THE PINCHAS SAPIR CENTER
FOR DEVELOPMENT
Tel Aviv University

המרכז לפיתוח ע"ש פנחס ספיר
ליד אוניברסיטת תל אביב
"עמותה רשומה"
580010221

THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

A Robust Wisdom of the Crowd

Daniel Fershtman

Discussion Paper No. 7-2026

A Robust Wisdom of the Crowd*

Daniel Fershtman†

June 2026

Abstract

This paper studies a model of outcome-based observational learning where a sequence of agents each choose among risky actions (e.g., which technological innovation to adopt), the outcome of which is a function of an unknown state. Agents observe past choices and outcomes, but are short-lived and do not internalize the informational value of their actions. The main result is a "robust wisdom of the crowd": a characterization of the set of information structures under which agents, without any intervention, all choose the socially optimal action, at any history, for any number of agents, and any prior. The results also shed light on the regulation of online artificial intelligence-based recommendation platforms.

Keywords: Observational learning; wisdom of the crowd; artificial intelligence recommendation platforms; technology adoption.

*I thank Arjada Bardhi, Jaden Chen, Eddie Dekel, Théo Durandard, Tristan Gagnon-Bartsch, Kevin He, Teddy Kim, Ilan Kremer, Teddy Mekonnen, James Peck, Ariel Rubinstein, Udayan Vaidya, Asher Wolinsky, and Boli Xu for helpful suggestions. Financial support from the Pinhas Sapir Center is gratefully acknowledged.

†Tel Aviv University. E-mail: danfershtman@gmail.com

1 Introduction

A sequence of agents each choose among risky actions (e.g., which technological innovation to adopt or which medical treatment to prescribe), the outcome of which is a function of an unknown state. Agents observe the history of decisions made by their predecessors as well as their outcomes. The agents, however, are short-lived and myopic and, therefore, do not internalize the informational value that their decision bears for future agents. When the number of agents is large, this informational value may be high. Can it nevertheless be the case that the agents, left to their own devices, each choose the socially optimal action, even when the number of agents is large and regardless of the history of past decisions and outcomes?

Consider the following problem. One of two technologies, α or β, can be used. Each technology yields either a success or a failure. Whether a technology is successful or not depends on the unknown state, which is either A or B. In particular, suppose that the prior probability that the state is A is equal to $p_0 = 1/2$, and that the following table describes the probability with which each of the technologies is successful, given each of the two possible states.

	state A	state B
action α	1	0
action β	0.6	0.6

Action α perfectly reveals the state, and action β is uninformative. There are two (Bayesian) agents, who care only about their own success. Agent 1 chooses first based on the prior, and agent 2 moves second, after observing agent 1's choice and its outcome.

In (Nash) equilibrium, both agents choose β: agent 1 chooses β because it has a higher expected probability of success given the prior 1/2. This is uninformative for agent 2, who chooses β as well. The resulting expected number of successes is $0.6 + 0.6 = 1.2$. Now suppose that there is a planner who chooses actions on behalf of the agents, with the goal of maximizing the expected number of overall successes. Rather than choosing β on behalf of agent 1, the planner can instead choose action α. A success (failure) would reveal that the state is A (B), in which case the planner would have agent 2 choose α (β). The expected number of successes in this case is $\frac{1}{2}(1 + 1) + \frac{1}{2}(0 + 0.6) = 1.3$. In other words, the equilibrium is not socially optimal.

Now suppose that instead of $1/2$, the prior is $p_0 = 0.4$. In this case, although the equilibrium outcome remains the same (both agents choose β), the planner no longer benefits from having agent 1 choose α.[1] Hence, in this case the equilibrium actions are in fact socially optimal.

Finally, suppose that the prior is $p_0 = 0.4$, but the number of agents is three instead of two. In equilibrium, all agents choose β. However, the informational value of having agent 1 choose α is now higher compared to the case of two agents, and the planner again finds this optimal.[2] Hence, in this case, the equilibrium is not socially optimal.

The example above demonstrates that the answer to the question of whether the equilibrium choices that agents make coincide with the socially optimal ones depends on the prior and the number of agents.

This paper considers the following question: Are there environments where the equilibrium is *always* socially optimal, regardless of the prior and number of agents? Such environments are of special interest because they imply that there is no need for a designer to intervene in the first place to achieve efficiency. Alternatively, suppose that the agents are not the ones who observe data about past outcomes, but instead follow the recommendations of an artificial intelligence (AI) based platform that has this data and whose goal is to maximize the overall expected number of successes of its users.[3] When recommending an action to an agent, the platform cares not only about the agent's outcome, but also about the information it generates towards more informed recommendations in the future. This potential misalignment of interests between the platform and its users has raised ethical concerns, particularly when it comes to healthcare, finance, employment, or even matchmaking recommendations.[4] For

[1] The resulting expected number of successes in this case is $0.4(1 + 1) + 0.6(0 + 0.6) = 1.16 < 1.2$.

[2] The expected number of successes in equilibrium is $0.6 + 0.6 + 0.6 = 1.8$, whereas the expected number of successes under the socially optimal policy is $0.4(1 + 1 + 1) + 0.6(0 + 0.6 + 0.6) = 1.92$.

[3] This is meant as a simplification. Clearly, such platforms typically seek to maximize profits rather than user welfare. While the platform's profit maximization problem is not modeled here, the forces studies in this paper would also arise in such a setting.

[4] AI recommendation systems often make non-optimal recommendations to gather data in order to improve future performance. While this may be acceptable in some contexts such as content personalization or product discovery, it raises ethical concerns in more sensitive domains such as healthcare, finance, or employment (Mittelstadt et al., 2016; Binns, 2018). It is even more of a concern when users are unaware that this might be the case. Whether such exploration is ethically acceptable hinges on transparency and fairness – standards that are often not met in practice (Floridi et al., 2018). The results in the current paper characterize when learning for future benefit does not come at the expense of current users, so that there is no such concern.

example, in 2014 OkCupid disclosed on its blog that it had run several A/B tests on users without their knowledge (see, e.g., Luca, 2014). These included recommending suboptimal matches by manipulating compatibility scores and even assigning matches without using its algorithm at all. The company framed this as standard industry practice, but the revelation sparked public criticism and raised ethical concerns about transparency and user manipulation. Given the prevalence of such examples, it is important to understand when such misalignment occurs.

The first of the paper's two main results characterizes the set of information structures such that agents, despite being short-lived and myopic, are all guaranteed to take the socially optimal action regardless of the history, initial prior, and number of agents. It is shown that each of the following two classes of information structures is sufficient to guarantee that the equilibrium is socially optimal in this robust sense: (i) *Symmetric* information structures, where the success probabilities are symmetric across the states – i.e., action α's (β's) probability of success in state A is the same as action β's (α's) probability of success in state B; (ii) Information structures with *competing actions*, where the success probabilities within each state sum to one. The result shows that when the information structure falls within one of these two classes, there is no need for intervention.[5] Alternatively, if it is a platform that observes the history of outcomes and makes recommendations to the agents - there is no concern that the platform's additional goal of generating information for the benefit of future recommendations will conflict with a current agent's best interest.

The second main result establishes that the characterization described above is sharp: Outside the set of symmetric or competing-actions information structures, there always exist a prior and number of agents such that the socially optimal actions differ from the equilibrium actions at some histories. Intervention is therefore needed in these cases to ensure efficiency. However, a continuity property implies that for information structures in the neighborhood of the set of symmetric or competing-actions information structures, the welfare loss is guaranteed to be "small". In other words, even if the environment does not fall exactly into one of the two classes but is sufficiently close, the need for intervention must be small.

The problem of characterizing the socially optimal policy in this model is new.

[5]A trivial case in which the equilibrium is always efficient is when one action dominates the other (i.e., has a higher probability of success) regardless of the state. This trivial case is ruled out by assuming that action α (β) has a higher probability of success in state A (B).

While at first glance it may resemble a multi-armed bandit problem, it is very different, for two reasons: First, and most importantly, the arms are not independent, as the payoffs from each of the "arms" depend on a common, unknown state. Second, the horizon is finite. For these reasons, the Gittins-index approach from the multi-armed bandit problem is not helpful here. In the absence of a convenient characterization of the optimal policy using indices, the analysis instead characterizes the properties of the planner's expected *advantage* from choosing one action relative to the other, as a function of the history. This expected advantage turns out to admit a convenient recursive representation which is central to the analysis, and is shown to be strictly increasing and continuous in the posterior belief, for every agent. These properties allow for a comparison with the advantage function under the equilibrium policy, which is the basis for the results.

Importantly, in contrast to settings with independent alternatives, by revealing information about the underlying state, the outcome of an action is typically informative about the future merits of *both* actions, rather than just the one that is chosen.[6] Under symmetric information structures, the properties described above imply that whenever an action has a higher expected probability of success given the posterior, its informational value must also be higher than that of the other action. Under competing-actions information structures, regardless of the observed history, the informational value is the same for both actions, as each action induces the same distribution over posteriors.

One natural application in the context of which symmetric and competing-actions information structures have an intuitive interpretation is that of medical treatment for a disease with two possible underlying types, A and B (e.g., two genetic variants of a virus or two molecular subtypes of a cancer). Physicians treat patients sequentially and must choose between two treatments, α and β, each of which can succeed or fail depending on the disease type. Treatment α is more effective in state A, while β is more effective in state B. The true state is unknown, and physicians observe past treatment choices and outcomes to update their beliefs before making a decision. In this context, symmetric information structures capture situations where each treatment is designed to be equally effective for its target disease type.[7] Competing-actions information

[6]The only exceptions are cases where one of the actions is completely uninformative.

[7]This is common in various fields of medicine where therapies are tailored to specific genetic profiles and calibrated to yield comparable outcomes across subtypes. For example, in non-small cell

structures reflect environments where the treatments are inversely effective, in the sense that if one has a high probability of success, the other must have a low one.[8] This arises, for example, when each treatment targets a distinct biological mechanism, only one of which is active in a given disease type. An additional application, to agricultural technology adoption, and the corresponding intuition for symmetric and competing-actions information structures, is discussed in Section 3.

The rest of the paper is organized as follows. The remainder of this section discusses the most related literature. Section 2 introduces the model. Section 3 contains all of the main results. Section 4 discusses an application of the model to AI-based recommendation platforms. Section 5 concludes and discusses several extensions. All proofs are relegated to the Appendix.

1.1 Related literature

The paper is closely related to a growing literature that takes a mechanism design perspective to incentivizing technology adoption; see Kremer, Mansour and Perry (2014), Che and Hörner (2018), and Glazer, Kremer and Perry (2021). This literature studies the problem of a principal who observes the actions (and/or outcomes) of a sequence of myopic agents, and chooses an incentive compatible recommendation for each agent that arrives. The principal's aim is to overcome the potential conflict between the action preferred by myopic agents and those that best serve future agents. This paper, instead, takes a step back and asks under what conditions there is a conflict between agents' own interests and the interest of the principal in the first place, and if such a conflict is indeed present, when is the corresponding welfare loss guaranteed to be small.[9]

lung cancer, Osimertinib and Crizotinib target distinct genetic mutations (EGFR and ALK) and are similarly effective when matched to the correct subtype. See Mok et al. (2017).

