RISK-SEEKING STRATEGIC ROBUSTNESS IN GAMES: A SUBSTITUTES-COMPLEMENTS RULE, AND WHERE IT BREAKS

Anonymous authors
Paper under double-blind review

ABSTRACT

Strategic risk aversion, formalized through a fictitious-adversary construction, has recently been shown to induce a "free lunch" in a shared-effort game: more robustness can strictly improve equilibrium welfare, in contrast to classical robust optimization. We study this phenomenon in a tractable class of two-player quadratic games, and introduce the natural risk-seeking counterpart of the construction, in which the fictitious adversary is replaced by a fictitious ally. Within this game class, we show that whether risk aversion, risk seeking, or neither attitude helps is governed by a single, classical structural property: whether the game exhibits strategic substitutes or complements, and that the losing attitude always strictly hurts welfare instead. We illustrate the theory on three examples: trading, a shared-effort aggregative game, and Bertrand price competition.

1 INTRODUCTION

A recent line of work formalizes strategic risk aversion for multi-agent learning: a player evaluates a fictitious adversary's ability to deviate from the true behavior of their opponents, within a bounded budget, and best-responds to the resulting worst case (Mazumdar et al., 2025; Lanzetti et al., 2025). This construction was introduced at ICLR (Mazumdar et al., 2025) as a way to make equilibrium computation in Markov games more tractable, and it has since been extended along two directions: to training generalizable collaborative agents, with empirical validation that includes a large-language-model collaboration task (Qu et al., 2026), and to dynamic, multi-stage linear quadratic games (Velasevic et al., 2026). Applied to a canonical shared-effort game, in which players jointly benefit from a common reward while each pays individually for their own effort, the construction yields a striking result: risk-averse players achieve strictly better equilibrium welfare than risk-neutral (Nash) players, a genuine free lunch, since ordinary robust optimization typically trades off performance for robustness (Qu et al., 2026; Ben-Tal et al., 2009).

Two questions follow. First, is this free-lunch result general, or specific to the one example it was proven for? Second, the existing literature considers risk aversion alone: is there a meaningful risk-seeking counterpart, built from a fictitious ally, and does it play a useful role in some other class of games? These questions matter beyond their immediate setting. Multi-agent learning systems (Zhang et al., 2021) are increasingly deployed in settings with strategic interaction, including systems of interacting learned or language-model agents on which the risk-averse construction has already been tested directly (Qu et al., 2026). Such systems are also increasingly observed to behave strategically and game-theoretically in their own right (Akata et al., 2025; Lorè & Heydari, 2024). A designer's ability to predict and steer which equilibrium such a system converges to, and to do so with a single scalar robustness parameter tunable at training or deployment time, is directly relevant to the safety and reliability of that deployment. Risk aversion supplies half of that design space; without a risk-seeking counterpart, and without knowing which of the two attitudes to reach for in a given strategic setting, that design space is incomplete.

This paper answers both questions. We define the risk-seeking construction, and show that whether risk aversion, risk seeking, or neither helps is governed by a single, classical structural property of the game: whether it exhibits strategic substitutes or complements, meaning, informally, that one player's best response falls as the other's action rises (substitutes) or rises with it (comple-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

ments). Concretely, we make three contributions. First, we define risk-seeking strategic robustness as the natural optimistic mirror of the existing risk-averse construction (Section 4). Second, we prove general results characterizing both attitudes: an existence theorem valid for arbitrary quadratic games, a general result on the qualitative shape of best-response dynamics, and a general directional theorem, holding for arbitrary game parameters subject only to two structural hypotheses on the base game, identifying which attitude improves equilibrium outcomes as a function of the substitutes/complements structure alone, and showing the losing attitude strictly hurts welfare wherever it is well posed (Section 5). Third, we illustrate the theory on three concrete games, grouped by which attitude benefits them, working out each in full; one of the three shows concretely where a specific tractable subclass's geometric pairing mechanism stops applying, while the underlying welfare-improving phenomenon the general theorem predicts persists (Section 6).

Summary of findings. At a high level, the mechanism we identify is Pigouvian (Pigou, 1920) in spirit: much as a Pigouvian tax or subsidy corrects a market outcome by directly offsetting an externality, the fictitious-adversary or fictitious-ally construction corrects a Nash inefficiency in the same direction the game's own strategic externality creates it, so the matching attitude, risk aversion under substitutes, risk seeking under complements, fixes exactly the distortion the other attitude would instead worsen. We also find that risk seeking is a structurally more fragile construction than risk aversion: it fails to be well posed, and in some regimes fails to admit an equilibrium at all, over a strictly wider range of parameters.

2 Related Work

Linear-quadratic and bilinear-quadratic games. The class of games we study sits in a long tradition of tractable models built on quadratic payoffs and linear best responses. Linear-quadratic games are the workhorse model of dynamic and differential game theory (Başar & Olsder, 1999), valued precisely because best responses and equilibria admit closed form even when the underlying strategic interaction is rich, and this tractability has made them a standard testbed for multi-agent reinforcement learning and control more broadly (Zhang et al., 2021; Velasevic et al., 2026). The static, symmetric, two-player payoff we formalize in Section 3 is the single-shot special case of this broader family, and it already nests several canonical strategic-interaction models as special cases of a single structural sign: Cournot quantity competition (Cournot, 1838), Bertrand price competition (Bertrand, 1883), and, once generalized beyond two players, congestion and routing games (Wardrop, 1952; Roughgarden & Tardos, 2002). Much of this game class's equilibrium and stability behavior is already well understood in the classical game-theory literature; we use it here to study the effect of risk-averse and risk-seeking robustification on equilibrium welfare, a question that, to our knowledge, has not previously been studied within it.

Robustness and risk-sensitivity in (multi-agent) reinforcement learning. The fictitious-adversary and fictitious-ally constructions we build on have roots in risk-sensitive control: Jacobson (1973) first showed that an exponential-cost linear-quadratic control problem is mathematically equivalent to a deterministic differential game against a worst-case disturbance, the same duality between a risk attitude and a fictitious opponent that our own constructions instantiate in a strategic, multi-player setting. This duality connects closely to robust optimization (Ben-Tal et al., 2009; Wiesemann et al., 2014) and to robust dynamic programming and robust MDPs (Iyengar, 2005; Nilim & El Ghaoui, 2005), both of which target robustness to environmental uncertainty in the transition dynamics, with finite-sample guarantees developed for the single-agent distributionally robust setting (Xu et al., 2023). In the multi-agent setting, this line of thinking has been extended to robust multi-agent reinforcement learning under model or state uncertainty (Zhang et al., 2020; 2021), and, closest to the present paper, to strategic robustness against an opponent's deviations (Mazumdar et al., 2025; Lanzetti et al., 2025), with follow-on work extending the construction to training collaborative agents (Qu et al., 2026) and to dynamic, multi-stage linear quadratic games (Velasevic et al., 2026). Computing a Nash equilibrium is intractable in the worst case: even in two-player general-sum matrix games, finding one is PPAD-hard (Daskalakis et al., 2009), placing it alongside a handful of fundamental problems with no known polynomial-time algorithm despite decades of effort. Motivated by this, Mazumdar et al. (2025) show that equipping agents with risk aversion, alongside bounded rationality, restores tractability of computation across *all* matrix and finite-horizon Markov games (Shapley, 1953), independent of the underlying game structure. Within the bilinear-quadratic

class we study, tractability of computation is not itself in question (it is inherited from the game class, above); our contribution is instead to determine which of the two attitudes, risk-averse or risk-seeking, yields a *better* equilibrium, a question their tractability result does not address, and we show the answer is governed entirely by the game's classical substitutes/complements structure. All of this prior work considers risk-averse robustness alone; we are not aware of a prior risk-seeking counterpart formalized and analyzed at this level of generality.

Strategic substitutes and complements. The classification of games by the sign of players' cross-partial derivatives, into strategic substitutes and complements, is a classical tool in economic theory (Fudenberg & Tirole, 1991) for characterizing the comparative statics of equilibria, including whether equilibria are efficient relative to the joint-welfare optimum (Bulow et al., 1985), and connects to the lattice-theoretic study of games with monotone best-response structure and their equilibrium and learning dynamics (Topkis, 1979; Milgrom & Roberts, 1990; Milgrom & Shannon, 1994). This classification is typically used to compare games with different payoff parameters or different numbers of players, holding the solution concept fixed at Nash equilibrium. Our directional result repurposes this same classification for a different question: holding the game fixed, it determines which of two robustness-adjusted solution concepts, the risk-averse or risk-seeking equilibrium, moves the outcome toward or away from efficiency.

Strategic behavior in learned and language-model multi-agent systems. A separate and more recent line of empirical work studies whether large language model agents, placed in classical game-theoretic settings, behave strategically and rationally in ways that resemble or depart from equilibrium play, and thus from the quantal-response and risk-averse behavior documented in the earlier behavioral-economics literature on human play (McKelvey & Palfrey, 1995; Goeree et al., 2003): Akata et al. (2025) let language models play finitely repeated 2×2 games against each other and against fixed strategies, finding persistent behavioral signatures that differ systematically across game types, Lorè & Heydari (2024) show that strategic behavior in this setting depends on both the underlying game structure and the contextual framing in which it is presented, and related work extends this to multi-party negotiation and simulated social settings (Park et al., 2023; Hua et al., 2024). This literature motivates our question from the opposite direction: it documents empirically that such systems already behave strategically and inconsistently across game types, while our contribution is a theoretical account of which robustness attitude a designer of such a system should target, and when, as a function of the same classical game structure these empirical studies vary over. Qu et al. (2026) already show that the risk-averse construction improves collaborative outcomes on a large-language-model collaboration task; our directional result indicates when the opposite, risk-seeking attitude would serve such a system better, though our formal results themselves stay proven strictly for the bilinear-quadratic game class of Definition 1.

3 PRELIMINARIES

We work throughout with a simple, tractable class of two-player games, general enough to host Cournot competition, Bertrand competition, congestion, and shared-effort games as special cases, and simple enough to admit closed-form equilibria under both risk attitudes.

Definition 1 (symmetric bilinear-quadratic game)**.** *Two players choose actions $a_1, a_2 \in \mathbb{R}$. Player i's payoff is*

$$u_i(a_i, a_j) = \alpha a_i + \beta a_i^2 + \gamma a_i a_j + \delta a_j^2 + \eta a_j, \qquad \beta < 0, \tag{1}$$

with the same $(\alpha, \beta, \gamma, \delta, \eta)$ for both players. The condition $\beta < 0$ makes player i's payoff strictly concave in a_i for any fixed opponent action, so each player's own maximization problem is well posed before any robustification. When $2\beta + \gamma \neq 0$, the game has a unique symmetric Nash equilibrium $a^{\mathrm{NE}} = -\alpha/(2\beta + \gamma)$.

