The Quiet Hand of Regulation: Harnessing Uncertainty and Disagreement

Daniel Andrei* Lorenzo Garlappi†

February 11, 2026

Abstract

We study economies with uncertainty and dispersed private information in which externalities arise because individual actions jointly determine an observable aggregate outcome. First, we characterize outcome-contingent "Coasean transfers" that restore efficiency by creating a synthetic market for the externality. Agents' deviations from the aggregate action are priced according to a schedule contingent on the realized outcome; the resulting prices coincide with equilibrium prices in a hypothetical competitive market for the externality and decentralize the team-efficient allocation. Next, we characterize the informational properties of Coasean transfers. Because the transfers' pricing schedule depends only on the realized outcome, dispersed private information is aggregated without communication, bargaining, or state verification, in contrast to standard quantity- and price-based policy instruments. Uncertainty and disagreement strengthen information aggregation. Finally, we study the political viability of Coasean transfers and illustrate the approach in the knowledge-spillover model of Romer (1986) and the tragedy-of-the-commons model of Gordon (1954).

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, Theodosios Dimopoulos, Vadim Di Pietro, Giada Durante, Paul Ehling, Jack Favilukis, Paolo Fulghieri, Andrea Gamba, Jon A. Garfinkel, Ron Giammarino, Paolo Giordani, Will Gornall, Vincent Grégoire, Peter G. Hansen, Christian Heyerdahl-Larsen, Philipp Illeditsch, Eric Jondeau, 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, Anton Tsoy, Sherry Xue, Maximilian Voigt, Masa Watanabe, Gregory Weitzner, Liyan Yang, 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, the Fall 2025 FTG meeting, the AFA 2026 meeting, the University of Alberta, and HEC Lausanne 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

An externality reflects a missing market (Meade, 1952; Arrow, 1969). When agents cannot trade the effects they impose on others, equilibrium allocations fail to internalize the social costs of individual actions. In economies without uncertainty and with complete information, the remedy is well understood: a regulator can impose a Pigouvian tax equal to the social cost (Pigou, 1920), or agents can bargain to efficiency under well-defined property rights (Coase, 1960). Both approaches operate by creating the missing market.

These remedies rely on informational conditions that are rarely satisfied. Pigouvian taxation requires knowledge of the social cost of the externality, which may be difficult or impossible for the regulator to obtain when information is dispersed across agents. As Weitzman (1974) formalized, uncertainty forces the regulator into a second-best choice between prices and quantities. Coasean bargaining becomes infeasible when agents are numerous. Even with complete information, Ellingsen and Paltseva (2016) show that voluntary negotiation breaks down in large groups due to strategic non-participation. Dispersed information further compounds this breakdown, rendering decentralized trade impractical.

We show that uncertainty and disagreement need not be obstacles to efficiency. Instead, they can be exploited through outcome-contingent transfers that create a *synthetic market* for the externality. The regulator need not observe the underlying fundamentals, nor do agents need to reveal their private information. Rather, the regulator commits to a price schedule contingent on an observable aggregate outcome, and agents act based on their own information. Conditioning transfers on the realized aggregate outcome induces agents to internalize the social value of their actions while dispersed private information is aggregated through the market equilibrium.

We consider an economy with a continuum of agents. The allocation is decentralized through a system of *Coasean transfers*: payments or charges to each agent based on their deviation from the cross-sectional average, priced at a rate contingent on the realized aggregate outcome. Specifically, agent i receives a transfer $(a_i - A)p(\Omega)$, where a_i denotes the individual action, A the average action, and $p(\Omega)$ the outcome-contingent price. These transfers resemble

futures contracts on the externality, in which agents take positions relative to the average and settlements occur ex post at the realized price. By pricing deviations from collective behavior, the transfer assigns social value to agents' private information.

We establish a formal market equivalence. Coasean transfers decentralize the same allocation that would arise if the externality were traded in a competitive market, a benchmark we refer to as the *Coasean ideal*. In this hypothetical market, agents supply socially beneficial actions, with the equilibrium price determined by society's marginal valuation of the aggregate outcome. The optimal pricing schedule $p(\Omega)$ replicates this price. Formally, it coincides with the Lindahl price of a public good (Lindahl, 1958; Foley, 1970; Bonnisseau, del Mercato, and Siconolfi, 2023). Under Lindahl pricing, agents face personalized per-unit prices reflecting their valuations, yet demand the same quantity of the public good. The resulting price ensures that each agent internalizes the social value of their marginal contribution. This equivalence clarifies why Coasean transfers are effective: they unify Pigouvian internalization of externalities with the Coasean logic of decentralized adjustment through a single pricing rule.

This equivalence yields three implications. First, the equilibrium allocation is *team-efficient*: it maximizes social welfare subject to dispersed private information that cannot be communicated to a center (Radner, 1962). Coasean transfers attain this frontier without reports, negotiation, or centralized information processing.

Second, the allocation satisfies a *separation property*. The regulator need only know society's valuation of the aggregate outcome and how that outcome responds to collective actions. The regulator need not observe the distribution of fundamental uncertainty, the structure of agents' private signals, or heterogeneity in agent types; all private information remains decentralized. As in a competitive market, the outcome-contingent price schedule acts as a sufficient statistic encoding the information relevant for efficient pricing.

Third, Coasean transfers strictly dominate both Pigouvian taxation and cap-and-trade, escaping the classical price-versus-quantity trade-off (Weitzman, 1974) by exploiting the very friction that activates it: dispersed private information. The mechanism nests these standard instruments and the first best as limiting cases. When private information vanishes

and agents rely solely on common priors, the transfers replicate the optimal tax or quantity cap. When information is perfect and agents observe the fundamental state, the allocation converges to the first best, whereas standard instruments do not. The welfare gain from Coasean transfers therefore measures the *social value of private information*: the surplus generated by aggregating dispersed knowledge that neither instrument can exploit.

That Coasean transfers extract the social value of private information highlights a broader point about information aggregation. Hayek (1945) argued that markets are indispensable because prices encode dispersed knowledge that no central authority could collect. Our analysis suggests a more nuanced interpretation. The schedule $p(\Omega)$ is not a market price formed through trading, but a rule fixed ex ante; yet it replicates the informational role of a market price. Private information is never centralized: it enters the aggregate outcome Ω through agents' actions and is reflected in equilibrium transfers. In this sense, regulation operates with a "Quiet Hand": it fixes a pricing rule, while decentralized actions embed dispersed information into the market equilibrium. Coasean transfers thus show that the aggregation property Hayek attributed to markets can be replicated by appropriately designed policy instruments.

Coasean transfers harness uncertainty and disagreement by reshaping equilibrium incentives. When society is averse to volatility in the aggregate outcome, the pricing rule induces strategic substitutability. Agents are rewarded for deviating from the average action and therefore place greater weight on their private signals. When uncertainty is high or agents distrust one another's information, this reliance intensifies, so that disagreement sharpens incentives and improves information aggregation in equilibrium.

These incentives extend to information acquisition. By tying payoffs to the aggregate outcome, Coasean transfers induce agents to expand their informational scope beyond purely private concerns. Agents have incentives to monitor risks that affect the social outcome, even when those risks are irrelevant for their standalone payoffs. In this way, Coasean transfers generate social value from information that would otherwise remain uncollected or unused.

We then assess the political viability of Coasean transfers by comparing them to the optimal Pigouvian tax and the optimal cap-and-trade scheme. The welfare difference decomposes

into efficiency gains and distributional transfers. Efficiency gains arise from two sources: individual flexibility, as agents tailor actions to private information rather than to a fixed tax or quantity cap, and collective adaptability, as the aggregate outcome adjusts to the realized state. These gains are strictly positive. Distributional transfers reflect heterogeneity in agents' valuations, creating winners and losers. When agents are homogeneous, these distributional effects vanish and Coasean transfers command unanimity.

Finally, to illustrate the generality of the analysis, we apply the framework to two canonical environments featuring opposite externalities. The first is a knowledge spillover environment, which generates a positive externality; the second is a tragedy of the commons environment, which generates a negative externality. We formalize these environments using the endogenous growth model of Romer (1986) and the fishery model of Gordon (1954). We extend both environments to allow for dispersed information and show that Coasean transfers implement team efficiency in each case. Thus, regardless of whether the externality is positive or negative, or whether strategic interactions exhibit complementarity or substitutability, the synthetic market created by Coasean transfers restores efficiency.

Literature. Our analysis builds on the coordination games literature (Morris and Shin, 2002; Angeletos and Pavan, 2007). We use this framework to formalize the creation of a missing market. This perspective unifies recent proposals for state-contingent instruments, such as the carbon securities in Lemoine (2024) or the indexed subsidies in Colombo, Femminis, and Pavan (2025). We show that outcome contingency is not merely a workaround for unobservable fundamentals, but the mechanism that implements Lindahl pricing. A closely related informational logic appears in Li (2023), who shows that, in a firm-level contracting environment, profit-sharing contracts aggregate dispersed private information by conditioning payoffs on a common outcome. Our transfer applies the same principle at the economy-wide level, treating externality resolution as a shared social surplus. Finally, our result that disagreement strengthens efficiency connects to work on heterogeneous beliefs (Banerjee, 2011; Dumas, Kurshev, and Uppal, 2009). We show that distrust, rather than undermining coordination, intensifies reliance on private signals and improves aggregate outcomes.

Our work departs from the literature on policy analysis in economies with dispersed information. Angeletos and Pavan (2009) characterize optimal tax wedges, extending the Pigouvian paradigm to environments in which information is privately held. That approach takes the tax instrument as given and asks how it should be adjusted to internalize externalities under informational constraints. We take a different perspective, starting from the observation that an externality reflects a missing market. The question is then which market is absent and how it can be decentralized. Coasean transfers provide an answer: they replicate the allocation that would arise if the externality were traded in a competitive market. The regulator need not observe agents' private information or the distribution of types. Information aggregation emerges endogenously in equilibrium, through agents' actions and the resulting aggregate outcome.

The use of market-based instruments to regulate externalities is a central theme in the general equilibrium literature on climate policy. Anderson and Duanmu (2025) embed quota and emission tax systems in a fully specified Arrow-Debreu economy and show that competitive equilibria exist and are constrained Pareto optimal. Their analysis provides a rigorous general equilibrium foundation for cap-and-trade systems in which market forces implement an emission target under complete information. Our analysis complements this work by extending the market-based logic to environments with dispersed private information. In such settings, we show that cap-and-trade pricing need not implement the team-efficient allocation, whereas outcome-contingent Coasean transfers do.

The rest of the paper proceeds as follows. Section 2 introduces the model. Section 3 derives the Coasean transfer, establishes its market equivalence, and analyzes the informational requirements of the approach. Section 4 shows how Coasean transfers harness uncertainty and disagreement. Section 5 characterizes the classical price-versus-quantity trade-off. Section 6 studies incentives for information acquisition. Section 7 addresses the political viability of Coasean transfers. Section 8 applies the framework to knowledge spillovers and common-pool resources. Section 9 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 $A \equiv \int_0^1 a_i di$. Agent i has a cost parameter $\eta_i > 0$ and a benefit parameter $\beta_i > 0$, which measures how sensitive the agent's private benefit is to an underlying fundamental $\theta \in \mathbb{R}$. The aggregate action and the fundamental, together with an exogenous shock $\nu \in \mathbb{R}$, determine the aggregate outcome Ω via the linear technology

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

where $\kappa > 0$ and $\Gamma(\theta, \nu)$ depends on the fundamental and the exogenous shock.

Agent i's utility consists of a private payoff component and a social-value component that depends on the aggregate outcome. Without loss of generality, all agents have a common endowment e. Agent i's utility is given by

$$u_i = \underbrace{e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i}_{\text{Private payoff}} + \underbrace{S_i(\Omega)}_{\text{Social value}} . \tag{2}$$

The term $\beta_i \theta a_i$ is the agent's private benefit from action a_i and depends on the fundamental θ. The term $\frac{\eta_i}{2} a_i^2$ is the convex private cost. The social value $S_i(\Omega)$ captures agent i's valuation of the aggregate outcome. We assume that $S_i(\cdot)$ is increasing and allow it to vary across agents.

We refer to a collection of actions $\{a_i\}_{i \in [0,1]}$ as an *allocation*. Expected social welfare, $W(\{a_i\})$, is the utilitarian sum of agents' payoffs:

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

where $\mathbb{E}[\cdot]$ denotes expectations taken with respect to the social planner's information set. Because Ω is common across agents, welfare depends only on the average social value $\bar{S}(\Omega) \equiv \int_0^1 S_i(\Omega) di$, which is increasing in Ω.

In this economy, an externality arises because agents ignore the social value of their actions when they maximize their private payoffs. The parameter κ in equation (1) governs the strength of this externality: it measures how the aggregate action translates into the socially

valued outcome Ω. Since $\kappa > 0$ and $S_i'(\Omega) > 0$, the model features a positive externality: agents do not fully internalize the social benefit of their actions. The case of a negative externality, with $S_i'(\Omega) < 0$, is analogous. For the remainder of the paper, we impose the following assumption on social preferences.

Assumption 1. *For each agent i, the social value function $S_i(\cdot)$ is twice continuously differentiable and strictly concave, i.e., $S_i''(\cdot) < 0$.*

Assumption 1 ensures that the planner's problem is well defined and that marginal social value is diminishing in the aggregate outcome Ω. Concavity implies aversion to volatility in Ω. As shown below, this property, together with $\kappa > 0$, generates strategic substitutability in agents' actions.

Timing is as follows. Nature draws the state (θ, ν) and the cross-sectional distribution of agent types $\{(\eta_i, \beta_i, S_i(\cdot))\}_{i\in[0,1]}$. At time $t = 0$, each agent i knows its type $(\eta_i, \beta_i, S_i(\cdot))$ and observes private and public information about (θ, ν). Let $\mathcal{I}_i$ denote agent i's information set, and define posterior means $m_{i,\theta} \equiv \mathbb{E}_i[\theta]$ and $m_{i,\nu} \equiv \mathbb{E}_i[\nu]$, where $\mathbb{E}_i[\cdot] \equiv \mathbb{E}[\cdot \mid \mathcal{I}_i]$. Agents then choose a_i simultaneously without observing the aggregate action A. At time $t = 1$, θ, ν, and Ω are realized and payoffs accrue.

2.1 Status Quo Allocation

We begin by characterizing the *status-quo* equilibrium. In this benchmark, agents are atomistic and maximize their private payoff, taking the social value $S_i(\Omega)$ as exogenous. Agent i chooses a_i to maximize

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

The first-order condition is $\eta_i a_i = \beta_i \mathbb{E}_i[\theta]$, so the optimal status-quo action is

$$a_i^{sq} = \frac{\beta_i}{\eta_i} m_{i,\theta}. \tag{3}$$

Aggregating across agents yields $A^{sq} = \int_0^1 \frac{\beta_i}{\eta_i} m_{i,\theta} di$. If types (η_i, β_i) are independent of posterior beliefs $m_{i,\theta}$, then

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

where $\bar{m}_\theta \equiv \int_0^1 m_{i,\theta} di$. 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 average ratio β_i / η_i. 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 the aggregate outcome $\Omega^{sq} = \kappa A^{sq} + \Gamma(\theta, \nu)$. We use the unconditional mean of this outcome as a reference benchmark, denoted μ_Ω^{sq}. By the Law of Iterated Expectations, $\mathbb{E}[\bar{m}_\theta] = \mathbb{E}[\theta]$, so

$$\mu_\Omega^{sq} \equiv \mathbb{E}[\Omega^{sq}] = \kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)]. \tag{4}$$

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

2.2 First Best Allocation

We next characterize the socially optimal allocation under complete information. A full-information planner observes the realized state (θ, ν) and the full profile of types $\{(\eta_i, \beta_i, S_i(\cdot))\}$, and then chooses an allocation $\{a_i^{fb}\}$ to maximize social welfare $W(\{a_i\})$ state by state. The following proposition characterizes the resulting first-best allocation.

Proposition 1 (First-best allocation)**.** *Under Assumption 1 and strictly convex individual cost functions, with $\eta_i > 0$ for all i, the unique first-best allocation $\{a_i^{fb}\}_{i \in [0,1]}$ satisfies, for each (θ, ν) and $i \in [0,1]$,*

$$a_i^{fb} = \frac{\beta_i}{\eta_i} \theta + \frac{\kappa}{\eta_i} \bar{S}'(\Omega^{fb}), \tag{5}$$

where $\bar{S}'(\Omega) \equiv \int_0^1 S_i'(\Omega) di$. The aggregate action $A^{fb} = \int_0^1 a_i^{fb} di$ is then the unique fixed point of the following mapping

$$A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H} \bar{S}' \left(\kappa A^{fb} + \Gamma(\theta, \nu) \right), \tag{6}$$

where $H \equiv \left(\int_0^1 \eta_i^{-1} di\right)^{-1}$ *and* $\bar{B}_H \equiv \int_0^1 \frac{\beta_i}{\eta_i} di$.

By Assumption 1 and $\eta_i > 0$, the first-best allocation is unique and characterized by the first-order condition (5). Aggregation yields the scalar fixed-point equation (6). Since $\bar{S}'(\cdot)$ is strictly decreasing, the associated fixed-point map is continuous and nonincreasing, implying a unique solution.

The first-best allocation maximizes utilitarian social welfare under complete information but is infeasible in the environment we study, as it requires the planner to observe agents' types and the realized state. This motivates the introduction of an efficiency benchmark that respects the informational constraints of the economy.

2.3 Team-Efficient Allocation

We now introduce our efficiency benchmark. Following Radner (1962) and Angeletos and Pavan (2007), we define efficiency as the allocation that solves a utilitarian planner's problem. The planner chooses decision rules $a_i(\cdot)$ mapping each agent's private information $\mathcal{I}_i$ into actions in order to maximize ex-ante expected welfare $W(\{a_i\})$. The planner is constrained by the informational structure: private information cannot be communicated or centralized. Formally, the problem is one of functional optimization over admissible decision rules.

This notion of efficiency differs from standard constrained-efficiency concepts in mechanism design (e.g., Mirrlees, 1971; Holmström and Myerson, 1983), which assume costless communication and focus on incentive constraints. In contrast, the planner here faces an *informational constraint* reflecting dispersed knowledge in the sense of Hayek (1945). The planner does not observe the state (θ, ν) and assigns to each agent i a decision rule that depends only on $\mathcal{I}_i$.[1]

Unlike individual agents, the planner internalizes the payoff externality generated by aggregate action A through the outcome Ω. Efficiency therefore requires each agent to

[1]This notion of efficiency is closer in spirit to constrained efficiency in general equilibrium with incomplete markets (e.g., Diamond, 1978; Stiglitz, 1982; Greenwald and Stiglitz, 1986; Davila, Hong, Krusell, and Ríos-Rull, 2012). In those settings, the planner takes market incompleteness as given but internalizes pecuniary externalities. Analogously, the planner here takes dispersed information as given but internalizes the payoff externality.

internalize the expected marginal externality. The next proposition characterizes the team-efficient allocation. Under concave social preferences and strictly convex individual costs, the planner's problem admits a unique solution. The optimal decision rules reflect the trade-off between implementability and internalization of the payoff externality generated by aggregate action.

Proposition 2 (Existence and Uniqueness of the Team-efficient Allocation)**.** *Under Assumption 1 and strictly convex individual cost functions, with $\eta_i > 0$ for all i, there exists a unique team-efficient allocation $\{a_i^{te}\}_{i\in[0,1]}$. This allocation satisfies*

$$a_i^{te} = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{\kappa}{\eta_i}\mathbb{E}_i[\bar{S}'(\Omega^{te})], \quad \textit{with} \quad \Omega^{te} = \kappa A^{te} + \Gamma(\theta,\nu), \tag{7}$$

where $A^{te} \equiv \int_0^1 a_i^{te} di$.

Relative to the status quo (3), the team-efficient allocation (7) shares the same first term, $\frac{\beta_i}{\eta_i} m_{i,\theta}$, reflecting the dependence of actions on agents' private beliefs. The difference arises from the additional term $\frac{\kappa}{\eta_i}\mathbb{E}_i[\bar{S}'(\Omega^{te})]$, which induces agents to internalize the expected marginal social value of their actions.

Relative to the first best, the team-efficient allocation differs because the planner cannot condition on the true state (θ,ν). Instead, optimal actions depend on agents' private estimates, so the team-efficient allocation is governed by the expected marginal social value conditional on $\mathcal{I}_i$, rather than by the realized marginal social value available in the first best. When agents observe (θ,ν) perfectly, the team-efficient allocation coincides with the first best.

2.4 Strategic Characterization of the Team-Efficient Allocation

To gain further insight into the structure of the team-efficient allocation, we make the following parametric assumption on social preferences.

Assumption 2 (Linear-Quadratic Social Value)**.** *The average marginal social value function $\bar{S}'(\Omega)$ is linear in Ω and takes the form*

$$\bar{S}'(\Omega) = \bar{s}_1 - \bar{s}_2(\Omega - \mu_\Omega^{sq}),$$

for some $\bar{s}_2 > 0$, where μ_Ω^{sq} is defined in equation (4).

Assumption 2 delivers a linear-quadratic structure that is standard in the dispersed-information literature (e.g., Angeletos and Pavan, 2007, 2009) and implies linear decision rules. This specification arises, for example, when each agent's marginal social value is linear, $S_i'(\Omega) = s_{1i} - s_{2i}(\Omega - \mu_\Omega^{sq})$, with $\bar{s}_1 \equiv \int_0^1 s_{1i} di$ and $\bar{s}_2 \equiv \int_0^1 s_{2i} di$.

The following proposition characterizes the first-best and efficiency-benchmark allocations under this specification.

