimal policies $\left(s, \frac{x_i}{\alpha_i} F_{-i}(s)\right)$ with scores in a common mixing interval $[0, \bar{s}]$ according to the CDF $F_i(s)$.[23]

The differential equations characterizing the equilibrium score CDFs arise intuitively from the developers' indifference condition over $[0, \bar{s}]$. The left hand side is $i's$ net marginal cost of crafting a higher-score policy given a fixed probability $F_{-i}(s)$ of winning the contest; the developer pays marginal cost $\alpha_i > 1$ for sure, but with probability $F_{-i}(s)$ his policy is chosen and he enjoys a marginal benefit of 1 (because he values quality). The right hand side represents $i's$ marginal ideological benefit of increasing his score. Doing so increases by $f_{-i}(s)$ the probability that his policy wins, which changes the ideological outcome from his opponent's optimal ideology $y_{-i}^*(s)$ at score s to his own optimal ideology $y_i^*(s)$ at score s.

[22]Profiles with "position-taking" are equilibria if the position-taking does not invite a deviation by $-k$ to negative scores; whether this is the case depends on k's score-CDF below 0 and the DM's outside options $\mathbb{O}$. When $(0, 0) \in \mathbb{O}$ the necessary conditions are also sufficient.

[23]Technically, the proposition does not state that the support interval is also bounded ($\bar{s} < \infty$), but this is later shown indirectly through the analytical equilibrium derivation.

B Closed Form Equilibrium Characterization

The first and most critical step in generating a unique closed form equilibrium characterization and analytically examining its properties is to use the coupled system of differential equations that characterize any pair of equilibrium score CDFs $(F_L(s), F_R(s))$ to derive a simple *functional relationship* that must hold between them.

Lemma B.1. *In any SPNE,* $\epsilon_i\left(F_{-i}\left(s\right)\right) = \epsilon_{-i}\left(F_i\left(s\right)\right)\ \forall s \geq 0$*, where*

$$\epsilon_i\left(p\right) = \int_p^1 \frac{|x_i|}{\alpha_i - q} dq = |x_i| \log\left(\frac{\alpha_i - p}{\alpha_i - 1}\right)$$

Proof: Rearranging the differential equation in score optimality yields $\frac{f_{-i}(s)\cdot|x_i|}{\alpha_i - F_{-i}(s)} = \frac{f_i(s)\cdot|x_{-i}|}{\alpha_{-i} - F_i(s)}$ $\forall s \in [0, \bar{s}] \to \int_s^{\bar{s}} \frac{f_{-i}(s)\cdot|x_i|}{\alpha_i - F_{-i}(s)} ds = \int_s^{\bar{s}} \frac{f_i(s)\cdot|x_{-i}|}{\alpha_{-i} - F_i(s)} ds\ \forall s \in [0, \bar{s}]$; a change of variables and the boundary condition $F_i\left(\bar{s}\right) = 1$ yields $\int_s^{\bar{s}} \frac{f_{-i}(s)\cdot|x_i|}{\alpha_i - F_{-i}(s)} ds = \int_{F_{-i}(s)}^1 \frac{|x_i|}{\alpha_i - q} dq = \epsilon_i\left(F_{-i}\left(s\right)\right)$. The relationship holds trivially for $s > \bar{s}$. QED

We refer to the property in Lemma B.1 as the *engagement equality.* To see why, observe that the decreasing function $\epsilon_i\left(p\right)$ captures i's relative willingness to deviate from a policy that wins with probability p to one that wins for sure (since the marginal ideological benefit of moving the ideological outcome in his direction is $|x_i|$, and the net marginal cost of increasing score on a policy winning the contest with probability q is $\alpha_i - q$). We call this function i's *engagement at probability* p. The engagement equality $\epsilon_i\left(F_{-i}\left(s\right)\right) = \epsilon_{-i}\left(F_i\left(s\right)\right)$ states that at every score $s \geq 0$ both developers must be equally engaged given the resulting probabilities of winning the contest, and therefore equally willing to deviate to the maximum score $\bar{s}$. It is easily verified that $\epsilon_i\left(1\right) = 0\ \forall i$ and $\epsilon_i\left(p\right)$ is strictly increasing in $|x_i|$ and decreasing in $\alpha_i\ \forall p \in [0, 1)$.

Usefully, the engagement equality implies a simple functional relationship between the developers' score CDFs that must hold in equilibrium regardless of their exact values. Letting

$$p_i\left(\epsilon\right) = \alpha_i - \left(\alpha_i - 1\right) e^{\frac{\epsilon}{|x_i|}}$$

denote the inverse of $\epsilon_i(p)$ (which is decreasing in p, increasing in $|x_i|$, and decreasing in α_i) equilibrium then requires that $F_i(s) = p_{-i}(\epsilon_i(F_{-i}(s)))\ \forall s \in [0, \bar{s}]$.

B.1 Identity of developer k and probabilities of participation

We first use the engagement equality to derive the identity of the sometimes-inactive developer k and the probability $F_k(0)$ that he is sometimes inactive, and perform comparative statics on $F_k(0)$.

Proposition B.1. *In equilibrium* $k \in \arg\min_i \{\epsilon_i(0)\}$ *and*

$$F_k(0) = p_{-k}(\epsilon_k(0)) = \alpha_{-k} - (\alpha_{-k} - 1)\left(\frac{\alpha_k}{\alpha_k - 1}\right)^{\left|\frac{x_k}{x_{-k}}\right|}.$$

The probability k is inactive $F_k(0)$ is decreasing in his distance from the DM $|x_k|$ and his opponent's quality costs α_{-k}, and increasing in his opponent's distance from the DM $|x_{-k}|$ and his own quality costs α_k. In addition, $\lim_{|x_k|\to 0}\{F_k(0)\} = \lim_{|x_{-k}|\to\infty}\{F_k(0)\} = \lim_{\alpha_k\to\infty}\{F_k(0)\} = \lim_{\alpha_{-k}\to 1}\{F_k(0)\} = 1$.

