The Quiet Hand of Regulation: Harnessing Uncertainty and Disagreement

Daniel Andrei*   Lorenzo Garlappi†

December 1, 2025

Abstract

Traditional Pigouvian and Coasean approaches to regulating externalities falter under uncertainty and disagreement, as they require precise knowledge of externality costs or frictionless bargaining. We develop a system of outcome-contingent "Coasean transfers" that leverage uncertainty and agents' heterogeneous information to achieve efficient outcomes without requiring disclosure of private information. These transfers operate through payments consisting of a quantity component (the gap between an agent's action and the market average) and a price component (tied to publicly observable aggregate outcomes). We prove a market-equivalence result: the optimal transfer pricing schedule corresponds to the equilibrium price in a hypothetical Coasean market for the externality. The equilibrium allocation under Coasean transfers is team efficient and strictly dominates traditional tools like Pigouvian taxes. These transfers are budget-balanced and informationally light, requiring only normative objectives, not private information or signal structures. They incentivize information acquisition, remain robust when agents distrust each other's information, and are politically viable, receiving ex-ante unanimous support.

JEL Classification Codes: D82, H23, Q54, Q58.
Keywords: Uncertainty, Disagreement, Externalities, Regulation, Mechanism Design, Information Aggregation

---

We would like to thank Markus Baldauf, Snehal Banerjee, Jan Bena, Jaroslav Borovička, Walid Busaba, Tolga Cenesizoglu, Vadim Di Pietro, Giada Durante, Paul Ehling, Jack Favilukis, Paolo Fulghieri, Jon A. Garfinkel, Ron Giammarino, Paolo Giordani, Will Gornall, Vincent Grégoire, Peter G. Hansen, Philipp Illeditsch, Ali Lazrak, Derek Lemoine, Fulin Li, Wenhao Li, Dongliang Lu, Arvind Mahajan, Aytek Malkhozov, George Malikov, Iwan Meier, Alessandro Pavan, Neil Pearson, Kevin Pisciotta, Filipp Prokopev, Daniel Schmidt, Petra Sinagl, Sorin Sorescu, Sherry Xue, Maximilian Voigt, Gregory Weitzner, and conference and seminar participants at Keio University, the KU Finance Conference, the BI Norwegian Business School, the University of Iowa Tippie College of Business, UBC, McGill, Ivey Business School, Texas A&M University, HEC Montréal, Concordia University, Michigan State University, the Adam Smith Workshops in Asset Pricing & Corporate Finance, the SFS Cavalcade North America 2025, the EFA Annual Meeting, and the Fall 2025 FTG meeting for their comments.

*McGill University, Desautels Faculty of Management, 1001 Sherbrooke Street West, Office 549, Montréal, Québec H3A 1G5, Canada daniel.andrei@mcgill.ca, danielandrei.info.

†University of British Columbia, Sauder School of Business, 2053 Main Mall, Vancouver V6T 1Z2, British Columbia, Canada, lorenzo.garlappi@sauder.ubc.ca, https://sites.google.com/site/lorenzogarlappi/.

1 Introduction

A central challenge in economics is the design of effective regulation for activities that generate externalities. This task becomes substantially more difficult in the presence of uncertainty and disagreement among economic agents. Traditional approaches, rooted in the work of Pigou (1920) and Coase (1960), often falter in such complex environments. Pigouvian taxes or subsidies, while theoretically appealing, require the regulator to possess precise knowledge of the externality's magnitude, such as the social cost of carbon. Such knowledge is unattainable under pervasive uncertainty and firm-specific private information. Similarly, Coasean bargaining relies on well-defined property rights and negligible transaction costs, conditions that are rarely met in practice, especially when dealing with a large number of agents holding private information.

The literature has responded to this challenge through two main approaches. The first approach, mechanism design, develops schemes to elicit and communicate private information to a central authority (e.g., Mirrlees, 1971; Holmström and Myerson, 1983; Dasgupta, Hammond, and Maskin, 1979). This approach assumes that dispersed information can be credibly transmitted to and processed by a regulator (Laffont and Tirole, 1986; Laffont, 1994). The second approach explores "second-best" policies that rely solely on publicly available information, such as uniform standards or taxes based on aggregate outcomes (e.g., Roberts and Spence, 1976; Weitzman, 1974; Lemoine, 2021). While conceptually sound, this approach still requires the regulator to act as a central calculator, solving a full model of the economy without access to firms' private knowledge. The regulation of externalities is thus confronted with a challenging trade-off: either complex mechanisms with potential fragility or simpler mechanisms with inherent inefficiency.

In this paper, we address this challenge by proposing a novel regulatory approach that *harnesses* uncertainty and disagreement to achieve efficiency *without* requiring the centralized collection of private information. Following Hayek (1945) and Radner (1962), we assume that information is fundamentally dispersed and cannot be communicated to a center. Rather than relying only on public information, our regulatory scheme actively harnesses dispersed private information as the very driver of efficiency. We implement this approach through a system of "Coasean transfers", simple outcome-contingent instruments in the spirit of Pigouvian taxes. We consider a setting where a continuum of agents undertake socially beneficial actions (e.g., investing in green technologies, contributing to public goods, enhancing cybersecurity). The effectiveness of these actions is uncertain, and agents possess heterogeneous private information about the underlying conditions. We show that the Coaesian transfers we propose operate as a "quiet hand" of regulation, in contrast to the invisible hand of laissez-faire, guiding agents toward efficiency without requiring disclosure of private information.

Coasean transfers operate through a simple rule with two components: (i) a *quantity component*, which depends on the gap between an agent's action and the observable average across all agents, and (ii) a *price component*, which is tied to a publicly observable aggregate outcome. This structure induces *strategic substitutability*: because agents are rewarded for their deviation from the average, they have an incentive to counteract the expected average action, which increases their reliance on private information. Formally, the Coasean transfer to agent i takes the form $(a_i - A) \times p(\Omega)$, where a_i is the agent's action, A is the cross-sectional average action, Ω is the observable aggregate outcome, and $p(\Omega)$ is a state-contingent price schedule. Depending on the sign of $(a_i - A)$ and $p(\Omega)$, this transfer can be positive (a subsidy) or negative (a tax). This structure resembles a financial derivative that pays agents based on their deviation from the market average, with payoffs contingent on realized aggregate outcomes. This outcome might include, for example, temperature change or total emissions (in environmental regulation), system stability (in financial or cybersecurity regulation), or aggregate productivity (in regulation that promotes research and knowledge creation, e.g., Romer 1986).

Our analysis builds on the canonical beauty-contest coordination game of Morris and Shin (2002). In this environment, the Coasean resolution of the externality amounts to the creation of a *synthetic competitive market*, the limit of a large Cournot game (Vives, 1988). In such a market, agents would supply their socially beneficial actions, while society's needs would determine the demand and, ultimately, the price. Agents behave as price takers, guided by an endogenous shadow price that reflects the true underlying outcome. The shadow price internalizes the externality and coordinates agents' actions in a way that aggregates private information, capturing the Coasean solution. Our contribution is to show that the Coasean transfer implements precisely this synthetic market.

We derive the optimal pricing function $p(\Omega)$ by requiring that an agent subject to the Coasean transfer faces the same incentives as an agent participating in the Coasean market. This yields a simple, closed-form solution: the optimal transfer price *equals* the social marginal value of the aggregate action. Under a linear-quadratic specification—a standard assumption in the dispersed information literature (e.g., Angeletos and Pavan, 2007, 2009)—this pricing function takes an intuitive form. The price consists of (i) a baseline value, reflecting society's marginal valuation of the aggregate activity, and (ii) an adjustment that decreases when the observed outcome exceeds expectations and increases when it falls short. This adjustment mechanism reflects society's aversion to volatility in aggregate outcomes, creating stronger incentives for action precisely when collective efforts are most needed. Given this equivalence to the Coasean market, the transfer $(a_i - A) \times p(\Omega)$ resembles the payoff of a futures contract.

Our main result is to show that the equilibrium allocation induced by the Coasean transfer is *team-efficient*, the best outcome attainable by a benevolent social planner under decentralized information (Radner, 1962; Angeletos and Pavan, 2007, 2009). This benchmark represents the efficiency frontier under decentralized information: no mechanism can do better without either violating incentive compatibility or requiring agents to reveal their information to a central authority. Formally, from a policy-design perspective, these Coasean transfers belong to the class of outcome- and aggregate-activity-contingent policies studied by Angeletos and Pavan (2009). Unlike Pigouvian taxes or cap-and-trade systems, which require precise ex-ante knowledge and are often suboptimal under uncertainty and disagreement, Coasean transfers reach this frontier without centralized information aggregation, negotiation, or pre-defined property rights, and is ex-post budget-balanced.

The equilibrium allocation induced by Coasean transfers strictly dominates both the unregulated status quo and the allocation based on public information alone (such as a Pigouvian tax). This improvement arises from two sources: *individual flexibility*, as agents tailor actions to their private signals, and *collective adaptability*, as aggregate behavior adjusts to the realized underlying conditions. Beyond static efficiency, Coasean transfers harness informational frictions that traditionally hinder regulation. Specifically, the mechanism exploits strategic substitutability: as uncertainty increases or agents distrust each other's information, they rely more heavily on their private signals, and the transfer's adaptive pricing schedule optimally aggregates this heightened sensitivity, thereby *strengthening* incentives and inducing a stronger collective response.

Crucially, Coasean transfers are not just efficient but also politically feasible. Behind a "veil of ignorance" (Rawls, 1971), all agents would *unanimously* vote for the transfers over any regime based only on public information. The reason is straightforward: Coasean transfers empower agents to act on their private information, while such public-information regimes do not. By the "value of information" principle (Blackwell, 1951, 1953; Marschak and Radner, 1972), having the flexibility to use private information in a concave decision problem can never reduce expected welfare and generically improves it. This unanimous ex-ante support addresses a fundamental challenge in regulatory design: unlike policies that create winners and losers, Coasean transfers make everyone better off *in expectation*. Once agents observe their private signals, however, individual interests and distributional concerns emerge. We characterize conditions under which political support for the transfer persists at this interim stage.

We illustrate the practical relevance of Coasean transfers through two applications that highlight its versatility across positive and negative externalities: firms investing in cybersecurity under uncertainty about aggregate network vulnerability (under-investment), and firms competing in an AI arms race where information acquisition imposes negative spillovers (over-investment). In both cases, the transfer's price function, $p(\Omega)$, is tailored to the sign of the externality: it is positive (a subsidy) to correct under-investment in cybersecurity and negative (a tax) to curb over-investment in AI technology.

Our paper makes three contributions. First, we prove a market-equivalence result: the optimal Coasean transfer pricing schedule equals the equilibrium price in the missing Coasean market. Building on the general existence result of Angeletos and Pavan (2009), this shows how the transfer bridges the Pigouvian goal of internalizing externalities and the Coasean principle of decentralized efficiency. Second, we show that this theoretical ideal is directly implementable, requiring the regulator to know only its normative objectives and how individual actions affect aggregate outcomes, not the distribution of private types, the signal structure, or the fundamental state. Third, we show that Coasean transfers harness uncertainty and even agent distrust, turning these traditional frictions into drivers of efficiency. We also prove that the transfers are politically viable: behind a veil of ignorance, the scheme would receive unanimous support. These findings have broad potential applications, ranging from environmental policy and innovation to financial stability, public health, and cybersecurity.

Literature. Our analysis builds on the global games literature on coordination under heterogeneous private information (Morris and Shin, 2002). We adopt a framework in which agents must anticipate the actions of others based on dispersed information. Unlike models where inefficiency arises solely from information externalities (e.g., Amador and Weill, 2010), our framework features a real payoff externality, which regulation directly addresses. Moreover, by inducing strategic substitutability, the Coasean transfer increases the marginal value of private information. This design feature actively leverages a key insight from the literature: when actions are substitutes, agents rely more on their private signals and less on public information, leading to more effective information aggregation (Hellwig and Veldkamp, 2009; Bergemann and Morris, 2016; Myatt and Wallace, 2012).

Our work is closely related to a literature on policy design for economies with dispersed information. Angeletos and Pavan (2009) show that taxes contingent on realized fundamentals and aggregate activity can restore efficiency, while Colombo, Femminis, and Pavan (2025) find that similar schedules can correct under-investment in information. We extend this literature by developing a solution for more general environments featuring significant agent heterogeneity and complex, unobservable fundamentals, settings that their instruments cannot generally solve. We prove a market-equivalence result that provides the economic intuition for *why* our proposed Coasean transfers, which condition only on the observable aggregate outcome Ω, are effective: they recreate the missing competitive market for the externality. The optimal pricing schedule equals the equilibrium price agents would face in that market, showing how dispersed private information can be harnessed without central disclosure, in the spirit of Hayek (1945).

Recent work on state-contingent policies to address externalities has advanced along two distinct lines: creating new market-based assets and designing corrective payment schedules. The first approach follows the vision of Shiller (1994) by creating new financial securities to manage societal risks. A recent example is Lemoine (2024), who suggests "carbon shares" whose market price internalizes the externality. The second, exemplified by Colombo et al. (2025), shows that a subsidy indexed to aggregate investment and fundamentals can correct for both real spillovers and under-investment in information. Our framework shows that these two paradigms are equivalent: a state-contingent Pigouvian schedule can be interpreted as the creation of a synthetic market for the externality. The Coasean transfer provides a simple regulatory scheme that makes this equivalence operational, implementing the allocation that would prevail if the missing market actually existed. This trend towards outcome-contingent regulation is already reflected in major policy innovations. For instance, the European Union's Market Stability Reserve adjusts the future supply of emission permits based on the current market surplus, creating a dynamic cap that automatically responds to realized outcomes.

Our model's structure, in which agents treat an aggregate variable as a public input, shares features with the endogenous growth framework of Romer (1986). We depart from that perfect-foresight setting by introducing asymmetric information about the state of the economy and designing a regulatory scheme to contend with it. This focus on designing incentive-compatible rules under uncertainty places our work at the intersection of endogenous growth and a classic literature on regulation. That tradition, exemplified by Roberts and Spence (1976), Kwerel (1977), and Montero (2008), has long sought to design mechanisms that function effectively when the regulator has imperfect information about firms. Our specific contribution is to show how regulation can create a missing market that leverages informational frictions as a source of efficiency.

The strategic interdependence in our model also connects it to the literature on heterogeneous beliefs and market dynamics (Keynes, 1964; Bikhchandani, Hirshleifer, and Welch, 1992; Morris and Shin, 2002; Veldkamp, 2011; Myatt and Wallace, 2008). We extend this idea, in the spirit of Banerjee (2011), by allowing agents to distrust one another's signals—a mechanism that introduces a risk component analogous to the "sentiment risk" studied by Dumas, Kurshev, and Uppal (2009). While this literature provides deep insights into market dynamics, our contribution is to use these insights as a foundation for policy design. We ask how a regulator can exploit these very strategic interactions as policy tools to align individual actions with societal goals, even amid agents' distrust.

Finally, our work contributes to the broader literature on uncertainty and economic decision-making. Since the seminal work of Knight (1921), uncertainty has often been viewed as a barrier to effective decisions (e.g., Bernanke, 1983; Rodrik, 1991; Dixit and Pindyck, 1994; Caballero and Pindyck, 1996). We challenge this perspective by demonstrating that, in the presence of informational frictions, uncertainty can be harnessed as a tool for achieving desirable societal goals. Our approach provides a market-based framework for designing robust regulatory policies that leverage, rather than suppress, dispersed information.

The rest of the paper proceeds as follows. Section 2 introduces the model and benchmark allocations, including the Coasean ideal. Section 3 derives the Coasean transfer, proves the market-equivalence result, presents two illustrative applications, and discusses the transfer's informational properties relative to tax and cap-and-trade solutions. Section 4 studies political viability. Section 5 shows how Coasean transfers harness uncertainty and disagreement. Section 6 analyzes incentives for information acquisition. Section 7 concludes.

2 Model

We consider an economy with a continuum of agents, $i \in [0,1]$, each choosing an action $a_i \in \mathbb{R}$. The aggregate action is $A \equiv \int_0^1 a_i di$. Each agent i is characterized by a private type (η_i, β_i), where $\eta_i > 0$ is a cost parameter and $\beta_i > 0$ captures the sensitivity of the agent's private benefit to an underlying fundamental $\theta \in \mathbb{R}$. An exogenous state $\nu \in \mathbb{R}$ also affects the aggregate outcome Ω, which is determined by the linear technology:

$$\Omega \equiv \kappa A + \Gamma(\theta, \nu), \tag{1}$$

where $\kappa > 0$ measures the aggregate contribution to the outcome, and $\Gamma(\theta, \nu)$ is a function driven by the fundamental and the exogenous state.

Agent i's utility consists of a private payoff component and a common social value component, $S(\Omega)$. With an endowment e, we write

$$u_i = \underbrace{\left(e - \frac{\eta_i}{2} a_i^2\right) + \beta_i \theta a_i}_{\text{Private component}} + S(\Omega). \tag{2}$$

The term $\beta_i \theta a_i$ represents the agent's private benefit from the action, which depends on the fundamental θ. The term $\frac{\eta_i}{2} a_i^2$ is the convex private cost. Expected social welfare, $W(\{a_i\})$, is the utilitarian aggregate of agent payoffs:

$$W(\{a_i\}) = \mathbb{E}\left[\int_0^1 \left(e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i\right) di + S(\Omega)\right]. \tag{3}$$

The timing is as follows. Nature draws the state (θ, ν) and the cross-sectional distribution of types $\{\eta_i, \beta_i\}_{i \in [0,1]}$. Agent i observes their own type (η_i, β_i) and forms posterior beliefs, $m_{i,\theta} \equiv \mathbb{E}_i[\theta]$ and $m_{i,\nu} \equiv \mathbb{E}_i[\nu]$, based on their information set, which we denote by $\mathcal{I}_i$, where $\mathbb{E}_i \equiv \mathbb{E}[\cdot | \mathcal{I}_i]$. Agents then choose a_i simultaneously, A and Ω are realized, and payoffs accrue.

Let us start by characterizing the *status-quo* equilibrium. In this benchmark, agents are atomistic and maximize their private payoff, taking the social value $S(\Omega)$ as exogenous to their decision. Agent i's problem is to choose a_i to maximize

$$\max_{a_i \in \mathbb{R}} \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2\right) + \beta_i \theta a_i\right]. \tag{4}$$

The first-order condition (FOC) is $\eta_i a_i = \beta_i \mathbb{E}_i[\theta]$, which yields the optimal status-quo action:

$$a_i^{sq} = \frac{\beta_i}{\eta_i} m_{i,\theta}. \tag{5}$$

Aggregating across agents gives the aggregate action $A^{sq} = \int_0^1 \frac{\beta_i}{\eta_i} m_{i,\theta} di$. We assume that agent types (η_i, β_i) are independent of their posterior beliefs $m_{i,\theta}$. This allows the integral to be factorized:

$$A^{sq} = \left(\int_0^1 \frac{\beta_i}{\eta_i} di\right) \left(\int_0^1 m_{i,\theta} di\right) = \bar{B}_H \bar{m}_\theta, \tag{6}$$

where $\bar{m}_\theta \equiv \int_0^1 m_{i,\theta} di$. We define $H \equiv \left(\int_0^1 \eta_i^{-1} di\right)^{-1}$ as the harmonic mean of the cost parameters and $\bar{B}_H \equiv \int_0^1 \frac{\beta_i}{\eta_i} di$ as the weighted average of benefit sensitivities. Since $\beta_i > 0$ and $\eta_i > 0$ for all i, both H and $\bar{B}_H$ are strictly positive.