[8] In bacterial infections, it is common that one antibiotic may be effective while another is not due to resistance. Likewise, in cardiology, the effectiveness of Clopidogrel or Ticagrelor depends on the CYP2C19 genotype. See Mega et al. (2009).

[9] Similarly, the model can also be viewed in the context of experimentation in a principal-agent setting with learning (see, e.g., Bergemann and Hege, 1998, 2005, Hörner and Samuelson, 2013, Guo, 2016, and Halac, Kartik and Liu, 2016) in which a myopic agent makes sequential decisions (e.g., selects projects) on behalf of a principal who is far-sighted. One can then ask – when does a conflict arise between the principal's preferred decisions and those of the agent? In this context, the paper offers, within the boundaries of this simple model, a characterization of the set of information structures under which there is no misalignment of incentives between the principal and the agent,

More broadly, the paper is also related to a large literature on observational learning studying herding and information cascades (e.g., Banerjee, 1992, Bikhchandani, Hirshleifer and Welch, 1992, Welch, 1992, Smith and Sørensen, 2000, and Wolitzky, 2018).[10] A key departure from this literature is that agents in the current model do not possess private information, and that the focus is on the efficiency of both short-run and long-run behavior, rather than the limit.

Another important departure from the literature mentioned above is that, in the current paper, agents learn by observing past actions and outcomes (Wolitzky, 2018, is a notable exception). There are many relevant economic applications that have this feature. For example, physicians choosing which medical treatments to prescribe may observe past treatment choices and their outcomes; farmers choosing which crop to grow or which fertilizer to use may be able to see their neighbors' choices and crop yields; firms choosing which manufacturing technology to adopt may observe other firms' production methods and output levels; see Wolitzky (2018) for an extensive discussion of a wide range of economic applications featuring "outcome-based" learning.[11] In the context of such applications, this paper asks: When can efficient technology adoption emerge (not just in the long run, but at all times) even when agents do not internalize the informational value of their choices?

2 A Model

There are $N > 1$ agents indexed by $i = 1, ..., N$, and two possible states of the world, $\omega \in \{A, B\}$. In each period i, agent i chooses among two risky actions, $\sigma_i \in \{\alpha, \beta\}$. Each action i yields either a success, $\theta_i = 1$, or a failure, $\theta_i = 0$.[12] The outcomes that each action generates are i.i.d. over time, with a probability of success that depends on the state ω. In particular, if the state is ω, the action α yields a success

and in the neighborhood of which such a misalignment is small.

[10] See Bikhchandani et al. (2024) for an excellent recent survey.

[11] Action-based learning may be a more appropriate model in settings in which actions correspond to consumption choices (e.g., choosing a restaurant or a hotel), so that outcomes correspond to utility realizations, which may be more difficult to observe. Outcome-based learning may be more natural in applications such as those described above, where actions correspond to choosing inputs in some form of production process.

[12] We follow much of the literature in assuming binary states and actions. The assumption that outcomes are binary, which follows Wolitzky (2018), greatly simplifies the analysis. See Section 5 for a discussion of how the results extend beyond the case of binary states.

with probability q^{α}_{ω} and the action β yields a success with probability q^{β}_{ω}. The state ω is unknown to all of the agents, and at period 0, the agents share a common prior $p_0 \in (0, 1)$, corresponding to the probability that the state is $\omega = A$.

	State A	State B
Action α	q^{α}_{A}	q^{α}_{B}
Action β	q^{β}_{A}	q^{β}_{B}

Figure 1: State-dependent success probabilities

The agents are short-lived and myopic. In particular, agent i's payoff is simply 1 if her action σ_i is successful, and 0 otherwise. Agents observe all of the actions and outcomes of their predecessors.[13] When agent 1 chooses an action, she does so based on the prior p_0. When agent $i \in \{2, ..., N\}$ chooses her action, she does so after observing the history $h_{i-1} = (\sigma^{i-1}, \theta^{i-1})$ of previous observations of actions $\sigma^{i-1} = (\sigma_1, ..., \sigma_{i-1})$ and outcomes $\theta^{i-1} = (\theta_1, ..., \theta_{i-1})$, updating her beliefs using Bayes' rule. For any $i \in \{1, ..., N\}$, denote by p_i the posterior belief that $\omega = A$. Denote the expected probability of success resulting from action $\sigma \in \{\alpha, \beta\}$ given a posterior p_i by

$$\mu_{\sigma}(p_i) = \mathbb{E}[q^{\sigma}_{\omega}|p_i].$$

Agents choose an action with the goal of maximizing their expected payoff. Consider the Nash equilibrium of this game. Since the agents are short-lived and myopic, each agent i finds it optimal to choose the action σ if her posterior p_{i-1} is such that[14]

$$\mu_{\sigma}(p_{i-1}) > \mu_{-\sigma}(p_{i-1}).$$

Let a *policy* be a rule specifying for every agent $i \in \{1, ..., N\}$ an action σ_i given their belief p_{i-1}. Denote the myopic policy that arises in equilibrium by π^E. The *efficient policy* is the policy that maximizes social welfare – the expected number of overall successes. Denote this socially optimal policy by π^W. We say that the policies π^E and π^W are identical, $\pi^E \equiv \pi^W$, if $\pi^E_i(p) = \pi^W_i(p)$ for any agent i and any belief p.

The central question we will be interested in is: Are the equilibrium policy π^E and

[13] Note that if it is common knowledge in period 0 that the probability that the state is A is p_0, then if outcomes are observable, the assumption that actions are also observable is without loss. Using the prior and outcomes, agents can back out the action that each previous agent chose.

[14] The action $-\sigma$ refers to β if $\sigma = \alpha$, and to α if $\sigma = \beta$.

the welfare-maximizing policy π^W necessarily in conflict? Can they be the same, and if so, when?

One obvious case in which $\pi^E \equiv \pi^W$ is when (i) in both states $\omega = A, B$, $q^\alpha_\omega \geq q^\beta_\omega$, or alternatively (ii) in both states, $q^\alpha_\omega \leq q^\beta_\omega$. In the first case, action α always weakly dominates action β regardless of the state, and in the second, action β always weakly dominates α regardless of the state. In particular, independently of the beliefs over the state, under both policies π^E and π^W, the action α is chosen at all periods in case (i), and the action β is chosen at all periods in case (ii).

To avoid these two trivial cases, we restrict attention throughout the paper to the case where $q^\alpha_A \geq q^\beta_A$ and $q^\alpha_B \leq q^\beta_B$. We refer to α as the *preferred action* in state A – meaning that α yields a success with higher probability in state A than B does – and to action β as the preferred action in state B.

3 Main results

The goal of this section is to fully characterize the set of information structures under which $\pi^E \equiv \pi^W$ for any number of agents.

An *information structure* is a tuple $(q^\alpha_A, q^\beta_A, q^\beta_B, q^\alpha_B) \in [0,1]^4$ such that $q^\alpha_A \geq q^\beta_A$ and $q^\alpha_B \leq q^\beta_B$. Denote by Q the set of all information structures and denote $\kappa_\omega \equiv q^\alpha_\omega + q^\beta_\omega$. The following special cases will play an important role in the analysis.

Definition 1 (**Symmetric and competing-actions information structures**).

1. An information structure $q \in Q$ such that $q^\alpha_A = q^\beta_B$ and $q^\alpha_B = q^\beta_A$ will be called a *symmetric* information structure. Let $S \subset Q$ denote the set of symmetric information structures.

2. An information structure $q \in Q$ where, for each state $\omega = A, B$, $\kappa_\omega = 1$ will be called a *competing-actions* information structure. Denote by $C \subset Q$ the set of such information structures.

Under a symmetric information structure, each of the actions is equally likely to yield a success in the state in which it is preferred, and equally likely to yield a success in the state in which the other action is preferred. This is typically the case when each action is tailored to a given state, and the actions are equally effective when applied in the state they are meant for. In the case of a competing-actions information structure, within each state, the actions α and β yield a success with complementary probabilities; that is, given each state, the probability that the preferred action yields a success is equal to the probability that the non-preferred action yields a failure. This typically arises in environments where the actions are inversely effective, in the sense that if one yields a high probability of success, the other yields a low one. Note that a symmetric information structure need not exhibit competing actions, and one with competing actions need not be symmetric. An information structure may of course satisfy both. The examples in Figures 2(a)-2(c) illustrate.

	A	B
α	2/3	1/2
β	1/2	2/3

Figure 2(a): Symmetric q

	A	B
α	3/4	1/5
β	1/4	4/5

Figure 2(b): Competing-actions q

	A	B
α	2/3	1/3
β	1/3	2/3

Figure 2(c): Symmetric and competing-actions q

To illustrate how these two classes of information structures can be interpreted, recall the application to medical treatments in the introduction, or consider the following application to agricultural technology adoption under environmental uncertainty. Consider farmers who sequentially choose between two types of fertilizer – α and β – to apply to a specific crop. The relative effectiveness of each fertilizer depends on an unknown soil condition or climate regime – state A or B. For example, one state may correspond to nitrogen-deficient soil and the other to phosphorus-deficient soil. Fertilizer α may be nitrogen-rich, while β is phosphate-rich, and each is more likely to result in a successful yield in the corresponding condition. Farmers observe the crop outcomes and fertilizer choices of others before making their own decisions, and update their beliefs about the underlying condition accordingly. Symmetric information structures correspond to settings where each fertilizer is designed to be equally effective under the soil condition it targets.[15] This reflects modern agricultural products optimized for specific nutrient deficiencies. Competing-actions information structures capture the scenario where a fertilizer effective under one condition is ineffective (or even counterproductive) under the other. For example, applying a nitrogen-heavy fertilizer to already nitrogen-saturated soil may inhibit crop growth, just as phosphate

[15] See, e.g., Ladha et al. (2005), which argues that different fertilizers target specific deficiencies with comparable yield potential under the right conditions.

overload can impair nutrient uptake in phosphorus-rich environments.[16]

3.1 Sufficiency

We first show that if an information structure is either symmetric or has competing actions (or both), then $\pi^E \equiv \pi^W$ for any number of agents.