Proof: Suppose $\epsilon_k(0) < \epsilon_{-k}(0)$; then $F_k(0) = 0$ and the engagement equality would imply $F_{-k}(0) < 0$, a contradiction. Since $F_i(0) = 0$ for some i we must have $F_{-k}(0) = 0$ and $F_k(0) = p_{-k}(\epsilon_k(0)) > 0$. Comparative statics and limit statements follow from previous observations on $\epsilon_i(\cdot)$ and $p_i(\cdot)$. QED

The sometimes-inactive developer is thus the one with the lowest engagement at probability 0 – that is, who is least willing to participate in the contest entirely.

B.2 Equilibrium Score CDFs

With the engagement equality and the identity of the sometimes-inactive developer k we may next characterize the equilibrium score CDFs $F_i(s)$ satisfying Proposition A.1, which are shown constructively to be unique.

Proposition B.2. *The unique score CDFs over $s \geq 0$ satisfying Proposition A.1 are* $F_i(s) =$

$p_{-i}(\epsilon(s))\ \forall i$, where $\epsilon(s)$ is the inverse of

$$s(\epsilon) = 2\int_{\epsilon}^{\epsilon_k(0)} \sum_j \frac{|x_j|}{\alpha_j} p_j(\hat{\epsilon})\, d\hat{\epsilon}.$$

The inverse score CDFs are $s_i(F_i) = s(\epsilon_{-i}(F_i))\ \forall i$, and the score targetted at each ideology is $s\left(\epsilon_i\left(\frac{y}{x_i/\alpha_i}\right)\right)$. The function $s(\epsilon)$ is strictly increasing in x_i and strictly decreasing in α_i $\forall \epsilon \in [0, \epsilon_k(0))$, and the maximum score is $\bar{s} = s(0)$.

Increasing a developer's extremism $|x_i|$ or decreasing his costs α_i first-order stochastically increases his own score CDF, but has ambiguous effects on his opponent's score CDF.

Proof: From the engagement equality $\epsilon_i(F_{-i}(s)) = \int_{F_{-i}(s)}^{1} \frac{|x_i|}{\alpha_i - q} dq = \epsilon(s)\ \forall i, s$ for some $\epsilon(s)$. We characterize the unique $\epsilon(s)$ implying score CDFs $F_i(s) = p_{-i}(\epsilon(s))$ and optimal ideologies $y_i(s) = \frac{x_i}{\alpha_i} p_i(\epsilon(s))$ that satisfy score optimality. First observe that $\epsilon'(s) = f_i(s)\epsilon'_{-i}(F_i(s)) = -\frac{f_{-i}(s)|x_i|}{\alpha_i - F_{-i}(s)}$. Next the differential equations may be rewritten as $\frac{\alpha_i - F_{-i}(s)}{f_{-i}(s)\cdot|x_i|} = 2\sum_j y_j(s)$. Substituting the preceding observations into both sides yields $\frac{1}{\epsilon'(s)} = -2\sum_j \frac{x_j}{\alpha_j} p_j(\epsilon(s))$, and rewriting in terms of the inverse $s(\epsilon)$ yields $s'(\epsilon) = -2\sum_j \frac{|x_j|}{\alpha_j} p_j(\epsilon)$. Lastly $\epsilon_k(F_{-k}(s)) = \epsilon(s)$ and $F_{-k}(0) = 0$ imply the boundary condition $s(\epsilon_k(0)) = 0$ so $s(\epsilon) = \int_{\epsilon}^{\epsilon_k(0)} -s'(\hat{\epsilon})\, d\hat{\epsilon} = 2\int_{\epsilon}^{\epsilon_k(0)} \sum_j \frac{|x_j|}{\alpha_j} p_j(\hat{\epsilon})\, d\hat{\epsilon}$. Now $s(\epsilon)$ is increasing in $|x_i|$ and decreasing in α_i given previous observations about $p_j(\hat{\epsilon})$. QED

The maximum score $\bar{s}$ thus changes continuously with the parameters of both developers even when one is dominant. This contrasts with the standard 2-player all pay contest, where the mixing interval is unaffected by the parameters of the stronger player. Increasing a developer's extremism $|x_i|$ or decreasing his costs α_i first-order stochastically increases his own score CDF, but has ambiguous effects on his opponent's score CDF. To see this, suppose that the always-active developer $-k$ becomes even more extreme or able. Then his opponent k becomes less likely to be active, but also the range of scores $[0, \bar{s}]$ over which he mixes *when* he is active increases. He thus has a higher probability of crafting a very high-score policy, even while he is simultaneously less likely to enter the contest.

B.3 Derivation of Strategies in Proposition 1

Finally we transform the preceding characterization of ideologically optimal policies and equilibrium score CDFs into the more intuitive characterization of equilibrium strategies provided in main text Proposition 1.

First, recall that a policy (s, y) has quality $q = \frac{y^2+s}{\lambda}$. Next, when a developer crafts a policy that is distance δ from the decisionmaker, its ideology is $i\delta$; consequently, the score at which developer i crafts policy $i\delta$ is $s\left(\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right)$ by Proposition B.2. Combining the preceding, the quality associated with a policy that is distance δ from the decisionmaker is $q_i(\delta) = \frac{(i\delta)^2+s\left(\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right)}{\lambda}$, which simplifies to the expression in the proposition.

Next, the probability distribution over the extremism of each developer's policy can be simply derived from the engagement equality as follows.

Proposition B.3. *Let $G_i(y) = \Pr(|y_i| \le \delta)$ denote the probability that i's policy is closer to the DM than δ. Then*

$$G_i(\delta) = p_{-i}\left(\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right) = \alpha_{-i} - (\alpha_{-i} - 1)\left(\frac{x_i - i\delta}{x_i - x_i/\alpha_i}\right)^{\left|\frac{x_i}{x_{-i}}\right|},$$

which is first-order stochastically increasing in i's extremism $|x_i|$, decreasing in his costs α_i, decreasing in his opponent's extremism $|x_{-i}|$, and increasing in his opponent's costs α_i.