Proposition 3. *Under Assumptions 1 and 2:*

(i) The unique first-best allocation $\{a_i^{fb}\}_{i\in[0,1]}$ satisfies

$$a_i^{fb} = \frac{\beta_i}{\eta_i}\theta + \frac{\kappa}{\eta_i}[\bar{s}_1 - \bar{s}_2(\Omega^{fb} - \mu_\Omega^{sq})], \tag{8}$$

where

$$\Omega^{fb} = \frac{\kappa \bar{B}_H \theta + (\kappa^2/H)\bar{s}_1 + (\kappa^2 \bar{s}_2/H)\mu_\Omega^{sq} + \Gamma(\theta, \nu)}{1 + \kappa^2 \bar{s}_2/H}. \tag{9}$$

(ii) The unique team-efficient allocation $\{a_i^{te}\}_{i\in[0,1]}$ satisfies the strategic form:

$$a_i^{te} = \underbrace{\left(1 - \frac{\alpha}{H}\right)\mathbb{E}_i[a_i^{fb}]}_{\textit{Private Motive}} + \underbrace{\frac{\alpha}{\eta_i}\mathbb{E}_i[A^{te}]}_{\textit{Strategic Motive}} - \underbrace{\frac{\alpha}{\eta_i}\left(\bar{B}_H - \frac{\beta_i}{H}\right)m_{i,\theta}}_{\textit{Heterogeneity Correction}}, \quad \textit{where} \quad \alpha \equiv -\kappa^2 \bar{s}_2 < 0 \tag{10}$$

Under the linear-quadratic specification, first-best actions in (8) are linear functions of the fundamentals. In the absence of the externality ($\kappa \to 0$), the first-best allocation collapses to the autarky benchmark, under which the social optimum coincides with private incentives.

The strategic form (10) decomposes the agent's action into three components. The first term anchors the action to the agent's expected first-best. The second term captures the response to the expected aggregate action, scaled by the agent's private cost η_i. The coefficient α measures the equilibrium degree of coordination; since $\alpha < 0$, actions are strategic substitutes, so agents optimally moderate their actions when they expect others to act more aggressively. The third term corrects for deviations from the population-average type $\bar{B}_H$.

This structure highlights two departures from Angeletos and Pavan (2007). First, the strategic motive is heterogeneous: agents with higher costs (higher η_i) are less responsive to expected aggregate action. Second, the heterogeneity-correction term vanishes only when an agent's type coincides with the population average ($\beta_i/H = \bar{B}_H$). For agents with above-average private benefits ($\beta_i/H > \bar{B}_H$), this term is negative (since $\alpha < 0$), attenuating the aggressiveness implied by their idiosyncratic type and pulling their action toward the team-efficient benchmark.

The last two terms in (10) reflect a single underlying force: society's aversion to volatility in the aggregate outcome, captured by $\alpha < 0$, that is, the concavity of $\bar{S}(\Omega)$. Because $\bar{S}'' < 0$, fluctuations in Ω are socially costly. To dampen such fluctuations, the planner induces agents to lean against the aggregate: when others are expected to act more aggressively, each agent optimally moderates their own action. This mechanism gives rise to strategic substitutability ($\alpha < 0$).

The heterogeneity correction arises from the same force interacting with convex private costs. Since marginal costs increase with action ($\eta_i > 0$), cross-sectional dispersion in actions is socially inefficient: producing a given aggregate A is cheaper when individual actions are more aligned. The correction term therefore compresses the distribution of actions toward the team-efficient average, reducing dispersion costs.

2.5 The Coasean Ideal: A Synthetic Market

The team-efficient allocation serves as the social planner's target. In a decentralized economy without intervention, however, this allocation will generally fail to arise, as atomistic agents do not internalize the social value of their actions. The objective is therefore to design transfers that implement the team-efficient outcome as a decentralized equilibrium. To motivate our mechanism, we begin with an idealized market environment in which the externality is perfectly priced. This construction creates the missing market that resolves what Arrow (1969, p. 58) calls the "difference of prices between buyer and seller."

In this idealized benchmark, the aggregate action A is priced at time 1, after the state is realized. At time 0, agents choose their actions while anticipating the time-1 price. Let the

state be $\omega = (\theta, \nu)$. Let $A^{mkt}(\omega)$ denote the aggregate action realized in state ω under this benchmark, and define

$$\Omega^{mkt}(\omega) \equiv \kappa A^{mkt}(\omega) + \Gamma(\omega).$$

We begin by determining pricing at time 1. Fix a realized state ω and consider a Lindahl (1958) pricing system for the aggregate action. Each agent i is charged a personalized per-unit price $p_i(\omega)$ for the aggregate action, reflecting their marginal valuation. Suppliers receive a common producer price $\hat{p}(\omega)$ per unit. Budget balance in state ω requires

$$\hat{p}(\omega) = \int_0^1 p_i(\omega) di. \tag{11}$$

Given $p_i(\omega)$, agent i chooses a desired level A of the aggregate action to maximize net benefit,

$$\max_A \ \{S_i(\kappa A + \Gamma(\theta, \nu)) - p_i(\omega) A\}.$$

An interior optimum satisfies

$$\kappa S_i'(\kappa A + \Gamma(\theta, \nu)) = p_i(\omega). \tag{12}$$

In a Lindahl equilibrium, the aggregate action is a public good: all agents face the same level $A^{mkt}(\omega)$. Evaluating (12) at $A^{mkt}(\omega)$ yields the personalized Lindahl prices,

$$p_i(\omega) = \kappa S_i'\left(\Omega^{mkt}(\omega)\right). \tag{13}$$

These prices are heterogeneous: agents with higher marginal valuations $S_i'(\Omega)$ pay more for the aggregate action. Substituting (13) into the budget-balance condition (11) yields the producer price,

$$\hat{p}(\omega) = \int_0^1 \kappa S_i'\left(\Omega^{mkt}(\omega)\right) di = \kappa \bar{S}'\left(\Omega^{mkt}(\omega)\right). \tag{14}$$

The producer price in (14) depends on ω only through the outcome $\Omega^{mkt}(\omega)$, since the average social value depends on ω only through Ω. We therefore write

$$\hat{p}(\Omega) \equiv \kappa \bar{S}'(\Omega). \tag{15}$$

At time 0, each atomistic agent i chooses a_i taking the mapping $\Omega \mapsto \hat{p}(\Omega)$ as given and forming expectations conditional on $\mathcal{I}_i$:

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

Because agents are atomistic, their choice does not affect Ω^{mkt} or $\hat{p}(\Omega^{mkt})$. The first-order condition yields

$$a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{1}{\eta_i} \mathbb{E}_i[\hat{p}(\Omega^{mkt})]. \tag{16}$$

Substituting (14) into (16) gives

$$a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{\kappa}{\eta_i} \mathbb{E}_i \left[\bar{S}'(\Omega^{mkt}) \right]. \tag{17}$$

Condition (17) coincides with the characterization of the team-efficient allocation in equation (7). Hence, in this benchmark, pricing the aggregate action at time 1 according to equation (15) implements the team-efficient allocation at time 0. We refer to this benchmark as the *Coasean ideal*.

This benchmark assumes an institution that, in each realized state, enforces the Lindahl pricing identities in (13)–(15). A Lindahl equilibrium achieves constrained Pareto efficiency: all agents agree on the aggregate level of the public good, and personalized prices align individual incentives with the social optimum. However, computing the individualized prices $p_i(\omega)$ requires observing each agent's marginal valuation $S_i'(\Omega)$, which reflects both preferences and beliefs. This informational requirement poses a fundamental implementation challenge (see, e.g., Anderson and Duanmu, 2025). The next section instead constructs simple outcome-contingent transfers that replicate the incentives implied by (14) *without* requiring the regulator to observe agent-specific valuations.

3 The Coasean Transfer

Section 2.5 introduced an idealized environment in which the externality is perfectly priced through a synthetic market, yielding the team-efficient allocation. We now show how to

replicate the outcome of that synthetic market in the original decentralized economy using a simple transfer scheme based only on observables. This scheme implements the team-efficient allocation without communication, bargaining, or state verification. We refer to it as the *Coasean transfer*.

The key idea is to make agents trade deviations from the aggregate action at a price that reflects the social value of the outcome. To this end, we assign each agent a transfer that is proportional to their deviation from the aggregate action and priced by a function of the realized outcome. Formally, the transfer to agent i is

$$T_i = (a_i - A)p(\Omega).$$

Under this transfer scheme, agent i solves:

$$\max_{a_i} \ \mathbb{E}_i \left[e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i + S_i(\Omega) + (a_i - A)p(\Omega) \right]. \tag{18}$$

Because the agent is atomless, they take A and Ω as given. The first-order condition yields:

$$a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{1}{\eta_i} \mathbb{E}_i[p(\Omega)]. \tag{19}$$

We now determine the pricing function $p(\Omega)$ that implements the desired allocation. Recall from Section 2.5 that, in the idealized Coasean market, the equilibrium producer price is given by $\hat{p}(\Omega) = \kappa \bar{S}'(\Omega)$. As shown in (17), behavior in that market satisfies

$$a_i = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{\kappa}{\eta_i} \mathbb{E}_i[\bar{S}'(\Omega)]. \tag{20}$$

To replicate the outcome of the Coasean market, the transfer mechanism must induce the same first-order condition for every agent and every belief. Comparing (19) and (20) therefore implies

$$p(\Omega) = \kappa \bar{S}'(\Omega)$$

for all realizations of Ω. Thus, the optimal pricing rule sets the transfer price equal to the average marginal social value. The following proposition formalizes the implementation result.

Proposition 4 (Coasean Transfer Implementation)**.** *Suppose Assumption 1 holds and agents*

are atomless. If the regulator sets transfers of the form

$$T_i = (a_i - A)\, p(\Omega),$$

with pricing rule $p(\Omega) = \kappa \bar{S}'(\Omega)$*, then the resulting decentralized allocation coincides with the team-efficient allocation.*

The transfer $T_i = (a_i - A)p(\Omega)$ operates like a futures contract: agent i takes a net position $(a_i - A)$ that settles ex post at the realized price $p(\Omega)$. This mechanism has several useful properties. It uses only observables, that is, individual actions a_i, the aggregate A, and the realized outcome Ω. It requires no reports of private beliefs or states. Finally, it is 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 Lindahl equilibrium with a single, outcome-indexed price.

The Coasean transfer aligns private incentives with the Coasean market by making each agent face the average marginal social value of their action. Effectively, the mechanism decentralizes the planner's functional optimization problem from Section 2.3: the transfer induces agents to voluntarily adopt the team-efficient decision rules $a_i(\cdot)$ without direct mandates. This result relies on three features of the environment: agents are atomless; the externality operates through the aggregate outcome; and the observable Ω is a sufficient statistic for the unobservable state $\omega = (\theta, \nu)$ relevant for pricing the externality.[2]

Proposition 4 implies that implementing the Coasean transfer does not require the regulator to observe agents' private information. The transfer $T_i = (a_i - A)p(\Omega)$ depends only on the pricing function $p(\Omega) = \kappa \bar{S}'(\Omega)$ and on the realized outcome Ω. In particular, the regulator need not know the distribution of states (θ, ν), the signal structure, or the cross-sectional distribution of private types (η_i, β_i, S_i).

The Coasean transfer therefore delivers a *separation result*, in the sense that policy design depends only on social preferences and technology, while all informational details remain decentralized. The regulator specifies only the pricing rule $p(\Omega)$, which depends solely on social preferences $\bar{S}(\cdot)$ and the technology parameter κ. All other information remains private.

[2]Formally, Ω is a sufficient statistic because the payoff-relevant effect of the state $\omega = (\theta, \nu)$ on welfare operates entirely through the realized outcome Ω. The average marginal social value $p(\Omega) = \kappa \bar{S}'(\Omega)$ depends on (θ, ν) only via Ω, so conditioning transfers on Ω captures all information needed to price the externality.

The mechanism is anonymous, budget-balanced, and requires no reports. Agents' private information aggregates endogenously into the outcome Ω, and pricing this single observable suffices to implement the team-efficient allocation.

This approach resolves the span problem that arises in policy mechanisms linear in fundamentals (e.g., Angeletos and Pavan, 2009). Such instruments fail when shocks, such as ν or nonlinear functions of θ, fall outside their span. Coasean transfers bypass this limitation by conditioning directly on the realized outcome Ω, which is a sufficient statistic embedding all payoff-relevant uncertainty.

The same logic addresses a broader concern raised by Radner (1982, Section 1.6), who argued that complete state-contingent markets are infeasible when the state space is too rich for a workable contract system. The Coasean transfer replaces the entire state space with a single outcome-indexed price $p(\Omega)$, thereby achieving the Coasean ideal with a feasible instrument.

Interpretation under Linear Marginal Social Value. We now specialize to the linear-quadratic case in Assumption 2. Although not essential for implementation, this restriction imposes linearity on the marginal social value function $\bar{S}'(\Omega)$ and yields a transparent interpretation of the Coasean transfer price. Using $\bar{S}'(\Omega) = \bar{s}_1 - \bar{s}_2(\Omega - \mu_\Omega^{sq})$, we can define two parameters. Let $\text{ORA} \equiv \bar{s}_2 = -\bar{S}''(\Omega)$ denote outcome risk aversion, the social aversion to fluctuations in Ω. Let $\text{MSV}^{sq} \equiv \bar{s}_1 = \bar{S}'(\mu_\Omega^{sq})$ denote the marginal social value at the status-quo benchmark. The Coasean transfer price is then

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

The price consists of a baseline component, κMSV^{sq}, adjusted for deviations of the realized outcome from the benchmark. The adjustment term reflects society's aversion to outcome volatility.

4 Harnessing Uncertainty and Disagreement

The Coasean transfer implements the team-efficient allocation by conditioning the transfer price $p(\Omega)$ on the realized outcome. We now study how this mechanism interacts with the informational environment—specifically, how uncertainty and disagreement shape the equilibrium allocation and how the mechanism remains robust to departures from rational expectations.

We specialize the environment to a standard benchmark. The following assumption imposes a simple parametric structure on the outcome process and the information structure. This specialization allows us to derive closed-form solutions for the team-efficient allocation and to characterize how uncertainty and disagreement shape equilibrium behavior.

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

$$\Gamma(\theta, \nu) = \theta + \nu.$$

Furthermore, the information structure is affine: for $z \in \{\theta, \nu\}$, the average posterior belief is linear in the true state,

$$\int_0^1 m_{i,z}\, di = \lambda_z z + (1 - \lambda_z)\mu_z, \tag{21}$$

where $\lambda_z \in [0, 1]$ measures the weight agents place on private signals (reflecting aggregate precision of information about z).

This linearity is standard under Bayesian updating. For example, it holds when agents observe signals of the form $y_i = z + \varepsilon_i$ and (z, ε_i) are jointly Gaussian, and it extends more generally to elliptical distributions.[3] The result is robust to the number and composition of public and private signals. Throughout the remainder of this section, we maintain Assumptions 2 and 3 and further assume that agent types (η_i, β_i, S_i) are independent of

[3] In the Gaussian case, suppose $y_i = z + \varepsilon_i$ with $z \sim \mathcal{N}(\mu_z, \sigma_z^2)$ and $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ independent. Then the posterior mean is $m_{i,z} = \lambda_z y_i + (1 - \lambda_z)\mu_z$, where $\lambda_z = \sigma_z^2/(\sigma_z^2 + \sigma_\varepsilon^2)$. By the law of large numbers, $\int_0^1 y_i\, di = z$, which implies condition (21).

posterior beliefs $m_{i,z}$. We first examine how uncertainty shapes the team-efficient allocation and then turn to the role of disagreement across agents.

4.1 Harnessing Uncertainty

We begin by characterizing the effect of aggregate uncertainty on the team-efficient allocation. The following proposition establishes that the fixed point in equation (10) admits a closed-form solution under Assumptions 2 and 3.

Proposition 5 (Team-Efficient Allocation under Linear-Quadratic-Affine Structure)**.** *Under Assumptions 2 and 3, the team-efficient aggregate action is* $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$*, with sensitivity coefficients:*

$$C_\theta = \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta}, \qquad C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu},$$

with $\alpha \equiv -\kappa^2\bar{s}_2$ *the efficient degree of coordination derived in Proposition 3. The constant* C_0 *does not affect the response to fundamentals and is provided in the proof.*

Proposition 5 establishes that the sensitivity of aggregate action to the fundamentals θ and ν depends endogenously on the information precision parameters λ_θ and λ_ν. These parameters capture the weight agents assign to private signals, which increases with prior uncertainty. The coefficients C_θ and C_ν are increasing in λ_θ and λ_ν, respectively, with $\partial|C_\theta|/\partial\lambda_\theta > 0$ whenever $H\bar{B}_H + \alpha/\kappa \neq 0$ and $\partial|C_\nu|/\partial\lambda_\nu > 0$.

Thus, the equilibrium aggregate response to fundamentals adjusts endogenously to informational conditions. Higher prior uncertainty leads agents to place greater weight on private signals, and outcome-contingent transfers aggregate these responses efficiently, yielding an allocation that becomes more sensitive to fundamentals as uncertainty rises.

4.2 Harnessing Disagreement

Thus far, the analysis has assumed rational expectations. We now study how disagreement about others' information affects the team-efficient allocation. We introduce *distrust*, whereby agents discount the informational content of others' actions due to uncertainty about their

beliefs. This approach builds on work on learning and higher-order beliefs (Banerjee, 2011) and is related to the sentiment-risk channel in Dumas et al. (2009).

We model distrust as a distortion in agents' beliefs about the responsiveness of aggregate action to fundamentals. For each fundamental $z \in \{\theta, \nu\}$, we introduce a trust parameter $\varphi_z \in [0, 1]$. The case $\varphi_z = 1$ corresponds to the rational-expectations benchmark, while $\varphi_z < 1$ captures distrust. Distrust enters the agent's problem through their perceived strategic motive, $\mathbb{E}_i[A^{te}]$, by attenuating the inferred sensitivity of aggregate action to z.

To formalize this distortion, consider agent i's beliefs about others' private signals regarding a fundamental $z \in \{\theta, \nu\}$. Agent i fully trusts their own signal y_i, but views other agents' signals y_j, for $j \neq i$, as less informative. Let y_j^i denote agent i's belief about agent j's signal. We assume that

$$y_j^i = \mu_z + \varphi_z(z - \mu_z) + \sqrt{1 - \varphi_z^2}\phi_i + \varepsilon_{y,j}. \tag{22}$$

Here, φ_z captures the correlation that agent i believes agent j's signal has with the true state z, relative to their own. When $\varphi_z = 1$, agent i believes others' signals are structured like their own. When $\varphi_z < 1$, agent i believes others' signals are distorted by an unobservable, mean-zero noise ϕ_i, which weakens their correlation with the truth.

The social planner evaluates welfare under the true distribution but takes as given that agents act on distrustful beliefs. This leads to a modified team-efficient allocation.

Proposition 6 (Team-Efficient Allocation under Distrust)**.** *Under Assumptions 2 and 3, with distrust parameterized by $(\varphi_\theta, \varphi_\nu)$, there is a unique team-efficient allocation. The individual action a_i^{te} satisfies the strategic form (10) from Proposition 3, but the expectation $\mathbb{E}_i$ reflects agent i's distrust. The aggregate action $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$, with sensitivity coefficients*

$$C_\theta = \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta\varphi_\theta}, \qquad C_\nu = \frac{(\alpha/\kappa)\lambda_\nu}{H - \alpha\lambda_\nu\varphi_\nu}, \quad \textit{with} \quad \alpha = -\kappa^2\bar{s}_2 < 0.$$

The constant C_0 does not affect responses to fundamentals and is provided in the proof.

Since $\alpha < 0$, the sensitivity coefficients $|C_\theta|$ and $|C_\nu|$ are strictly decreasing in the trust parameters φ_θ and φ_ν. Proposition 6 therefore implies that distrust increases aggregate

responsiveness to fundamentals. With strategic substitutability, agents who discount others' information underestimate aggregate adjustment and optimally respond more strongly to their own signals. The resulting equilibrium aggregates these responses, leading to greater sensitivity to fundamentals than under rational expectations.

A key implication of Proposition 6 is the robustness of Coasean transfers to distrust. As established by Lemma 1 in the proof of Proposition 6, the standard transfer scheme, $T_i = (a_i - A)p(\Omega)$, implements the team-efficient allocation without requiring the regulator to observe or parameterize the degree of distrust, that is, the trust parameters φ_θ and φ_ν. The pricing function $p(\Omega) = \kappa \bar{S}'(\Omega)$ depends only on the social objective and the technology, and is therefore invariant to agents' beliefs. While the instrument itself is fixed, equilibrium behavior adjusts endogenously. As distrust increases, agents perceive aggregate action to be less responsive and optimally increase the weight placed on their own private signals. As a result, aggregate responsiveness to fundamentals is amplified. Even in the limit of vanishing trust ($\varphi_z \to 0$ for $z \in \{\theta, \nu\}$), the equilibrium allocation remains sensitive to the state.

5 The Price-vs.-Quantity Trade-off

The preceding analysis showed that Coasean transfers harness uncertainty and disagreement by conditioning prices on the realized outcome. We now show that conventional instruments cannot replicate this property. When the regulator is restricted to setting either a fixed tax or a fixed quantity cap, the same uncertainty that Coasean transfers exploit forces a second-best choice between stabilizing prices and stabilizing quantities. We characterize the optimal tax and the optimal quantity cap in our environment and show that the welfare ranking between them depends on the curvature of private costs relative to social benefits—a version of the Weitzman (1974) price-quantity trade-off. Coasean transfers strictly dominate both instruments by escaping this trade-off entirely. Finally, we show that price discovery in a competitive permit market does not restore efficiency.

5.1 Optimal Price Instrument: Pigouvian Tax/Subsidy

We begin by contrasting the Coasean transfer with the standard Pigouvian tax or subsidy. Suppose the regulator is restricted to a single, uniform tax or subsidy τ on individual actions a_i, set ex ante and independent of realized outcomes. The regulator chooses τ to internalize the expected marginal externality and maximize social welfare.

