Productive Policy Competition and Asymmetric Extremism*

Alexander V. Hirsch†

September 4, 2026

Abstract

Viable policies must be developed by individuals and groups with the expertise and willingness to do so. We develop a model in which competing policy-motivated developers 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 – where the decisionmaker is always harmed by asymmetries due to the discouragement effect – and results from the developers' policy motivations. I also introduce and solve a broader class of *generalized policy contests* that nests both the policy-development model and the all-pay contest, and develop new methods for solving asymmetric contests with spillovers. Broadly, my analysis provides a novel rationale for why extreme actors may 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 initially *developed*. In his sweeping work on the policy process, 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 on the challenges of effective policy development as follows:

> Just attending to all the technical details of putting together a real proposal takes a lot of time. There's tremendous detail in the work. 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)). 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 I 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, I 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,

[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.

whether such beneficial effects extend to asymmetric extremism; and indeed, it is reasonable to 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 my 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.

My 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.* My 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 my 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 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 I also use my model to examine the effect of asymmetries in ability at crafting high quality policies. My 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.

My 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. My results also yield an important lesson for the design of effective political institutions; that there is really no way to allocate ostensibly "non-partisan" resources to develop policy (like expert staff and budgets) without having "ideological" consequences (Reynolds (2020)). When a developer becomes endowed with greater ability in my 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, I use two distinct approaches to explore robustness and develop the intuition for my novel results about the potential benefits of asymmetries.

First, I solve the main model by introducing and solving a larger class of models that I term *generalized policy contests*. Since the classic all-pay contest is also a special case of the generalized policy contest, this solution yields a clean analytic characterization of how my model's two key differences with the all-pay contest – that the developers intrinsically value quality, and that they care about about the endogenous ideological outcome – separately affect equilibrium strategies and payoffs. The analysis also introduces new techniques for solving asymmetric contests with spillovers.

Second, after the main analysis and substantive model results, I consider a series of additional 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. I show that with this change the benefits of asymmetric extremism and ability vanish, which illustrates the importance of the developers' ability to strategically choose the extremism of their policy for my results. In the second, the developers receive a fixed payoff when losing, so that they only care about policy when they win; I again show that the benefits of asymmetric extremism and ability vanish, illustrating the importance of policy motivations for my results. After analyzing these two variants, I return to the main model and flesh out only the *combination* of the preceding features yields a benefit from asymmetries.

In the final variant, the players' preferences over policy are linear but the costs of generating quality are quadratic; this is to verify that our results are not an artifact of combining an exponentially-growing value for policy concessions with a linear cost of generating the quality used to extract them. While equilibria must be computed numerically and certain properties change, our main finding that the DM broadly benefits from asymmetric extremism and ability persists; the reason is that it remains true in this variant that each developer reacts to less activity from his opponent with more ideological aggression, thereby magnifying the opponent's cost of losing.

Related Literature

My work connects to several large literatures across political science and economics.

Policy Expertise and Policy Development I first relate to an extensive literature studying policy expertise and the development of high quality policies. 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* (e.g. Gilligan and Krehbiel (1989)). 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.

My 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 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

[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 classical 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."

resemble entry deterrence models of market competition.

The intended empirical domain of my 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 (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).

Political Contests My model also relates to a literature studying political contests, usually applied to lobbying and elections. Foundational work by Tullock (1980) modelled lobbying as a process whereby competing groups exert wasteful effort to increase the 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 the groups may be asymmetric – because one values winning more, dislikes effort less, or both.[5] This model is now known as

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

the *all-pay contest* because it generalizes the all-pay auction format, in which the prize is awarded to the highest bidder but all bids are paid (Baye, Kovenock and de Vries (1996); Siegel (2009)). The two player all-pay contest exhibits a phenomenon widely known in the contest literature as the *discouragement effect*. As the stronger bidder becomes yet stronger (either because he values the prize more, or dislikes effort less), the weaker bidder bids less, anticipating a likely loss. However, the stronger bidder *also* bids (weakly) less, anticipating less competition. Consequently, any asymmetry between the bidders reduces overall effort, thereby harming anyone who benefits from that effort; even if the asymmetry is due to one bidder valuing the prize *more* or disliking effort *less*.

My 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 my developers' payoffs from winning and losing are *endogenous* to the policies they develop because they are policy motivated rather than merely "rent-seeking." In the terminology of Baye, Kovenock and de Vries (2012), my contest model 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. This property reverses the standard discouragement effect, and yields novel implications about the consequences of asymmetries in political settings with ideologically motivated actors.

Finally, the literature on political contests includes several Downsian election models in which candidates choose both ideological platforms and costly valence-generating campaign spending, much like the policy developers in my model choosing ideology and quality (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 models study candidates sequentially choosing platforms and then spending, spending and then platforms, or choosing in a pre-determined order. Such models feature strategic forces that are closely related to an earlier generation of Downsian models that studied candidates with exogenous valence advantages (e.g. Groseclose (2001), Aragones and Palfrey (2002)). Indeed, my developers' attempts to parlay quality investments into support for more extreme policies is akin to the force underlying the influential "marginality hypothesis" (Groseclose (2001) p. 863), whereby an incumbent will "parlay the advantage into a policy position that is closer to her ideal point." (For an extensive discussion see Hirsch (2023)).

The Policy Development Model

Two policy developers, labelled -1 (left) and 1 (right), craft competing policies for consideration by a decisionmaker (henceforth 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 DM is located at $x_D = 0$, the left developer is distance $|x_{-1}| > 0$ to her left, and the right developer is distance $|x_1| > 0$ to her right. A developer's distance $|x_i| = ix_i$ from the DM captures 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 stage, 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 policy development. $\alpha_i = \frac{a_i}{\lambda}$ 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. I assume that $\alpha_i > 1$ for both developers, implying that neither would invest in quality for its own sake. In the second stage the DM chooses a final policy. This may be a policy crafted by the developers, or any other policy from a set of (possibly empty) outside options $\mathbb{O}$ that only includes policies weakly worse for the DM than her ideal point with 0-quality (and may be empty).[6]

The Monopolist's Problem and the Score

To clarify incentives I first briefly review the model with a single developer, i.e. a "monopolist" (see Hitt, Volden and Wiseman (2017) and Hirsch and Shotts (2018)). I also introduce a critical quantity for the subsequent analysis called a policy's *score*. 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 above the bold curve is the set of policies the DM is willing to implement in lieu of (y_0, q_0), which denotes the best policy she can implement on her own (i.e., her "outside option"). I temporarily assume that this outside option is weakly better than the DM's ideal with 0-quality $(0, 0)$.

A monopolist will choose his policy $(y_i, q_i) \in \mathbf{R} \times \mathbf{R}_{\geq 0}$ to maximize

$$-a_i q_i + \mathbf{1}_{\substack{\lambda q_i - y_i^2 \\ \geq \lambda q_0 - y_0^2}} \cdot \left(\lambda q_i - (y_i - x_i)^2\right) + \left(1 - \mathbf{1}_{\substack{\lambda q_i - y_i^2 \\ \geq \lambda q_0 - y_0^2}}\right) \cdot \left(\lambda q_0 - (y_0 - x_i)^2\right), \tag{1}$$

where $\mathbf{1}_S$ is an indicator variable for whether the statement S is true. Specifically, he pays the up-front cost $a_i q_i$ of generating his policy's quality regardless of what the DM does. If his policy is at least as good for the DM as her outside option (i.e. $\lambda q_i - y_i^2 \geq \lambda q_0 - y_0^2$) then she will implement it, and the developer will receive policy utility $\lambda q_i - (y_i - x_i)^2$. Otherwise she will implement her outside option, and the developer will receive policy utility $\lambda q_0 - (y_0 - x_i)^2$. Intuitively, the developer must choose both *whether* to craft a new policy that the DM is actually willing to implement (since otherwise any effort on quality will be wasted), and if so exactly *which* policy (y_i, q_i) to craft.

To facilitate the analysis, I henceforth reparameterize policies (y_i, q_i) in terms of their ideology y_i and the *utility they give the DM* s_i; I call the latter a policy's *score*. Now observe that $s_i = \lambda q_i - y_i^2$; thus, a policy with score s_i and ideology y_i must have quality $q_i = \frac{s_i + y_i^2}{\lambda}$. Next recall that $\alpha_i = \frac{a_i}{\lambda}$ is the marginal cost of generating quality weighted by the marginal benefit. Then substituting into

[6] Hirsch and Shotts (2015) assume that this set is just the DM's ideal with zero quality.

(1) and simplifying yields the easier reparameterized objective function

$$-\alpha_i\left(s_i + y_i^2\right) + \mathbf{1}_{s_i \geq s_0} \cdot (s_i + 2x_i y_i) + (1 - \mathbf{1}_{s_i \geq s_0}) \cdot (s_0 + 2x_i y_0) - x_i^2 \tag{2}$$

A monopolist thus competes in a sort of contest – albeit a very predictable one. He makes a multidimensional "bid" (s_i, y_i) that consists of his policy's score s_i and ideology y_i; but to "win" he need only craft a policy with a score s_i that exceeds the score s_0 of the DM's outside option. Winning yields utility $V_i(s_i, y_i) = s_i + 2x_i y_i$ based on his own policy, while losing yields utility $V_i(s_0, y_0) = s_0 + 2x_i y_0$ based on the outside option. Holding ideology fixed, the developer enjoys a linear benefit from a higher score outcome because it implies higher quality. Holding score fixed, a right-wing (left-wing) developer enjoys a linear benefit from a rightward (leftward) policy movement weighted by his ideological extremism $|x_i|$. Finally, the constant $-x_i^2$ may be ignored.

How will a monopolist choose the optimal policy (s_i^*, y_i^*)? First observe that conditional on choosing a winning score ($s_i \geq s_0$) the derivative of the developer's objective with respect to score is $-\alpha_i + 1 < 0$; thus, there is no benefit to choosing a winning score above the minimum needed to win. *If* the developer targets a winning score (i.e., crafts a policy the DM is willing to implement), he will therefore target the lowest one ($s_i^* = s_0$). Further, he will choose its ideology y_i^* by trading off the linear benefit $2x_i y_i^*$ of moving the outcome in his desired direction (holding score fixed at $s_i^* = s_0$) against the quadratic cost $\alpha_i y_i^2$ of compensating the DM for a more extreme ideology with additional quality. Differentiating with respect to y_i and setting equal to 0 yields the optimal monopoly ideology

$$y_i^* = \frac{x_i}{\alpha_i}, \tag{3}$$

which is a weighted average (by $\frac{1}{\alpha_i}$) of the developer and DM's ideal ideologies. A monopolist who chooses to be active thus exploits his monopoly access to the technology of policy development to exercise what Hirsch and Shotts (2012) term *informal agenda power*, investing in only barely enough quality to make the DM willing to accept a policy that partially serves his ideological interests.

Finally, the developer will choose to be active and craft this optimal policy if and only if the

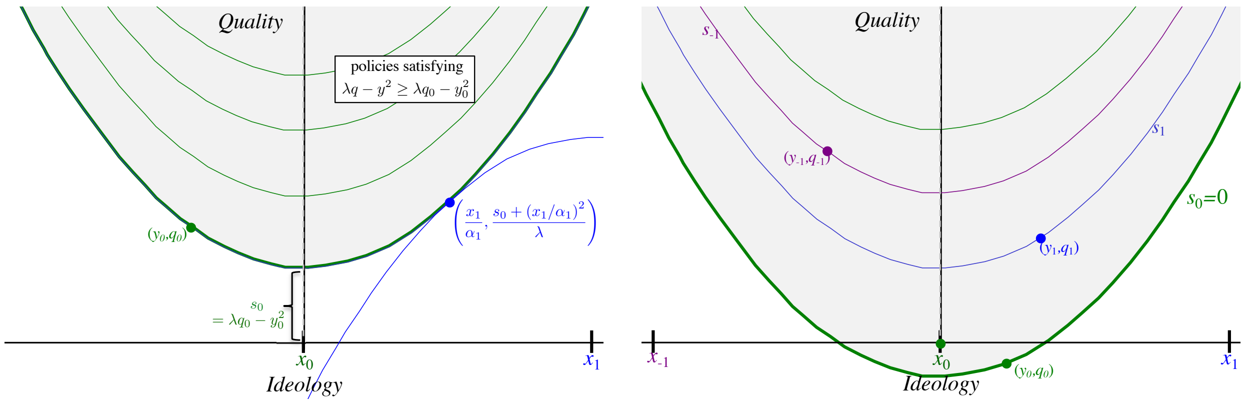

Figure 1: *The Developer's Problem.* The left panel depicts the problem faced by a "monopolist" when the DM's outside option is (y_0, q_0); the green curves are the DM's indifference curves, the shaded region depicts the policies the DM weakly prefers to her outside option, and the blue curve depicts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

utility from doing so exceeds the utility from simply accepted DM's outside option (s_0, y_0). The outside option thus functions as both a *constraint on* and *motive for* policy development; the reason is that the developer must live under the outside option if he fails to craft an acceptable alternative. While all outside options on the same score curve s_0 are equally difficult to beat, those further to the left are worse for the developer, and so more strongly incentivize policy development.

The Competitive Problem

With a little reflection, it is clear that two competing policy developers cannot be crafting predictable policies in equilibrium – that is, there cannot be a pure strategy equilibrium of the competitive model. A developer who anticipated a specific policy from his ideological opponent would try to craft a policy that serves his ideological interests but is just barely better than the opponent's policy (in the eyes of the DM); but then his losing opponent would would choose to deviate and pursue the same strategy against this new alternative, and so on. This "bidding up" dynamic closely resembles what occurs in a classic model of political competition that I will review shortly; the *all pay contest* (Siegel (2009)). In the all pay contest, two bidders simultaneously exert costly effort to win a prize, which falls to the bidder who exerts the most effort. Equilibrium effort levels in the all pay contest cannot be predictable because a certain loser would drop out, a certain winner would reduce his effort a little bit, and in a tie both would slightly increase effort to win.

The competitive game is depicted in the right panel of Figure 1. Each developer simultaneously chooses his policy (s_i, y_i), and the DM implements either the one with the highest score (i.e., on the highest indifference curve in Figure 1) or her best outside option (s_0, y_0). In the Appendix I show that the developers' strategies must take the following general form in any Nash equilibrium.[7]

Lemma 1. *In any equilibrium, each developer's strategy may be written as*

- *A smooth* ***cumulative distribution function*** $F_i(s)$ *with full support over a nonempty finite interval* $[0, \bar{s}]$ *that describes the probability developer* i *crafts a policy with score* $\leq s$
- *A smooth* ***ideology function*** $y_i(s)$ *that describes the ideology developer* i *targets when crafting a policy with score* s

Furthermore, the opponent of at least one developer k *must always be* ***active****, in the sense of always crafting a positive score (and hence positive quality) policy (i.e.,* $\exists k$ *s.t.* $F_{-k}(0) = 0$*).*

Given the preceding, a developer i's expected utility from crafting (s_i, y_i) with $s_i > 0$ is thus:

$$-\underbrace{\alpha_i\left(s_i + y_i^2\right)}_{\text{quality cost}} + \underbrace{F_{-i}\left(s_i\right)}_{\text{Pr win}} \underbrace{\left(s_i + 2x_i y_i\right)}_{\text{utility from winning}} + \underbrace{\int_{s_i}^{\bar{s}} \left(s_{-i} + 2x_i y_{-i}\left(s\right)\right) f_{-i}\left(s_{-i}\right) ds_{-i}}_{\text{expected utility when losing}} - x_i^2. \tag{4}$$

The expression resembles equation 2, with two key differences. First, the probability i's policy (s_i, y_i) wins is $F_{-i}(s_i)$ (the probability his opponent crafts a lower score policy) rather than $\mathbf{1}_{s_i \geq s_0}$ (whether his score exceeds the DM's known outside option). Second, i's utility when losing depends on the *average higher-score policy* $E\left[\left(s_{-i}, y_{-i}\left(s_{-i}\right)\right) | s_{-i} > s_i\right]$ that his opponent crafts, which is endogenous. Summarizing, the competitive policy development model is effectively a multidimensional contest in

[7] See Appendix B for proofs of a more general model subsuming the policy development model.

which each player chooses a "bid" (s_i, y_i) at up-front cost $\alpha_i(s_i + y_i^2)$ and receives final utility $V_i(s_w, y_w) = s_w + 2x_i y_w$, where w denotes the identity of the highest-score player (i.e. the ***w***inner).

Reviewing the two-player all-pay contest

Before characterizing equilibrium, it is now helpful to step back and review the classic two-player all pay contest (Hillman and Riley (1989); Siegel (2009)). In that model, each bidder chooses up-front effort $s_i \geq 0$ at a marginal cost $\alpha_i \geq 0$, a decisionmaker selects the bidder exerting the most effort as the winner, and each bidder has a fixed individual-specific win payoff normalized to $2|x_i|$. As in my model, the unique equilibrium is in mixed strategies described by univariate CDFs $F_i(s_i)$ with full support over a nonempty interval finite $[0, \bar{s}]$, and at least one bidder's opponent is always active ($\exists k$ s.t. $F_{-k}(0) = 0$). Bidder i's expected utility from exerting effort $s_i > 0$ is thus:

$$-\alpha_i s_i + F_{-i}(s_i) \cdot 2|x_i| \tag{5}$$

Comparing equations 4 and 5 reveals two key differences between my model and the all pay contest. First, in my model each developer enjoys a *direct benefit* from the score s_w of the winning bidder; it is thus *as if* the winner's effort has intrinsic value (albeit less than the marginal cost). Second, each developer can effectively *raise the stakes* of the contest by choosing an ideology y_i that is further in his desired ideological direction, which increases the value of winning for both players. Both differences constitute what Baye, Kovenock and de Vries (2012) term *rank order spillovers* – the strategy of the first ranked bidder (the winner) "spills over" onto the utility of the second ranked bidder (the loser) because it is the policy he must live under.