Proof: Developer i's ideology at score s is $y_i^*(s) = \frac{x_i}{\alpha_i}F_{-i}(s)$ (from ideological optimality), so $F_{-i}(s_i^*(y)) = \frac{y}{x_i/\alpha_i}$ where $s_i^*(y)$ is the inverse of $y_i^*(s)$. That is, the probability $-i$ crafts a policy with score $\le s_i^*(y)$ is $\frac{y}{x_i/\alpha_i}$. Now the probability $G(\delta)$ that i crafts a policy closer to the DM than y is $F_i(s_i^*(i\delta))$, which is $= p_{-i}(\epsilon_i(F_{-i}(s_i^*(i\delta)))) = p_{-i}\left(\epsilon_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right)$ from the engagement equality. Comparative statics are straightforward. QED

C Additional Quantities and Comparative Statics

In this section we calculate and examine the general properties of additional equilibrium quantities; these propositions form the basis for the main-text propositions that study properties of the model in the special cases of pure asymmetric extremism and pure asymmetric ability.

C.1 Probabilities of Victory

We next use the engagement equality to derive the developers' probabilities of victory.

Proposition C.1. *In equilibrium the probability developer k loses the contest is*

$$\int_0^1 p_{-k}\left(\epsilon_k\left(p\right)\right) dp = \int_0^1 \left(\alpha_{-k} - \left(\alpha_{-k} - 1\right)\left(\frac{\alpha_k - p}{\alpha_k - 1}\right)^{\left|\frac{x_k}{x_{-k}}\right|}\right) dp$$

which is decreasing in his distance from the DM $|x_k|$ and his opponent's quality costs α_{-k}, and increasing in his opponent's distance from the DM $|x_{-k}|$ and his own quality costs α_k.

Proof: The probability k loses the contest is $\int_0^{\bar{s}} f_{-k}\left(s\right) F_k\left(s\right) ds$; applying the engagement equality this is $\int_0^{\bar{s}} p_{-k}\left(\epsilon_k\left(F_{-k}\left(s\right)\right)\right) f_{-k}\left(s\right) ds$, and applying a change of variables of $F_{-k}\left(s\right)$ for p (recalling $F_{-k}(0) = 0$) yields the result. QED

The probability k loses thus obeys the same comparative statics as his probability of inactivity. Somewhat paradoxically, he becomes *less* likely to win when his preferences are closer to the DM or his opponent's are more distant. More intuitively, he becomes more likely to win if he is more able or his opponent less able.

C.2 Conditions for First-Order Stochastic Dominance

In the standard asymmetric two-player all-pay contest there is always an unambiguously weaker player, who makes bids that are first-order stochastically worse for the DM. In the present contest, in contrast, there may be no unambiguously weaker player in this sense.

Proposition C.2. *Developer i is dominated ($F_{-i}\left(s\right) < F_i\left(s\right)\ \forall s \in \left(0, \bar{s}\right)$) i.f.f. he is less engaged at every probability p ($\epsilon_i\left(p\right) < \epsilon_{-i}\left(p\right)\ \forall p \in \left(0, 1\right)$). Equivalently, he is dominated i.f.f. both $\int_0^1 \frac{|x_i|}{\alpha_i - q} dq \leq \int_0^1 \frac{|x_{-i}|}{\alpha_{-i} - q} dq$ and $\frac{|x_i|}{\alpha_i - 1} \leq \frac{|x_{-i}|}{\alpha_{-i} - 1}$, where the latter condition is stronger than the former i.f.f. i has a cost advantage.*

Proof: Lemma B.1 and the engagement function $\epsilon_i(p)$ strictly decreasing when $p \in [0,1)$ immediately implies $sign(\epsilon_{-k}(F_{-k}(s)) - \epsilon_k(F_{-k}(s))) = sign(F_k(s) - F_{-k}(s))\ \forall s \in [0,\bar{s})$, which straightforwardly yields the first statement. Now let $\delta(p) = \epsilon_{-k}(p) - \epsilon_k(p)$, so $\delta(0) \geq 0 = \delta(1)$. We argue $\delta'(1) \leq 0$ is necessary and sufficient. For necessity, $\delta'(1) > 0 = \delta(1) \to \delta(p) < 0$ in a neighborhood below 1. For sufficiency, it is easily verified that $\delta'(p) = \frac{|x_k|}{\alpha_k - p} - \frac{|x_{-k}|}{\alpha_{-k} - p}$ crosses 0 at most once when the developers are asymmetric; thus $\delta(0) \geq 0 = \delta(1) \geq \delta'(0)$ implies $\delta(p)$ strictly quasi-concave over $[0,1]$ and $\delta(p) > \min\{\delta(0), \delta(1)\} \geq 0$ for $p \in (0,1)$.

We last argue $\delta(0) \geq 0$ and $\alpha_k > \alpha_{-k} \to \delta'(1) < 0$. Observe that $\alpha_k < \alpha_k$ and $\delta'(0) = \frac{x_k}{\alpha_k} - \frac{x_{-k}}{\alpha_{-k}} \leq 0 \to \delta'(1) = \frac{|x_k|}{\alpha_k}\left(\frac{1}{1-1/\alpha_k}\right) - \frac{|x_{-k}|}{\alpha_{-k}}\left(\frac{1}{1-1/\alpha_{-k}}\right) < 0$. If $\delta'(0) \leq 0$ we are done; if $\delta'(0) > 0$ then $\delta'(1) \geq 0 \to \delta'(p) > 0\ \forall p \in [0,1) \to \delta(1) > 0$, a contradiction. QED

Clearly, a developer k who is *both* less extreme ($|x_k| \leq |x_{-k}|$) and less able ($\alpha_k \geq \alpha_{-k}$) (with one strict) satisfies both conditions and is therefore dominated. However, when one developer is more extreme while the other is more able, then lower engagement at probability 0 is necessary but not sufficient for the more able developer to be dominated.

C.3 Developer Payoffs

Using Proposition B.2, the developers' equilibrium payoffs are as follows.

Proposition C.3. *Developer i's equilibrium utility is* $\Pi_i^*(\bar{s}; F^*) = -\left(1 - \frac{1}{\alpha_i}\right) x_i^2 - (\alpha_i - 1)\bar{s}$, *which is decreasing in his own costs* α_i *as well as either players' extremism* $|x_j|\,\forall j$, *and increasing in his opponent's costs* α_{-i}.