Definition 2 (substitutes and complements)**.** *The classical strategic type of the game is the sign of the cross-partial $\gamma = \partial^2 u_i / \partial a_i \partial a_j$. The ordinary (non-robustified) best-response curve, from the first-order condition $\alpha + 2\beta a_i + \gamma a_j = 0$, is $a_i = -(\alpha + \gamma a_j)/(2\beta)$, with slope $\gamma/(2|\beta|)$. The game exhibits strategic substitutes if $\gamma < 0$ (downward-sloping best responses) and strategic complements if $\gamma > 0$ (upward-sloping).*

Having fixed the baseline game and its substitutes/complements classification, we turn to this paper's central construction: robustifying each player's belief about their opponent. The risk-averse version

replaces the opponent's action with a fictitious adversary's report, chosen adversarially within a bounded budget.

Definition 3 (risk-averse robust utility)**.** *For a robustness weight $\lambda > 0$, player i's risk-averse robustified utility replaces the opponent's action with a fictitious adversary's report d, chosen to minimize player i's payoff subject to a quadratic penalty, weighted by $1/\lambda$, for deviating from the truth:*

$$u_i^{\mathrm{ave}}(a_i; a_j) = \min_d \left[\alpha a_i + \beta a_i^2 + \gamma a_i d + \delta d^2 + \eta d + \frac{1}{\lambda}(d - a_j)^2 \right]. \tag{2}$$

A pair $(a_1^\star, a_2^\star)$ is a risk-averse equilibrium if each $a_i^\star$ maximizes $u_i^{\mathrm{ave}}(\cdot\,; a_j^\star)$.

This construction, and its interpretation as a fictitious adversary constrained by a quadratic penalty, mirrors the KL-constrained construction of Mazumdar et al. (2025) and the optimal-transport-based construction of Lanzetti et al. (2025). λ is the robustness weight: a larger λ permits larger deviations and more robustification, while $\lambda \to 0$ forces the fictitious report to coincide with the truth and recovers Nash.

4 Risk-Seeking Strategic Robustness

The risk-averse construction of Definition 3 assumes the worst case: a fictitious adversary who wants to hurt player i, constrained not to deviate too far from the truth. The natural mirror assumes the best case instead: a fictitious ally, equally constrained, who reports whichever nearby action helps player i the most.

Definition 4 (risk-seeking robust utility)**.** *For a robustness weight $\lambda > 0$, player i's risk-seeking robustified utility replaces the opponent's action with a fictitious ally's report d, chosen to maximize player i's payoff subject to the same quadratic deviation penalty, weighted by $1/\lambda$:*

$$u_i^{\mathrm{seek}}(a_i; a_j) = \max_d \left[\alpha a_i + \beta a_i^2 + \gamma a_i d + \delta d^2 + \eta d - \frac{1}{\lambda}(d - a_j)^2 \right]. \tag{3}$$

A pair $(a_1^\star, a_2^\star)$ is a risk-seeking equilibrium if each $a_i^\star$ maximizes $u_i^{\mathrm{seek}}(\cdot\,; a_j^\star)$.

The sign of the penalty term is forced by the requirement that deviation remain costly regardless of which direction the fictitious player optimizes. Under a minimization, a deviation cost must be added; under a maximization, the identical cost must be subtracted. Both constructions are the same idea, a bounded fictitious perturbation of the truth, applied through the arithmetic appropriate to whichever direction is being optimized.

Call the optimization over d inside each definition the *inner problem*, and call each player's own choice of a_i (the maximization in the equilibrium clause of Definitions 3 and 4) the *outer problem*.

Remark 1 (risk aversion is minimax, risk seeking is a joint maximization)**.** *For risk aversion, the inner and outer problems compose into $\max_{a_i} \min_d[\cdot]$, the same minimax structure underlying the robust-control constructions of Mazumdar et al. (2025) and Lanzetti et al. (2025). For risk seeking, they compose into $\max_{a_i} \max_d[\cdot]$ instead, simply a joint maximization over (a_i, d), carrying no adversarial tension between the two problems. This is exactly why Theorem 1's risk-averse half needs only one well-posedness condition while the risk-seeking half needs two.*

5 Theoretical Results

Following the standard theory of linear-quadratic games (Başar & Olsder, 1999), we first check well-posedness of each player's own optimization problem (Section 5.1) and a shared fact about best-response dynamics (Section 5.2). We then treat the two attitudes in turn, existence of a symmetric equilibrium through convergence of best-response dynamics, for risk-aversion (Section 5.3) and then risk-seeking (Section 5.4). Section 5.5 then brings the two attitudes together and determines which one raises welfare, as a function of the base game's own substitutes/complements structure alone. Full proofs of the two theorems below, and of every proposition, are given in Appendix A.

5.1 WELL-POSEDNESS

Call the coefficient of a_i^2 in a robustified utility, once the inner problem's optimal d has been eliminated, its *effective curvature*: the direct analogue of β in the unrobustified game of Definition 1, and negative exactly when player i's own problem is concave in a_i, the same role $\beta < 0$ plays there.

Theorem 1 (existence, both attitudes)**.** *The risk-averse construction is well posed if and only if* $1 + \delta\lambda > 0$*, in which case the effective curvature of the robustified utility in* a_i *is* $\beta^{\text{ave}}_{\text{eff}}(\lambda) = \beta - \dfrac{\gamma^2\lambda}{4(1+\delta\lambda)}$*. The risk-seeking construction is well posed if and only if* $1-\delta\lambda > 0$ *and* $\beta^{\text{seek}}_{\text{eff}}(\lambda) = \beta + \dfrac{\gamma^2\lambda}{4(1-\delta\lambda)} < 0$*; given the first condition, the second is equivalent to the single closed-form threshold* $\lambda\big(\delta+\gamma^2/(4|\beta|)\big) < 1$*. When well posed, the reaction curve is linear,* $a_i = K(\lambda)+L(\lambda)a_j$*, with* δ *and the sign of the* γ^2 *term flipping between the risk-averse and risk-seeking cases (full expressions equation 6, equation 7 in Appendix A.1).*

Proof. See Appendix A.1. □

The risk-averse half of Theorem 1 sharpens Lanzetti et al. (2025)'s existence guarantee for their optimal-transport construction on general concave games, from conditions mirroring Rosen's for Nash equilibrium to the exact closed-form threshold $1+\delta\lambda > 0$, together with the effective curvature on either side. The risk-seeking half has no prior analogue to sharpen, since the robustness literature (Section 2) is risk-averse only.

5.2 BEST-RESPONSE DYNAMICS

We track best-response dynamics through the scalar recursion $a_{t+1} = K(\lambda)+L(\lambda)a_t$, the trajectory of a single player's action under sequential (alternating) best responses. When $L(\lambda) \neq 1$, this recursion has a unique fixed point, $a = K(\lambda)/(1 - L(\lambda))$, a symmetric equilibrium of the game; when $L(\lambda) = 1$ instead, it has none, and no symmetric equilibrium exists. The recursion converges to its fixed point exactly when $|L(\lambda)| < 1$, and diverges when $|L(\lambda)| > 1$.

Proposition 1 (sign of the reaction-curve slope)**.** *Whenever the risk-averse or risk-seeking construction is well posed,* $\operatorname{sign}(L(\lambda)) = \operatorname{sign}(\gamma)$*, independent of* $\alpha, \beta, \delta, \eta$*, and* λ*. Consequently, the scalar best-response recursion above is oscillatory under substitutes (*$\gamma < 0$*, alternating on either side of the equilibrium) and monotonic under complements (*$\gamma > 0$*), for any starting point of that recursion.*

Proof. See Appendix A.2. □

Since $\operatorname{sign}(L(\lambda))$ never changes, only one of the two boundary values $L = \pm 1$ is ever reachable, -1 under substitutes and 1 under complements. Next, we analyze convergence of best-response dynamics, splitting by attitude, risk aversion first (Propositions 2 and 3) and then risk seeking (Section 5.4, Propositions 4 and 5), and within each, by strategic type.

5.3 RISK-AVERSION

Both propositions turn on how $L^{\text{ave}}(\lambda)$ moves across the well-posed domain; its closed form is equation 6 (Appendix A.1). Write $\bar{\lambda}$ for the domain's right endpoint ($\bar{\lambda} = \infty$ if $\delta \geq 0$, and $\bar{\lambda} = -1/\delta$ if $\delta < 0$), $L^{\text{ave}}_0 := L^{\text{ave}}(0) = \gamma/(2|\beta|)$ for $L^{\text{ave}}(\lambda)$'s value as $\lambda \to 0$ (the Nash limit), and $L^{\text{ave}}_{\bar{\lambda}} := \lim_{\lambda\to\bar{\lambda}} L^{\text{ave}}(\lambda)$ for its value at the other end of the domain ($L^{\text{ave}}_{\bar{\lambda}} = 0$ if $\delta \geq 0$, $L^{\text{ave}}_{\bar{\lambda}} = -2\delta/\gamma$ if $\delta < 0$). $L^{\text{ave}}(\lambda)$ is strictly monotone on $(0, \bar{\lambda})$ (Lemma 1, Appendix A.3), so it sweeps out exactly the open interval between L^{ave}_0 and $L^{\text{ave}}_{\bar{\lambda}}$, reaching any value strictly between them once and no value outside that interval. Both propositions below ask whether that swept interval reaches $\operatorname{sign}(\gamma)$: -1 under substitutes, 1 under complements, the sign $L^{\text{ave}}(\lambda)$ itself always carries (Proposition 1). This is a condition on β, γ, δ alone:

$$(|\gamma| - 2|\beta|)\left(\delta + \frac{|\gamma|}{2}\right) > 0 \tag{4}$$

holds exactly when $\mathrm{sign}(\gamma)$ lies strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\mathrm{ave}}$ (Appendix A.5 derives this equivalence).