The all-pay contest is simple to solve because the bidders' stakes are fixed. Since every bid $[0, \bar{s}]$ in the common mixing interval must be optimal, each must maximize equation 5, implying that

$$f_k(s) = \frac{1}{2|x_{-k}|/\alpha_{-k}} \quad \text{and} \quad f_{-k}(s) = \frac{1}{2|x_k|/\alpha_k} \quad \forall s \in [0, \bar{s}]. \tag{6}$$

Combined with $F_{-k}(0) = 0$ (some bidder k has an opponent who is always active) and $F_{-k}(\bar{s}) = F_k(\bar{s}) = 1$ (both CDFs reach 1 at the same top score) it then follows that $F_{-k}(s) = \frac{s}{2|x_k|/\alpha_k}$; in words, the always-active bidder $-k$ must randomize uniformly over an interval $\left[0, \frac{2|x_k|}{\alpha_k}\right]$ whose upper bound is his opponent k's *ratio of prize value to marginal costs* $\frac{2|x_k|}{\alpha_k}$. This ratio is also the highest score bidder k is willing to choose to win before he prefers to simply drop out, which Siegel (2009) terms bidder k's *reach*; I henceforth denote this critical quantity as $r_i = \frac{2|x_i|}{\alpha_i}$. Next, since the CDF $F_k(s)$ of the sometimes-inactive bidder k must equal 1 at the same top score $\bar{s} = r_k$, and also have a slope $f_k(s)$ equal to reciprocal $\frac{1}{r_{-k}}$ of bidder $-k$'s reach, the full equilibrium strategies must be

$$F_{-k}(s) = \frac{s}{r_k} \quad \text{and} \quad F_k(s) = \left(1 - \frac{r_k}{r_{-k}}\right) + \frac{s}{r_{-k}},$$

which can be written as $F_i(s) = 1 - \frac{r_k - s}{r_i}\ \forall i$. Finally, the bidder k with an always-active opponent must be the one with the lowest reach ($k = \arg\min_j \{r_j\}$) since otherwise $F_k(0) < 0$, a contradiction.

The bidders' equilibrium utilities are just equation 5 with the highest score $\bar{s} = r_k$ plugged in, so

$$U_i^* = -\alpha_i r_k + 2|x_i| = 2|x_i|\left(1 - \frac{r_k}{r_i}\right)$$

The weaker bidder k's equilibrium utility U_k^* is thus fixed at 0 regardless of the stronger bidder's characteristics, while the stronger bidder earns "rents" that are a function of his reach advantage $1 - \frac{r_k}{r_{-k}}$. Lastly, the DM's equilibrium utility (assuming that she only benefits from the effort of the winner) is simply *the expected maximum score*, which has a CDF of $\prod_j F_j(s)$. Thus

$$U_{DM}^* = \int_0^{\bar{s}} s \cdot \frac{\partial}{\partial s}\left(\prod_j F_j(s)\right) \cdot ds = \int_0^{\bar{s}} \left(1 - \prod_j F_j(s)\right) ds,$$

where the second equality follows from integration by parts and rearranging.

Several crucial properties of the all-pay contest are now worth noting to later contrast with my model. First, the bidders' equilibrium strategies depend *only* on the profile of reaches – there is thus an isomorphism between a bidder valuing the prize more and disliking effort less. Second, if the two bidders have the same reach ($r_{-1} = r_1 = r$), then the DM's expected utility $U_{DM}^* = \int_0^r \left(1 - \left(\frac{s}{r}\right)^2\right) ds$ is strictly increasing in that reach. By implication, in a *symmetric* all-pay contest ($|x_{-1}| = |x_1| = x$

and $\alpha_{-1} = \alpha_1 = \alpha$), the decisionmaker benefits from both a *common* increase in the value of the prize x or a common decrease in the cost of effort α; both foster more intense competition.

Finally, the two-player all-pay contest exhibits a phenomenon known as the *discouragement effect.* As the stronger bidder becomes yet stronger (r_{-k} increases), the weaker bidder becomes strictly less active, but the stronger bidder *also* becomes (weakly) less active (anticipating less competition). Thus, *any* asymmetry between the bidders reduce the provision of effort and harms a decisionmaker who benefits from it – even if it is due to one bidder valuing the prize more or disliking effort less.

To see this property, consider the effect of an increase in r_{-k}; from equation 6 the CDF of the weaker bidder k must become *flatter* ($f_k(s)$ decreases $\forall s$) to preserve the stronger bidder's willingness to mix over a common score interval $[0, \bar{s}]$. But since both bidders' score CDFs must equal 1 at the same top score $\bar{s}$, this can only be accomplished by strictly increasing the weaker bidder's probability of inactivity $F_k(0)$, so that $F_k(s)$ *first-order stochastically decreases* (i.e., increase $\forall s \in [0, \bar{s}]$). This is the first part of the discouragement effect; as a stronger bidder becomes yet stronger, his opponent becomes less active. What then happens to the stronger bidder's score CDF $F_{-k}(s)$? The answer is *nothing at all.* The reason is simple: the stronger bidder's score CDF is determined *only* by need to keep the weaker bidder willing to target scores in $[0, \bar{s}]$, and neither the weaker bidder's prize value $2|x_k|$ nor marginal costs α_k have changed. This yields the following.

Observation 1. *In the two-player all pay contest, asymmetries always harm the decisionmaker.*

Worth emphasizing is that the stronger bidder cannot simply become "motivated enough" to overcome the negative impact of the discouragement effect; in equilibrium his growing desire to win has no effect on his actual participation, even as it pushes his opponent out of the contest.

Deriving equilibrium and the generalized policy contest

To both solve the policy development model and compare with the all pay contest, I nest both models into a larger class of models that I term *generalized policy contests.* This section is not required to read and interpret the main results, but is helpful for the later discussion on robustness.

Definition 1. *In a generalized policy contest, each bidder $i \in \{-1, 1\}$ chooses strategy (s_i, y_i) at up-front cost $\alpha_i(s_i + y_i^2)$ and receives final utility $-\alpha_i(s_i + y_i^2) + V_i(\mathbf{1}_{i=w}, s_w, y_w)$, where*

$$V_i(\rho_i, s_w, y_w) = \gamma_i s_w + 2x_i \left((1-\theta) \cdot i\rho_i + \theta \cdot y_w\right),$$

***w**inner denotes the highest-scoring bidder (i.e. $w = \arg\max_j\{s_j\}$), and $\gamma_i < \alpha_i$. Special cases include the all pay contest ($\theta = \gamma_1 = \gamma_{-1} = 0$) and the policy development model ($\theta = \gamma_1 = \gamma_{-1} = 1$).*

The generalized policy contest adds parameters $(\theta, \gamma_1, \gamma_{-1})$ that build a bridge between the policy development model and the all pay contest. θ captures the extent to which the prize has an endogenous ideological component $2x_i y_w$ vs. a fixed value of winning $2|x_i|$, while γ_i captures the extent to which bidder i intrinsically values the winner's effort.

In the Appendix I show that equilibrium strategies also take the form in Lemma 1; a bidder i's expected utility from choosing strategy (s_i, y_i) is thus

$$-\underbrace{\alpha_i\left(s_i + y_i^2\right)}_{\text{quality cost}} + \underbrace{F_{-i}\left(s_i\right)}_{\text{Pr win}}\underbrace{\left(\gamma_i s_i + 2x_i\left(i\left(1-\theta\right) + \theta y_i\right)\right)}_{\text{utility from winning}} + \underbrace{\int_{s_i}^{\bar{s}} \left(\gamma_i s_{-i} + 2x_i \cdot \theta y_{-i}\left(s\right)\right) f_{-i}\left(s_{-i}\right) ds_{-i}}_{\text{expected utility when losing}} \quad (7)$$

To solve this model, first consider bidder i's choice of ideology $y_i\left(s_i\right)$ holding score fixed at s_i. This choice resembles that of a monopolist choosing his policy's ideology, in that there is an up-front quadratic cost of a more extreme ideology, and an ideological benefit conditional on winning. Differentiating with respect to ideology y_i and setting equal to 0 yields the optimal score-s ideology

$$y_i(s) = F_{-i}(s)\frac{x_i}{\alpha_i}\theta. \quad (8)$$

This differs from the monopolist's optimal ideology (in eqn. 3) in two ways. First, it depends on the degree θ to which the prize has an ideological component. Second, it is weighted by the endogenous probability $F_{-i}\left(s_i\right)$ that i's opponent $-i$ crafts a lower score policy. Thus, a bidder in the competitive setting *behaves more ideologically aggressively* at a score s when he believes that a score-s policy is *more likely to win* (higher $F_i(s)$). Correspondingly, should one bidder participate less in the contest, the other will strategically respond by becoming more ideologically aggressive, thereby raising the stakes of the contest for both bidders. This force crucially distinguishes my model from a classic all-pay contest; and the more the prize has an ideological component (higher θ), the stronger it is.

Next consider a bidder's optimal choice of score s_i. Substituting equation 8 into equation 7, differentiating with respect to s_i, observing that $y_{-1}(0) \leq 0 \leq y_{-i}(0)$ (i.e., the bidders choose ideologies on opposite sides of the decisionmaker), and simplifying yields:

$$-\left(\alpha_i - \gamma_i F_{-i}(s_i)\right) + f_{-i}(s_i) \cdot 2\left|x_i\right| \cdot \left((1-\theta) + \theta \sum_{j \in \{-1,1\}} \left|y_j(s_i)\right|\right) \tag{9}$$

A bidder's incentive to target a marginally higher score (assuming both target optimal ideologies) thus has two main components. The first is a *net marginal cost* $\alpha_i - \gamma_i F_{-i}(s_i)$ of generating a higher score; this in turn consists of the up-front marginal cost of effort α_i, net of the intrinsic score benefit γ_i he enjoys if his bid wins (with probability $F_{-i}(s_i)$). A bidder intrinsically valuing quality ($\gamma_i > 0$) thus acts as a sort of "force multiplier" that effectively reduces his marginal cost of effort; but the degree to which it does so depends on the exact score and opponent's strategy. The second is a marginal increase $f_{-i}(s_i)$ in the probability that bidder i wins, which gives him a direct winning benefit of $(1-\theta)$, but also a policy benefit of changing the ideological outcome from $y_{-i}(s)$ (on the opposite side of the DM) to $y_i(s)$ (on his side of the DM). The generalized policy contest thus exhibits *common score-specific stakes* $(1-\theta) + \theta \sum_{j \in \{-1,1\}} |y_j(s)|$ that depend on the endogenous ideologies that the bidders target at each score. Each bidder i then scales these stakes by $2|x_i|$.

Identifying the weaker bidder Siegel (2009) shows that the reach $r_i = \frac{2|x_i|}{\alpha_i}$ has special significance in a general class of all-pay contests; it both identifies the weaker bidder k, *and* can be used to directly calculate equilibrium payoffs.[8] This approach cannot be used to identify the weaker bidder in the generalized policy contest, however, because the bidders' payoff from losing is endogenous;

[8]With two bidders, payoffs must be what they get from winning for sure by playing k's reach.

there is thus no clearly defined "highest score" that they are willing to choose. I can nevertheless derive a comparable quantity to identify the weaker bidder that I term the *relative reach* $\tilde{r}_i$.

Although the stakes in the generalized policy contest are both endogenous and score-specific, the model becomes tractable because these stakes are *common across the two bidders* up to a bidder-specific scalar multiple $2|x_i|$. Formally, observe that equilibrium requires every positive score in the common support of the bidders' score CDFs to be optimal (so that they are willing to randomize among them). Thus, equation 9 must be equal to 0 for all i and $s \in (0, \bar{s}]$, i.e.,

$$\frac{2|x_i| \cdot f_{-i}(s)}{\alpha_i - \gamma_i F_{-i}(s)} = \left((1-\theta) + \theta \sum_{j \in \{-1,1\}} |y_j(s)|\right)^{-1} \quad \forall i \text{ and } s \in (0, \bar{s}]. \tag{10}$$

That is, at every positive score in the common support, a bidder's ratio of marginal policy benefits to net marginal costs must equal the reciprocal of the common stakes at that score. But then these ratios must *also* be equal to each other for all $s \in (0, \bar{s}]$, implying that

$$\int_0^{\bar{s}} \frac{2|x_i| \cdot f_{-i}(s)}{\alpha_i - \gamma_i F_{-i}(s)} ds = \int_0^{\bar{s}} \frac{2|x_{-i}| \cdot f_i(s)}{\alpha_{-i} - \gamma_{-i} F_i(s)} ds \tag{11}$$

Now observing that $\int_0^s \left(\frac{2x \cdot f(s)}{\alpha - \gamma F(s)}\right) ds = \int_{F(0)}^{F(\bar{s})} \left(\frac{2x}{\alpha - \gamma F}\right) dF$ via a change of variables from s to F, and also that $F_{-i}(\bar{s}) = F_i(\bar{s}) = 1$ (both CDFs equal 1 at the same top score $\bar{s}$), equation 11 simplifies to

$$\int_{F_{-i}(0)}^{1} \frac{2|x_i|}{\alpha_i - \gamma_i F} dF = \int_{F_i(0)}^{1} \frac{2|x_{-i}|}{\alpha_{-i} - \gamma_{-i} F} dF.$$

The preceding is a simple functional relationship between the bidders' probabilities of inactivity $F_i(0)$ which can be used to identify the weaker bidder k as follows.

Lemma 2. *Term* $\tilde{r}_i = \int_0^1 \frac{2|x_i|}{\alpha_i - \gamma_i F} dF$ *bidder* i*'s **relative reach**. The weaker bidder* k *with an always-active opponent* ($F_{-k}(0) = 0$) *is the one with the lowest relative reach* ($k = \arg\min_j \{\tilde{r}_j\}$)*.*

Proof: Suppose that $\tilde{r}_k < \tilde{r}_{-k}$ (k has a strictly lower relative reach) but $F_k(0) = 0$ (k is always active). Then $\int_{F_k(0)}^1 \frac{2|x_{-k}|}{\alpha_{-k} - \gamma_{-k} F} dF = \int_0^1 \frac{2|x_{-k}|}{\alpha_{-k} - \gamma_{-k} F} dF = \int_{F_{-k}(0)}^1 \frac{2|x_k|}{\alpha_k - \gamma_k F} dF$; but since $\int_0^1 \frac{2|x_{-k}|}{\alpha_{-k} - \gamma_{-k} F} dF > \int_0^1 \frac{2|x_k|}{\alpha_k - \gamma_k F} dF$ it must be that $F_{-k}(0) < 0$, a contradiction. QED

Bidder i's relative reach is strictly increasing in his intrinsic valuation of effort γ_i, and equals the standard reach $\frac{2|x_i|}{\alpha_i}$ when he does not ($\gamma_i = 0$). Intrinsically valuing effort thus effectively increases a bidder's strength by reducing his marginal cost of effort. This characterization further implies that the presence of endogenous ideological stakes does not affect the bidders' relative reaches. However, it does preclude using the relative reach to directly calculate equilibrium payoffs.

Equilibrium characterization and payoffs I now fully characterize equilibrium and payoffs (see Appendix for details). The key step is to observe that the same property used to derive the relative reach – that the endogenous score-specific stakes are common between the bidders up to a scalar multiple – also implies the following relationship between the bidders' full equilibrium score CDFs:[9]

$$\tilde{r}_k\left(F_{-k}\left(s\right)\right) = \tilde{r}_{-k}\left(F_k\left(s\right)\right) \quad \forall s \in [0, \bar{s}] \quad \text{where} \quad \tilde{r}_i\left(W_i\right) = \int_{W_i}^{1} \frac{2\left|x_i\right|}{\alpha_i - \gamma_i F} dF. \tag{12}$$

I term $\tilde{r}_i(W_i)$ bidder i's *residual reach at win probability* W_i, since it effectively captures how much of bidder i's relative reach is "left over" when he is already winning with probability W_i (bidder i's relative reach $\tilde{r}_i = \tilde{r}_i(0)$ is just his residual reach at probability 0). Equilibrium thus requires that at every score $s \in [0, \bar{s}]$ in the common support, the bidders' win probabilities $W_k = F_{-k}(s)$ and $W_{-k} = F_k(s)$ yield the same residual reach. The equilibrium characterization can thus be completed with a strictly decreasing function $\tilde{\mathbf{r}}(\mathbf{s})$ that connects this common residual reach to the score, i.e.,

$$\tilde{\mathbf{r}}(s) = \tilde{r}_k(F_{-k}(s)) = \tilde{r}_{-k}(F_k(s)) \quad \forall s \in [0, \bar{s}]. \tag{13}$$

$\tilde{\mathbf{r}}(s)$ must begin at the weaker bidder k's total relative reach (i.e., $\tilde{\mathbf{r}}(0) = \tilde{r}_k(0) = \tilde{r}_k$, since at the bottom score he surely loses) and end at 0 (i.e., $\tilde{\mathbf{r}}(\bar{s}) = 0$, since at the top score he surely wins).

The *inverse* of this function $\mathbf{s}(r)$ (that is, the score associated with each equilibrium level of the residual reach) is then simply characterized as follows. First, it must be a decreasing function from

[9] This relationship is derived using identical steps as in deriving the relative reach, except substituting in an arbitrary score $s \in [0, \bar{s}]$ for the lowest score $s = 0$.

$\tilde{\mathbf{s}}(0) = \bar{s}$ to $\tilde{\mathbf{s}}(\tilde{r}_k) = 0$. Next, its derivative satisfies

$$\tilde{\mathbf{r}}'(s) = \tilde{r}_i'\left(F_{-i}\left(s\right)\right) f_{-i}\left(s\right) = -\frac{2\left|x_i\right| f_i\left(s\right)}{\alpha_i - \gamma_i F_{-i}\left(s\right)} = -\left((1-\theta) + \theta \sum_{j\in\{-1,1\}} \left|y_j\left(s\right)\right|\right)^{-1},$$

where the first two equalities follow from equation 13 and the third follows from equation 10. Finally the preceding may be written in terms of the derivative of the inverse $\mathbf{s}'(r)$; substituting in the optimal ideologies $y_i\left(s\right) = F_{-i}(s)\frac{x_i}{\alpha_i}\theta$ yields

$$\mathbf{s}'(r) = -\left((1-\theta) + \theta^2 \sum_{j\in\{-1,1\}} F_{-j}(s\left(r\right))\frac{|x_j|}{\alpha_j}\right).$$

In words, as the common residual reach decreases, the associated score must increase at a rate equal to the common score-specific stakes. With the previously stated boundary conditions and this derivative, a complete integral characterization of $\mathbf{s}(r)$ and equilibrium is as follows.