Given a uniform tax τ on a_i and a lump-sum rebate $R = \tau A$ of aggregate tax revenue, each agent i takes τ and aggregates as given and chooses a_i to maximize expected utility

$$\max_{a_i} \mathbb{E}_i \left[e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i + S_i(\Omega) - \tau a_i + R \right]. \tag{23}$$

The agent's best response to a given τ is

$$a_i^{tax}(\tau) = \frac{\beta_i m_{i,\theta} - \tau}{\eta_i}. \tag{24}$$

The regulator chooses the tax rate τ to maximize ex-ante welfare, taking into account how agents optimally respond to the tax. Formally, the optimal Pigouvian tax solves

$$\max_{\tau} \mathbb{E} \left[\int_0^1 \left(e - \frac{\eta_i}{2} a_i^{tax}(\tau)^2 + \beta_i \theta a_i^{tax}(\tau) + S_i(\Omega(\tau)) \right) di \right],$$

where $\Omega(\tau) \equiv \kappa A^{tax}(\tau) + \Gamma(\theta, \nu)$ is the resulting outcome and $A^{tax}(\tau) \equiv \int_0^1 a_i^{tax}(\tau) di$ is the corresponding aggregate action. The following proposition characterizes the optimal Pigouvian tax in closed form.

Proposition 7 (Optimal Pigouvian Tax/Subsidy)**.** *Consider a uniform tax or subsidy τ per unit of action. Under Assumptions 1 and 2, and assuming agent types $(\beta_i, \eta_i, S_i(\cdot))$ are independent of posterior beliefs $m_{i,z}$, $z \in \{\theta, \nu\}$, the unique welfare-maximizing Pigouvian rate τ^* is*

$$\tau^* = \frac{-\kappa \bar{s}_1 + \kappa \bar{s}_2 \left(\kappa \bar{B}_H \, \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq} \right)}{1 + \kappa^2 \bar{s}_2 / H}. \tag{25}$$

The optimal Pigouvian rate satisfies $\tau^* = -\kappa \mathbb{E}[\bar{S}'(\Omega(\tau^*))]$. It is a tax when $\bar{S}'(\Omega) < 0$ and a subsidy when $\bar{S}'(\Omega) > 0$.

Proposition 7 highlights the informational burden of Pigouvian taxation. Implementing the optimal rate τ^* requires knowledge of primitives $(\bar{s}_1, \bar{s}_2, \kappa)$, public priors, and cross-sectional aggregates summarized by $(H, \bar{B}_H)$. By contrast, the Coasean transfer requires neither priors nor cross-sectional information, relying only on the pricing function $\kappa \bar{S}'(\cdot)$ and the realized outcome Ω.

The next corollary shows that Coasean transfers reduce to the optimal Pigouvian instrument in the absence of private information.

Corollary 7.1. *Suppose agents have no private information about (θ, ν) beyond the public prior, so that posterior beliefs coincide with priors. Under Assumptions 1 and 2, the allocation implemented by the Coasean transfer coincides with the allocation induced by the optimal uniform Pigouvian tax or subsidy τ^* in Proposition 7.*

Importantly, Coasean transfers deliver this Pigouvian outcome without requiring the regulator to know the cross-sectional aggregates $(H, \bar{B}_H)$ or the priors needed to compute τ^*. When private information is absent, the uniform tax or subsidy emerges endogenously as a special case of the Coasean transfer, with rate $\tau^{te} = -\kappa \mathbb{E}[\bar{S}'(\Omega^{te})]$.

5.2 Optimal Quantity Instrument: Cap-and-Trade

We next consider the optimal quantity-based policy instrument, the natural counterpart to the Pigouvian tax analyzed in Section 5.1. Suppose the regulator is restricted to setting a single aggregate cap A^{cap} on total action A, chosen ex ante on the basis of public information. The regulator's problem is to select the cap that maximizes social welfare.

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 aggregates as given and chooses a_i to maximize expected utility

$$\max_{a_i} \mathbb{E}_i \left[e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i + S_i(\Omega) - P a_i + R \right]. \tag{26}$$

The agent's best response to a given price P is

$$a_i^{cap}(P) = \frac{\beta_i m_{i,\theta} - P}{\eta_i}. \tag{27}$$

For a given cap A^{cap}, the competitive permit price $P(A^{cap})$ clears the market by equating aggregate demand to the fixed supply, so that the induced allocation aggregates to A^{cap}. The regulator then chooses A^{cap} to maximize ex-ante welfare, anticipating that the permit market implements this allocation. The following proposition provides the closed-form solution.

Proposition 8 (Optimal cap-and-trade quantity)**.** *Under Assumptions 1 and 2, and independence of types and posterior beliefs, the unique welfare-maximizing cap is*

$$A^{cap*} = \frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa\bar{s}_1 - \kappa\bar{s}_2\Big(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\Big)}{H + \kappa^2\bar{s}_2}. \tag{28}$$

The expected permit price under cap-and-trade equals the optimal Pigouvian rate from Proposition 7, $\mathbb{E}[P(A^{cap*})] = \tau^*$*, and the expected aggregate action under Pigouvian taxation equals the optimal cap,* $\mathbb{E}[A(\tau^*)] = A^{cap*}$*.*

Propositions 7 and 8 show that Pigouvian taxation and cap-and-trade are equivalent in expectation. In particular, the expected permit price under cap-and-trade equals the optimal Pigouvian tax, and the expected aggregate action under Pigouvian taxation equals the optimal cap. Although both instruments share this equivalence in expectations—and the same informational burden for the regulator—they differ in their state-by-state behavior. Under cap-and-trade, aggregate action is fixed at A^{cap*} while the permit price varies with agents' information. Under Pigouvian taxation, the tax rate τ^* is fixed while aggregate action responds to realized information. The regulator therefore faces a choice between fixing quantities and letting prices adjust, or fixing prices and letting quantities adjust. This is precisely the price-quantity trade-off emphasized by Weitzman (1974), which we now characterize in our setting.

5.3 Prices *vs.* Quantities

We compare expected welfare under Pigouvian taxation and cap-and-trade. Pigouvian taxes let aggregate action respond to private information, increasing outcome volatility and reducing welfare. Cap-and-trade fixes aggregate action, but convex private costs render permit price volatility costly. The following proposition quantifies this trade-off.

Proposition 9 (Prices *vs.* Quantities)**.** *Under Assumptions 1, 2, and 3, the expected welfare difference between cap-and-trade and Pigouvian taxation is*

$$\mathbb{E}[W^{cap}] - \mathbb{E}[W^{tax}] = \underbrace{-H\bar{B}_H^2\left(1 - \frac{\lambda_\theta}{2}\right)\sigma_\theta^2\lambda_\theta}_{\text{Loss from Price Volatility}} + \underbrace{\frac{\bar{s}_2\kappa\bar{B}_H}{2}(2 + \kappa\bar{B}_H\lambda_\theta)\sigma_\theta^2\lambda_\theta}_{\text{Gain from Outcome Stabilization}}. \tag{29}$$

The first term captures the private loss from price volatility under cap-and-trade. With convex private costs, permit price fluctuations inherent to cap-and-trade lower expected private payoffs. The second term captures the social gain from outcome stabilization. By fixing quantity, cap-and-trade eliminates the amplification of fundamental shocks that arises when agents respond to private information. The price-vs-quantity trade-off therefore pits H (the curvature of private costs) against $\bar{s}_2$ (the curvature of social value).

Both terms are proportional to $\sigma_\theta^2\lambda_\theta$. When agents have no private information about θ ($\lambda_\theta = 0$), the instruments deliver the same welfare and the trade-off disappears. Dispersed information is thus the friction that activates the trade-off. Weitzman (1974) analyzed the limiting case $\lambda_\theta = 1$, where firms observe the state perfectly after the regulator acts. Our analysis extends the trade-off to environments with dispersed information.

Coasean transfers escape this trade-off by exploiting the very friction that generates it. Because they implement the unconstrained team-efficient allocation, Coasean transfers strictly dominate both standard instruments.[4] Agents' informational advantage over the regulator, captured by $\lambda_\theta > 0$, is what outcome-contingent transfers exploit to implement the efficient allocation, without requiring a choice between price stability and quantity stability.

5.4 Can Price Revelation Restore Efficiency?

A natural objection is that cap-and-trade may also exploit agents' informational advantage through price discovery. In a rational-expectations equilibrium, permit prices aggregate private signals and may reveal information about fundamentals. Because the price clears the market by equating aggregate demand to the fixed cap, it reveals the average signal driving

[4] In the proof of Proposition 10, Lemma 2 formally derives the welfare difference between Coasean transfers and the optimal Pigouvian tax. It shows that the gain is strictly positive whenever agents possess private information ($\lambda_\theta > 0$ or $\lambda_\nu > 0$). A similar result holds when comparing with cap-and-trade.

that demand. This raises the possibility that price discovery could restore efficiency and make cap-and-trade the preferred instrument.

We show that this is not the case. Even under the most favorable assumptions for cap-and-trade—namely that the permit price perfectly reveals the fundamental state θ—a fundamental limitation remains. Price discovery can improve allocative efficiency by reallocating activity toward agents with lower costs or higher benefits, but it does not address the central problem in our environment, which is the pricing of the externality.

We formalize this argument by comparing Coasean transfers with a cap-and-trade regime in which the permit price fully reveals θ. Full price revelation alters agents' actions but does not change the regulator's ex-ante problem. In particular, the optimal cap A^{cap*} from Proposition 8 remains unchanged and, as shown in the proof of Proposition 10, the equivalence $\mathbb{E}[P(A^{cap*})] = \tau^*$ continues to hold. The only effect of price discovery is therefore on expected welfare. The following proposition characterizes this welfare comparison.

Proposition 10 (Coasean Transfers vs. Cap-and-Trade with Full Price Revelation)**.** *Under Assumptions 1, 2, and 3, the expected welfare difference between Coasean transfers and a cap-and-trade regime with full price revelation is*

$$\mathbb{E}[W^{ct}] - \mathbb{E}[W^{cap}] = \underbrace{\frac{\alpha^2 \lambda_\nu \sigma_\nu^2}{2\kappa^2(H - \alpha\lambda_\nu)}}_{\textit{Exogenous Stabilization}} + \underbrace{\frac{\lambda_\theta(\alpha + \kappa H \bar{B}_H)^2 \sigma_\theta^2}{2\kappa^2(H - \alpha\lambda_\theta)}}_{\textit{Fundamental Stabilization}} - \underbrace{\frac{1}{2}(1 - \lambda_\theta)\Delta_{het}\sigma_\theta^2}_{\textit{Allocative Efficiency Loss}}, \quad (30)$$

where $\Delta_{het} \equiv \overline{B^2}_H - H\bar{B}_H^2 \geq 0$, $\overline{B^2}_H \equiv \int_0^1 (\beta_i^2/\eta_i)\, di$, $\bar{B}_H \equiv \int_0^1 (\beta_i/\eta_i)\, di$ *and* $H \equiv \left(\int_0^1 \eta_i^{-1} di\right)^{-1}$.

The first two terms in (30) are strictly positive and favor Coasean transfers. They capture the social gains from letting the aggregate outcome adjust to the exogenous shock ν and the fundamental θ. By pricing the externality directly, Coasean transfers induce agents to respond to these shocks, adjusting the aggregate outcome. Cap-and-trade fixes the aggregate ex ante, and price discovery cannot overcome this rigidity.

The third term is negative and favors cap-and-trade. It captures the value of perfect information about θ, which lets agents allocate the fixed aggregate burden A efficiently across heterogeneous types. This advantage arises entirely from the private component of the payoff

function (2) and is unrelated to the social objective of pricing the externality. If agents are homogeneous ($\Delta_{het} = 0$), the advantage disappears, and Coasean transfers strictly dominate.

Proposition 10 clarifies the distinction between informational efficiency and social efficiency. Price revelation under cap-and-trade improves the cross-sectional allocation of effort, but the permit price P reflects the scarcity of the aggregate constraint rather than the marginal social value of the outcome, $\bar{S}'(\Omega)$. In this sense, cap-and-trade prices the constraint, not the externality.

6 Incentives for Information Acquisition

Thus far, the analysis has taken the information structure as given and shown that outcome-contingent Coasean transfers implement the team-efficient allocation under uncertainty and disagreement. We now study how the same mechanism affects agents' incentives to acquire information. Coasean transfers operate along two margins. First, for fundamentals that directly affect private payoffs (such as θ), the dependence of transfers on the aggregate outcome introduces an additional strategic motive—forecasting the state-contingent price, which strengthens existing incentives to learn. Second, by tying payoffs to the aggregate outcome Ω, the mechanism creates incentives to acquire information about fundamentals that do not directly enter private payoffs (such as ν). As a result, agents optimally acquire information about all fundamentals relevant for social welfare.

6.1 Acquiring Information about the Fundamental θ

We first analyze incentives to acquire information about the fundamental θ. To obtain closed-form comparisons, we maintain the linear-quadratic-affine structure in Assumptions 2 and 3. In the status quo, agents acquire information about θ only to improve the private payoff term $\beta_i \theta a_i$. A Pigouvian tax, being fixed ex ante, does not affect this incentive, as it shifts actions but does not alter their information sensitivity.

Under Coasean transfers, optimal actions depend on agents' forecasts of the state-contingent price $\mathbb{E}_i[p(\Omega)]$. Information about θ therefore affects payoffs both directly and through the induced price response. The next proposition characterizes when Coasean

transfers strengthen incentives to acquire information about θ relative to the status quo.

Proposition 11 (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*

$$\beta_i < \frac{-\alpha H(1+\kappa \bar{B}_H \lambda_\theta)}{2\kappa(H-\alpha\lambda_\theta)}. \tag{31}$$

Under strategic substitutability ($\alpha < 0$), Coasean transfers introduce a strategic incentive to respond to anticipated aggregate behavior. Agents who perceive weaker aggregate adjustment optimally increase the responsiveness of their own actions, raising the marginal value of information. When the private benefit parameter β_i is sufficiently large, the private payoff motive dominates and the condition in (31) fails. Thus, Coasean transfers strengthen incentives to acquire information precisely when strategic considerations outweigh private incentives.

6.2 Acquiring Information about the exogenous shock ν

We next analyze incentives to acquire information about the exogenous state ν. In the status quo, optimal actions depend only on agents' beliefs about θ and are independent of ν. A Pigouvian tax likewise leaves actions invariant to ν, as it shifts the level of actions without affecting their informational sensitivity. In both cases, information about ν has zero marginal value.

Under Coasean transfers, optimal actions depend on agents' forecasts of the state-contingent price $\mathbb{E}_i[p(\Omega)]$. Because the aggregate outcome Ω depends on ν, information about ν affects payoffs through the price response. As a result, acquiring information about ν becomes payoff-relevant. The next proposition characterizes agents' incentives to acquire information about ν under Coasean transfers relative to the status quo.

Proposition 12 (Information Acquisition about ν)**.** *Under Assumptions 2 and 3, Coasean transfers create a strictly positive incentive to acquire information about ν.*

The logic parallels the case of θ, with one important difference. For ν, there is no private payoff motive to offset the strategic incentive induced by Coasean transfers. In the status quo,

ν does not enter private payoffs, and agents optimally ignore it. Under Coasean transfers, equilibrium prices vary with the aggregate outcome to stabilize social welfare, reflecting society's aversion to volatility, $\bar{s}_2 > 0$. Because the aggregate outcome depends on ν, the state-contingent price inherits this dependence. Anticipating the transfer therefore requires agents to forecast ν, rendering information about ν payoff-relevant. In the absence of any countervailing private incentive, the strategic motive operates without attenuation, implying a strictly positive value of information for all agents.

This comparison highlights a qualitative difference in how Coasean transfers affect information acquisition. For fundamentals such as θ that already enter private payoffs, the mechanism operates on the intensive margin by strengthening existing learning incentives. For fundamentals such as ν that matter only through the aggregate outcome, the mechanism operates on the extensive margin by creating learning incentives where none exist in the status quo. By linking payoffs to the aggregate outcome Ω, outcome-contingent transfers align private information acquisition with the full set of fundamentals relevant for the team-efficient allocation.

7 Political Viability

This section studies the political viability of Coasean transfers by comparing them to the optimal Pigouvian tax from Section 5.1 and the optimal cap-and-trade scheme from Section 5.2. To unify the analysis, let $\hat{P}$ denote the benchmark marginal price ($\hat{P} = \tau^*$ under taxation, $\hat{P} = P$ under cap-and-trade) and let $\hat{a}_i$ and $\hat{A}$ denote the corresponding individual and aggregate actions. An agent supports Coasean transfers over the benchmark when the expected welfare difference, conditional on $\mathcal{I}_i$, is positive. The following proposition decomposes this welfare difference into an efficiency gain and a distributional component.

Proposition 13 (Coasean Transfers versus Standard Instruments)**.** *Fix agent i with type $(\beta_i, \eta_i, S_i(\cdot))$ under Assumption 2. Let a_i^{ct} and A^{ct} denote individual and aggregate actions under the Coasean transfer, and let $\hat{a}_i$ and $\hat{A}$ denote the corresponding actions under the benchmark regulation. Define $\Delta a_i := a_i^{ct} - \hat{a}_i$ and $\Delta A := A^{ct} - \hat{A}$. The agent's expected payoff*

difference admits the decomposition

$$\mathbb{E}_i[u_i^{ct} - \hat{u}_i] = \Delta_i^{\text{eff}} + \Delta_i^{\text{dist}}, \tag{32}$$

where

$$\Delta_i^{\text{eff}} \equiv \frac{\eta_i}{2}(\Delta a_i)^2 + \frac{\kappa^2}{2}\bar{s}_2\,\mathbb{E}_i[(\Delta A)^2] \tag{33}$$

is the efficiency gain and

$$\Delta_i^{\text{dist}} \equiv \frac{\kappa^2}{2}(\bar{s}_2 - s_{2i})\,\mathbb{E}_i[(\Delta A)^2] + \kappa\mathbb{E}_i[(S_i'(\hat{\Omega}) - \bar{S}'(\hat{\Omega}))\Delta A] + \mathbb{E}_i[(\hat{a}_i - \hat{A})(p(\Omega^{ct}) + \hat{P})] \tag{34}$$

is the distributional component, which vanishes under preference homogeneity.

The decomposition in (32) reveals a central result: Coasean transfers generate strictly positive efficiency gains (equation (33)), while heterogeneity induces zero-sum transfers (equation (34)). If agents are homogeneous ($s_{2i} = \bar{s}_2$, $S_i'(\Omega) = \bar{S}'(\Omega)$, $\eta_i = \eta$, and $\beta_i = \beta$), the distributional terms vanish. In this case, Coasean transfers Pareto-dominate both benchmarks and command unanimous support.

The efficiency gain Δ_i^{eff} arises from two sources. The term $\frac{\eta_i}{2}(\Delta a_i)^2$ reflects individual flexibility: relative to the benchmark instrument, the Coasean transfer lets agents adjust their actions more efficiently in response to private information. The term $\frac{\kappa^2}{2}\bar{s}_2\mathbb{E}_i[(\Delta A)^2]$ reflects collective adaptability: aggregate action responds to the realized state, reducing outcome volatility that is socially costly.

The distributional component Δ_i^{dist} arises from preference heterogeneity and has three sources. The term $\frac{\kappa^2}{2}(\bar{s}_2 - s_{2i})\mathbb{E}_i[(\Delta A)^2]$ reflects risk reallocation: transfers price aggregate volatility at the social average $\bar{s}_2$, so agents with below-average curvature ($s_{2i} < \bar{s}_2$) earn a surplus by bearing risk for more risk-averse agents. The term $\kappa\mathbb{E}_i[(S_i'(\hat{\Omega}) - \bar{S}'(\hat{\Omega}))\Delta A]$ reflects valuation heterogeneity, capturing differences between an agent's marginal valuation of the aggregate outcome and the population average. The term $\mathbb{E}_i[(\hat{a}_i - \hat{A})(p(\Omega^{ct}) + \hat{P})]$ reflects an agent's net position in the externality: agents with $\hat{a}_i > \hat{A}$ sell excess contributions, while those with $\hat{a}_i < \hat{A}$ buy them, with gains or losses determined by the realized Coasean price relative to the benchmark price $-\hat{P}$.

8 Applications

We illustrate the generality of the Coasean transfer by applying it to two canonical environments with externalities. The first is the endogenous growth model of Romer (1986), which features positive knowledge spillovers. The second is the fishery model of Gordon (1954), which features a negative congestion externality. In both cases, we extend the original formulations to allow for dispersed information about fundamentals. These extensions let us study how outcome-contingent transfers implement the team-efficient allocation in settings where the original models abstract from informational frictions.

8.1 Positive Externalities: Knowledge Spillovers

We extend the two-period endogenous growth model of Romer (1986) to allow for dispersed information about productivity. A continuum of firms $i \in [0, 1]$ lives for two periods. In period 1, each firm is endowed with e units of the consumption good and chooses investment $k_i \in \mathbb{R}$ in knowledge capital. Aggregate investment is $K = \int_0^1 k_i \, di$. Investment entails a convex adjustment cost, so period-1 consumption is $c_{1i} = e - \frac{c}{2}k_i^2$. In period 2, output is $c_{2i} = (\theta + \gamma K)k_i$, where $\gamma > 0$ measures the knowledge spillover: each firm's investment raises all firms' productivity through the aggregate stock K. Firm i maximizes total net output, $c_{1i} + c_{2i}$, so its payoff is[5]

$$u_i(k_i, K, \theta) = e - \frac{c}{2}k_i^2 + (\theta + \gamma K)k_i.$$

We assume $c > 2\gamma$ to ensure a well-defined equilibrium by making the private cost sufficiently convex relative to the spillover. This setting maps to our baseline framework with $a_i = k_i$, $A = K$, and $\kappa = \gamma$. The key difference is that the externality enters multiplicatively through production rather than additively through a separable social value function. Despite this difference, the same market-equivalence logic applies, and Coasean transfers implement the team-efficient allocation.

[5] Romer (1986) models the cost of knowledge investment as foregone consumption, $c_1 = \bar{e} - k$. We instead assume a convex adjustment cost, $\frac{c}{2}k_i^2$, which is standard in investment models and delivers an interior optimum. The linear term $-k_i$ can be absorbed into the endowment normalization without changing the equilibrium structure or the externality.