Theorem 1 (Sufficiency)**.** *If $q \in S \cup C$ then $\pi^E \equiv \pi^W$ for any number of agents $N \geq 1$.*

One way to think about Theorem 1 is as a *robust wisdom of the crowd*. When the information structure is symmetric or has competing actions, agents left to their own devices, despite being myopic, will each take the efficient action given any prior and observed history of outcomes and actions, regardless of the number of agents.[17]

While the equilibrium policy is trivial to derive, the problem of characterizing the socially optimal policy in this model does not admit a tractable solution.[18] To make progress, the analysis below instead takes an indirect approach of studying the properties of the expected *advantage* from choosing one action relative to the other under the efficient policy, as a function of the history. As will be shown below, this expected advantage admits a convenient recursive representation, which will be crucial to all of the results. In particular, this recursive representation, as well as several other properties of the advantage function, will allow for a comparison with the advantage function under the equilibrium policy. This comparison will serve as the basis for the results.

[16] See Fageria (2016), which argues that imbalanced fertilization – e.g., excessive nitrogen or phosphorus – often reduces yields when misaligned with soil nutrient profiles.

[17] The term "robust" here is meant to indicate that the result is not sensitive to the prior or the number of agents.

[18] Note that this problem differs from the multi-armed bandit problem since the arms are not independent (the payoffs from each of the "arms" depend on a common, unknown state) and also given that the horizon is finite. As a result, the standard Gittins-index approach is not helpful here.

The Advantage Function and its properties

Given a belief p_{i-1}, if agent i chooses the action $\sigma \in \{A, B\}$, the updated posterior belief p_i is

$$p_i(p_{i-1}, \sigma, 1) = \frac{p_{i-1} q_A^\sigma}{p_{i-1} q_A^\sigma + (1 - p_{i-1}) q_B^\sigma} \tag{1}$$

if i's action results in a success ($\theta_i = 1$), and

$$p_i(p_{i-1}, \sigma, 0) = \frac{p_{i-1}(1 - q_A^\sigma)}{p_{i-1}(1 - q_A^\sigma) + (1 - p_{i-1})(1 - q_B^\sigma)} \tag{2}$$

if i's action results in a failure ($\theta_i = 0$). Denote by $p_i(p_{i-1}, \sigma)$ the (stochastic) next-period posterior when the current belief is p_{i-1} and the action σ is chosen.

Denote by $V_i^W(p_{i-1})$ the expected number of future successes under the socially optimal policy π^W, starting from period $i \in \{1, ..., N+1\}$, when the belief is p_{i-1}. Similarly, denote by $V_i^E(p_{i-1})$ the expected number of successes under the equilibrium policy π^E. The period $i = 1$ refers to the period at which none of the agents have made their choice yet (and agent 1 is about to make their choice), and the period $i = N + 1$ refers to the time at which all agents have already moved.

By definition, $V_{N+1}^W(p_N) = 0$ for all p_N. For any $i = 1, ..., N$,

$$V_i^W(p_{i-1}) = \max\left\{\mu_\alpha(p_{i-1}) + \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1}, \alpha))\right], \mu_\beta(p_{i-1}) + \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1}, \beta))\right]\right\}.$$

Note that, in particular, when only one agent has not moved,

$$V_N^W(p_{N-1}) = \max\left\{\mu_\alpha(p_{N-1}), \mu_\beta(p_{N-1})\right\}.$$

Consider now the advantage of choosing the action α relative to β. For any period $i = 1, ..., N$, denote by $\phi_i^W(p_{i-1})$ the difference between (i) the expected future payoff when the action α is chosen in period i and (if there are additional remaining periods) the optimal policy π^W is followed thereafter, and (ii) the expected future payoff when the action β is chosen in period i and (again, if there are additional remaining periods) the optimal policy π^W is followed thereafter. The functions $\phi^W = (\phi_i^W(\cdot))_{i=1}^N$ will be referred to as the *advantage functions*.

Clearly, when it is agent N's turn to choose an action, the advantage of choosing

α over β is given by the difference in the expected one-shot payoff under the two actions,

$$\phi_N^W(p_{N-1}) = \mu_\alpha(p_{N-1}) - \mu_\beta(p_{N-1}).$$

For any agent $i \in \{1, ..., N-1\}$, the advantage is

$$\phi_i^W(p_{i-1}) = \mu_\alpha(p_{i-1}) + \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1}, \alpha))\right] - \left(\mu_\beta(p_{i-1}) + \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1}, \beta))\right]\right).$$

Using the advantage functions $\phi^W = (\phi_i^W(\cdot))_{i=1}^N$, the optimal policy π^W can be described as follows.

Lemma 1 (Optimal policy using ϕ^W)**.** *When it is agent i's turn to choose and p is the current belief that the state is A, the socially optimal action (i.e., the action prescribed by π^W) is α if $\phi_i^W(p) > 0$, and β if $\phi_i^W(p) < 0$.*[19]

By deriving a number of key properties of ϕ^W, we can further characterize the socially optimal policy with the goal of understanding how it compares to the equilibrium policy π^E. In particular, it will be shown below that the advantage of choosing α relative to β is a *strictly* increasing and continuous function of the current belief p, and that, moreover, for each agent $i = 1, ..., N$ there exists a unique threshold belief $\hat{p}_{i-1} \in (0, 1)$ such that $\phi^W(\hat{p}_{i-1}) = 0$.

The following representation of $\phi^W = (\phi_i^W(\cdot))_{i=1}^N$ is a key step in establishing these properties, and will also be used in the subsequent results.

Lemma 2 (Recursive representation of ϕ^W)**.** *For any $i = 1, ..., N$,*

$$\phi_i^W(p_{i-1}) = \mathbb{E}[max\{\phi_{i+1}^W(p_i(p_{i-1}, \alpha)), 0\}] - \mathbb{E}[max\{-\phi_{i+1}^W(p_i(p_{i-1}, \beta)), 0\}].$$

We can now establish the following properties of $\phi_i^W(\cdot)$, which together with Lemma 2 underlie the main results.

Lemma 3 (Properties of ϕ^W)**.** *For each $i \in \{1, ..., N\}$,*

1. *$\phi_i^W(p_{i-1})$ is a strictly increasing, continuous function of p_{i-1}, and*

[19] When $\phi_i^W(p) = 0$, the action can be chosen arbitrarily.

2. *there exists a unique belief* $\hat{p}_{i-1} \in (0,1)$ *such that* $\phi_i^W(\hat{p}_{i-1}) = 0$, *and such that* $\phi_i^W(p_{i-1}) < 0$ *if* $p_{i-1} < \hat{p}_{i-1}$, *and* $\phi_i^W(p_{i-1}) > 0$ *if* $p_{i-1} > \hat{p}_{i-1}$.

The proof (in the Appendix) establishes these properties by induction, starting with the last agent and moving backward. First, for the last agent to move (agent N), the advantage function $\phi_N^W(p_{N-1})$ as a function of the belief p_{N-1} is easy to derive. It can then be verified directly that it is indeed continuous, strictly increasing, and that there exists a unique belief

$$\hat{p}_{N-1} = \frac{q_B^\beta - q_B^\alpha}{q_A^\alpha - q_A^\beta + q_B^\beta - q_B^\alpha} \in (0,1) \tag{3}$$

for which $\phi_N^W(\hat{p}_{N-1}) = 0$ and such that for $p_{N-1} < \hat{p}_{N-1}$, $\phi_N^W(p_{N-1}) < 0$, and for $p_{N-1} > \hat{p}_{N-1}$, $\phi_N^W(p_{N-1}) > 0$. Next, assuming all of the properties stated in the lemma hold for agents $i \in \{N-k, ..., N\}$, where $k \in \{0, ..., N-2\}$, the proof shows that the properties hold for agent $N-k-1$. The key challenge is in establishing the *strict* monotonicity of ϕ_{N-k-1}^W. Using Lemma 2, ϕ_{N-k-1}^W can be written as

$$\begin{aligned} \phi_{N-k-1}^W(p_{N-k-2}) = \mathbb{E}[\max\{\phi_{N-k}^W(p_{N-k-1}(p_{N-k-2},\alpha)),0\}] \\ - \mathbb{E}[\max\{-\phi_{N-k}^W(p_{N-k-1}(p_{N-k-2},\beta)),0\}]. \end{aligned} \tag{4}$$

Each of the two components on the RHS of (4) is *weakly* increasing in p_{N-k-2}. The first component is equal to zero for beliefs $p_{N-k-2} \in [0,x]$ and strictly increasing for beliefs $p_{N-k-2} \in (x,1]$. The second component is strictly increasing for beliefs $p_{N-k-2} \in [0,y)$ and constant at zero for $p_{N-k-2} \in [y,1]$. From the induction assumption, there exists a unique belief $\hat{p}_{N-k-1} \in (0,1)$ such that $\phi_{N-k}^W(\hat{p}_{N-k-1}) = 0$, and such that $\phi_{N-k}^W(p_{N-k-1}) < 0$ if $p_{N-k-1} < \hat{p}_{N-k-1}$, and $\phi_{N-k}^W(p_{N-k-1}) > 0$ if $p_{N-k-1} > \hat{p}_{N-k-1}$. This property is used to derive the precise values of x and y and argue that $y > x$, which establishes that overall the sum of the two components on the RHS of (4) is strictly increasing, and in turn that $\phi_{N-k-1}^W(\cdot)$ is strictly increasing. The continuity of $\phi_{N-k-1}^W(\cdot)$ then follows from the induction assumption, Lemma 2, and the monotone convergence theorem. The existence and uniqueness of $\hat{p}_{N-k-2} \in (0,1)$ is then immediate.

Together, Lemmas 1 and 3 imply the following characterization of the socially optimal policy π^W using threshold beliefs.

Lemma 4 (Socially optimal policy in terms of ϕ^W)**.** *For every agent $i \in \{1, ..., N\}$, there exists a unique belief $\hat{p}_{i-1} \in (0, 1)$ such that the socially optimal policy π^W prescribes action β whenever $p_{i-1} < \hat{p}_{i-1}$, and action α whenever $p_{i-1} > \hat{p}_{i-1}$.*

Under the socially optimal policy, the last agent (agent N) is prescribed her myopically optimal action: α if $\phi_N^W(p_{N-1}) = \mu_\alpha(p_{N-1}) - \mu_\beta(p_{N-1}) > 0$, and β if $\phi_N^W(p_{N-1}) < 0$. In particular, as shown in the proof of Lemma 3, action α is prescribed whenever $p_{N-1} > \hat{p}_{N-1}$, and action β whenever $p_{N-1} < \hat{p}_{N-1}$, where $\hat{p}_{N-1}$ is given by (3).

In equilibrium, all agents behave as if they were the last agent. Denote by $\hat{p} \equiv \hat{p}_{N-1}$ the *myopic belief-threshold*. The following result, which follows from Lemma 4, states that the condition that for all agents the policy π^W specifies the same threshold belief $\hat{p}$ – below which β is chosen and above which α is chosen – is equivalent to $\pi^W \equiv \pi^E$.