Lemma 3. *In a Nash equilibrium of the generalized policy contest, the strategy of each bidder $i \in \{-1,1\}$ is $y_i\left(s\right) = F_{-i}(s)\frac{x_i}{\alpha_i}\theta$ and $F_i\left(s\right) = W_{-i}\left(\tilde{\mathbf{r}}\left(s\right)\right)$, where*

- *$W_i\left(r_i\right)$ is the inverse of*

$$\tilde{r}_i\left(W_i\right) = \int_{W_i}^{1} \frac{2\left|x_i\right|}{\alpha_i - \gamma_i F} dF$$

- *$\tilde{\mathbf{r}}\left(s\right)$ is the inverse of*

$$\mathbf{s}\left(r\right) = \int_{r}^{\tilde{r}_k(0)} \left((1-\theta) + \theta^2 \sum_{j\in\{-1,1\}} W_j\left(t\right)\frac{|x_j|}{\alpha_j}\right) dt$$

It is easily verified that this characterization reduces to the standard all-pay contest equilibrium when $\gamma_i = 0\ \forall i$ and $\theta = 0$, since then $r_i = \tilde{r}_i = \frac{2|x_i|}{\alpha_i}$ and $\tilde{r}_i\left(W_i\right) = \left(1 - W_i\right)\tilde{r}_i \iff W_i\left(r\right) = 1 - \frac{r}{\tilde{r}_i}$ and $\mathbf{s}\left(r\right) = \tilde{r}_k - r \iff \tilde{\mathbf{r}}\left(s\right) = \tilde{r}_k - s$, which together imply $F_i\left(s\right) = W_{-i}\left(\tilde{\mathbf{r}}\left(s\right)\right) = 1 - \frac{\tilde{r}_k - s}{\tilde{r}_i}$. Further, the complete equilibrium characterization cleanly separates the effect of the bidders intrinsically valuing quality γ_i (which only directly affects the definition of the residual reach functions $\tilde{r}_i\left(W_i\right)$) from the effect of the contest's endogenous ideological stakes θ (which only directly affects the equilibrium relationship $\mathbf{s}\left(r\right)$ between the score and the residual reach).

Lemma 3 can be used to straightforwardly calculate equilibrium payoffs as follows. For the bidders, just plug in the top score $\bar{s} = \mathbf{s}(0)$ into equation 7 and simplify, which yields

$$U_i^* = -(\alpha_i - \gamma_i)\,\bar{s} + 2x_i\left((1-\theta)\,i + \theta^2 \frac{x_i}{2\alpha_i}\right) \text{ with } \bar{s} = (1-\theta)\,\tilde{r}_k + \theta \int_0^{\tilde{r}_k} \sum_j W_j(r)\,\frac{|x_j|}{\alpha_j}\theta \cdot dr$$

The payoff of the weaker bidder k is thus directly affected by characteristics of the stronger bidder $-k$, in contrast to the all-pay contest. Formally, a bidder's payoff is decreasing in the maximum score $\bar{s}$, which in the generalized policy contest is a function of *both* bidders' characteristics via the endogenous ideological stakes $\sum_{j\in\{-1,1\}} W_j(r)\,\frac{|x_j|}{\alpha_j}$. Moreover, the maximum score $\bar{s}$ does not depend simply on bidder k's total relative reach $\tilde{r}_k$ (as in the all-pay contest), but rather on an integral over the endogenous ideological stakes for all possible values $r \in [0, \tilde{r}_k]$ of the residual reach.

Finally, to calculate the DM's equilibrium utility, recall that $U_{DM}^* = \int_0^{\bar{s}}\left(1 - \prod_j F_j(s)\right) ds$ from the all-pay contest. A change of variables from score $s \in [0, \bar{s}]$ to residual reach $r \in [0, \tilde{r}_k]$ then yields $\int_{\tilde{\mathbf{r}}(\bar{s})}^{\tilde{\mathbf{r}}(0)}\left(1 - \prod_j F_j(\mathbf{s}(r))\right)\left(-\frac{1}{\tilde{\mathbf{r}}'(\mathbf{s}(r))}\right) dr$. Simplifying using previously described equalities then yields

$$U_{DM}^* = \int_0^{\tilde{r}_k}\left(1 - \prod_j W_j(r)\right)\left((1-\theta) + \theta^2 \sum_j W_j(r)\,\frac{|x_j|}{\alpha_j}\right) dr \tag{14}$$

The first term in the integral $1 - \prod_j W_j(r)$ reflects the impact of the discouragement effect; as the stronger bidder $-k$ becomes stronger, his necessary win probability $W_{-k}(r)$ at each possible level of the residual reach $r \in [0, \tilde{r}_k]$ increases, implying less activity from the weaker bidder k. The second term $(1-\theta) + \theta^2 \sum_j W_j(r)\,\frac{|x_j|}{\alpha_j}$ reflects the impact of the endogenous ideological stakes. As the stronger bidder $-k$ becomes more extreme (higher $|x_{-k}|$) or capable (lower α_{-k}), he chooses more extreme ideologies ceteris paribus; this raises the contest stakes and thus the score $\mathbf{s}(r)$ (i.e., DM's utility) associated with each equilibrium level of the residual reach $r \in [0, \tilde{r}_k]$. The balance of these two effects determines whether asymmetries harm or help the DM overall.

Results

I now use the preceding general characterization to derive a more accessible characterization of equilibrium in the policy development model (see Appendix for details).

In equilibrium, each developer randomizes smoothly over a continuum of policies with ideologies between their own ideal point and the DM. A developer's equilibrium strategy can be written as two functions $(q_i(\delta), G_i(\delta))$ that describe: (1) the level of quality $q_i(\delta)$ developer i produces when crafting a policy whose ideology is *distance* δ from the DM, and (2) a *cumulative distribution function* $G_i(\delta)$ describing the probability developer i crafts a policy with ideology closer to the DM than δ.

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

$$\tilde{r}_i\left(W_i\right) = \int_{W_i}^{1} \frac{2|x_i|}{\alpha_i - F} dF = 2|x_i| \log\left(\frac{\alpha_i - W_i}{\alpha_i - 1}\right).$$[10]

Let k denote the developer with the smallest value of $\tilde{r}_i(0)$, and let $W_i(r_i) = \alpha_i - (\alpha_i - 1)\, e^{\frac{1}{2|x_i|} r}$ denote the inverse of $\tilde{r}_i\left(W_i\right)$.

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

- *The probability $\Pr(|y_i| \leq \delta)$ that developer i targets an y_i closer to the DM than δ is*

$$G_i\left(\delta\right) = W_{-i}\left(\tilde{r}_i\left(\frac{\delta}{|x_i|/\alpha_i}\right)\right) = \alpha_{-i} - \left(\alpha_{-i} - 1\right)\left(\frac{|x_i| - \delta}{|x_i| - |x_i|/\alpha_i}\right)^{\left|\frac{x_k}{x_{-k}}\right|}$$

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

$$s_i\left(\delta\right) = \int_{\tilde{r}_i\left(\frac{\delta}{|x_i|/\alpha_i}\right)}^{\tilde{r}_k(0)} \left(\sum_{j \in \{-1,1\}} W_j\left(r\right) \frac{|x_j|}{\alpha_j}\right) dr.$$[11]

[10] This is just the residual reach function in Lemma 3 for the policy development model.

[11] This is just the score function in Lemma 3 for the policy development model, with the residual reach generated by i's implied win probability $W_i = \frac{\delta}{|x_i|/\alpha_i}$ when his optimal policy is $i\delta$ plugged in.

Although the equilibrium strategies are straightforward to express, they are somewhat hard to interpret from the equations alone. I therefore describe the structure of equilibrium and key properties with the aid of an example.[12]

Figure 2 depicts equilibrium strategies and outcomes, comparing when the developers are symmetric (the left panels) to when they are equally capable but the right developer is more extreme (the right panels). Each developer i randomizes over an interval of ideologies whose outer bound is the ideology $\frac{x_i}{\alpha_i}$ he would target as a monopolist (which is a weighted average of his ideal point and the DMs; see eqn. 3). The top row depicts the policies (i.e. ideology and quality) that the left and right developers randomize over using dark curves (the DM's indifference curves are in thin gray); note that more extreme ideologies within a mixed strategy are always paired with both higher quality and higher scores (i.e., they sit on higher indifference curves). The middle row depicts the densities (PDFs) of the left and right developer's ideologies using dark curves; the density of the final ideological outcome is depicted with gray dashed lines.[13] Note that in the asymmetric case, the left developer also sometimes chooses to craft *no* new policy; this is illustrated in the right first-row panel by the dot at the origin (the DM's ideal point with zero quality), while the height of the vertical segment in the right second-row panel depicts the probability that this occurs. Finally, the bottom row depicts the developers' equilibrium score CDFs using dark curves, as well as the CDF of the maximum score policy (i.e., the DM's final utility) with a dashed curve; in the asymmetric case the left developer is sometimes inactive with probability $F_{-1}(0) > 0$.

[12]These properties are also implied by the implicit characterization in Hirsch and Shotts (2015).

[13]To understand why the density of developer i's ideology $\frac{\partial}{\partial \delta}\left(G_i\left(\delta\right)\right)|_{\delta=|y|} = g_i\left(|y|\right)$, observe that the CDF of developer 1's ideology is $\Pr(y_1 \leq y) = G(y) = 1 \cdot G(1 \cdot y)$ with $y \geq 0$ while the CDF of developer -1's ideology is $\Pr(y_{-1} \leq y) = 1 - G(|y|) = 1 + (-1)G((-1)y)$ with $y \leq 0$. Hence $\Pr(y_i \leq y) = \frac{1-i}{2} + iG(iy)$ for $i \in \{-1, 1\}$ and the derivative (i.e. density) is $i^2 g_i(iy) = g_i(|y|)$.

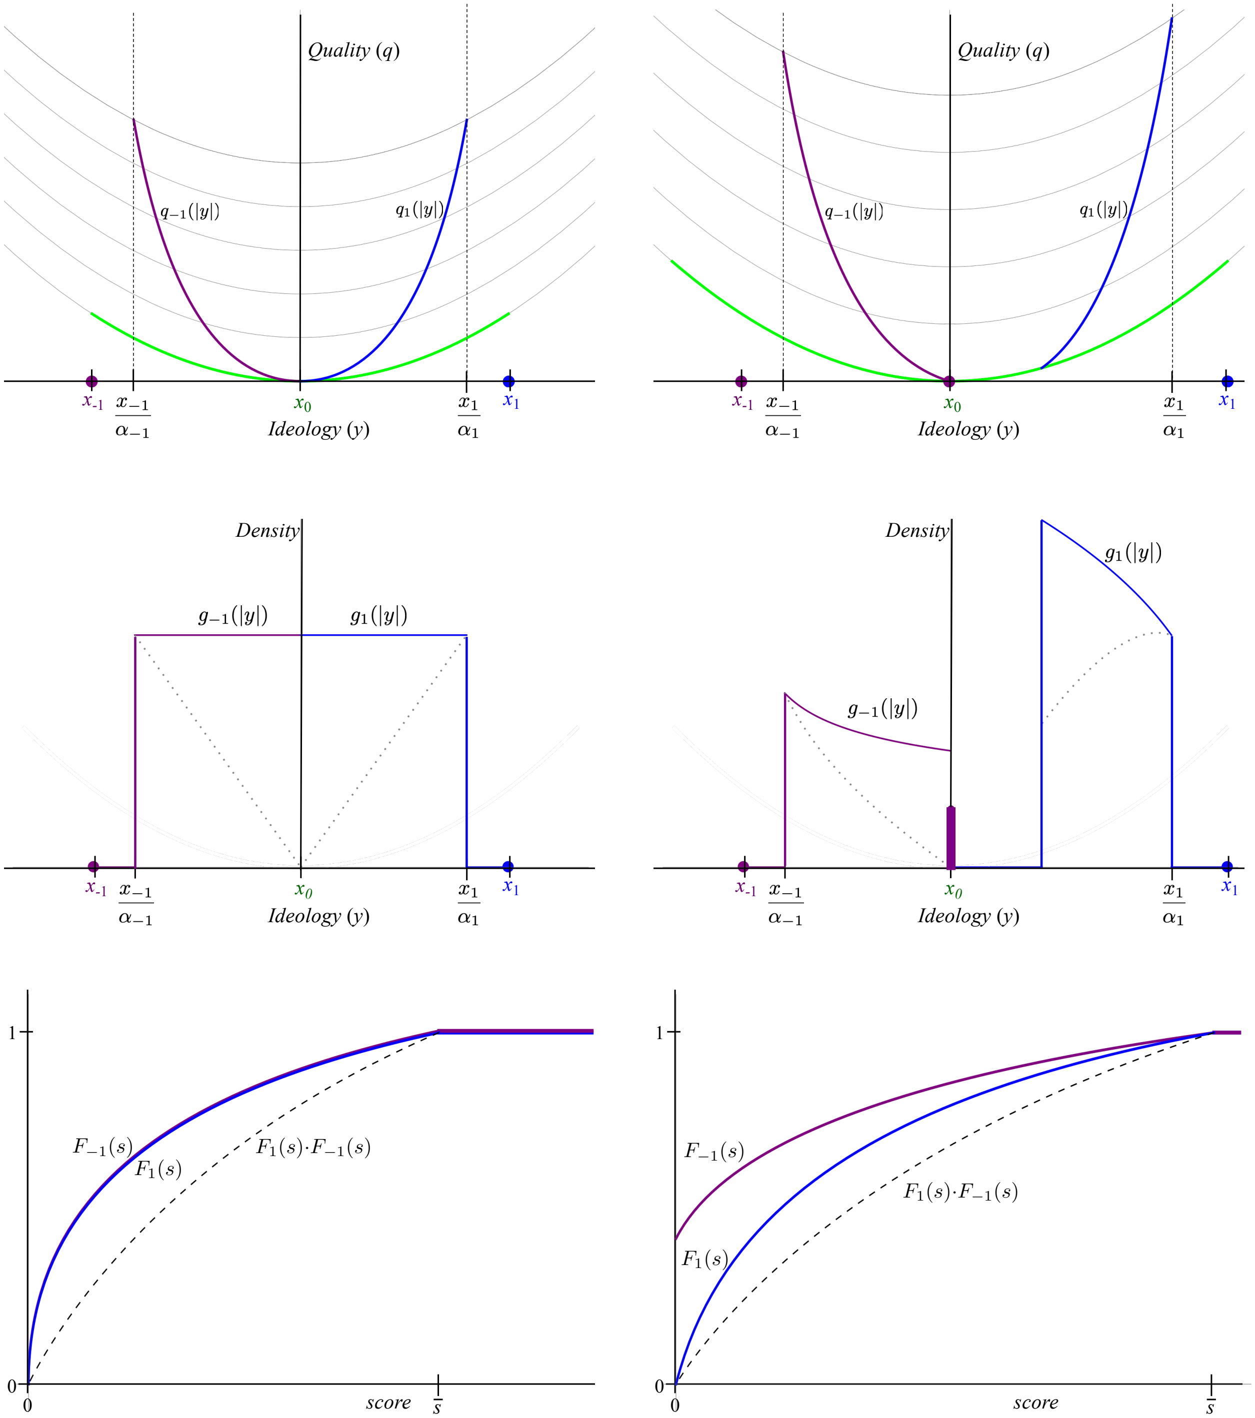

Figure 2: *Equilibrium with Symmetric Developers (left) vs. Asymmetric Extremism (right).*

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 new policy is strictly better for the DM than $(0, 0)$ (her ideal point with 0-quality) (in the first row of 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. The other developer (in the example the left developer, and more generally developer k in Proposition 1) is only *sometimes* active; with strictly positive probability he develops nothing. Overall, the pattern of participation resembles that of the classic all-pay contest; in both models there cannot be a strictly positive probability of a "tie" at non-participation since both would then want to deviate to a slightly more competitive strategy.

When a developer chooses to craft a new policy, its ideology always strictly diverges from the DM's ideal (in the first row of Figure 2, all positive-quality policies have ideologies that differ from the DM's ideal). 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 if they could craft policy as a "monopolist" (Hirsch and Shotts (2018)) – both invest in enough quality to (over)compensate the DM for her ideological losses from their respective policies, but not enough to compensate each other.

The equilibrium also exhibits a variety of inefficiencies. Except in the special case of symmetric developers who are equally extreme and capable, 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 process due to risk aversion (its distribution is depicted by the dashed grey lines in the second row of Figure 2). Finally, because the developers must make their quality investments before they know which policy will be chosen, all of the effort invested to improve the quality of the losing policy is wasted.

The Politics of Asymmetric Extremism

I now examine 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 I 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 is also first-order stochastically higher quality and better for the DM*
- *the extremist's policy is strictly more likely to be chosen*

Recall that 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 despite being better aligned with the DM. The extremist also develops a first-order stochastically more extreme policy than the moderate. Surprisingly, however, his policy is actually strictly more likely to be chosen than the moderate's, because it is so much higher quality so as to be first-order stochastically better for the DM *despite* its greater extremism. What explains the extremist's dominance of the policy process despite his more ideologically-extreme policy? First, the extremist is advantaged because his extremism makes him *more motivated* to invest in enough quality to compensate the DM for a more extreme policy. Second, in my model the extremist is still free to *strategically* moderate his policy if doing so is necessary to protect his interests.

I next consider what happens to equilibrium policies as one developer *becomes* more extreme.

Proposition 3. *If developer i becomes intrinsically more extreme (higher $|x_i|$), his* ***own strategy*** *and his* ***opponent's strategy*** *change as follows:*

(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 DM, 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 first-order stochastic change to his policy's quality or appeal to the DM, 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 likelihood of participation and his opponent's likelihood of 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 my model shares with the all-pay contest – that *balance* in the participants' motivation and abilities maximizes the likelihood that both players participate in the contest.