The status-quo allocation $\{a_i^{sq}\}$ induces an aggregate outcome $\Omega^{sq} = \kappa A^{sq} + \Gamma(\theta, \nu)$. We define the unconditional expectation of this outcome as our reference benchmark, μ_Ω^{sq}. By the Law of Iterated Expectations, $\mathbb{E}[\bar{m}_\theta] = \mathbb{E}[\theta]$, so

$$\mu_\Omega^{sq} \equiv \mathbb{E}[\Omega^{sq}] = \kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)]. \tag{7}$$

This benchmark depends only on the primitive parameters of the economy.

2.1 First Best

We first establish the socially optimal allocation under complete information. A full-information planner observes the realized state (θ, ν) and the full profile of types $\{\eta_i, \beta_i\}$ and chooses an allocation $\{a_i^{fb}\}$ to maximize social welfare $W(\{a_i\})$ state by state.

Assumption 1 (Social Value)**.** *The social value function $S \in C^2$ is concave: $S''(\cdot) \leq 0$. For results requiring uniqueness, we assume strong concavity: there exists $\underline{\rho} > 0$ such that $-S'' \geq \underline{\rho}$.*

This assumption ensures a well-behaved planner's problem with diminishing marginal social value.

Proposition 1 (First-best)**.** *Under Assumption 1 and $\eta_i > 0$, the unique first-best allocation $\{a_i^{fb}\}$ satisfies, for each (θ, ν):*

$$\eta_i a_i^{fb} = \beta_i \theta + \kappa S'(\Omega^{fb}), \tag{8}$$

where $\Omega^{fb} = \kappa A^{fb} + \Gamma(\theta, \nu)$ and $A^{fb} = \int_0^1 a_i^{fb} di$. Equivalently, the aggregate action A^{fb} is the unique solution to

$$A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H} S'(\kappa A^{fb} + \Gamma(\theta, \nu)). \tag{9}$$

The welfare function $W(\{a_i\})$ is strictly concave in $\{a_i\}$ (given $\eta_i > 0$ and Assumption 1), so the maximizer is unique and characterized by the FOC (8). Aggregating this condition across i yields the scalar fixed-point equation for A^{fb}. Since $\kappa S'(\cdot)$ is nonincreasing ($S'' \leq 0$), the fixed-point map is continuous and nonincreasing, guaranteeing a unique solution.

The first-best allocation is Pareto optimal and serves as our theoretical efficiency benchmark. This allocation is, however, unattainable as it requires the planner to observe agents' private types $\{\eta_i, \beta_i\}$ and the true state (θ, ν).

2.2 Team-Efficient Equilibrium

We next characterize the *team-efficient* equilibrium (Radner, 1962), which serves as the constrained-efficient benchmark under decentralized information. This allocation is the solution to a utilitarian planner's problem of choosing decision rules, $a_i(\mathcal{I}_i)$, for each agent, contingent on their private information $\mathcal{I}_i$, to maximize ex-ante expected welfare $W(\{a_i\})$.

This allocation represents the *constrained-efficient* outcome—the best result possible given that information is privately observed. It serves as a theoretical upper bound for what any *incentive-compatible* mechanism can achieve, a point we return to in Section 3.2.5.

The planner's optimization must respect the information constraints: the rule for agent i can only depend on $\mathcal{I}_i$, and the planner must account for the externality imposed by each a_i on the aggregate. The resulting FOC for each agent i is:

$$\eta_i a_i^{te} = \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\Omega^{te})], \qquad \Omega^{te} = \kappa A^{te} + \Gamma(\theta, \nu). \tag{10}$$

This condition defines the team-efficient allocation. It differs from the status-quo condition (5) by the addition of the second term, $\kappa \mathbb{E}_i[S'(\Omega^{te})]$. The first term, $\beta_i m_{i,\theta}$, ensures the rule is implementable, as the agent's action must be contingent on their private belief $m_{i,\theta}$. The second term reflects the utilitarian planner's objective, forcing the agent to internalize the expected marginal social value of their action, conditional on their information.

Proposition 2 (Existence and Uniqueness of the Team-efficient Equilibrium)**.** *Under Assumption 1 and strictly convex individual costs ($\eta_i > 0$ for all i), there exists a unique team-efficient equilibrium $\{a_i^{te}\}_{i \in [0,1]}$. This equilibrium is characterized by the unique aggregate action A^{te} that satisfies the fixed-point condition $A^{te} = \Phi(A^{te})$, where $\Phi(A)$ is the continuous* *and nonincreasing aggregate mapping:*

$$\Phi(A) = \int_0^1 \frac{\beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\kappa A + \Gamma(\theta, \nu))]}{\eta_i} di. \tag{11}$$

The welfare function $W(\{a_i\})$ is strictly concave in the full allocation $\{a_i\}$ because individual cost functions are strictly convex ($\eta_i > 0$) and $S(\cdot)$ is concave (Assumption 1). This ensures a unique maximizer $\{a_i^{te}\}$. The FOC (10) defines the individual action $a_i(A)$ as a nonincreasing function of a posited aggregate action A, since $\kappa S'(\cdot)$ is nonincreasing ($S'' \leq 0$). The aggregate mapping $\Phi(A) = \int a_i(A) di$ inherits this property. As $\Phi(A)$ is a continuous, nonincreasing function, standard fixed-point arguments guarantee that a unique A^{te} exists.

2.3 Strategic Characterization

To gain further insight into the structure of the team-efficient equilibrium, we now impose a parametric assumption that allows us to express agents' equilibrium actions as best responses, yielding a closed-form solution. We refer to this as the *strategic characterization* of the equilibrium.

Assumption 2 (Linear-Quadratic Social Value)**.** *The marginal social value function, $S'(\Omega)$, is linear in the aggregate outcome Ω, implying a constant second derivative $S''(\Omega) = -s_2$ for some $s_2 > 0$. We normalize this function by centering it around the expected status-quo outcome, μ_Ω^{sq}:*

$$S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq}). \tag{12}$$

This linear-quadratic specification is standard in the dispersed information literature (e.g., Angeletos and Pavan, 2007, 2009) and ensures that the equilibrium decision rules are linear.

2.3.1 First Best

Under Assumption 2, the first-best allocation derived in Proposition 1 simplifies to a closed-form linear function of the fundamentals.

Proposition 3 (First Best under Assumption 2)**.** *The unique first-best allocation is:*

$$a_i^{fb} = \frac{\beta_i \theta + \kappa[s_1 - s_2(\Omega^{fb} - \mu_\Omega^{sq})]}{\eta_i}, \tag{13}$$

where the state-contingent outcome Ω^{fb} *is*

$$\Omega^{fb} = \frac{\Gamma(\theta, \nu) + \kappa \bar{B}_H \theta + (\kappa^2/H)s_1 + (\kappa^2 s_2/H)\mu_\Omega^{sq}}{1 + \kappa^2 s_2/H}. \tag{14}$$

This closed-form expression characterizes the optimal allocation as a linear function of the fundamentals. In the limit of no externality ($\kappa \to 0$), the solution collapses to the autarky benchmark, where the social optimum coincides with private incentives.

2.3.2 Team-Efficient Equilibrium

We now characterize the team-efficient allocation under the linear-quadratic assumption. The equilibrium exhibits a strategic structure that decomposes the agent's decision rule and reveals the precise role of heterogeneity.

Proposition 4 (Team-efficient strategic form under Assumption 2)**.** *Let* a_i^{fb} *be the first-best action. Define the degree of coordination as* $\alpha \equiv -\kappa^2 s_2$*. The team-efficient action* a_i^{te} *satisfies the strategic form:*

$$a_i^{te} = \underbrace{\left(1 - \frac{\alpha}{H}\right) \mathbb{E}_i[a_i^{fb}] + \frac{\alpha}{\eta_i} \mathbb{E}_i[A^{te}]}_{\text{Heterogeneous Strategic Form}} \underbrace{- \frac{\alpha}{\eta_i} \left(\bar{B}_H - \frac{\beta_i}{H}\right) m_{i,\theta}}_{\text{Type-Correction Term}}. \tag{15}$$

The strategic form (15) provides a clear decomposition of incentives and highlights two key departures from the standard Angeletos and Pavan (2007) framework. First, the strategic motive itself is heterogeneous. The agent's response to the aggregate expectation, $\frac{\alpha}{\eta_i}\mathbb{E}_i[A^{te}]$, is scaled by their private cost η_i. Agents with higher costs (higher η_i) are less responsive to the strategic motive, a feature absent in the homogeneous-response benchmark.

Second, and central to our analysis, is the new *Type-Correction Term*. This term is zero only if an agent's type happens to match the population average ($\beta_i/H = \bar{B}_H$). For an agent with an above-average private benefit ($\beta_i/H > \bar{B}_H$), the correction term becomes negative (since $\alpha < 0$). This "curbs the enthusiasm" implied by their idiosyncratic type, pulling their action from their private first-best back toward the team optimum.

2.4 The Coasean Ideal: A Synthetic Market

Before we introduce our mechanism, let us first conduct a thought experiment. Consider an idealized, complete market where the externality is perfectly priced. In this synthetic market, all uncertainty is spanned. We define a *state of the world*, $z = (\theta, \nu)$, as the realization of the underlying fundamentals. The good traded in this market is a state-contingent claim: a contract for *one unit of action*, deliverable if and only if state z occurs.

In this market, a competitive price schedule $\hat{p}(z)$ emerges. The demand side establishes this price as the social marginal value of the action in that state: $\hat{p}(z) = \kappa S'(\Omega(A, z))$, where $\Omega(A, z) = \kappa A + \Gamma(\theta, \nu)$. On the supply side, atomless agents act as price-takers. Each agent i chooses a_i to maximize their expected payoff, which now includes the market revenue $\hat{p}(z)a_i$:

$$\max_{a_i} \ \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) + \hat{p}(z) a_i \right]. \tag{16}$$

As atomless agents, they take $S(\Omega)$ as given ($\partial S(\Omega)/\partial a_i = 0$). The agent's FOC is therefore

$$\mathbb{E}_i[-\eta_i a_i + \beta_i \theta + \hat{p}(z)] = 0 \quad \Rightarrow \quad \eta_i a_i = \beta_i \mathbb{E}_i[\theta] + \mathbb{E}_i[\hat{p}(z)]. \tag{17}$$

In equilibrium, the price $\hat{p}(z)$ must equal the social marginal value $\kappa S'(\Omega)$. Substituting this condition, $\mathbb{E}_i[\hat{p}(z)] = \kappa \mathbb{E}_i[S'(\Omega)]$, into the agent's FOC (17) yields:

$$\eta_i a_i = \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\Omega^{mkt})], \quad \text{where } \Omega^{mkt} = \kappa A^{mkt} + \Gamma(\theta, \nu). \tag{18}$$

This allocation rule is identical to the team-efficient rule (10). The synthetic market thus perfectly implements the constrained-efficient benchmark—a form of the First Welfare Theorem under private information. This represents the *Coasean ideal*.

Such a market is, however, moot in practice. It would require a complete set of state-contingent contracts for every possible realization of $z = (\theta, \nu)$, which is infeasible due to prohibitive transaction costs and complexity. While this idealized market is infeasible, its equilibrium provides the market-based target for the mechanism we develop next.

3 The Coasean Transfer

We now show how to replicate the synthetic market's equilibrium—and thus implement the team-efficient allocation—using a simple transfer scheme based only on observables, which we call the *Coasean transfer*.

We design a transfer T_i for agent i that is a function of their deviation from the aggregate action, $(a_i - A)$, priced by a function $p(\Omega)$ that depends only on the observable outcome. The transfer is $T_i = (a_i - A)p(\Omega)$. An agent i under this scheme solves:

$$\max_{a_i} \ \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) + (a_i - A)p(\Omega) \right]. \tag{19}$$

As the agent is atomless, they take A and Ω as given in their optimization. The FOC for the agent under this scheme is:

$$\mathbb{E}_i[-\eta_i a_i + \beta_i \theta + p(\Omega)] = 0 \quad \Rightarrow \quad \eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega)]. \tag{20}$$

We now find the optimal pricing function $p(\Omega)$. Recall from Section 2.4 that the FOC for the idealized Coasean market is $\eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[\hat{p}(z)]$, with $\hat{p}(z) = \kappa S'(\Omega)$. The target allocation—the Coasean ideal—is therefore characterized by the rule:

$$\eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[\kappa S'(\Omega)]. \tag{21}$$

For the transfer mechanism (20) to induce the same behavior as the Coasean market (21), the conditions must be identical for any agent belief. This requires $p(\Omega) = \kappa S'(\Omega)$ for all states. The optimal price is the social marginal value.

Proposition 5 (Coasean Transfer Implementation)**.** *Suppose Assumption 1 holds and agents are atomless. If the regulator sets the transfer $T_i = (a_i - A)p(\Omega)$ with $p(\Omega) = \kappa S'(\Omega)$, then the induced equilibrium allocation coincides with the team-efficient allocation.*

The transfer $T_i = (a_i - A)p(\Omega)$ operates like a futures contract, where agent i takes a net position $(a_i - A)$ that settles ex-post at the realized market price $p(\Omega)$. This mechanism has several useful properties: it uses only observables—individual actions a_i, the aggregate A, and the realized outcome Ω; it requires no reports of private beliefs or states; finally, it is exactly budget-balanced ex-post: $\int_0^1 T_i di = \int_0^1 (a_i - A)p(\Omega)di = p(\Omega)\int_0^1 (a_i - A)di = 0$. The Coasean transfer thus replicates the Arrow-Debreu market with a single, outcome-indexed price.

The Coasean transfer aligns private incentives with the Coasean market by forcing each agent to face the social marginal value of their action. Because the Coasean market achieves team efficiency, the Coasean transfer implements the team-efficient allocation. In this environment, there is no efficiency loss from private information: the incentive-compatible outcome coincides with the team-efficient benchmark.[1] This result holds because agents are atomless, the externality operates through the aggregate outcome, and the observable Ω is a sufficient statistic for the unobservable states (θ, ν) needed to price the externality.[2]

Interpretation under Linear-Quadratic Preferences. Under Assumption 2, the optimal pricing function $p(\Omega) = \kappa S'(\Omega)$ becomes transparent. Using $S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq})$, we can define two economic parameters. Let $\text{ORA} \equiv s_2 = -S''(\Omega)$ be the social aversion to outcome fluctuations, or *Outcome Risk Aversion*. Let $\text{MSV}^{sq} \equiv s_1 = S'(\mu_\Omega^{sq})$ be the *Marginal Social Value* at the status-quo benchmark. The Coasean transfer price is:

$$p(\Omega) = \kappa \left[\text{MSV}^{sq} - \text{ORA}(\Omega - \mu_\Omega^{sq})\right]. \tag{22}$$

The optimal price is a baseline value, κMSV^{sq}, adjusted for deviations from the benchmark. The adjustment, $-\kappa\text{ORA}(\Omega - \mu_\Omega^{sq})$, reflects society's aversion to outcome volatility.

3.1 Illustrative Examples

We illustrate the versatility of the Coasean transfer through two complementary examples. First, we analyze cybersecurity, a classic case of a positive externality leading to under-investment. Second, we examine a "business-stealing" AI arms race, where a negative externality drives over-investment. Together, these examples show how the transfer endoge-

[1] An outcome is *incentive-compatible* if, under the mechanism's transfers, each agent's prescribed action is exactly the action that maximizes their own expected payoff given their information. In this mechanism, $T_i = (a_i - A)p(\Omega)$, and the FOC (20) with $p(\Omega) = \kappa S'(\Omega)$ reproduces the team-efficient rule (10). The allocation is thus incentive-compatible.

[2] *Sufficient* means the hidden states (θ, ν) matter for the mechanism only through the realized outcome Ω. The social marginal value $p(\Omega(A, z)) = \kappa S'(\Omega(A, z))$ depends on z only via Ω. Conditioning transfers on Ω therefore captures all payoff-relevant information.

nously creates the correct positive price (a subsidy) or negative price (a tax) to align private incentives with social efficiency.

3.1.1 Cybersecurity Investment

We analyze a stylized model of cybersecurity investment. This is a classic case of a positive externality: an agent's defensive efforts improve their own security but also contribute to network-wide security, a benefit they do not privately capture. The setting illustrates how the Coasean transfer solves the resulting under-investment and coordination failure.

Consider an economy where the agent's action a_i is their investment in cybersecurity. To map this to our general model, we assume homogeneity by setting $\eta_i = 1$ and $\beta_i = 1$ for all agents. The agent's private payoff from equation (2) is thus $(e - \frac{1}{2}a_i^2) + \theta a_i$, where θ represents the baseline intensity or private value of fending off cyber threats.

The key externality is that aggregate investment, $A = \int_0^1 a_i di$, creates the public good of "network security." We model this with the simplest possible technology: $\kappa = 1$ and $\Gamma(\theta, \nu) = 0$, so the aggregate outcome is $\Omega = A$.

Society values this network security. We specify a social value function $S(\Omega)$ that matches our linear-quadratic assumption (Assumption 2):

$$S(\Omega) = s_1\Omega - \frac{s_2}{2}(\Omega - \mu_\Omega^{sq})^2, \tag{23}$$

where $s_1 > 0$ is the baseline social value of security and $s_2 > 0$ captures social aversion to outcome volatility (risk). This implies $S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq})$.

In the status quo, each agent i maximizes their private payoff. From equation (5), the optimal action is:

$$a_i^{sq} = \frac{\beta_i}{\eta_i} m_{i,\theta} = m_{i,\theta}. \tag{24}$$

This action is independent of the aggregate A. Agents only respond to their private assessment of the threat θ and completely ignore both the positive externality s_1 and the social risk s_2.

The team-efficient allocation, by contrast, requires agents to coordinate. The optimal degree of coordination, defined as $\alpha \equiv -\kappa^2 s_2$, is:

$$\alpha = -1^2 \cdot s_2 = -s_2. \tag{25}$$

Strategic substitutability ($\alpha < 0$) is optimal here because society is averse to outcome volatility ($s_2 > 0$). If all agents use the same information, their actions are highly correlated, leading to a volatile aggregate outcome Ω. Strategic substitutability provides an incentive for agents to use their idiosyncratic private information, which reduces this costly correlation. The unregulated market provides no such incentive, while the optimum requires it.