Lemma 5 (Equilibrium efficiency and belief thresholds)**.** *The equilibrium policy and the socially optimal policy are identical (i.e., $\pi^E \equiv \pi^W$) if and only if $\hat{p}_{i-1} = \hat{p}$ for all agents $i \in \{1, ..., N\}$.*

The proof of Theorem 1 uses Lemma 5 together with Lemmas 2 and 3. To understand why these properties of ϕ^W imply that, under symmetric information structures, the equilibrium and efficient policies always coincide, note first that for such information structures, from (3), $\hat{p} = 1/2$. Furthermore, under symmetric information structures, for any agent i, the advantage function $\phi_i^W(\cdot)$ is symmetric around $1/2$, i.e., $\phi_i^W(p) = \phi_i^W(1-p)$. Hence, from Lemma 3, for any agent i, the advantage function must cross zero at a belief of $1/2$, i.e., $\phi_i^W(1/2) = 0$. That is, $\hat{p}_{i-1} = 1/2 = \hat{p}$. It follows from Lemma 5 that $\pi^E \equiv \pi^W$, for any prior and any number of agents.

Intuitively, we can think of the value of choosing an action as a balance between its immediate expected probability of success and its informational value. Importantly, however, and in contrast to environments with independent alternatives, by revealing information about the underlying state, the outcome of an action is typically informative about the future merits of *both* actions, rather than just the one chosen (and typically both success and failure are informative). The reason why Theorem 1 holds for symmetric information structures is that the action corresponding to the more likely state based on the current posterior is necessarily also the action that has higher informational value (i.e., more value for the purpose of learning the state). In particular, when the posterior is above (below) 1/2, so that state A (B) is more likely, action α (β) has a higher expected probability of success *and also* necessarily has a higher informational value. Hence, even though the informational value of an agent's decision may be very high, as is the case when there are many subsequent agents who will learn from it, it must always align with the agent's myopic interest.

Next, consider the case of competing-action information structures. To show that Theorem 1 holds for competing-actions information structures, the proof first shows that

$$\phi_N^W(p_{N-1}) = \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left(p_{N-1} - \hat{p}\right),$$

and that, for such information structures,

$$\mathbb{E}\left[\left|\phi_N^W\left(p_{N-1}(p_{N-2}, \sigma)\right)\right|\right] = -\mathbb{E}\left[\left|\phi_N^W\left(p_{N-1}(p_{N-2}, -\sigma)\right)\right|\right].$$

The proof then decomposes $\phi_{N-1}^W(p_{N-2})$ using Lemma 2, and uses these observations to show that for any belief p and any agent i, $\phi_i^W(p) = (p - \hat{p})\,(q_B^\beta - q_B^\alpha)/\hat{p}$. In particular, this means that for all agents, $\phi_i^W(\hat{p}) = 0$, which in turn implies that $\hat{p}_{i-1} = \hat{p}$. The result then follows from Lemma 5.

The intuition is the following. A key feature of such information structures is that

$$p_i(p_{i-1}, \sigma, 1) = p_i(p_{i-1}, -\sigma, 0). \tag{5}$$

That is, for any agent i and belief p_{i-1}, the posterior belief after agent i chooses one action σ and succeeds is the same as the posterior after agent i chooses the other action, $-\sigma$, and fails. In particular, regardless of the posterior, the informational value of α is always the same as that of β, as both actions induce the same distribution over posteriors. Since the informational value of both actions is the same, they differ only in their immediate expected probability of success.

Continuity

Importantly, Theorem 1 is informative not only about environments that fall exactly within the cases of symmetric information structures or information structures with competing actions. Continuity of ϕ_i^W and ϕ_i^E guarantee that for information structures "close" to $S \cup C$, the potential welfare loss from following the equilibrium policy instead of the socially optimal one must be "small".

Formally, given a set $X \subseteq Q$, let $B(X, \delta)$ denote the set of elements $q \in Q$ that are distanced less than $\delta > 0$ from some element in X. Denote by $V_i^W(p; q)$ and $V_i^E(p; q)$ the expected number of successes under the socially optimal policy π^W and the equilibrium policy π^E, respectively, starting from period i, when the belief is p, and given the information structure q.

Proposition 1 (Continuity)**.** *Fix a number of agents $N \geq 1$, any agent i, and a belief $p \in (0, 1)$. Let $\varepsilon > 0$. There exists $\delta > 0$ such that for all $\hat{q} \in B(S \cup C, \delta)$, $|V_i^W(p; \hat{q}) - V_i^E(p; \hat{q})| < \varepsilon$.*

Hence, even if $q \notin S \cup C$, if q is sufficiently close to $S \cup C$ then, despite the potential for a departure of the equilibrium policy from the socially optimal one, the resulting welfare loss is guaranteed to be small.

3.2 Necessity

We have shown that the condition that $q \in S \cup C$ is sufficient for the equilibrium and socially optimal policies to be the same for any prior p_0 and number of agents $N \geq 1$. We now turn to necessity, and show that this characterization is sharp for all $N \geq 3$.

Before doing so, it will be useful to study the case of two agents, $N = 2$. Under the efficient policy ϕ^W, agent 2 is prescribed α when $p_1 > \hat{p}_1$, and β for $p_1 < \hat{p}_1$. From Lemma 5, the equilibrium of the two-agent game maximizes welfare if and only if $\hat{p}_0 = \hat{p}$. Recall that $\kappa_\omega \equiv q_\omega^\alpha + q_\omega^\beta$. The following result gives a necessary and sufficient condition for when $\pi^E \equiv \pi^W$.

Proposition 2 (Characterization for $N = 2$)**.** *Assume $N = 2$. Then $\pi^E \equiv \pi^W$ if and only if $\kappa_A = \kappa_B$.*

The proof uses Lemma 2 to derive $\phi_1^W(\hat{p})$. Since, by definition, $\phi_1^W(\hat{p}_0) = 0$, if $\phi_1^W(\hat{p}) = 0$ as well then $\hat{p}_0 = \hat{p}_1 = \hat{p}$ and (by Lemma 5) the equilibrium policy π^E and the socially optimal policy π^W always choose the same actions. The proof shows that this boils down to whether $\kappa_A = \kappa_B$ or not.

Note that if an information structure is symmetric or has competing actions, it also satisfies $\kappa_A = \kappa_B$, but the converse is not true. That is, the condition that $q \in S \cup C$ is sufficient, but not necessary for π^E and π^W to be identical when $N = 2$.

Turning to the case of $N \geq 3$, the following result shows that the sufficient condition that $q \in S \cup C$ is also necessary. In other words, if the information structure is not symmetric and does not have competing actions, the equilibrium policy and the socially optimal policy need not agree.

Theorem 2 (Necessity)**.** *If $N \geq 3$ and $q \notin S \cup C$ then π^W and π^E differ. That is, for some priors and observed histories, the two policies specify different actions.*

Together, Theorems 1 and 2 yield a sharp characterization of the set of information structures under which, for any $N \geq 3$, the equilibrium and efficient policies are always the same, given any prior and observed history.[20]

Corollary 1 (Characterization for all $N \geq 3$)**.** *Let $N \geq 3$. Then $\pi^E \equiv \pi^W$ if and only if $q \in S \cup C$.*

To establish Theorem 2, first note that the condition that $\kappa_A = \kappa_B$ is necessary not only when $N = 2$, but for any $N \geq 3$. This is because, when this condition is not satisfied, just as Proposition 2 implies agent 1 can choose an inefficient action when $N = 2$, it also implies agent $N - 1$ can choose an inefficient action when $N \geq 3$. We will therefore assume from now on that

$$\kappa_A = \kappa_B \equiv \kappa.$$

Consider the cases $\kappa = 1$, $\kappa < 1$, and $\kappa > 1$ separately. Recall that the case of $\kappa = 1$ corresponds to the case of competing actions. We now consider the cases $\kappa < 1$ and $\kappa > 1$. Denote

$$\hat{\phi}(p) = \frac{(q_B^\beta - q_B^\alpha)(p - \hat{p})}{\hat{p}},$$

[20] Note that the results state that for any $N \geq 3$, π^W and π^E are the same *for any prior* p_0 if and only if $q \in S \cup C$. For a *given* prior, there are potentially additional information structures, outside of $S \cup C$, under which the two policies π^W and π^E agree. However, it can be shown that for information structures outside of the set $S \cup C$, for any given prior there exists a number of agents N such that $\pi^W \not\equiv \pi^E$.

and let

$$x_1^\alpha = \frac{\hat{p}q_B^\alpha}{\hat{p}q_B^\alpha + q_A^\alpha(1-\hat{p})} \quad , \quad x_2^\alpha = \frac{\hat{p}(1-q_B^\alpha)}{\hat{p}(1-q_B^\alpha) + (1-q_A^\alpha)(1-\hat{p})}$$
$$x_1^\beta = \frac{\hat{p}q_B^\beta}{\hat{p}q_B^\beta + q_A^\beta(1-\hat{p})} \quad , \quad x_2^\beta = \frac{\hat{p}(1-q_B^\beta)}{(1-q_A^\beta)(1-\hat{p}) + \hat{p}(1-q_B^\beta)}.$$

We start by fully characterizing $\phi_{N-1}^W(p)$ for any belief p.

Lemma 6 (Characterizing $\phi_{N-1}^W(\cdot)$)**.** *If $\kappa < 1$, then $x_1^\alpha < x_2^\beta < x_2^\alpha < x_1^\beta$ and*

$$\phi_{N-1}^W(p) = \begin{cases} \hat{\phi}(p) & \text{if } \ p < x_1^\alpha \\ \hat{\phi}(p) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[q_A^\alpha(1-\hat{p})p - \hat{p}q_B^\alpha(1-p)\right] & \text{if } \ x_1^\alpha < p < x_2^\beta \\ \kappa\hat{\phi}(p) & \text{if } \ x_2^\beta < p < x_2^\alpha \\ \hat{\phi}(p) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[pq_A^\beta(1-\hat{p}) - q_B^\beta\hat{p}(1-p)\right] & \text{if } \ x_2^\alpha < p < x_1^\beta \\ \hat{\phi}(p) & \text{if } \ p > x_1^\beta \end{cases}.$$

Similarly, if $\kappa > 1$, then $x_2^\beta < x_1^\alpha < x_1^\beta < x_2^\alpha$, and

$$\phi_{N-1}^W(p) = \begin{cases} \hat{\phi}(p) & \text{if } \ p < x_2^\beta \\ \hat{\phi}(p) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[\hat{p}(1-q_B^\beta)(1-p) - p(1-q_A^\beta)(1-\hat{p})\right] & \text{if } \ x_2^\beta < p < x_1^\alpha \\ \kappa\hat{\phi}(p) & \text{if } \ x_1^\alpha < p < x_1^\beta \\ \hat{\phi}(p) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[\hat{p}(1-q_B^\alpha)(1-p) - (1-q_A^\alpha)(1-\hat{p})p\right] & \text{if } \ x_1^\beta < p < x_2^\alpha \\ \hat{\phi}(p) & \text{if } \ p > x_2^\alpha \end{cases}.$$

Figure 3 illustrates the advantage functions ϕ_{N-1}^W (grey line) for several examples, when $\kappa < 1$ (left figures) and $\kappa > 1$ (right figures), and for the case of symmetric information structures (bottom figures) and asymmetric ones (top figures), comparing each to the corresponding advantage functions under the equilibrium policy (dotted line). Consistent with Proposition 2, in each of the examples $\kappa_A = \kappa_B = \kappa$, and hence the equilibrium and socially optimal policies are identical for agent $N-1$; this can be seen by the fact that in each of the examples the dotted line and the grey line both intersect with the horizontal axis at the same belief, $\hat{p}$.