As in Romer's original model, aggregate investment raises the social marginal product of knowledge. We introduce uncertainty by assuming that productivity $\theta \sim \mathcal{N}(\mu, \sigma_\theta^2)$ is unobserved, with μ large enough that $\theta > 0$ with high probability. Each firm observes a private signal $y_i = \theta + \varepsilon_i$, where $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ is i.i.d. across firms. Let $m_i \equiv \mathbb{E}[\theta \mid y_i] = (1-\lambda)\mu + \lambda y_i$, where $\lambda = \sigma_\theta^2/(\sigma_\theta^2 + \sigma_\varepsilon^2) \in (0,1)$ measures signal precision.

Lemma 4 in the Appendix characterizes the status quo and team-efficient allocations. In the status quo, firms take aggregate investment as given, yielding the first-order condition $ck_i = m_i + \gamma \mathbb{E}_i[K]$. Under team efficiency, the planner internalizes the spillover, leading to $ck_i = m_i + 2\gamma \mathbb{E}_i[K]$. Both can be written as

$$k_i = \frac{1}{c} m_i + \alpha \mathbb{E}_i[K],$$

where $\alpha^{sq} = \gamma/c$ and $\alpha^{te} = 2\gamma/c$. The optimal degree of coordination α^{te} is positive, reflecting strategic complementarity: higher expected aggregate investment raises each firm's marginal productivity. Coordination is weaker in the status quo because firms do not internalize spillovers.

Consider a regulator who imposes the transfer $T_i(k_i, K) = (k_i - K)p(K)$ with price $p(K) = \gamma K$. This price equals the Lindahl price obtained by aggregating marginal valuations,

$$\int_0^1 \frac{\partial u_i}{\partial K} di = \int_0^1 \gamma k_i di = \gamma K.$$

Under this transfer, firm i maximizes $\mathbb{E}_i[u_i + T_i]$, yielding

$$k_i = \frac{1}{c} m_i + \frac{2\gamma}{c} \mathbb{E}_i[K].$$

Thus, Coasean transfers implement the team-efficient allocation.

The allocation implemented by Coasean transfers nests standard Pigouvian correction and the first best. When information is uninformative ($\lambda \to 0$), firms rely on the prior and the team-efficient allocation converges to $K^{te} = \mu/(c - 2\gamma)$, which coincides with the outcome under the optimal Pigouvian subsidy. When information is perfect ($\lambda \to 1$), firms observe θ and the team-efficient allocation equals the first best $K^{te} = \theta/(c - 2\gamma)$.

A Pigouvian subsidy fixed ex ante cannot reach this frontier. As $\lambda \to 1$, it induces an

aggregate response coefficient $1/(c-\gamma)$, which is strictly smaller than the first-best coefficient $1/(c-2\gamma)$. Coasean transfers therefore interpolate between fixed Pigouvian correction and state-contingent first-best efficiency, with dispersed information determining the degree of responsiveness.

8.2 Negative Externalities: Common-Pool Resources

We next consider a canonical negative externality: the open-access fishery model of Gordon (1954). In Gordon's model, aggregate harvesting effort reduces the productivity of the fish stock, generating a crowding externality that leads to overexploitation under open access.

Let e_i denote the effort exerted by agent i, and let aggregate effort be $E = \int_0^1 e_i \, di$. Total catch depends on aggregate effort and on the size of the fish stock, which we write as $a = \bar{a} + \theta$, where θ captures fluctuations in biological productivity. The aggregate production function takes the form

$$L = \frac{c\, a\, E}{1 + cbE},$$

with $c, b > 0$. This specification captures diminishing returns to aggregate effort due to stock depletion: as E rises, congestion lowers average productivity.

To bring the model into our linear-quadratic framework and introduce dispersed information, we linearize the production function around low congestion, $cbE \ll 1$, and assume Gaussian uncertainty, $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. This yields the approximation (derived in the Appendix)

$$L(E, \theta) \approx \alpha E + \beta \theta E - \gamma E^2, \tag{35}$$

where $\alpha \equiv c\bar{a}$, $\beta \equiv c$, and $\gamma \equiv c^2 \bar{a} b > 0$. The term $-\gamma E^2$ captures the crowding externality: higher aggregate effort reduces marginal productivity for all agents.

Under open access, each agent receives the average product $L/E = \alpha + \beta\theta - \gamma E$. With convex effort costs $\frac{\phi}{2} e_i^2$, agent i's payoff is

$$\pi_i = e_i(\alpha + \beta\theta - \gamma E) - \frac{\phi}{2} e_i^2.$$

Agents take E as given and equate private marginal cost to the average product, failing

to internalize the effect of their effort on aggregate productivity. As in Gordon's original analysis, this leads to excessive and overly volatile aggregate effort.

The structure of this environment is isomorphic to the Romer application, with one key difference: the spillover parameter is negative. Aggregate effort reduces, rather than increases, individual marginal returns, implying strategic substitutability instead of complementarity. The team-efficient allocation internalizes the crowding term $-\gamma E$, resulting in lower and less volatile aggregate effort. In this setting, the Coasean transfer takes the form of a state-contingent tax

$$p(E) = -\gamma E,$$

applied to an agent's deviation from the average effort. This price coincides with the marginal external cost of aggregate effort. Agents forecast E using their private information about θ and optimally reduce effort in anticipation of the tax. As a result, the Coasean transfer implements the team-efficient allocation by pricing congestion in the commons.

Taken together with the Romer application, this example illustrates that outcome-contingent Coasean transfers restore efficiency in both positive and negative externality environments. The sign of the externality determines whether the mechanism operates through a subsidy or a tax, but the underlying market-equivalence logic is unchanged.

9 Conclusion

Externality regulation under dispersed information poses a fundamental challenge. Pigouvian instruments require centralized knowledge that the regulator cannot possess, while Coasean bargaining fails when agents are numerous. This tension has traditionally forced a trade-off between simplicity and efficiency.

This paper resolves the tension by recognizing that an externality is a missing market. The solution is to create that market through outcome-contingent Coasean transfers. Pricing agents' deviations from the cross-sectional average at a rate contingent on the realized aggregate outcome replicates the allocation that would arise if the externality were traded competitively. The regulator need only observe the outcome and know society's valuation of

it; information about fundamentals, signals, and agent types remains dispersed.

The analysis reveals a broader principle. Uncertainty and disagreement, often viewed as obstacles to policy implementation, can instead facilitate efficient decentralization. When agents face greater uncertainty or hold divergent beliefs, they place greater weight on private information, and outcome-contingent pricing aggregates these responses into the team-efficient allocation. Informational frictions therefore strengthen, rather than weaken, decentralized implementation.

The appropriate response to uncertainty is thus not to centralize information, but to design instruments that decentralize its use. Coasean transfers achieve this by mapping observable outcomes into prices. More generally, when aggregate outcomes summarize payoff-relevant uncertainty, contingent pricing can substitute for direct observability of fundamentals.

References

Anderson, R. M. and H. Duanmu (2025). Cap-and-trade and carbon tax meet arrow–debreu. *Econometrica 93*(2), 357–393.

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.

Arrow, K. J. (1969). The organization of economic activity: Issues pertinent to the choice of market versus nonmarket allocation. In *The Analysis and Evaluation of Public Expenditures: The PPB System*, Volume 1, pp. 47–64. Washington, D.C.: U.S. Government Printing Office. Joint Economic Committee, 91st Congress, 1st Session.

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

Bonnisseau, J.-M., E. L. del Mercato, and P. Siconolfi (2023). Existence of an equilibrium in arrowian markets for consumption externalities. *Journal of Economic Theory 209*, 105638.

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*.

Davila, J., J. H. Hong, P. Krusell, and J.-V. Ríos-Rull (2012). Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica 80*(6), 2431–2467.

Diamond, P. A. (1978). The role of a stock market in a general equilibrium model with technological uncertainty. In *Uncertainty in Economics*, pp. 209–229. Elsevier.

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.

Ellingsen, T. and E. Paltseva (2016). Confining the coase theorem: contracting, ownership, and free-riding. *The Review of Economic Studies 83*(2), 547–586.

Foley, D. K. (1970). Lindahl's solution and the core of an economy with public goods. *Econometrica 38*(1), 66–72.

Gelfand, I. M. and S. V. Fomin (1963). *Calculus of Variations*. Englewood Cliffs, N.J.: Prentice-Hall. Revised English edition.

Gordon, H. S. (1954). The economic theory of a common-property resource: The fishery. *Journal of Political Economy 62*(2), 124–142.

Greenwald, B. C. and J. E. Stiglitz (1986). Externalities in economies with imperfect information and incomplete markets. *The Quarterly Journal of Economics 101*(2), 229–264.

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

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

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

Li, J. (2023). Information aggregation via contracting. *The Journal of Finance 78*(2), 935–965.

Lindahl, E. (1958). Just taxation—a positive solution. In *Classics in the theory of public finance*, pp. 168–176. Springer.

Meade, J. E. (1952). External economies and diseconomies in a competitive situation. *Economic Journal 62*(245), 54–67.

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

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

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.

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

Stiglitz, J. E. (1982). The inefficiency of the stock market equilibrium. *The Review of Economic Studies 49*(2), 241–261.

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

A Appendix

Proof of Proposition 1

The social planner maximizes ex-post welfare state by state. Fix the state (θ, ν). The objective function

$$W(\{a_i\}, \theta, \nu) = \int_0^1 \left(e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i + S_i \left(\kappa \int_0^1 a_j dj + \Gamma(\theta, \nu) \right) \right) di \tag{A1}$$

is strictly concave in the actions $\{a_i\}_{i \in [0,1]}$ because the private cost functions are strictly convex ($\eta_i > 0$) and each social value function $S_i(\cdot)$ is strictly concave in total activity (Assumption 1). Thus, the first-best allocation exists, is unique, and is characterized by the first-order condition. Differentiating the objective function (A1) with respect to a_i yields

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

where $\bar{S}(\Omega) \equiv \int_0^1 S_j(\Omega) dj$.

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

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

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 (6). To establish the uniqueness of A^{fb}, define the mapping $\Psi(A) \equiv \bar{B}_H \theta + \frac{\kappa}{H} \bar{S}'(\kappa A + \Gamma(\theta, \nu))$. Since $S_i''(\cdot) < 0$ for all i, we have $\bar{S}''(\cdot) < 0$, so the function $\kappa \bar{S}'(\cdot)$ is strictly decreasing, which implies that $\Psi(A)$ is strictly decreasing in A. Because the left-hand side of (6) is strictly increasing in A while the right-hand side is strictly decreasing, there exists a unique fixed point $A^{fb} = \Psi(A^{fb})$. The individual actions $\{a_i^{fb}\}$ are then uniquely determined by the first-order condition (A2). □

Proof of Proposition 2

The utilitarian planner chooses decision rules $a_i(\mathcal{I}_i)$ to maximize expected social welfare

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

$$\text{subject to:} \quad a_i(\cdot) \text{ is } \mathcal{I}_i\text{-measurable for each } i; \quad A = \int_0^1 a_i di; \quad \Omega = \kappa A + \Gamma(\theta, \nu).$$

The planner's problem is a functional optimization problem. Unlike the first-best benchmark, in which the planner selects state-contingent actions $a_i \in \mathbb{R}$, the planner here chooses measurable decision rules $a_i : \mathcal{I}_i \to \mathbb{R}$. We characterize the optimum using a variational argument, following the approach in Davila et al. (2012). In particular, optimality requires that the functional derivative of the planner's objective vanish for all admissible $\mathcal{I}_i$-measurable perturbations.

Because Ω is common across agents, aggregate social value satisfies $\int_0^1 S_i(\Omega)\, di = \bar{S}(\Omega)$. By Assumption 1, each $S_i(\cdot)$ is strictly concave, and hence so is $\bar{S}(\cdot)$. Since Ω is an affine function of the allocation $\{a_i\}$, the composition $\bar{S}(\kappa A + \Gamma(\theta, \nu))$ is strictly concave in $\{a_i\}$. The private component of welfare is strictly concave in a_i for each agent because $\eta_i > 0$. It follows that, for any realization of (θ, ν), instantaneous welfare is strictly concave in $\{a_i\}$. Taking expectations preserves strict concavity, so the expected welfare functional $W(\{a_i(\cdot)\})$ is strictly concave in the decision rules. A unique global maximizer therefore exists.

Let $\lambda(\theta, \nu)$ denote the Lagrange multiplier associated with the constraint $A - \int_0^1 a_i di = 0$ in state (θ, ν). The Lagrangian is:

$$\Lambda = \mathbb{E}\left[\int_0^1 \left(e - \frac{\eta_i}{2} a_i^2 + \beta_i \theta a_i\right) di + \bar{S}(\Omega) + \lambda(\theta, \nu)\left(A - \int_0^1 a_i di\right)\right].$$

The first-order condition with respect to the aggregate variable A in state (θ, ν) is:

$$\lambda(\theta, \nu) = -\kappa \bar{S}'(\Omega).$$

The planner chooses the decision rule $a_i(\cdot)$ to maximize the Lagrangian. Because a_i must be $\mathcal{I}_i$-measurable, the optimal rule must satisfy the first-order condition pointwise conditional on $\mathcal{I}_i$. To see this, consider a perturbation $a_i(\cdot) + \varepsilon h(\cdot)$, where $h(\cdot)$ is an arbitrary $\mathcal{I}_i$-measurable function. The first-order condition requires the derivative $\frac{d\Lambda}{d\varepsilon}$ to be zero at the optimum (i.e., when $\varepsilon = 0$):

$$\left.\frac{d\Lambda}{d\varepsilon}\right|_{\varepsilon=0} = \mathbb{E}\left[\left(-\eta_i a_i + \beta_i \theta - \lambda(\theta, \nu)\right) h(\mathcal{I}_i)\right] = 0,$$

or $\mathbb{E}\left[\mathbb{E}\left[-\eta_i a_i + \beta_i \theta - \lambda(\theta, \nu) \mid \mathcal{I}_i\right] h(\mathcal{I}_i)\right] = 0$ by the Law of Iterated Expectations. Since this must hold for all $\mathcal{I}_i$-measurable functions h, the fundamental lemma of the calculus of variations (see, e.g., Gelfand and Fomin, 1963, Section 1.3, Lemma 1) implies:

$$\mathbb{E}\left[-\eta_i a_i + \beta_i \theta - \lambda(\theta, \nu) \mid \mathcal{I}_i\right] = 0.$$

Substituting the expression for the Lagrange multiplier $\lambda(\theta, \nu) = -\kappa \bar{S}'(\Omega)$ we have:

$$-\eta_i a_i + \beta_i \mathbb{E}[\theta \mid \mathcal{I}_i] + \mathbb{E}[\kappa \bar{S}'(\Omega) \mid \mathcal{I}_i] = 0.$$

Using the notation $m_{i,\theta} \equiv \mathbb{E}[\theta \mid \mathcal{I}_i]$ and $\mathbb{E}_i[\cdot] \equiv \mathbb{E}[\cdot \mid \mathcal{I}_i]$, and rearranging terms, we obtain equation (7):

$$a_i^{te} = \frac{\beta_i}{\eta_i} m_{i,\theta} + \frac{\kappa}{\eta_i} \mathbb{E}_i[\bar{S}'(\Omega^{te})]. \tag{A3}$$

To establish the existence and uniqueness of the equilibrium aggregate A^{te}, we appeal to the properties of the planner's objective function. The welfare functional is:

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

Under Assumption 1, each social value function $S_i(\cdot)$ is strictly concave, and hence so is $\bar{S}(\cdot)$. Because the outcome Ω is affine in the allocation $\{a_i\}$, the composition $\bar{S}(\Omega)$ is strictly concave in $\{a_i\}$. The private component $-\frac{\eta_i}{2} a_i^2$ is strictly concave for each agent since $\eta_i > 0$. It follows that the welfare functional $W(a_i(\cdot))$ is strictly concave on the convex set of admissible decision rules, defined as $\mathcal{I}_i$-measurable, square-integrable functions.

A strictly concave functional admits at most one maximizer. Existence follows because the feasible set is closed and the objective is coercive due to the quadratic cost term. The team-efficient allocation a_i^{te} is therefore unique, and the associated aggregate outcome $A^{te}(\theta, \nu) = \int_0^1 a_i^{te}\, di$ is uniquely determined. □

Proof of Proposition 3

Part (i) From Proposition 1, the first-best allocation satisfies the condition $\eta_i a_i^{fb} = \beta_i \theta + \kappa \bar{S}'(\Omega^{fb})$. Substituting the linear average marginal social value function from Assumption 2, $\bar{S}'(\Omega) = \bar{s}_1 - \bar{s}_2(\Omega - \mu_\Omega^{sq})$, yields the expression for the individual action:

$$a_i^{fb} = \frac{\beta_i}{\eta_i} \theta + \frac{\kappa}{\eta_i} [\bar{s}_1 - \bar{s}_2(\Omega^{fb} - \mu_\Omega^{sq})].$$

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:

$$A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H} [\bar{s}_1 - \bar{s}_2(\Omega^{fb} - \mu_\Omega^{sq})].$$

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

$$\Omega^{fb} = \kappa \bar{B}_H \theta + \frac{\kappa^2}{H}[\bar{s}_1 - \bar{s}_2(\Omega^{fb} - \mu_\Omega^{sq})] + \Gamma(\theta, \nu).$$

Solving for Ω^{fb} yields the closed-form expression stated in equation (9). Uniqueness is guaranteed by the strict concavity of the objective function ($\bar{S}''(\cdot) = -\bar{s}_2 < 0$).

Part (ii) The first-order conditions for the first-best (fb) allocation in (A2) and the team-efficient (te) allocation in (A3) are:

$$\eta_i a_i^{fb} = \beta_i \theta + \kappa \bar{S}'(\Omega^{fb}), \tag{A4}$$

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

Taking the conditional expectation $\mathbb{E}_i[\cdot]$ of (A4) and subtracting it from (A5) yields

$$\eta_i(a_i^{te} - \mathbb{E}_i[a_i^{fb}]) = \kappa \mathbb{E}_i\big[\bar{S}'(\Omega^{te}) - \bar{S}'(\Omega^{fb})\big]. \tag{A6}$$

Under Assumption 2, the average marginal social value is linear with slope $-\bar{s}_2$. Consequently,

$$\kappa \mathbb{E}_i\big[\bar{S}'(\Omega^{te}) - \bar{S}'(\Omega^{fb})\big] = -\kappa \bar{s}_2 \mathbb{E}_i\big[\Omega^{te} - \Omega^{fb}\big] = -\kappa^2 \bar{s}_2 \mathbb{E}_i\big[A^{te} - A^{fb}\big].$$

Using the definition $\alpha \equiv -\kappa^2 \bar{s}_2$, equation (A6) becomes

$$\eta_i(a_i^{te} - \mathbb{E}_i[a_i^{fb}]) = \alpha \mathbb{E}_i\big[A^{te} - A^{fb}\big]. \tag{A7}$$

From Proposition 1, $A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H}\bar{S}'(\Omega^{fb})$, where, using equation (A4), we obtain

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

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

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

Substituting this expression into (A7) and rearranging terms yields the strategic form (10):

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

Proof of Proposition 4

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

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

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

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

Rearranging terms yields the optimal decision rule for the agent:

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

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

Proof of Proposition 5

We first establish the coefficients for the first-best aggregate action A^{fb}. From Proposition 1, A^{fb} satisfies the fixed-point equation (6): $A^{fb} = \bar{B}_H \theta + \frac{\kappa}{H} \bar{S}'(\kappa A^{fb} + \Gamma(\theta, \nu))$. Under Assumptions 2 and 3, this becomes the linear equation $HA^{fb} = H\bar{B}_H \theta + \kappa[\bar{s}_1 - \bar{s}_2(\kappa A^{fb} + \theta + \nu - \mu_\Omega^{sq})]$. Using $\alpha \equiv -\kappa^2 \bar{s}_2$, substituting $-\kappa \bar{s}_2 = \alpha/\kappa$, and collecting terms on A^{fb}, we obtain $A^{fb}(H - \alpha) = (H\bar{B}_H + \alpha/\kappa)\theta + (\alpha/\kappa)\nu + \kappa \bar{s}_1 - (\alpha/\kappa)\mu_\Omega^{sq}$. Thus, the first-best aggregate action is $A^{fb} = J_0 + J_\theta \theta + J_\nu \nu$, with coefficients

$$J_0 \equiv \frac{\kappa \bar{s}_1 - (\alpha/\kappa)\mu_\Omega^{sq}}{H - \alpha}, \quad J_\theta \equiv \frac{H\bar{B}_H + \alpha/\kappa}{H - \alpha}, \quad \text{and} \quad J_\nu \equiv \frac{\alpha/\kappa}{H - \alpha}. \tag{A8}$$

For the team-efficient aggregate action A^{te}, we postulate the linear form $A^{te} = C_0 + C_\theta \theta + C_\nu \nu$. From equation (A7) in the proof of Proposition 3, 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

$$A^{te} - \int_0^1 \mathbb{E}_i[a_i^{fb}]di = \frac{\alpha}{H}\int_0^1 (\mathbb{E}_i[A^{te}] - \mathbb{E}_i[A^{fb}])di. \tag{A9}$$

Under the affine information structure of Assumption 3, for any linear variable $X = K_0 + K_\theta\theta + K_\nu\nu$, the cross-sectional average of expectations is $\int_0^1 \mathbb{E}_i[X]di = K_0 + K_\theta\bar{m}_\theta + K_\nu\bar{m}_\nu$, where $\bar{m}_\theta = (1-\lambda_\theta)\mu_\theta + \lambda_\theta\theta$, and $\bar{m}_\nu = (1-\lambda_\nu)\mu_\nu + \lambda_\nu\nu$. 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_0, J_θ, J_ν) as A^{fb}, but applied to the average posterior beliefs. Specifically, write $a_i^{fb} = c_{0i} + c_{\theta i}\theta + c_{\nu i}\nu$, where the coefficients depend on agent i's type. Then $\mathbb{E}_i[a_i^{fb}] = c_{0i} + c_{\theta i}m_{i,\theta} + c_{\nu i}m_{i,\nu}$. Integrating:

$$\int_0^1 \mathbb{E}_i[a_i^{fb}]\,di = \int_0^1 c_{0i}\,di + \left(\int_0^1 c_{\theta i}\,di\right)\bar{m}_\theta + \left(\int_0^1 c_{\nu i}\,di\right)\bar{m}_\nu = J_0 + J_\theta\bar{m}_\theta + J_\nu\bar{m}_\nu,$$

where $\int_0^1 c_{zi}di = J_z$ for $z \in \{0, \theta, \nu\}$. Applying this to equation (A9), we equate the coefficients on the fundamental θ:

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

Rearranging to solve for C_θ gives

$$C_\theta = J_\theta\lambda_\theta\frac{H-\alpha}{H-\alpha\lambda_\theta}.$$

Substituting the expression for J_θ from (A8), we obtain

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

Similarly, from (A9), equating coefficients for ν gives

$$C_\nu - J_\nu\lambda_\nu = \frac{\alpha}{H}(C_\nu\lambda_\nu - J_\nu\lambda_\nu),$$

which, using (A8), leads to

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

Finally, from (A9), equating the constant terms yields

$$C_0 = J_0 + \frac{HJ_\theta(1-\lambda_\theta)\mu_\theta}{H-\alpha\lambda_\theta} + \frac{HJ_\nu(1-\lambda_\nu)\mu_\nu}{H-\alpha\lambda_\nu}.$$

To verify that uncertainty strengthens responsiveness, consider the magnitude $|C_\theta|$. Let $N \equiv |H\bar{B}_H + \alpha/\kappa|$ and $D \equiv -\alpha > 0$ (since $\alpha = -\kappa^2\bar{s}_2 < 0$). Then $|C_\theta| = \frac{N\lambda_\theta}{H+D\lambda_\theta}$. Differentiating with respect to λ_θ yields $\frac{\partial|C_\theta|}{\partial\lambda_\theta} = \frac{NH}{(H+D\lambda_\theta)^2}$, which is strictly positive provided $N > 0$. The same logic applies to $|C_\nu|$ with $N = |\alpha/\kappa| > 0$. Thus, as prior uncertainty rises (increasing λ_θ or λ_ν), the aggregate action becomes strictly more sensitive to the state. □

Proof of Proposition 6

We solve for the aggregate action A^{te} implied by the team-efficient allocation $\{a_i^{te}\}_{i\in[0,1]}$ when agents distrust others' signals. The derivation of the strategic form in Part (ii) of Proposition 3 carries over unchanged: agent i holds correct beliefs about the fundamentals (θ, ν), and distrust affects only the forecast of A^{te}. Thus,

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

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

$$A^{te} = \left(1 - \frac{\alpha}{H}\right)\int_0^1 \mathbb{E}_i[a_i^{fb}]di + \frac{\alpha}{H}\int_0^1 \mathbb{E}_i[A^{te}]di. \tag{A10}$$

The first term, $\int_0^1 \mathbb{E}_i[a_i^{fb}]\,di$, is unaffected by distrust because agents hold undistorted beliefs about the fundamentals. From the proof of Proposition 5, $A^{fb} = J_0 + J_\theta\theta + J_\nu\nu$, with coefficients given in equation (A8). Write $a_i^{fb} = c_{0i} + c_{\theta i}\theta + c_{\nu i}\nu$, where the coefficients depend on agent i's type. Then $\mathbb{E}_i[a_i^{fb}] = c_{0i} + c_{\theta i}m_{i,\theta} + c_{\nu i}m_{i,\nu}$. Integrating and using independence of types and signals:

$$\int_0^1 \mathbb{E}_i[a_i^{fb}]\,di = J_0 + J_\theta\bar{m}_\theta + J_\nu\bar{m}_\nu, \tag{A11}$$

where $\bar{m}_\theta = \lambda_\theta\theta + (1-\lambda_\theta)\mu_\theta$ and $\bar{m}_\nu = \lambda_\nu\nu + (1-\lambda_\nu)\mu_\nu$.

For the second term in (A10), $\int_0^1 \mathbb{E}_i[A^{te}]di$, we must explicitly account for how agent i forecasts the actions of other agents $j \neq i$ under distrust. We first conjecture the linear form

$$A^{te} = C_0 + C_\theta\theta + C_\nu\nu.$$

Because $A^{te} = \int_0^1 a_j^{te} dj$, the individual equilibrium strategy a_j^{te} must be linear in the agent's private beliefs $m_{j,\theta}$ and $m_{j,\nu}$. Focusing on the fundamental θ, we note that the average belief is $\bar{m}_\theta = \int_0^1 m_{j,\theta} dj = \lambda_\theta \theta + (1-\lambda_\theta)\mu_\theta$. For the aggregate action to load on θ with coefficient C_θ, the individual action a_j^{te} must load on the belief $m_{j,\theta}$ with coefficient C_θ/λ_θ. The same applies to ν. Thus, we can write the individual strategy as:

$$a_j^{te} = K_0 + \frac{C_\theta}{\lambda_\theta} m_{j,\theta} + \frac{C_\nu}{\lambda_\nu} m_{j,\nu},$$

where K_0 is a constant determined by the requirement that the strategy aggregates to C_0. Since $\int_0^1 a_j^{te} dj = A^{te}$, K_0 must satisfy

$$C_0 = K_0 + \frac{C_\theta}{\lambda_\theta}(1-\lambda_\theta)\mu_\theta + \frac{C_\nu}{\lambda_\nu}(1-\lambda_\nu)\mu_\nu. \tag{A12}$$

Now consider agent i's expectation of agent j's action, denoted by $\mathbb{E}_i[a_j^{te}]$. This depends on agent i's expectation of j's belief:

$$\mathbb{E}_i[a_j^{te}] = K_0 + \frac{C_\theta}{\lambda_\theta}\mathbb{E}_i[m_{j,\theta}] + \frac{C_\nu}{\lambda_\nu}\mathbb{E}_i[m_{j,\nu}]. \tag{A13}$$

Recall that $m_{j,z} = (1-\lambda_z)\mu_z + \lambda_z y_j$ for $z \in \{\theta, \nu\}$. The distrust assumption in equation (22) implies that agent i views agent j's signal as having a dampened correlation with the truth. Specifically, $\mathbb{E}_i[y_j^i] = \mu_z + \varphi_z(m_{i,z} - \mu_z)$. Substituting this into the expectation of beliefs:

$$\begin{aligned}
\mathbb{E}_i[m_{j,z}^i] &= (1-\lambda_z)\mu_z + \lambda_z \mathbb{E}_i[y_j^i] \\
&= (1-\lambda_z)\mu_z + \lambda_z\big[\mu_z + \varphi_z(m_{i,z} - \mu_z)\big] \\
&= \mu_z + \lambda_z \varphi_z (m_{i,z} - \mu_z).
\end{aligned}$$

Substituting this back into the expression (A13) for $\mathbb{E}_i[a_j^{te}]$:

$$\begin{aligned}
\mathbb{E}_i[a_j^{te}] &= K_0 + \frac{C_\theta}{\lambda_\theta}\Big(\mu_\theta + \lambda_\theta\varphi_\theta(m_{i,\theta} - \mu_\theta)\Big) + \frac{C_\nu}{\lambda_\nu}\Big(\mu_\nu + \lambda_\nu\varphi_\nu(m_{i,\nu} - \mu_\nu)\Big) \\
&= \underbrace{\left[K_0 + C_\theta\mu_\theta\left(\frac{1}{\lambda_\theta} - \varphi_\theta\right) + C_\nu\mu_\nu\left(\frac{1}{\lambda_\nu} - \varphi_\nu\right)\right]}_{\text{Constant} \equiv K_0'} + C_\theta\varphi_\theta m_{i,\theta} + C_\nu\varphi_\nu m_{i,\nu}.
\end{aligned}$$

Substituting K_0 from equation (A12) into the definition of K_0' yields

$$K_0' = C_0 + C_\theta\mu_\theta(1-\varphi_\theta) + C_\nu\mu_\nu(1-\varphi_\nu).$$

We aggregate the expectations to find agent i's forecast of the aggregate action:

$$\mathbb{E}_i[A^{te}] = \mathbb{E}_i\left[\int_0^1 a_j^{te} dj\right] = \int_0^1 \mathbb{E}_i[a_j^{te}]dj = K_0' + C_\theta\varphi_\theta m_{i,\theta} + C_\nu\varphi_\nu m_{i,\nu}. \tag{A14}$$

Note that since $\mathbb{E}_i[a_j^{te}]$ is independent of j, the integral over j simply returns the integrand.

We can now integrate (A14) over the population i (where $\int_0^1 m_{i,z}di = \bar{m}_z$):

$$\int_0^1 \mathbb{E}_i[A^{te}]di = K_0' + C_\theta\varphi_\theta\bar{m}_\theta + C_\nu\varphi_\nu\bar{m}_\nu. \tag{A15}$$

Substitute (A11) and (A15) into the fixed-point equation (A10) to solve for the coefficients:

$$C_0 + C_\theta\theta + C_\nu\nu = \left(1 - \frac{\alpha}{H}\right)\left(J_0 + J_\theta\bar{m}_\theta + J_\nu\bar{m}_\nu\right) + \frac{\alpha}{H}\left(K_0' + C_\theta\varphi_\theta\bar{m}_\theta + C_\nu\varphi_\nu\bar{m}_\nu\right).$$

Using $\bar{m}_z = (1-\lambda_z)\mu_z + \lambda_z z$, matching coefficients on θ and solving for C_θ, we obtain:

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

Substituting J_θ from (A8):

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

Similarly, matching coefficients on ν and substituting J_ν from (A8) implies:

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

Finally, matching constant terms yields:

$$C_0 = J_0 + J_\theta\mu_\theta + J_\nu\mu_\nu + \frac{J_\theta(\alpha-H)\lambda_\theta\mu_\theta}{H - \alpha\lambda_\theta\varphi_\theta} + \frac{J_\nu(\alpha-H)\lambda_\nu\mu_\nu}{H - \alpha\lambda_\nu\varphi_\nu},$$

with J_0, J_θ, and J_ν given in (A8). Setting $\varphi_\theta = \varphi_\nu = 1$ recovers the coefficients in Proposition 5. Since $\alpha < 0$, $|C_\theta|$ is decreasing in φ_θ and $|C_\nu|$ is decreasing in φ_ν. □

The following lemma establishes that the Coasean transfer implements the team-efficient allocation under distrust.

Lemma 1. *The Coasean transfer scheme $T_i = (a_i - A)p(\Omega)$ with the pricing rule $p(\Omega) = \kappa\bar{S}'(\Omega)$ implements the team-efficient allocation characterized in Proposition 6, regardless of the degree of distrust parameterized by φ_θ and φ_ν.*

Proof. Consider the problem of agent i who distrusts others' signals, facing the Coasean transfer $T_i = (a_i - A)p(\Omega)$. The agent chooses a_i to maximize their subjective expected payoff:

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

where the expectation $\mathbb{E}_i$ is computed using the distrustful beliefs about others' signals defined in equation (22). Since agents are atomless, they treat the aggregate action A and the outcome Ω (and thus the price $p(\Omega)$) as independent of their individual choice a_i. The first-order condition with respect to a_i yields:

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

This condition matches the general form of the team-efficient equilibrium derived in Proposition 2: $\eta_i a_i^{te} = \beta_i m_{i,\theta} + \kappa \mathbb{E}_i[\bar{S}'(\Omega^{te})]$. In the present context, the operator $\mathbb{E}_i[\cdot]$ reflects the agent's subjective (distrustful) beliefs defined in equation (22). If the Coasean transfer sets the price equal to the marginal social value, $p(\Omega) = \kappa \bar{S}'(\Omega)$, the agent's optimization condition becomes identical to the team-efficient condition. Thus, Coasean transfers implement the team-efficient equilibrium subject to the constraint of agent distrust. □

Proof of Proposition 7

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

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

Since agents are atomless, they treat the aggregates as exogenous. This yields the decision rule

$$a_i(\tau) = \frac{\beta_i m_{i,\theta} - \tau}{\eta_i}. \tag{A16}$$

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

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

where $\Omega(\tau) = \kappa A(\tau) + \Gamma(\theta, \nu)$, and $A(\tau)$ is the aggregate action:

$$A(\tau) = \bar{B}_H \bar{m}_\theta - \frac{\tau}{H}, \tag{A17}$$

with $\bar{m}_\theta \equiv \int_0^1 m_{i,\theta} di$. Defining the average social value $\bar{S}(\Omega) \equiv \int_0^1 S_i(\Omega) di$, we obtain:

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

Because the best response $a_i(\tau)$ in (A16) is linear in τ, the private-cost term $-\frac{\eta_i}{2} a_i(\tau)^2$ is strictly concave in τ. Since $\Omega(\tau)$ is linear in τ and $\bar{S}$ is strictly concave, $\bar{S}(\Omega(\tau))$ is also strictly concave in τ, so the objective is strictly concave. Differentiating with respect to τ and using (A16), we have

$$\frac{d\mathbb{E}[W(\tau)]}{d\tau} = \mathbb{E}\left[\int_0^1 \left(-\eta_i a_i(\tau) + \beta_i \theta\right)\left(-\frac{1}{\eta_i}\right) di + \bar{S}'(\Omega(\tau)) \frac{\partial \Omega(\tau)}{\partial \tau}\right]. \tag{A18}$$

Using the individual first-order condition (A16), $-\eta_i a_i(\tau) + \beta_i \theta = \tau + \beta_i(\theta - m_{i,\theta})$, and noting that $\mathbb{E}[\theta - m_{i,\theta}] = 0$, the integral term in (A18) reduces to $-\tau/H$. Since $\Omega(\tau) = \kappa A(\tau) + \Gamma(\theta, \nu)$ is affine in τ with slope $-\kappa/H$, it follows that

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

Setting this derivative to zero characterizes the optimal tax/subsidy as

$$\tau^* = -\kappa \mathbb{E}[\bar{S}'(\Omega(\tau^*))]. \tag{A19}$$

Under Assumption 2, the linearity of the average marginal social value implies $-\tau^* = \kappa \bar{s}_1 - \kappa \bar{s}_2 (\mathbb{E}[\Omega(\tau^*)] - \mu_\Omega^{sq})$. Using $\mathbb{E}[\Omega(\tau^*)] = \kappa \mathbb{E}[A(\tau^*)] + \mathbb{E}[\Gamma(\theta, \nu)]$ and (A17), which gives $\mathbb{E}[A(\tau)] = \bar{B}_H \mathbb{E}[\theta] - \tau/H$, yields (25):

$$\tau^* = \frac{-\kappa \bar{s}_1 + \kappa \bar{s}_2 \left(\kappa \bar{B}_H \mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq}\right)}{1 + \kappa^2 \bar{s}_2 / H}. \tag{A20}$$

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

Proof of Corollary 7.1

With uninformative private information, the team-efficient first-order condition (7) becomes

$$\eta_i a_i^{te} = \beta_i \mathbb{E}[\theta] + \kappa \mathbb{E}[\bar{S}'(\Omega^{te})].$$

Define the constant wedge $\tau^{te} \equiv -\kappa \mathbb{E}[\bar{S}'(\Omega^{te})]$. Then

$$a_i(\tau^{te}) = \frac{\beta_i \mathbb{E}[\theta] - \tau^{te}}{\eta_i},$$

which coincides with the private best response under a uniform tax (or subsidy) τ^{te} in (A16) when private information is uninformative. Aggregating and using $\Omega^{te} = \Omega(\tau^{te}) = \kappa A(\tau^{te}) + \Gamma(\theta, \nu)$ yields:

$$\tau^{te} = -\kappa \mathbb{E}[\bar{S}'(\Omega(\tau^{te}))],$$

which coincides with the fixed-point condition (A19). By Proposition 7, $\tau^{te} = \tau^*$, so the induced allocations coincide. □

Proof of Proposition 8

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 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 acts as a price-taker, choosing a demand schedule $a_i(P)$ to maximize expected utility:

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

where $\Omega = \kappa A + \Gamma(\theta, \nu)$. Since $S_i(\Omega)$, P, and R do not depend on a_i, the first-order condition yields the agent's optimal demand:

$$a_i(P) = \frac{\beta_i m_{i,\theta} - P}{\eta_i}. \tag{A21}$$

Aggregating across agents, using independence of types and posterior beliefs, and defining $\bar{m}_\theta \equiv \int_0^1 m_{i,\theta} di$, we obtain the aggregate demand under price P:

$$A(P) = \int_0^1 a_i(P) di = \bar{B}_H \bar{m}_\theta - \frac{P}{H}.$$

Given the cap A, the market clearing condition $A(P) = A$ determines the unique permit price:

$$P(A) = H \bar{B}_H \bar{m}_\theta - HA. \tag{A22}$$

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 is

$$\mathbb{E}[W(A)] = \mathbb{E}\Big[\int_0^1 \Big(e - \frac{\eta_i}{2}a_i(A)^2 + \beta_i\theta a_i(A) + S_i(\Omega)\Big)di\Big].$$

By separating the social value term and defining the average $\bar{S}(\Omega) \equiv \int_0^1 S_i(\Omega)di$, we obtain:

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

Consider the private payoff component of welfare, $W_{priv}(A) \equiv \int_0^1 (e - \frac{\eta_i}{2}a_i(A)^2 + \beta_i\theta a_i(A))di$. We differentiate this term with respect to A. From the individual best response (A21) $a_i(P(A)) = (\beta_i m_{i,\theta} - P(A))/\eta_i$ and the price sensitivity from (A22) $dP(A)/dA = -H$, 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:

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

Thus, the first-order condition for the regulator is:

$$\frac{d\mathbb{E}[W(A)]}{dA} = H\bar{B}_H\mathbb{E}[\theta] - HA + \kappa\mathbb{E}[\bar{S}'(\kappa A + \Gamma(\theta,\nu))] = 0,$$

which under Assumption 2 becomes:

$$0 = H\bar{B}_H\mathbb{E}[\theta] - HA^{cap*} + \kappa\Big[\bar{s}_1 - \bar{s}_2\big(\kappa A^{cap*} + \mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\big)\Big]. \tag{A23}$$

Solving for A^{cap*} gives the unique welfare-maximizing cap in equation (28):

$$A^{cap*} = \frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa\bar{s}_1 - \kappa\bar{s}_2\big(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\big)}{H + \kappa^2\bar{s}_2}.$$

Since $\bar{S}''(\Omega) = -\bar{s}_2 < 0$, expected welfare is strictly concave in A, so A^{cap*} is unique.

Equivalence in expectation. We establish the two equalities stated in the text:

$$\mathbb{E}[P(A^{cap*})] = \tau^* \quad \text{and} \quad \mathbb{E}[A(\tau^*)] = A^{cap*} \tag{A24}$$

For the first equality, compute first the expected permit price by applying the law of iterated expectations to (A22):

$$\mathbb{E}[P(A)] = H\bar{B}_H\mathbb{E}[\theta] - HA,$$

then substitute the optimal cap A^{cap*}:

$$\mathbb{E}[P(A^{cap*})] = H\bar{B}_H\mathbb{E}[\theta] - H\left[\frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa\bar{s}_1 - \kappa\bar{s}_2(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq})}{H + \kappa^2\bar{s}_2}\right].$$

Simplifying yields:

$$\mathbb{E}[P(A^{cap*})] = \frac{-\kappa\bar{s}_1 + \kappa\bar{s}_2\left(\kappa\bar{B}_H\mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\right)}{1 + \kappa^2\bar{s}_2/H} = \tau^*,$$

where the second equality follows from equation (A20) in the proof of Proposition 7.

For the second equality, note that, from equation (A17), $A(\tau^*) = \bar{B}_H\bar{m}_\theta - \tau^*/H$. Using the expression for $P(A)$ in equation (A22) and $\mathbb{E}[\bar{m}_\theta] = \mathbb{E}[\theta]$ we have

$$\mathbb{E}[P(A^{cap*})] = \tau^* = H\bar{B}_H\mathbb{E}[\theta] - HA^{cap*},$$

so $\tau^*/H = \bar{B}_H\mathbb{E}[\theta] - A^{cap*}$ and we obtain

$$A(\tau^*) = \bar{B}_H\left(\bar{m}_\theta - \mathbb{E}[\theta]\right) + A^{cap*}.$$

Taking expectations we obtain $\mathbb{E}[A(\tau^*)] = A^{cap*}$. □

Proof of Proposition 9

Total expected welfare is the sum of a private component and a social component,

$$\mathbb{E}[W] = \mathbb{E}[W_{priv}] + \mathbb{E}[W_{soc}],$$

where $\mathbb{E}[W_{priv}] = \mathbb{E}\left[\int_0^1 \left(e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i\right) di\right]$ and $\mathbb{E}[W_{soc}] = \mathbb{E}[\bar{S}(\Omega)]$. We analyze each in turn.

Private component. Under the tax and cap regimes, individual actions are $a_i^{tax} = (\beta_i m_{i,\theta} - \tau^*)/\eta_i$ and $a_i^{cap} = (\beta_i m_{i,\theta} - P)/\eta_i$, respectively; see (A16) and (A21). The difference is $\Delta a_i = (\tau^* - P)/\eta_i = \delta/\eta_i$, where $\delta \equiv \tau^* - P$. We evaluate the welfare difference by performing a second-order expansion of the private payoff W_{priv}^{cap} around the tax allocation a_i^{tax}. Since private payoffs are quadratic, this expansion is exact:

$$W_{priv}^{cap} - W_{priv}^{tax} = \int_0^1 (-\eta_i a_i^{tax} + \beta_i\theta)\Delta a_i\, di - \frac{1}{2}\int_0^1 \eta_i(\Delta a_i)^2\, di.$$

Substituting $\Delta a_i = \delta/\eta_i$ and using the agent's first-order condition (A16) under the Pigouvian

regime, $-\eta_i a_i^{tax} + \beta_i\theta = \tau^* + \beta_i(\theta - m_{i,\theta})$, yields

$$W_{priv}^{cap} - W_{priv}^{tax} = \frac{\delta\tau^*}{H} + \delta\bar{B}_H(\theta - \bar{m}_\theta) - \frac{\delta^2}{2H}. \tag{A25}$$

Under Assumption 3, the forecast error is $\theta - \bar{m}_\theta = (1-\lambda_\theta)(\theta - \mu_\theta)$ and $\bar{m}_\theta - \mu_\theta = \lambda_\theta(\theta - \mu_\theta)$. We first characterize the moments of $\delta = \tau^* - P$. Since $\mathbb{E}[P] = \tau^*$ by (A24), we have $\mathbb{E}[\delta] = 0$. Using the equilibrium cap-and-trade permit price (A22), $P = H\bar{B}_H\bar{m}_\theta - HA^{cap*}$, δ is

$$\delta = \mathbb{E}[P] - P = -H\bar{B}_H(\bar{m}_\theta - \mu_\theta) = -H\bar{B}_H\lambda_\theta(\theta - \mu_\theta). \tag{A26}$$

This implies that the variance of δ is $\mathrm{Var}(\delta) = H^2\bar{B}_H^2\lambda_\theta^2\sigma_\theta^2 = \mathrm{Var}(P)$.