The Coasean transfer restores the missing strategic link. The optimal pricing function from equation (22) is:

$$p(\Omega) = \kappa\left[s_1 - s_2(\Omega - \mu_\Omega^{sq})\right] = s_1 - s_2(\Omega - \mu_\Omega^{sq}). \tag{26}$$

The agent's FOC under the Coasean transfer, from equation (20), is $\eta_i a_i^{ct} = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega^{ct})]$. Substituting the parameters for this example yields:

$$1 \cdot a_i^{ct} = 1 \cdot m_{i,\theta} + \mathbb{E}_i\left[s_1 - s_2(\Omega^{ct} - \mu_\Omega^{sq})\right]. \tag{27}$$

Since $\Omega^{ct} = A^{ct}$, this simplifies to $a_i^{ct} = m_{i,\theta} + s_1 - s_2(\mathbb{E}_i[A^{ct}] - \mu_\Omega^{sq})$. The resulting equilibrium has a strategic interaction of $-s_2$, which matches the team-efficient α and restores efficiency.

This application offers direct policy insights. If society is risk-neutral ($s_2 = 0$), the optimal degree of coordination is zero ($\alpha = 0$), and the Coasean transfer acts as a simple Pigouvian subsidy ($p(\Omega) = s_1$). If society is risk-averse ($s_2 > 0$), the optimal interaction is strategic substitutability ($\alpha = -s_2 < 0$). The Coasean transfer correctly implements this by inducing strategic substitutability to reduce aggregate volatility. This distinction echoes real-world debates—for instance, between imposing uniform cybersecurity standards, which can heighten correlated risks (Kunreuther and Heal, 2003), and fostering a *diversity of defense* to build resilience (Jajodia, Ghosh, Swarup, Wang, and Wang, 2011).

3.1.2 AI Arms Race and Business Stealing

We now model an AI arms race—a "business-stealing" externality where firms invest in a socially wasteful activity to capture market share from rivals. The action a_i is an investment in data and AI. The private payoff is $(e - \frac{\eta_i}{2}a_i^2) + \beta_i\theta a_i$, where $\beta_i\theta a_i$ represents the private benefit from business stealing and θ is the total stealable profit pool.

This aggregate investment, $A = \int a_i di$, is a zero-sum game representing a socially wasteful arms race. It thus creates a negative externality, which we model as $S'(\Omega) < 0$. For simplicity, we set the outcome as $\Omega = \kappa A$.

In the status quo, agents optimize their private payoff, $\eta_i a_i^{sq} = \beta_i \mathbb{E}_i[\theta]$. Because the social cost $\kappa S'(\Omega) < 0$ is ignored, this leads to over-investment: $A^{sq} > A^{te}$.

The team-efficient allocation requires agents to internalize this social cost. The FOC is:

$$\eta_i a_i^{te} = \beta_i \mathbb{E}_i[\theta] + \mathbb{E}_i[p(\Omega^{te})]. \tag{28}$$

The optimal price for the Coasean transfer is $p(\Omega) = \kappa S'(\Omega)$. Because $S' < 0$, the optimal price is negative, $p(\Omega) < 0$. By forcing agents to internalize this negative expected price $\mathbb{E}_i[p(\Omega)] < 0$, the mechanism lowers the total marginal return on investment and solves the over-investment problem.

The Coasean transfer $T_i = (a_i - A)p(\Omega)$ acts as an arms-control tax. An agent who escalates $(a_i > A)$ pays a tax, as $(a_i - A) > 0$ and $p(\Omega) < 0$ implies $T_i < 0$. Conversely, an agent who disarms $(a_i < A)$ receives a subsidy, as $(a_i - A) < 0$ and $p(\Omega) < 0$ implies $T_i > 0$. This example complements the cybersecurity application, showing the Coasean transfer solves over-investment (via a negative price) just as it solves under-investment (via a positive price).

3.2 Discussion

We now discuss the properties of the Coasean transfer. We first show that its implementation is "light on information" for the regulator, whereas a uniform Pigouvian policy carries a significant informational burden. This burden, we argue, forces the Pigouvian regulator into the classic Weitzman (1974) prices-versus-quantities dilemma, a second-best choice that the Coasean transfer bypasses. We conclude by placing our results in a mechanism-design context, showing that the Coasean transfer implements the informational second-best.

3.2.1 A Separation Result

It is immediate from Proposition 5 that the regulator does not need the information that agents possess. The optimal Coasean transfer, $T_i = (a_i - A)p(\Omega)$, requires knowledge of only the function $p(\Omega) = \kappa S'(\Omega)$. It does not require the regulator to know the distribution of states (θ, ν), the agents' signal structure, or the cross-sectional distribution of private types (η_i, β_i).

Coasean transfers thus create a clear separation of information. The regulator need only specify the price function $p(\Omega)$, which relies solely on social preferences $S(\cdot)$ and the technology κ. All other information is private to the agents. The mechanism remains anonymous, budget-balanced, and requires no reports. Agents' private information is aggregated endogenously into the outcome Ω, and by pricing this single observable, the regulator implements the team-efficient allocation.

This approach solves the "span problem" common to policy mechanisms that are linear in fundamentals (e.g., Angeletos and Pavan, 2009). Such instruments are ineffective when shocks—like ν or nonlinear functions of θ—lie outside the span of the chosen instruments. The Coasean transfer bypasses this problem. By conditioning directly on the outcome Ω, it uses the sufficient statistic that already embeds all payoff-relevant uncertainty.

This addresses the classic observation, going back to Radner (1982), that complete Arrow–Debreu markets are infeasible when the underlying microstate space is too fine to be represented in a feasible system of contracts.[3] The Coasean transfer replaces this space with a single, outcome-indexed price $p(\Omega)$, achieving the Coasean ideal with a feasible instrument.

3.2.2 Optimal Price Instrument (Pigouvian Tax/Subsidy)

We now contrast the Coasean transfer mechanism with the standard Pigouvian approach. We consider a regulator constrained to a single, uniform tax or subsidy τ on the agent's action a_i, set ex-ante based on public information. The regulator's problem is to choose the welfare-maximizing τ that internalizes the expected marginal externality.

[3] See Radner (1982, Section 1.6), who argues that the Arrow–Debreu model requires "a complete system of insurance and future markets" that is "too complex, detailed, and refined to have any practical significance."

Given a uniform tax τ on a_i and a lump-sum rebate R of aggregate tax revenues, each agent i takes τ and the aggregates as given and chooses a_i to maximize expected utility

$$\max_{a_i} \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) - \tau a_i + R \right]. \tag{29}$$

The agent's best response to a given τ is $a_i^{tax} = (\beta_i m_{i,\theta} - \tau)/\eta_i$. The regulator then chooses τ to maximize ex-ante welfare $\mathbb{E}[W]$, knowing agents will follow this rule. The resulting FOC yields the intuitive Pigouvian condition that the optimal tax must equal the negative of the expected marginal social value, $\tau^* = -\kappa \mathbb{E}[S'(\Omega(\tau^*))]$. The following proposition provides the closed-form solution to this fixed-point problem.

Proposition 6 (Optimal Pigouvian Tax/Subsidy)**.** *Consider a uniform tax or subsidy τ. Under Assumptions 1 and 2, and independence of types and signals, the unique welfare-maximizing tax or subsidy τ^* is given by:*

$$\tau^* = \frac{-\kappa s_1 + \kappa s_2 \left(\kappa \bar{B}_H \, \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq} \right)}{1 + \kappa^2 s_2 / H}. \tag{30}$$

This τ^* satisfies the classic Pigouvian condition $\tau^* = -\kappa \mathbb{E}[S'(\Omega(\tau^*))]$, implying the tax is a subsidy if the expected marginal social value is positive.

The crucial insight lies in what is required to compute this τ^*. To set the optimal tax/subsidy, the regulator must know: all primitives (s_1, s_2, κ); the public priors $(\mathbb{E}[\theta], \mathbb{E}[\Gamma(\theta, \nu)], \mu_\Omega^{sq})$; and, most important, the regulator must have structural knowledge of the economy's cross-section, captured by the cost and benefit coefficients $(H, \bar{B}_H)$. This is the informational burden of the Pigouvian approach. It stands in contrast to the Coasean transfer, which, as discussed in Section 3.2.1, dispenses with all priors and all cross-sectional aggregates, requiring only the function $\kappa S'(\cdot)$ and the realized outcome Ω.

3.2.3 Optimal Quantity Instrument (Cap-and-Trade)

We now derive the optimal *quantity*-based policy instrument, which serves as the natural counterpart to the Pigouvian tax analyzed in Section 3.2.2. We consider a regulator constrained to a single, aggregate cap A^{cap} on the total action A, set ex-ante based on public information. The regulator's problem is to choose the welfare-maximizing cap.

Given a cap A^{cap} and a competitive permit price P, with permit revenues rebated lump-sum as $R = PA$, each agent i takes P, R, and the aggregates as given and chooses a_i to maximize expected utility

$$\max_{a_i} \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) - P a_i + R \right]. \tag{31}$$

The agent's best response to a given price P is $a_i^{cap}(P) = (\beta_i m_{i,\theta} - P)/\eta_i$. For a given cap A^{cap}, the competitive permit price $P(A^{cap})$ clears the market by equating aggregate demand for permits to the fixed supply, so that the resulting allocation of individual actions aggregates to A^{cap}. The regulator then chooses A^{cap} to maximize ex-ante welfare $\mathbb{E}[W(A^{cap})]$, knowing that the permit market implements this allocation. The resulting FOC has the standard form that expected marginal social benefit equals expected marginal social cost of relaxing the cap. The following proposition provides the closed-form solution to this problem.

Proposition 7 (Optimal cap-and-trade quantity)**.** *Under Assumptions 1 and 2, and independence of types and signals, the unique welfare-maximizing cap is*

$$A^{cap*} = \frac{H \bar{B}_H \mathbb{E}[\theta] + \kappa s_1 - \kappa s_2 \Big(\mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq} \Big)}{H + \kappa^2 s_2}. \tag{32}$$

The optimal cap satisfies the standard condition that expected marginal social benefit equals expected marginal social cost.

This expression mirrors the Pigouvian solution in Proposition 6. In fact, the two instruments are duals. The equilibrium permit price $P(A)$ is determined by the aggregate marginal private value of the action, $P(A) = H(\bar{B}_H \mathbb{E}[\theta] - A)$. Substituting the optimal cap A^{cap*} into this expression yields a permit price exactly equal to the optimal tax τ^*. Thus, the regulator achieves the same outcome regardless of the instrument. In both cases, however, the regulator must know the structural parameters $(H, \bar{B}_H)$ and the public priors $(\mathbb{E}[\theta], \mathbb{E}[\Gamma(\theta, \nu)], \mu_\Omega^{sq})$. Neither the optimal tax nor the optimal cap adapts to the realized state (θ, ν), and both are therefore second-best under uncertainty. Coasean transfers, by contrast, adjusts the price to the realized outcome Ω and bypass this informational limitation.

3.2.4 The Weitzman (1974) Tradeoff

The duality between taxes and caps holds *only* if the regulator possesses the structural knowledge assumed in Propositions 6 and 7. When the regulator is uncertain about the economy's cross-section—specifically the aggregate cost slope H or the benefit parameter $\bar{B}_H$—this equivalence vanishes. Such uncertainty forces a choice between fixing the price (leaving the quantity volatile) or fixing the quantity (leaving the marginal cost volatile). This dilemma is the essence of the Weitzman (1974) "prices-versus-quantities" tradeoff. The regulator must select the instrument that is most robust to error, a second-best choice determined by the relative slopes of the marginal benefit and marginal cost curves.

To frame this tradeoff, we define the marginal social cost (MSC) and marginal social benefit (MSB) schedules, which are derived from the aggregate components of the agent's payoff in equation (2). The MSC schedule, $\text{MSC}(A) = HA$, captures the aggregate private cost of action (the $\frac{\eta_i}{2}a_i^2$ term) and has a slope of $\text{MSC}'(A) = H$.[4] The MSB schedule captures the aggregate private and social benefits (the $\beta_i\theta a_i$ and $S(\Omega)$ terms). Under Assumption 2, the MSB curve has a constant slope, $\text{MSB}'(A) = -\kappa^2 s_2$. This slope $\text{MSB}'(A)$ is precisely the degree of coordination α defined in Proposition 4. The regulator's problem is that they are uncertain about the true H (the cost slope) and about the true state $(\theta, \Gamma, \bar{B}_H)$ (which shifts the benefit intercept). This leads to the classic result.

Proposition 8 (Weitzman prices-versus-quantities)**.** *In a neighborhood of the efficient quantity, the cost side has slope* $\text{MSC}'(A) = H$ *and the benefit side has slope* $\text{MSB}'(A) = -\kappa^2 s_2$*. The locally welfare-superior instrument is*

$$\textit{Choose} \begin{cases} \textit{price (tax/subsidy)} & \textit{if } \mathbb{E}[H] > \kappa^2 s_2, \\ \textit{quantity (cap)} & \textit{if } \mathbb{E}[H] < \kappa^2 s_2. \end{cases} \tag{33}$$

Thus, if $\mathbb{E}[H] > \kappa^2 s_2$, the regulator minimizes expected deadweight loss by setting a price (tax/subsidy). Conversely, if $\mathbb{E}[H] < \kappa^2 s_2$, the regulator prefers setting a quantity (cap). This condition recovers the fundamental result in Weitzman (1974) (equation 20). In his notation, prices are preferred if $C'' + B'' > 0$, where B'' and C'' are the second derivatives of

[4] Details of these derivations are in Appendix A.8.

the total benefit and cost functions, that is, the slopes of the marginal benefit and marginal cost curves. Identifying his cost slope C'' with our expected cost slope $\mathbb{E}[H]$ and his benefit slope B'' with $-\kappa^2 s_2$ yields precisely the condition $\mathbb{E}[H] > \kappa^2 s_2$.

The tradeoff's logic is about "choosing the safest way to be wrong." A steep cost slope ($\mathbb{E}[H]$ is large) makes quantity mistakes highly costly, as a small over-shoot causes a significant increase in marginal costs; a price is safer as it caps the marginal cost. Conversely, a steep benefit slope ($\kappa^2 s_2$ is large) makes price mistakes similarly costly, as a small under-shoot causes a significant loss of social benefit; a quantity is safer as it locks in the high-value actions.

This contrast highlights the informational advantage of the Coasean transfer. The Weitzman tradeoff is a consequence of using an information-blind instrument—a single price τ^* or quantity A^{cap*}—set ex-ante. The Coasean transfer is not a single price, but a state-contingent *price schedule*, $p(\Omega)$, which bypasses the dilemma. It does not require the regulator to estimate the uncertain parameters $(H, \bar{B}_H, \theta, \Gamma(\theta, \nu))$. Instead, it uses the realized outcome Ω as a sufficient statistic to compute the correct marginal price ex-post. The Weitzman tradeoff is thus the second-best dilemma the regulator must solve only in the absence of a mechanism, like the Coasean transfer, that can condition on observables to implement the efficient allocation.

Uncertainty breaks the equivalence between Pigou and Coase: it forces the Pigouvian regulator into the second-best Weitzman dilemma and renders the complete Coasean market infeasible. The Coasean transfer restores this link. The function $p(\Omega) = \kappa S'(\Omega)$ acts as a *state-contingent Pigouvian tax* that adjusts automatically to the realized state Ω. It thus achieves the efficient allocation of the ideal Coasean market, but through a feasible, Pigouvian-style instrument.

3.2.5 Mechanism-Design Interpretation

We conclude this discussion by placing our result in a mechanism-design context. The team-efficient allocation is the *informational second-best*—it maximizes expected welfare given agents' private information, but ignoring incentive constraints. We show that our mechanism implements this benchmark exactly. To formalize this, we consider a Bayesian game where a strategy $a_i(\mathcal{I}_i)$ maps the information set $\mathcal{I}_i$ to an action, and a Bayesian Nash equilibrium is a strategy profile in which every agent maximizes their expected utility given the aggregate behavior of others. Because agents are atomless and the objective is strictly concave, the team-efficient rule constitutes the unique best response for every agent, thereby forming a Bayesian Nash equilibrium.

Proposition 9 (The Coasean transfer implements the informational second-best)**.** *Consider the mechanism in which agents choose actions a_i and transfers are*

$$T_i = (a_i - A)\, p(\Omega), \qquad p(\Omega) = \kappa S'(\Omega), \qquad \Omega = \kappa A + \Gamma(\theta, \nu). \tag{34}$$

With atomless agents and an observable Ω, this mechanism:

1. *implements the team-efficient allocation in a Bayesian Nash equilibrium;*
2. *is exactly budget-balanced ex post;*
3. *is ex-ante individually rational whenever the team-efficient allocation is ex-ante individually rational relative to the unregulated status quo.*

The implementation is straightforward. An agent i under the Coasean transfer solves the problem described in equation (19). As the agent is atomless, they take A and Ω as given. Their FOC is

$$-\eta_i a_i + \beta_i \mathbb{E}_i[\theta] + \mathbb{E}_i[p(\Omega)] = 0. \tag{35}$$

By setting $p(\Omega) = \kappa S'(\Omega)$, this FOC becomes identical to the team-efficient FOC (10). Strict concavity ensures this solution is unique; thus, the mechanism implements the team-efficient allocation as the unique Bayesian Nash equilibrium. Exact budget balance, $\int_0^1 T_i\, di = 0$, holds ex post by construction. From an ex-ante perspective, expected transfers are zero; thus, the mechanism satisfies the participation constraint whenever the efficient allocation dominates the unregulated status quo.

In standard mechanism design, a gap often exists between the informational second-best and the best incentive-compatible (IC) allocation. Our result shows this IC gap closes entirely.

This is possible because agents are atomless and the observable outcome Ω is a sufficient statistic for the unobservables. A single, outcome-indexed price is all that is needed to align private and social margins. The Coasean transfer thus attains the highest welfare possible under private information, while remaining anonymous, requiring no reports, and balancing the budget.

Corollary 9.1. *The Coasean transfer attains the informational second-best exactly:*

$$\textit{First-best} > \textit{Team-efficient} = \textit{Best IC allocation}. \tag{36}$$

The usual incentive-compatibility gap vanishes.

Our result is an existence and implementation statement under the maintained assumptions (a continuum of agents, a verifiable outcome Ω, and an externality that operates through Ω). We do not take a position on performance outside this baseline. Analyzing departures such as finitely many agents or non-contractible outcomes is beyond the scope of this paper.

4 Political Viability

This section analyzes the political viability of the Coasean transfer. We first examine the ex-ante case, where agents choose a policy from behind a Rawls (1971) "veil of ignorance," before their private information is realized. We show that the Coasean transfer receives unanimous support, a result that connects to the classical value-of-information principle (Blackwell, 1951; Marschak and Radner, 1972). We then analyze the more substantive interim problem, which accounts for the distributional concerns that emerge after agents observe their private signals.