Proposition 2 (substitutes: $\gamma < 0$)**.** ***Existence.*** *$L_0^{\mathrm{ave}}, L_{\bar{\lambda}}^{\mathrm{ave}} \leq 0$ always, so 1 is never reached: a symmetric risk-averse equilibrium exists unconditionally over the whole well-posed domain.* ***Convergence.*** *Dynamics converge on $(0, \bar{\lambda})$ if equation 4 fails and $|\gamma| < 2|\beta|$. If equation 4 holds, $|L^{\mathrm{ave}}(\lambda)|$ crosses 1 exactly once, at $\lambda_{-1} := \dfrac{-2(\gamma + 2|\beta|)}{\gamma^2 + 4|\beta|\delta}$, and dynamics converge on $(0, \lambda_{-1})$ if $|\gamma| < 2|\beta|$, or on $(\lambda_{-1}, \bar{\lambda})$ if $|\gamma| \geq 2|\beta|$ instead. Dynamics diverge at every other λ in the well-posed domain (Appendix A.5 gives the full case analysis).*

Proof. See Appendix A.5. □

Proposition 3 (complements: $\gamma > 0$)**.** ***Existence.*** *If equation 4 fails, existence holds throughout, since the swept interval never reaches 1. If equation 4 holds instead, the symmetric fixed-point equation degenerates exactly once, at $\lambda_1 := \dfrac{2(\gamma - 2|\beta|)}{\gamma^2 + 4|\beta|\delta}$: a symmetric risk-averse equilibrium exists at every other λ in the well-posed domain.* ***Convergence.*** *Dynamics converge on $(0, \bar{\lambda})$ if equation 4 fails and $|\gamma| < 2|\beta|$, or on $(0, \lambda_1)$ if equation 4 holds and $|\gamma| < 2|\beta|$, or on $(\lambda_1, \bar{\lambda})$ if equation 4 holds and $|\gamma| \geq 2|\beta|$ instead. Dynamics diverge at every other λ in the well-posed domain (Appendix A.6 gives the full case analysis).*

Proof. See Appendix A.6. □

5.4 RISK-SEEKING

Section 5.3's two propositions both located their convergent side by the same argument (Appendix A.4): given a target crossing value, exactly one of the domain's two endpoint limits has magnitude below 1, and monotonicity carries that safety to the adjacent piece of the domain. The same argument applies to risk-seeking, with one difference: when $\bar{\lambda}^{\mathrm{seek}}$ (Theorem 1), the risk-seeking analogue of $\bar{\lambda}$, is finite, $K^{\mathrm{seek}}(\lambda), L^{\mathrm{seek}}(\lambda)$ blow up there instead of approaching a finite limit, so the far endpoint is $\pm\infty$ rather than a finite value, unlike the risk-averse case. An infinite endpoint is never the safe one, which is the source of the extra case below.

Proposition 4 (substitutes: $\gamma < 0$)**.** *Write $\bar{\lambda}^{\mathrm{seek}} := 1/\big(\delta + \gamma^2/(4|\beta|)\big)$ if $\delta > -\gamma^2/(4|\beta|)$, and $\bar{\lambda}^{\mathrm{seek}} := \infty$ otherwise (Theorem 1), mirroring $\bar{\lambda}$'s definition for risk aversion. $L_0^{\mathrm{seek}} = \gamma/(2|\beta|)$, the far limit is $-\infty$ if $\bar{\lambda}^{\mathrm{seek}}$ is finite and 0 otherwise, and set $\lambda_{-1}^{\mathrm{seek}} := \dfrac{2(2|\beta| + \gamma)}{\gamma^2 + 4|\beta|\delta}$.* ***Existence.*** *Existence holds throughout the well-posed domain except exactly at $\lambda_{-1}^{\mathrm{seek}}$, whenever that threshold falls inside it.* ***Convergence.*** *If $|\gamma| < 2|\beta|$, dynamics converge on $(0, \lambda_{-1}^{\mathrm{seek}})$ when $\bar{\lambda}^{\mathrm{seek}}$ is finite, or on all of $(0, \bar{\lambda}^{\mathrm{seek}})$ when it is infinite. If instead $|\gamma| \geq 2|\beta|$, dynamics converge on $(\lambda_{-1}^{\mathrm{seek}}, \bar{\lambda}^{\mathrm{seek}})$ when $\bar{\lambda}^{\mathrm{seek}}$ is infinite, or nowhere when it is finite. Dynamics diverge at every other λ in the well-posed domain.*

Proof. See Appendix A.7. □

Proposition 5 (complements: $\gamma > 0$)**.** *With $\bar{\lambda}^{\mathrm{seek}}$ as above: $L_0^{\mathrm{seek}} = \gamma/(2|\beta|)$, the far limit is $+\infty$ if $\bar{\lambda}^{\mathrm{seek}}$ is finite and 0 otherwise, and set $\lambda_1^{\mathrm{seek}} := \dfrac{2(2|\beta| - \gamma)}{\gamma^2 + 4|\beta|\delta}$.* ***Existence.*** *Existence holds throughout the well-posed domain except exactly at $\lambda_1^{\mathrm{seek}}$, whenever that threshold falls inside it. The curve $\delta/|\beta| = -(\gamma/|\beta|)^2/4$ marks exactly where $\bar{\lambda}^{\mathrm{seek}}$ transitions from finite to unbounded (Appendix A.9).* ***Convergence.*** *If $|\gamma| < 2|\beta|$, dynamics converge on $(0, \lambda_1^{\mathrm{seek}})$ when $\bar{\lambda}^{\mathrm{seek}}$ is finite, or on all of $(0, \bar{\lambda}^{\mathrm{seek}})$ when it is infinite. If instead $|\gamma| \geq 2|\beta|$, dynamics converge on $(\lambda_1^{\mathrm{seek}}, \bar{\lambda}^{\mathrm{seek}})$ when $\bar{\lambda}^{\mathrm{seek}}$ is infinite, or nowhere when it is finite. Dynamics diverge at every other λ in the well-posed domain.*

Proof. See Appendix A.8. □

5.5 DIRECTION: WHICH ATTITUDE HELPS

Sections 5.3 and 5.4 settle, at each λ, whether a symmetric equilibrium exists and whether best-response dynamics converge to it. Theorem 2 below asks a third, independent question: assuming the equilibrium exists, is it actually good for the players? This does not require convergence: the theorem compares the equilibrium's welfare against Nash's wherever the construction is well posed (Theorem 1) and the equilibrium exists, whether or not best-response play from an arbitrary starting point would actually reach it. That is the free-lunch question from Section 1 (Qu et al., 2026): whether robustifying an already well-behaved game can also raise the welfare it delivers, and if so, for which attitude. We measure welfare by $W(a) := u_1(a, a) + u_2(a, a) = 2(\alpha + \eta)a + 2(\beta + \gamma + \delta)a^2$, the sum of both players' payoffs under symmetric play.

Theorem 2 (direction, general)**.** *Suppose $2\beta+\gamma < 0$ and $\beta+\gamma+\delta < 0$. Then, for arbitrary α, δ, η: risk-averse play strictly improves equilibrium welfare over Nash if and only if $\gamma < 0$ (substitutes); risk-seeking play strictly improves it if and only if $\gamma > 0$ (complements). The losing attitude actively hurts: wherever it is well posed, it strictly worsens welfare relative to Nash, monotonically, with no improving phase. Moreover, the winning attitude's equilibrium action $a^\star(\lambda)$ is itself strictly monotonic in λ throughout the whole well-posed domain, with no interior singularity to interrupt it.*

Proof. See Appendix A.9. □

Note that $\beta + \gamma + \delta < 0$ makes W strictly concave, giving it the single best point $a^\dagger = -(\alpha + \eta)/(2(\beta+\gamma+\delta))$ the proof compares Nash against; without it, W is convex or affine, and its critical point is a minimum, or does not exist at all, leaving nothing for Nash to be compared to.

Remark 2 (when the equilibrium reaches the optimum)**.** *Since the winning attitude's $a^\star(\lambda)$ is monotonic and W is strictly concave with unique maximizer $a^\dagger$, welfare $W(a^\star(\lambda))$ is unimodal in λ. If $a^\star(\lambda)$'s monotonic path reaches $a^\dagger$ at some $\lambda^\star$ within the well-posed domain, welfare strictly increases on $(0, \lambda^\star)$ and strictly decreases after it, and the robustified equilibrium equals the joint-welfare optimum exactly at $\lambda = \lambda^\star$; the three worked examples of Section 6 are all of this kind, with $\lambda^\star$ marked on each welfare panel. If $a^\star(\lambda)$ never reaches $a^\dagger$ within the domain instead, welfare increases throughout it with no interior maximum, and the equilibrium approaches the joint optimum without ever reaching it.*

Corollary 1 ($\delta = \eta = 0$: Nash over- and under-provision)**.** *Suppose in addition $\delta = \eta = 0$ and $\alpha > 0$. Under Theorem 2's hypotheses, Nash strictly over-provides relative to the joint optimum ($a^{\mathrm{NE}} > a^\dagger$) under substitutes ($\gamma < 0$) and strictly under-provides ($a^{\mathrm{NE}} < a^\dagger$) under complements ($\gamma > 0$), recovering Bulow et al. (1985)'s classical picture; in each case, the welfare-improving attitude Theorem 2 identifies is exactly the one that pulls the equilibrium toward $a^\dagger$.*

Proof. See Appendix A.10. □

6 EXAMPLES

Each example gives the payoff, its classification under Definition 2, the attitude identified by Theorem 2 where applicable, and a figure showing the equilibrium action and net welfare against λ; all three sit exactly where Figure 1 places them.

6.1 TRADING AND CONGESTION

Two funds trade a common asset with linear price impact $p(Q) = A - bQ$, a linear-quadratic reduction of Cournot quantity competition (Cournot, 1838), or equivalently, two agents route flow across a shared resource with affine latency $c(F) = c_0 + c_1 F$, in the tradition of Wardrop-style congestion models and their price of anarchy (Wardrop, 1952; Roughgarden & Tardos, 2002); the two reduce to the identical payoff up to relabeling, the instance of equation 1 with $u_i(a_i, a_j) = Aa_i - ba_i^2 - ba_i a_j$, i.e. $(\alpha, \beta, \gamma, \delta, \eta) = (A, -b, -b, 0, 0)$. This is a substitutes game ($\gamma = -b < 0$) with $\delta = \eta = 0$, so Theorem 2 applies directly: risk aversion helps, and Proposition 2's unconditional guarantee for substitutes games means existence needs no separate check. With $A = 8, b = 1$: Nash $a^{\mathrm{NE}} = 8/3$,

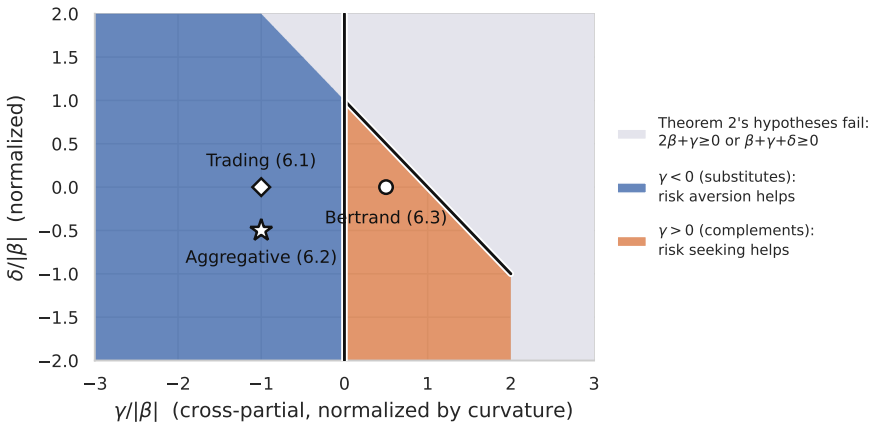

Figure 1: Phase diagram of Theorem 2, in normalized coordinates $(\gamma/|\beta|, \delta/|\beta|)$. Grey: the base game fails one of Theorem 2's two hypotheses (not Theorem 1's λ-dependent construction well-posedness, a different, λ-dependent condition unrelated to this plot). Blue ($\gamma < 0$): risk aversion helps. Orange ($\gamma > 0$): risk seeking helps instead. The three worked examples of Section 6 sit exactly where the theorem predicts.

joint-optimal $a^\dagger = 2$; Nash over-provides relative to the joint optimum, exactly the strategic-substitutes case of Bulow et al. (1985). The risk-averse equilibrium $a^\star(\lambda) = 2A/(6b + b^2\lambda)$ decreases monotonically from a^{NE} at $\lambda = 0$ toward 0 as λ increases, pulling play back toward $a^\dagger$; net welfare peaks exactly there, at $\lambda^\star = 2$, and returns to the Nash level at $\lambda_{\mathrm{thresh}} = 6$.