Taking expectations of the welfare difference in equation (A25), the first term vanishes because $\mathbb{E}[\delta] = 0$. Using (A26), the expected cross-term in (A25) is

$$\begin{aligned}\mathbb{E}\left[\delta\bar{B}_H(\theta - \bar{m}_\theta)\right] &= \mathbb{E}\left[(-H\bar{B}_H\lambda_\theta(\theta - \mu_\theta))\bar{B}_H(1-\lambda_\theta)(\theta - \mu_\theta)\right] \\ &= -H\bar{B}_H^2\lambda_\theta(1-\lambda_\theta)\sigma_\theta^2.\end{aligned}$$

The expected quadratic term in (A25) is

$$\mathbb{E}\left[-\frac{\delta^2}{2H}\right] = -\frac{1}{2H}\mathrm{Var}(\delta) = -\frac{H\bar{B}_H^2\lambda_\theta^2}{2}\sigma_\theta^2.$$

Combining these results yields the expected private welfare loss:

$$\begin{aligned}\mathbb{E}[W_{priv}^{cap}] - \mathbb{E}[W_{priv}^{tax}] &= -H\bar{B}_H^2\lambda_\theta(1-\lambda_\theta)\sigma_\theta^2 - \frac{H\bar{B}_H^2\lambda_\theta^2}{2}\sigma_\theta^2 \\ &= -H\bar{B}_H^2\left(1 - \frac{\lambda_\theta}{2}\right)\sigma_\theta^2\lambda_\theta.\end{aligned} \tag{A27}$$

Noting that $\mathrm{Var}(P) = \mathrm{Var}(\delta) = H^2\bar{B}_H^2\lambda_\theta^2\sigma_\theta^2$, this loss can be interpreted as $-\frac{2-\lambda_\theta}{2H\lambda_\theta}\mathrm{Var}(P)$, showing that the expected private welfare loss is driven by price volatility in the permit market.

Social component. The outcomes are $\Omega^{cap} = \kappa A^{cap*} + \Gamma(\theta,\nu)$ and $\Omega^{tax} = \kappa A(\tau^*) + \Gamma(\theta,\nu)$. From equation (A24), $\mathbb{E}[A(\tau^*)] = A^{cap*}$, so the expected outcomes are identical: $\mathbb{E}[\Omega^{cap}] = \mathbb{E}[\Omega^{tax}]$.

Under Assumption 2, the average social value function is quadratic with constant second derivative $\bar{S}''(\Omega) = -\bar{s}_2$. We use an exact second-order Taylor expansion around the mean $\mathbb{E}[\Omega]$ to express the expected social value:

$$\mathbb{E}[\bar{S}(\Omega)] = \mathbb{E}\left[\bar{S}(\mathbb{E}[\Omega]) + \bar{S}'(\mathbb{E}[\Omega])(\Omega - \mathbb{E}[\Omega]) - \frac{\bar{s}_2}{2}(\Omega - \mathbb{E}[\Omega])^2\right] = \bar{S}(\mathbb{E}[\Omega]) - \frac{\bar{s}_2}{2}\mathrm{Var}(\Omega), \tag{A28}$$

leading to

$$\mathbb{E}[W_{soc}^{cap}] - \mathbb{E}[W_{soc}^{tax}] = \left(\bar{S}(\mathbb{E}[\Omega^{cap}]) - \frac{\bar{s}_2}{2}\text{Var}(\Omega^{cap})\right) - \left(\bar{S}(\mathbb{E}[\Omega^{tax}]) - \frac{\bar{s}_2}{2}\text{Var}(\Omega^{tax})\right)$$

$$= \frac{\bar{s}_2}{2}\left[\text{Var}(\Omega^{tax}) - \text{Var}(\Omega^{cap})\right]. \tag{A29}$$

Under Assumption 3, $\Gamma(\theta, \nu) = \theta + \nu$. Cap-and-trade fixes the aggregate action at A^{cap*}, so $\text{Var}(\Omega^{cap}) = \sigma_\theta^2 + \sigma_\nu^2$. Under the Pigouvian tax, equation (A17), $A(\tau^*) = \bar{B}_H \bar{m}_\theta - \tau^*/H$, and $\bar{m}_\theta = (1 - \lambda_\theta)\mu_\theta + \lambda_\theta \theta$ imply

$$\Omega^{tax} = (1 + \kappa \bar{B}_H \lambda_\theta)\theta + \nu + \kappa \bar{B}_H (1 - \lambda_\theta)\mu_\theta - \frac{\kappa \tau^*}{H},$$

which gives $\text{Var}(\Omega^{tax}) = (1 + \kappa \bar{B}_H \lambda_\theta)^2 \sigma_\theta^2 + \sigma_\nu^2$. The difference in variances is then

$$\text{Var}(\Omega^{tax}) - \text{Var}(\Omega^{cap}) = \kappa \bar{B}_H \lambda_\theta (2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2.$$

Substituting in equation (A29) yields

$$\mathbb{E}[W_{soc}^{cap}] - \mathbb{E}[W_{soc}^{tax}] = \frac{\bar{s}_2 \kappa \bar{B}_H}{2}(2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2 \lambda_\theta. \tag{A30}$$

Total welfare. Combining the private (A27) and social (A30) components yields equation (29):

$$\mathbb{E}[W^{cap}] - \mathbb{E}[W^{tax}] = -H\bar{B}_H^2 \left(1 - \frac{\lambda_\theta}{2}\right)\sigma_\theta^2 \lambda_\theta + \frac{\bar{s}_2 \kappa \bar{B}_H}{2}(2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2 \lambda_\theta. \qquad \square$$

Proof of Proposition 10

The proof has two parts. The first characterizes equilibrium under cap-and-trade with price revelation and shows that the optimal cap is unchanged. The second derives the welfare difference between the Coasean transfer system and cap-and-trade with price revelation.

Part 1: Cap-and-Trade with Price Discovery. We first characterize the competitive permit price for a given cap A under a Rational Expectations Equilibrium (REE). We then determine the optimal 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. Agents observe P and the cap A, and thus understand that the price contains information about θ. Conjecture that the equilibrium price takes the linear form:

$$P = a_\theta \theta + a_A A.$$

Given this conjecture, each agent infers $\theta = (P - a_A A)/a_\theta$ and chooses the action

$$a_i(P) = \frac{\beta_i \theta - P}{\eta_i}. \tag{A31}$$

Aggregating across agents yields

$$A = \bar{B}_H \theta - \frac{P}{H},$$

and thus $a_\theta = H\bar{B}_H$ and $a_A = -H$. The fully revealing equilibrium permit price is then

$$P(A) = H\bar{B}_H \theta - HA. \tag{A32}$$

The cap A is chosen ex ante to maximize expected social welfare. The regulator does not directly observe θ. With budget-balanced permit revenues, expected welfare is:

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

As in the proof of Proposition 8, consider the private payoff component, $W_{priv}(A)$, and differentiate with respect to A. Substituting the first-order condition $\beta_i \theta - \eta_i a_i = P(A)$, the derivative of the private component equals the competitive permit price:

$$\frac{dW_{priv}}{dA} = \int_0^1 (\beta_i \theta - \eta_i a_i) \frac{da_i}{dA} di = P(A) \int_0^1 \frac{da_i}{dA} di = P(A).$$

To see that $\int_0^1 \frac{da_i}{dA} di = 1$, note from (A31) that a_i depends on A only through the price $P(A)$ obtained in (A32). Thus, $\frac{da_i}{dA} = \frac{-1}{\eta_i} \frac{dP}{dA} = \frac{H}{\eta_i}$. Integrating across agents yields $\int_0^1 \frac{da_i}{dA} di = H \int_0^1 \frac{1}{\eta_i} di = 1$. Thus, the regulator's first-order condition is:

$$\frac{d\mathbb{E}[W(A)]}{dA} = \mathbb{E}\big[P(A) + \kappa \bar{S}'(\kappa A + \Gamma(\theta, \nu))\big] = 0.$$

From (A32), we compute the expected permit price:

$$\mathbb{E}[P(A)] = H\bar{B}_H \mathbb{E}[\theta] - HA. \tag{A33}$$

The expected price expression (A33) depends on $\mathbb{E}[\theta]$, just as in the cap-and-trade case without learning from prices, and similarly for the expected marginal social value. Therefore, the optimality condition is identical to the one derived in the proof of Proposition 8, equation (A23):

$$0 = H\bar{B}_H \mathbb{E}[\theta] - HA^{cap*} + \kappa\Big[\bar{s}_1 - \bar{s}_2\big(\kappa A^{cap*} + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq}\big)\Big].$$

Thus, the optimal cap A^{cap*} remains as in Proposition 8, equation (28).

We further verify the equalities

$$\mathbb{E}[P(A^{cap*})] = \tau^* \quad \text{and} \quad \mathbb{E}[A(\tau^*)] = A^{cap*}. \tag{A34}$$

For the first equality, substitute the optimal cap A^{cap*} into the expected inverse demand (A33):

$$\mathbb{E}[P(A^{cap*})] = H\bar{B}_H\mathbb{E}[\theta] - H\left[\frac{H\bar{B}_H\mathbb{E}[\theta] + \kappa\bar{s}_1 - \kappa\bar{s}_2(\mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq})}{H + \kappa^2\bar{s}_2}\right].$$

Simplifying yields the expression for τ^* derived in Proposition 7:

$$\mathbb{E}[P(A^{cap*})] = \frac{-\kappa\bar{s}_1 + \kappa\bar{s}_2\left(\kappa\bar{B}_H\mathbb{E}[\theta] + \mathbb{E}[\Gamma(\theta,\nu)] - \mu_\Omega^{sq}\right)}{1 + \kappa^2\bar{s}_2/H} = \tau^*.$$

For the second equality, note that, from equation (A17), $A(\tau^*) = \bar{B}_H\bar{m}_\theta - \tau^*/H$. Using the expression for $P(A)$ in equation (A32) and $\mathbb{E}[\bar{m}_\theta] = \mathbb{E}[\theta]$ we have

$$\mathbb{E}[P(A^{cap*})] = \tau^* = H\bar{B}_H\mathbb{E}[\theta] - HA^{cap*},$$

so $\tau^*/H = \bar{B}_H\mathbb{E}[\theta] - A^{cap*}$ and we obtain

$$A(\tau^*) = \bar{B}_H\left(\bar{m}_\theta - \mathbb{E}[\theta]\right) + A^{cap*}.$$

Taking expectations we obtain $\mathbb{E}[A(\tau^*)] = A^{cap*}$. This is a cross-regime comparison between cap-and-trade with learning and Pigouvian taxation without learning. This result holds because the aggregate action is linear in the average belief $\bar{m}_\theta$. Since $\mathbb{E}[\bar{m}_\theta] = \mathbb{E}[\theta]$ by the Law of Iterated Expectations, the differing information structures do not affect the expected aggregate action.

Finally, since $\Omega = \kappa A + \Gamma(\theta,\nu)$, we also note that

$$\mathbb{E}[\Omega^{cap}] = \mathbb{E}[\Omega^{tax}].$$

Part 2: Welfare analysis. We evaluate the welfare difference $\mathbb{E}[W^{ct}] - \mathbb{E}[W^{cap}]$ by decomposing it into two parts relative to the optimal Pigouvian tax benchmark $\mathbb{E}[W^{tax}]$:

$$\mathbb{E}[W^{ct}] - \mathbb{E}[W^{cap}] = \underbrace{\left(\mathbb{E}[W^{ct}] - \mathbb{E}[W^{tax}]\right)}_{\text{Lemma 2}} - \underbrace{\left(\mathbb{E}[W^{cap}] - \mathbb{E}[W^{tax}]\right)}_{\text{Lemma 3}}$$

where W^{cap} denotes welfare under a cap-and-trade regime with full price revelation. Lemma 2 below confirms that Coasean transfers strictly dominate the Pigouvian tax whenever agents hold private information. Lemma 3 extends the Weitzman (1974) prices-versus-quantities trade-off of

Proposition 9 to the cross-regime comparison between the cap-and-trade regime with full price revelation and Pigouvian taxation.

Lemma 2. *The welfare gain of Coasean transfers relative to the Pigouvian tax is*

$$\mathbb{E}[W^{ct}] - \mathbb{E}[W^{tax}] = \frac{\alpha^2 \lambda_\theta (1 + \kappa \bar{B}_H \lambda_\theta)^2 \sigma_\theta^2}{2\kappa^2 (H - \alpha\lambda_\theta)} + \frac{\alpha^2 \lambda_\nu \sigma_\nu^2}{2\kappa^2 (H - \alpha\lambda_\nu)}. \tag{A35}$$

Proof. As in the proof of Proposition 9, we decompose welfare into private and social components, $W = W_{priv} + W_{soc}$, and analyze each in turn. The individual actions under each regime are

$$a_i^{ct} = \frac{\beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega^{ct})]}{\eta_i} \quad \text{and} \quad a_i^{tax} = \frac{\beta_i m_{i,\theta} - \tau^*}{\eta_i}. \tag{A36}$$

Define the action difference $\Delta a_i \equiv a_i^{ct} - a_i^{tax} = \delta_i/\eta_i$, where $\delta_i \equiv \mathbb{E}_i[p(\Omega^{ct})] + \tau^*$ represents agent i's perceived price deviation from the fixed tax.

A second-order (exact) Taylor expansion of W_{priv}^{ct} around a_i^{tax} leads to:

$$W_{priv}^{ct} - W_{priv}^{tax} = \int_0^1 \left(-\eta_i a_i^{tax} + \beta_i \theta\right) \Delta a_i di - \frac{1}{2} \int_0^1 \eta_i (\Delta a_i)^2 di.$$

Using the Pigouvian first-order condition (A16), $-\eta_i a_i^{tax} + \beta_i \theta = \tau^* + \beta_i(\theta - m_{i,\theta})$, and substituting $\Delta a_i = \delta_i/\eta_i$, we obtain

$$W_{priv}^{ct} - W_{priv}^{tax} = \frac{\tau^*}{H} \int_0^1 \delta_i di + \bar{B}_H \int_0^1 \delta_i (\theta - m_{i,\theta}) di - \frac{1}{2H} \int_0^1 \delta_i^2 di. \tag{A37}$$

We first show that the expected aggregate actions coincide, i.e., $\mathbb{E}[A^{ct}] = \mathbb{E}[A^{tax}]$. Aggregating the Coasean transfer actions a_i^{ct} from (A36) yields $A^{ct} = \bar{B}_H \bar{m}_\theta + \frac{\kappa}{H} \int_0^1 \mathbb{E}_i[\bar{S}'(\Omega^{ct})] di$. Taking unconditional expectations and using the linearity of $\bar{S}'(\cdot)$ leads to

$$\mathbb{E}[A^{ct}] = \bar{B}_H \mu_\theta + \frac{\kappa}{H} \left(\bar{s}_1 - \bar{s}_2(\kappa \mathbb{E}[A^{ct}] + \mathbb{E}[\Gamma(\theta, \nu)] - \mu_\Omega^{sq})\right). \tag{A38}$$

Similarly, aggregating a_i^{tax} from (A36) yields $A^{tax} = \bar{B}_H \bar{m}_\theta - \tau^*/H$, which together with the Pigouvian condition $\tau^* = -\kappa \mathbb{E}[\bar{S}'(\Omega^{tax})] = -\mathbb{E}[p(\Omega^{tax})]$ yields an equation for $\mathbb{E}[A^{tax}]$ identical to (A38). Uniqueness implies $\mathbb{E}[A^{ct}] = \mathbb{E}[A^{tax}]$. Since Ω is linear in A and $p(\Omega)$ is linear in Ω, we also have $\mathbb{E}[p(\Omega^{ct})] = \mathbb{E}[p(\Omega^{tax})]$. Therefore, $\mathbb{E}[\delta_i] = \mathbb{E}[p(\Omega^{ct})] + \tau^* = \mathbb{E}[p(\Omega^{ct})] - \mathbb{E}[p(\Omega^{tax})] = 0$.

To evaluate the expectation of the second term in (A37), we note that δ_i is $\mathcal{I}_i$-measurable and $\mathbb{E}_i[\theta - m_{i,\theta}] = 0$. Thus $\mathbb{E}[\delta_i(\theta - m_{i,\theta})] = \mathbb{E}[\delta_i \cdot \mathbb{E}_i[\theta - m_{i,\theta}]] = 0$.

Consequently, the expected private welfare difference (A37) reduces to

$$\mathbb{E}[W^{ct}_{priv}] - \mathbb{E}[W^{tax}_{priv}] = -\frac{1}{2H}\text{Var}(\delta_i).$$

Using the price function $p(\Omega) = \kappa[\bar{s}_1 - \bar{s}_2(\Omega - \mu^{sq}_\Omega)]$ and the outcome $\Omega^{ct} = \kappa A^{ct} + \theta + \nu$, the deviation is $\delta_i = \kappa\bar{s}_1 - \kappa\bar{s}_2(\mathbb{E}_i[\Omega^{ct}] - \mu^{sq}_\Omega) + \tau^*$, implying $\text{Var}(\delta_i) = \kappa^2\bar{s}_2^2\text{Var}(\mathbb{E}_i[\Omega^{ct}])$.
By Proposition 4, the Coasean transfer implements the team-efficient allocation, hence $A^{ct} = A^{te}$. As shown in the proof of Proposition 5, the team-efficient aggregate action is linear, $A^{te} = C_0 + C_\theta\theta + C_\nu\nu$. Thus, the conditional expectation of the outcome is $\mathbb{E}_i[\Omega^{ct}] = \kappa C_0 + (1 + \kappa C_\theta)m_{i,\theta} + (1 + \kappa C_\nu)m_{i,\nu}$. Since individual posteriors satisfy $\text{Var}(m_{i,z}) = \lambda_z\sigma_z^2$ for $z \in \{\theta, \nu\}$ and $\bar{s}_2 = -\alpha/\kappa^2$, we have

$$\text{Var}(\delta_i) = \frac{\alpha^2}{\kappa^2}\left[(1 + \kappa C_\theta)^2\lambda_\theta\sigma_\theta^2 + (1 + \kappa C_\nu)^2\lambda_\nu\sigma_\nu^2\right].$$

Substituting this variance into the expected private welfare difference $\mathbb{E}[W^{ct}_{priv} - W^{tax}_{priv}]$ yields

$$\mathbb{E}[W^{ct}_{priv}] - \mathbb{E}[W^{tax}_{priv}] = -\frac{1}{2H}\frac{\alpha^2}{\kappa^2}\left[(1 + \kappa C_\theta)^2\lambda_\theta\sigma_\theta^2 + (1 + \kappa C_\nu)^2\lambda_\nu\sigma_\nu^2\right].$$

Next, we evaluate the social welfare difference. Equation (A28) shows that $\mathbb{E}[\bar{S}(\Omega)] = \bar{S}(\mathbb{E}[\Omega]) - \frac{\bar{s}_2}{2}\text{Var}(\Omega)$. Expected outcomes are identical across regimes, $\mathbb{E}[\Omega^{ct}] = \mathbb{E}[\Omega^{tax}]$, thus:

$$\mathbb{E}[W^{ct}_{soc}] - \mathbb{E}[W^{tax}_{soc}] = \frac{\bar{s}_2}{2}\left[\text{Var}(\Omega^{tax}) - \text{Var}(\Omega^{ct})\right].$$

Under Pigouvian taxation, substituting $A^{tax} = \bar{B}_H[\lambda_\theta\theta + (1 - \lambda_\theta)\mu_\theta] - \tau^*/H$ into Ω^{tax} yields $\text{Var}(\Omega^{tax}) = (1 + \kappa\bar{B}_H\lambda_\theta)^2\sigma_\theta^2 + \sigma_\nu^2$. Under Coasean transfers, substituting $A^{ct} = C_0 + C_\theta\theta + C_\nu\nu$ into Ω^{ct} yields $\text{Var}(\Omega^{ct}) = (1 + \kappa C_\theta)^2\sigma_\theta^2 + (1 + \kappa C_\nu)^2\sigma_\nu^2$.
Summing the private and social components, substituting $\bar{s}_2 = -\alpha/\kappa^2$, and using the expressions for C_θ and C_ν from the proof of Proposition 5, yields the decomposition in (A35). □

Lemma 3. *The welfare difference between cap-and-trade (with full price revelation) and the Pigouvian tax is:*

$$\mathbb{E}[W^{cap}] - \mathbb{E}[W^{tax}] = \frac{\sigma_\theta^2}{2}\left[(1 - \lambda_\theta)\overline{B^2}_H - H\bar{B}_H^2\right] + \frac{\bar{s}_2\kappa\bar{B}_H}{2}(2 + \kappa\bar{B}_H\lambda_\theta)\sigma_\theta^2\lambda_\theta. \tag{A39}$$

Proof. Under the Pigouvian tax τ^*, the agent relies on the private signal $m_{i,\theta}$. Under cap-and-trade, the price P reveals θ. Therefore the agents optimize using the true value of the fundamental θ.