Proof: A developer's equilibrium utility is straightforward since $(\bar{s}, y_i^*(\bar{s}))$ is in the support of their strategy and wins for sure. Comparative statics of a developer i's parameters on his opponent $-i$'s utility, as well as of x_i on his own utility, follow immediately from previously-shown statics on $\bar{s} = s(\epsilon)$. Taking the derivative with respect to α_i, substituting in $\frac{\partial}{\partial \alpha_i}\left(\frac{p_i(\epsilon)}{\alpha_i}\right) = \frac{x_i p_i'(\epsilon)}{(\alpha_i - 1)\alpha_i^2}$, $\frac{\partial \epsilon_k(0)}{\partial \alpha_k} = -\frac{|x_k|}{\alpha_k(\alpha_k - 1)}$, $-\frac{p_i'(\epsilon) x_i}{\alpha_i - p_i(\epsilon)} = 1$, performing a change of variables, and rearranging the expression yields

$-\mathbf{1}_{i=k}\cdot 2\int_0^{\epsilon_k(0)} \frac{|x_{-k}|}{\alpha_{-k}} \left(p_{-k}(\epsilon) - \left(\alpha_k \log\left(\frac{\alpha_k}{\alpha_k-1}\right)\right)^{-1} p_{-k}(\epsilon_k(0))\right) d\epsilon - \left(\frac{|x_i|}{\alpha_i}\right)^2 \left(1 + 2\int_{p_i(\epsilon_k(0))}^1 \left(\frac{\alpha_i p}{\alpha_i - p} - 1\right)\right) dp.$ The first term is negative since $p_{-k}(\epsilon) > p_{-k}(\epsilon_k(0))$ for $\epsilon < \epsilon_k(0)$ and $\frac{1}{\alpha_k} < \int_0^1 \frac{1}{\alpha_k - p} dp = \log\left(\frac{\alpha_k}{\alpha_k - 1}\right)$. The second term is also negative since $1 + 2\int_{p_i(\epsilon_k(0))}^1 \left(\frac{\alpha_i p}{\alpha_i - p} - 1\right) > (1 - p_i(\epsilon_k(0))) + 2\int_{p_i(\epsilon_k(0))}^1 \left(\frac{\alpha_i p}{\alpha_i} - 1\right) = \int_{p_i(\epsilon_k(0))}^1 (2p-1)\, dp \geq 0$. QED

A developer's equilibrium utility has two components. The first $-\left(1 - \frac{1}{\alpha_i}\right) x_i^2$ is his utility if he could craft a policy as a "monopolist" (and the DM's outside option included $(0,0)$). The second $-(\alpha_i - 1)\bar{s}$ is the cost generated by competition, which forces him to craft a policy that leaves the DM strictly better off than the best "free" policy $(0,0)$ in order to maintain influence. This competition cost is increasing in i's marginal cost α_i of generating quality (holding $\bar{s}$ fixed) as well as the maximum score $\bar{s}$, which in turn is increasing in both developers' ideological extremism and decreasing in their costs everywhere in the parameter space. A developer is thus strictly harmed when his competitor becomes more extreme or able. This is distinct from all pay contests without spillovers (Siegel (2009)), where the equilibrium utility of the "sometimes inactive" player is pinned at his fixed value for losing.

A developer also worse off when his *own* preference become more distant from the decisionmaker. Finally, a developer is worse off when his costs of producing quality increase – even though there is a countervailing effect of reducing the intensity of competition (and indeed, the competition cost $(\alpha_i - 1)\bar{s}$ alone is not generically monotonic in α_i).

C.4 Decisionmaker Payoffs

Lastly, again using Proposition B.2 the DM's equilibrium utility and the developers' average scores (which bound the DM's utility from below) are as follows.

Proposition C.4. *The DM's equilibrium utility is* $U_{DM}^* = \int_{\epsilon_k(0)}^0 s(\epsilon) \cdot \frac{\partial}{\partial \epsilon}\left(\prod_j p_j(\epsilon)\right) d\epsilon =$

$$2\int_0^{\epsilon_k(0)} \left(1 - \prod_j p_j(\epsilon)\right) \cdot \left(\sum_j \frac{|x_j|}{\alpha_j} p_j(\epsilon)\right) d\epsilon$$

Developer i's average score is $E[s_i] = \int_{\epsilon_k(0)}^{0} s(\epsilon) \cdot \frac{\partial}{\partial \epsilon}(p_{-i}(\epsilon))\, d\epsilon =$

$$2\int_0^{\epsilon_k(0)} (1 - p_{-i}(\epsilon)) \cdot \left(\sum_j \frac{|x_j|}{\alpha_j} p_j(\epsilon)\right) d\epsilon$$

Proof: $F_i(s)\,F_{-i}(s)$ is the CDF of $\max\{s_i, s_{-i}\}$ so the DM's utility is $\int_0^{\bar{s}} s \cdot \frac{\partial}{\partial s}\left(\prod_j F_j(s)\right) ds =$ $\int_0^{\bar{s}} s \cdot \frac{\partial}{\partial s}\left(\prod_j p_j(\epsilon(s))\right) ds$. A change of variables from s to ϵ yields the first expression and integration by parts and rearranging yields the second. Nearly identical steps yield i's average score. QED

Direct comparative statics on the DM's utility U^*_{DM} are difficult because changing a developer's parameters has mixed effects on his opponent's score CDF. We thus consider two special cases; breaking symmetry, and the limiting cases of extreme imbalance.

Proposition C.5. *When the developers are symmetric ($|x_i| = |x_{-i}|$ and $\alpha_i = \alpha_{-i}$), the DM's utility is locally increasing eithers' extremism or ability.*