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 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 DM* despite being first-order stochastically more extreme, an effect that contrasts starkly with the asymmetric all-pay contest (where the strategy of the stronger player does not change as he becomes even stronger). Further, his opponent becomes strictly less likely to be active (as in the all pay contest), but also crafts a first-order stochastically better policy for the DM *conditional on* being active (in contrast to the all-pay contest). These effects result from the 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 my model, the 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.[14] In my model, in contrast, it causes the stronger player to behave more ideologically aggressively, which raises the ideological stakes of the contest and reinvigorates the weaker player's desire to participate. To prevent this reinvigoration in equilibrium, an increasingly extreme developer must therefore *also* craft an increasingly appealing policy to the DM that is more difficult to beat.

With these unusual effects in mind, I last examine the consequences of a developer i becoming more extreme for the other players' welfare. I 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*

The effect of unilateral extremism on a competitor's welfare is thus unambiguous; despite the greater

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

quality of the extremist's policy, the moderate is harmed. The reason is that this greater quality is insufficient to compensate the moderate for the policy's greater extremism. While intuitive, this effect also differs from the classic 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 DM's welfare is as follows.

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

- *If the developers are symmetrically capable and extreme ($|x_i| = |x_{-i}|, \alpha_i = \alpha_{-i}$) and developer i becomes intrinsically more extreme, the DM'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 DM's utility approaches infinity*

Characterizing the precise local effect of a more-extreme developer is difficult due to countervailing forces of the extremist's greater activity and the moderate's greater discouragement. However, the broader relationship is simple and striking; the DM strongly benefits from unilateral extremism. If the developers begin symmetrically capable and extreme and one becomes a bit more extreme, the DM benefits despite the resulting imbalance in participation. And as a developer becomes increasingly extreme, the DM becomes increasingly better off, even though his competitor also becomes vanishingly likely to participate. Unilateral extremism thus strongly benefits the DM, 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 DM due to the discouragement effect (Hillman and Riley (1989)). In my model, part of the discouragement effect is still present – the weaker player (the moderate) becomes less active due to 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 positive attributes are sufficient enough to outweigh the cost to the DM of the moderate's discouragement.[15]

The Politics of Asymmetric Ability

I next 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 I call the more capable developer "the expert" and the other "the amateur." My main finding is that asymmetric ability and extremism are *observationally equivalent*.[16]

Proposition 6. *If the developers are equally extreme but i is more skilled ($|x_i| = |x_{-i}|, \alpha_i < \alpha_{-i}$),*

- *the expert always develops a new policy, while the amateur only sometimes does so*

[15] 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) DMs are harmed by asymmetries because of the discouragement effect – my model is a notable exception. However, if there is enough "noise" in the outcome, then the DM 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 DM 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.

[16] Note that there is not an exact isomorphism between extremism and ability in the policy development model (unlike the all pay contest) because the developers intrinsically value quality, so an investment in quality is only partially "all-pay." Formally, in the generalized policy contest solved in Lemma 3, extremism and ability only become isomorphic for player i when he places no intrinsic value on the winner's effort ($\gamma_i = 0$). If neither player values the winner's effort then my results about the benefits of asymmetries remain, but the equilibrium characterization simplifies substantially.

- *the expert's policy is first-order stochastically more extreme than the amateur's policy, but is also first-order stochastically higher quality and better for the DM*

- *the expert's policy is strictly more likely to be chosen*

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 (first-order stochastically) when 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.[17]

The observational equivalence between asymmetric extremism and ability also extends to most of the consequences of one developer becoming more skilled.

Proposition 7. *If developer i becomes more skilled (lower α_i), his* ***own strategy*** *and his* ***opponent's strategy*** *change as follows:*

(Own strategy)

- *if he previously did not always develop a policy, he becomes strictly more likely to do so*

[17]Note that there is not an exact isomorphism between extremism and ability in the policy development model (unlike the all pay contest) because the developers intrinsically value quality, so an investment in quality is only partially "all-pay." Formally, in the generalized policy contest solved in Lemma 3, extremism and ability only become isomorphic for player i when he places no intrinsic value on the winner's effort ($\gamma_i = 0$). If neither player values the winner's effort then my results about the benefits of asymmetries remain, but the equilibrium characterization simplifies substantially.

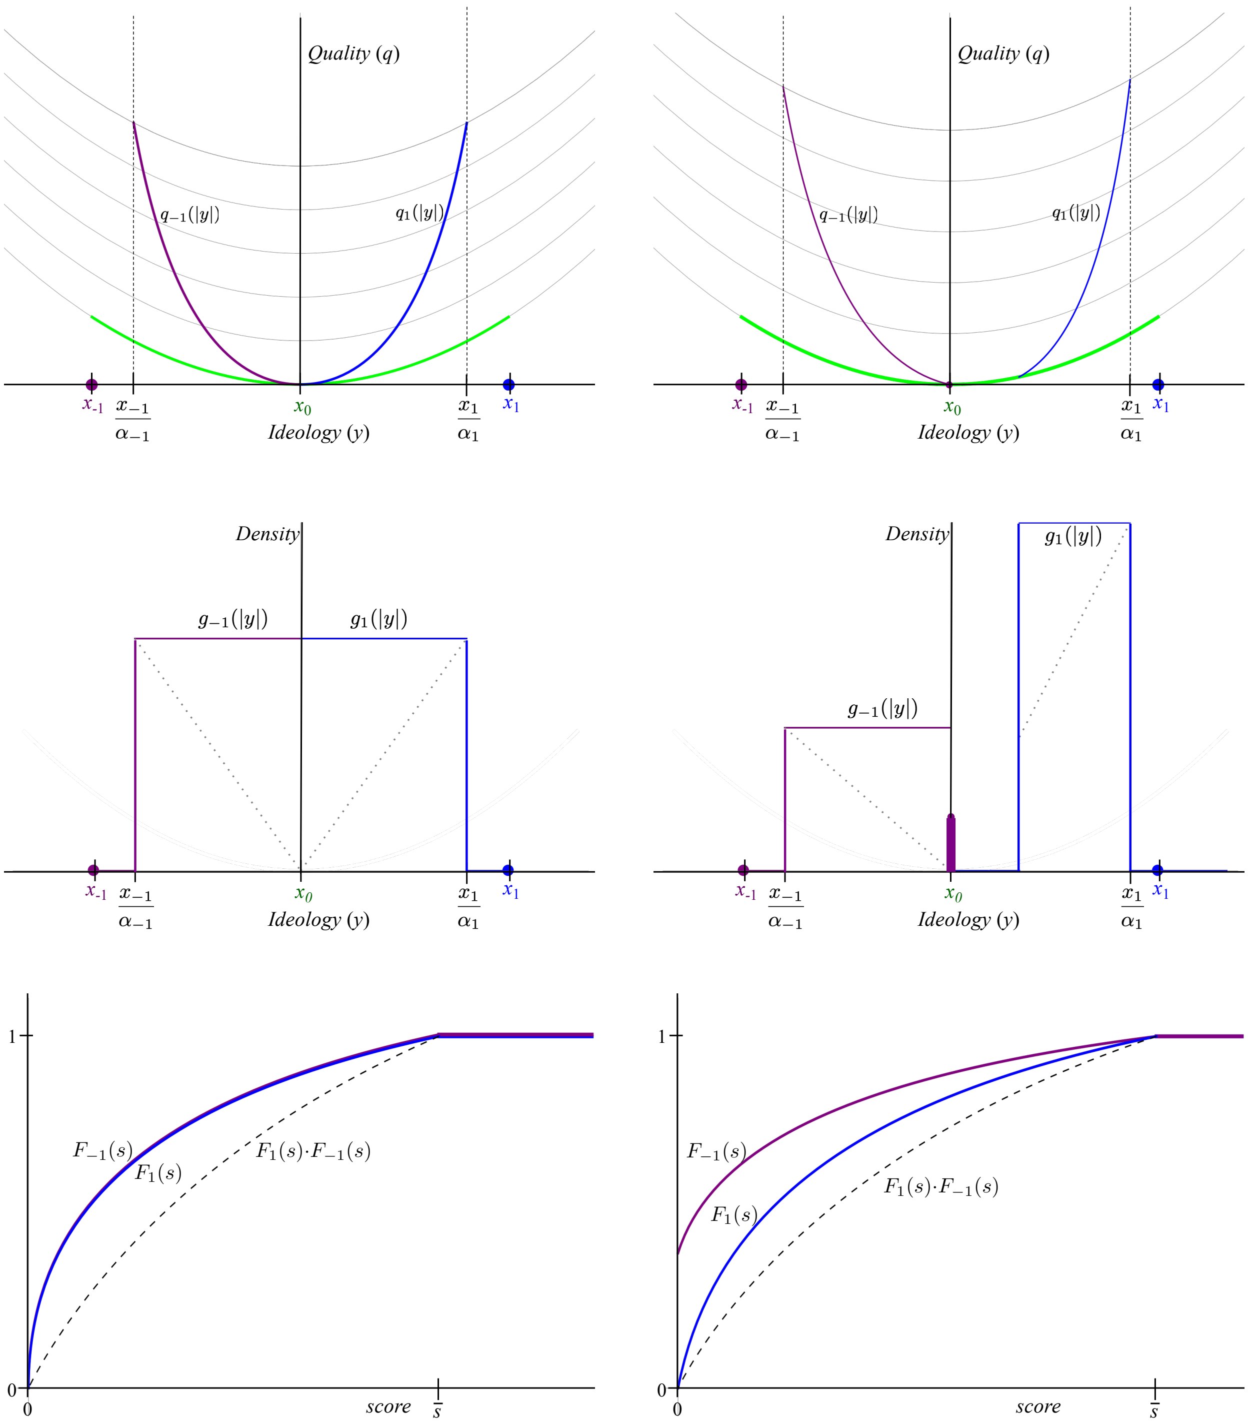

Figure 3: *Equilibrium with Symmetric Developers vs. Asymmetric Ability*

- *his policy becomes first-order stochastically more extreme*
- *his policy becomes first-order stochastically higher quality, better for the DM, 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 first-order stochastic change to his policy's quality or appeal to the DM, but his policy becomes strictly less likely to be chosen*

Finally, 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 DM. 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 latter property 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 conceptually 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 exploit areas of agreement to gain support from the DM rather than the opposition.

I last examine the DM's welfare.

Proposition 8. *If the developers begin symmetric ($|x_i| = |x_{-i}|, \alpha_i = \alpha_{-i}$) and developer i unilaterally becomes more skilled ($\alpha_i \to 1$), the DM's utility (i) locally increases, and (ii) approaches a finite limit strictly greater than her utility under symmetry provided $\alpha_{-i} \geq \underline{\alpha} \approx 1.0347$.*

The DM thus benefits when a developer becomes unilaterally more skilled, even though he crafts a more extreme policy and his competitor participates less. Finally, 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 of participating, but the DM still benefits from his potential participation (in the sense that her utility is bounded away from her utility under monopoly). The DM's utility doesn't increase unboundedly, however; rather, it approaches a strictly positive value that is generally higher than her utility under symmetry.

Strategic Extremism, Policy Spillovers, and Robustness

To sharpen the intuition for why asymmetries can benefit the DM in my model, I last consider several variants and compare them to the main model (see Appendix E for details). For the subsequent results I say that *asymmetries always harm the DM* if she becomes strictly worse off whenever one of two symmetric developers becomes differentially extreme or capable.

Strategic Extremism A key property of my model is that the developers strategically choose the extremism of their policy. To see this, consider a variant where each developer $i \in \{-1, 1\}$ may only target a fixed ideology $y_i = iy \neq 0$ (further assume that the DM has no outside option).[18] For analytical clarity I reparameterize the score so that $s(y_i, q_i) = \lambda q_i$; that is, the score is now the utility *gain* that the DM gets from a developer i's policy with a quality investment, which fully

[18] 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)).

determines the DM's preferred policy when the developers are constrained to craft equally extreme policies. Then developer i's utility from policy (s_i, iy) simplifies to

$$-\alpha_i s_i + F_{-i}(s_i) \cdot (s_i + 4x_i y \cdot i) + \int_{s_i}^{\infty} s_{-i} f_{-i}(s_{-i})\, ds_{-i} - (x_i + iy)^2 . \tag{15}$$

Equation 15 reveals that this variant is nearly isomorphic to a standard all-pay contest (with player i's prize value equal to $4|x_i|y$) with the caveat that it is as if effort is productive. While this model does not fit into the all-pay contest framework of Siegel (2009) (since the payoff from losing is endogenous), it is a special case of my generalized policy contest with no endogenous ideological component ($\theta = 0$). Applying equation 14, the DM's equilibrium expected utility simplifies to:

$$U_{DM}^{*} = \int_0^{\tilde{r}_k} \left(1 - \prod_j W_j(r)\right) dr,$$

where the residual reach functions are now $\tilde{r}_i(W_i) = \int_{W_i}^{1} \frac{4|x_i|y}{\alpha_i - F} dF$. This expression straightforwardly decreases when the stronger developer $-k$ becomes more extreme (higher $|x_{-k}|$) or less capable (lower α_{-k}), because in the absence of an endogenous ideological component, only the terms capturing the impact of the discouragement effect remain. The result is the following.

Proposition 9. *If the developers are constrained to craft policies with fixed ideologies $y_i = iy \neq 0$ and the DM has no outside options, then asymmetries always harm the DM.*[19]

Policy Spillovers Another crucial property of my model is that the developers care about policy when they lose. To see this, consider a variant where each developer can still choose the ideology of his policy, but receives a fixed payoff from losing.[20] Then the objective function of a developer i is:

$$-(\alpha_i - F_{-i}(s))\, s + F_{-i}(s) \cdot 2x_i y_i - \alpha_i y_i^2 - x_i^2 .$$

[19] This variant easily generalizes to allow a developer to value quality at $\lambda \gamma_i q_i$, so that the residual reach function is $\tilde{r}_i(W_i) = \int_{W_i}^{1} \frac{4|x_i|y}{\alpha_i - \gamma_i F} dF$, which is strictly increasing in γ_i. An interesting implication is that the DM is harmed by an asymmetry arising from one developer valuing quality *more*.

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

Maximizing first with respect to y_i, the optimal ideology to target is $y_i(s) = F_{-i}(s)\frac{x_i}{\alpha_i}$ (as in the policy development model). As in the generalized policy contest, the developers' score CDFs must be continuous and strictly increasing over a common interval $[0, \bar{s}]$ with $F_{-k}(0) = 0$ for some k. Differentiating w.r.t. s and setting equal to 0 yields that they must satisfy the differential equations

$$f_k(s) = \frac{\alpha_{-k} - F_k(s)}{2x_{-k}y_{-k}(s)} \quad \text{and} \quad f_{-k}(s) = \frac{\alpha_k - F_{-k}(s)}{2x_k y_k(s)} \;\; \forall s \in [0, \bar{s}]. \tag{16}$$

This variant more closely resembles my model, in that each developer i's ideological stakes $2x_i y_i(s)$ are endogenous; but this is nevertheless insufficient to yield a benefits from asymmetries. The reason is that these stakes depend 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, although the stronger developer $-k$ will still become more ideologically aggressive when the weaker developer k starts to participate less, this aggression does not raise developer k's *own* stakes (since he does not care about how he loses) and pull him back in. I again recover the discouragement effect in full.

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

Main Model I now return to the main model, and reexamine how the *combination* of strategic extremism and policy spillovers breaks the discouragement effect. Starting from equation 10 in the analysis of the generalized policy contest and considering the special case of the policy-development model ($\gamma_{-1} = \gamma_1 = \theta$), the developers' score CDFs must satisfy

$$f_k(s) = \left(\frac{2|x_{-k}|\sum_j |y_j(s)|}{\alpha_{-k} - F_k(s)}\right)^{-1} \quad \text{and} \quad f_{-k}(s) = \left(\frac{2|x_k|\sum_j |y_j(s)|}{\alpha_k - F_{-k}(s)}\right)^{-1} \quad \forall s \in [0, \bar{s}]$$

A developer's ideological stakes $2|x_i|\sum_j |y_j(s)|$ at a score s are now endogenous to both his own policy $y_i(s)$ and his opponent's policy $y_{-i}(s)$. Now 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 $2|x_k| \sum_j |y_j(s)|$ 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 *own* participation $F_k(s)$. Intuitively, should the weaker developer k start to drop out of the contest, the stronger developer $-k$ will react with more ideological aggression (higher $|y_{-k}(s)|$), thereby raising the stakes to k and *pulling him back in.* To ensure that the weaker developer remains willing to target scores in $[0, \bar{s}]$ (i.e., to keep him from rushing all the way 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 $\forall s$). That is, the stronger developer *must craft better policies for the DM* to incentivize his weaker opponent to participate less, despite the fact that doing so will trigger more ideological aggression. The combined effect for the DM is beneficial, despite increasingly extreme policies being crafted by an increasingly dominant developer.

Quadratic Costs and Linear Benefits In the main model an increasingly extreme developer enjoys exponentially-increasing gains from ideological concessions, but pays only a linear cost to craft the quality used to extract them. This raises the question of whether our results about the benefits of asymmetries are sensitive to these functional forms. To verify this is not the case, I last suppose that 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 modified 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 4). 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 DM for ideological movements holding score fixed becomes $\alpha_i (s + y)$ (instead of $2\alpha_i y$), which depends directly on the score s being targetted. Intuitively, with quadratic quality costs it becomes more expensive to compensate the DM 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 differs in an in an interesting way from our main model; 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. However, it shares the crucial property that that a developer becomes more ideologically aggressive ($|y_i(s)|$ increases) when his opponent participates less (higher $F_{-i}(s)$).