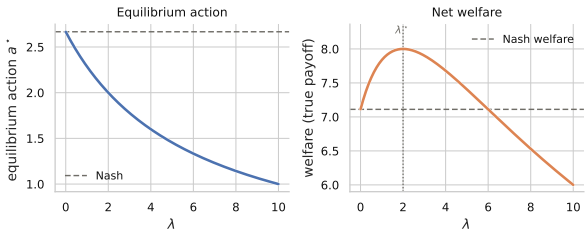

Figure 2: Trading/Cournot (affine): equilibrium action and welfare versus λ.

6.2 A SHARED-EFFORT AGGREGATIVE GAME

A shared-reward, private-cost aggregative game, following Qu et al. (2026)'s original example, gives a substitutes instance ($\gamma < 0$) with $\delta \neq 0$: the $\delta = \eta = 0$ pull-toward-zero picture of Section 6.1 no longer literally applies, and the equilibrium action is in fact pulled *away* from 0 as λ increases. Both of Theorem 2's hypotheses still hold, however, so the theorem still applies and still correctly predicts that risk aversion strictly improves welfare over Nash. The classical over/under-provision correspondence (Corollary 1) does not survive though: despite substitutes ($\gamma < 0$), Nash *under*-provides here ($a^{\mathrm{NE}} = 4/3 < a^\dagger = 1.6$), the reverse of the $\delta = \eta = 0$ picture, even as Theorem 2 continues to identify risk aversion correctly as welfare-improving. The free-lunch phenomenon of Qu et al. (2026) persists, net welfare rising from the Nash level to a peak at an intermediate λ before falling back toward the well-posedness ceiling. Appendix B gives the full derivation, closed forms, and figure.

6.3 RISK-SEEKING EXAMPLE: BERTRAND COMPETITION

Two firms set prices $p_i \geq 0$ for differentiated products (Bertrand, 1883), in the differentiated-duopoly formulation of Singh & Vives (1984): demand for firm i's product is $q_i = a - bp_i + dp_j$

($d > 0$), both firms face the same constant marginal cost c, and firm i's payoff is its own profit,

$$u_i(p_i, p_j) = (p_i - c)(a - bp_i + dp_j), \qquad (\alpha, \beta, \gamma, \delta, \eta) = (a + bc, -b, d, 0, -cd), \tag{5}$$

the instance of equation 1 obtained by expanding and reading off coefficients. This gives a complements instance ($\gamma = d > 0$) with $\eta \neq 0$, to which Theorem 2's general closed form still applies directly. With $a = 10, b = 2, d = 1, c = 1$: Nash $p^{\mathrm{NE}} = 4$, joint-profit-maximizing $p^{\dagger} = 5.5$, and the risk-seeking equilibrium works out to the closed form $p^{\star}(\lambda) = (24 - \lambda)/(6 - \lambda)$. Risk-seeking play raises the equilibrium price monotonically from p^{NE} toward $p^{\dagger}$ as λ rises from 0, reproducing the same welfare hump found under risk aversion in Section 6.1 but for the opposite attitude, peaking exactly at $\lambda^{\star} = 2$; because the two firms here are on the same side of the market, this improvement in joint profit acts as an uncoordinated substitute for collusion, at the expense of the consumers left out of the game. Admissibility (non-negative demand $q^{\star}(\lambda) = a - (b - d)p^{\star}(\lambda)$) binds before the price formula's own singularity at $\lambda_{\mathrm{pole}} = 6$, at $\lambda_{\mathrm{admiss}} = 4$, past which prices are real-valued but describe no real market. Appendix B gives the full derivation.

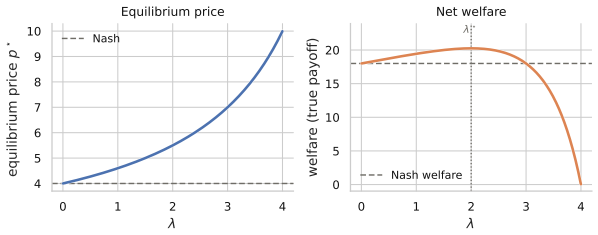

Figure 3: Bertrand (risk-seeking), $\lambda \leq \lambda_{\mathrm{admiss}}$: price and welfare versus λ.

7 Discussion

We set out to extend strategic robustness to the risk-seeking case and to determine, for either attitude, when robustifying an already well-behaved game raises the welfare it delivers. We defined risk-seeking play as the optimistic mirror of the existing risk-averse construction, a fictitious ally in place of a fictitious adversary, and carried it through the same stages of analysis. Well-posedness is far more delicate for risk seeking, failing well before risk aversion's own construction would; we then characterized exactly when a symmetric equilibrium exists and best-response dynamics converge to it, for both attitudes and game types, showed which attitude actually raises welfare as a function of the game's classical substitutes/complements structure alone, with the losing attitude actively hurting, not merely failing to help, and pinned down exactly when the robustified equilibrium reaches the welfare optimum itself, not only moves toward it. We further showed that a classical correspondence between Nash's inefficiency and this same structure, long treated as automatic, holds only on a narrower subclass. Trading, an aggregative game, and Bertrand competition instantiate every regime this theory predicts.

Two things follow beyond the results themselves. First, the mechanism is Pigouvian: the fictitious report a robustified player best-responds to deviates from the truth by a multiple of the same marginal externality one player's action imposes on the other, so the correction it induces in the Nash condition offsets that externality. This is why the matching attitude, risk aversion under substitutes, risk seeking under complements, is the one that helps, and the mismatched attitude amplifies the distortion instead. Second, that structure makes risk seeking the more fragile attitude to deploy, since its well-posedness ceiling binds far more readily than risk aversion's.

A further implication is worth flagging. The classical over- and under-provision correspondence of Bulow et al. (1985) is often treated as an immediate corollary of the strategic-substitutes/complements distinction; Corollary 1 shows this holds only on the tractable $\delta = \eta = 0$ subclass, with the aggregative game giving a concrete case outside it where the correspondence reverses. The two questions are governed by different quantities in general, coinciding only inside that subclass.

THE USE OF LARGE LANGUAGE MODELS

Large language models were used to aid and polish the writing of this paper, and to draft parts of several sections; all mathematical claims and derivations were derived by the authors.

REFERENCES

Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *Nature Human Behaviour*, 9:1380–1390, 2025.

Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory*. SIAM, 2nd edition, 1999.

Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

Joseph Bertrand. Théorie mathématique de la richesse sociale. *Journal des Savants*, pp. 499–508, 1883.

Jeremy I. Bulow, John D. Geanakoplos, and Paul D. Klemperer. Multimarket oligopoly: Strategic substitutes and complements. *Journal of Political Economy*, 93(3):488–511, 1985.

Antoine Augustin Cournot. *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette, 1838.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

Jacob K. Goeree, Charles A. Holt, and Thomas R. Palfrey. Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, 45(1):97–113, 2003.

Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, et al. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990*, 2024.

Garud N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2): 257–280, 2005.

David Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games. *IEEE Transactions on Automatic Control*, 18(2): 124–131, 1973.

Nicolas Lanzetti, Silvia Fricker, Saverio Bolognani, Florian Dörfler, and Dario Paccagnan. Strategically robust game theory via optimal transport. *arXiv preprint arXiv:2507.15325*, 2025.

Nunzio Lorè and Babak Heydari. Strategic behavior of large language models and the role of game structure versus contextual framing. *Scientific Reports*, 14(1):18490, 2024.

Eric Mazumdar, Kishan Panaganti, and Laixi Shi. Tractable multi-agent reinforcement learning through behavioral economics. In *International Conference on Learning Representations*, 2025.

Richard D. McKelvey and Thomas R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.

Paul Milgrom and John Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990.

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.

Arnab Nilim and Laurent El Ghaoui. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, 2023.

Arthur C. Pigou. *The Economics of Welfare*. Macmillan, 1920.

Chang Qu, Yue Zhang, Nicolas Lanzetti, and Eric Mazumdar. Training generalizable collaborative agents via strategic risk aversion. *arXiv preprint arXiv:2602.21515*, 2026.

Tim Roughgarden and Éva Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10): 1095–1100, 1953.

Nirvikar Singh and Xavier Vives. Price and quantity competition in a differentiated duopoly. *RAND Journal of Economics*, 15(4):546–554, 1984.

Donald M. Topkis. Equilibrium points in nonzero-sum n-person submodular games. *SIAM Journal on Control and Optimization*, 17(6):773–787, 1979.

Bojan Velasevic, Nicolas Lanzetti, and Eric Mazumdar. Strategically robust linear quadratic dynamic games. *arXiv preprint arXiv:2604.22318*, 2026.

J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers*, 1(3):325–362, 1952.

Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.

Zaiyan Xu, Kishan Panaganti, and Dileep Kalathil. Improved sample complexity bounds for distributionally robust reinforcement learning. In *Proceedings of The 26th International Conference on Artificial Intelligence and Statistics*, 2023.

Kaiqing Zhang, Tao Sun, Yunzhe Tao, Sahika Genc, Sunil Mallya, and Tamer Başar. Robust multi-agent reinforcement learning with model uncertainty. *Advances in Neural Information Processing Systems*, 33:10571–10583, 2020.

Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. Multi-agent reinforcement learning: A selective overview of theories and algorithms. In *Handbook of Reinforcement Learning and Control*, pp. 321–384. Springer, 2021.

A Deferred Proofs and Supporting Derivations

This appendix gives, in order, the full proof of Theorem 1 (Appendix A.1), whose reaction curve everything afterward builds on; the proof of Proposition 1 (Appendix A.2); a shared lemma on the monotonicity of $L(\lambda)$ (Appendix A.3) and a shared safe-endpoint fact (Appendix A.4) that all four propositions' proofs use, followed by those four proofs themselves (Appendices A.5, A.6, A.7, and A.8); the third of these also derives $L^{\mathrm{seek}}(\lambda)$'s own monotonicity, the risk-seeking analogue of Appendix A.3, which the fourth then reuses; and the full proof of Theorem 2, including its monotonicity claim (Appendix A.9), which also uses the reaction curve.

A.1 Proof of Theorem 1: the reaction curve and equilibrium closed form

This subsection proves Theorem 1 in full: the well-posedness conditions, the effective curvature, the reaction curve $K(\lambda), L(\lambda)$, and, as their direct corollary, the symmetric equilibrium $a^\star(\lambda)$, worked out for both attitudes in turn.