The individual actions are $a_i^{tax} = (\beta_i m_{i,\theta} - \tau^*)/\eta_i$ and $a_i^{cap} = (\beta_i \theta - P)/\eta_i$. We define the action difference $\Delta a_i = a_i^{cap} - a_i^{tax}$. Let $\xi_i \equiv \theta - m_{i,\theta}$ denote the agent's forecast error under the tax, and let $\delta \equiv \tau^* - P$ denote the price surprise under the cap. We can express the difference as $\Delta a_i = (\beta_i \xi_i + \delta)/\eta_i$. As before,

$$W_{priv}^{cap} - W_{priv}^{tax} = \int_0^1 (-\eta_i a_i^{tax} + \beta_i \theta)\Delta a_i di - \frac{1}{2}\int_0^1 \eta_i (\Delta a_i)^2 di.$$

Using the agent's first-order condition under the tax, $-\eta_i a_i^{tax} + \beta_i m_{i,\theta} - \tau^* = 0$, we substitute $-\eta_i a_i^{tax} + \beta_i \theta = \tau^* + \beta_i \xi_i$. The integrand becomes $(\tau^* + \beta_i \xi_i)\left(\frac{\beta_i \xi_i + \delta}{\eta_i}\right) - \frac{\eta_i}{2}\left(\frac{\beta_i \xi_i + \delta}{\eta_i}\right)^2$. Expanding the terms yields

$$W_{priv}^{cap} - W_{priv}^{tax} = \int_0^1 \frac{1}{\eta_i}\left[\tau^*(\beta_i \xi_i + \delta) + \frac{1}{2}\beta_i^2 \xi_i^2 - \frac{1}{2}\delta^2\right] di.$$

Since $\mathbb{E}[\xi_i] = \mathbb{E}[\delta] = 0$, the linear terms $\tau^*(\beta_i \xi_i + \delta)$ vanish in expectation (see (A34)). The expected private welfare difference is then

$$\mathbb{E}[W_{priv}^{cap}] - \mathbb{E}[W_{priv}^{tax}] = \int_0^1 \frac{1}{2\eta_i}\beta_i^2 \mathrm{Var}(\xi_i) di - \int_0^1 \frac{1}{2\eta_i}\mathrm{Var}(\delta) di.$$

Under standard Bayesian updating, individual posteriors satisfy $m_{i,\theta} = (1 - \lambda_\theta)\mu_\theta + \lambda_\theta(\theta + \varepsilon_i)$ with ε_i i.i.d. mean-zero. By the law of total variance, $\mathrm{Var}(\xi_i) = (1 - \lambda_\theta)\sigma_\theta^2$. For δ, recall that $P = H\bar{B}_H \theta - HA$ and $\tau^* = \mathbb{E}[P]$. Thus $\delta = -H\bar{B}_H(\theta - \mathbb{E}[\theta])$, implying $\mathrm{Var}(\delta) = H^2 \bar{B}_H^2 \sigma_\theta^2$. Substituting these variances and defining $\overline{B^2}_H \equiv \int_0^1 (\beta_i^2/\eta_i) di$, the private component is:

$$\mathbb{E}[W_{priv}^{cap}] - \mathbb{E}[W_{priv}^{tax}] = \frac{1}{2}(1 - \lambda_\theta)\sigma_\theta^2 \overline{B^2}_H - \frac{H\bar{B}_H^2}{2}\sigma_\theta^2.$$

We next consider the social component. As before, welfare differences are driven entirely by outcome variances:

$$\mathbb{E}[W_{soc}^{cap}] - \mathbb{E}[W_{soc}^{tax}] = \frac{\bar{s}_2}{2}\left[\mathrm{Var}(\Omega^{tax}) - \mathrm{Var}(\Omega^{cap})\right].$$

Under Pigouvian taxation, substituting the aggregate action $A^{tax} = \bar{B}_H \bar{m}_\theta - \tau^*/H$ into the outcome equation yields $\Omega^{tax} = (1 + \kappa \bar{B}_H \lambda_\theta)\theta + \nu + \text{const}$. Thus, $\mathrm{Var}(\Omega^{tax}) = (1 + \kappa \bar{B}_H \lambda_\theta)^2 \sigma_\theta^2 + \sigma_\nu^2$. Under cap-and-trade, the aggregate action is fixed at A^{cap*}, so $\Omega^{cap} = \kappa A^{cap*} + \theta + \nu$, implying $\mathrm{Var}(\Omega^{cap}) = \sigma_\theta^2 + \sigma_\nu^2$. Computing the difference in variances yields:

$$\mathrm{Var}(\Omega^{tax}) - \mathrm{Var}(\Omega^{cap}) = \kappa \bar{B}_H \lambda_\theta (2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2.$$

Substituting this back into the welfare difference yields the final social component:

$$\mathbb{E}[W_{soc}^{cap}] - \mathbb{E}[W_{soc}^{tax}] = \frac{\bar{s}_2 \kappa \bar{B}_H}{2}(2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2 \lambda_\theta.$$

Combining the private and social components yields the expression (A39) in the lemma. □

Derivation of equation (30). Subtracting the welfare difference in Lemma 3 from that in Lemma 2 yields the total difference $\mathbb{E}[W^{ct}] - \mathbb{E}[W^{cap}] = (\mathbb{E}[W^{ct}] - \mathbb{E}[W^{tax}]) - (\mathbb{E}[W^{cap}] - \mathbb{E}[W^{tax}])$.
We group the resulting terms by source. The term involving σ_ν^2 comes solely from Coasean transfers (Lemma 2) and corresponds to the exogenous stabilization gain:

$$\frac{\alpha^2 \lambda_\nu \sigma_\nu^2}{2\kappa^2(H - \alpha\lambda_\nu)}.$$

The term involving heterogeneity comes solely from the private component of cap-and-trade (Lemma 3) and corresponds to the allocative efficiency loss:

$$-\frac{1}{2}(1 - \lambda_\theta)(\overline{B^2}_H - H\bar{B}_H^2)\sigma_\theta^2.$$

The remaining terms involving σ_θ^2 combine to form the fundamental stabilization gain. Specifically, subtracting the stabilization term in Lemma 3 from the θ-term in Lemma 2 yields:

$$\frac{\alpha^2 \lambda_\theta (1 + \kappa \bar{B}_H \lambda_\theta)^2}{2\kappa^2(H - \alpha\lambda_\theta)}\sigma_\theta^2 - \frac{\bar{s}_2 \kappa \bar{B}_H}{2}(2 + \kappa \bar{B}_H \lambda_\theta)\sigma_\theta^2 \lambda_\theta.$$

Using $\bar{s}_2 = -\alpha/\kappa^2$ and simplifying the algebraic fraction results in:

$$\frac{\lambda_\theta(\alpha + \kappa H \bar{B}_H)^2 \sigma_\theta^2}{2\kappa^2(H - \alpha\lambda_\theta)}.$$

Summing these three components yields the final expression in equation (30). □

Proof of Proposition 11

The agent's objective functions in the four regimes are given by:

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

$$\text{Pigouvian Tax (eq. (23))}: \quad u_i^{tax} = \mathbb{E}_i\left[e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i + S_i(\Omega) + (a_i - A)(-\tau)\right],$$

$$\text{Cap-and-Trade (eq. (26))}: \quad u_i^{cap} = \mathbb{E}_i\left[e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i + S_i(\Omega) + (a_i - A)(-P(A))\right],$$

$$\text{Coasean Transfer (eq. (18))}: \quad u_i^{ct} = \mathbb{E}_i\left[e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i + S_i(\Omega) + (a_i - A)p(\Omega)\right].$$

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

$$u_i = \mathbb{E}_i\left[-\frac{\eta_i}{2}a_i^2 + a_i M_i\right] + C_i, \tag{A40}$$

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 $C_i = e + \mathbb{E}_i[S_i(\Omega)]$; in the Pigouvian Tax regime, $M_i = \beta_i\theta - \tau$ and $C_i = e + \mathbb{E}_i[S_i(\Omega) + R]$; in the Cap-and-Trade regime, $M_i = \beta_i\theta - P$ and $C_i = e + \mathbb{E}_i[S_i(\Omega) + R]$; and in the Coasean Transfer regime, $M_i = \beta_i\theta + p(\Omega)$ and $C_i = e + \mathbb{E}_i[S_i(\Omega) - Ap(\Omega)]$.

The first-order condition is $\eta_i a_i = \mathbb{E}_i[M_i]$. 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, η_i is a deterministic parameter. Considering the terms that depend on a_i, we have:

$$\mathbb{E}\left[-\frac{\eta_i}{2}a_i^2 + a_i M_i\right] = -\frac{\eta_i}{2}\mathbb{E}[a_i^2] + \mathbb{E}[a_i M_i].$$

By the Law of Iterated Expectations, $\mathbb{E}[a_i M_i] = \mathbb{E}[a_i\mathbb{E}_i[M_i]]$. Using the first-order condition $\mathbb{E}_i[M_i] = \eta_i a_i$, we have $\mathbb{E}[a_i M_i] = \mathbb{E}[a_i(\eta_i a_i)] = \eta_i\mathbb{E}[a_i^2]$. Substituting this back in (A40) yields:

$$\mathbb{E}[u_i] = -\frac{\eta_i}{2}\mathbb{E}[a_i^2] + \eta_i\mathbb{E}[a_i^2] + \mathbb{E}[C_i] = \frac{\eta_i}{2}\mathbb{E}[a_i^2] + \mathbb{E}[C_i]. \tag{A41}$$

Let the optimal action be linear in the posterior belief: $a_i = K_i m_{i,\theta} + Z_i$, where Z_i is uncorrelated with $m_{i,\theta}$, a conjecture that we will verify in all cases. Using $\mathbb{E}[m_{i,\theta}] = \mu_\theta$, the expected squared action is:

$$\mathbb{E}[a_i^2] = K_i^2(\mathrm{Var}(m_{i,\theta}) + \mu_\theta^2) + \underbrace{2K_i\mu_\theta\mathbb{E}[Z_i] + \mathbb{E}[Z_i^2]}_{\text{Terms independent of } m_{i,\theta}}, \tag{A42}$$

where we used the property of variance, $\mathbb{E}[m_{i,\theta}^2] = \text{Var}(m_{i,\theta}) + \mu_\theta^2$. The agent chooses information precision to maximize this value. Since $\text{Var}(m_{i,\theta})$ increases with information precision, the marginal benefit of acquiring information is directly proportional to the coefficient K_i^2. By (A41) and (A42), the comparison across regulatory systems reduces to comparing squared responsiveness K_i^2.

Status Quo: The agent maximizes $\max_{a_i} \mathbb{E}_i[e - \frac{\eta_i}{2}a_i^2 + \beta_i\theta a_i]$. From the first-order condition we obtain:

$$a_i = \frac{\beta_i}{\eta_i} m_{i,\theta}.$$

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

$$K_i^{sq} = \frac{\beta_i}{\eta_i}.$$

Pigouvian Tax: The agent faces a constant tax τ set ex-ante. From the first-order condition we obtain:

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

Since τ is constant, the derivative remains unchanged:

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

Cap-and-Trade: From Proposition 8, the agent chose price-contingent actions:

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

In the absence of learning from prices, equation (A22) shows that the market-clearing price is linear in the aggregate belief $\bar{m}_\theta$, i.e., $P = H\bar{B}_H\bar{m}_\theta - HA$. To identify K_i^{cap}, we compute the coefficient in the linear projection of a_i onto $m_{i,\theta}$:

$$K_i^{cap} = \frac{\text{Cov}(a_i, m_{i,\theta})}{\text{Var}(m_{i,\theta})}.$$

Starting from the demand schedule and the price equation, we have:

$$\text{Cov}(a_i, m_{i,\theta}) = \frac{1}{\eta_i}\left[\beta_i \text{Var}(m_{i,\theta}) - H\bar{B}_H \text{Cov}(m_{i,\theta}, \bar{m}_\theta)\right]. \tag{A43}$$

Since $\bar{m}_\theta = \lambda_\theta\theta + (1-\lambda_\theta)\mu_\theta$, $\text{Cov}(\bar{m}_\theta, m_{i,\theta}) = \lambda_\theta \text{Cov}(\theta, m_{i,\theta})$. To evaluate $\text{Cov}(\theta, m_{i,\theta})$, we apply

the law of total covariance with respect to $\mathcal{I}_i$:

$$\text{Cov}(\theta, m_{i,\theta}) = \mathbb{E}[\text{Cov}(\theta, m_{i,\theta} \mid \mathcal{I}_i)] + \text{Cov}(\mathbb{E}[\theta \mid \mathcal{I}_i], \mathbb{E}[m_{i,\theta} \mid \mathcal{I}_i]).$$

The first term vanishes because $m_{i,\theta}$ is $\mathcal{I}_i$-measurable and thus constant conditional on $\mathcal{I}_i$. For the second term, $\mathbb{E}[\theta \mid \mathcal{I}_i] = m_{i,\theta}$ and $\mathbb{E}[m_{i,\theta} \mid \mathcal{I}_i] = m_{i,\theta}$. Therefore, $\text{Cov}(\theta, m_{i,\theta}) = \text{Var}(m_{i,\theta})$, and thus $\text{Cov}(\bar{m}_\theta, m_{i,\theta}) = \lambda_\theta \text{Var}(m_{i,\theta})$. Substituting back in equation (A43), we have $\text{Cov}(a_i, m_{i,\theta}) = \frac{1}{\eta_i}\left[\beta_i - H\bar{B}_H\lambda_\theta\right]\text{Var}(m_{i,\theta})$. Therefore the responsiveness coefficient is

$$K_i^{cap} = \frac{\beta_i - H\bar{B}_H\lambda_\theta}{\eta_i}.$$

For most agents, cap-and-trade dampens the incentive to acquire information relative to the status quo. The incentive to acquire information under cap-and-trade exceeds that under the status quo if and only if $(K_i^{cap})^2 > (K_i^{sq})^2$, or equivalently:

$$(\beta_i - H\bar{B}_H\lambda_\theta)^2 > \beta_i^2,$$

which simplifies to:

$$\beta_i < \frac{H\bar{B}_H\lambda_\theta}{2}.$$

Coasean Transfer: We derive K_i^{ct} directly from the agent's first-order condition (20) under the Coasean transfer:

$$\eta_i a_i = \beta_i m_{i,\theta} + \mathbb{E}_i[p(\Omega)]. \tag{A44}$$

Using the price $p(\Omega) = \kappa[\bar{s}_1 - \bar{s}_2(\Omega - \mu_\Omega^{sq})]$ and the definition $\alpha \equiv -\kappa^2\bar{s}_2$, we can express the expected price as $\mathbb{E}_i[p(\Omega)] = \kappa\bar{s}_1 + \frac{\alpha}{\kappa}(\mathbb{E}_i[\Omega] - \mu_\Omega^{sq})$. Substituting the outcome $\Omega = \kappa A^{ct} + \theta + \nu$ in equation (A44) yields:

$$\eta_i a_i = \beta_i m_{i,\theta} + \kappa\bar{s}_1 - \frac{\alpha}{\kappa}\mu_\Omega^{sq} + \frac{\alpha}{\kappa}\mathbb{E}_i[\kappa A^{ct} + \theta + \nu].$$

Using $A^{ct} = C_0 + C_\theta\theta + C_\nu\nu$ and taking expectations we obtain:

$$\eta_i a_i = \left[\kappa\bar{s}_1 + \alpha C_0 - \frac{\alpha}{\kappa}\mu_\Omega^{sq}\right] + \left[\beta_i + \alpha\left(C_\theta + \frac{1}{\kappa}\right)\right]m_{i,\theta} + \left[\alpha\left(C_\nu + \frac{1}{\kappa}\right)\right]m_{i,\nu}. \tag{A45}$$

Hence,

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

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

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

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

$$\alpha Z\left(2\beta_i + \alpha Z\right) > 0, \tag{A46}$$

where $Z \equiv (C_\theta + 1/\kappa)$. Substituting C_θ from Proposition 5:

$$Z = \frac{(\bar{B}_H H + \alpha/\kappa)\lambda_\theta}{H - \alpha\lambda_\theta} + \frac{1}{\kappa} = \frac{H - \alpha\lambda_\theta + \kappa\bar{B}_H H\lambda_\theta + \alpha\lambda_\theta}{\kappa(H - \alpha\lambda_\theta)} = \frac{H(1 + \kappa\bar{B}_H\lambda_\theta)}{\kappa(H - \alpha\lambda_\theta)}. \tag{A47}$$

Since $H, \kappa, \bar{B}_H > 0$, $\lambda_\theta \in [0, 1]$, and $\alpha < 0$, $Z > 0$. Given $Z > 0$ and $\alpha < 0$, for the inequality (A46) to hold, the term $2\beta_i + \alpha Z$ must be strictly negative, implying:

$$\beta_i < -\frac{\alpha}{2}Z.$$

Substituting the expression for Z from (A47) yields the condition (31) stated in the proposition. □

Proof of Proposition 12

As shown in the proof of Proposition 11, the agent's ex-ante expected payoff in any linear-quadratic regime takes the form in equation (A41), $\mathbb{E}[u_i] = \frac{\eta_i}{2}\mathbb{E}[a_i^2] + \mathbb{E}[C_i]$. 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, Pigouvian Tax, and Cap-and-Trade. In the status quo, the optimal action is $a_i^{sq} = \beta_i m_{i,\theta}/\eta_i$ (see equation (3)). Under a Pigouvian tax and under cap-and-trade, optimal actions are $a_i^{tax} = (\beta_i m_{i,\theta} - \tau)/\eta_i$ and $a_i^{cap} = (\beta_i m_{i,\theta} - P)/\eta_i$, respectively (see equations (24) and (27)). In all three regimes, actions depend only on posterior beliefs about the fundamental θ and fixed parameters, and are independent of beliefs about the noise term ν. It follows that $K_{i,\nu}^{sq} = K_{i,\nu}^{tax} = K_{i,\nu}^{cap} = 0$, so the marginal value of information about ν is zero.

Coasean Transfer. The individual action under the Coasean transfer satisfies equation (A45). The responsiveness to the belief about ν is therefore identified directly as:

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

To determine if this is non-zero, we substitute C_ν from Proposition 5:

$$K_{i,\nu}^{ct} = \frac{\alpha}{\eta_i}\frac{H}{\kappa(H - \alpha\lambda_\nu)}.$$

Since $H > 0$ and $\alpha < 0$, the denominator is strictly positive, so $(K_{i,\nu}^{ct})^2 > 0$. Because the agent's optimal action responds to $m_{i,\nu}$, the marginal value of information regarding ν is strictly positive. □

Proof of Proposition 13

We compare the equilibrium payoffs under the Coasean transfer regime (denoted by ct) and the benchmark regulation (denoted by $\hat{\cdot}$), characterized by the marginal price $\hat{P}$ and induced actions $\hat{a}_i, \hat{A}$. Note that $\hat{P}$ acts as a random variable within the expectation operator $\mathbb{E}_i$, covering both the deterministic Pigouvian tax case ($\hat{P} = \tau^*$) and the stochastic cap-and-trade case ($\hat{P} = P$).
Under the Coasean transfer, the agent's ex-post payoff is

$$u_i^{ct} = e - \frac{\eta_i}{2}(a_i^{ct})^2 + \beta_i\theta a_i^{ct} + S_i(\Omega^{ct}) + (a_i^{ct} - A^{ct})p(\Omega^{ct}),$$

where the pricing function is $p(\Omega) = \kappa\bar{S}'(\Omega)$. The first-order condition is

$$\mathbb{E}_i[-\eta_i a_i^{ct} + \beta_i\theta + p(\Omega^{ct})] = 0. \tag{A48}$$

Under the benchmark regulation, the agent pays $\hat{P}$ per unit of action and receives a lump-sum rebate proportional to the aggregate $\hat{A}$, such that the net transfer is $-\hat{P}(\hat{a}_i - \hat{A})$. The payoff is

$$\hat{u}_i = e - \frac{\eta_i}{2}(\hat{a}_i)^2 + \beta_i\theta\hat{a}_i + S_i(\hat{\Omega}) + (\hat{a}_i - \hat{A})(-\hat{P}).$$

The first-order condition in the benchmark regime is $-\eta_i\hat{a}_i + \beta_i m_{i,\theta} - \hat{P} = 0$, which implies

$$\eta_i\hat{a}_i = \beta_i m_{i,\theta} - \hat{P}.$$

Let $\Delta u_i = u_i^{ct} - \hat{u}_i$, and define the differences in allocations as $\Delta a_i = a_i^{ct} - \hat{a}_i$ and $\Delta A = A^{ct} - \hat{A}$:

$$\Delta u_i = -\frac{\eta_i}{2}[(a_i^{ct})^2 - (\hat{a}_i)^2] + \beta_i\theta\Delta a_i + S_i(\Omega^{ct}) - S_i(\hat{\Omega}) + (a_i^{ct} - A^{ct})p(\Omega^{ct}) + (\hat{a}_i - \hat{A})\hat{P}. \tag{A49}$$

We proceed by expanding the payoff difference Δu_i. The difference in the private cost term is

$$-\frac{\eta_i}{2}[(a_i^{ct})^2 - (\hat{a}_i)^2] = -\eta_i \hat{a}_i \Delta a_i - \frac{\eta_i}{2}(\Delta a_i)^2. \tag{A50}$$

The difference in social value, using the exact second-order expansion for the quadratic function $S_i(\cdot)$ around $\hat{\Omega}$, is

$$S_i(\Omega^{ct}) - S_i(\hat{\Omega}) = S_i'(\hat{\Omega})(\kappa \Delta A) - \frac{s_{2i}}{2}(\kappa \Delta A)^2. \tag{A51}$$

We next decompose the net transfer difference Δu_i in (A49). The Coasean transfer is $(a_i^{ct} - A^{ct})p(\Omega^{ct})$ and the benchmark transfer is $-\hat{P}(\hat{a}_i - \hat{A})$. Substituting $a_i^{ct} = \hat{a}_i + \Delta a_i$ and $A^{ct} = \hat{A} + \Delta A$, we write

$$(a_i^{ct} - A^{ct})p(\Omega^{ct}) = (\hat{a}_i - \hat{A})p(\Omega^{ct}) + (\Delta a_i - \Delta A)p(\Omega^{ct}).$$