Proof: First differentiating the DM's utility U^*_{DM} with respect to $|x_{-k}|$ and applying symmetry yields $\frac{2}{\alpha}\int_0^{\epsilon(0)} \left(\left(1 - 3(p(\epsilon))^2\right) \cdot x\frac{\partial p(\epsilon)}{\partial x} + \left(1 - (p(\epsilon))^2\right) p(\epsilon)\right) d\epsilon$ which is $\geq 2\frac{x}{\alpha}\int_0^{\epsilon(0)} \left(1 - 3(p(\epsilon))^2\right) \frac{\partial p(\epsilon)}{\partial x} d\epsilon$. Now substituting $\frac{\partial p(\epsilon)}{\partial x} = -\log\left(\frac{\alpha - p(\epsilon)}{\alpha - 1}\right) p'(\epsilon)$ and a change of variables yields $2\frac{x}{\alpha}\int_0^1 \left(1 - 3p^2\right) \log\left(\frac{\alpha - p}{\alpha - 1}\right) dp =$ $2\frac{x}{\alpha}\int_0^1 \left(\frac{p - p^3}{\alpha - p}\right) dp > 0$. Next differentiating U^*_{DM} w.r.t. α_{-k} and applying symmetry yields $2x\int_0^{\epsilon(0)} \left(\left(1 - (p(\epsilon))^2\right) \frac{\partial}{\partial \alpha}\left(\frac{p(\epsilon)}{\alpha}\right) - \frac{2}{\alpha}(p(\epsilon))^2 \frac{\partial p(\epsilon)}{\partial \alpha}\right) d\epsilon$. Finally, substituting $\frac{\partial}{\partial \alpha}\left(\frac{p(\epsilon)}{\alpha}\right) = \frac{x}{(\alpha-1)\alpha^2} p'(\epsilon)$, $\frac{\partial p(\epsilon)}{\partial \alpha} = -\left(\frac{1 - p(\epsilon)}{\alpha - 1}\right)$, $-\frac{p'(\epsilon)x}{\alpha - p(\epsilon)} = 1$, rearranging the expression, and another change of variables yields $\frac{2x^2}{(\alpha-1)\alpha^2}\int_0^1 \left(2p^2\left(\frac{\alpha - \alpha p}{\alpha - p}\right) - \left(1 - p^2\right)\right) dp < 0$. QED

The DM thus strictly benefits locally if developers between the players is broken by one becoming more extreme or able – even though the other also becomes less active. The effect of extreme asymmetries is as follows.

Proposition C.6. *The DM's utility exhibits the following limiting behavior*

$$0 = \lim_{\alpha_i \to \infty} U^*_{DM} = \lim_{x_i \to 0} U^*_{DM} < \lim_{x_i \to \infty} U^*_{DM} = \infty$$

and $\lim_{\alpha_i \to 1} U^*_{DM} = 2x_k \int_0^1 \left(\frac{1-p}{\alpha_k - p}\right) \cdot \left(\frac{x_k}{\alpha_k} p + x_{-k}\right) dp$, *which is strictly increasing in* x_k *and strictly decreasing in* α_k.

Proof: Observe that $E[s_{-k}] \leq U^*_{DM} \leq \bar{s}$. For the first two limiting statements it is easily verified that $\bar{s} \to 0$ as $\alpha_k \to \infty$ or $x_k \to 0$. For the third limiting statement observe that $E[s_{-k}] \geq \frac{|x_{-k}|}{\alpha_{-k}} p_{-k}(\epsilon_k(0)) \cdot 2\int_0^{\epsilon_k(0)} (1 - p_k(\epsilon))\, d\epsilon$ which $\to \infty$ as $|x_{-k}| \to \infty$ since the first term $\to \infty$ and the remaining terms are non-decreasing. For the fourth limiting statement, using the definition in Proposition C.4 and that $\lim_{\alpha_{-k} \to 1} \{p_{-k}(\epsilon)\} = 1\ \forall \epsilon \in [0, \epsilon_k(0)]$ yields a limit of $2\int_0^{\epsilon_k(0)} (1 - p_k(\epsilon)) \cdot \left(\frac{x_k}{\alpha_k} p_k(\epsilon) + x_{-k}\right) d\epsilon$. Observing that $-\frac{p'_k(\epsilon) x_k}{\alpha_k - p_k(\epsilon)} = 1$, substituting into the expression, and applying a change of variables yields the expression, which straightforwardly obeys the stated comparative statics. QED

If an extreme imbalance is the result of one developer's incompetence *or* ideological moderation, the DM's utility approaches 0, her utility if $-i$ were a "monopolist" (and the DM's outside options included $(0, 0)$). (Developer $-i$'s utility also approaches his utility if he were a monopolist). However, if extreme imbalance is the result of one developer's greater ability to produce quality (specifically, if his marginal cost of producing quality approaches its intrinsic value), then the DM's utility is bounded away from 0. In this case the DM strictly benefits from the potential for competition, even though actual competition is almost never observed (since $F_{-i}(0) = F_k(0)$ approaches 1). Finally, unilateral ideological extremism benefits the decisionmaker in a strong sense; the DM can achieve arbitrarily high utility with a developer whose preferences are sufficiently *distant* from her own.

D Main Text Propositions

In this Appendix we describe where to locate the results collated in main text Propositions 2-8 in the general analysis contained in Appendices B-C.

Proposition 2 To see the first bullet point, observe that Proposition B.1 on activity implies that the moderate is the sometimes-inactive developer k and is inactive with strictly positive probability.

To see the second bullet point, first observe that the greater (first-order stochastic) extremism of the extremist's policy is an implication of the ideology comparative statics stated in Proposition B.3, which states that as a developer becomes unilaterally more extreme his policy's ideology becomes more extreme and his opponent's policy's ideology simultaneously becomes more moderate. Next observe that the greater (first-order stochastic) overall appeal to the decisionmaker of the extremist's policy follows from the necessary and sufficient conditions for score-dominance in Proposition C.2 – a developer being more extreme and able with at least one strict is a sufficient condition for score dominance. Finally, the statement on quality is an immediate implication of the extremist crafting a more ideologically extreme but also higher score policy (first order stochastically).

Lastly, the third bullet point is an immediate implication of score-dominance.

Proposition 3 The first bullet point under both "own strategy" and "opponent's strategy" follow from Proposition B.1 on activity. The second bullet point under both "own strategy" and "opponent's strategy" follow from Proposition B.3 on ideology. The third bullet point under "own strategy" is a joint implication of the ideology comparative statics in Proposition B.3 and the comparative statics on "own score" in Proposition B.2. The third bullet point under "opponent strategy" also follows from Proposition B.2 and the subsequent discussion.

Proposition 4 Follows immediately from Proposition C.3 characterizing the developers' payoffs.

Proposition 5 The first bullet point just restates Proposition C.5. The second bullet point follows from Proposition B.1 (on activity) and Proposition C.6 (on the decisionmaker's welfare).