As in the main 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

$$f_k(s) = \left(\frac{(x_{-k} - k)\, y_{-k}(s) - (x_{-k} + k)\, y_k(s)}{\alpha_{-k}(s - k y_{-k}(s)) - F_k(s)}\right)^{-1} \text{ and } f_{-k}(s) = \left(\frac{(x_k + k)\, y_k(s) - (x_k - k)\, y_{-k}(s)}{\alpha_k (s + k y_k(s)) - F_{-k}(s)}\right)^{-1} \tag{17}$$

for all $s \in [0, \bar{s}]$. This lacks the simplicity of the main model; both because the optimal ideology might be the "corner" solution of the DM's ideal, and because the score s also enters the differential equations directly (instead of indirectly via the score CDFs). It thus does not appear to admit a closed form solution in the asymmetric case (equilibria are computed numerically – see Appendix for details). Nevertheless, it retains all the features that generate a benefit from asymmetries; as the stronger developer $-k$ becomes yet stronger the weaker developer k must become discouraged, which then elicits more ideological aggression from the stronger developer and pulls the weaker developer back in. 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 DM despite the discouragement of his competitor, potentially benefitting the DM.

Figure 4 depicts an example in which asymmetric extremism benefits the DM; as before the left panels depict symmetric developers, while the right panels depict a more extreme right developer. The first row depict the policies the developers randomize over using dark curves (note the DM's indifference curves are now linear); positive-quality policies no longer necessarily diverge from the DM's ideal, nor are higher score policies within a mixed strategy necessarily more extreme. The second row depicts the developers' equilibrium score CDFs using dark curves, as well as the CDF of the maximum score policy (i.e., the DM's final utility) with a dashed curve. A more-extreme right developer crafts a more ideologically extreme policy that is also better for the DM (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 DM's equilibrium utility (i.e., a rightward shift in the dashed curve).

Numerical results broadly illustrate a benefit from asymmetries in the quadratic-linear variant. Figure 5 is a contour plot of the *net benefit* to the DM 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 DM benefits from the right developer's greater extremism, whereas in the red regions she is harmed. When the left developer is relatively moderate the DM 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 DM always benefits from a more-extreme right developer x_1. Numerical results for the effects of asymmetric ability are broadly similar – the DM 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 DM's utility is large.

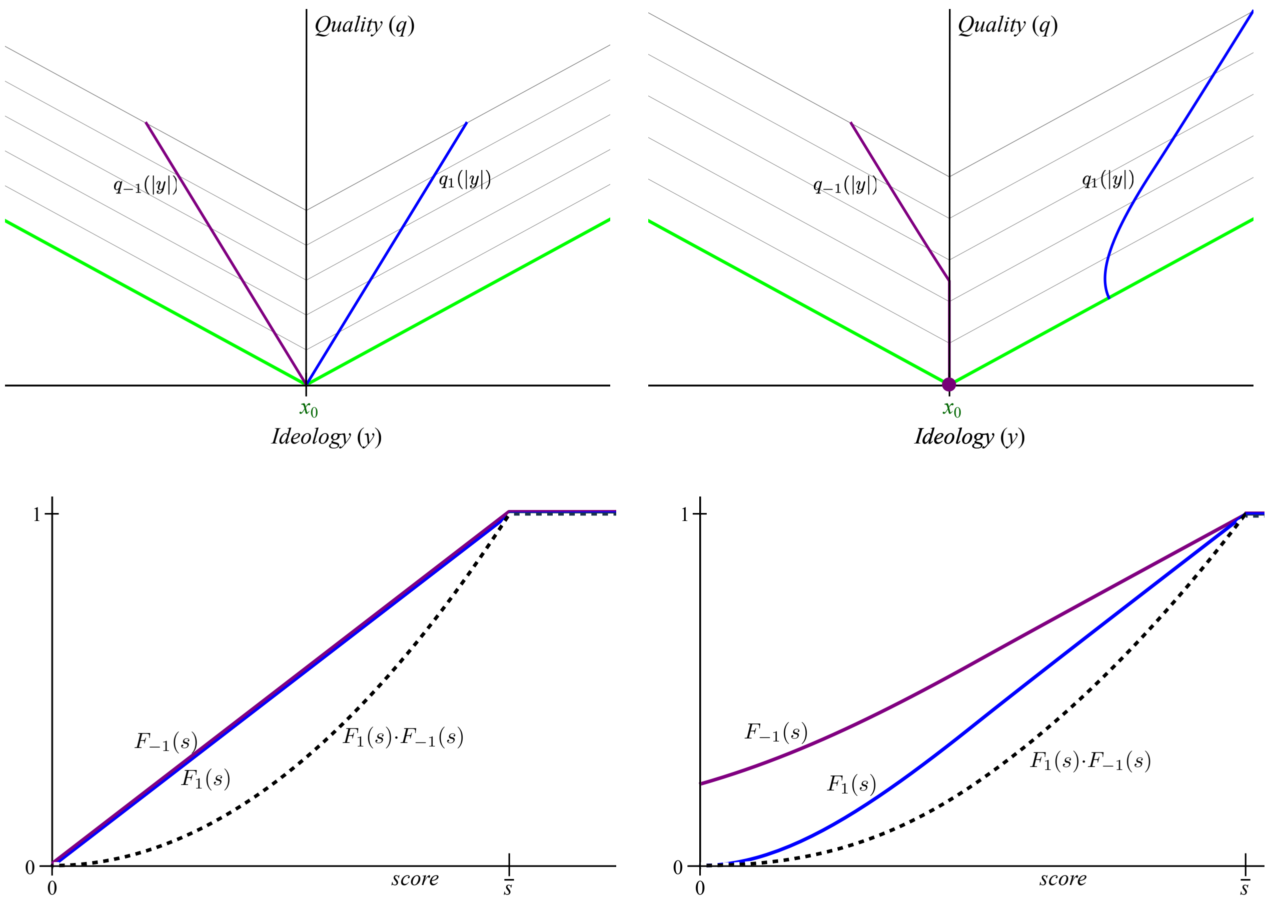

Figure 4: *Equilibrium with Symmetric vs. Asymmetric Extremism (Linear Ideological Preferences and Quadratic Costs of Quality)*

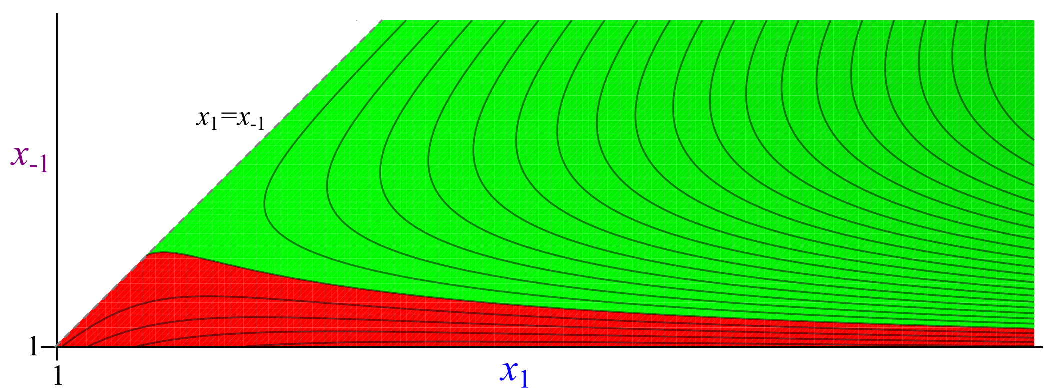

Figure 5: *Net benefit to DM of right developer being more extreme than left developer*

Discussion and Conclusion

This paper studies policy development by competing developers who differ in their ideological extremism and/or ability at crafting high quality policies. I 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. My model thus provides a novel rationale for how ideological extremism may come to dominate policymaking (as in Osborne, Rosenthal and Turner (2000)) that applies even when policy competition is *productive* rather than driven by 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 my findings, it is worth briefly remarking on the boundaries of my 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. I conjecture that this effect would similarly extend to the benefits of asymmetries. Indeed, the driving force of my 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 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 my model requires that such destructive means of policy influence be absent, relatively costly to employ (as compared to the productive means I study), or prohibited by either formal rules or shared norms of governance.

Finally, my 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 a 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.[21] 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

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

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? 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? I 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.

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.

Gilligan, Thomas W., and Keith Krehbiel. 1989. “Asymmetric Information and Legislative Rules with a Heterogeneous Committee.” *American Journal of Political Science*, 33(2): 459–490.

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 derives the contest formulation of the policy development model, and introduces the generalized policy contest subsuming both that model and the all pay contest. Appendix B is a preliminary analysis of the generalized policy contest concluding with a statement of necessary and sufficient conditions for equilibrium. Appendix C derives the closed-form characterization of the equilibrium of the generalized policy contest and 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 A-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 my results.

A The Model(s)

We begin with a (slightly) more general formulation of the policy development 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 Contest Formulation

We observe that the policy development model is just 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 2.

1. *Player i's utility for policy (s, y) is $U_i\left(y + X_D, \frac{s+y^2}{\lambda}\right) = x_i^2 + V_i(s, y)$, where*

$$V_i(s, y) = s + 2x_i y$$

and $x_i = X_i - X_D$ 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.2 The Generalized Policy Contest

We now introduce the generalized policy contest, which nests the policy development model introduced in the preceding subsection.

Definition 3. *In the generalized policy contest*

1. *Two bidders* $i \in \{-1, 1\}$ *simultaneously choose bids* (s_i, y_i) *at up-front cost* $c_i(s_i, y_i) = \alpha_i \left(s_i + y_i^2\right)$*; every bid must have nonnegative cost* $c_i(s_i, y_i) \geq 0$.

2. *The DM chooses a winner* $w \in \{-1, 0, 1\}$*, where* (s_0, y_0) *is a reserve bid satisfying* $s_0 \leq 0$.

Final utilities in the game are as follows

- *The DM's utility for choosing winner* w *is* s_w

- *Bidder* i*'s final utility from outcome* $(w, \mathbf{s}, \mathbf{y})$ *is* $-c_i(s_i, y_i) + V_i(\mathbf{1}_{i=w}, s_w, y_w)$*, where*

$$V_i(\rho_i, s_w, y_w) = \gamma_i s_w + 2x_i((1-\theta) \cdot i\rho_i + \theta y_w)$$

The generalized policy contest introduces three additional continuous parameters $(\gamma_1, \gamma_{-1}, \theta)$ to build a family of models that includes both the policy development model and the all pay contest.

B Preliminary Analysis of the Generalized Policy Contest

In the generalized policy contest, a bidder's pure strategy (s_i, y_i) is a two-dimensional element of $\mathbb{B} \equiv \left\{(s, y) \in R^2 \mid s + y^2 \geq 0\right\}$ (bids with nonnegative up-front costs). 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.[22]

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 bid (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 any bid with $s_i > 0$ 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 bids. Now

$$\bar{\Pi}_i(s_i, y_i; \sigma_{-i}) = -\alpha_i\left(s_i + y_i^2\right) + F_{-i}(s_i) \cdot V_i(1, s_i, y_i) + \int\limits_{s_{-i} > s_i} V_i(0, s_{-i}, y_{-i})\, d\sigma_{-i}. \tag{A.1}$$

[22] 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 first term is the up-front cost of the bid. The second term is the probability i wins times his utility for winning. The third term is i's utility should he lose to $-i$, which requires integrating over all the bids 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 4. *For any score $s_i > 0$ where $F_{-i}(\cdot)$ has no atom, $(s_i, y_i^*(s_i))$ with $y_i^*(s_i) = F_{-i}(s_i) \cdot \frac{x_i}{\alpha_i}\theta$, is bidder i's strictly best score-s_i bid.*

Proof: Straightforward. QED

Lemma 4 states that at almost every score $s_i > 0$, bidder i's unique best ideology is just a weighted average of the bidder's and DM's weights on ideological shifts $\left(\left(1 - \frac{1}{\alpha_i}\right) \cdot 0 + \frac{1}{\alpha_i} \cdot x_i = \frac{x_i}{\alpha_i}\right)$, multiplied by both the probability $F_{-i}(s_i)$ i's opponent makes a lower-score bid and the degree θ to which the prize has an endogenous ideological component. Note that i's optimal ideology does not depend directly on his opponent $-i$'s ideologies, since a bid'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 $F_{-i}(s_i)$ that bidder i wins the contest, since i's utility conditional on winning is additively separable in score and ideology.

The second lemma establishes that at least one of the bidders is always *active*, in the sense of choosing a bid with a strictly positive score (note that all strictly-positive score bids have strictly positive up-front costs).

Lemma 5. *In equilibrium $F_j(0) > 0$ for at most one j.*

Proof: Suppose not, so $F_i(0) > 0 \; \forall i$ in some equilibrium. Let U_i^* denote bidder i's equilibrium utility; by the definition of equilibrium this can be achieved by mixing according to his strategy *conditional on* crafting a score-$s_i \leq 0$ policy. Now conditional on *both* bidders crafting score-$s \leq 0$ policies, let $\bar{y}^0$ denote the expected ideological outcome, $\bar{s}^0 \leq 0$ denote the expected score outcome, and $\bar{\rho}_i^0$ denote i's expected probability of winning. Using that $V_i(\rho_i, s_w, y_w)$ is linear in all three parameters, we then have

$$U_i^* = -\alpha_i \cdot E\left[s_i + y_i^2 | s_i \leq 0\right] + F_{-i}(s_i) \cdot V_i\left(\bar{\rho}_i^0, \bar{s}^0, \bar{y}^0\right) + \int\limits_{s_{-i} > s_i} V_i(0, s_{-i}, y_{-i})\, d\sigma_{-i}.$$

If $\theta = 0$ (there is no endogenous ideological component to the prize) then then $\exists j$ s.t. $\bar{\rho}_j^0 < 1$ so $V_j\left(\bar{\rho}_j^0, \bar{s}^0, \bar{y}^0\right) < V_j(1, 0, 0)$, implying $U_j^* < \bar{\Pi}_j(0, 0; \sigma_{-j}) = \bar{\Pi}_j\left(0, y_j^*(0); \sigma_{-j}\right)$ so this j has a strictly profitable deviation. If $\theta > 1$ (there is an endogenous ideological component to the prize) then since $x_{-1} < 0 < x_1$ there $\exists j$ s.t. $V_j\left(\bar{\rho}_j^0, \bar{s}^0, \bar{y}^0\right) \leq V_j(1, 0, 0)$ implying $U_j^* \leq \bar{\Pi}_j(0, 0; \sigma_{-j}) < \bar{\Pi}_j\left(0, y_j^*(0); \sigma_{-j}\right)$ with the last inequality strict since $F_{-j}(0) > 0$ by assumption, so again this j has a strictly profitable deviation. QED

The third Lemma shows there is 0 probability of ties at a strictly positive score in equilibrium.

Lemma 6. *In equilibrium there is 0-probability of a tie at scores $s > 0$.*

Proof: Suppose not, so each bidder's strategy generates an atom of size $p_j^{\hat{s}} > 0$ at some $\hat{s} > 0$. Let $\bar{y}_j^{\hat{s}}$ denote j's expected ideology conditional on making a score $\hat{s}$ bid, let $\bar{y}^{\hat{s}}$ denote the expected ideological outcome conditional on *both* bidders making a score-$\hat{s}$ bid (which depends on the DM's tie breaking rules), and let $\bar{\rho}_j^0$ denote bidder j's expected probability of winning conditional on both bidders making a score-$\hat{s}$ bid (which again depends on the DM's tie breaking rules). Since bidder i achieves his equilibrium utility U_i^* by mixing according to his strategy conditional on making a score-$\hat{s}$ bid, we have (again using that $V_i(\rho_i, s_w, y_w)$ is linear) that $U_i^* =$

$$-\alpha_i\left(\hat{s} + \bar{y}_i^{\hat{s}} + Var_{\hat{s}}[y_j]\right) + \left(F_{-i}(\hat{s}) - p_{-i}^{\hat{s}}\right) V_i\left(1, \hat{s}, \bar{y}_i^{\hat{s}}\right) + p_{-i}^{\hat{s}} V_i\left(\bar{\rho}_i^{\hat{s}}, \hat{s}, \bar{y}^{\hat{s}}\right) + \int_{s_{-i} > s_i} V_i(0, s_{-i}, y_{-i})\, d\sigma_{-i},$$

where $Var_{\hat{s}}[y_i]$ denotes the variance of i's choice of ideology conditional on making a score $\hat{s}$ bid.

Suppose first that $\theta = 0$ (there is no endogenous ideological component to the prize). Then $\exists j$ s.t. $\bar{p}_j^0 < 1$. Since $V_j(1, \hat{s}, y) = V_j(1, \hat{s}, 0)$ when $\theta = 0$ and $Var_{\hat{s}}[y_j] \geq 0$ we have $U_j^* <$

$$-\alpha_j \hat{s} + F_{-j}(s) V_j(1, \hat{s}, 0) + \int_{s_{-j} > s_j} V_j(0, s_{-j}, y_{-j})\, d\sigma_{-j} = \bar{\Pi}_j(0, 0; \sigma_{-j})$$

so j has a strictly profitable deviation.