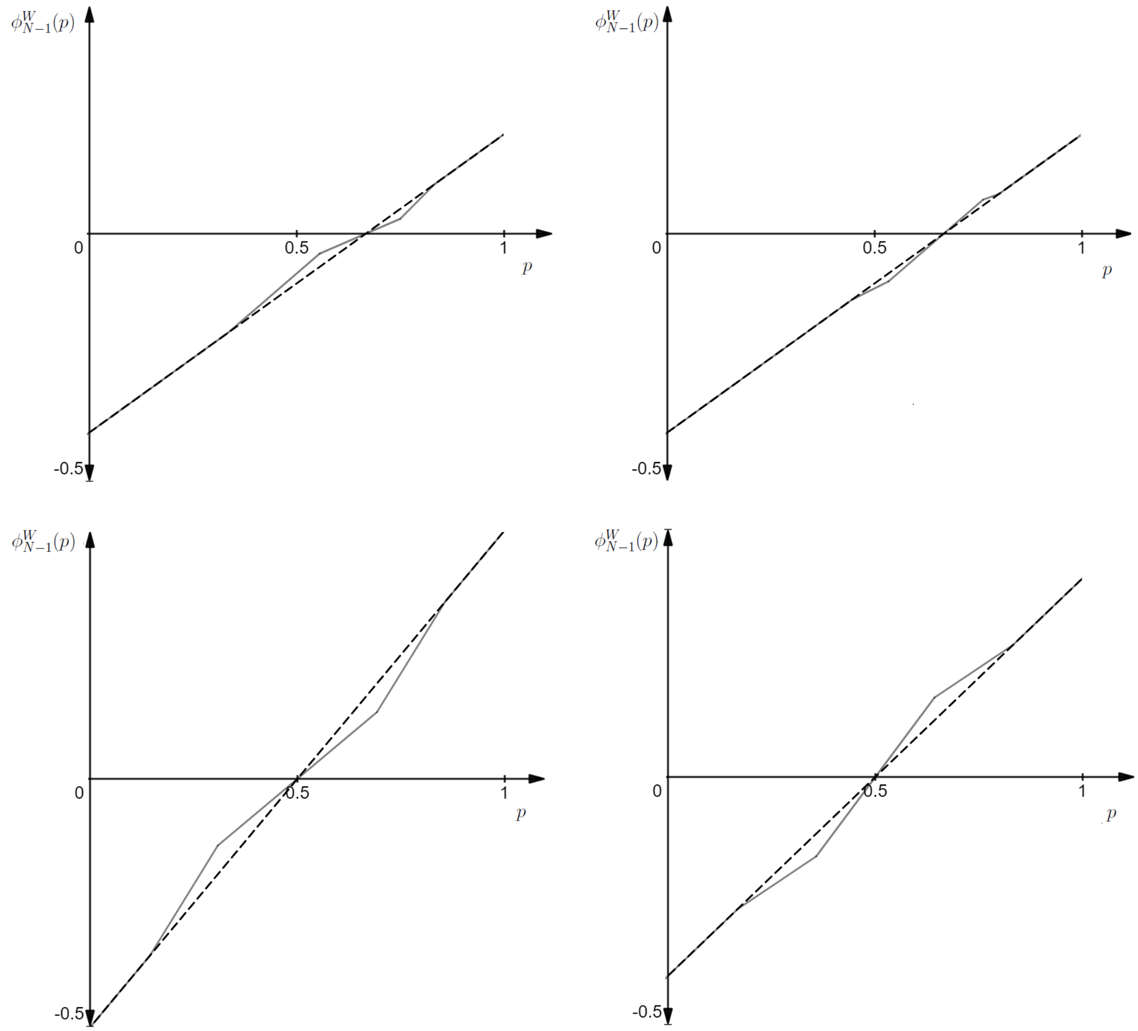

Figure 3: The advantage function ϕ^W_{N-1} for agent $N-1$ under the socially optimal policy (grey line). The figures correspond to the information structures $(q_A^\alpha, q_A^\beta, q_B^\alpha, q_B^\beta) = (0.4, 0.2, 0.1, 0.5)$ (top left), $(q_A^\alpha, q_A^\beta, q_B^\alpha, q_B^\beta) = (0.7, 0.5, 0.4, 0.8)$ (top right), $(q_A^\alpha, q_A^\beta, q_B^\alpha, q_B^\beta) = (0.6, 0.1, 0.1, 0.6)$ (bottom left), and $(q_A^\alpha, q_A^\beta, q_B^\alpha, q_B^\beta) = (0.9, 0.5, 0.5, 0.9)$ (bottom right). In particular, note that for the figures on the left (right), $\kappa < 1$ ($\kappa > 1$), and that the examples corresponding to the bottom figures correspond to symmetric information structures. The dotted line represents the advantage function under the equilibrium policy.

Using this Lemma, the proof of Theorem 2 derives $\phi^W_{N-2}(\hat{p})$ by determining for each possible posterior $p_{N-2}(p_{N-3}, \sigma, \theta)$ where it falls within the regions determined by $x^\alpha_1, x^\alpha_2, x^\beta_1, x^\beta_2$ (the ordering of which itself depends on the value of κ). For instance, in the Appendix it is shown that for $\kappa < 1$,

$$\phi^W_{N-2}(\hat{p}) = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B) \times \begin{cases} \kappa - 1 & \text{if } q^\alpha_A + q^\alpha_B > 1 \\ \kappa - (q^\alpha_A + q^\alpha_B) & \text{if } 2\kappa - 1 < q^\alpha_A + q^\alpha_B < 1 \\ 1 - \kappa & \text{if } q^\alpha_A + q^\alpha_B < 2\kappa - 1 \text{ and } \kappa > \frac{1}{2}. \end{cases}$$

In particular, this implies that it is only possible that $\phi^W_{N-2}(\hat{p}) = 0$ when $\kappa - (q^\alpha_A + q^\alpha_B) = 0$, where the latter implies $q^\beta_A = q^\alpha_B$ and $q^\alpha_A = q^\beta_B$. In other words, if $\kappa < 1$, then it is only possible that $\phi^W_{N-2}(\hat{p}) = 0$ under a symmetric information structure. This property also holds for any information structure with $\kappa > 1$. Theorem 2 then follows from Lemma 5.

4 AI recommendation platforms

AI recommendation platforms often face a trade-off between offering users the best-known option based on existing data and trying potentially suboptimal options to improve future performance. While this may be harmless in environments like entertainment, it raises serious ethical concerns in more sensitive domains such as healthcare, credit, hiring, and matchmaking. In these contexts, individuals may receive inferior recommendations not because it benefits them, but because their responses help the platform learn and improve over time. For example, in 2014, OkCupid publicly revealed on its official blog (titled "We Experiment On Human Beings!") that it had conducted several A/B tests on users without informing them (see, e.g., Luca, 2014). These included recommending suboptimal matches[21] and sometimes pairing users without algorithmic input. OkCupid defended these actions as standard practice in online platforms, prompting significant public backlash and raising ethical concerns about user manipulation and transparency.

Beyond the OkCupid example, such practices are often carried out without the knowledge or consent of users, undermining their autonomy and potentially increasing

[21]This was done by manipulating compatibility scores – showing low-match users as highly compatible – to learn about user behavior.

the risk of unjustified harm (Mittelstadt et al., 2016; Binns, 2018). Unlike formal experiments where participants are informed and consent is obtained, most platforms do not disclose such experimentation or seek user permission, even when the consequences may be significant. This has led to calls for transparency, consent mechanisms, and context-specific limits on such experimentation (Floridi et al., 2018; Cowgill et al., 2020). An important question this debate has raised is: In what scenarios is there a need for regulation and when is there no cause for such concerns?

The observational learning model above can serve as a modest step toward a better understanding of this question. In the model of Section 2, a sequence of short-lived agents observe the past actions and outcomes of their predecessors before choosing an action. Now consider instead an environment in which agents do not observe data on past outcomes themselves but instead follow the recommendations of an AI platform with access to this information. The platform's objective is to maximize the overall expected number of successes across users.[22] As a result, when it recommends an action to a given agent, it weighs not only the expected outcome for that agent but also the value of the information that the agent's response provides for improving future recommendations.

The results in the previous sections show that when the information structure is either symmetric (i.e., each of the actions performs equally well when adopted in the state it is meant for) or has competing actions (i.e., where in each state, the actions are inversely effective: the more effective is one action, the less effective is the other) then there is no misalignment of interests between the platform and its users. Furthermore, in the neighborhood of the set of symmetric or competing-actions information structures, the welfare loss is relatively small. Further away from these cases, the need for regulation may be more significant.

5 Concluding remarks and extensions

This paper has developed a simple model of observational learning to study when technology adoption is efficient without the need for outside intervention. A sequence of agents each choose among risky technologies, the outcomes of which are a function

[22]Clearly, such platforms are generally driven by profit maximization rather than user welfare. The platform's profit maximization problem is not modeled explicitly here, but a similar tension would arise in such a setting.

of an unknown state. Agents learn about the state from the past decisions of their predecessors, as well as their outcomes. The agents are short-lived, and hence do not internalize the informational value of their choice for future agents (which can be very large if there are many agents remaining), potentially yielding a conflict between equilibrium and efficient technology adoption.

When the information structure is symmetric (i.e., environments where each of the actions performs equally well when adopted in the state it is meant for), or when the information structure features competing actions (i.e., environments with inversely effective actions, in the sense that in any given state, if one action yields a high probability of success, the other yields a low one), technology adoption is "robustly" efficient: given any prior, any number of agents, and any history of decisions and outcomes, agents will always choose the action that is socially optimal. In such cases, there is no need to incentivize agents – they make the socially optimal choice even without internalizing their externality. Similarly, when the information structure is close to one of these two classes, the inefficiency must be small. The characterization is sharp in that for information structures outside of these two sets there is always a prior and number of agents such that technology adoption is inefficient.

Theorems 1 and 2 continue to hold in an environment with discounting and an infinite sequence of agents. Furthermore, consider an extension of the model to the case of more than two states. Under competing-actions information structures – i.e., if for each possible state ω, $q^\alpha_\omega + q^\beta_\omega = 1$ – for any prior and any number of agents, the equilibrium and socially optimal policies coincide.[23] The following figure illustrates.