Proposition 6 Follows from Propositions B.1, B.3, and C.2 according to a nearly identical argument as in the proof of Proposition 2.

Proposition 7 The first statement follows from Propositions B.1-B.3 by a nearly identical argument as in the proof of Proposition 3. The second statement follows immediately from Proposition C.3.

Proposition 8 The first bullet point just restates Proposition C.5. The second bullet point follows from Proposition B.1 (on activity) and Proposition C.6 (on the decisionmaker's welfare).

E Model Variants and Robustness

In this Appendix we consider several variants of the model to examine robustness.

E.1 Fixed Policies

In this subsection we consider a variant of the model in which the developers are constrained to target a fixed ideology $y_i = iy$ with $y > 0$ and the decisionmaker has no outside option. Borrowing from the main analysis, it is easily verified that a developer i's expected utility when he crafts a policy (s, y_i) where his opponent has no atom or a tie would be broken in his favor is equal to:

$$-\left(\alpha_i - F_{-i}\left(s\right)\right)s + F_{-i}\left(s\right)\cdot\left(2x_i y_i - x_i^2\right) - \alpha_i\left(y_i\right)^2 + \int_s^{\bar{s}}\left(s_{-i} + 2x_i y_{-i} - x_i^2\right)f_{-i}\left(s_{-i}\right)ds_{-i} \quad \text{(E.1)}$$

Necessary and Sufficient Conditions for Equilibrium Using a similar series of steps as in Appendix A.1 it is straightforward to show that any equilibrium must take a similar form as in the main model; a pair of score CDFs $F_i\left(s\right)$ $\forall i$ that are continuously increasing over a common interval $[\underline{s}, \bar{s}]$ with $\underline{s} = -y^2$ satisfying (i) $F_k\left(\underline{s}\right) > 0$ for at most one k, (ii) $F_i\left(\bar{s}\right) = 1$ $\forall i$, and (iii) the following equality (which is eqn. E.1 differentiated w.r.t. s and set equal to 0) $\forall s \in [\underline{s}, \bar{s}]$ and $\forall i$:

$$\alpha_i - F_{-i}\left(s\right) = f_{-i}\left(s\right)\cdot 4\left|x_i\right|y \quad \text{(E.2)}$$

Deriving Equilibrium This system may be solved directly with conventional methods since the score CDFs are only dependent via the boundary conditions, but we do so using the methods and notation of the main model so as to preserve comparability with main model equilibrium quantities.

First, rearranging we have $\frac{f_{-i}(s)|x_i|}{\alpha_i - F_{-i}(s)} = \frac{1}{4y}$ $\forall i$, implying $\frac{f_{-i}(s)|x_i|}{\alpha_i - F_{-i}(s)} = \frac{f_i(s)|x_{-i}|}{\alpha_{-i} - F_i(s)}$, so applying the boundary condition $F_i\left(\bar{s}\right) = 1$ yields the engagement equality $\epsilon_i\left(F_{-i}\left(s\right)\right) = \epsilon_{-i}\left(F_i\left(s\right)\right)$ $\forall s \geq \underline{s}$ with

$\epsilon_i(p) = \int_p^1 \frac{|x_i|}{\alpha_i - q}$ as in the main model. Thus it remains true that the sometimes-inactive developer $k \in \arg\min_i \{\epsilon_i(0)\}$ and $F_k(\underline{s}) = p_{-k}(\epsilon_k(0))$.

To solve for the engagement function $\epsilon(s) = \epsilon_i(F_{-i}(s)) = \epsilon_{-i}(F_i(s))$ satisfying the boundary conditions $\epsilon(\underline{s}) = \epsilon_k(0)$ and $\epsilon(\bar{s}) = 0$, recall that $\epsilon'(s) = -\frac{f_{-i}(s)|x_i|}{\alpha_i - F_{-i}(s)}$ so that $-\frac{f_{-i}(s)|x_i|}{\alpha_i - F_{-i}(s)} = -\frac{1}{4y} \iff \epsilon'(s) = -\frac{1}{4y}$. We then have $\epsilon(s) = \frac{\bar{s}-s}{4y}$ with $\underline{s} = -y^2$ and $\bar{s} = \underline{s} + 4y \cdot \epsilon_k(0)$, and the following characterization of the unique equilibrium score CDFs:

$$F_k(s) = p_{-k}\left(\epsilon_k(0) - \frac{s-\underline{s}}{4y}\right) \quad \text{and} \quad F_{-k}(s) = p_k\left(\epsilon_k(0) - \frac{s-\underline{s}}{4y}\right) \tag{E.3}$$

Proof of Proposition 9 Although this variant does not fit precisely into the "all-pay contest" framework of Siegel (2009) because the developers still care about the quality of the policy they lose to, it nevertheless exhibits most of the characteristic properties of the standard two player asymmetric all pay contest (in contrast to the main model); this is because $\epsilon(s) = \frac{4y \cdot \epsilon_k(0) - (s-\underline{s})}{4y}$ depends only on the characteristics of the weaker player k.

First, evaluating eqn. E.1 for the weaker player k at $s = \bar{s}$ yields her equilibrium expected utility

$$-(\alpha_k - 1)(\underline{s} + 4y \cdot \epsilon_k(0)) + 2|x_k| y - x_k^2 - \alpha_k y^2$$

Thus, as in the standard 2-player asymmetric all pay contest, the equilibrium utility of the weaker player k is *invariant to* the characteristics of the stronger player.

Second, it is clear from the characterization of the equilibrium score CDF $F_{-k}(s)$ of the stronger player $-k$ in eqn. E.3 that the score CDF of the stronger player is *invariant to her own characteristics* (x_{-k}, α_{-k}), and depends entirely on the characteristics of her weaker competitor (x_k, α_k).

Third, it is easily verified from the characterization of the equilibrium score CDF $F_k(s)$ of the weaker player k in eqn. E.3 that at any $s \in [\underline{s}, \bar{s})$ we have $F_k(s)$ strictly increasing in x_{-k} and strictly decreasing in α_{-k}; that is, the weaker player's score CDF is first order stochastically decreasing in the stronger player's extremism and increasing in the weaker player's cost. It is also easily verified that $F_k(s) \to 1$ $\forall s \in [0, \bar{s})$ as $x_{-k} \to \infty$ or $\alpha_{-k} \to 1$. Thus, the discouragement effect is present in the model, and as one developer becomes arbitrarily extreme or capable the other developer becomes almost always inactive. Combining these observations with the previous observation that the score CDF of the stronger player $-k$ is invariant to her own characteristics immediately yields that the decisionmaker is harmed by asymmetries, since $-k$ becoming more extreme or capable decreases k's score CDF but has no effect on her own.