4.1 Ex-Ante Unanimity

We first study the choice of regime from an ex-ante perspective, assuming society decides to regulate to internalize the externality. Let the *internalized* utility of agent i be denoted by $\tilde{u}_i(a, z)$. This is precisely the objective that a social planner would use when choosing a for agent i, holding $A_{-i}(z)$ fixed—i.e., this objective internalizes the social value term $S(\cdot)$ at the individual level:

$$\tilde{u}_i(a, z) = \left(e - \frac{\eta_i}{2}a^2\right) + \beta_i\theta a + S\Big(\kappa[a + A_{-i}(z)] + \Gamma(z)\Big). \tag{37}$$

Given $\eta_i > 0$ and Assumption 1, $\tilde{u}_i$ is strictly concave in a_i. We compare two mechanisms that agents vote on a priori: a *public-information regime* (such as one based on a Pigouvian tax) in which the agent i's action is a_i^{pub} is restricted to depend only on public information $\mathcal{I}^{pub}$, and the *Coasean transfer* regime, which is based on the agent's finer information set $\mathcal{I}_i$ and, as shown in Proposition 5, implements the team-efficient allocation a_i^{te}.

Proposition 10 (Ex-Ante Unanimity)**.** *For any agent i, the Coasean transfer regime is (weakly) preferred to the public-information regime:*

$$\mathbb{E}\left[\tilde{u}_i(a_i^{te}, z)\right] \geq \mathbb{E}\left[\tilde{u}_i(a_i^{pub}, z)\right]. \tag{38}$$

The inequality is strict if the agent's private information $\mathcal{I}_i$ is genuinely informative for the agent's decision.

The intuition rests on the classic value of information principle (Blackwell, 1951; Marschak and Radner, 1972). Because the Coasean transfer aligns the agent's private incentives with the internalized utility $\tilde{u}_i$, the agent effectively faces a single maximization problem under both regimes. The difference lies solely in the constraints: the Coasean regime allows the agent to choose a policy contingent on the finer information set $\mathcal{I}_i$, whereas the public regime restricts the policy to the coarser set $\mathcal{I}^{pub}$. Since $\mathcal{I}^{pub} \subset \mathcal{I}_i$, the set of feasible decision rules under the public regime is a strict subset of those available under the Coasean regime. Consequently, the agent can achieve a higher expected payoff by tailoring their action to their private information.

Therefore, at a constitutional stage, the Coasean transfer regime receives unanimous support. Note that this holds even if agents are fully heterogeneous and know their own types (β_i, η_i) when voting. The "veil of ignorance" need only obscure the future realized signals, not the agent's own identity.

The natural next question—politically and economically—is whether agents would still support the Coasean transfer at the interim stage, or *after* observing their private information.

4.2 Interim Result: Coasean Transfers vs. Optimal Pigouvian Tax

We now analyze the incentives and distributional conflicts that emerge at the *interim stage*—that is, after agents have observed their private signals but before the aggregate outcome is realized. In this analysis, all expectations are conditional on the agent's private information set $\mathcal{I}_i$ (denoted by $\mathbb{E}_i[\cdot]$), distinct from the unconditional ex-ante expectations $\mathbb{E}[\cdot]$ used in the previous section.

We compare the Coasean transfer to the simpler, non-contingent Pigouvian tax τ derived in Section 3.2.2. The comparison reveals a clean decomposition of the welfare difference for each agent into a pure efficiency gain and a distributional transfer.

Proposition 11 (Interim Comparison, Coasean Transfer vs. Pigouvian Tax)**.** *Fix agent i with type (β_i, η_i) under Assumption 2. Let a_i^{ct} and A^{ct} denote the individual and aggregate actions under the Coasean transfer, and a_i^{tax} and A^{tax} denote the corresponding actions under the optimal Pigouvian tax. Define the differences as $\Delta a_i := a_i^{ct} - a_i^{tax}$ and $\Delta A := A^{ct} - A^{tax}$. The agent's interim expected payoff difference is:*

$$\mathbb{E}_i[u_i^{ct} - u_i^{tax}] = \underbrace{\frac{1}{2}\eta_i(\Delta a_i)^2}_{\textit{Individual flexibility}} + \underbrace{\frac{1}{2}\kappa^2 s_2 \mathbb{E}_i[(\Delta A)^2]}_{\textit{Collective adaptability}} + \underbrace{\mathbb{E}_i[(a_i^{tax} - A^{tax})(p(\Omega^{ct}) + \tau)]}_{\textit{Distributional transfer}}. \quad (39)$$

This identity is derived from an exact second-order expansion, using the agents' FOCs from each regime. The first two terms in the expression represent the pure efficiency gain. As a sum of squared terms (given $\eta_i > 0, s_2 > 0$), this gain is always non-negative and is strictly positive whenever the Coasean transfer's flexible allocation differs from the Pigouvian allocation. It captures the value of allowing individual actions to respond to private information (*individual flexibility*) and of allowing the aggregate action to adapt to the realized state (*collective adaptability*). Note that the first term does not require an expectation operator because agents know their own action adjustment Δa_i at the interim stage.

The third term is the net transfer difference, which captures the distributional conflict. An agent i supports the Coasean transfer only if this term is not so negative that it outweighs the guaranteed efficiency gain. This reveals the precise nature of the political opposition.

We can analyze this transfer term by identifying an agent's position relative to the aggregate. To fix ideas, consider a positive externality where the optimal Pigouvian instrument is a subsidy (mathematically, $\tau < 0$).

An "overperformer" ($a_i^{tax} > A^{tax}$) is an agent who contributes more than the average. Under the fixed Pigouvian regime, this agent receives a net subsidy payment proportional to $|\tau|$. Under the Coasean transfer, they are a net seller facing the adaptive price $p(\Omega^{ct})$. This agent benefits from a high price and would oppose the Coasean transfer only if they expect the adaptive price $p(\Omega^{ct})$ to fall significantly below the fixed subsidy level $|\tau|$. If the price drops low enough (even while remaining positive), the loss in transfer revenue can outweigh the efficiency gain.

Conversely, an "underperformer" ($a_i^{tax} < A^{tax}$) is a net buyer. Under the Pigouvian regime, they receive a smaller subsidy payment. This agent would oppose the Coasean transfer only if they expect the adaptive price to rise significantly above the fixed benchmark $|\tau|$, increasing their effective costs by more than the efficiency gain. Interim opposition, therefore, comes not from the mechanism's inefficiency—it is strictly more efficient—but from agents who expect the adaptive price to move unfavorably relative to the fixed position they held under the Pigouvian tax.

5 Harnessing Uncertainty and Disagreement

A central challenge for regulation is that agents and regulators are uncertain about the true state of the world. This uncertainty is typically viewed as a barrier, forcing regulators into second-best, information-blind instruments. We now show that the Coasean transfer inverts this problem: it is designed to *harness* this uncertainty, turning it from a barrier into a catalyst for an efficient, information-driven response.

The intuition is a direct consequence of Bayesian decision-making. As agents' signals become more informative relative to the prior, the "pull of the prior" weakens, forcing them to rely more on their private signals. Their actions become more sensitive to their information. The Coasean transfer, by implementing the team-efficient solution, aligns incentives such that this heightened sensitivity aggregates to the optimal outcome.

To derive a closed-form solution and analyze these effects, we now impose a standard parametric and informational structure.

Assumption 3 (Additive Outcome and Affine Information Structure)**.** *The exogenous component $\Gamma(\theta, \nu)$ is additive and linear in the fundamental state θ and the noise term ν:*

$$\Gamma(\theta, \nu) = \theta + \nu. \tag{40}$$

Furthermore, the information structure is affine, such that the average posterior belief is a linear function of the true state:

$$\int_0^1 m_{i,\theta} di = \lambda_\theta \theta + (1 - \lambda_\theta)\mu_\theta, \tag{41}$$

and similarly for ν, where $\lambda_\theta, \lambda_\nu \in [0, 1]$ parameterize the aggregate quality (or "precision") of information.

This linearity is a standard feature of Bayesian updating with Gaussian variables, and extends more generally to the broader class of elliptical distributions. The property holds regardless of the number or composition of public and private signals. This structure, combined with Assumption 2, ensures the strategic-form problem is linear and yields a tractable, closed-form solution for the equilibrium coefficients.

5.1 Harnessing Uncertainty

Under these assumptions, we can solve the fixed-point problem from Proposition 4, equation (15), for the team-efficient aggregate action A^{te} in closed form. We express the solution using the degree of coordination $\alpha \equiv -\kappa^2 s_2$.

Proposition 12 (Team-Efficient Action under LQ-Affine Structure)**.** *Under Assumptions 2 and 3, the team-efficient aggregate action $A^{te} = C_0 + C_\theta \theta + C_\nu \nu$ is given by the coefficients:*

$$C_0 = \frac{\kappa s_1 - (\alpha/\kappa)\mu_\Omega^{sq}}{H - \alpha}, \qquad C_\theta = \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta}, \qquad C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu}. \tag{42}$$

This proposition is derived by solving the linear system of equations that emerges from the strategic form (15) combined with the affine information structure. The result shows how the economy's aggregate response to the fundamentals θ and ν is an endogenous function of information precision. The "harnessing" of uncertainty is now clear. The coefficient C_θ measures the aggregate action's sensitivity to the fundamental θ. The magnitude of this sensitivity, $|C_\theta|$, is an increasing function of the information precision λ_θ (i.e., $\partial|C_\theta|/\partial\lambda_\theta > 0$), provided the private and social motives are not perfectly offset (generically, $|H\bar{B}_H + \alpha/\kappa| > 0$). Similarly, the magnitude of the sensitivity to the noise state ν, $|C_\nu|$, is also strictly increasing in its precision λ_ν (i.e., $\partial|C_\nu|/\partial\lambda_\nu > 0$).

In short, the Coasean transfer harnesses uncertainty by optimally aggregating agents' information-driven actions. As agents' prior uncertainty increases (the "pull of the prior" weakens), they are forced to rely more on their private signals. The Coasean transfer's price mechanism correctly aggregates this heightened sensitivity, ensuring the collective response to the fundamental is as strong and as efficient as their information quality $(\lambda_\theta, \lambda_\nu)$ permits.

5.2 Harnessing Disagreement

Thus far, our analysis has assumed agents have rational expectations. We now relax this by introducing *distrust*, where agents question the reliability of others' information. This builds on the literature examining differing beliefs about others' information (Banerjee, 2011) and is conceptually related to the "sentiment risk" channel in Dumas et al. (2009).

We model distrust as an agent's belief that the aggregate action is less responsive to the fundamentals than it truly is. We formalize this with parameters $\varphi_\theta, \varphi_\nu \in [0, 1]$, where $\varphi = 1$ is the rational-expectations benchmark (full trust) and $\varphi < 1$ represents distrust. This modification enters the agent's problem through their perception of the strategic motive, $\mathbb{E}_i[A^{te}]$.

This belief can be micro-founded. To fix ideas, consider how agent i perceives the private signals of others regarding θ. While agent i trusts their own signal y_i, they may believe other agents' signals y_j ($j \neq i$) are less informative. Agent i's *belief* about j's signal, y_j^i, is:

$$y_j^i = \mu_\theta + \varphi_\theta(\theta - \mu_\theta) + \sqrt{1 - \varphi_\theta^2}\phi_i + \varepsilon_{y,j}. \tag{43}$$

Here, the trust parameter φ_θ captures the correlation agent i believes j's signal has with the true state θ, relative to their own. When $\varphi_\theta = 1$, agent i believes others' signals are structured just like their own. When $\varphi_\theta < 1$, i believes others' signals are distorted by idiosyncratic noise ϕ_i and are thus less correlated with the state.

When agents solve their optimization problem under this form of distrust, the resulting team-efficient aggregate action is modified as follows.

Proposition 13 (Team-Efficient Action under Distrust)**.** *Under Assumptions 2 and 3, and agent distrust parameterized by* $(\varphi_\theta, \varphi_\nu)$*, the team-efficient aggregate action* $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$ *is given by the coefficients:*

$$C_0 = \frac{\kappa s_1 - (\alpha/\kappa)\mu_\Omega^{sq}}{H - \alpha}, \qquad C_\theta = \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta\varphi_\theta}, \qquad C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu\varphi_\nu}. \tag{44}$$

This result shows how the Coasean transfer harnesses disagreement. The trust parameters φ_θ and φ_ν appear only in the denominators. As distrust increases (φ_θ or φ_ν fall toward zero), the denominators decrease, which in turn *increases* the magnitude of the response coefficients, $|C_\theta|$ and $|C_\nu|$.

The mechanism works via strategic substitutability ($\alpha = -s_2\kappa^2 < 0$). When agents distrust others ($\varphi < 1$), they perceive the aggregate action as being less responsive to the state. Given that actions are substitutes, their own optimal response—implemented by the Coasean transfer—is to rely more heavily on their own private information. This increases their action's sensitivity to the fundamental, strengthening the aggregate response and harnessing the disagreement.

It is important to note that this result concerns the equilibrium decision rules. The optimality of the Coasean transfer *itself* remains robust. The team-efficient allocation is defined by maximizing true social welfare, based on the true objective structure of the economy, not agents' subjective beliefs. The Coasean transfer implements this true allocation even under distrust.

6 Incentives for Information Acquisition

We now analyze how the Coasean transfer alters agents' incentives to acquire information. We show that the mechanism operates on two distinct margins. First, for fundamentals that already affect private payoffs (like θ), the Coasean transfer introduces a new strategic motive—the need to forecast the state-contingent price—which amplifies the existing incentive to learn. Second, and more fundamentally, the Coasean transfer creates an incentive to learn about states that agents would otherwise find irrelevant (like ν). By linking payoffs to the aggregate outcome Ω, the mechanism forces agents to "expand their interests" and acquire information about all fundamentals that drive social welfare.

6.1 Acquiring Information about θ

We first analyze the incentive to acquire information about θ. To derive a closed-form comparison of these incentives, we work under the linear-quadratic-affine framework (Assumptions 2 and 3).

In the status quo, an agent has a private motive to learn θ to better inform the private benefit term, $\beta_i \theta a_i$. A standard Pigouvian tax τ, being a constant set ex-ante, does not change this motive; it shifts the level of the action but not its sensitivity to information.

The Coasean transfer, by contrast, introduces a new, strategic motive. The agent's FOC (20) shows that the optimal action now also depends on the forecast of the state-contingent price, $\mathbb{E}_i[p(\Omega^{ct})]$. This creates a trade-off between the private benefit sensitivity (β_i) and the price sensitivity, which is governed by the degree of coordination α.

Proposition 14 (Incentives for Information Acquisition about θ)**.** *The incentive to acquire information about θ is strictly stronger under the Coasean transfer than in the status quo if and only if:*

$$\beta_i < \frac{-\alpha H(1 + \kappa \bar{B}_H \lambda_\theta)}{2\kappa(H - \alpha\lambda_\theta)}. \tag{45}$$

Under strategic substitutability ($\alpha < 0$), the Coasean transfer creates a new strategic motive for information acquisition: the desire to differentiate. Agents profit by acting against the expected aggregate action, which effectively means betting against the fundamental θ. (A high θ drives up the aggregate action A, depressing the transfer price $p(\Omega)$ and incentivizing the agent to free-ride by doing less.) This strategic motive generally amplifies the incentive to learn. However, if an agent's private benefit β_i is extremely high, their primary goal is simply to scale their action with θ, regardless of what others do. In such cases, the private motive dominates, and the additional value of differentiation becomes negligible relative to the status quo. The condition shows that the Coasean transfer strengthens incentives to learn precisely when the strategic value of differentiation dominates this private motive, i.e., for agents with a small enough β_i.

6.2 Acquiring Information about ν

We now analyze the incentive to acquire information about the exogenous state ν. This analysis reveals a more fundamental impact of the Coasean transfer: it makes agents "expand their interests."

In the status quo, an agent's optimal action $a_i^{sq} = \beta_i m_{i,\theta}/\eta_i$ depends only on their private benefit, which is tied to θ. The state ν has no bearing on their private optimization. Similarly, a fixed Pigouvian tax τ only adds a constant to this rule, leaving the agent's decision independent of ν. In both regimes, information about ν has zero instrumental value.

The Coasean transfer, however, forces the agent to forecast the state-contingent price $\mathbb{E}_i[p(\Omega^{ct})]$. This price is a function of the aggregate outcome, Ω^{ct}, which itself depends on ν via the term $\Gamma(\theta, \nu)$. To choose their optimal action, agents *must* now forecast ν. Information that was previously worthless acquires a strictly positive economic value.

Proposition 15 (Incentives for Information Acquisition about ν)**.** *Under Assumptions 2 and 3, the Coasean transfer creates a strictly positive incentive to acquire information about ν, whereas the marginal value of such information is zero in the status quo or under a Pigouvian tax.*

The intuition directly parallels the logic for θ, but with a key difference: regarding ν, there is no private motive to crowd out the strategic incentive. In the status quo, agents ignore ν entirely because it does not affect their private benefits. Under the Coasean transfer, however, the price fluctuates to stabilize aggregate outcomes (driven by society's aversion to volatility, $s_2 > 0$). This fluctuation activates the pure strategic motive—the desire to differentiate—described above. Because the price varies with ν, agents must forecast it to anticipate the transfer. Since there is no opposing private pressure to scale actions with ν, this strategic desire to differentiate operates unopposed, generating a strictly positive value for information regardless of the agent's type.

This result highlights a qualitative shift in incentives. Regarding θ, the Coasean transfer operates on the *intensive margin*, amplifying an existing incentive. Regarding ν, the Coasean transfer operates on the *extensive margin*, creating a new incentive where none existed before. This "expanding interest" is a direct consequence of indexing transfers to the aggregate outcome Ω. By doing so, the Coasean transfer endogenously incentivizes agents to monitor *any* fundamental risk factor that drives that outcome, not just the factors that affect their private bottom line.

7 Conclusion

When uncertainty and disagreement are pervasive, traditional regulatory tools for addressing externalities struggle. Pigouvian taxes require precise knowledge of externality costs, while Coasean bargaining demands frictionless negotiation. We propose an alternative: a system of outcome-contingent transfers, which we call "Coasean transfers," that works with these frictions rather than against them. The transfer resembles a Pigouvian tax in its simplicity but achieves Coasean efficiency through decentralized coordination. Each agent receives a payment tied to their deviation from the market average, with the payment rate determined by realized aggregate outcomes. Our main result shows that this transfer recreates a synthetic competitive Coasean market for the externality, with the optimal pricing schedule corresponding to the equilibrium price of that market. This market-equivalence bridges the Pigouvian and Coasean approaches: the transfer operates as a simple outcome-contingent instrument in the spirit of a Pigouvian tax, but implements the decentralized efficiency of Coasean bargaining.

The equilibrium allocation under Coasean transfers is team-efficient, requiring the regulator to know only its normative objectives and how individual actions affect aggregate outcomes. It needs neither the distribution of private information nor the realized fundamental state. This allocation strictly dominates both the unregulated status quo and traditional tools such as Pigouvian taxes. Moreover, the transfer incentivizes agents to acquire information and remains robust even when agents distrust each other's signals, turning these traditional obstacles into sources of stronger incentives. The transfer is also politically viable: if brought to a vote under a veil of ignorance, it would enjoy unanimous support.