Suppose next that $\theta > 0$ (there is an endogenous ideological component to the prize). Since $x_{-1} < 0 < x_1$ we have $V_j\left(\hat{s}, \bar{y}^{\hat{s}}\right) \leq V_j\left(\hat{s}, 0\right)$ for at least one j, so (using also that $\bar{\rho}_j^{\hat{s}} \leq 1$ and $Var_{\hat{s}}\left[y_j\right] \geq 0$) we have $U_j^* \leq$

$$\begin{aligned}
& -\alpha_j\left(\hat{s}+\bar{y}_j^{\hat{s}}\right)+\left(F_{-j}(\hat{s})-p_{-j}^{\hat{s}}\right) V_j\left(1, \hat{s}, \bar{y}_j^{\hat{s}}\right)+p_{-j}^{\hat{s}} V_j(1, \hat{s}, 0)+\int_{s_{-j}>s_j} V_j\left(0, s_{-j}, y_{-j}\right) d \sigma_{-j} \\
= & -\alpha_j\left(\hat{s}+\bar{y}_j^{\hat{s}}\right)+F_{-j}(\hat{s}) \cdot V_j\left(1, \hat{s},\left(1-\frac{p_{-j}^{\hat{s}}}{F_{-j}(\hat{s})}\right) \bar{y}_j^{\hat{s}}\right)+\int_{s_{-j}>s_j} V_j\left(0, s_{-j}, y_{-j}\right) d \sigma_{-j} .
\end{aligned}$$

If $\bar{y}_j^{\hat{s}} = 0$ then the preceding is equal to $\bar{\Pi}_j\left(\hat{s}, 0; \sigma_{-j}\right)$, which in turn is $< \bar{\Pi}_j\left(\hat{s}, y_j^*(\hat{s}); \sigma_{-j}\right)$ (since $F_{-j}(\hat{s}) > 0$) so j has a strictly profitable deviation. If $\bar{y}_j^{\hat{s}} \neq 0$ then the preceding is strictly less than $\left(1-\frac{p_{-j}^{\hat{s}}}{F_{-j}(\hat{s})}\right) \bar{\Pi}_j\left(\hat{s}, \bar{y}_j^{\hat{s}}; \sigma_{-j}\right)+\left(\frac{p_{-j}^{\hat{s}}}{F_{-j}(\hat{s})}\right) \bar{\Pi}_j\left(\hat{s}, 0; \sigma_{-j}\right)$, which is j's utility if he were to instead bid $(\hat{s}, 0)$ with probability $\frac{p_{-j}}{F_{-j}(\hat{s})}$ and bid $\left(\hat{s}, \bar{y}_j^{\hat{s}}\right)$ with the remaining probability (and always win ties). Thus j again has a strictly profitable deviation. QED

Lemmas 4 – 6 jointly imply that in equilibrium, bidder i can compute his expected utility *as if* his opponent only makes bids of the form $\left(s_{-i}, y_{-i}^*\left(s_{-i}\right)\right)$. The utility from *any* bid $\left(s_i, y_i\right)$ with $s_i > 0$ where $-i$ has no atom (or a tie would be broken in i's favor) is therefore

$$\bar{\Pi}_i^*\left(s_i, y_i; F\right)=-\alpha_i\left(s_i+y_i^2\right)+F_{-i}\left(s_i\right) \cdot V_i\left(1, s_i, y_i\right)+\int_{s_i}^{\infty} V_i\left(0, s_{-i}, y_{-i}^*\left(s_{-i}\right)\right) d F_{-i}. \tag{A.2}$$

Bidder i's utility from the *best* bid with score s_i is $\bar{\Pi}_i^*\left(s_i, y_i^*\left(s_i\right); F\right)$, henceforth denoted $\bar{\Pi}_i^*\left(s_i; F\right)$.

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

Lemma 7. *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 $\bar{\Pi}_i\left(\hat{s}, y_i; F\right) - \bar{\Pi}_i\left(s, y_i; F\right) = -\left(\alpha_i - \gamma_i F_{-i}(\hat{s})\right) \cdot (\hat{s} - s) < 0$, 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 bidder has support just below, so $F_i(s) < F_i(\hat{s})\ \forall i,\ s < \hat{s}$ would imply both bidders 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 11. *Necessary conditions for SPNE are as follows:*

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

   (a) weakly negative score $s_i \leq 0$ *and* 0*-cost* $\left(c_i(s_i, y_i) = 0 \iff s_i + y_i^2 = 0\right)$

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

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 least one bidder* 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* $s \in [0, \bar{s}]$,

$$\frac{2\,|x_i| \cdot f_{-i}(s)}{\alpha_i - \gamma_i F_{-i}(s)} = \left((1-\theta) + \theta \sum_{j \in \{-1,1\}} |y_j(s)|\right)^{-1}$$

The above and $F_i(s) = 0\ \forall i, s < 0$ *are also 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 7) 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 7. (*Ideological Optimality*) At most one bidder k has $F_k(0) > 0$, so so $F_{-k}(0) = 0$. Bids with $s_k \le 0$ lose for sure and never influence a tie, and therefore must must have zero up-front cost $c_k(s_k, y_k) = 0$ with probability 1, yielding property (a). Atomless score CDFs $\forall s > 0$ implies $(s, y_i^*(s))$ is the strictly best score-s bid (by Lemma 4), yielding property (b). (*Sufficiency*) Necessary conditions imply all bids $(s, y_i^*(s))$ with $s \in (0, \bar{s}]$ yield a constant U_i^*. $F_{-k}(0) = 0$ implies k's strictly best score-0 bid 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 bidders' 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 $\Pi_i^*(s; F) - \Pi_i^*(\bar{s}; F) = -(\alpha_i - \gamma_i)(s - \bar{s}) < 0$. To see k cannot profitably deviate to $s_k \le 0$, $F_{-k}(0) = 0$ implies such bids lose and never influence a tie, and so yield utility $\le U_k^*$. To see $-k$ cannot profitably deviate to some (s_{-k}, y_{-k}) with $s_{-k} \le 0$, observe all such bids result in utility $\le \max\left\{\bar{\Pi}_{-k}(0, 0; F), \bar{\Pi}_{-k}(0, y_{-k}; F)\right\}$ (since conditional on $s_k \le 0$ such bids win with probability ≤ 1, result in a score outcome of $s = 0$, and an ideological outcome of either $y_k^*(0) = 0$ or y_{-k}); such bids thus yield utility $\le U_{-k}^*$. QED

Preliminary Observations about Equilibria Proposition 11 implies that all equilibria have a familiar form. At least one bidder (henceforth labelled k) has an opponent $-k$ who is always active ($F_{-k}(0) = 0$). Thus, competition not only strictly benefits the DM in expectation, but with probability 1. Bidder 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 the highest-score bid with $0-$cost ($s_k = 0, y_k = 0$), or any other 0-cost bid $c_k(s_k, y_k) = 0$ with strictly lower score $s_k < 0$; we interpret such costless bids with non-centrist ideologies (which are certain to lose in equilibrium) as position taking. However, any equilibrium with position taking is payoff-equivalent to one without it; we therefore focus on the latter for comparative statics.[23] When a bidder i is active, he mixes smoothly

[23] Bid profiles with position-taking by bidder k 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

over the ideologically-optimal policies $\left(s, F_{-i}\left(s\right) \frac{x_i}{\alpha_i}\theta\right)$ with scores in a common mixing interval $[0, \bar{s}]$ according to the CDF $F_i\left(s\right)$.[24] The differential equations characterizing the equilibrium score CDFs arise intuitively from the bidders' indifference condition over $[0, \bar{s}]$ as described in the main text.

C Equilibrium of the Generalized Policy Contest

The most critical step in deriving a unique closed form equilibrium characterization is to use the coupled system of differential equations characterizing a pair of equilibrium score CDFs $(F_{-1}(s), F_1(s)$ in Proposition 11 to derive a simple functional relationship that must hold between them.

Lemma 8. *In any SPNE,* $\tilde{r}_i\left(F_{-i}\left(s\right)\right) = \tilde{r}_{-i}\left(F_i\left(s\right)\right)\ \forall s \geq 0$, *where*

$$\tilde{r}_i\left(W_i\right) = \int_{W_i}^{1} \frac{2\left|x_i\right|}{\alpha_i - \gamma_i F} dF = \frac{2\left|x_i\right|}{\gamma_i} \log\left(\frac{\alpha_i - \gamma_i W_i}{\alpha_i - \gamma_i}\right)$$

Proof: Rearranging the differential equation in score optimality yields $\frac{2|x_i|f_{-i}(s)}{\alpha_i-\gamma_i F_{-i}(s)} = \frac{2|x_{-i}|f_i(s)\cdot}{\alpha_{-i}-\gamma_{-i}F_i(s)}$ $\forall s \in [0, \bar{s}] \to \int_{\hat{s}}^{\bar{s}} \frac{2|x_i|f_{-i}(s)}{\alpha_i-\gamma_i F_{-i}(s)} d\hat{s} = \int_{s}^{\bar{s}} \frac{2|x_{-i}|f_i(s)\cdot}{\alpha_{-i}-\gamma_{-i}F_i(s)} d\hat{s}\ \forall \hat{s} \in [0, \bar{s}]$; a change of variables and the boundary condition $F_i\left(\bar{s}\right) = 1$ yields $\int_{\hat{s}}^{\bar{s}} \frac{2|x_i|f_{-i}(s)}{\alpha_i-\gamma_i F_{-i}(s)} d\hat{s} = \int_{F_{-i}(s)}^{1} \frac{2|x_i|}{\alpha_i-\gamma_i F} dF = \tilde{r}_i\left(F_{-i}\left(s\right)\right)$. The relationship holds trivially for $s > \bar{s}$. QED

We refer to the property in Lemma 8 as the *residual reach equality.* To see why, observe that the decreasing function $\tilde{r}_i\left(W_i\right)$ captures how much of i's *relative* willingness to deviate to a score that wins with probability 1 is "left over" when he is already winning with probability W_i (since the relative marginal ideological benefit of moving the ideological outcome in his direction is $|x_i|$, and the relative net marginal cost of increasing score on a policy winning the contest with probability F is $\alpha_i - F$). We call $\tilde{r}_i\left(W_i\right)$ i's (relative) *residual reach at win probability* W_i. The residual reach equality $\tilde{r}_i\left(F_{-i}\left(s\right)\right) = \tilde{r}_{-i}\left(F_i\left(s\right)\right)$ states that at every score $s \geq 0$, both bidders must have the same amount

the DM's reserve bid (s_0, y_0). When $(s_0, y_0) = (0, 0)$ the necessary conditions are also sufficient.

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

of their relative reach "left over" given the implied win probabilities $W_i = F_{-i}(s)$ and $W_{-i} = F_i(s)$. It is easily verified that $\tilde{r}_i(0) > \tilde{r}_i(1) = 0$.

Usefully, the residual reach equality implies a simple method for going back and forth between the bidders' score CDFs in any equilibrium. Let

$$W_i(r_i) = \frac{\alpha_i}{\gamma_i} - \left(\frac{\alpha_i}{\gamma_i} - 1\right) e^{\frac{\gamma_i}{2|x_i|} r}$$

denote the inverse of $\tilde{r}_i(W_i)$ (which is decreasing in W_i); the residual reach equality then implies that $F_{-i}(s) = W_i(\tilde{r}_{-i}(F_i(s)))$ for all $s \in [0, \bar{s}]$.

For later proofs, it is also helpful to note the straightforward comparative statics of the residual reach function $\tilde{r}_i(\cdot)$ and its inverse $W_i(\cdot)$ in the underlying parameters.

Lemma 9. *Player i's residual reach function $\tilde{r}_i(W_i) = \int_{W_i}^{1} \frac{2|x_i|}{\alpha_i - \gamma_i F}$ is strictly increasing in $(|x_i|, \gamma_i)$ and strictly decreasing in α_i $\forall W_i \in [0, 1)$. The inverse residual reach function $W_i(r_i)$ shares these comparative statics $\forall r_i \in (0, \tilde{r}_i(0)]$.*

Proof: Comparative statics of $\tilde{r}_i(W_i)$ are straightforward from the definition. To see the inverse obeys the same comparative statics in some arbitrary parameter β_i observe that

$$\frac{\partial W_i(r_i; \beta_i)}{\partial \beta_i} = -\left(\frac{\partial \tilde{r}_i(W_i(r_i; \beta_i); \beta_i)}{\partial W_i}\right)^{-1} \left(\frac{\tilde{r}_i(W_i(r_i; \beta_i); \beta_i)}{\partial \beta_i}\right)$$

and $\frac{\partial \tilde{r}_i(W_i(r_i; \beta_i); \beta_i)}{\partial W_i} < 0$. QED.

C.1 Implications of the residual reach equality

The most useful implication of the residual reach equality is the identity of the "weaker" bidder k who has an always active opponent (i.e., $F_{-k}(0) = 0$).

Proposition 12. *In equilibrium $k \in \arg\min_j \{\tilde{r}_j(0)\}$*

Proof: Suppose $\tilde{r}_k(0) < \tilde{r}_{-k}(0)$; then $F_k(0) = 0$ and the residual reach equality would imply $F_{-k}(0) < 0$, a contradiction. Since $F_j(0) = 0$ for some j we must have $F_{-k}(0) = 0$ and $F_k(0) = W_{-k}(\tilde{r}_k(0)) > 0$. QED

The bidder k with an always active opponent (i.e., $F_{-k}(0) = 0$) is therefore the one with the lowest residual reach $\tilde{r}_i(0)$ at probability 0 – that is, with the lowest total relative reach. We henceforth call this quantity simply the relative reach and denote as $\tilde{r}_i = \tilde{r}_i(0)$.

C.1.1 Probabilities of participation and victory

The residual reach equality is also used to derive the probability $F_k(0)$ that bidder k is inactive, as well as the probability that he loses the contest, and perform comparative statics on both.

Proposition 13. *In equilibrium bidder k is inactive with probability $F_k(0) =$*

$$W_{-k}(\tilde{r}_k(0)) = \frac{\alpha_{-k}}{\gamma_{-k}} - \left(\frac{\alpha_{-k}}{\gamma_{-k}} - 1\right)\left(\frac{1}{1-(\alpha_k/\gamma_k)^{-1}}\right)^{\frac{|x_k|}{\gamma_k}\Big/\frac{|x_{-k}|}{\gamma_{-k}}}$$

and loses the contest with probability $\int_0^{\bar{s}} f_{-k}(s)\, F_k(s)\, ds =$

$$\int_0^1 W_{-k}(\tilde{r}_k(W))\, dW = \int_0^1 \left(\frac{\alpha_{-k}}{\gamma_{-k}} - \left(\frac{\alpha_{-k}}{\gamma_{-k}} - 1\right)\left(\frac{1-(\alpha_k/\gamma_k)^{-1}\cdot W}{1-(\alpha_k/\gamma_k)^{-1}}\right)^{\frac{|x_k|}{\gamma_k}\Big/\frac{|x_{-k}|}{\gamma_{-k}}}\right) dW$$

Both probabilities are decreasing in his own extremism $|x_k|$ and his opponent's costs α_{-k}, and increasing in his opponent's extremism $|x_{-k}|$ and his own costs α_k. Further, both approach 1 as $|x_k| \to 0$, $|x_{-k}| \to \infty$, $\alpha_k \to \infty$, or $\alpha_{-k} \to 1$.

Proof: The first statement is immediately implied by the residual reach equality. For the second statement the probability k loses the contest is $\int_0^{\bar{s}} f_{-k}(s)\, F_k(s)\, ds$; applying the residual reach equality this is $\int_0^{\bar{s}} W_{-k}(\tilde{r}_k(F_{-k}(s)))\, f_{-k}(s)\, ds$, and applying a change of variables from s to $F_{-k}(s)$ yields $\int_{F_{-k}(0)}^{F_{-k}(\bar{s})} W_{-k}(\tilde{r}_k(F_{-k})) \cdot dF_{-k}$, which simplifies to the stated expression (recalling that $F_{-k}(0) = 0$ and $F_{-k}(\bar{s}) = 1$). Finally, comparative statics and limit statements follow directly from previous observations on $\tilde{r}_i(\cdot)$ and $W_i(\cdot)$. QED

The probability that the weaker bidder k is inactive ($F_k(0)$) thus obeys the same comparative statics as his overall probability of losing the contest. Somewhat paradoxically, he becomes *less* likely to win when he is less extreme or his opponent is more extreme. More intuitively, he becomes more likely to win if he is more capable or his opponent is less capable.

C.1.2 Strength and Conditions for First-Order Stochastic Dominance

In the standard asymmetric two-player all-pay contest there is an unambiguously weaker bidder in two senses; he is both both less likely to bid, and also makes bids that are first-order stochastically worse for the DM. In the present contest there is always a weaker bidder in the first sense, but not necessarily in the second one; the reason is that cost advantages become magnified at higher scores when a bidder intrinsically values effort, since he is more likely to win and enjoy the intrinsic benefits of that effort. We thus have the following.

Proposition 14. *Developer i is dominated ($F_{-i}(s) < F_i(s)\ \forall s \in (0, \bar{s})$) i.f.f. he has a lower residual reach at every win probability W (i.e., $\tilde{r}_i(W) < \tilde{r}_{-i}(W)\ \forall W \in [0, 1)$).*

Proof: Lemma 8 and the residual reach function $\tilde{r}_i(W_i)$ being strictly decreasing when $W_i \in [0, 1)$ implies $sign\left(\tilde{r}_{-k}(F_{-k}(s)) - \tilde{r}_k(F_{-k}(s))\right) = sign\left(F_k(s) - F_{-k}(s)\right)\ \forall s \in [0, \bar{s})$, which straightforwardly yields the first statement; this argument also applies to the generalized policy contest. QED

Clearly, a bidder k who is *both* less extreme ($|x_k| \leq |x_{-k}|$) and less capable ($\alpha_k \geq \alpha_{-k}$) (with one strict) will be dominated if both intrinsically value score equally ($\gamma_k = \gamma_{-k} > 0$). However, when one bidder is cost advantaged while the other is more ideologically extreme, it is possible for the cost-advantaged bidder to sometimes be inactive but nevertheless have a higher probability of targetting the highest scores (and therefore not be dominated in a first order stochastic sense). Lastly, if neither bidder intrinsically values effort ($\gamma_i = \gamma_{-i} = 0$) it is easily verified that $\tilde{r}_i(W) = \frac{2|x_i|}{\alpha_i}(1 - W_i)$ so that $\tilde{r}_i(W) < \tilde{r}_i(W)\ \forall W \iff \frac{2|x_i|}{\alpha_i} = \tilde{r}_i < \frac{2|x_{-i}|}{\alpha_{-i}} = \tilde{r}_{-i}$; that is, there *is* an unambiguously weaker player (as in the standard all pay contest) determined in the same manner as in a standard all pay contest with prizes $2|x_j|$ and marginal costs α_j.

C.1.3 Distribution over ideologies

We last use the residual reach equality to derive the probability distribution over the *ideology* of each bidder's bids as follows.