The total transfer difference is therefore

$$(a_i^{ct} - A^{ct})p(\Omega^{ct}) + (\hat{a}_i - \hat{A})\hat{P} = (\hat{a}_i - \hat{A})(p(\Omega^{ct}) + \hat{P}) + (\Delta a_i - \Delta A)p(\Omega^{ct}). \tag{A52}$$

Substituting (A50), (A51), and (A52) into (A49), we obtain

$$\begin{aligned} \Delta u_i = {} & (-\eta_i \hat{a}_i + \beta_i \theta + p(\Omega^{ct}))\Delta a_i - \frac{\eta_i}{2}(\Delta a_i)^2 \\ & + \left(S_i'(\hat{\Omega})\kappa - p(\Omega^{ct})\right)\Delta A - \frac{s_{2i}}{2}\kappa^2(\Delta A)^2 \\ & + (\hat{a}_i - \hat{A})(p(\Omega^{ct}) + \hat{P}). \end{aligned} \tag{A53}$$

Substituting $a_i^{ct} = \hat{a}_i + \Delta a_i$ into (A48) and using $\mathcal{I}_i$-measurability of Δa_i yields

$$\mathbb{E}_i[-\eta_i \hat{a}_i + \beta_i \theta + p(\Omega^{ct})] = \eta_i \Delta a_i. \tag{A54}$$

Taking the conditional expectation of the terms involving Δa_i in (A53) and using (A54) yields:

$$\mathbb{E}_i[(-\eta_i \hat{a}_i + \beta_i \theta + p(\Omega^{ct}))\Delta a_i] - \frac{\eta_i}{2}(\Delta a_i)^2 = \Delta a_i \cdot \eta_i \Delta a_i - \frac{\eta_i}{2}(\Delta a_i)^2 = \frac{\eta_i}{2}(\Delta a_i)^2.$$

Since $\Omega = \kappa A + \Gamma(\theta, \nu)$, the outcome difference is $\Omega^{ct} - \hat{\Omega} = \kappa \Delta A$. Under Assumption 2, the price difference is linear in the aggregate difference:

$$p(\Omega^{ct}) - p(\hat{\Omega}) = \kappa[\bar{S}'(\Omega^{ct}) - \bar{S}'(\hat{\Omega})] = -\kappa^2 \bar{s}_2 \Delta A. \tag{A55}$$

For the ΔA terms in the second line of (A53), we use $p(\Omega^{ct}) = \kappa \bar{S}'(\Omega^{ct})$ together with (A55) to

obtain

$$\bar{S}'(\Omega^{ct}) = \bar{S}'(\hat{\Omega}) - \bar{s}_2\kappa\Delta A \quad \text{and} \quad p(\Omega^{ct}) = \kappa\bar{S}'(\hat{\Omega}) - \kappa^2\bar{s}_2\Delta A.$$

Thus, the expectation of the second line in (A53) is

$$\mathbb{E}_i\left[\left(S_i'(\hat{\Omega})\kappa - p(\Omega^{ct})\right)\Delta A - \frac{s_{2i}}{2}\kappa^2(\Delta A)^2\right] = \kappa\mathbb{E}_i[(S_i'(\hat{\Omega}) - \bar{S}'(\hat{\Omega}))\Delta A] + \kappa^2\left(\bar{s}_2 - \frac{s_{2i}}{2}\right)\mathbb{E}_i[(\Delta A)^2].$$

Combining all terms, we obtain agent i's expectation of (A53):

$$\begin{aligned}\mathbb{E}_i[\Delta u_i] = {} & \frac{\eta_i}{2}(\Delta a_i)^2 + \frac{\kappa^2}{2}\bar{s}_2\mathbb{E}_i[(\Delta A)^2] \\ & + \frac{\kappa^2}{2}(\bar{s}_2 - s_{2i})\mathbb{E}_i[(\Delta A)^2] + \kappa\mathbb{E}_i[(S_i'(\hat{\Omega}) - \bar{S}'(\hat{\Omega}))\Delta A] \\ & + \mathbb{E}_i[(\hat{a}_i - \hat{A})(p(\Omega^{ct}) + \hat{P})]. \quad \square\end{aligned}$$

Equilibrium in the Romer (1986) Knowledge Spillover Model

The following lemma characterizes the unique linear equilibrium for any first-order condition of the form

$$ck_i = m_i + \delta\mathbb{E}_i[K], \tag{A56}$$

where $\delta > 0$, $K = \int_0^1 k_i\,di$, and $m_i \equiv \mathbb{E}[\theta \mid y_i] = \mathbb{E}_i[\theta]$. Under the Gaussian information structure in the text, $m_i = (1-\lambda)\mu + \lambda y_i$ with $\lambda \in [0,1]$. We will subsequently show that the status quo regime generates this form with $\delta = \gamma$, and the team-efficient regime generates it with $\delta = 2\gamma$.

Lemma 4. *Assume $c > \delta$. There exists a unique, globally stable, symmetric linear equilibrium*

$$k_i = \phi_0 + \phi_1 m_i, \tag{A57}$$

characterized by

$$\phi_0 = \frac{\delta(1-\lambda)\mu}{(c-\delta)(c-\delta\lambda)} \quad \textit{and} \quad \phi_1 = \frac{1}{c-\delta\lambda}.$$

Moreover,

$$K(\theta) = \phi_0 + \phi_1\big((1-\lambda)\mu + \lambda\theta\big). \tag{A58}$$

Proof. Conjecture (A57). Aggregating across i yields

$$K(\theta) = \int_0^1 (\phi_0 + \phi_1 m_i)di = \phi_0 + \phi_1 \int_0^1 m_i di.$$

Because $y_i = \theta + \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, $\int_0^1 y_i di = \theta$ a.s., so $\int_0^1 m_i di = (1-\lambda)\mu + \lambda\theta$, giving (A58). Firm i's forecast of K is

$$\mathbb{E}_i[K] = \mathbb{E}_i\Big[\phi_0 + \phi_1\big((1-\lambda)\mu + \lambda\theta\big)\Big] = \phi_0 + \phi_1\big((1-\lambda)\mu + \lambda m_i\big), \tag{A59}$$

where we have used $\mathbb{E}_i[\theta] = m_i$.

Substituting (A57) and (A59) into (A56) and matching coefficients, we obtain

$$\phi_0 = \frac{\delta\phi_1(1-\lambda)\mu}{c-\delta} = \frac{\delta(1-\lambda)\mu}{(c-\delta)(c-\delta\lambda)}, \quad \text{and} \quad \phi_1 = \frac{1}{c-\delta\lambda}.$$

To verify that the solution is globally stable, consider the best-response mapping from conjectured coefficients $(\tilde{\phi}_0, \tilde{\phi}_1)$ to actual coefficients (ϕ_0, ϕ_1). Suppose agent i believes all other agents are using the strategy defined by the coefficients $\tilde{\phi}_0$ and $\tilde{\phi}_1$. Agent i forecasts the aggregate as $\mathbb{E}_i[K] = \tilde{\phi}_0 + \tilde{\phi}_1((1-\lambda)\mu + \lambda m_i)$. Substituting this into the first-order condition $ck_i = m_i + \delta\mathbb{E}_i[K]$ generates a best response k_i with updated coefficients $\frac{\delta}{c}\tilde{\phi}_0 + \frac{\delta(1-\lambda)\mu}{c}\tilde{\phi}_1$ and $\frac{1}{c} + \frac{\delta\lambda}{c}\tilde{\phi}_1$. This defines the map $\mathcal{T} : \mathbb{R}^2 \to \mathbb{R}^2$:

$$\mathcal{T}(\tilde{\phi}_0, \tilde{\phi}_1) = \left(\frac{\delta}{c}\tilde{\phi}_0 + \frac{\delta(1-\lambda)\mu}{c}\tilde{\phi}_1, \quad \frac{1}{c} + \frac{\delta\lambda}{c}\tilde{\phi}_1\right).$$

The Jacobian matrix $J_\mathcal{T}$ represents the sensitivity of the updated coefficients to the conjecture:

$$J_\mathcal{T} = \begin{pmatrix} \frac{\partial\phi_0}{\partial\tilde{\phi}_0} & \frac{\partial\phi_0}{\partial\tilde{\phi}_1} \\ \frac{\partial\phi_1}{\partial\tilde{\phi}_0} & \frac{\partial\phi_1}{\partial\tilde{\phi}_1} \end{pmatrix} = \begin{pmatrix} \frac{\delta}{c} & \frac{\delta(1-\lambda)\mu}{c} \\ 0 & \frac{\delta\lambda}{c} \end{pmatrix}.$$

Since the matrix is upper-triangular, its eigenvalues are the diagonal entries: δ/c and $\delta\lambda/c$. Because $\lambda \in [0,1]$ and $c > \delta$, both eigenvalues are strictly less than one. This implies that the mapping is a contraction: iterated best responses converge to the unique fixed point (ϕ_0, ϕ_1). Thus, the equilibrium is unique and globally stable if $c > \delta$. □

Status quo equilibrium. In the status quo, firm i solves

$$\max_{k_i} \ \mathbb{E}_i\Big[e - \frac{c}{2}k_i^2 + (\theta + \gamma K)k_i\Big],$$

taking K as given. The first-order condition is

$$ck_i = m_i + \gamma \mathbb{E}_i[K].$$

This is (A56) with $\delta = \gamma$. Lemma 4 then implies

$$k_i^{sq} = \frac{\gamma(1-\lambda)\mu}{(c-\gamma)(c-\gamma\lambda)} + \frac{1}{c-\gamma\lambda} m_i \quad \text{and} \quad K^{sq}(\theta) = \frac{c(1-\lambda)\mu}{(c-\gamma)(c-\gamma\lambda)} + \frac{\lambda}{c-\gamma\lambda}\theta.$$

Team-efficient equilibrium. Ex-post welfare in state θ is

$$\int_0^1 \Big[e - \frac{c}{2}k_i^2 + (\theta + \gamma K)k_i\Big] di = e + \theta K + \gamma K^2 - \frac{c}{2}\int_0^1 k_i^2 di, \quad K = \int_0^1 k_i di.$$

The team planner chooses $\mathcal{I}_i$-measurable decision rules $k_i(\cdot)$ to maximize ex-ante welfare subject to the constraint $K = \int_0^1 k_i(\mathcal{I}_i)di$:

$$\max_{\{k_i(\mathcal{I}_i)\}} \mathbb{E}\Big[e + \theta K + \gamma K^2 - \frac{c}{2}\int_0^1 k_i^2 di\Big].$$

Let $\Lambda(\theta)$ be the multiplier on $K - \int_0^1 k_i di = 0$. The Lagrangian is

$$\mathcal{L} = \mathbb{E}\Big[e + \theta K + \gamma K^2 - \frac{c}{2}\int_0^1 k_i^2 di + \Lambda(\theta)\Big(K - \int_0^1 k_i di\Big)\Big].$$

The pointwise first-order condition with respect to K yields

$$\Lambda(\theta) = -(\theta + 2\gamma K(\theta)). \tag{A60}$$

By the same variational argument as in Proposition 2, the optimal rule $k_i(\mathcal{I}_i)$ satisfies

$$\mathbb{E}_i[-ck_i(\mathcal{I}_i) - \Lambda(\theta)] = 0. \tag{A61}$$

Substituting (A60) into (A61) and using $\mathbb{E}_i[\theta] = m_i$ yields

$$ck_i^{te} = m_i + 2\gamma \mathbb{E}_i[K^{te}]. \tag{A62}$$

This is (A56) with $\delta = 2\gamma$. Lemma 4 then implies

$$k_i^{te} = \frac{2\gamma(1-\lambda)\mu}{(c-2\gamma)(c-2\gamma\lambda)} + \frac{1}{c-2\gamma\lambda} m_i \quad \text{and} \quad K^{te}(\theta) = \frac{c(1-\lambda)\mu}{(c-2\gamma)(c-2\gamma\lambda)} + \frac{\lambda}{c-2\gamma\lambda}\theta, \tag{A63}$$

which is well-defined when $c > 2\gamma$.

Implementation with Coasean transfers. Consider the transfer schedule $T_i(k_i, K) = (k_i - K)p(K)$ with price $p(K) = \gamma K$. This price corresponds to the Lindahl pricing schedule, obtained by

aggregating individual marginal valuations of the aggregate stock (holding private actions k_i fixed):

$$\hat{p}(K) = \int_0^1 \frac{\partial u_i}{\partial K}\, di = \int_0^1 \gamma k_i\, di = \gamma K.$$

Firm i maximizes $\mathbb{E}_i[(\theta + \gamma K)k_i - \frac{c}{2}k_i^2 + (k_i - K)\gamma K]$. Because the firm is atomistic, it takes K as given. Thus, the first-order condition is

$$ck_i = m_i + 2\gamma \mathbb{E}_i[K],$$

which coincides with (A62). Hence the Coasean transfer implements the team-efficient allocation.

Limit cases. From (A63), as $\lambda \to 0$ (no private information),

$$K^{te}(\theta) \to \frac{\mu}{c - 2\gamma}. \tag{A64}$$

As $\lambda \to 1$ (perfect private information),

$$K^{te}(\theta) \to \frac{\theta}{c - 2\gamma}. \tag{A65}$$

Implementation with Pigouvian tax/subsidy. Consider a regulator who sets a constant subsidy τ ex ante, with lump-sum revenue $R = -\tau K$ redistributed. Firm i solves

$$\max_{k_i}\ \mathbb{E}_i\Big[e - \frac{c}{2}k_i^2 + (\theta + \gamma K)k_i + \tau k_i + R\Big],$$

taking K and R as given. The first-order condition is

$$ck_i = m_i + \gamma \mathbb{E}_i[K] + \tau.$$

Conjecture $k_i = \psi_0 + \psi_1 m_i$. By the same steps as in Lemma 4, matching coefficients yields

$$\psi_1 = \frac{1}{c - \gamma\lambda}, \quad \psi_0 = \frac{\gamma(1-\lambda)\mu}{(c-\gamma)(c-\gamma\lambda)} + \frac{\tau}{c-\gamma}. \tag{A66}$$

The aggregate response is

$$K^{subsidy}(\tau, \theta) = \frac{c(1-\lambda)\mu}{(c-\gamma)(c-\gamma\lambda)} + \frac{\lambda\theta}{c-\gamma\lambda} + \frac{\tau}{c-\gamma}. \tag{A67}$$

Ex-ante welfare is defined as $W(\tau) = \mathbb{E}[\int_0^1 (e + (\theta + \gamma K(\tau, \theta))k_i - \frac{c}{2}k_i^2)di]$. Using $K = \int_0^1 k_i di$ and

the identity $\int_0^1 k_i^2 di = K^2 + \int_0^1 (k_i - K)^2 di$, we decompose welfare as follows:

$$\begin{aligned} W(\tau) &= \mathbb{E}\left[e + \theta K(\tau,\theta) + \gamma K(\tau,\theta)^2 - \frac{c}{2}\left(K(\tau,\theta)^2 + \int_0^1 (k_i - K)^2 di\right)\right] \\ &= \mathbb{E}\left[e + \theta K(\tau,\theta) - \frac{c-2\gamma}{2} K(\tau,\theta)^2\right] - \frac{c}{2}\mathbb{E}\left[\mathrm{Var}(k_i)\right], \end{aligned}$$

where $\mathrm{Var}(k_i)$ denotes cross-sectional variance across firms. In the linear equilibrium $k_i = \psi_0 + \psi_1 m_i$, the cross-sectional variance is $\mathrm{Var}(k_i) = \psi_1^2 \mathrm{Var}(m_i)$. Since $\psi_1 = (c - \gamma\lambda)^{-1}$ and $\mathrm{Var}(m_i) = \lambda^2 \sigma_\varepsilon^2$ are independent of τ, the subsidy affects only the mean capital stock, not its dispersion. Consequently, the regulator's problem reduces to maximizing the first term, $\mathbb{E}[e + \theta K(\tau,\theta) - \frac{c-2\gamma}{2} K(\tau,\theta)^2]$. The first-order condition with respect to τ yields $\mathbb{E}[(\theta - (c-2\gamma)K)\frac{1}{c-\gamma}] = 0$, and thus

$$\mathbb{E}[K^{subsidy}] = \frac{\mu}{c - 2\gamma},$$

which, using (A67), implies the optimal subsidy

$$\tau^* = \frac{\mu\gamma}{c - 2\gamma}.$$

Substituting τ^* into the equilibrium coefficients (A66) gives

$$k_i^{subsidy} = \frac{\mu\gamma(2-\lambda)}{(c-\gamma\lambda)(c-2\gamma)} + \frac{m_i}{c-\gamma\lambda} \quad \text{and} \quad K^{subsidy}(\theta) = \frac{\mu(c(1-\lambda)+\gamma\lambda)}{(c-\gamma\lambda)(c-2\gamma)} + \frac{\lambda\theta}{c-\gamma\lambda}.$$

Limit Cases. Consider now the behavior of the Pigouvian solution in the limits of information precision. First, as private information becomes uninformative ($\lambda \to 0$), firms rely solely on the prior ($m_i \to \mu$). The allocation converges to:

$$k_i^{subsidy} \to \frac{\mu}{c - 2\gamma} \quad \text{and} \quad K^{subsidy} \to \frac{\mu}{c - 2\gamma}.$$

This is the team-efficient allocation in equation (A64). Second, as information becomes perfect ($\lambda \to 1$), firms observe the fundamental ($m_i \to \theta$). The allocation converges to:

$$\begin{aligned} k_i^{subsidy} &\to \frac{\mu\gamma}{(c-\gamma)(c-2\gamma)} + \frac{\theta}{c-\gamma}, \\ K^{subsidy}(\theta) &\to \frac{\mu\gamma}{(c-\gamma)(c-2\gamma)} + \frac{\theta}{c-\gamma}. \end{aligned}$$

This deviates from the team-efficient allocation in equation (A65). First, the intercept remains dependent on the prior μ because the subsidy is fixed ex ante. Second, and most importantly, the

aggregate responsiveness to the fundamental θ remains $\frac{1}{c-\gamma}$, exactly as in the status quo. This is strictly less than the team-efficient responsiveness, $\frac{1}{c-2\gamma}$. A fixed Pigouvian subsidy corrects the level of investment on average ($\mathbb{E}[K^{subsidy}(\theta)] = \mathbb{E}[K^{te}(\theta)] = \frac{\mu}{c-2\gamma}$), but fails to correct the marginal incentive to respond to information. Thus, even when agents have perfect information, the Pigouvian solution is inefficient. □

Equilibrium in the Gordon (1954) Fishery Model

The original production function in Gordon (1954) is $L = caE/(1 + cbE)$, where a is the natural population, c is the production coefficient, and b is the depletion coefficient. We first linearize the model for small congestion ($cbE \ll 1$) using the first-order Taylor expansion $(1 + x)^{-1} \approx 1 - x$:

$$L \approx caE(1 - cbE) = caE - c^2abE^2.$$

We define the state θ as the deviation of the population from its mean $\bar{a}$, such that $a = \bar{a} + \theta$, where $\theta \sim \mathcal{N}(0, \sigma_\theta^2)$. Substituting this into the linear term and approximating $a \approx \bar{a}$ in the quadratic crowding term (consistent with standard LQG frameworks to avoid cubic terms) yields:

$$L \approx c(\bar{a} + \theta)E - c^2\bar{a}bE^2 = c\bar{a}E + c\theta E - c^2\bar{a}bE^2.$$

Mapping the physical parameters to the linear-quadratic coefficients $\alpha \equiv c\bar{a}$, $\beta \equiv c$, and $\gamma \equiv c^2\bar{a}b$, we obtain the approximation (35):

$$L(E, \theta) \approx \alpha E + \beta\theta E - \gamma E^2.$$

Agents receive private noisy signals $y_i = \theta + \varepsilon_i$ with $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$, and update beliefs to form the posterior mean $m_i \equiv \mathbb{E}[\theta \mid y_i] = \lambda y_i$, where $\lambda \equiv \sigma_\theta^2/(\sigma_\theta^2 + \sigma_\varepsilon^2)$ is the signal-to-noise ratio.

Under open access, agents maximize $\pi_i = e_i(\alpha + \beta\theta - \gamma E) - \frac{\phi}{2}e_i^2$. This problem is isomorphic to the Romer (1986) model in Section 8.1, with the negative externality term $-\gamma E$ replacing the positive spillover term γK. Conjecture $e_i = a_0 + a_1 m_i$. Using the first-order condition $\phi e_i = \alpha + \beta m_i - \gamma\mathbb{E}_i[E]$ and $\bar{m} = \int_0^1 m_i di = \lambda\theta$, aggregating, and matching coefficients, the symmetric linear equilibrium is:

$$e_i^{sq} = \frac{\alpha}{\phi + \gamma} + \frac{\beta}{\phi + \gamma\lambda}m_i \quad \text{and} \quad E^{sq} = \frac{\alpha}{\phi + \gamma} + \frac{\beta\lambda}{\phi + \gamma\lambda}\theta.$$

Since the team-efficient problem is isomorphic to the Romer model, the solution follows from the

same fixed-point technique as in Lemma 4:

$$e_i^{te} = \frac{\alpha}{\phi + 2\gamma} + \frac{\beta}{\phi + 2\gamma\lambda} m_i \quad \text{and} \quad E^{te} = \frac{\alpha}{\phi + 2\gamma} + \frac{\beta\lambda}{\phi + 2\gamma\lambda}\theta.$$

Because $\gamma > 0$, the intercept and slope coefficients in the status quo are strictly larger than in the team-efficient allocation, implying that open access leads to inefficiently high average effort and excessive sensitivity to the stock. As in the Romer (1986) model, the efficient allocation is implemented by the Coasean transfer scheme $T_i = (e_i - E)p(E)$, where the price $p(E) = -\gamma E$ represents the aggregated marginal congestion cost (the Lindahl price):

$$\hat{p}(E) = \int_0^1 \frac{\partial \pi_i}{\partial E}\, di = \int_0^1 (-\gamma e_i)\, di = -\gamma E.$$

Under this transfer, agent i maximizes $\mathbb{E}_i[\pi_i + T_i]$. Taking E as given, the first-order condition is $\phi e_i = \alpha + \beta m_i - 2\gamma \mathbb{E}_i[E]$, which implements e_i^{te}. □