These results suggest a broader lesson for policy design in complex environments. Rather than seeking to overcome informational frictions such as uncertainty and disagreement, effective policy can harness them productively. By creating the "missing markets" for externalities, Coasean transfers can achieve decentralized coordination without central knowledge, even amid uncertainty and disagreement.

References

Aliprantis, C. D. and K. C. Border (2006). *Infinite dimensional analysis: a hitchhiker's guide*. Springer.

Amador, M. and P.-O. Weill (2010). Learning from prices: Public communication and welfare. *Journal of Political Economy 118*(5), 866–907.

Angeletos, G.-M. and A. Pavan (2007). Efficient use of information and social value of information. *Econometrica 75*(4), 1103–1142.

Angeletos, G.-M. and A. Pavan (2009). Policy with dispersed information. *Journal of the European Economic Association 7*(1), 11–60.

Banerjee, S. (2011). Learning from prices and the dispersion in beliefs. *The Review of Financial Studies 24*(9), 3025–3068.

Bergemann, D. and S. Morris (2016). Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review 106*(5), 586–591.

Bernanke, B. S. (1983). Irreversibility, uncertainty, and cyclical investment. *The Quarterly Journal of Economics 98*(1), 85–106.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Blackwell, D. (1951). Comparison of experiments. In J. Neyman (Ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, CA, pp. 93–102. University of California Press.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 265–272.

Caballero, R. J. and R. S. Pindyck (1996). Uncertainty, investment, and industry evolution. *International Economic Review 37*(3), 641–662.

Coase, R. H. (1960). The problem of social cost. *The Journal of Law and Economics 3*(1), 1–44.

Colombo, L., G. Femminis, and A. Pavan (2025). Investment subsidies with spillovers and endogenous private information: Why pigou got it all right. *Available at SSRN 5165242*.

Dasgupta, P., P. Hammond, and E. Maskin (1979). The implementation of social choice rules: Some general results on incentive compatibility. *The Review of Economic Studies 46*(2), 185–216.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under uncertainty.* Princeton University Press.

Dumas, B., A. Kurshev, and R. Uppal (2009). Equilibrium portfolio strategies in the presence of sentiment risk and excess volatility. *The Journal of Finance 64*(2), 579–629.

Hayek, F. A. (1945). The use of knowledge in society. *American Economic Review 35*(4), 519–530.

Hellwig, C. and L. Veldkamp (2009). Knowing what others know: Coordination motives in information acquisition. *The Review of Economic Studies 76*(1), 223–251.

Holmström, B. and R. B. Myerson (1983). Efficient and durable decision rules with incomplete information. *Econometrica: Journal of the Econometric Society*, 1799–1819.

Jajodia, S., A. K. Ghosh, V. Swarup, C. Wang, and X. S. Wang (Eds.) (2011). *Moving Target Defense: Creating Asymmetric Uncertainty for Cyber Threats*, Volume 54 of *Advances in Information Security.* New York: Springer.

Keynes, J. M. (1964). *General Theory of Employment, Interest, and Money.* Harbinger Books. Reprint of 1936 edition.

Knight, F. H. (1921). *Risk, Uncertainty, and Profit.* Houghton Mifflin.

Kunreuther, H. and G. Heal (2003). Interdependent security. *Journal of Risk and Uncertainty 26*(2-3), 231–249.

Kwerel, E. (1977). To tell the truth: Imperfect information and optimal pollution control. *The Review of Economic Studies 44*(3), 595–601.

Laffont, J.-J. (1994). Regulation of pollution with asymmetric information. In *Nonpoint source pollution regulation: Issues and analysis*, pp. 39–66. Springer.

Laffont, J.-J. and J. Tirole (1986). Using cost observation to regulate firms. *Journal of Political Economy 94*(3, Part 1), 614–641.

Lemoine, D. (2021). The climate risk premium: how uncertainty affects the social cost of carbon. *Journal of the Association of Environmental and Resource Economists 8*(1), 27–57.

Lemoine, D. (2024). Informationally efficient climate policy: Designing markets to measure and price externalities. Technical report, National Bureau of Economic Research.

Marschak, J. and R. Radner (1972). *Economic Theory of Teams*. New Haven, CT: Yale University Press.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *Review of Economic Studies 38*(2), 175–208.

Montero, J.-P. (2008). A simple auction mechanism for the optimal allocation of the commons. *American Economic Review 98*(1), 496–518.

Morris, S. and H. S. Shin (2002). Social value of public information. *American Economic Review 92*(5), 1521–1534.

Myatt, D. and C. Wallace (2008). On the sources and value of information: Public announcements and macroeconomic performance. Working paper, Department of Economics, University of Oxford.

Myatt, D. P. and C. Wallace (2012). Endogenous information acquisition in coordination games. *The Review of Economic Studies 79*(1), 340–374.

Pigou, A. C. (1920). *The Economics of Welfare* (1st ed.). London: Macmillan & Co. Pp. xxxvi + 976. 8vo.

Radner, R. (1962). Team decision problems. *The Annals of Mathematical Statistics 33*(3), 857–881.

Radner, R. (1982). Equilibrium under uncertainty. *Handbook of mathematical economics 2*, 923–1006.

Rawls, J. (1971). *A Theory of Justice* (Original ed.). Cambridge, MA: Belknap Press of Harvard University Press.

Roberts, M. J. and M. Spence (1976). Effluent charges and licenses under uncertainty. *Journal of Public Economics 5*(3-4), 193–208.

Rodrik, D. (1991). Policy uncertainty and private investment in developing countries. *Journal of Development Economics 36*(2), 229–242.

Romer, P. M. (1986). Increasing returns and long-run growth. *Journal of Political Economy 94*(5), 1002–1037.

Shiller, R. J. (1994). *Macro markets: creating institutions for managing society's largest economic risks*. OUP Oxford.

Veldkamp, L. L. (2011). *Information choice in macroeconomics and finance*. Princeton University Press.

Vives, X. (1988). Aggregation of information in large cournot markets. *Econometrica*, 851–876.

Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies 41*(4), 477–491.

A Appendix

A.1 Proof of Proposition 1

The social planner maximizes ex-post welfare state by state. Fix the state (θ, ν). The objective function is strictly concave in the allocation $\{a_i\}_{i\in[0,1]}$ because the private cost functions are strictly convex ($\eta_i > 0$) and the social value function $S(\cdot)$ is concave (Assumption 1). Thus, the first-best allocation exists, is unique, and is characterized by the FOCs. Differentiating the objective in equation (3) with respect to a_i and using (3) yields

$$-\eta_i a_i^{fb} + \beta_i\theta + \kappa S'(\Omega^{fb}) = 0. \tag{A1}$$

Rearranging terms yields condition (8):

$$a_i^{fb} = \frac{\beta_i\theta + \kappa S'(\Omega^{fb})}{\eta_i}. \tag{A2}$$

Integrating over the population $i \in [0,1]$, we obtain the aggregate condition:

$$A^{fb} = \int_0^1 \frac{\beta_i\theta}{\eta_i} di + \kappa S'(\Omega^{fb}) \int_0^1 \frac{1}{\eta_i} di = \bar{B}_H\theta + \frac{\kappa}{H} S'(\Omega^{fb}), \tag{A3}$$

where $\bar{B}_H \equiv \int_0^1 (\beta_i/\eta_i)di$ and $H \equiv (\int_0^1 \eta_i^{-1} di)^{-1}$. Substituting $\Omega^{fb} = \kappa A^{fb} + \Gamma(\theta,\nu)$ yields equation (9). To establish the uniqueness of A^{fb}, define the mapping $\Psi(A) \equiv \bar{B}_H\theta + \frac{\kappa}{H} S'(\kappa A + \Gamma(\theta,\nu))$. Since $S''(\cdot) \le 0$, the function $\kappa S'(\cdot)$ is nonincreasing, which implies that $\Psi(A)$ is nonincreasing in A. Because the left-hand side of (9) is strictly increasing in A while the right-hand side is nonincreasing, there exists a unique fixed point $A^{fb} = \Psi(A^{fb})$. The individual actions $\{a_i^{fb}\}$ are then uniquely determined by the FOC. □

A.2 Proof of Proposition 2

The utilitarian planner maximizes expected social welfare

$$W(\{a_i\}) = \mathbb{E}\left[\int_0^1 \left(e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i\right) di + S(\kappa A + \Gamma(\theta,\nu))\right], \qquad A = \int_0^1 a_i di. \tag{A4}$$

The private component of the utility function is strictly concave in a_i because $\eta_i > 0$. By Assumption 1, $S(\cdot)$ is concave. Since $\Omega = \kappa A + \Gamma(\theta,\nu)$ is an affine transformation of the allocation $\{a_i\}$, the composition $S(\kappa A + \Gamma)$ preserves concavity. Consequently, the aggregate objective $W(\{a_i\})$ is strictly concave with respect to the allocation $\{a_i\}_{i\in[0,1]}$.

Strict concavity implies that the optimization problem admits a unique global maximizer. Let $\{a_i^{te}\}_{i\in[0,1]}$ denote this unique allocation. The FOC with respect to a_i is

$$\frac{\partial W}{\partial a_i} = -\eta_i a_i^{te} + \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\Omega^{te})] = 0. \tag{A5}$$

Rearranging terms yields condition (10):

$$\eta_i a_i^{te} = \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\Omega^{te})]. \tag{A6}$$

To establish existence and uniqueness, we examine the monotonicity of the best responses. For a fixed aggregate action A, define the conditional individual response function as

$$a_i(A) = \frac{\beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\kappa A + \Gamma(\theta,\nu))]}{\eta_i}. \tag{A7}$$

Under Assumption 1, $S''(\cdot) \leq 0$, which implies that the marginal social value function $\kappa S'(\cdot)$ is nonincreasing. Since $\kappa > 0$, the term $\kappa \mathbb{E}_i[S'(\kappa A + \Gamma(\theta,\nu))]$ is nonincreasing in A. It follows that $a_i(A)$ is nonincreasing in A for every agent i.

Aggregating these individual responses yields the mapping $\Phi(A) = \int_0^1 a_i(A)di$. The continuity of $S'(\cdot)$ and the integrability of the parameters ensure that $\Phi(A)$ is continuous. Moreover, because each component $a_i(A)$ is nonincreasing, the aggregate function $\Phi(A)$ is nonincreasing.

The equilibrium aggregate action A^{te} is characterized by the fixed-point condition $A^{te} = \Phi(A^{te})$. Define $G(A) \equiv \Phi(A) - A$. Since $\Phi(\cdot)$ is continuous and nonincreasing, $G(\cdot)$ is continuous and strictly decreasing. Monotonicity of $\Phi(\cdot)$ implies that for any $A > 0$, $\Phi(A) \leq \Phi(0)$, and thus $G(A) \leq \Phi(0) - A$. Taking the limit yields $\lim_{A\to\infty} G(A) = -\infty$. Similarly, for any $A < 0$, monotonicity implies $\Phi(A) \geq \Phi(0)$, and thus $G(A) \geq \Phi(0) - A$, ensuring $\lim_{A\to-\infty} G(A) = \infty$. By the Intermediate Value Theorem, there exists A^{te} such that $G(A^{te}) = 0$, and the strict monotonicity of $G(\cdot)$ ensures that A^{te} is unique. □

A.3 Proof of Proposition 3

From Proposition 1, the first-best allocation satisfies the condition $\eta_i a_i^{fb} = \beta_i \theta + \kappa S'(\Omega^{fb})$. Substituting the linear marginal social value function from Assumption 2, $S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq})$, yields the expression for the individual action:

$$a_i^{fb} = \frac{\beta_i \theta + \kappa[s_1 - s_2(\Omega^{fb} - \mu_\Omega^{sq})]}{\eta_i}. \tag{A8}$$

Aggregating across agents and using the definitions $\bar{B}_H \equiv \int_0^1 (\beta_i/\eta_i)di$ and $H \equiv (\int_0^1 \eta_i^{-1} di)^{-1}$, we obtain the aggregate action:

$$A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H}[s_1 - s_2(\Omega^{fb} - \mu_\Omega^{sq})]. \tag{A9}$$

The aggregate outcome is defined as $\Omega^{fb} = \kappa A^{fb} + \Gamma(\theta, \nu)$. Substituting the expression for A^{fb} into this definition yields:

$$\Omega^{fb} = \Gamma(\theta, \nu) + \kappa \bar{B}_H \theta + \frac{\kappa^2}{H}[s_1 - s_2(\Omega^{fb} - \mu_\Omega^{sq})]. \tag{A10}$$

This is a linear equation in Ω^{fb}. Collecting the terms involving Ω^{fb} on the left-hand side results in:

$$\Omega^{fb}\left(1 + \frac{\kappa^2 s_2}{H}\right) = \Gamma(\theta, \nu) + \kappa \bar{B}_H \theta + \frac{\kappa^2 s_1}{H} + \frac{\kappa^2 s_2}{H}\mu_\Omega^{sq}. \tag{A11}$$

Solving for Ω^{fb} yields the closed-form expression stated in equation (14). Uniqueness is guaranteed by the strict concavity of the objective function ($S'' = -s_2 < 0$). □

A.4 Proof of Proposition 4

The FOCs for the first-best (FB) and team-efficient (TE) allocations (eqs. (A1) and (A6)) are:

$$\eta_i a_i^{fb} = \beta_i \theta + \kappa S'(\Omega^{fb}), \tag{A12}$$

$$\eta_i a_i^{te} = \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[S'(\Omega^{te})]. \tag{A13}$$

Taking the conditional expectation $\mathbb{E}_i[\cdot]$ of (A12) and subtracting it from (A13) yields

$$\eta_i(a_i^{te} - \mathbb{E}_i[a_i^{fb}]) = \kappa \mathbb{E}_i\big[S'(\Omega^{te}) - S'(\Omega^{fb})\big]. \tag{A14}$$

Under Assumption 2, the marginal social value is linear with slope $-s_2$. Consequently, the difference in marginal social values is proportional to the difference in aggregate actions:

$$\kappa \mathbb{E}_i\big[S'(\Omega^{te}) - S'(\Omega^{fb})\big] = -\kappa s_2 \mathbb{E}_i\big[\Omega^{te} - \Omega^{fb}\big] = -\kappa^2 s_2 \mathbb{E}_i\big[A^{te} - A^{fb}\big]. \tag{A15}$$

Using the definition $\alpha \equiv -\kappa^2 s_2$, the difference equation becomes

$$\eta_i(a_i^{te} - \mathbb{E}_i[a_i^{fb}]) = \alpha \mathbb{E}_i\big[A^{te} - A^{fb}\big]. \tag{A16}$$

To expand $\mathbb{E}_i[A^{fb}]$, we express the first-best aggregate action A^{fb} using the individual first-best condition. From Proposition 1, $A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H} S'(\Omega^{fb})$, where, from equation (A12), $\kappa S'(\Omega^{fb}) =$

$\eta_i a_i^{fb} - \beta_i \theta$. Hence,

$$A^{fb} = \bar{B}_H \theta + \frac{1}{H}(\eta_i a_i^{fb} - \beta_i \theta) = \theta \left(\bar{B}_H - \frac{\beta_i}{H} \right) + \frac{\eta_i}{H} a_i^{fb}. \tag{A17}$$

Taking the conditional expectation $\mathbb{E}_i[\cdot]$ leads to

$$\mathbb{E}_i[A^{fb}] = m_{i,\theta} \left(\bar{B}_H - \frac{\beta_i}{H} \right) + \frac{\eta_i}{H} \mathbb{E}_i[a_i^{fb}]. \tag{A18}$$

Substituting this expression back into (A16) and rearranging terms:

$$\eta_i a_i^{te} = \eta_i \mathbb{E}_i[a_i^{fb}] + \alpha \mathbb{E}_i[A^{te}] - \alpha \left[m_{i,\theta} \left(\bar{B}_H - \frac{\beta_i}{H} \right) + \frac{\eta_i}{H} \mathbb{E}_i[a_i^{fb}] \right]. \tag{A19}$$

Grouping the terms involving $\mathbb{E}_i[a_i^{fb}]$ and dividing by η_i yields the strategic form (15):

$$a_i^{te} = \left(1 - \frac{\alpha}{H} \right) \mathbb{E}_i[a_i^{fb}] + \frac{\alpha}{\eta_i} \mathbb{E}_i[A^{te}] - \frac{\alpha}{\eta_i} \left(\bar{B}_H - \frac{\beta_i}{H} \right) m_{i,\theta}. \quad \square \tag{A20}$$

A.5 Proof of Proposition 5

Consider the problem of agent i under the Coasean transfer $T_i = (a_i - A)\Omega)$. The agent chooses a_i to maximize the expected payoff:

$$\max_{a_i} \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) + (a_i - A) p(\Omega) \right]. \tag{A21}$$

Since agents are atomless, they treat the aggregate action A and the outcome Ω as independent of their individual choice a_i. Consequently, $\partial \Omega / \partial a_i = 0$ and $\partial p(\Omega) / \partial a_i = 0$. The FOC with respect to a_i is:

$$\mathbb{E}_i[-\eta_i a_i + \beta_i \theta + p(\Omega)] = 0. \tag{A22}$$

Rearranging terms yields the optimal decision rule for the agent:

$$\eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega)]. \tag{A23}$$

By Proposition 2, the team-efficient allocation $\{a_i^{te}\}$ is characterized by the condition $\eta_i a_i^{te} = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega^{te})]$. If the Coasean transfer sets the price equal to the social marginal value $p(\Omega) = \kappa S'(\Omega)$ for all possible outcomes, the agent's condition under the transfer is identical to the team-efficient condition. Thus, the equilibrium allocation induced by the Coasean transfer coincides with the unique team-efficient allocation. $\square$

A.6 Proof of Proposition 6

The regulator determines the optimal uniform tax τ by maximizing expected welfare subject to agents' best responses. Given a uniform tax τ and a lump-sum rebate $R = \tau A$, agent i maximizes their expected utility:

$$\max_{a_i} \mathbb{E}_i \left[\left(e - \frac{\eta_i}{2} a_i^2 \right) + \beta_i \theta a_i + S(\Omega) - \tau a_i + R \right]. \tag{A24}$$

Since agents are atomless, they treat the aggregates as exogenous. The FOC is $-\eta_i a_i + \beta_i m_{i,\theta} - \tau = 0$, which yields the decision rule

$$a_i(\tau) = \frac{\beta_i m_{i,\theta} - \tau}{\eta_i}. \tag{A25}$$

Aggregating across agents and using the independence of types and signals implies that the expected aggregate action is

$$\mathbb{E}[A(\tau)] = \bar{B}_H \mathbb{E}[\theta] - \frac{\tau}{H}. \tag{A26}$$

From equation (1), the outcome is $\Omega(\tau) = \kappa A(\tau) + \Gamma(\theta, \nu)$. Therefore, the sensitivity of the expected outcome to the tax is

$$\frac{\partial \mathbb{E}[\Omega(\tau)]}{\partial \tau} = -\frac{\kappa}{H}. \tag{A27}$$