Finally, when the developers are symmetrically extreme and capable ($|x_{-1}| = |x_1| = x$ and $\alpha_{-1} = \alpha_1 = \alpha$), the unique symmetric equilibrium score CDF $F(s)$ is:

$$F(s) = p\left(\epsilon(0) - \frac{s - \underline{s}}{4y}\right) \text{ with } \epsilon(p) = \int_p^1 \frac{x}{\alpha - q} dq$$

which is clearly first-order stochastically increasing in extremism x and decreasing in costs α. Thus, this variant without spillovers exhibits the benefit of greater *symmetric extremism and ability* in Hirsch and Shotts (2015), but does not exhibit the benefit of greater *asymmetric extremism and ability* in the main model. QED

E.2 No Policy Spillovers

In this subsection we examine a variant of the model that lacks "rank order spillovers," in which the developers only care about policy *when they win*. For algebraic simplicity, we assume that if they lose they receive utility "as if" the policy $(0, 0)$ is implemented. Borrowing from the main analysis, it is easily verified that a developer's i's expected utility when he develops a policy (s, y) with $s \geq 0$ where either his opponent has no atom or a tie would be broken in his favor is equal to:

$$-(\alpha_i - F_{-i}(s))\, s + F_{-i}(s) \cdot 2x_i y - \alpha_i y^2 - x_i^2 \tag{E.4}$$

Necessary and Sufficient Conditions for Equilibrium Using a similar series of steps as in Appendix A.1 it is straightforward to show that any equilibrium must take an identical form as in the main model; a pair of score CDFs $F_i(s)$ $\forall i$ that are continuously increasing over a common interval $[0, \bar{s}]$ satisfying (i) $F_k(0) > 0$ for at most one k, (ii) $F_i(\bar{s}) = 1$ $\forall i$, (iii) $y_i(s) = F_{-i}(s) \frac{x_i}{\alpha_i}$, and (iv) the following expression (which is eqn. E.4 with $y_i(s)$ substituted in) constant $\forall s \in [0, \bar{s}]$ and $\forall i$:

$$-(\alpha_i - F_{-i}(s))\, s + [F_{-i}(s)]^2 \frac{x_i^2}{\alpha_i} - x_i^2 \tag{E.5}$$

Deriving Equilibrium For each developer any $s \in [0, \bar{s}]$ and $\bar{s}$ must yield the same utility, i.e.

$$-(\alpha_i - F_{-i}(s))\, s + [F_{-i}(s)]^2 \frac{x_i^2}{\alpha_i} - x_i^2 = -(\alpha_i - F_{-i}(s))\, s + [F_{-i}(s)]^2 \frac{x_i^2}{\alpha_i} - x_i^2$$

Applying $F_j(\bar{s}) = 1\ \forall i$ and simplifying yields the implicit characterization:

$$\frac{x_i^2}{\alpha_i(\alpha_i - 1)} \cdot \left(1 - [F_{-i}(s)]^2\right) = \bar{s} - \left(\frac{\alpha_i - F_{-i}(s)}{\alpha_i - 1}\right) \cdot s \tag{E.6}$$

It is easily verified that for a given value of $\bar{s}$ this expression uniquely defines a continuously increasing $F_{-i}(s)\ \forall s \in [0, \bar{s}]$; it remains only to identify the value of $\bar{s}$ that will satisfy the boundary condition $F_k(0) > 0$ for at most one k. Letting $\bar{s}_i = \frac{x_i^2}{\alpha_i(\alpha_i - 1)}$, it is easily verified that the boundary condition at 0 is satisfied if and only if $k \in \arg\min_i \left\{\frac{x_i^2}{\alpha_i(\alpha_i-1)}\right\}$ and $\bar{s} = \bar{s}_k$; thus equilibrium is unique. Finally, substituting into E.6 and simplifying yields a characterization of the unique score CDFs for the always active developer $-k$:

$$x_k^2 \cdot [F_{-k}(s)]^2 = \alpha_k(\alpha_k - F_{-k}(s)) \cdot s \tag{E.7}$$

and the sometimes inactive developer k:

$$x_{-k}^2 \cdot \left(1 - [F_k(s)]^2\right) = \alpha_{-k}\left((\alpha_{-k} - 1)\, \bar{s}_k - (\alpha_{-k} - F_k(s))\, s\right) \tag{E.8}$$

Proof of Proposition 10 Although the variant without spillovers does not fit precisely into the "all-pay contest" framework of Siegel (2009) due to its multidimensionality, it nevertheless exhibits most of the characteristic properties of the standard two player asymmetric all pay contest (in contrast to the main model). First, evaluating eqn. E.5 for the weaker player k (that is, the one with the lower value of $\bar{s}_i = \frac{x_i^2}{\alpha_i(\alpha_i-1)}$) at $s = 0$ yields her equilibrium expected utility $-x_k^2$. Thus, as in the standard 2-player asymmetric all pay contest, the equilibrium utility of the weaker player k is *invariant to* the characteristics of the stronger player.

Second, it is clear from the characterization of the equilibrium score CDF $F_{-k}(s)$ of the stronger player $-k$ in eqn. E.7 that the score CDF of the stronger player is *invariant to her own characteristics* (x_{-k}, α_{-k}), and depends entirely on the characteristics of her weaker competitor (x_k, α_k).

Third, it is easily verified from the implicit characterization of the equilibrium score CDF $F_k(s)$ of the weaker player k in eqn. E.8 that at any $s \in [0, \bar{s})$ we have $F_k(s)$ strictly increasing in x_{-k} and strictly decreasing in α_{-k}; that is, the weaker player's score CDF is first order stochastically decreasing in the stronger player's extremism and increasing in the weaker player's cost. It is also easily verified that $F_k(s) \to 1 \; \forall s \in [0, \bar{s})$ as $x_{-k} \to \infty$ or $\alpha_{-k} \to 1$. Thus, the discouragement effect is present in the model, and as one developer becomes arbitrarily extreme or capable the other developer becomes almost always inactive. Combining these observations with the previous observation that the score CDF of the stronger player $-k$ is invariant to her own characteristics immediately yields that the decisionmaker is harmed by asymmetries, since $-k$ becoming more extreme or capable decreases k's score CDF but has no effect on her own.