	A	B	C
α	q^α_A	q^α_B	q^α_C
β	q^β_A	q^β_B	q^β_C

Figure 4. Three states, two actions.

When each of the columns in this table sums to one, the equilibrium and socially optimal policies always coincide. Note that in this case there is no immediate definition of symmetric information structures though, and no immediate analogue of Theorem 2.

[23]As with two states, there are also the trivial information structures where one action has a higher probability of success than the others for all states, in which case the equilibrium and efficient policies are also the same.

A Appendix

Proof of Lemma 2. For any $i = 1, ..., N$, $\phi_i^W(p_{i-1})$ can be written as follows:

$$
\begin{aligned}
\phi_i^W(p_{i-1}) &= \mu_\alpha(p_{i-1}) - \mu_\beta(p_{i-1}) + \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1},\alpha))\right] - \mathbb{E}\left[V_{i+1}^W(p_i(p_{i-1},\beta))\right] \\
&= \mu_\alpha(p_{i-1}) - \mu_\beta(p_{i-1}) \\
&\quad + \mathbb{E}\Bigg[\max\Big\{\mu_\alpha(p_i(p_{i-1},\alpha)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\alpha)\right)\right] \\
&\qquad\qquad , \ \mu_\beta(p_i(p_{i-1},\alpha)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\beta)\right)\right]\Big\}\Bigg] \\
&\quad - \mathbb{E}\Bigg[\max\Big\{\mu_\alpha(p_i(p_{i-1},\beta)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\alpha)\right)\right] \\
&\qquad\qquad , \ \mu_\beta(p_i(p_{i-1},\beta)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\beta)\right)\right]\Big\}\Bigg] \\
&= \mu_\alpha(p_{i-1}) - \mu_\beta(p_{i-1}) \\
&\quad + \mathbb{E}\Bigg[\max\Big\{\mu_\alpha(p_i(p_{i-1},\alpha)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\alpha)\right)\right] \\
&\qquad\qquad - \mu_\beta(p_i(p_{i-1},\alpha)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\beta)\right)\right] \ , \ 0\Big\}\Bigg] \\
&\quad + \mu_\beta(p_i(p_{i-1},\alpha)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\beta)\right)\right] \\
&\quad - \mathbb{E}\Bigg[\max\Big\{0 \ , \ \mu_\beta(p_i(p_{i-1},\beta)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\beta)\right)\right] \\
&\qquad\qquad - \mu_\alpha(p_i(p_{i-1},\beta)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\alpha)\right)\right]\Big\}\Bigg] \\
&\quad - \mu_\alpha(p_i(p_{i-1},\beta)) + \mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\alpha)\right)\right] \\
&= \mathbb{E}\left[\max\left\{\phi_{i+1}^W\left(p_i(p_{i-1},\alpha)\right), 0\right\}\right] - \mathbb{E}\left[\max\left\{-\phi_{i+1}^W\left(p_i(p_{i-1},\beta)\right), 0\right\}\right] \\
&\quad + \left(\mu_\alpha(p_{i-1}) + \mathbb{E}\left[\mu_\beta(p_i(p_{i-1},\alpha))\right]\right) - \left(\mu_\beta(p_{i-1}) + \mathbb{E}\left[\mu_\alpha(p_i(p_{i-1},\beta))\right]\right) \\
&\quad + \mathbb{E}\left[\mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\beta)\right)\right]\right] - \mathbb{E}\left[\mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\alpha)\right)\right]\right].
\end{aligned}
$$

Since

$$\mu_\alpha(p_{i-1}) + \mathbb{E}\left[\mu_\beta(p_i(p_{i-1},\alpha))\right] = \mu_\beta(p_{i-1}) + \mathbb{E}\left[\mu_\alpha(p_i(p_{i-1},\beta))\right]$$

and

$$\mathbb{E}\left[\mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\alpha),\beta)\right)\right]\right] = \mathbb{E}\left[\mathbb{E}\left[V_{i+2}^W\left(p_{i+1}(p_i(p_{i-1},\beta),\alpha)\right)\right]\right],$$

this completes the proof. ■

Proof of Lemma 3. When it is agent N's turn to move,

$$\begin{aligned}
\phi_N^W(p_{N-1}) &= \mu_\alpha(p_{N-1}) - \mu_\beta(p_{N-1}) \\
&= p_{N-1}q_A^\alpha + (1-p_{N-1})q_B^\alpha - \left(p_{N-1}q_A^\beta + (1-p_{N-1})q_B^\beta\right) \\
&= q_B^\alpha - q_B^\beta + p_{N-1}\left(q_A^\alpha - q_B^\alpha - q_A^\beta + q_B^\beta\right).
\end{aligned}$$

Clearly, $\phi_N^W(p_{N-1})$ is a continuous function of p_{N-1}, and since $q_A^\alpha > q_A^\beta$ and $q_B^\alpha < q_B^\beta$, it is also strictly increasing, which proves part 1 for agent N. Note that $\phi_N^W(0) = q_B^\alpha - q_B^\beta < 0$ and that $\phi_N^W(1) = q_A^\alpha - q_A^\beta > 0$. Hence, there exists a unique belief p_{N-1} for which $\phi_N^W(p_{N-1}) = 0$. In particular, this threshold belief is given by:

$$\hat{p}_{N-1} = \frac{q_B^\beta - q_B^\alpha}{q_A^\alpha - q_A^\beta + q_B^\beta - q_B^\alpha} \in (0,1).$$

For $p_{N-1} < \hat{p}_{N-1}$, $\phi_N^W(p_{N-1}) < 0$, and for $p_{N-1} > \hat{p}_{N-1}$, $\phi_N^W(p_{N-1}) > 0$. We have therefore established the result for agent N.

We now prove the result for the remaining agents by induction. Suppose the properties in parts 1 and 2 hold for agents $i \in \{N-k, ..., N\}$, where $k \in \{0, ..., N-2\}$. We will now show that they also hold for agent $N-k-1$.

By the induction assumption, $\phi_{N-k}^W(p_{N-k-1})$ is strictly increasing in p_{N-k-1}. From equations (1) and (2), it can be verified that for each $\sigma \in \{\alpha,\beta\}$ and $\theta \in \{0,1\}$, $p_{N-k-1}(p_{N-k-2},\sigma,\theta)$ is a strictly increasing function of p_{N-k-2}. It follows that, for each $\theta \in \{0,1\}$. both $\phi_{N-k}^W(p_{N-k-1}(p_{N-k-2},\alpha,\theta))$ and $\phi_{N-k}^W(p_{N-k-1}(p_{N-k-2},\beta,\theta))$ are also strictly increasing functions of p_{N-k-2}. We need to show that ϕ_{N-k-1} is strictly increasing in p_{N-k-2}.

From Lemma 2,

$$\begin{aligned}\phi^W_{N-k-1}(p_{N-k-2}) = \mathbb{E}[\max\{\phi^W_{N-k}\left(p_{N-k-1}(p_{N-k-2},\alpha)\right),0\}] \\ - \mathbb{E}[\max\{-\phi^W_{N-k}\left(p_{N-k-1}(p_{N-k-2},\beta)\right),0\}].\end{aligned}$$

From the observations in the previous paragraph, each of the two components on the right hand side of this equation is *weakly* increasing in p_{N-k-2}. To show that, overall, the sum of the two components is strictly increasing, we now characterize the regions in which each of the components is strictly increasing.

The expectation

$$\mathbb{E}[\max\{\phi^W_{N-k}\left(p_{N-k-1}(p_{N-k-2},\alpha)\right),0\}] \tag{6}$$

is equal to zero for beliefs $p_{N-k-2} \in [0,x]$ and strictly increasing for beliefs $p_{N-k-2} \in (x,1]$. To derive this x, first note that from (1) and (2),

$$p_{N-k-1}(p_{N-k-2},\sigma,1) = \frac{p_{N-k-2}q^\sigma_A}{p_{N-k-2}q^\sigma_A + (1-p_{N-k-2})q^\sigma_B} \tag{7}$$

and

$$p_{N-k-1}(p_{N-k-2},\sigma,0) = \frac{p_{N-k-2}(1-q^\sigma_A)}{p_{N-k-2}(1-q^\sigma_A) + (1-p_{N-k-2})(1-q^\sigma_B)}. \tag{8}$$

From the induction assumption, there exists a unique belief $\hat{p}_{N-k-1} \in (0,1)$ such that $\phi^W_{N-k}(\hat{p}_{N-k-1}) = 0$, and such that $\phi^W_{N-k}(p_{N-k-1}) < 0$ if $p_{N-k-1} < \hat{p}_{N-k-1}$, and $\phi^W_{N-k}(p_{N-k-1}) > 0$ if $p_{N-k-1} > \hat{p}_{N-k-1}$. Hence, (6) is strictly increasing if and only if either

$$p_{N-k-1}(p_{N-k-2},\alpha,1) = \frac{p_{N-k-2}q^\alpha_A}{p_{N-k-2}q^\alpha_A + (1-p_{N-k-2})q^\alpha_B} > \hat{p}_{N-k-1}$$

or

$$p_{N-k-1}(p_{N-k-2},\alpha,0) = \frac{p_{N-k-2}(1-q^\alpha_A)}{p_{N-k-2}(1-q^\alpha_A) + (1-p_{N-k-2})(1-q^\alpha_B)} > \hat{p}_{N-k-1}.$$

Rearranging these inequalities, we conclude that (6) is strictly increasing if and only if either

$$p_{N-k-2} > \frac{q_B^\alpha \hat{p}_{N-k-1}}{q_B^\alpha \hat{p}_{N-k-1} + q_A^\alpha (1 - \hat{p}_{N-k-1})}$$

or

$$p_{N-k-2} > \frac{(1 - q_B^\alpha)\hat{p}_{N-k-1}}{(1 - q_B^\alpha)\hat{p}_{N-k-1} + (1 - q_A^\alpha)\,(1 - \hat{p}_{N-k-1})}.$$

Hence,

$$x = \min\left\{ \frac{q_B^\alpha \hat{p}_{N-k-1}}{q_B^\alpha \hat{p}_{N-k-1} + q_A^\alpha (1 - \hat{p}_{N-k-1})}, \frac{(1 - q_B^\alpha)\hat{p}_{N-k-1}}{(1 - q_B^\alpha)\hat{p}_{N-k-1} + (1 - q_A^\alpha)\,(1 - \hat{p}_{N-k-1})} \right\}.$$