The regulator chooses τ to maximize ex-ante expected welfare $\mathbb{E}[W(\tau)]$. Because tax revenues are redistributed via the rebate ($R = \tau A$), the transfer terms cancel out in the aggregate. Expected welfare is:

$$\mathbb{E}[W(\tau)] = \mathbb{E} \left[\int_0^1 \left(e - \frac{\eta_i}{2} a_i(\tau)^2 + \beta_i \theta a_i(\tau) \right) di + S(\Omega(\tau)) \right]. \tag{A28}$$

Since the best response $a_i(\tau)$ is linear in τ, the private cost term $-\frac{\eta_i}{2} a_i(\tau)^2$ is strictly concave in τ. Furthermore, because the aggregate outcome $\Omega(\tau)$ is linear in τ, the composite function $S(\Omega(\tau))$ preserves the concavity of $S(\cdot)$. The sum of these terms yields a strictly concave objective function. Differentiating with respect to τ, and noting that $\partial a_i / \partial \tau = -1/\eta_i$, yields

$$\frac{d\mathbb{E}[W(\tau)]}{d\tau} = \mathbb{E} \left[\int_0^1 (-\eta_i a_i(\tau) + \beta_i \theta) \left(-\frac{1}{\eta_i} \right) di + S'(\Omega(\tau)) \frac{\partial \Omega(\tau)}{\partial \tau} \right]. \tag{A29}$$

We can simplify the integral term by substituting the agent's FOC, $-\eta_i a_i(\tau) + \beta_i \theta = \tau + \beta_i(\theta - m_{i,\theta})$. Since $\mathbb{E}[\theta - m_{i,\theta}] = 0$, the integral in equation (A29) is equal to $-\tau/H$. The welfare derivative thus simplifies to

$$\frac{d\mathbb{E}[W]}{d\tau} = -\frac{\tau}{H} - \frac{\kappa}{H}\mathbb{E}[S'(\Omega)], \tag{A30}$$

where the last term follows from equation (A27). Setting this derivative to zero characterizes the optimal tax as

$$\tau^* = -\kappa\mathbb{E}[S'(\Omega(\tau^*))]. \tag{A31}$$

Under Assumption 2, we substitute the linear marginal social value function to obtain the condition

$$-\tau^* = \kappa s_1 - \kappa s_2 \left(\mathbb{E}[\Omega(\tau^*)] - \mu_\Omega^{sq}\right). \tag{A32}$$

Substituting the expression for $\mathbb{E}[\Omega(\tau^*)] = \mathbb{E}[\kappa A(\tau^*) + \Gamma(\theta, \nu)]$ and using (A26) yields

$$-\tau^* = \kappa s_1 - \kappa s_2 \left(\kappa \bar{B}_H \mathbb{E}[\theta] - \frac{\kappa\tau^*}{H} + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq}\right). \tag{A33}$$

Grouping the terms involving τ^* yields

$$-\tau^* \left(1 + \frac{\kappa^2 s_2}{H}\right) = \kappa s_1 - \kappa s_2 \left(\kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq}\right). \tag{A34}$$

which gives the closed-form solution stated in equation (30):

$$\tau^* = \frac{-\kappa s_1 + \kappa s_2 \left(\kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq}\right)}{1 + \kappa^2 s_2 / H}. \tag{A35}$$

Uniqueness follows from the global concavity of the welfare function. □

A.7 Proof of Proposition 7

We derive the welfare-maximizing cap A^{cap} in two steps. First, we characterize the competitive permit price for a given cap A. Second, we determine the optimal cap that maximizes expected welfare.

Market equilibrium for a given cap. Consider a cap-and-trade system with a fixed supply of permits A. The regulator auctions these permits at a market-clearing price P and redistributes the revenue as a lump-sum rebate $R = PA$. Each atomistic agent i takes P, R, and the aggregate outcome Ω as given and chooses a_i to maximize expected utility:

$$\max_{a_i} \mathbb{E}_i\Big[\Big(e - \frac{\eta_i}{2}a_i^2\Big) + \beta_i\theta a_i + S(\Omega) - Pa_i + R\Big], \tag{A36}$$

where $\Omega = \kappa A + \Gamma(\theta, \nu)$. Since $S(\Omega)$ and R do not depend on a_i, the FOC with respect to a_i is

$$-\eta_i a_i + \beta_i m_{i,\theta} - P = 0, \tag{A37}$$

which yields the agent's demand for permits (and hence for the action):

$$a_i(P) = \frac{\beta_i m_{i,\theta} - P}{\eta_i}. \tag{A38}$$

Aggregating across agents and using independence of types and signals, we obtain the expected aggregate action under price P:

$$\mathbb{E}[A(P)] = \int_0^1 \mathbb{E}[a_i(P)]\, di = \bar{B}_H \mathbb{E}[\theta] - \frac{P}{H}, \tag{A39}$$

where $\bar{B}_H \equiv \int_0^1 (\beta_i/\eta_i) di$ and $H \equiv (\int_0^1 \eta_i^{-1} di)^{-1}$. Given the cap A, the competitive permit price $P(A)$ is determined by market clearing,

$$\mathbb{E}[A(P(A))] = A, \tag{A40}$$

which implies a unique equilibrium price

$$P(A) = H\big(\bar{B}_H \mathbb{E}[\theta] - A\big). \tag{A41}$$

This expression is the aggregate inverse demand for permits.

Regulator's choice of the cap. Given a cap A, the permit market described above generates a competitive equilibrium with aggregate action A and outcome $\Omega = \kappa A + \Gamma(\theta, \nu)$. Because permit revenues $P(A)A$ are rebated lump-sum, they cancel out in aggregate welfare. Expected welfare as a function of the cap is therefore

$$\mathbb{E}[W(A)] = \mathbb{E}\Big[\int_0^1 \Big(e - \frac{\eta_i}{2} a_i(A)^2 + \beta_i \theta a_i(A)\Big) di + S\big(\kappa A + \Gamma(\theta, \nu)\big)\Big], \tag{A42}$$

where $\{a_i(A)\}_{i \in [0,1]}$ denotes the competitive allocation induced by the cap A through the permit price $P(A)$.

We simplify expected welfare by explicitly evaluating the private payoff component. Let $W_{priv}(A) = \int_0^1 (e - \frac{\eta_i}{2} a_i(A)^2 + \beta_i \theta a_i(A)) di$. We differentiate this term with respect to A. From the individual best response $a_i(P(A)) = (\beta_i m_{i,\theta} - P(A))/\eta_i$ and the price sensitivity $dP(A)/dA = -H$, the sensitivity of the individual action is $\frac{da_i}{dA} = \frac{da_i}{dP}\frac{dP}{dA} = (-\frac{1}{\eta_i})(-H) = \frac{H}{\eta_i}$. The derivative of private welfare is:

$$\frac{dW_{priv}}{dA} = \int_0^1 (\beta_i \theta - \eta_i a_i) \frac{da_i}{dA} di = H\theta \int_0^1 \frac{\beta_i}{\eta_i} di - H \int_0^1 a_i di = H \bar{B}_H \theta - HA. \tag{A43}$$

Thus

$$\frac{d\mathbb{E}[W]}{dA} = H\bar{B}_H\mathbb{E}[\theta] + \kappa\mathbb{E}\big[S'(\kappa A + \Gamma(\theta,\nu))\big] - HA. \tag{A44}$$

Using the inverse demand (A41), we can rewrite this derivative as

$$\frac{d\mathbb{E}[W]}{dA} = \mathbb{E}\big[P(A) + \kappa S'(\kappa A + \Gamma(\theta,\nu))\big]. \tag{A45}$$

The welfare-maximizing cap A^{cap*} therefore satisfies the condition

$$\mathbb{E}\big[P(A^{cap*}) + \kappa S'(\kappa A^{cap*} + \Gamma(\theta,\nu))\big] = 0, \tag{A46}$$

which equates the expected marginal private value of relaxing the cap (the permit price) to the expected marginal social cost.

Under Assumption 2, we have

$$S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq}), \tag{A47}$$

with $s_2 > 0$. Substituting $\Omega = \kappa A + \Gamma(\theta,\nu)$, taking expectations, and using linearity gives

$$\mathbb{E}\big[S'(\kappa A + \Gamma(\theta,\nu))\big] = s_1 - s_2\big(\kappa A + \mathbb{E}[\Gamma] - \mu_\Omega^{sq}\big). \tag{A48}$$

Substituting the above expression and equation (A41) into the FOC (A46) yields

$$0 = H\big(\bar{B}_H\mathbb{E}[\theta] - A\big) + \kappa\Big[s_1 - s_2\big(\kappa A + \mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\big)\Big]. \tag{A49}$$

Collecting the terms in A on the left-hand side,

$$\big(H + \kappa^2 s_2\big)A = H\bar{B}_H\mathbb{E}[\theta] + \kappa s_1 - \kappa s_2\big(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\big). \tag{A50}$$

Solving for A gives the unique welfare-maximizing cap in equation (32):

$$A^{cap*} = \frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa s_1 - \kappa s_2(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq})}{H + \kappa^2 s_2}, \tag{A51}$$

Since $S''(\Omega) = -s_2 < 0$, expected welfare is strictly concave in A, so A^{cap*} is unique.

Equivalence of optimal price and quantity instruments. We confirm the duality between the instruments by substituting the optimal cap A^{cap*} into the equilibrium inverse demand function, $P(A) = H(\bar{B}_H\mathbb{E}[\theta] - A)$. This substitution yields

$$P(A^{cap*}) = H\bar{B}_H\mathbb{E}[\theta] - H\left[\frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa s_1 - \kappa s_2(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq})}{H + \kappa^2 s_2}\right]. \tag{A52}$$

Multiplying the first term by $(H + \kappa^2 s_2)/(H + \kappa^2 s_2)$ and simplifying the numerator yields

$$P(A^{cap*}) = \frac{-\kappa s_1 + \kappa s_2(\kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq})}{1 + \kappa^2 s_2/H}, \tag{A53}$$

which is exactly the optimal Pigouvian tax τ^* derived in Proposition 6. □

A.8 Proof of Proposition 8

We derive the ranking of instruments by mapping the model's primitives into the Weitzman (1974) framework. This requires characterizing the slopes of the marginal social cost and marginal social benefit schedules with respect to the aggregate action A.

First, we characterize the Total Social Cost, $\text{TSC}(A)$, defined as the minimum private cost required to implement a given aggregate action A in a given state. The allocation problem is:

$$\text{TSC}(A) = \min_{\{a_i\}} \int_0^1 \frac{\eta_i}{2} a_i^2 di \quad \text{subject to} \quad \int_0^1 a_i di = A. \tag{A54}$$

Let λ denote the Lagrange multiplier. The FOC $\eta_i a_i = \lambda$ implies $a_i = \lambda/\eta_i$. Aggregating yields $A = \lambda/H$, so the shadow price is $\lambda = HA$. Substituting the optimal allocation $a_i(A) = HA/\eta_i$ back into the objective yields $\text{TSC}(A) = \frac{1}{2}HA^2$. The Marginal Social Cost is $\text{MSC}(A) \equiv \partial_A \text{TSC}(A) = HA$. Thus, the slope of the marginal cost curve is determined by the aggregate cost parameter:

$$\text{MSC}'(A) = H. \tag{A55}$$

Next, we characterize the Total Social Benefit, $\text{TSB}(A)$. The total social benefit is:

$$\text{TSB}(A) = \int_0^1 \beta_i \theta a_i(A) di + S(\kappa A + \Gamma(\theta, \nu)). \tag{A56}$$

Assuming cost-efficient allocation, we can substitute $a_i(A) = HA/\eta_i$. Then the integral term becomes $\theta H \bar{B}_H A$, which is linear in A. The Marginal Social Benefit is:

$$\text{MSB}(A) \equiv \frac{\partial \text{TSB}}{\partial A} = \theta H \bar{B}_H + \kappa S'(\Omega). \tag{A57}$$

Under Assumption 2, $S'(\Omega)$ is linear with slope $-s_2$. Differentiating $\text{MSB}(A)$ with respect to A yields the slope:

$$\text{MSB}'(A) = 0 + \kappa(-s_2)\kappa = -\kappa^2 s_2. \tag{A58}$$

Since $s_2 > 0$ (concavity of social value), the MSB curve is downward sloping.

Define the comparative advantage of prices over quantities as

$$\Delta \equiv \mathbb{E}\big[(\mathrm{TSB}(A^{\mathrm{price}}) - \mathrm{TSC}(A^{\mathrm{price}})) - (\mathrm{TSB}(A^{\mathrm{quantity}}) - \mathrm{TSC}(A^{\mathrm{quantity}}))\big], \tag{A59}$$

where A^{price} and A^{quantity} denote the aggregate actions implemented, in a given state, by an optimally chosen price (tax/subsidy) and quantity (cap) instrument, respectively.

Under Assumption 2, $\mathrm{TSC}(A)$ and $\mathrm{TSB}(A)$ are twice continuously differentiable and quadratic in A. Uncertainty about the cross-section enters marginal costs only through the aggregate cost parameter H, so that $\mathrm{MSC}'(A) = H$. Randomness in $(\theta, \Gamma(\theta,\nu), \bar{B}_H)$ shifts only the intercept of $\mathrm{MSB}(A)$ and does not affect $\mathrm{MSC}'(A)$ or $\mathrm{MSB}'(A)$. We follow Weitzman (1974) in assuming that the cost shock H is independent of the benefit-side shocks $(\theta, \Gamma(\theta,\nu), \bar{B}_H)$ and that cost uncertainty is small, so his local quadratic approximation applies. Under these conditions, equation (20) in Weitzman (1974) delivers a local ranking of price and quantity instruments.

Let A^* denote, for each realization of the state, the efficient aggregate action that maximizes $\mathrm{TSB}(A) - \mathrm{TSC}(A)$. Weitzman's equation (20) implies that, for sufficiently small cost uncertainty,

$$\mathrm{sign}(\Delta) = \mathrm{sign}\big(\mathbb{E}[\mathrm{MSC}'(A^*)] + \mathrm{MSB}'(A^*)\big). \tag{A60}$$

In our environment the slopes are constant in A and given by equations (A55) and (A58), so that

$$\mathbb{E}[\mathrm{MSC}'(A^*)] + \mathrm{MSB}'(A^*) = \mathbb{E}[H] - \kappa^2 s_2. \tag{A61}$$

Therefore

$$\Delta > 0 \iff \mathbb{E}[H] > \kappa^2 s_2. \tag{A62}$$

When $\Delta > 0$ the expected surplus under the price instrument is larger than under the quantity instrument, so the regulator prefers a price (tax/subsidy); when $\Delta < 0$ the regulator prefers a quantity (cap). This establishes the ranking in Proposition 8. □

A.9 Proof of Proposition 9

Consider the mechanism in which each agent i chooses an action a_i and receives the transfer

$$T_i = (a_i - A)p(\Omega), \qquad p(\Omega) = \kappa S'(\Omega), \qquad \Omega = \kappa A + \Gamma(\theta,\nu), \tag{A63}$$

where $A = \int_0^1 a_j dj$. A (Bayesian) strategy for agent i is a measurable decision rule $a_i(\mathcal{I}_i)$ mapping her information $\mathcal{I}_i$ into an action.

Given a profile of strategies of the other agents, the induced aggregate A and outcome Ω are random variables that do not depend on a_i, because agent i has measure zero ($\partial A/\partial a_i = 0$ and $\partial \Omega/\partial a_i = 0$). Agent i chooses a_i to maximize

$$\mathbb{E}_i\Big[\Big(e - \frac{\eta_i}{2}a_i^2\Big) + \beta_i\theta a_i + S(\Omega) + (a_i - A)p(\Omega)\Big]. \tag{A64}$$

Taking A and Ω as given, the FOC with respect to a_i is

$$\mathbb{E}_i\big[-\eta_i a_i + \beta_i\theta + p(\Omega)\big] = 0 \quad \implies \quad \eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega)]. \tag{A65}$$

We substitute the Coasean pricing rule $p(\Omega) = \kappa S'(\Omega)$ into (A65):

$$\eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[\kappa S'(\Omega)]. \tag{A66}$$

This condition is mathematically identical to the condition characterizing the team-efficient allocation derived in Proposition 2 (Equation 10). Since the agent's objective function is strictly concave in a_i (given $\eta_i > 0$), this FOC characterizes the unique best response. Consequently, the team-efficient allocation is the unique Bayesian Nash equilibrium of the mechanism.

For any profile of actions $\{a_i\}$ and any realization of Ω, the total transfer is:

$$\int_0^1 T_i di = \int_0^1 (a_i - A)p(\Omega)di = p(\Omega)\Big(\int_0^1 a_i di - A\Big) = 0. \tag{A67}$$

Thus, the mechanism is exactly budget-balanced ex post in every state.

Finally, we establish ex-ante individual rationality. Let U_i^{te} denote agent i's ex-ante expected utility under the team-efficient allocation, and U_i^{ct} denote the utility under the mechanism. Because the mechanism is budget-balanced in every state ($\int T_i di = 0$), the ex-ante expected transfer is zero. Consequently, the mechanism yields the same expected payoff as the team-efficient allocation: $\mathbb{E}[U_i^{ct}] = \mathbb{E}[U_i^{te}]$. Therefore, if the team-efficient allocation satisfies the individual rationality constraint relative to the unregulated status quo ($\mathbb{E}[U_i^{te}] \geq \bar{U}$), the Coasean transfer mechanism satisfies it as well. □

A.10 Proof of Proposition 10

Let the *internalized* utility of agent i be denoted by $\tilde{u}_i(a, z)$. This is precisely the objective that a social planner would use when choosing a for agent i, holding $A_{-i}(z)$ fixed—i.e., this objective internalizes the social value term $S(\cdot)$ at the individual level. Writing the aggregate action as $a + A_{-i}(z)$ makes explicit how agent i's action enters the social value term:

$$\tilde{u}_i(a, z) = \Big(e - \frac{\eta_i}{2}a^2\Big) + \beta_i\theta a + S\big(\kappa[a + A_{-i}(z)] + \Gamma(z)\big). \tag{A68}$$

Given Assumption 1, $S(\cdot)$ is concave. The private cost term $-\eta_i a^2/2$ is strictly concave in a since $\eta_i > 0$. Because the composition of the concave social value function with an affine mapping in a preserves concavity, and the private benefit $\beta_i \theta a$ is linear in a, the internalized utility $\tilde{u}_i(\cdot, z)$ is strictly concave in a for every realized state z.

We compare the expected internalized utility of agent i under two regimes characterized by different information sets. Let $\mathcal{I}^{pub}$ and $\mathcal{I}_i$ denote the σ-algebras generated by public information and agent i's information, respectively, with $\mathcal{I}^{pub} \subset \mathcal{I}_i$. The public-information regime (such as a Pigouvian tax) implements a policy $a_i^{pub}(z)$ that is measurable with respect to $\mathcal{I}^{pub}$. The Coasean transfer regime implements the team-efficient allocation $a_i^{te}(z)$, which is measurable with respect to $\mathcal{I}_i$.