Proposition 15. *Let $G_i(\delta) = \Pr(|y_i| \leq \delta)$ denote the probability that i makes a bid with ideology less extreme than δ. Then*

$$G_i(\delta) = W_{-i}\left(\tilde{r}_i\left(\frac{i\delta}{\theta x_i/\alpha_i}\right)\right) = \frac{\alpha_{-i}}{\gamma_{-i}} - \left(\frac{\alpha_{-i}}{\gamma_{-i}} - 1\right)\left(\frac{\theta|x_i| - \delta\gamma_i}{\theta|x_i| - \frac{\theta|x_i|}{\alpha_i}\gamma_i}\right)^{\frac{|x_i|}{\gamma_i}\Big/\frac{|x_{-i}|}{\gamma_{-i}}}$$

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: Bidder i's ideology at score s is $y_i(s) = F_{-i}(s)\frac{x_i}{\alpha_i}\theta$, so $F_{-i}(s_i(y)) = \frac{y}{\theta x_i/\alpha_i}$ where $s_i(y)$ is the inverse of $y_i(s)$. That is, the probability $-i$ crafts a policy with score $\leq s_i(y)$ is $\frac{y}{\theta x_i/\alpha_i}$. Now the probability $G(\delta)$ that i makes a bid with ideology less extreme than y is $F_i(s_i(i\delta))$, which is $= W_{-i}(\tilde{r}_i(F_{-i}(s_i(i\delta)))) = W_{-i}\left(\tilde{r}_i\left(\frac{i\delta}{\theta x_i/\alpha_i}\right)\right)$ from the residual reach equality. Comparative statics are then straightforward from properties of $r_i(\cdot)$ and $W_i(\cdot)$. QED

The ideological comparative statics of the policy development model thus extend to the generalized policy contest, except the edge case of no endogenous ideological component ($\theta = 0$).

C.2 Equilibrium Score CDFs and Payoffs

With the residual reach equality and the identity of the sometimes-inactive bidder k, we may now fully characterize equilibrium score CDFs $F_i(s)$ satisfying Proposition 11, which are shown constructively to be unique over over $s \geq 0$.

Proposition 16. *The unique score CDFs over $s \geq 0$ satisfying Proposition 11 are $F_i(s) = W_{-i}(\tilde{\mathbf{r}}(s))$ $\forall i$, where $\tilde{\mathbf{r}}(s)$ is the inverse of*

$$\mathbf{s}(r) = \int_r^{\tilde{r}_k(0)} \left((1-\theta) + \theta^2 \sum_{j\in\{-1,1\}} W_j(t)\frac{|x_j|}{\alpha_j}\right) dt.$$

The inverse score CDFs are $s_i(F_i) = \mathbf{s}(\tilde{r}_{-i}(F_i))\ \forall i$*, and the score targetted at each ideology is* $\mathbf{s}\left(\tilde{r}_i\left(\frac{y}{x_i/\alpha_i}\right)\right)$*. The function* $\mathbf{s}(r)$ *is strictly increasing in* x_i *and strictly decreasing in* α_i $\forall r \in [0, \tilde{r}_k(0))$*. Finally the maximum score is* $\bar{s} = \mathbf{s}(0)$*, which also obeys these comparative statics.*

Increasing a bidder'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 residual reach equality $\tilde{r}_i(F_{-i}(s)) = \int_{F_{-i}(s)}^{1} \frac{2|x_i|}{\alpha_i - \gamma_i F} dF = \tilde{\mathbf{r}}(s)\ \forall i, s$ for some $\tilde{\mathbf{r}}(s)$. We characterize the unique $\tilde{\mathbf{r}}(s)$ implying score CDFs $F_i(s) = W_{-i}(\tilde{\mathbf{r}}(s))$ and optimal ideologies $y_i(s) = W_i(\tilde{\mathbf{r}}(s)) \frac{x_i}{\alpha_i}\theta$ that satisfy score optimality. First observe that $\tilde{\mathbf{r}}'(s) = \tilde{r}_i'(F_{-i}(s)) f_{-i}(s) = -\frac{2|x_i| f_i(s)}{\alpha_i - \gamma_i F_{-i}(s)}$. Next we rewrite the the differential equations characterizing the equilibrium score CDFs so that we have $-\frac{2|x_i| f_i(s)}{\alpha_i - \gamma_i F_{-i}(s)} = -\left((1-\theta) + \theta \sum_{j \in \{-1,1\}} |y_j(s)|\right)^{-1}$. Substituting the preceding observations into both sides yields $\frac{1}{\tilde{\mathbf{r}}'(s)} = -\left((1-\theta) + \theta^2 \sum_j \frac{|x_j|}{\alpha_j} W_j(\tilde{\mathbf{r}}(s))\right)$, and rewriting in terms of the inverse $\mathbf{s}(\tilde{r})$ yields $\mathbf{s}'(\tilde{r}) = -\left((1-\theta) + \theta^2 \sum_j \frac{|x_j|}{\alpha_j} W_j(\tilde{r})\right)$. Lastly $\tilde{r}_k(F_{-k}(s)) = \tilde{\mathbf{r}}(s)$ and $F_{-k}(0) = 0$ imply the boundary condition $\mathbf{s}(\tilde{r}_k(0)) = 0$ so $\mathbf{s}(\tilde{r}) = \int_r^{\tilde{r}_k(0)} -\mathbf{s}'(\hat{r})\, d\hat{r} = \int_r^{\tilde{r}_k(0)} \sum_j \left((1-\theta) + \theta^2 \frac{|x_j|}{\alpha_j} W_j(t)\right) dt$. Now $\mathbf{s}(r)$ is increasing in $|x_i|$ and decreasing in α_i given previous observations about $W_j(t)$. QED

The maximum score $\bar{s}$ thus changes continuously with the parameters of both bidders, except in the edge case of no endogenous ideological component ($\theta = 0$). This contrasts starkly with the standard 2-player all pay contest, where the mixing interval is unaffected by the parameters of the stronger bidder. Increasing a bidder'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 bidder $-k$ 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 choosing a very high score , even while he is simultaneously less likely to be active.

C.2.1 Developer Payoffs

Using Proposition 16 we may calculate the bidders' equilibrium payoffs as follows.

Proposition 17. *Developer i's equilibrium utility*

$$\Pi_i^*\left(\bar{s};F^*\right)=2x_i\left(\left(1-\theta\right)i+\theta^2\frac{x_i}{2\alpha_i}\right)-\left(\alpha_i-\gamma_i\right)\bar{s}$$

is decreasing in his opponent's extremism and ability, and increasing in his own ability.

Proof: A bidder's equilibrium utility is straightforward since $\left(\bar{s},y_i^*\left(\bar{s}\right)\right)$ is in the support of their strategy and wins for sure. Comparative statics of a bidder i's utility on his opponent $-i$'s parameters follow immediately from previously-shown statics on $\bar{s}=s\left(0\right)$.

For comparative statics in own ability α_i, first differentiate with respect to α_i which yields:

$$\frac{\partial\left(\Pi_i^*\left(\bar{s};F^*\right)\right)}{\partial\alpha_i}=-\theta^2\left(\frac{|x_i|}{\alpha_i}\right)^2-\bar{s}-\left(\alpha_i-\gamma_i\right)\frac{\partial\bar{s}\left(\alpha_i\right)}{\partial\alpha_i}$$

Recall $\frac{1}{W_i'(r)}=r_i'\left(W_i\left(r\right)\right)=-\frac{2|x_i|}{\alpha_i-\gamma_iW_i(r)}$; using a change of variables from t to $W_i\left(t\right)$ we have

$$\begin{aligned}\bar{s} &= \int_0^{\tilde{r}_k(0)}\left(\left(1-\theta\right)+\theta^2\left(\frac{W_i\left(t\right)}{\alpha_i}\left|x_i\right|+\frac{W_{-i}\left(t\right)}{\alpha_{-i}}\left|x_{-i}\right|\right)\right)dt\\ &= \frac{2\left|x_i\right|}{\alpha_i}\int_{W_i(\tilde{r}_k(0))}^1\left(\left(1-\theta\right)+\theta^2\left(W\frac{|x_i|}{\alpha_i}+W_{-i}\left(r_i\left(W\right)\right)\frac{|x_{-i}|}{\alpha_{-i}}\right)\right)\left(\frac{\alpha_i}{\alpha_i-\gamma_iW}\right)dW\end{aligned}$$

Next observe that

$$\begin{aligned}\frac{\partial\bar{s}\left(\alpha_i\right)}{\partial\alpha_i} &= \int_0^{\tilde{r}_k(0)}\theta^2\left|x_i\right|\frac{\partial}{\partial\alpha_i}\left(\frac{W_i\left(t\right)}{\alpha_i}\right)dt+\mathbf{1}_{i=k}\cdot\frac{\partial\tilde{r}_i\left(0\right)}{\alpha_i}\cdot\left(\left(1-\theta\right)+\theta^2\cdot W_{-i}\left(\tilde{r}_i\left(0\right)\right)\frac{|x_{-i}|}{\alpha_{-i}}\right)\\ &= \frac{2\left|x_i\right|}{\alpha_i\left(\alpha_i-\gamma_i\right)}\int_0^{\tilde{r}_k(0)}\theta^2\frac{|x_i|}{\alpha_i}W_i'\left(t\right)dt-\mathbf{1}_{i=k}\cdot\frac{2\left|x_i\right|}{\alpha_i\left(\alpha_i-\gamma_i\right)}\cdot\left(\left(1-\theta\right)+\theta^2\cdot W_{-i}\left(\tilde{r}_i\left(0\right)\right)\frac{|x_{-i}|}{\alpha_{-i}}\right)\\ &= -\left(\frac{1}{\alpha_i-\gamma_i}\right)\frac{2\left|x_i\right|}{\alpha_i}\left(\int_{W_i(\tilde{r}_k(0))}^1\theta^2\frac{|x_i|}{\alpha_i}dW+\mathbf{1}_{i=k}\cdot\left(\left(1-\theta\right)+\theta^2\cdot W_{-i}\left(\tilde{r}_k\left(0\right)\right)\frac{|x_{-i}|}{\alpha_{-i}}\right)\right)\\ &= -\left(\frac{1}{\alpha_i-\gamma_i}\right)\frac{2\left|x_i\right|}{\alpha_i}\left(\mathbf{1}_{i=k}\cdot\left(1-\theta\right)+\theta^2\left(\left(1-W_i\left(\tilde{r}_k\left(0\right)\right)\right)\cdot\frac{|x_i|}{\alpha_i}+\frac{|x_{-i}|}{\alpha_{-i}}W_{-i}\left(\tilde{r}_k\left(0\right)\right)\right)\right)\end{aligned}$$

where the second equality follows from $\frac{\partial}{\partial\alpha_i}\left(\frac{W_i(r)}{\alpha_i}\right)=\frac{2|x_i|W_i'(r)}{(\alpha_i-\gamma_i)\alpha_i^2}$ and $\frac{\partial\tilde{r}_i(0)}{\partial\alpha_i}=-\frac{2|x_i|}{\alpha_i(\alpha_i-\gamma_i)}$, the third equality follows from a change of variables from t to $W_i\left(t\right)$, and the fourth equality follows from the observation that $W_{-i}\left(\tilde{r}_k\left(0\right)\right)=0$ when $i=-k$.

Now combining and factoring out $\frac{2|x_i|}{\alpha_i}$ we have that $\frac{\partial\left(\Pi_i^*(\bar{s};F^*)\right)}{\partial\alpha_i}$ shares a sign with

$$-\theta^2\frac{|x_i|}{\alpha_i}\cdot\frac{1}{2}-\int_{W_i(\tilde{r}_k(0))}^{1}\left((1-\theta)+\theta^2\left(W\frac{|x_i|}{\alpha_i}+W_{-i}\left(r_i\left(W\right)\right)\frac{|x_{-i}|}{\alpha_{-i}}\right)\right)\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}\right)dW$$
$$+\left(\mathbf{1}_{i=k}\cdot(1-\theta)+\theta^2\left(\left(1-W_i\left(\tilde{r}_k\left(0\right)\right)\right)\cdot\frac{|x_i|}{\alpha_i}+\frac{|x_{-i}|}{\alpha_{-i}}W_{-i}\left(\tilde{r}_k\left(0\right)\right)\right)\right)$$

We then rearrange terms into a more usable expression:

$$-\theta^2\frac{|x_i|}{\alpha_i}\cdot\left(\int_{W_i(\tilde{r}_k(0))}^{1}\left(\frac{\alpha_i W}{\alpha_i-\gamma_i W}-1\right)dW+\frac{1}{2}\right)$$
$$-\theta^2\frac{|x_{-i}|}{\alpha_{-i}}\cdot\left(\int_{W_i(\tilde{r}_k(0))}^{1}W_{-i}\left(r_i\left(W\right)\right)\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}\right)dW-W_{-i}\left(\tilde{r}_k\left(0\right)\right)\right)$$
$$-\left(1-\theta\right)\cdot\left(\int_{W_i(\tilde{r}_k(0))}^{1}\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}\right)dW-\mathbf{1}_{i=k}\right)$$

Now either (a) $i=-k$ and $W_i\left(\tilde{r}_k\left(0\right)\right)\geq 0$ and $W_{-i}\left(\tilde{r}_k\left(0\right)\right)=\mathbf{1}_{i=k}=0$, or (b) $i=k$ and $W_i\left(\tilde{r}_k\left(0\right)\right)=0$ and $W_{-i}\left(\tilde{r}_k\left(0\right)\right)=W_{-i}\left(\tilde{r}_i\left(0\right)\right)$. Thus, the preceding may be rewritten more simply as:

$$-\theta^2\frac{|x_i|}{\alpha_i}\cdot\left(\int_{W_i(\tilde{r}_k(0))}^{1}\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}W-\frac{1}{2}\right)dW+\frac{1}{2}W_i\left(\tilde{r}_k\left(0\right)\right)\right)$$
$$-\theta^2\frac{|x_{-i}|}{\alpha_{-i}}\cdot\int_{W_i(\tilde{r}_k(0))}^{1}\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}\cdot W_{-i}\left(r_i\left(W\right)\right)-W_{-i}\left(\tilde{r}_i\left(0\right)\right)\right)dW$$
$$-\left(1-\theta\right)\cdot\int_{W_i(\tilde{r}_k(0))}^{1}\left(\frac{\alpha_i}{\alpha_i-\gamma_i W}-\mathbf{1}_{i=k}\right)dW$$

It then suffices to show that all three terms inside the parentheses are ≥ 0. The third term is ≥ 0 since $\frac{\alpha_i}{\alpha_i-\gamma_i W}\geq 1$; the second term is ≥ 0 since $\frac{\alpha_i}{\alpha_i-\gamma_i W}\geq 1$ and $W_{-i}\left(r_i\left(W\right)\right)\geq W_{-i}\left(\tilde{r}_i\left(0\right)\right)$; the first term is ≥ 0 since since $\frac{\alpha_i}{\alpha_i-\gamma_i W}\geq 1$ and $\int_{W_i(\tilde{r}_k(0))}^{1}\left(W-\frac{1}{2}\right)dW\geq 0$. QED

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

Finally, a bidder is strictly worse off when his costs 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.2.2 Decisionmaker Payoffs

Lastly, again using Proposition 16, the DM's equilibrium utility and the bidders' average scores (which bound the DM's utility from below) may be calculated as follows.

Proposition 18. *The DM's equilibrium utility is* $U^*_{DM} = E\left[\max\left\{s_1, s_{-1}\right\}\right] = \int_0^{\bar{s}} s \frac{\partial}{\partial s}\left(\prod_j F_j\left(s\right)\right) ds =$

$$U^*_{DM} = \int_0^{\tilde{r}_k(0)} \left(1 - \prod_j W_j\left(r\right)\right) \cdot \left((1-\theta) + \theta^2 \sum_j W_j\left(r\right) \frac{|x_j|}{\alpha_j}\right) dr$$

Developer i's average score is $E\left[s_i\right] = \int_0^{\bar{s}} s \frac{\partial}{\partial s}\left(F_i\left(s\right)\right) ds =$

$$\int_0^{\tilde{r}_k(0)} \left(1 - W_{-i}\left(r\right)\right) \cdot \left((1-\theta) + \theta^2 \sum_j W_j\left(r\right) \frac{|x_j|}{\alpha_j}\right) dr$$

Proof: $F_i\left(s\right) F_{-i}\left(s\right)$ is the CDF of $\max\{s_i, s_{-i}\}$ so the DM's utility is $\int_0^{\bar{s}} s \cdot \frac{\partial}{\partial s}\left(\prod_j F_j\left(s\right)\right) ds =$ $\int_0^{\bar{s}}\left(1 - \prod_j F_j\left(s\right)\right) ds$ from integration by parts and rearranging. A change of variables from score $s \in [0, \bar{s}]$ to residual reach $r \in [0, \tilde{r}_k]$ then yields $\int_{\tilde{\mathbf{r}}(\bar{s})}^{\tilde{\mathbf{r}}(0)}\left(1 - \prod_j F_j\left(\mathbf{s}\left(r\right)\right)\right)\left(-\frac{1}{\tilde{\mathbf{r}}'(\mathbf{s}(r))}\right) dr$; substituting $\tilde{\mathbf{r}}\left(0\right) = 0$ and $\tilde{\mathbf{r}}\left(\bar{s}\right) = \tilde{r}_k\left(0\right)$ and $\prod_j F_{-j}\left(\mathbf{s}\left(r\right)\right) = W_j\left(r\right)$ and $\frac{1}{\tilde{\mathbf{r}}'(\mathbf{s}(r))} = \mathbf{s}\left(r\right)$ yields the result.

Developer i's average score is $E\left[s_i\right] = \int_0^{\bar{s}} s \frac{\partial}{\partial s}\left(F_i\left(s\right)\right) ds = \int_0^{\bar{s}}\left(1 - F_i\left(s\right)\right) ds$ from integration by parts. A change of variables yields $\int_{\tilde{\mathbf{r}}(\bar{s})}^{\tilde{\mathbf{r}}(0)}\left(1 - F_i\left(\mathbf{s}\left(r\right)\right)\right)\left(-\frac{1}{\tilde{\mathbf{r}}'(\mathbf{s}(r))}\right) dr = \int_0^{\tilde{r}_k(0)}\left(1 - W_{-j}\left(r\right)\right) \cdot \left(-\mathbf{s}'\left(r\right)\right) dr$, which is the desired expression. QED

The first term under the integral in the DM's utility captures the impact of the discouragement effect, while the second term captures the impact of the endogenous ideological stakes.