Finally, when the developers are symmetrically extreme and capable ($|x_{-1}| = |x_1| = x$ and $\alpha_{-1} = \alpha_1 = \alpha$), the unique symmetric equilibrium score CDF $F(s)$ is characterized by the equation:

$$x^2 \cdot [F(s)]^2 = \alpha (\alpha - F(s)) \cdot s$$

which is clearly first-order stochastically increasing in extremism x and decreasing in costs α. Thus, the variant without spillovers exhibits the benefit of greater *symmetric extremism and ability* in Hirsch and Shotts (2015), but does not exhibit the benefit of greater *asymmetric extremism and ability* in the main model. QED

E.3 Linear-Quadratic Preferences

In this subsection we consider a variant of the model in which the players' utilities are *linear* in ideology ($U_D(y, q) = \lambda q - |y|$ and $U_i(y, q) = \lambda q + x_i y$ for $i \in \{-1, 1\}$ where $|x_i| > 1$), while the costs of crafting quality are *quadratic* ($c_i(q) = \frac{a_i}{2} q^2$). With these functional forms a score-s policy with ideology y has quality $q = \frac{s+|y|}{\lambda}$, and developer i's expected utility when he crafts a policy (s, y) at a score where his opponent has no atom or a tie would be broken in his favor is equal to:

$$-\frac{\alpha_i}{2}(s+|y|)^2+F_{-i}(s)(s+|y|+x_iy)+\int_s^{\bar{s}}(s+|y_{-i}(s)|+x_iy_{-i}(s))f_{-i}(s_{-i}ds_{-i}), \tag{E.9}$$

where $\alpha_i = \frac{a_i}{\lambda^2}$.

Necessary and Sufficient Conditions for Equilibrium Using a similar series of steps as in Appendix A.1 it is straightforward to show that any equilibrium must take a similar form as in the main model; a pair of score CDFs $F_i(s)\ \forall i$ that are continuously increasing over a common interval $[0,\bar{s}]$ satisfying (i) $F_k(0)>0$ for at most one k, (ii) $F_i(\bar{s})=1\ \forall i$, (iii) $y_i(s)=i\cdot\max\left\{\frac{1+|x_i|}{\alpha_i}F_{-i}(s)-s,0\right\}$, and (iv) equation E.9 constant over $[0,\bar{s}]$ when $y_i(s)$ substituted in. To derive the required condition for (iv) we substitute $y_i(s)$ in, differentiate w.r.t. s, and set $=0$. Observe that we may also first differentiate w.r.t. s and then substitute in $y_i(s)$ (by the envelope theorem if $y_i(s)\neq 0$ and since $y_i'(s)=0$ if $y_i(s)=0$), so the required condition is then

$$\alpha_i(s+i\cdot y_i(s))-F_{-i}(s)=f_{-i}(s)\cdot((x_i+i)\cdot y_i(s)-(x_i-i)\cdot y_{-i}(s))$$

The model does not appear to admit a closed form solution in the asymmetric case; we derive symmetric equilibrium and compute asymmetric equilibria numerically.

Symmetric Model

We first analytically characterize equilibrium when developers are symmetric, $x_i=ix$ and $\alpha_i=\alpha$. Imposing symmetric parameters yields the system of differential equations:

$$\alpha(s+y(F_{-i}(s);s))-F_{-i}(s)=f_{-i}(s)\cdot[(x+1)\cdot y(F_{-i}(s);s)+(x-1)\cdot y(F_i(s);s)] \tag{E.10}$$

with $y(F;s)=\max\left\{\frac{1+x}{\alpha}F-s,0\right\}$. We conjecture a symmetric solution ($F_i(s)=F(s)\ \forall s,i$), which yields the single differential equation

$$\alpha(s+y(F(s),s))-F(s)=f(s)\cdot 2x\cdot y(F(s);s) \tag{E.11}$$

with $F(0) = 0$. We solve by conjecturing a particular form, solving, and verifying that it satisfies the necessary properties and boundary condition. Suppose $F(s) = \beta s$ and $\frac{1+x}{\alpha}\beta s - s \geq 0 \; \forall s \geq 0$ so that the optimal ideology $y(s) > 0 \; \forall s > 0$. Substituting into the differential equation yields:

$$\alpha \left(s + \left(\frac{1+x}{\alpha} \beta s - s \right) \right) - \beta s = \beta \cdot 2x \cdot \left(\frac{1+x}{\alpha} \beta s - s \right)$$

which holds $\forall s$ (since the terms cancel) provided $\beta = \frac{3\alpha}{2(1+x)}$. We last verify $\frac{1+x}{\alpha} F(s) - s = \frac{1+x}{\alpha} \left(\frac{3\alpha}{2(1+x)} \right) s - s = \frac{1}{2} s > 0 \; \forall s > 0$ which clearly holds, so this is indeed the solution. It is self evident that this variant exhibits the benefits of greater symmetric extremism and ability in Hirsch and Shotts (2015).

Asymmetric Model

We compute solutions numerically for the asymmetric model in Mathematica. For simplicity we restrict attention to parameter configurations in which the identity of the always-active developer $-k$ is unambiguous (because he is weakly more extreme and capable with at least one strict). The score CDFs are pinned down uniquely from the differential equations and lower boundary condition that $F_{-k}(0) = 0$ up to the value of the sometimes-active developer's probability of inactivity $F_k(0) > 0$; to derive equilibria we search numerically for a value of $F_k(0)$ that satisfies the upper boundary condition that $F_k(\bar{s}) = F_{-k}(\bar{s}) = 1$ (that is, that both CDFs reach 1 at the same top score $\bar{s}$, which is endogenously determined). Equilibria appear to be unique for all parameter values considered, and our numerically-computed equilibria appear to converge to approach the unique analytically computed equilibrium as developer parameters approach symmetry.