Consider the conditional maximization problem for the action of agent i (under the internalized objective) given the information set $\mathcal{I}_i$, defined as

$$\max_{a \in \mathbb{R}} \mathbb{E}\big[\tilde{u}_i(a, z) \mid \mathcal{I}_i\big]. \tag{A69}$$

Let the objective function be denoted by $f(z, a)$, which represents the expected utility *conditional on the information* $\mathcal{I}_i$ *observed in state* z. We observe that $f(z, a)$ satisfies the Carathéodory conditions: for each fixed a, the mapping $z \mapsto f(z, a)$ is $\mathcal{I}_i$-measurable (since the conditional expectation varies with z through the information set), and for each fixed z, the mapping $a \mapsto f(z, a)$ is continuous in a due to the continuity of $\tilde{u}_i$. Furthermore, because the quadratic cost term dominates the linear and concave benefit terms, $\tilde{u}_i(a, z) \to -\infty$ as $|a| \to \infty$. This coercivity ensures that the maximizer lies within a compact interval for every z.

By the Measurable Maximum Theorem (see Aliprantis and Border, 2006, Thm. 18.19), there exists a unique $\mathcal{I}_i$-measurable selector $a_i^\star(z)$ that maximizes the conditional expected utility almost surely. (Note that $\mathcal{I}_i$-measurability implies that $a_i^\star(z)$ depends on the realized state z only through the information set $\mathcal{I}_i$). Since the Coasean transfer implements the team-efficient allocation, we identify $a_i^\star(z) = a_i^{te}(z)$.

Now consider any alternative policy $b(z)$ that is measurable with respect to the coarser public information set $\mathcal{I}^{pub}$. Since $\mathcal{I}^{pub} \subset \mathcal{I}_i$, $b(z)$ is also $\mathcal{I}_i$-measurable. By the definition of optimality under $\mathcal{I}_i$, it must be that

$$\mathbb{E}\big[\tilde{u}_i(a_i^{te}(z), z) \mid \mathcal{I}_i\big] \geq \mathbb{E}\big[\tilde{u}_i(b(z), z) \mid \mathcal{I}_i\big] \quad \text{a.s.} \tag{A70}$$

Taking unconditional expectations on both sides and applying the Law of Iterated Expectations yields

$$\mathbb{E}\big[\tilde{u}_i(a_i^{te}(z), z)\big] \geq \mathbb{E}\big[\tilde{u}_i(b(z), z)\big]. \tag{A71}$$

Setting $b(z) = a_i^{pub}(z)$, which is the optimal policy under the public-information constraint, establishes the result

$$\mathbb{E}\big[\tilde{u}_i(a_i^{te}(z), z)\big] \geq \mathbb{E}\big[\tilde{u}_i(a_i^{pub}(z), z)\big]. \tag{A72}$$

This dominance is a direct implication of the value-of-information principle (Blackwell, 1951; Marschak and Radner, 1972): optimizing a concave objective with the strictly finer information set $\mathcal{I}_i$ cannot yield a lower expected payoff than optimizing with the coarser set $\mathcal{I}^{pub}$. The inequality is strict if $\mathcal{I}_i$ strictly refines $\mathcal{I}^{pub}$ in a way that is payoff-relevant, such that the change in conditional beliefs about the state z alters the unique optimizer on a set of positive probability. This dominance holds pointwise for any type (β_i, η_i), ensuring ex-ante unanimous support for the Coasean transfer regime. □

A.11 Proof of Proposition 11

We compare the equilibrium payoffs under the Coasean transfer regime (denoted by ct) and the Pigouvian tax regime (denoted by tax). Under the Coasean transfer, the agent's ex-post payoff is

$$u_i^{ct} = \left(e - \frac{\eta_i}{2}(a_i^{ct})^2\right) + \beta_i \theta a_i^{ct} + S(\Omega^{ct}) + (a_i^{ct} - A^{ct})p(\Omega^{ct}), \tag{A73}$$

where the pricing function is $p(\Omega) = \kappa S'(\Omega)$. Since agents are atomless, the FOC is

$$\mathbb{E}_i\big[-\eta_i a_i^{ct} + \beta_i \theta + p(\Omega^{ct})\big] = 0. \tag{A74}$$

Under the Pigouvian regime, the agent faces a uniform tax τ and receives a lump-sum rebate proportional to the aggregate A^{tax}, such that the net transfer is $-\tau(a_i^{tax} - A^{tax})$. The payoff is

$$u_i^{tax} = \left(e - \frac{\eta_i}{2}(a_i^{tax})^2\right) + \beta_i \theta a_i^{tax} + S(\Omega^{tax}) + (a_i^{tax} - A^{tax})(-\tau). \tag{A75}$$

The FOC in the tax regime is $-\eta_i a_i^{tax} + \beta_i m_{i,\theta} - \tau = 0$, which implies

$$\eta_i a_i^{tax} = \beta_i m_{i,\theta} - \tau. \tag{A76}$$

Let $\Delta u_i = u_i^{ct} - u_i^{tax}$. We define the differences in allocations as $\Delta a_i = a_i^{ct} - a_i^{tax}$ and $\Delta A = A^{ct} - A^{tax}$. Under Assumption 2, the marginal social value is linear, $S'(\Omega) = s_1 - s_2(\Omega - \mu_\Omega^{sq})$, which implies

$S''(\Omega) = -s_2$. Consequently, the price difference is linear in the aggregate difference:

$$p(\Omega^{ct}) - p(\Omega^{tax}) = \kappa[S'(\Omega^{ct}) - S'(\Omega^{tax})] = -\kappa^2 s_2 \Delta A. \tag{A77}$$

We proceed by expanding the payoff difference Δu_i. The difference in the private cost term is $-\frac{\eta_i}{2}[(a_i^{ct})^2 - (a_i^{tax})^2] = -\eta_i a_i^{tax}\Delta a_i - \frac{\eta_i}{2}(\Delta a_i)^2$. The difference in social value, using the exact second-order expansion for the quadratic function $S(\cdot)$ around Ω^{tax}, is $S(\Omega^{ct}) - S(\Omega^{tax}) = S'(\Omega^{tax})(\kappa\Delta A) - \frac{s_2}{2}(\kappa\Delta A)^2$.

Next, we decompose the difference in net transfers, $(a_i^{ct} - A^{ct})p(\Omega^{ct}) - (a_i^{tax} - A^{tax})(-\tau)$. Substituting $a_i^{ct} = a_i^{tax} + \Delta a_i$, $A^{ct} = A^{tax} + \Delta A$, and $p(\Omega^{ct}) = p(\Omega^{tax}) - \kappa^2 s_2 \Delta A$, we obtain:

$$(a_i^{tax} - A^{tax})(p(\Omega^{tax}) + \tau) + p(\Omega^{tax})(\Delta a_i - \Delta A) - \kappa^2 s_2 (a_i^{ct} - A^{ct})\Delta A.$$

Adding the difference in private benefits $\beta_i\theta\Delta a_i$ and collecting all terms, we note that the social value term $\kappa S'(\Omega^{tax})\Delta A = p(\Omega^{tax})\Delta A$ cancels with the transfer term $-p(\Omega^{tax})\Delta A$. The total difference is:

$$\begin{aligned}\Delta u_i = {} & (a_i^{tax} - A^{tax})(p(\Omega^{tax}) + \tau) \\ & + \big[-\eta_i a_i^{tax} + \beta_i\theta + p(\Omega^{tax})\big]\Delta a_i \\ & - \kappa^2 s_2 (a_i^{tax} - A^{tax})\Delta A \\ & - \frac{\eta_i}{2}(\Delta a_i)^2 - \kappa^2 s_2 \Delta a_i \Delta A + \frac{\kappa^2 s_2}{2}(\Delta A)^2.\end{aligned} \tag{A78}$$

The first term represents the change in the value of the baseline position $(a_i^{tax} - A^{tax})$. The second line captures the first-order effect of the individual adjustment Δa_i, evaluated at the quantities and prices of the tax regime. The remaining lines capture the second-order and interaction effects.

To simplify (A78), we use the FOCs to eliminate the linear term in Δa_i. Evaluating the expectation of the Coasean FOC (A74) and substituting $a_i^{ct} = a_i^{tax} + \Delta a_i$ and $p(\Omega^{ct}) = p(\Omega^{tax}) - \kappa^2 s_2 \Delta A$, we obtain

$$\begin{aligned}0 &= \mathbb{E}_i\big[-\eta_i(a_i^{tax} + \Delta a_i) + \beta_i\theta + p(\Omega^{tax}) - \kappa^2 s_2 \Delta A\big] \\ &= \mathbb{E}_i\big[-\eta_i a_i^{tax} + \beta_i\theta + p(\Omega^{tax})\big] - \eta_i \Delta a_i - \kappa^2 s_2 \mathbb{E}_i[\Delta A].\end{aligned} \tag{A79}$$

Rearranging this equation provides an expression for the first-order terms:

$$\mathbb{E}_i\big[-\eta_i a_i^{tax} + \beta_i\theta + p(\Omega^{tax})\big] = \eta_i \Delta a_i + \kappa^2 s_2 \mathbb{E}_i[\Delta A]. \tag{A80}$$

We now take the conditional expectation $\mathbb{E}_i[\cdot]$ of the expanded payoff difference. We substitute the expression for the linear first-order terms derived above into the term multiplying Δa_i. Note that Δa_i is known to agent i, so it can be factored out of the expectation. Specifically, the term involving Δa_i becomes:

$$\begin{aligned} \mathbb{E}_i[[-\eta_i a_i^{tax} + \beta_i\theta + p(\Omega^{tax})]\Delta a_i] - \frac{\eta_i}{2}(\Delta a_i)^2 \\ = \Delta a_i \left(\eta_i \Delta a_i + \kappa^2 s_2 \mathbb{E}_i[\Delta A]\right) - \frac{\eta_i}{2}(\Delta a_i)^2 \\ = \frac{\eta_i}{2}(\Delta a_i)^2 + \kappa^2 s_2 \Delta a_i \mathbb{E}_i[\Delta A]. \end{aligned} \tag{A81}$$

This substitution yields a cross term $\kappa^2 s_2 \Delta a_i \mathbb{E}_i[\Delta A]$, which cancels exactly with the expectation of the interaction term $-\kappa^2 s_2 \Delta a_i \Delta A$ in (A78). Substituting this back into the expected payoff difference, and noting that the cross terms cancel, we are left with:

$$\begin{aligned} \mathbb{E}_i[\Delta u_i] = \frac{1}{2}\eta_i(\Delta a_i)^2 + \frac{1}{2}\kappa^2 s_2 \mathbb{E}_i[(\Delta A)^2] \\ + \mathbb{E}_i\left[(a_i^{tax} - A^{tax})(p(\Omega^{tax}) + \tau)\right] - \mathbb{E}_i\left[\kappa^2 s_2 (a_i^{tax} - A^{tax})\Delta A\right]. \end{aligned} \tag{A82}$$

Finally, we combine the last two terms. Using the price relation from (A77), we can write $p(\Omega^{tax}) - \kappa^2 s_2 \Delta A = p(\Omega^{ct})$. Thus:

$$\mathbb{E}_i\left[(a_i^{tax} - A^{tax})(p(\Omega^{tax}) - \kappa^2 s_2 \Delta A + \tau)\right] = \mathbb{E}_i\left[(a_i^{tax} - A^{tax})(p(\Omega^{ct}) + \tau)\right]. \tag{A83}$$

This yields the result stated in equation (39):

$$\mathbb{E}_i[\Delta u_i] = \frac{1}{2}\eta_i(\Delta a_i)^2 + \frac{1}{2}\kappa^2 s_2 \mathbb{E}_i[(\Delta A)^2] + \mathbb{E}_i\left[(a_i^{tax} - A^{tax})(p(\Omega^{ct}) + \tau)\right]. \tag{A84}$$

Since $\eta_i > 0$ and $s_2 > 0$, the first two terms are non-negative, representing the efficiency gains from individual flexibility and collective adaptability, while the final term represents the expected net transfer difference. □

A.12 Proof of Proposition 12

We first establish the coefficients for the first-best aggregate action A^{fb}. From Proposition 1, A^{fb} satisfies the fixed-point equation (9): $A^{fb} = \bar{B}_H\theta + \frac{\kappa}{H}S'(\kappa A^{fb} + \Gamma(\theta,\nu))$. Under Assumptions 2 and 3, this becomes the linear equation $HA^{fb} = H\bar{B}_H\theta + \kappa[s_1 - s_2(\kappa A^{fb} + \theta + \nu - \mu_\Omega^{sq})]$. Using $\alpha \equiv -\kappa^2 s_2$ and substituting $-\kappa s_2 = \alpha/\kappa$, we collect terms on A^{fb} to obtain $A^{fb}(H - \alpha) = (H\bar{B}_H + \alpha/\kappa)\theta + (\alpha/\kappa)\nu + \kappa s_1 - (\alpha/\kappa)\mu_\Omega^{sq}$. Thus, the first-best aggregate action is $A^{fb} = J_\theta\theta + J_\nu\nu + J_0$, with coefficients

$$J_\theta \equiv \frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}, \quad J_\nu \equiv \frac{\alpha/\kappa}{H-\alpha}, \quad \text{and} \quad J_0 \equiv \frac{\kappa s_1 - (\alpha/\kappa)\mu_\Omega^{sq}}{H-\alpha}. \tag{A85}$$

We now solve for the team-efficient aggregate action A^{te}, postulating the linear form $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$. Recall from the proof of Proposition 4 (equation (A16)) that the individual actions satisfy $\eta_i(a_i^{te} - \mathbb{E}_i[a_i^{fb}]) = \alpha\mathbb{E}_i[A^{te} - A^{fb}]$. Dividing by η_i and integrating over the population $i \in [0,1]$ yields

$$A^{te} - \int_0^1 \mathbb{E}_i[a_i^{fb}]di = \frac{\alpha}{H}\int_0^1 (\mathbb{E}_i[A^{te}] - \mathbb{E}_i[A^{fb}])di. \tag{A86}$$

We compute the aggregate expectations using the affine information structure from Assumption 3. For any linear variable $X = K_\theta\theta + K_\nu\nu + K_0$, the cross-sectional average of expectations is $\int_0^1 \mathbb{E}_i[X]di = K_\theta(\lambda_\theta\theta + (1-\lambda_\theta)\mu_\theta) + K_\nu(\lambda_\nu\nu + (1-\lambda_\nu)\mu_\nu) + K_0$. Also, since the first-best action a_i^{fb} is linear in the state, and agent types are independent of signals, the aggregate expectation $\int_0^1 \mathbb{E}_i[a_i^{fb}]di$ inherits the same coefficients J as A^{fb}, but applied to the average posterior beliefs. Applying this to the difference equation above, we equate the coefficients on the fundamental θ:

$$C_\theta - J_\theta\lambda_\theta = \frac{\alpha}{H}(C_\theta\lambda_\theta - J_\theta\lambda_\theta). \tag{A87}$$

Rearranging to solve for C_θ gives

$$C_\theta\left(1 - \frac{\alpha\lambda_\theta}{H}\right) = J_\theta\lambda_\theta\left(1 - \frac{\alpha}{H}\right) \quad \Longrightarrow \quad C_\theta = J_\theta\lambda_\theta\frac{H-\alpha}{H-\alpha\lambda_\theta}. \tag{A88}$$

Substituting the expression for J_θ, the term $(H-\alpha)$ in the numerator cancels with the denominator of J_θ, yielding

$$C_\theta = \frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}\lambda_\theta\frac{H-\alpha}{H-\alpha\lambda_\theta} = \frac{(H\bar{B}_H + \alpha/\kappa)\lambda_\theta}{H-\alpha\lambda_\theta}. \tag{A89}$$

Similarly, equating coefficients for ν gives $C_\nu(1-\alpha\lambda_\nu/H) = J_\nu\lambda_\nu(1-\alpha/H)$, which leads to

$$C_\nu = J_\nu\lambda_\nu\frac{H-\alpha}{H-\alpha\lambda_\nu} = \frac{\alpha/\kappa}{H-\alpha}\lambda_\nu\frac{H-\alpha}{H-\alpha\lambda_\nu} = \frac{(\alpha/\kappa)\lambda_\nu}{H-\alpha\lambda_\nu}. \tag{A90}$$

Finally, equating the constant terms: constants are known, so expectations are identity.

$$C_0 - J_0 = \frac{\alpha}{H}(C_0 - J_0) \quad \Longrightarrow \quad C_0\left(1-\frac{\alpha}{H}\right) = J_0\left(1-\frac{\alpha}{H}\right). \tag{A91}$$

Since $\alpha/H \neq 1$, this implies $C_0 = J_0 = \frac{\kappa s_1 - (\alpha/\kappa)\mu_\Omega^{sq}}{H-\alpha}$.

To verify the comparative statics with respect to information precision, consider the magnitude $|C_\theta|$. Let $N \equiv |H\bar{B}_H + \alpha/\kappa|$ and $D \equiv -\alpha > 0$ (since $\alpha = -\kappa^2 s_2 < 0$). Then $|C_\theta| = \frac{N\lambda_\theta}{H+D\lambda_\theta}$. Differentiating with respect to λ_θ yields $\frac{\partial |C_\theta|}{\partial \lambda_\theta} = \frac{N(H+D\lambda_\theta)-N\lambda_\theta D}{(H+D\lambda_\theta)^2} = \frac{NH}{(H+D\lambda_\theta)^2}$, which is strictly positive provided $N > 0$. The same logic applies to $|C_\nu|$ with $N = |\alpha/\kappa| > 0$. Thus, the aggregate action becomes strictly more sensitive to the state as information precision λ_θ or λ_ν increases. □

A.13 Proof of Proposition 13

We solve for the team-efficient aggregate action A^{te} under the condition that agents distrust the signals of others. We conjecture and verify the linear form $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$.