Proof. **Risk aversion.** Fix a_j and solve the inner minimization of Definition 3. The objective $\alpha a_i + \beta a_i^2 + \gamma a_i d + \delta d^2 + \eta d + \frac{1}{\lambda}(d - a_j)^2$ is quadratic in d, with first-order condition

$$\gamma a_i + 2\delta d + \eta + \frac{2}{\lambda}(d - a_j) = 0.$$

This is a genuine minimum exactly when the coefficient of d^2, $2\delta + 2/\lambda$, is positive, i.e. $1 + \delta\lambda > 0$: this is where the risk-averse well-posedness condition of Theorem 1 comes from. On that domain, solving for d gives

$$d^\star(a_i, a_j) = \frac{a_j - \frac{\lambda}{2}(\gamma a_i + \eta)}{1 + \delta\lambda}.$$

Substituting $d^\star$ back into equation 2 turns $u_i^{\text{ave}}(a_i; a_j)$ into an explicit quadratic in a_i alone, with a_j now a parameter. Differentiating with respect to a_i and collecting terms, the coefficient of a_i in the resulting first-order condition is exactly $2\beta_{\text{eff}}^{\text{ave}}(\lambda)$, confirming that $\beta_{\text{eff}}^{\text{ave}}(\lambda) = \beta - \gamma^2\lambda/\big(4(1 + \delta\lambda)\big)$ from the theorem statement is indeed the curvature governing this outer problem. Since $\lambda > 0$ and $1 + \delta\lambda > 0$ on this domain, $\gamma^2\lambda/(4(1 + \delta\lambda)) \geq 0$, so $\beta_{\text{eff}}^{\text{ave}}(\lambda) \leq \beta < 0$ automatically: the outer maximization is always well posed once the inner one is, which is why the risk-averse half of Theorem 1 has only the single stated condition, unlike the risk-seeking half below. Setting the first-order condition to zero and solving for a_i gives the reaction curve

$$a_i = K^{\text{ave}}(\lambda) + L^{\text{ave}}(\lambda)\, a_j,$$

$$K^{\text{ave}}(\lambda) = \frac{-2\alpha - \lambda(2\alpha\delta - \eta\gamma)}{4\beta + \lambda(4\beta\delta - \gamma^2)}, \qquad L^{\text{ave}}(\lambda) = \frac{-2\gamma}{4\beta + \lambda(4\beta\delta - \gamma^2)}. \tag{6}$$

Imposing the symmetric fixed point $a_i = a_j = a$, i.e. $a = K^{\text{ave}}(\lambda) + L^{\text{ave}}(\lambda)\, a$, and solving for a (valid whenever $L^{\text{ave}}(\lambda) \neq 1$) gives

$$a^\star_{\text{ave}}(\lambda) = \frac{K^{\text{ave}}(\lambda)}{1 - L^{\text{ave}}(\lambda)} = \frac{-2\alpha - \lambda(2\alpha\delta - \eta\gamma)}{4\beta + 2\gamma + \lambda(4\beta\delta - \gamma^2)},$$

the closed-form risk-averse equilibrium.

Risk seeking. The derivation is identical, now with $d^\star$ chosen to maximize instead of minimize, which flips the sign of the penalty term in equation 3 throughout. This sign flip has the same effect on the first-order condition as substituting $\lambda \mapsto -\lambda$ everywhere in the risk-averse derivation above, since the two constructions differ only in the sign multiplying $(d - a_j)^2$: every quantity above carries over under that substitution. The effective curvature becomes $\beta_{\text{eff}}^{\text{seek}}(\lambda) = \beta + \gamma^2\lambda/\big(4(1 - \delta\lambda)\big)$, well posed exactly when this is negative. This is a genuine second condition, unlike the risk-averse case: $\gamma^2\lambda/(4(1 - \delta\lambda)) \geq 0$ pushes the effective curvature toward 0, so negativity is no longer automatic once $1 - \delta\lambda > 0$ holds alone. Clearing denominators (valid since $1 - \delta\lambda > 0$) turns $\beta + \gamma^2\lambda/(4(1 - \delta\lambda)) < 0$ into $4\beta(1 - \delta\lambda) + \gamma^2\lambda < 0$; dividing by $4\beta < 0$, which reverses the inequality, gives the single closed-form threshold $\lambda(\delta + \gamma^2/(4|\beta|)) < 1$ stated in Theorem 1. The reaction curve is

$$K^{\text{seek}}(\lambda) = \frac{-2\alpha - \lambda(\eta\gamma - 2\alpha\delta)}{4\beta + \lambda(\gamma^2 - 4\beta\delta)}, \qquad L^{\text{seek}}(\lambda) = \frac{-2\gamma}{4\beta + \lambda(\gamma^2 - 4\beta\delta)}, \tag{7}$$

and imposing the symmetric fixed point exactly as above gives

$$a^\star_{\text{seek}}(\lambda) = \frac{K^{\text{seek}}(\lambda)}{1 - L^{\text{seek}}(\lambda)} = \frac{-2\alpha - \lambda(\eta\gamma - 2\alpha\delta)}{4\beta + 2\gamma + \lambda(\gamma^2 - 4\beta\delta)},$$

the closed-form risk-seeking equilibrium. □

A.2 Proof of Proposition 1: sign of the reaction-curve slope

Proof. The denominator of $L(\lambda)$ is, in either case, a product of a positive term, $4(1+\delta\lambda)$ for risk aversion and $4(1-\delta\lambda)$ for risk seeking, positive exactly by Theorem 1's well-posedness condition ($1+\delta\lambda>0$, respectively $1-\delta\lambda>0$), and the effective curvature, negative by well-posedness, hence always negative; since $\lambda>0$, the numerator has sign $-\mathrm{sign}(\gamma)$, giving $\mathrm{sign}(L(\lambda))=\mathrm{sign}(\gamma)$. For the dynamics, the linear recursion $a_{t+1}=K(\lambda)+L(\lambda)a_t$ satisfies $a_t-a^\star=L(\lambda)^t(a_0-a^\star)$ for any $a_0\neq a^\star$: the sign of $L(\lambda)^t$ alternates with t exactly when $L(\lambda)<0$ and is constant when $L(\lambda)>0$, regardless of $|L(\lambda)|$; separately, $|L(\lambda)^t|\to 0$ iff $|L(\lambda)|<1$ and $|L(\lambda)^t|\to\infty$ iff $|L(\lambda)|>1$. □

A.3 Monotonicity of the reaction-curve slope $L^{\mathrm{ave}}(\lambda)$

Proposition 2 and 3 both turn on the same fact about how $L^{\mathrm{ave}}(\lambda)$ moves across the well-posed domain. Proved once here, it is used directly in each proposition's own proof below.

Lemma 1. *On the well-posed domain, $L^{\mathrm{ave}}(\lambda)=2\gamma/D(\lambda)$ where $D(\lambda):=4|\beta|(1+\delta\lambda)+\gamma^2\lambda>0$. $L^{\mathrm{ave}}(\lambda)$ is strictly monotone in λ, and, as λ ranges over the domain, takes exactly every value strictly between $L_0^{\mathrm{ave}}:=\gamma/(2|\beta|)$ and $L_{\bar\lambda}^{\mathrm{ave}}$, its limits at the two ends of the domain, and no value outside that open interval.*

Proof. **Step 1: rewrite the denominator.** By equation 6, $L^{\mathrm{ave}}(\lambda)=-2\gamma/\big(4\beta+\lambda(4\beta\delta-\gamma^2)\big)$. Since $\beta<0$, write

$$D(\lambda):=-\big(4\beta+\lambda(4\beta\delta-\gamma^2)\big)=-4\beta-4\beta\delta\lambda+\gamma^2\lambda=4|\beta|(1+\delta\lambda)+\gamma^2\lambda,$$

using $-\beta=|\beta|$, so $L^{\mathrm{ave}}(\lambda)=2\gamma/D(\lambda)$.

Step 2: $D(\lambda)>0$ throughout the domain. On the well-posed domain, $1+\delta\lambda>0$ (Theorem 1), so $4|\beta|(1+\delta\lambda)>0$; separately $\gamma^2\lambda\geq 0$ since $\lambda>0$. $D(\lambda)$ is a sum of a strictly positive and a nonnegative term, hence strictly positive, so $L^{\mathrm{ave}}(\lambda)$ is finite and well defined everywhere on the domain.

Step 3: monotonicity. Differentiating, $D'(\lambda)=4|\beta|\delta+\gamma^2$, and since the numerator 2γ is constant in λ,

$$L'(\lambda)=\frac{-2\gamma D'(\lambda)}{D(\lambda)^2}=\frac{-2\gamma(\gamma^2+4|\beta|\delta)}{D(\lambda)^2}.$$

$D(\lambda)^2>0$ throughout by Step 2, so the sign of $L'(\lambda)$ is fixed for every λ in the domain by the constant $-\gamma(\gamma^2+4|\beta|\delta)$: it never changes sign, so $L^{\mathrm{ave}}(\lambda)$ is strictly monotone ($\gamma\neq 0$ throughout the paper, since $\gamma=0$ is neither substitutes nor complements).

Step 4: endpoint values. As $\lambda\to 0$: $D(0)=4|\beta|$, so $L_0^{\mathrm{ave}}=2\gamma/(4|\beta|)=\gamma/(2|\beta|)$. As $\lambda\to\bar\lambda$: if $\delta\geq 0$, $\bar\lambda=\infty$ and $D(\lambda)\to\infty$ (the $\gamma^2\lambda$ term alone diverges), so $L_{\bar\lambda}^{\mathrm{ave}}=\lim 2\gamma/D(\lambda)=0$; if $\delta<0$, $\bar\lambda=-1/\delta$ and

$$D(-1/\delta)=4|\beta|\Big(1+\delta\cdot\big(-\tfrac{1}{\delta}\big)\Big)+\gamma^2\Big(-\tfrac{1}{\delta}\Big)=4|\beta|\cdot 0-\frac{\gamma^2}{\delta}=-\frac{\gamma^2}{\delta},$$

positive since $\delta<0$, consistent with Step 2; hence $L_{\bar\lambda}^{\mathrm{ave}}=2\gamma/(-\gamma^2/\delta)=-2\delta/\gamma$.