The second component in the decomposition of $\phi_{N-k-1}^W(p_{N-k-2})$, i.e., expectation

$$-\mathbb{E}[\max\{-\phi_{N-k}^W\left(p_{N-k-1}(p_{N-k-2}, \beta)\right), 0\}], \tag{9}$$

is strictly increasing for beliefs $p_{N-k-2} \in [0, y)$ and constant at zero for $p_{N-k-2} \in [y, 1]$. To derive y, note that the expectation (9) is equal to zero if and only if both

$$p_{N-k-1}(p_{N-k-2}, \beta, 1) = \frac{p_{N-k-2} q_A^\beta}{p_{N-k-2} q_A^\beta + (1 - p_{N-k-2}) q_B^\beta} < \hat{p}_{N-k-1}$$

and

$$p_{N-k-1}(p_{N-k-2}, \beta, 0) = \frac{p_{N-k-2}(1 - q_A^\beta)}{p_{N-k-2}(1 - q_A^\beta) + (1 - p_{N-k-2})(1 - q_B^\beta)} < \hat{p}_{N-k-1}.$$

Rearranging these inequalities, we conclude that (9) is strictly increasing if and only if

$$p_{N-k-2} < \frac{q_B^\beta \hat{p}_{N-k-1}}{q_B^\beta \hat{p}_{N-k-1} + q_A^\beta\,(1 - \hat{p}_{N-k-1})}$$

and

$$p_{N-k-2} < \frac{(1 - q_B^\beta)\hat{p}_{N-k-1}}{(1 - q_B^\beta)\hat{p}_{N-k-1} + (1 - q_A^\beta)(1 - \hat{p}_{N-k-1})}.$$

Hence,

$$y = \max\left\{ \frac{q_B^\beta \hat{p}_{N-k-1}}{q_B^\beta \hat{p}_{N-k-1} + q_A^\beta\,(1-\hat{p}_{N-k-1})}, \frac{(1-q_B^\beta)\hat{p}_{N-k-1}}{(1-q_B^\beta)\hat{p}_{N-k-1} + (1-q_A^\beta)(1-\hat{p}_{N-k-1})} \right\}.$$

We have therefore argued that both (6) and (9) are weakly increasing, that (6) is strictly increasing in the interval $(x, 1]$, and that (9) is strictly increasing in the interval $[0, y)$. To see that $y > x$, note that

$$y \geq \frac{q_B^\beta \hat{p}_{N-k-1}}{q_B^\beta \hat{p}_{N-k-1} + q_A^\beta\,(1-\hat{p}_{N-k-1})} > \frac{q_B^\alpha \hat{p}_{N-k-1}}{q_B^\alpha \hat{p}_{N-k-1} + q_A^\alpha (1-\hat{p}_{N-k-1})} \geq x,$$

where the strict inequality holds as it can be rearranged to obtain $q_B^\beta q_A^\alpha > q_B^\alpha q_A^\beta$ (which holds since $q_A^\alpha > q_B^\alpha$ and $q_B^\beta > q_A^\beta$).

It therefore follows that the sum of the two components (6) and (9) is always strictly increasing, and hence from Lemma 2, $\phi_{N-k-1}^W(p_{N-k-2})$ is strictly increasing in p_{N-k-2}. Using this monotonicity, Lemma 2, and the continuity of $\phi_{N-k}^W(p_{N-k-1})$ (from the induction assumption), the monotone convergence theorem yields the continuity of $\phi_{N-k-1}^W(p_{N-k-2})$. Finally, from the induction assumption, $\phi_{N-k}^W(0) < 0$ and $\phi_{N-k}^W(1) > 0$. Hence, since $\phi_{N-k-1}^W(0) = \phi_{N-k}^W(0) < 0$ and $\phi_{N-k-1}^W(1) = \phi_{N-k}^W(1) > 0$, there exists a unique belief $\hat{p}_{N-k-2} \in (0,1)$ such that $\phi_{N-k-1}^W(\hat{p}_{N-k-2}) = 0$, and such that $\phi_{N-k-1}^W(p_{N-k-2}) < 0$ if $p_{N-k-2} < \hat{p}_{N-k-2}$, and $\phi_{N-k-1}^W(p_{N-k-2}) > 0$ if $p_{N-k-2} > \hat{p}_{N-k-2}$. ■

Proof of Theorem 1. Starting with symmetric information structures, it is easy to verify that for any $i \in \{1, ..., N\}$, $\phi_i^W(\cdot)$ is symmetric around 1/2, i.e., $\phi_i^W(p) = \phi_i^W(1-p)$. Hence, from Lemma 3, for any $i \in \{1, ..., N\}$, $\phi_i^W(1/2) = 0$ and $\hat{p}_{i-1} = 1/2$. It follows from Lemma 5 that $\pi^E \equiv \pi^W$.

Turning to competing-actions information structures, first note that

$$\begin{aligned}
\phi_N^W(p_{N-1}) &= \mu_\alpha(p_{N-1}) - \mu_\beta(p_{N-1}) \\
&= (p_{N-1}q_A^\alpha + (1-p_{N-1})q_B^\alpha) - (p_{N-1}(1-q_A^\alpha) + (1-p_{N-1})(1-q_B^\alpha)) \\
&= p_{N-1}\left(q_A^\alpha - q_A^\beta + q_B^\beta - q_B^\alpha\right) - \left(q_B^\beta - q_B^\alpha\right) \\
&= \left(q_A^\alpha - q_A^\beta + q_B^\beta - q_B^\alpha\right)(p_{N-1} - \hat{p})
\end{aligned}$$

$$= \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \left(p_{N-1} - \hat{p} \right).$$

Hence,

$$\phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 1) \right) = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \left(\frac{p_{N-2} q_A^\sigma}{p_{N-2} q_A^\sigma + (1 - p_{N-2}) q_B^\sigma} - \hat{p} \right)$$

and

$$\phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 0) \right) = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \left(\frac{p_{N-2}(1 - q_A^\sigma)}{p_{N-2}(1 - q_A^\sigma) + (1 - p_{N-2})(1 - q_B^\sigma)} - \hat{p} \right).$$

Multiplying $\phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 1) \right)$ by the probability of success given p_{N-2} and the action σ,

$$\begin{aligned} &(p_{N-2} q_A^\sigma + (1 - p_{N-2}) q_B^\sigma) \cdot \phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 1) \right) \\ &\quad = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \left(p_{N-2} q_A^\alpha (1 - \hat{p}) - \hat{p} q_B^\alpha \left(1 - p_{N-2} \right) \right). \end{aligned}$$

Analogously, multiplying $\phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 0) \right)$ by the probability of failure yields

$$\begin{aligned} &(p_{N-2}(1 - q_A^\sigma) + (1 - p_{N-2})(1 - q_B^\sigma)) \cdot \phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 0) \right) \\ &\quad = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \left(p_{N-2}(1 - q_A^\sigma)(1 - \hat{p}) - \hat{p}(1 - q_B^\sigma)(1 - p_{N-2}) \right). \end{aligned}$$

Similarly,

$$\begin{aligned} &(p_{N-2} q_A^\sigma + (1 - p_{N-2}) q_B^\sigma) \cdot \left| \phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 1) \right) \right| \\ &\quad = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \cdot \left| p_{N-2} q_A^\sigma - \hat{p}(p_{N-2} q_A^\sigma + (1 - p_{N-2}) q_B^\sigma) \right| \end{aligned}$$

and

$$\begin{aligned} &(p_{N-2}(1 - q_A^\sigma) + (1 - p_{N-2})(1 - q_B^\sigma)) \cdot \left| \phi_N^W \left(p_{N-1}(p_{N-2}, \sigma, 1) \right) \right| \\ &\quad = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \cdot \left| p_{N-2}(1 - q_A^\sigma)(1 - \hat{p}) - \hat{p}(1 - q_B^\sigma)(1 - p_{N-2}) \right|. \end{aligned}$$

Calculating the expectation of $\phi_N^W(p_{N-1}(p_{N-2},\sigma))$, we have

$$\mathbb{E}\left[\phi_N^W(p_{N-1}(p_{N-2},\sigma))\right] = \frac{q_B^\beta - q_B^\alpha}{\hat{p}}(p_{N-2}-\hat{p}). \tag{10}$$

Replacing p_{N-2} with $\hat{p}$ yields $\mathbb{E}\left[\phi_N^W(p_{N-1}(\hat{p},\sigma))\right] = 0$. Similarly,

$$\begin{aligned}\mathbb{E}\left[|\phi_N^W(p_{N-1}(p_{N-2},\sigma))|\right] = \frac{q_B^\beta - q_B^\alpha}{\hat{p}} \times \Big[&|p_{N-2}q_A^\alpha(1-\hat{p}) - \hat{p}(1-p_{N-2})q_B^\alpha| \\ &+ |p_{N-2}(1-q_A^\sigma)(1-\hat{p}) - \hat{p}(1-q_B^\sigma)(1-p_{N-2})|\Big].\end{aligned}$$

Recall that for $q \in C$, (5) holds. Hence,

$$\mathbb{E}\left[|\phi_N^W(p_{N-1}(p_{N-2},\sigma))|\right] = -\mathbb{E}\left[|\phi_N^W(p_{N-1}(p_{N-2},-\sigma))|\right]. \tag{11}$$

Decomposing $\phi_{N-1}^W(p_{N-2})$ using Lemma 2,

$$\begin{aligned}\phi_{N-1}^W(p_{N-2}) &= \mathbb{E}[\max\{\phi_N^W(p_{N-1}(p_{N-2},\alpha)),0\}] - \mathbb{E}[\max\{-\phi_N^W(p_{N-1}(p_{N-2},\beta)),0\}] \\ &= \frac{1}{2}\mathbb{E}\left[\phi_N^W(p_{N-1}(p_{N-2},\alpha)) + |\phi_N^W(p_{N-1}(p_{N-2},\alpha))|\right] \\ &\quad - \frac{1}{2}\mathbb{E}\left[-\phi_N^W(p_{N-1}(p_{N-2},\beta)) + |\phi_i^W(p_{N-1}(p_{N-2},\beta))|\right] \\ &= \frac{q_B^\beta - q_B^\alpha}{\hat{p}}(p_{N-2}-\hat{p}),\end{aligned}$$

where the last equality follows from (10) and (11).

We have therefore shown that for any belief p, $\phi_{N-1}^W(p) = \phi_N^W(p)$. The exact same argument above can be repeated to show that for any p, $\phi_{N-2}^W(p) = \phi_{N-1}^W(p)$. Hence, for all $i \in \{1,...,N\}$, $\phi_i^W(p) = (p-\hat{p})(q_B^\beta - q_B^\alpha)/\hat{p}$. In particular, this implies that for all $i \in \{1,...,N\}$, $\phi_i^W(\hat{p}) = 0$, which in turn implies that for all $i \in \{1,...,N\}$, $\hat{p}_{i-1} = \hat{p}$. The result therefore follows from Lemma 5. ■

Proof of Proposition 1. First, recall that for any i and p,

$$\phi_i^E(p;q) = \mu_\alpha(p;q) - \mu_\beta(p;q) = q_B^\alpha - q_B^\beta + p\left(q_A^\alpha - q_B^\alpha - q_A^\beta + q_B^\beta\right).$$

For $\phi_i^W(p;q)$ the proof is by induction. Note first that

$$\phi_N^W(p_{N-1};q) = \mu_\alpha(p_{N-1};q) - \mu_\beta(p_{N-1};q) = q_B^\alpha - q_B^\beta + p_{N-1}\left(q_A^\alpha - q_B^\alpha - q_A^\beta + q_B^\beta\right).$$

Suppose continuity holds for agents $i \in \{N-k,...,N\}$, where $k \in \{0,...,N-2\}$. From Lemma 2, the continuity of $\phi_{N-k}^W(p_{N-k-1};q)$ implies the continuity of $\phi_{N-k-1}^W(p_{N-k-2};q)$. Hence $\phi_i^W(p;q)$ is continuous in q for every i.