Recall the strategic form of the team-efficient action from Proposition 4:

$$\eta_i a_i^{te} = \eta_i \left(1 - \frac{\alpha}{H}\right) \mathbb{E}_i[a_i^{fb}] + \alpha \mathbb{E}_i[A^{te}] - \alpha \left(\bar{B}_H - \frac{\beta_i}{H}\right) m_{i,\theta}. \tag{A92}$$

Aggregating across agents yields the fixed-point equation for the aggregate action:

$$A^{te} = \left(1 - \frac{\alpha}{H}\right) \int_0^1 \mathbb{E}_i[a_i^{fb}]di + \frac{\alpha}{H} \int_0^1 \mathbb{E}_i[A^{te}]di. \tag{A93}$$

The first term involves the expectation of the first-best action. Since the first-best allocation depends on the true state and agents form rational expectations about their own payoffs (distrust only applies to the *strategic* component involving others), the aggregation of these expectations remains the same as in the benchmark case with no distrust. Specifically, for the coefficient on θ:

$$\text{Coeff}_\theta \left[\int_0^1 \mathbb{E}_i[a_i^{fb}]di\right] = J_\theta \lambda_\theta. \tag{A94}$$

The second term involves the expectation of the aggregate action A^{te}. Under distrust, agent i believes that the aggregate action is less responsive to the fundamentals than it truly is. Recall from equation (43) that agent i perceives the signals of others to be loaded on θ with coefficient φ_θ, whereas true signals have a loading of 1. Since individual actions are linear in signals, the aggregation of these perceived signals implies that agent i perceives the aggregate response to θ to be damped by φ_θ. Thus, if the true coefficient is C_θ, the perceived coefficient is $C_\theta\varphi_\theta$. Formally, since the equilibrium is linear, we isolate the component of the aggregate action that depends on θ. Agent i calculates the expectation of this component as $\mathbb{E}_i[C_\theta\varphi_\theta\theta] = C_\theta\varphi_\theta m_{i,\theta}$. Aggregating these expectations across all agents yields:

$$\int_0^1 \mathbb{E}_i[C_\theta \varphi_\theta \theta]di = C_\theta \varphi_\theta \int_0^1 m_{i,\theta} di. \tag{A95}$$

Using the affine information structure from Assumption 3, the average belief is $\int_0^1 m_{i,\theta} di = \lambda_\theta \theta + (1-\lambda_\theta)\mu_\theta$. Substituting this into the integral shows that the perceived aggregate response to the fundamental scales with $\lambda_\theta \varphi_\theta$. Multiplying by the strategic weight α/H, the coefficient on θ for the second term is:

$$\text{Coeff}_\theta \left[\frac{\alpha}{H} \int_0^1 \mathbb{E}_i[A^{te}]di \right] = \frac{\alpha}{H} C_\theta \varphi_\theta \lambda_\theta. \tag{A96}$$

Substituting these coefficients into the aggregate equation (A93) gives the condition for C_θ:

$$C_\theta = \left(1 - \frac{\alpha}{H}\right) J_\theta \lambda_\theta + \frac{\alpha}{H} C_\theta \lambda_\theta \varphi_\theta. \tag{A97}$$

Rearranging to solve for C_θ:

$$C_\theta \left(1 - \frac{\alpha \lambda_\theta \varphi_\theta}{H}\right) = J_\theta \lambda_\theta \left(1 - \frac{\alpha}{H}\right). \tag{A98}$$

Using the definition of the first-best coefficient $J_\theta = \frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}$, we have:

$$C_\theta \frac{H - \alpha\lambda_\theta\varphi_\theta}{H} = \left(\frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}\right) \lambda_\theta \frac{H-\alpha}{H} = \frac{(H\bar{B}_H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta\varphi_\theta}. \tag{A99}$$

The derivation for C_ν follows similar steps. The first-best coefficient is $J_\nu = \frac{\alpha/\kappa}{H-\alpha}$. The distrust parameter φ_ν damps the perceived strategic response to ν. The condition for C_ν is similar to (A97):

$$C_\nu \left(1 - \frac{\alpha \lambda_\nu \varphi_\nu}{H}\right) = J_\nu \lambda_\nu \left(1 - \frac{\alpha}{H}\right). \tag{A100}$$

Solving for C_ν yields:

$$C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu\varphi_\nu}. \tag{A101}$$

Finally, for the constant term C_0, the distrust mechanism does not apply because C_0 is a known constant, not a random variable subject to private information or signal extraction. Therefore, the condition is identical to the rational expectations case:

$$C_0\left(1-\frac{\alpha}{H}\right)=J_0\left(1-\frac{\alpha}{H}\right)\implies C_0=J_0=\frac{\kappa s_1-(\alpha/\kappa)\mu_\Omega^{sq}}{H-\alpha}.\quad\square \tag{A102}$$

A.14 Proof of Proposition 14

We first establish why the incentive to acquire information is determined by the square of the action's responsiveness. The agent's objective functions in the three regimes are given by:

$$\text{Status Quo (eq. (2)):}\qquad u_i^{sq}=\mathbb{E}_i\left[\left(e-\frac{\eta_i}{2}a_i^2\right)+\beta_i\theta a_i+S(\Omega)\right], \tag{A103}$$

$$\text{Pigouvian Tax (eq. (29)):}\qquad u_i^{tax}=\mathbb{E}_i\left[\left(e-\frac{\eta_i}{2}a_i^2\right)+\beta_i\theta a_i+S(\Omega)-\tau a_i+R\right], \tag{A104}$$

$$\text{Coasean Transfer (eq. (19)):}\qquad u_i^{ct}=\mathbb{E}_i\left[\left(e-\frac{\eta_i}{2}a_i^2\right)+\beta_i\theta a_i+S(\Omega)+(a_i-A)p(\Omega)\right]. \tag{A105}$$

In all three regimes, agent i chooses a_i to maximize an objective of the general form:

$$u_i=\mathbb{E}_i\left[-\frac{\eta_i}{2}a_i^2+a_iM_i\right]+C_i \tag{A106}$$

where M_i is the regime-specific marginal benefit (dependent on θ and potentially Ω) and C_i contains constant terms independent of a_i. Specifically: in the Status Quo regime, $M_i=\beta_i\theta$ and $C_i=e+\mathbb{E}_i[S(\Omega)]$; in the Pigouvian Tax regime, $M_i=\beta_i\theta-\tau$ and $C_i=e+\mathbb{E}_i[S(\Omega)+R]$; and in the Coasean Transfer regime, $M_i=\beta_i\theta+p(\Omega)$ and $C_i=e+\mathbb{E}_i[S(\Omega)-Ap(\Omega)]$.

The first-order condition is $\eta_ia_i=\mathbb{E}_i[M_i]$. To evaluate the ex-ante expected payoff $\mathbb{E}[u_i]$, we substitute the optimal action back into the objective. Because we evaluate the payoff for a fixed agent i, the parameter η_i is constant with respect to the expectation. Considering the terms that depend on a_i, we have:

$$\mathbb{E}\left[-\frac{\eta_i}{2}a_i^2+a_iM_i\right]=-\frac{\eta_i}{2}\mathbb{E}[a_i^2]+\mathbb{E}[a_iM_i] \tag{A107}$$

By the Law of Iterated Expectations, $\mathbb{E}[a_iM_i]=\mathbb{E}[a_i\mathbb{E}_i[M_i]]$. Using the first-order condition $\mathbb{E}_i[M_i]=\eta_ia_i$, this term becomes $\mathbb{E}[a_i(\eta_ia_i)]=\eta_i\mathbb{E}[a_i^2]$. Substituting this back yields:

$$\mathbb{E}[u_i]=-\frac{\eta_i}{2}\mathbb{E}[a_i^2]+\eta_i\mathbb{E}[a_i^2]+\mathbb{E}[C_i]=\frac{\eta_i}{2}\mathbb{E}[a_i^2]+\text{Constant} \tag{A108}$$

Let the optimal action be linear in the posterior belief: $a_i=K_im_{i,\theta}+Z_i$, where Z_i is uncorrelated with $m_{i,\theta}$. The expected squared action is:

$$\mathbb{E}[a_i^2] = K_i^2 \mathbb{E}[m_{i,\theta}^2] + \text{terms independent of information} \tag{A109}$$

Using the property of variance, $\mathbb{E}[m_{i,\theta}^2] = \text{Var}(m_{i,\theta}) + (\mathbb{E}[m_{i,\theta}])^2$. The agent chooses information precision to maximize this value. Since $\text{Var}(m_{i,\theta})$ increases with information precision (as the posterior belief moves more away from the prior), the marginal benefit of acquiring information is directly proportional to the coefficient K_i^2. Thus, to compare incentives, we must compare the squared responsiveness K_i^2 across regimes.

Status Quo: The agent maximizes $\max_{a_i} \mathbb{E}_i[(e - \frac{\eta_i}{2} a_i^2) + \beta_i \theta a_i]$. The FOC is:

$$\eta_i a_i = \beta_i \mathbb{E}_i[\theta] \implies a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} \tag{A110}$$

The responsiveness is the derivative with respect to the signal $m_{i,\theta}$:

$$K_i^{sq} = \frac{\beta_i}{\eta_i} \tag{A111}$$

Pigouvian Tax: The agent faces a constant tax τ set ex-ante. The FOC becomes:

$$\eta_i a_i = \beta_i \mathbb{E}_i[\theta] - \tau \implies a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} - \frac{\tau}{\eta_i} \tag{A112}$$

Since τ is constant and does not depend on the specific realization of the agent's signal $m_{i,\theta}$, the derivative remains unchanged:

$$K_i^{tax} = \frac{\beta_i}{\eta_i} = K_i^{sq} \tag{A113}$$

Coasean Transfer: We derive K_i^{ct} by differentiating the strategic form of the team-efficient action (Proposition 4) with respect to $m_{i,\theta}$.

$$a_i^{ct} = \left(1 - \frac{\alpha}{H}\right) \mathbb{E}_i[a_i^{fb}] + \frac{\alpha}{\eta_i} \mathbb{E}_i[A^{ct}] - \frac{\alpha}{\eta_i} \left(\bar{B}_H - \frac{\beta_i}{H}\right) m_{i,\theta} \tag{A114}$$

Differentiating yields:

$$K_i^{ct} = \left(1 - \frac{\alpha}{H}\right) \frac{\partial \mathbb{E}_i[a_i^{fb}]}{\partial m_{i,\theta}} + \frac{\alpha}{\eta_i} \frac{\partial \mathbb{E}_i[A^{ct}]}{\partial m_{i,\theta}} - \frac{\alpha}{\eta_i} \left(\bar{B}_H - \frac{\beta_i}{H}\right) \tag{A115}$$

We evaluate the partial derivatives. For the First-Best action, recall from Proposition 3 that $a_i^{fb} = \eta_i^{-1}(\beta_i \theta + \kappa[s_1 - s_2(\Omega^{fb} - \mu_\Omega^{sq})])$. Using $\alpha \equiv -\kappa^2 s_2$, the term $-\kappa s_2$ corresponds to α/κ. The aggregate outcome Ω^{fb} is linear in θ, with the aggregate action sensitivity J_θ defined in the proof of Proposition 12. Thus, the sensitivity of the outcome is $1 + \kappa J_\theta$. The derivative is:

$$\frac{\partial \mathbb{E}_i[a_i^{fb}]}{\partial m_{i,\theta}} = \frac{1}{\eta_i}\left[\beta_i + \frac{\alpha}{\kappa}(1 + \kappa J_\theta)\right] = \frac{1}{\eta_i}\left(\beta_i + \frac{\alpha}{\kappa} + \alpha J_\theta\right) \tag{A116}$$

For the aggregate action, under the affine structure, $A^{ct} = C_\theta \theta + \ldots$, so $\mathbb{E}_i[A^{ct}] = C_\theta m_{i,\theta} + \ldots$, which implies:

$$\frac{\partial \mathbb{E}_i[A^{ct}]}{\partial m_{i,\theta}} = C_\theta \tag{A117}$$

Substituting these derivatives back into the expression for K_i^{ct} and factoring out $1/\eta_i$:

$$K_i^{ct} = \frac{1}{\eta_i}\left[\left(1 - \frac{\alpha}{H}\right)\left(\beta_i + \frac{\alpha}{\kappa} + \alpha J_\theta\right) + \alpha C_\theta - \alpha\left(\bar{B}_H - \frac{\beta_i}{H}\right)\right] \tag{A118}$$

We simplify the terms inside the square brackets. First, we group the terms involving β_i:

$$\left(1 - \frac{\alpha}{H}\right)\beta_i + \frac{\alpha\beta_i}{H} = \beta_i - \frac{\alpha\beta_i}{H} + \frac{\alpha\beta_i}{H} = \beta_i \tag{A119}$$

Next, we group the remaining constant terms:

$$\text{Const} = \left(1 - \frac{\alpha}{H}\right)\left(\frac{\alpha}{\kappa} + \alpha J_\theta\right) - \alpha \bar{B}_H \tag{A120}$$

Substituting the definition $J_\theta = \frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}$ and using $1 - \frac{\alpha}{H} = \frac{H-\alpha}{H}$:

$$\text{Const} = \frac{H-\alpha}{H}\alpha\left(\frac{1}{\kappa} + \frac{H\bar{B}_H + \alpha/\kappa}{H-\alpha}\right) - \alpha\bar{B}_H \tag{A121}$$

$$= \frac{\alpha(H-\alpha)}{H}\left(\frac{H - \alpha + \kappa H\bar{B}_H + \alpha}{\kappa(H-\alpha)}\right) - \alpha\bar{B}_H \tag{A122}$$

$$= \frac{\alpha}{H}\left(\frac{H(1+\kappa\bar{B}_H)}{\kappa}\right) - \alpha\bar{B}_H \tag{A123}$$

$$= \frac{\alpha}{\kappa} + \alpha\bar{B}_H - \alpha\bar{B}_H = \frac{\alpha}{\kappa} \tag{A124}$$

Combining the simplified β_i term, the constant term α/κ, and the aggregate response αC_θ, we obtain:

$$K_i^{ct} = \frac{1}{\eta_i}\left(\beta_i + \alpha C_\theta + \frac{\alpha}{\kappa}\right) = \frac{1}{\eta_i}\left[\beta_i + \alpha\left(C_\theta + \frac{1}{\kappa}\right)\right] \tag{A125}$$

The incentive is strictly higher under CT if and only if $(K_i^{ct})^2 > (K_i^{sq})^2$:

$$\left(\frac{\beta_i + \alpha(C_\theta + 1/\kappa)}{\eta_i}\right)^2 > \left(\frac{\beta_i}{\eta_i}\right)^2. \tag{A126}$$

Since $\eta_i > 0$, this inequality is equivalent to:

$$\left[\alpha\left(C_\theta + \frac{1}{\kappa}\right)\right] \cdot \left[2\beta_i + \alpha\left(C_\theta + \frac{1}{\kappa}\right)\right] > 0. \tag{A127}$$

We determine the sign of the term $Z \equiv (C_\theta + 1/\kappa)$. Substituting C_θ from Proposition 12:

$$Z = \frac{1}{\kappa} + \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta} = \frac{H - \alpha\lambda_\theta + \kappa\bar{B}_H H\lambda_\theta + \alpha\lambda_\theta}{\kappa(H - \alpha\lambda_\theta)} = \frac{H(1 + \kappa\bar{B}_H\lambda_\theta)}{\kappa(H - \alpha\lambda_\theta)}. \tag{A128}$$

Since $H, \kappa, \bar{B}_H > 0$, $\lambda_\theta \in [0, 1]$, and $\alpha < 0$, the numerator and denominator are strictly positive. Thus $Z > 0$. Given $Z > 0$ and $\alpha < 0$, the first bracket $[\alpha Z]$ is strictly negative. For the inequality to hold, the second bracket must also be strictly negative:

$$2\beta_i + \alpha Z < 0 \iff \beta_i < \frac{-\alpha}{2}Z. \tag{A129}$$

Substituting the expression for Z derived above yields the condition stated in the proposition:

$$\beta_i < \frac{-\alpha}{2}\left[\frac{H(1 + \kappa\bar{B}_H\lambda_\theta)}{\kappa(H - \alpha\lambda_\theta)}\right] = \frac{-\alpha H(1 + \kappa\bar{B}_H\lambda_\theta)}{2\kappa(H - \alpha\lambda_\theta)}. \quad \square \tag{A130}$$

A.15 Proof of Proposition 15

As established in the proof of Proposition 14, the agent's ex-ante expected payoff in any linear-quadratic regime is convex in the optimal action, taking the form in equation (A108), $\mathbb{E}[u_i] = \text{Constant} + \frac{\eta_i}{2}\mathbb{E}[a_i^2]$. If the optimal action is linear in the posterior belief $m_{i,\nu}$, such that $a_i = K_{i,\nu}m_{i,\nu} + Z_i$ (with Z_i uncorrelated with $m_{i,\nu}$), then the marginal value of information regarding ν is proportional to $(K_{i,\nu})^2\text{Var}[m_{i,\nu}]$. The value of information is strictly positive if and only if the action responsiveness $K_{i,\nu} \equiv \frac{\partial a_i}{\partial m_{i,\nu}}$ is non-zero.

Status Quo and Pigouvian Tax. In the status quo, the optimal action is $a_i^{sq} = \frac{\beta_i}{\eta_i}m_{i,\theta}$. Under a Pigouvian tax, the action is $a_i^{tax} = \frac{\beta_i m_{i,\theta} - \tau}{\eta_i}$. In both cases, the action depends only on the posterior regarding the fundamental θ and fixed parameters. The action is mathematically independent of the posterior regarding the noise term ν. Thus, $K_{i,\nu}^{sq} = K_{i,\nu}^{tax} = 0$, and the marginal value of information about ν is zero.

Coasean Transfer. The agent's FOC under the Coasean transfer is

$$\eta_i a_i^{ct} = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega^{ct})]. \tag{A131}$$

Using $p(\Omega) = \kappa S'(\Omega)$ and $\Omega = \kappa A + \theta + \nu$, the derivative of the expected price with respect to $m_{i,\nu}$ is:

$$\frac{\partial}{\partial m_{i,\nu}} \mathbb{E}_i[p(\Omega^{ct})] = \kappa S''(\Omega) \frac{\partial \mathbb{E}_i[\Omega]}{\partial m_{i,\nu}} = -\kappa s_2 \left(\kappa \frac{\partial \mathbb{E}_i[A^{ct}]}{\partial m_{i,\nu}} + 1 \right). \tag{A132}$$

Substituting $\alpha \equiv -\kappa^2 s_2$, we can write $-\kappa s_2 = \alpha/\kappa$. From Proposition 12, the aggregate action is fully defined as $A^{ct} = C_0 + C_\theta \theta + C_\nu \nu$. The agent's expectation is therefore $\mathbb{E}_i[A^{ct}] = C_0 + C_\theta m_{i,\theta} + C_\nu m_{i,\nu}$. Since C_0 and $m_{i,\theta}$ are independent of $m_{i,\nu}$, the partial derivative is $\frac{\partial \mathbb{E}_i[A^{ct}]}{\partial m_{i,\nu}} = C_\nu$. The responsiveness $K_{i,\nu}^{ct}$ is therefore:

$$\eta_i K_{i,\nu}^{ct} = \frac{\alpha}{\kappa}(\kappa C_\nu + 1). \tag{A133}$$

Using the expression $C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu}$ from Proposition 12, we have:

$$1 + \kappa C_\nu = 1 + \kappa \left(\frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu} \right) = 1 + \frac{\alpha\lambda_\nu}{H - \alpha\lambda_\nu} = \frac{H - \alpha\lambda_\nu + \alpha\lambda_\nu}{H - \alpha\lambda_\nu} = \frac{H}{H - \alpha\lambda_\nu}. \tag{A134}$$

Since $H > 0$ and $\alpha < 0$, the denominator is strictly positive, so the term is non-zero. Thus, $K_{i,\nu}^{ct} \neq 0$. Because the agent's optimal action responds to $m_{i,\nu}$, the marginal value of information regarding ν is strictly positive. □