Direct comparative statics on the DM's utility U^*_{DM} are difficult because of the contrasting impact of the discouragement effect (which reduces the DM's utility when a bidder becomes more extreme or capable ceteris paribus) and the endogenous ideological stakes (which increases the DM's utility when a bidder becomes more extreme or capable ceteris paribus). However, because one developer must always be active, the impact of the discouragement effect is effectively bounded, while the impact of endogenous ideological stakes is not; this leads to simple and striking limiting properties of the DM's utility as follows.

Proposition 19. *When the generalized policy contest exhibits endogenous ideological stakes ($\theta > 0$), 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 $\lim_{\alpha_{-k} \to \gamma_{-k}} U^*_{DM} = \int_0^1 \left(\frac{2|x_k|(1-W)}{\alpha_k - \gamma_k W}\right) \cdot \left((1-\theta) + \theta^2 \left(\frac{|x_{-k}|}{\gamma_{-k}} + W \frac{|x_k|}{\alpha_k}\right)\right) dW$*, 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 $-\frac{W_i'(r) \cdot 2|x_i|}{\alpha_i - \gamma_i W_i(r)} = 1$. Substituting into $E[s_{-k}]$ and performing a change of variables from r to $W_k(r)$ yields:

$$E[s_{-k}] = \int_0^1 (1 - W_k) \cdot \left((1-\theta) + \theta^2 \left(W_k \frac{|x_k|}{\alpha_k} + W_{-k}(r_k(W_k)) \frac{|x_{-k}|}{\alpha_{-k}}\right)\right) \cdot \left(\frac{2|x_k|}{\alpha_k - \gamma_k W_k}\right) dW_k$$

Now for $|x_k|$ sufficiently high we will have $\frac{2|x_{-k}|}{\alpha_{-k} - \gamma_{-k} W} > \frac{2|x_k|}{\alpha_k - \gamma_k W}$ $\forall W \in [0,1]$, implying that $W_{-k}(r_k(W_k)) > W_k(r_k(W_k))$ $\forall W_k \in [0,1)$, so for a sufficiently extreme $-k$ his expected score is bounded below by

$$\int_0^1 (1 - W_k) \cdot \left((1-\theta) + \theta^2 \cdot W_k \cdot \left(\frac{|x_k|}{\alpha_k} + \frac{|x_{-k}|}{\alpha_{-k}}\right)\right) \cdot \left(\frac{2|x_k|}{\alpha_k - \gamma_k W_k}\right) dr$$

which in turn is bounded below by

$$\begin{aligned} & \left(\frac{2\left|x_k\right|}{\alpha_k}\right) \cdot \int_0^1 \left(1 - W_k\right) \cdot \left((1-\theta) + \theta^2 \cdot W_k \cdot \left(\frac{\left|x_k\right|}{\alpha_k} + \frac{\left|x_{-k}\right|}{\alpha_{-k}}\right)\right) \cdot dW_k \\ = \; & \frac{\left|x_k\right|}{\alpha_k}\left((1-\theta) + \frac{\theta^2}{3} \cdot \left(\frac{\left|x_k\right|}{\alpha_k} + \frac{\left|x_{-k}\right|}{\alpha_{-k}}\right)\right) \end{aligned}$$

This expression straightforwardly tends to ∞ as $\left|x_{-k}\right| \to \infty$ when $\theta > 0$.

For the fourth limiting statement, the definition in Proposition 18 and $\lim_{\alpha_{-k} \to \gamma_{-k}} \left\{W_{-k}(r)\right\} = 1$ $\forall r \in \left[0, \tilde{r}_k(0)\right]$ yields a limit of

$$\int_0^{\tilde{r}_k(0)} \left(1 - W_k(r)\right) \cdot \left((1-\theta) + \theta^2 \left(\frac{\left|x_{-k}\right|}{\gamma_{-k}} + W_k(r) \frac{\left|x_k\right|}{\alpha_k}\right)\right) dr$$

Again using that $-\frac{W_i'(r) \cdot 2\left|x_i\right|}{\alpha_i - \gamma_i W_i(r)} = 1$ and performing a change of variables from r to $W_k(r)$ yields the desired expression, which straightforwardly obeys the desired 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 Proofs for Main Text Propositions

In this Appendix we prove the main text Proposition using a mixture of general results in Appendices A-C and original proofs.

Lemma 1 This is an implication of the general necessary conditions for an SPNE of the generalized policy contest in Proposition 11, of which the policy development model is a special case. QED

Lemma 2 This is an implication of Lemma 8 and Proposition 12. QED

Lemma 3 This is a joint implication of the general equilibrium characterization in Proposition 11 and the characterization of equilibrium score CDFs in Proposition 16. QED.

Proposition 1 In the special case of the generalized policy contest that is the policy development model ($\gamma_{-1} = \gamma_1 = \theta = 1$), the residual reach function $\tilde{r}_i(W_i)$ and inverse $W_i(r_i)$ is from Lemma 8, the identity of the sometimes-active developer k is from Proposition 12, k's probability of inactivity $F_k(0)$ is from Proposition 13, and the probability distribution over ideologies $G_i(\delta)$ is from Proposition 15. To derive the quality produced when developer i crafts a policy that is distance δ from the DM, 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 $\mathbf{s}\left(\tilde{r}_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right)$ by Proposition 16. Combining the preceding, the quality associated with a policy that is distance δ from the decisionmaker is $q_i(\delta) = \frac{(i\delta)^2+\mathbf{s}\left(\tilde{r}_i\left(\frac{i\delta}{x_i/\alpha_i}\right)\right)}{\lambda}$, which simplifies to the expression in the proposition. QED

Proposition 2 To see the first bullet point, observe that Propositions 12 and 13 imply 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 15, 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 14 – 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 Propositions 12 and 13. The second bullet point under both "own strategy" and "opponent's strategy" follows from Proposition 15 on ideology. The third bullet point under "own strategy" is a joint implication of the ideology comparative statics in Proposition 15 and the comparative statics on "own score" in Proposition 16. The third bullet point under "opponent strategy" also follows from Proposition 16 and the subsequent discussion.

Proposition 4 Follows immediately from Proposition 17 characterizing the developers' payoffs.

Proposition 5 The second bullet point follows from Proposition 13 (on activity) and Proposition 19 (on the decisionmaker's welfare).

To see the first bullet point, differentiating U^*_{DM} w.r.t. $|x_{-k}|$ in the special case of the policy development model and applying symmetry yields $\frac{2}{\alpha}\int_0^{\tilde{r}(0)}\left(\left(1-3\left(W\left(r\right)\right)^2\right)\cdot x\frac{\partial W(r)}{\partial x}+\left(1-\left(W\left(r\right)\right)^2\right)W\left(r\right)\right)dr$ which is $\geq 2\frac{x}{\alpha}\int_0^{\tilde{r}(0)}\left(1-3\left(W\left(r\right)\right)^2\right)\frac{\partial W(r)}{\partial x}dr$. Now substituting $\frac{\partial W(r)}{\partial x}=-\log\left(\frac{\alpha-W(r)}{\alpha-1}\right)W'\left(r\right)$ and a change of variables yields $2\frac{x}{\alpha}\int_0^1\left(1-3W^2\right)\log\left(\frac{\alpha-W}{\alpha-1}\right)dW=2\frac{x}{\alpha}\int_0^1\left(\frac{W-W^3}{\alpha-W}\right)dW>0$. QED

Proposition 6 Follows from Propositions 13 - 15 according to a nearly identical argument as in the proof of Proposition 2.

Proposition 7 The first statement follows from Propositions 13-15 by a nearly identical argument as in the proof of Proposition 3. The second statement follows immediately from Proposition 17.

Proposition 8 To see the first bullet point, differentiating U^*_{DM} w.r.t. α_{-k} in the special case of the policy development model and applying symmetry yields $2x\int_0^{\tilde{r}(0)}\left(\left(1-(W(r))^2\right)\frac{\partial}{\partial\alpha}\left(\frac{W(r)}{\alpha}\right)-\frac{2}{\alpha}(W(r))^2\frac{\partial W(r)}{\partial\alpha}\right)dr$. Then substituting $\frac{\partial}{\partial\alpha}\left(\frac{W(r)}{\alpha}\right)=\frac{x}{(\alpha-1)\alpha^2}W'(r)$, $\frac{\partial W(r)}{\partial\alpha}=-\left(\frac{1-W(r)}{\alpha-1}\right)$, $-\frac{W'(r)x}{\alpha-W(r)}=1$, rearranging the expression, and another change of variables yields $\frac{2x^2}{(\alpha-1)\alpha^2}\int_0^1\left(2W^2\left(\frac{\alpha-\alpha W}{\alpha-W}\right)-\left(1-W^2\right)\right)dW<0$.

To see the second bullet point, applying Proposition 18 for the policy development model ($\gamma_i=\theta=1\ \forall i$) the DM's equilibrium utility under symmetry ($|x_{-k}|=|x_k|$ and $\alpha_{-k}=\alpha_k$) is:

$$U_{DM}^{sym}=\int_0^{\tilde{r}_k(0)}\left(1-[W_k(r)]^2\right)\cdot\left(2W_k(r)\frac{|x_k|}{\alpha_k}\right)dr$$

Using $\frac{1}{W_k'(r)}=r_k'(W_k(r))=-\frac{2|x_k|}{\alpha_k-W_k(r)}$, changing variables from r to $W_k(r)$, and rearranging yields:

$$U_{DM}^{sym}=2|x_k|^2\int_0^1\left(\frac{1-W^2}{\alpha_k-W}\right)\cdot\left(\frac{2W}{\alpha_k}\right)dW$$

From Proposition 19 the DM's limiting utility as $\alpha_{-k}\to\gamma_{-k}=1$ is:

$$U_{DM}^{\lim}=\int_0^1\left(\frac{2|x_k|(1-W)}{\alpha_k-W}\right)\cdot\left(W\frac{|x_k|}{\alpha_k}+|x_k|\right)dW=2|x_k|^2\int_0^1\left(\frac{1-W}{\alpha_k-W}\right)\cdot\left(\frac{W}{\alpha_k}+1\right)dW$$

Thus $U_{DM}^{\lim}>U_{DM}^{sym}\iff$

$$\int_0^1\left(\frac{1-W}{\alpha_k-W}\right)\cdot\left(\frac{W}{\alpha_k}+1\right)dW>\int_0^1\left(\frac{1-W^2}{\alpha_k-W}\right)\cdot\left(\frac{2W}{\alpha_k}\right)dW$$

which simplifies to

$$\int_0^1(1-W)\left(1-\frac{2W^2}{\alpha_k-W}\right)dW=\frac{5}{6}-\alpha_k(2\alpha_k-1)+2(\alpha_k-1)\alpha_k^2\log\left(\frac{\alpha_k}{\alpha_k-1}\right)>0$$

The final expression is clearly increasing in α_k; by Mathematica it crosses 0 at ≈ 1.0347. QED

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 policy development 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. For notational clarity denote the developer's ideals in this variant as $\hat{x}_i$.

We first show this variant is isomorphic to the generalized policy contest with $\theta = 0$, $\gamma_i = 1\ \forall i$, and $x_i = 2\hat{x}_i y$. In this variant of the generalized policy contest, a non-zero choice of idelogy $y_i \neq 0$ carries a strict up-front cost and no actual or potential benefit, so the optimal choice of ideology is $y_i^*(s) = 0\ \forall i, s$. The problem then effectively reduces to the bidder's optimal choice of score, where $c_i(s_i, 0) = \alpha_i s_i$ and $V_i(\rho_i, s_w, 0) = s_w + 2x_i \cdot i\rho_i$.

Alternatively, in the fixed policy variant the ideology of developer i's policy is exogenously equal to iy. Redefining the score as $s(y_i, q_i) = \lambda q$ (that is, equal to the portion of the DM's utility derived from quality), a policy with score s_i has quality $q_i = \frac{s_i}{\lambda}$, and the score still fully determines the identity of the winner. Thus, player i's up-front cost from choosing score s_i is $c_i(s_i) = a \cdot \frac{s_i}{\lambda} = \alpha_i s_i$, and her utility $V_i(\mathbf{1}_{i=w}, s_w, 0)$ from player w winning is:

$$\begin{aligned}
V_i(\mathbf{1}_{i=w}, s_w, 0) &= \lambda \cdot \frac{s_i}{\lambda} - (\hat{x}_i - wy)^2 = s_i - (\hat{x}_i - wy)^2 + (\hat{x}_i + iy)^2 - (\hat{x}_i + iy)^2 \\
&= s_i + 4y\hat{x}_i \cdot i\mathbf{1}_{i=w} - 2y^2 - (x_i + iy)^2
\end{aligned}$$

But this is just the outcome utility from the generalized policy contest with $2x_i = 4y\hat{x}_i$ less a constant; hence the models are isomorphic and equilibrium and payoffs are as characterized in Appendix C with $\theta = 0$, $\gamma_i = 1$, and $2x_i = 4y\hat{x}_i$.

Proof of Proposition 9 Although this variant does not fit into the general all-pay contest framework of Siegel (2009) because the developers 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). Indeed, any variant of the generalized policy contest with $\theta = 0$ (the absence of an ideological component to the prize) does so.

First, by Proposition 16 we have $\mathbf{s}(r) = \tilde{r}_k(0) - r$ (identical to the standard all pay contest), so $\bar{s} = \tilde{r}_k(0)$ and $\tilde{\mathbf{r}}(s) = \tilde{r}_k(0) - s$. By Proposition 17 the weaker player k's expected utility is then:

$$4|\hat{x}_k| y - (\alpha_k - 1) \cdot \tilde{r}_k(0)$$

which is invariant to the characteristics of the stronger player.

Further, by Proposition 16 the equilibrium score CDF $F_{-k}(s)$ of the stronger player $-k$ is

$$F_{-k}(s) = W_k(\tilde{\mathbf{r}}(s)) = W_k(\tilde{r}_k(0) - s),$$

which is invariant to her own characteristics (because it is exclusively determined by the need to keep the weaker player k indifferent over the common mixing interval $s \in [0, \tilde{r}_k(0)]$). The score CDF $F_k(s)$ of the weaker player k is:

$$F_k(s) = W_{-k}(\tilde{r}_k(0) - s)$$

which is straightforwardly increasing in the stronger player's extremism $|x_{-k}|$ and decreasing in his costs α_k via the effect of these parameters on $W_{-k}(\cdot)$. Thus the discouragement effect is present in full and asymmetries harm the DM's utility. More specifically the DM's equilibrium utility is

$$U^*_{DM} = \int_0^{\tilde{r}_k(0)} (1 - W_k(r) W_{-k}(r)) \cdot dr,$$

which is straightforwardly decreasing in $|\hat{x}_{-k}|$ and increasing in α_{-k} via $W_{-k}(r)$.

Finally, when the developers are symmetrically extreme and capable ($|\hat{x}_{-1}| = |\hat{x}_1| = \hat{x}$ and $\alpha_{-1} = \alpha_1 = \alpha$) they have a common CDF $F(s) =$

$$\begin{aligned} W(\tilde{r}(0) - s) &= \alpha - (\alpha - 1) e^{\frac{4\hat{x}y \log\left(\frac{\alpha}{\alpha-1}\right) - s}{4\hat{x}y}} \\ &= \alpha\left(1 - e^{-\frac{s}{4\hat{x}y}}\right) \end{aligned}$$

which is straightforwardly increasing in α and decreasing in x. Thus, the common score CDF is first order stochastically increasing in extremism and ability, so the DM benefits from symmetric increases in either extremism or ability, as in Hirsch and Shotts (2015) and the standard all pay contest. 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:

$$-\left(\alpha_i - F_{-i}\left(s\right)\right) s + F_{-i}\left(s\right) \cdot 2x_i y - \alpha_i y^2 - x_i^2 \tag{E.1}$$

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

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

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

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

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

$$\frac{x_i^2}{\alpha_i\left(\alpha_i - 1\right)} \cdot \left(1 - \left[F_{-i}\left(s\right)\right]^2\right) = \bar{s} - \left(\frac{\alpha_i - F_{-i}\left(s\right)}{\alpha_i - 1}\right) \cdot s \tag{E.3}$$

It is easily verified that for a given value of $\bar{s}$ this expression uniquely defines a continuously increasing $F_{-i}\left(s\right) \forall s \in [0, \bar{s}]$; 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.3 and simplifying yields a characterization of the unique score CDFs for the always active developer $-k$:

$$x_k^2 \cdot [F_{-k}(s)]^2 = \alpha_k (\alpha_k - F_{-k}(s)) \cdot s \tag{E.4}$$

and the sometimes inactive developer k:

$$x_{-k}^2 \cdot \left(1 - [F_k(s)]^2\right) = \alpha_{-k} \left((\alpha_{-k} - 1) \bar{s}_k - (\alpha_{-k} - F_k(s)) s \right) \tag{E.5}$$

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.2 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.4 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.5 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:

$$-\frac{\alpha_i}{2}(s+|y|)^2 + F_{-i}(s)(s+|y|+x_i y) + \int_s^{\bar{s}} (s+|y_{-i}(s)| + x_i y_{-i}(s)) f_{-i}(s_{-i} ds_{-i}), \qquad \text{(E.6)}$$

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

Necessary and Sufficient Conditions for Equilibrium Using a similar series of steps as in Appendix B 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.6 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:

$$\alpha (s + y(F_{-i}(s); s)) - F_{-i}(s) = f_{-i}(s) \cdot [(x+1) \cdot y(F_{-i}(s); s) + (x-1) \cdot y(F_i(s); s)] \tag{E.7}$$

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

$$\alpha (s + y(F(s), s)) - F(s) = f(s) \cdot 2x \cdot y(F(s); s) \tag{E.8}$$

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.