Step 5: the swept range. $L^{\mathrm{ave}}(\lambda)$ is continuous on the domain (Step 2: $D(\lambda)\neq 0$ throughout) and strictly monotone (Step 3) on the connected, open interval $(0,\bar\lambda)$, so by the intermediate value theorem it takes every value strictly between its two limits L_0^{ave} and $L_{\bar\lambda}^{\mathrm{ave}}$ exactly once, and, being monotone, no value outside that open interval anywhere in the domain. □

A.4 A safe-endpoint fact shared by all four propositions

Propositions 2, 3, 4, and 5 all locate their convergent region by the same short argument: let $\ell(\lambda)$ be continuous and strictly monotone on a connected domain $(0,\Lambda)$, $\Lambda\in(0,\infty]$, with limits $\ell_0:=\lim_{\lambda\to 0}\ell(\lambda)$ and $\ell_\Lambda:=\lim_{\lambda\to\Lambda}\ell(\lambda)$ (each finite or infinite), both sharing a fixed sign, and let $T\in\{-1,1\}$ share that sign. Suppose T lies strictly between ℓ_0 and ℓ_Λ. Since ℓ maps $(0,\Lambda)$ bijectively and monotonically onto the open interval between ℓ_0 and ℓ_Λ (Lemma 1's argument), ℓ crosses T exactly once, at $\lambda_T:=\ell^{-1}(T)\in(0,\Lambda)$. Because ℓ_0 and ℓ_Λ share T's sign and T

sits strictly between them, exactly one of the two already has magnitude below 1; call it the *safe endpoint*. Monotonicity carries that safety all the way to λ_T: $|\ell(\lambda)| < 1$ throughout the piece of the domain adjacent to the safe endpoint, and $|\ell(\lambda)| > 1$ throughout the other piece, the one adjacent to the endpoint already above 1 in magnitude.

The four proofs below instantiate this with $\ell = L^{\text{ave}}$, $\Lambda = \bar{\lambda}$ (always finite-valued) for the two risk-averse propositions, and $\ell = L^{\text{seek}}$, $\Lambda = \bar{\lambda}^{\text{seek}}$ (whose far limit can be infinite) for the two risk-seeking ones.

A.5 Proof of Proposition 2: substitutes ($\gamma < 0$)

Proof. A symmetric risk-averse equilibrium exists at λ exactly when $L^{\text{ave}}(\lambda) \neq 1$, and dynamics converge exactly when $|L^{\text{ave}}(\lambda)| < 1$ and diverge exactly when $|L^{\text{ave}}(\lambda)| > 1$ (Section 5.2). By Lemma 1, $L^{\text{ave}}(\lambda)$ takes any given target value somewhere in the domain exactly when that value lies strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$, and never otherwise.

Existence. $L_0^{\text{ave}} = \gamma/(2|\beta|) < 0$. For $L_{\bar{\lambda}}^{\text{ave}}$: if $\delta \geq 0$, $L_{\bar{\lambda}}^{\text{ave}} = 0$; if $\delta < 0$, $L_{\bar{\lambda}}^{\text{ave}} = -2\delta/\gamma$, and since $-2\delta > 0$ while $\gamma < 0$, this is negative too. Either way $L_0^{\text{ave}}, L_{\bar{\lambda}}^{\text{ave}} \leq 0$, so 1 (positive) is never strictly between two numbers that are both at most 0. By Lemma 1, $L^{\text{ave}}(\lambda)$ therefore never reaches 1, and a symmetric risk-averse equilibrium exists for every λ in the well-posed domain.

Convergence. By Lemma 1, $L^{\text{ave}}(\lambda)$ is monotone and sweeps exactly the open interval between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$. Two outcomes follow, according to where -1 sits relative to that interval; we first show the betweenness condition itself reduces to equation 4.

-1 lies strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$ exactly when $(L_0^{\text{ave}} + 1)(L_{\bar{\lambda}}^{\text{ave}} + 1) < 0$. If $\delta \geq 0$, $L_{\bar{\lambda}}^{\text{ave}} = 0$, so $L_{\bar{\lambda}}^{\text{ave}} + 1 = 1 > 0$ always, and the product is negative exactly when $L_0^{\text{ave}} + 1 < 0$, i.e. $\gamma < -2|\beta|$, i.e. $|\gamma| > 2|\beta|$; since $\delta \geq 0$ makes $\delta + |\gamma|/2 > 0$ too, this is exactly $|\gamma| > 2|\beta|$ combined with $\delta + |\gamma|/2 > 0$, i.e. equation 4. If $\delta < 0$, $L_{\bar{\lambda}}^{\text{ave}} = -2\delta/\gamma$, so $L_{\bar{\lambda}}^{\text{ave}} + 1 = (\gamma - 2\delta)/\gamma$, and

$$(L_0^{\text{ave}} + 1)(L_{\bar{\lambda}}^{\text{ave}} + 1) = \frac{\gamma + 2|\beta|}{2|\beta|} \cdot \frac{\gamma - 2\delta}{\gamma}.$$

Since $2|\beta| > 0$ and $\gamma < 0$, this is negative exactly when $(\gamma + 2|\beta|)(\gamma - 2\delta) > 0$; substituting $\gamma + 2|\beta| = -(|\gamma| - 2|\beta|)$ and $\gamma - 2\delta = -2(\delta + |\gamma|/2)$ (both valid since $\gamma < 0$) turns this into $2(|\gamma| - 2|\beta|)(\delta + |\gamma|/2) > 0$, again equation 4. So in either case, -1 lies strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$ exactly when equation 4 holds.

If equation 4 holds, -1 lies strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$, so the safe-endpoint fact (Appendix A.4) applies with $\ell = L^{\text{ave}}$, $\Lambda = \bar{\lambda}$, $T = -1$: $L^{\text{ave}}(\lambda)$ crosses -1 exactly once, at the λ solving $L^{\text{ave}}(\lambda) = -1$, i.e. $D(\lambda) = -2\gamma$:

$$\begin{aligned} 4|\beta|(1 + \delta\lambda) + \gamma^2\lambda = -2\gamma &\iff \lambda\big(4|\beta|\delta + \gamma^2\big) = -2\gamma - 4|\beta| \\ &\iff \lambda_{-1} = \frac{-2\gamma - 4|\beta|}{\gamma^2 + 4|\beta|\delta} = \frac{-2(\gamma + 2|\beta|)}{\gamma^2 + 4|\beta|\delta}, \end{aligned}$$

and $|L^{\text{ave}}(\lambda)| < 1$ throughout the piece of the domain adjacent to whichever of $L_0^{\text{ave}}, L_{\bar{\lambda}}^{\text{ave}}$ is the safe endpoint: $|L_0^{\text{ave}}| < 1$ exactly when $|\gamma| < 2|\beta|$, independent of λ, giving convergence below λ_{-1} and divergence, still oscillating (Proposition 1), beyond it in that case; if instead $|\gamma| \geq 2|\beta|$, then L_0^{ave} itself already has magnitude at least 1 (the Nash best response oscillates without settling even before any robustification), so $L_{\bar{\lambda}}^{\text{ave}}$ is the safe endpoint instead, giving divergence below λ_{-1} and convergence beyond it, robustification repairing what Nash could not.

If instead equation 4 fails, -1 does not lie strictly between L_0^{ave} and $L_{\bar{\lambda}}^{\text{ave}}$ (the safe-endpoint fact's premise fails), so the entire swept interval lies on one side of -1: either $L_0^{\text{ave}}, L_{\bar{\lambda}}^{\text{ave}} \in (-1, 0)$, so $|L^{\text{ave}}(\lambda)| < 1$ throughout and dynamics converge for every λ; or $L_0^{\text{ave}}, L_{\bar{\lambda}}^{\text{ave}} \leq -1$, so $|L^{\text{ave}}(\lambda)| \geq 1$ throughout and dynamics diverge for every λ, a Nash best response already too steep for any amount of robustification to repair. □

A.6 Proof of Proposition 3: complements ($\gamma > 0$)

Proof. As in the substitutes case, a symmetric risk-averse equilibrium exists at λ exactly when $L^{\text{ave}}(\lambda) \neq 1$, and dynamics converge exactly when $|L^{\text{ave}}(\lambda)| < 1$ and diverge exactly when $|L^{\text{ave}}(\lambda)| > 1$ (Section 5.2). Here $L_0^{\text{ave}} = \gamma/(2|\beta|) > 0$ (Proposition 1), so $L^{\text{ave}}(\lambda) > 0$ throughout the domain (Lemma 1: it is monotone and never crosses 0 to reach a negative value on the way from L_0^{ave} to $L_{\bar\lambda}^{\text{ave}}$, both nonnegative under this sign of γ), so the only boundary either question can reach is $L^{\text{ave}}(\lambda) = 1$; there is no separate $L^{\text{ave}}(\lambda) = -1$ crossing to locate, unlike under substitutes.

By Lemma 1, 1 lies strictly between L_0^{ave} and $L_{\bar\lambda}^{\text{ave}}$ exactly when $L^{\text{ave}}(\lambda)$ reaches 1 somewhere in the domain, at the single point where $L^{\text{ave}}(\lambda) = 1$. The same computation as in the substitutes proof (Appendix A.5), with $\gamma > 0$ in place of $\gamma < 0$ and target 1 in place of -1, reduces this betweenness condition to the same inequality equation 4: $(L_0^{\text{ave}} - 1)(L_{\bar\lambda}^{\text{ave}} - 1) < 0$ simplifies, via $\gamma - 2|\beta| = |\gamma| - 2|\beta|$ and $2\delta + \gamma = 2(\delta + |\gamma|/2)$ (both valid since $\gamma > 0$), to $(|\gamma| - 2|\beta|)(\delta + |\gamma|/2) > 0$, so equation 4 again governs exactly when a crossing exists. Solving explicitly:

$$\begin{aligned} L^{\text{ave}}(\lambda) = \frac{2\gamma}{D(\lambda)} = 1 &\iff D(\lambda) = 2\gamma \\ &\iff 4|\beta|(1 + \delta\lambda) + \gamma^2\lambda = 2\gamma \\ &\iff \lambda\big(4|\beta|\delta + \gamma^2\big) = 2\gamma - 4|\beta|, \end{aligned}$$

so

$$\lambda_1 = \frac{2\gamma - 4|\beta|}{\gamma^2 + 4|\beta|\delta} = \frac{2(\gamma - 2|\beta|)}{\gamma^2 + 4|\beta|\delta},$$

matching the threshold in the proposition statement, valid when equation 4 holds. Monotonicity (Lemma 1) ensures λ_1 is the only point where $L^{\text{ave}}(\lambda) = 1$; at every other λ in the well-posed domain, $L^{\text{ave}}(\lambda) \neq 1$, so the fixed-point equation has a unique solution and a symmetric equilibrium exists: existence fails only at the single point λ_1 itself, where the equation degenerates, not on either side of it.

Since 1 lies strictly between L_0^{ave} and $L_{\bar\lambda}^{\text{ave}}$ here, the safe-endpoint fact (Appendix A.4) applies with $\ell = L^{\text{ave}}$, $\Lambda = \bar\lambda$, $T = 1$: $|L^{\text{ave}}(\lambda)| < 1$ throughout the piece of the domain adjacent to the safe endpoint. Since $|L_0^{\text{ave}}| < 1$ exactly when $|\gamma| < 2|\beta|$ (independent of δ): if $|\gamma| < 2|\beta|$, L_0^{ave} is the safe endpoint, giving convergence below λ_1 and divergence beyond it; if instead $|\gamma| \geq 2|\beta|$, then L_0^{ave} itself already has magnitude at least 1 (the Nash best response already diverges before any robustification), so $L_{\bar\lambda}^{\text{ave}}$ is the safe endpoint instead, giving divergence below λ_1 and convergence beyond it.

If instead equation 4 fails, 1 does not lie strictly between L_0^{ave} and $L_{\bar\lambda}^{\text{ave}}$ (the safe-endpoint fact's premise fails), so the entire swept interval lies on one side of 1: either $L_0^{\text{ave}}, L_{\bar\lambda}^{\text{ave}} \in (0, 1)$, so $L^{\text{ave}}(\lambda)$ never reaches 1 (existence holds throughout) and $|L^{\text{ave}}(\lambda)| < 1$ throughout (dynamics converge throughout); or $L_0^{\text{ave}}, L_{\bar\lambda}^{\text{ave}} \geq 1$, so $L^{\text{ave}}(\lambda)$ again never reaches exactly 1 on the open swept interval (existence still holds throughout), but $L^{\text{ave}}(\lambda) \geq 1$ throughout, generically > 1 in the interior (dynamics diverge throughout), a Nash best response already too steep for any amount of robustification to repair. □

A.7 Proof of Proposition 4: risk-seeking, substitutes ($\gamma < 0$)

Proof. By equation 7, $L^{\text{seek}}(\lambda) = 2\gamma/D^{\text{seek}}(\lambda)$ where $D^{\text{seek}}(\lambda) := 4|\beta| - \lambda(\gamma^2 + 4|\beta|\delta)$. When $\gamma^2 + 4|\beta|\delta \leq 0$ ($\bar\lambda^{\text{seek}}$ infinite), $D^{\text{seek}}(\lambda) \geq 4|\beta| > 0$ throughout $(0, \infty)$, increasing without bound, so $L^{\text{seek}}(\lambda) \to 0$ as $\lambda \to \infty$; when $\gamma^2 + 4|\beta|\delta > 0$ ($\bar\lambda^{\text{seek}}$ finite), $D^{\text{seek}}(\lambda)$ decreases linearly from $4|\beta|$ to 0 exactly at $\bar\lambda^{\text{seek}} = 4|\beta|/(\gamma^2 + 4|\beta|\delta)$, forcing $K^{\text{seek}}(\lambda), L^{\text{seek}}(\lambda)$ to blow up there, unlike the risk-averse case, so $L^{\text{seek}}(\lambda) \to \text{sign}(\gamma) \cdot \infty$ as $\lambda \uparrow \bar\lambda^{\text{seek}}$. Differentiating, $dL^{\text{seek}}/d\lambda = 2\gamma(\gamma^2 + 4|\beta|\delta)/D^{\text{seek}}(\lambda)^2$, whose sign is fixed throughout the domain, depending only on the (constant) signs of γ and $\gamma^2 + 4|\beta|\delta$; so $L^{\text{seek}}(\lambda)$ is strictly monotone on $(0, \bar\lambda^{\text{seek}})$, from $L_0^{\text{seek}} = \gamma/(2|\beta|)$ to the far limit stated in the proposition. This is the analogue of Lemma 1 for L^{seek}, letting the safe-endpoint fact (Appendix A.4) apply with $\ell = L^{\text{seek}}$.

Case $|\gamma| < 2|\beta|$. Then $|L_0^{\text{seek}}| < 1$, so L_0^{seek} is safe whenever a comparison is needed. If $\bar\lambda^{\text{seek}}$ is finite, the far limit $-\infty$ is unsafe, so the safe-endpoint fact gives a unique crossing solving

$L^{\text{seek}}(\lambda) = -1$:

$$\frac{2\gamma}{D^{\text{seek}}(\lambda)} = -1 \iff D^{\text{seek}}(\lambda) = -2\gamma$$
$$\iff 4|\beta| - \lambda(\gamma^2 + 4|\beta|\delta) = -2\gamma$$
$$\iff \lambda_{-1}^{\text{seek}} = \frac{4|\beta| + 2\gamma}{\gamma^2 + 4|\beta|\delta} = \frac{2(2|\beta| + \gamma)}{\gamma^2 + 4|\beta|\delta},$$

matching the proposition; convergence holds on the piece adjacent to L_0^{seek}, i.e. below $\lambda_{-1}^{\text{seek}}$, divergence beyond it. If instead $\bar{\lambda}^{\text{seek}}$ is infinite, the far limit 0 is also safe, so -1 does not lie strictly between the two endpoints and the safe-endpoint fact's premise fails: no crossing exists, and $|L^{\text{seek}}(\lambda)| < 1$ throughout, since monotonicity confines every swept value to the interval $(-1, 0)$ between the two safe endpoints.

Case $|\gamma| \geq 2|\beta|$. Then $|L_0^{\text{seek}}| \geq 1$, so L_0^{seek} is itself unsafe. If $\bar{\lambda}^{\text{seek}}$ is finite, the far limit $(-\infty)$ is unsafe too, so no crossing exists inside the domain and $|L^{\text{seek}}(\lambda)| \geq |L_0^{\text{seek}}| \geq 1$ throughout (monotonicity only carries $L^{\text{seek}}(\lambda)$ further from 0 than L_0^{seek} as λ increases toward the pole). If instead $\bar{\lambda}^{\text{seek}}$ is infinite, the far limit 0 is safe, so the same equation above gives the same threshold $\lambda_{-1}^{\text{seek}}$, now with L_0^{seek} on the unsafe side: divergence holds below $\lambda_{-1}^{\text{seek}}$ (adjacent to the unsafe endpoint L_0^{seek}) and convergence beyond it (adjacent to the safe endpoint 0).

In every sub-case, $L^{\text{seek}}(\lambda) = -1$ only at $\lambda_{-1}^{\text{seek}}$ when that threshold falls inside $(0, \bar{\lambda}^{\text{seek}})$, and never otherwise (the case split above shows $L^{\text{seek}}(\lambda)$ stays fixed on one side of -1 or the other away from it), so a symmetric risk-seeking equilibrium exists throughout the well-posed domain except exactly at that one point. □

A.8 Proof of Proposition 5: risk-seeking, complements ($\gamma > 0$)

Proof. Identical to the substitutes proof above with $\gamma > 0$ throughout and $T = 1$ in place of $T = -1$: $L_0^{\text{seek}} = \gamma/(2|\beta|) > 0$, monotone as established there, with far limit $+\infty$ when $\bar{\lambda}^{\text{seek}}$ is finite, 0 otherwise. Solving $L^{\text{seek}}(\lambda) = 1$:

$$\frac{2\gamma}{D^{\text{seek}}(\lambda)} = 1 \iff D^{\text{seek}}(\lambda) = 2\gamma$$
$$\iff 4|\beta| - \lambda(\gamma^2 + 4|\beta|\delta) = 2\gamma$$
$$\iff \lambda_{1}^{\text{seek}} = \frac{4|\beta| - 2\gamma}{\gamma^2 + 4|\beta|\delta} = \frac{2(2|\beta| - \gamma)}{\gamma^2 + 4|\beta|\delta},$$

matching the proposition. Geometrically, a reaction line whose slope exceeds 1 can only meet the 45° line at a negative action, so the equilibrium is real but economically inadmissible on the far side of λ_1^{seek} before $\bar{\lambda}^{\text{seek}}$ itself is reached, exactly as worked out numerically in Section 6.3. The same case split as the substitutes proof applies with 1 in place of -1: if $|\gamma| < 2|\beta|$ and $\bar{\lambda}^{\text{seek}}$ is finite, convergence below λ_1^{seek}, divergence beyond; if $|\gamma| < 2|\beta|$ and $\bar{\lambda}^{\text{seek}}$ is infinite, convergence throughout; if $|\gamma| \geq 2|\beta|$ and $\bar{\lambda}^{\text{seek}}$ is finite, divergence throughout; if $|\gamma| \geq 2|\beta|$ and $\bar{\lambda}^{\text{seek}}$ is infinite, divergence below λ_1^{seek} and convergence beyond. As in the substitutes case, a symmetric risk-seeking equilibrium exists throughout the well-posed domain except exactly at λ_1^{seek}, when that threshold falls inside it. □

A.9 Proof of Theorem 2: the sign identity and the winning attitude's monotonicity

With $a^\star_{\text{ave}}(\lambda)$ and $a^\star_{\text{seek}}(\lambda)$ established in Appendix A.1, this subsection proves Theorem 2 in full, including the monotonicity claim, by differentiating them, comparing signs, and then locating each construction's singularity. The strategy is to show that risk aversion moves the equilibrium in the direction that raises welfare exactly when it moves it toward $a^\dagger$, the welfare-maximizing action (this equivalence itself follows below, from the strict concavity of W); so it suffices to determine the sign of $a^\star_{\text{ave}}(\lambda) - a^{\text{NE}}$ as λ increases from 0 and compare it against the sign of $a^\dagger - a^{\text{NE}}$. Both turn out to be governed by the same single quantity B, defined below, which is what makes the comparison possible in closed form; monotonicity then follows once we show the winning attitude's own construction carries no interior singularity, so that this sign never changes.

Proof. Restating both closed forms together, for reference below,

$$a^\star_{\text{ave}}(\lambda) = \frac{-2\alpha - \lambda(2\alpha\delta - \eta\gamma)}{4\beta + 2\gamma + \lambda(4\beta\delta - \gamma^2)}, \qquad a^\star_{\text{seek}}(\lambda) = \frac{-2\alpha - \lambda(\eta\gamma - 2\alpha\delta)}{4\beta + 2\gamma + \lambda(\gamma^2 - 4\beta\delta)}, \tag{8}$$

the second obtained from the first by $\lambda \mapsto -\lambda$; both reduce to a^{NE} as $\lambda \to 0$. Differentiating $a^\star_{\text{ave}}(\lambda)$ with respect to λ and clearing denominators gives

$$\frac{da^\star_{\text{ave}}}{d\lambda} = \frac{-2\gamma\big(\alpha(2\delta+\gamma) - \eta(2\beta+\gamma)\big)}{\big(4\beta + 2\gamma + \lambda(4\beta\delta - \gamma^2)\big)^2} = \frac{-2\gamma B}{\big(4\beta + 2\gamma + \lambda(4\beta\delta - \gamma^2)\big)^2},$$
$$B := \alpha(2\delta+\gamma) - \eta(2\beta+\gamma).$$

The denominator is a square and is nonzero away from the risk-averse pole, so $\operatorname{sign}(da^\star_{\text{ave}}/d\lambda) = -\operatorname{sign}(\gamma B)$ everywhere $a^\star_{\text{ave}}$ is defined; the same computation for $a^\star_{\text{seek}}(\lambda)$, using $\lambda \mapsto -\lambda$, gives $\operatorname{sign}(da^\star_{\text{seek}}/d\lambda) = \operatorname{sign}(\gamma B)$. Separately, subtracting $a^\dagger = -(\alpha+\eta)/(2(\beta+\gamma+\delta))$ from $a^{\text{NE}} = -\alpha/(2\beta+\gamma)$ and clearing denominators gives

$$a^{\text{NE}} - a^\dagger = \frac{-B}{2(2\beta+\gamma)(\beta+\gamma+\delta)},$$

the identical constant B appearing in both places, unaffected by the reparametrization.

The direction claim follows: under the stated hypotheses, $\operatorname{sign}(a^{\text{NE}} - a^\dagger) = -\operatorname{sign}(B)$, so risk aversion moves the equilibrium toward $a^\dagger$ as λ increases from 0 exactly when $-\operatorname{sign}(\gamma B) = \operatorname{sign}(B)$, i.e. $\gamma < 0$ (the case $B = 0$ gives $a^{\text{NE}} = a^\dagger$ identically, where the claim is vacuous); strict concavity of W then converts "moves toward $a^\dagger$" into "strictly increases welfare," and continuity in λ extends this from a local statement at $\lambda = 0$ to the entire connected well-posed domain. The risk-seeking case follows symmetrically from the opposite sign of $da^\star_{\text{seek}}/d\lambda$.

For monotonicity, this same constant-sign derivative already does most of the work: for the winning attitude, γ and B carry a fixed relationship to each other throughout, so the sign above is genuinely constant, not just at $\lambda = 0$, wherever $a^\star(\lambda)$ is defined. What remains is to check that $a^\star(\lambda)$ is in fact defined throughout, i.e. that its denominator never vanishes inside the well-posed domain. The denominator of $a^\star_{\text{seek}}(\lambda)$ above vanishes at $\lambda_{\text{pole}} = 2(2\beta+\gamma)/(4\beta\delta - \gamma^2)$, the reciprocal of $M_{\text{pole}} = (4\beta\delta - \gamma^2)/(2(2\beta+\gamma))$; this recovers exactly Proposition 5's singularity, present whenever $\gamma > 0$ and $\delta > -\gamma^2/4$. By the same computation, the denominator of $a^\star_{\text{ave}}(\lambda)$ vanishes at

$$\lambda^{\text{ave}}_{\text{pole}} = \frac{2(2\beta+\gamma)}{\gamma^2 - 4\beta\delta} = -\lambda_{\text{pole}}.$$

Whether this point lies inside the risk-averse construction's well-posed domain $\lambda \in (0, \infty)$, $1 + \delta\lambda > 0$ (Theorem 1), is a parameter-space question unaffected by the choice of M or λ as coordinate, since the two describe the identical set of admissible points; solving (equivalently, in either coordinate) gives that it lies inside if and only if $\delta < -\gamma/2$, obtained from $2\beta + \gamma < 0$ fixing the direction of the resulting inequality.

Since this new risk-averse singularity requires $\gamma > 0$ (as $\delta < -\gamma/2$ can only bind for $\gamma > 0$ within the well-posed region, where it must also satisfy $\delta < 1 - \gamma$), and this theorem identifies risk aversion as the winning attitude only when $\gamma < 0$, the two conditions $\gamma > 0$ and $\gamma < 0$ are mutually exclusive: the new risk-averse singularity and the region where risk aversion wins are disjoint by construction, and the singularity can only ever afflict risk aversion in exactly the region where it is already the losing attitude. The same disjointness holds, a fortiori, for risk-seeking's own singularity relative to the region where risk-seeking is losing ($\gamma < 0$), since that singularity likewise requires $\gamma > 0$. So the winning attitude's own construction is never the one carrying the singularity: it is defined without interruption across the entire well-posed domain, and a function with a constant-sign derivative on a connected interval is strictly monotonic on it. □

A.10 Proof of Corollary 1: Nash over- and under-provision

Proof. At $\delta = \eta = 0$, $a^{\text{NE}} = -\alpha/(2\beta+\gamma)$ (Definition 1) and $a^\dagger = -\alpha/(2(\beta+\gamma))$, so

$$a^{\text{NE}} - a^\dagger = \frac{-\alpha\gamma}{2(\beta+\gamma)(2\beta+\gamma)}.$$

Under Theorem 2's hypotheses, $2\beta + \gamma < 0$ and $\beta + \gamma < 0$ (the latter being $\beta + \gamma + \delta < 0$ at $\delta = 0$), so the denominator is positive, and $\operatorname{sign}(a^{\text{NE}} - a^\dagger) = -\operatorname{sign}(\alpha\gamma) = -\operatorname{sign}(\gamma)$ given $\alpha > 0$. □

B Additional Worked Examples

This appendix gives the full derivations for the aggregative game and the Bertrand example summarized in Sections 6.2 and 6.3.

B.1 A shared-effort aggregative game, in full

Two agents exert effort $a_i \geq 0$ toward a shared goal with reward $R(s) = \frac{H}{2}s^2 + hs$ on the sum $s = a_i + a_j$ (following Qu et al., 2026, $R(a) = -\frac{1}{2}\|\sum a_i - \bar{a}\|_H^2$ with $H \succ 0$; our H absorbs the sign, so $H < 0$ here and $h = -H\bar{a}$), and private cost $\frac{\rho}{2}a_i^2$, the instance of equation 1 with

$$u_i(a_i, a_j) = \frac{H-\rho}{2}a_i^2 + Ha_ia_j + \frac{H}{2}a_j^2 + ha_i + ha_j, \qquad (\alpha, \beta, \gamma, \delta, \eta) = \left(h, \frac{H-\rho}{2}, H, \frac{H}{2}, h\right). \tag{9}$$

This satisfies $\gamma = 2\delta$ and $\alpha = \eta$: both players share the identical reward $R(a_i + a_j)$, each paying only their own private cost, the structure underlying Qu et al. (2026)'s original example. This is a substitutes game ($\gamma = H < 0$), so Proposition 2's unconditional guarantee again means existence needs no separate check; but $\delta = H/2 \neq 0$: the $\delta = \eta = 0$ subclass's geometric picture of pulling toward 0 does not literally apply. Applying Theorem 1 directly, the sign of $da^\star/d\lambda$ is governed here by $-\mathrm{sign}(Hh\rho)$ (the constant factor of 2 in the general expression does not affect the sign), a rule distinct from $\mathrm{sign}(\gamma)$: with $H = -1, \rho = 1, h = \bar{a} = 4$, the Nash action is $a^{\mathrm{NE}} = 4/3$, and the action is pulled *away* from 0 as λ increases, the opposite direction from Section 6.1, converging to $-\eta/(4\delta) = 2$ as λ approaches the well-posedness ceiling $\bar{\lambda} = 1/|\delta| = 2$. Yet this game satisfies both of Theorem 2's hypotheses ($2\beta + \gamma = -3 < 0$, $\beta + \gamma + \delta = -2.5 < 0$), so the general theorem applies and correctly predicts risk aversion still strictly improves welfare over Nash, toward the joint optimum $a^\dagger = 1.6$: the welfare statement survives even though its pull-toward-zero mechanism does not. This game also satisfies $\gamma = 2\delta$, the knife-edge case of Proposition 2: its convergence threshold coincides with the existence boundary, so best-response dynamics converge throughout.

The free-lunch phenomenon nonetheless persists, exactly as in Qu et al. (2026)'s own Example 4.2, on which this example is based: the per-player payoff $u_i(a, a)$, matching that paper's own reward convention, rises from the Nash level (6.22) to a peak (6.40) at an intermediate $\lambda = 1$, distinct from both the Nash limit and the ceiling value, before falling back toward the ceiling (6.00 as $\lambda \to 2^-$); the summed welfare $W(a)$ of Section 5.5 is exactly twice each value. This matches Qu et al. (2026)'s own closed-form reward for this example, verified independently in their native parametrization.

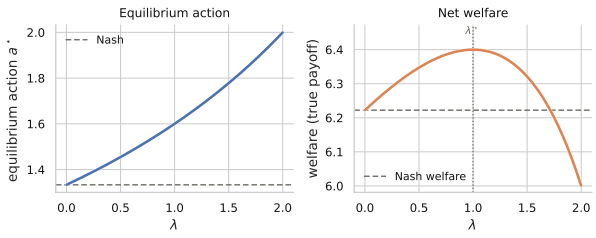

Figure 4: Aggregative game: action pulled away from 0 (left); welfare peaks at intermediate λ (right).

B.2 Bertrand competition, in full

Two firms set prices $p_i \geq 0$ for differentiated products, with demand, cost, and payoff given in Section 6.3 (Equation 5). The mapping to Definition 1 there drops the additive constant $-ac$ produced by expanding equation 5: it affects neither best responses nor equilibria, but it does affect absolute payoff values, so the welfare figure below is computed from equation 5 directly rather than from the constant-free $(\alpha, \beta, \gamma, \delta, \eta)$ form. Here $\gamma = d > 0$: a complements game. Although $\eta \neq 0$, Theorem 2's general closed form equation 8 is proven for arbitrary δ, η and applies directly here (with $\delta = 0$), giving the exact price path $p^\star(\lambda) = (24 - \lambda)/(6 - \lambda)$ at $a = 10, b = 2, d = 1, c = 1$; the (K, L) machinery of Theorem 1 gives the identical result, and Proposition 5 applies regardless of η. Nash $p^{\mathrm{NE}} = 4$, joint-profit-maximizing $p^\dagger = 5.5$ (Bertrand famously under-prices relative to this

benchmark), a well-posedness ceiling $\bar{\lambda}^{\text{seek}} = 4|\beta|/(4|\beta|\delta + \gamma^2) = 8$, and an interior singularity at $\lambda_{\text{pole}} = 6$ (Proposition 5), strictly inside that ceiling.

$\lambda < \lambda_{\text{pole}} = 6$ is necessary for a well-posed, real-valued equilibrium but not sufficient for an economically admissible one: demand $q^\star(\lambda) = a - (b - d)p^\star(\lambda)$ turns negative before $p^\star(\lambda)$ diverges, requiring the strictly lower threshold $\lambda \leq \lambda_{\text{admiss}} = 4$ (where $q^\star = 0$); for $\lambda \in (\lambda_{\text{admiss}}, \lambda_{\text{pole}})$ the price formula above is real-valued but describes negative demand, with no economic meaning. Restricting to $\lambda \leq \lambda_{\text{admiss}}$, which still contains the welfare peak at $\lambda^\star = 2$, risk-seeking price rises monotonically from the Nash price as λ rises, hits $p^\dagger$ exactly at $\lambda^\star = 2$, and net welfare follows the identical hump shape found in Section 6.1, reproduced by the opposite risk attitude, exactly as Theorem 2 predicts for a complements game. As throughout the paper, "welfare" here denotes the sum of the two players' own payoffs; for two firms that sum is joint producer profit, so, unlike in Sections 6.1–6.2, where the two players are not on opposite sides of a market, risk-seeking robustness in this example functions as an uncoordinated substitute for collusion, raising the firms' joint profit at the expense of consumers left out of the game itself. Figure 3 in Section 6.3 plots price and welfare against λ.