The result therefore follows from the continuity of $\phi_i^E(p;q)$ and $\phi_i^W(p;q)$ in q, together with Theorem 1. ■

Proof of Proposition 2. We use Lemma 2 to derive $\phi_1^W(\hat{p})$. As in the proof of Theorem 1, $\phi_2^W(p_1) = (q_B^\beta - q_B^\alpha)(p_1 - \hat{p})/\hat{p}$. Following similar steps as in the proof of Theorem 1,

$$\phi_2^W(p_1(p_0,\sigma,1)) = \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left(\frac{p_0 q_A^\sigma}{p_0 q_A^\sigma + (1-p_0)q_B^\sigma} - \hat{p}\right)$$

and

$$\phi_2^W(p_1(p_0,\sigma,0)) = \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left(\frac{p_0(1-q_A^\sigma)}{p_0(1-q_A^\sigma) + (1-p_0)(1-q_B^\sigma)} - \hat{p}\right).$$

Calculating the expectation of $\phi_2^W(p_1(p_0,\sigma))$ yields

$$\mathbb{E}\left[\phi_2^W(p_1(p_0,\sigma))\right] = \frac{q_B^\beta - q_B^\alpha}{\hat{p}}(p_0 - \hat{p}),$$

and replacing p_0 with $\hat{p}$ yields $\mathbb{E}\left[\phi_2^W(p_1(\hat{p},\sigma))\right] = 0$. Similarly, $\mathbb{E}\left[|\phi_2^W(p_1(p_0,\sigma))|\right]$ is equal to

$$\frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[|p_0 q_A^\sigma - \hat{p}(p_0 q_A^\sigma + (1-p_0)q_B^\sigma)| + |p_0(1-q_A^\sigma)(1-\hat{p}) - \hat{p}(1-q_B^\sigma)(1-p_0)|\right],$$

and, replacing p_0 with $\hat{p}_1$,

$$\mathbb{E}\left[|\phi_2^W(p_1(\hat{p},\sigma))|\right] = 2(q_B^\beta - q_B^\alpha)(1-\hat{p}) \cdot |q_A^\sigma - q_B^\sigma|.$$

We can now calculate $\phi_1^W(\hat{p})$:

$$\begin{aligned}
\phi_1^W(\hat{p}) &= \mathbb{E}[\max\{\phi_2^W(p_1(\hat{p},\alpha)),0\}] - \mathbb{E}[\max\{-\phi_2^W(p_1(\hat{p},\beta)),0\}] \\
&= \frac{1}{2}\mathbb{E}\left[\phi_2^W(p_1(\hat{p},\alpha)) + |\phi_2^W(p_1(\hat{p},\alpha))|\right] - \frac{1}{2}\mathbb{E}\left[-\phi_2^W(p_1(\hat{p},\beta)) + |\phi_2^W(p_1(\hat{p},\beta))|\right] \\
&= \frac{1}{2}\mathbb{E}\left[|\phi_2^W(p_1(\hat{p},\alpha))|\right] - \frac{1}{2}\mathbb{E}\left[|\phi_2^W(p_1(\hat{p},\beta))|\right] \\
&= \left(q_B^\beta - q_B^\alpha\right)(1-\hat{p}) \cdot \left(|q_A^\alpha - q_B^\alpha| - |q_A^\beta - q_B^\beta|\right).
\end{aligned}$$

It is now easy to see that $\phi_1^W(\hat{p}) = 0$ if and only if $\kappa_A = \kappa_B$. The result follows from Lemma 5 since $\phi^W(\hat{p}) = 0$ if and only if $\hat{p}_0 = \hat{p}$. ■

Proof of Lemma 6. First note that $x_1^\alpha < x_2^\beta$ if and only if $x_2^\alpha < x_1^\beta$ if and only if $\kappa < 1$, and that $x_2^\beta < x_2^\alpha$ if and only if $(1-q_A^\alpha)(1-q_B^\beta) < (1-q_B^\alpha)(1-q_A^\beta)$, where the latter inequality holds by the assumption that action α is preferred in state A and action β in state B. The respective ordering over the $x_1^\alpha, x_2^\alpha, x_1^\beta, x_2^\beta$ for each of the two cases, $\kappa < 1$ and $\kappa > 1$ follows.

Decomposing $\phi_{N-1}^W(p)$ using Lemma 2, we have

$$\begin{aligned}
&\phi_{N-1}^W(p) \\
&= \mathbb{E}[\max\{\phi_N^W(p_{N-1}(p,\alpha)),0\}] - \mathbb{E}[\max\{-\phi_N^W(p_{N-1}(p,\beta)),0\}] \\
&= \frac{1}{2}\mathbb{E}\left[\sum_{\sigma=\alpha,\beta}\phi_N^W(p_{N-1}(p,\sigma)) + |\phi_N^W(p_{N-1}(p,\alpha))| - |\phi_N^W(p_{N-1}(p,\beta))|\right] \\
&= \hat{\phi}(p) + \frac{q_B^\beta - q_B^\alpha}{2\hat{p}}\Big[|p\left(q_A^\alpha(1-\hat{p}) + \hat{p}q_B^\alpha\right) - \hat{p}q_B^\alpha| \qquad (12) \\
&\quad + |p\left((1-q_A^\alpha)(1-\hat{p}) + \hat{p}(1-q_B^\alpha)\right) - \hat{p}(1-q_B^\alpha)| - |p(q_A^\beta(1-\hat{p}) + \hat{p}q_B^\beta) - \hat{p}q_B^\beta| \\
&\quad - |p((1-q_A^\beta)(1-\hat{p}) + \hat{p}(1-q_B^\beta)) - \hat{p}(1-q_B^\beta)|\Big].
\end{aligned}$$

Writing out (12) for each of the cases $\kappa < 1$ and $\kappa > 1$, and each of the respective regions of beliefs yields the result. ■

Proof of Theorem 2. We derive $\phi_{N-2}^W(\hat{p})$ by rewriting it as follows:

$$\begin{aligned}
&\phi_{N-2}^W(\hat{p}) \\
&= \mathbb{E}[\max\{\phi_{N-1}^W(p_{N-2}(\hat{p},\alpha)),0\}] - \mathbb{E}[\max\{-\phi_{N-1}^W(p_{N-2}(\hat{p},\beta)),0\}] \qquad (13)
\end{aligned}$$

$$= \frac{1}{2}\mathbb{E}\left[\sum_{\sigma\in\{\alpha,\beta\}} \phi^W_{N-1}\left(p_{N-2}(\hat{p},\sigma)\right) + |\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)| - |\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)|\right].$$

Suppose the belief of agent $N-2$ is $\hat{p}$. The relevant posterior beliefs that can arise after agent $N-2$ moves are:

$$p_{N-2}(p_{N-3},\sigma,1) = \frac{\hat{p}q^\sigma_A}{\hat{p}q^\sigma_A + (1-\hat{p})q^\sigma_B}\ ,\ p_{N-2}(p_{N-3},\sigma,0) = \frac{\hat{p}(1-q^\sigma_A)}{\hat{p}(1-q^\sigma_A) + (1-\hat{p})(1-q^\sigma_B)}. \tag{14}$$

To derive $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\sigma)\right)$, it must first be determined in which regions among the beliefs $x^\alpha_1, x^\alpha_2, x^\beta_1, x^\beta_2$ these posteriors fall.

Note that since we are assuming $q^\alpha_\omega + q^\beta_\omega = \kappa$, it must be the case that $q^\alpha_A > q^\alpha_B$ and $q^\beta_B > q^\beta_A$ (otherwise, $q^\beta_B > q^\alpha_B > q^\alpha_A > q^\beta_A$ or $q^\alpha_A > q^\beta_A > q^\beta_B > q^\alpha_B$, respectively, contradicting $q^\alpha_\omega + q^\beta_\omega = \kappa$).

Assume $\kappa < 1$ and recall that in this case $x^\alpha_1 < x^\beta_2 < x^\alpha_2 < x^\beta_1$.

Consider the posterior $p_{N-2}(p_{N-3},\alpha,1)$. It can be verified that $p_{N-2}(p_{N-3},\alpha,1) > x^\beta_2$ if and only if $q^\alpha_A(1-q^\beta_A) > (1-q^\beta_B)q^\alpha_B$, which must hold since $q^\alpha_A > q^\alpha_B$ and $q^\beta_B > q^\beta_A$. Hence, $p_{N-2}(p_{N-3},\alpha,1) > x^\beta_2$. Next, it can easily be verified that $p_{N-2}(p_{N-3},\alpha,1) < x^\alpha_2$ if and only if $q^\alpha_A(1-q^\alpha_A) < q^\alpha_B(1-q^\alpha_B)$, which can be rearranged as $q^\alpha_A - q^\alpha_B < (q^\alpha_A - q^\alpha_B)(q^\alpha_A + q^\alpha_B)$. The latter is equivalent to $1 < q^\alpha_A + q^\alpha_B$ since $q^\alpha_A > q^\alpha_B$. Moreover, $p_{N-2}(p_{N-3},\alpha,1) < x^\beta_1$ if and only if $q^\alpha_A q^\beta_A < q^\beta_B q^\alpha_B$, which is equivalent to $q^\alpha_A(\kappa - q^\alpha_A) < q^\alpha_B(\kappa - q^\alpha_B)$, and in turn, to $q^\alpha_A + q^\alpha_B > \kappa$.

Hence, we have shown that if $q^\alpha_A + q^\alpha_B > 1$ then $x^\beta_2 < p_{N-2}(p_{N-3},\alpha,1) < x^\alpha_2$; if $\kappa < q^\alpha_A + q^\alpha_B < 1$ then $x^\alpha_2 < p_{N-2}(p_{N-3},\alpha,1) < x^\beta_1$; and if $q^\alpha_A + q^\alpha_B < \kappa$ then $p_{N-2}(p_{N-3},\alpha,1) > x^\beta_1$. This implies that:

$$\begin{aligned}
&\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,1)\right)\\
&= \begin{cases}
\kappa\hat{\phi}(p_{N-2}(\hat{p},\alpha,1)) & \text{if } q^\alpha_A + q^\alpha_B > 1,\ \kappa < 1\\
\hat{\phi}(p_{N-2}(\hat{p},\alpha,1)) + \frac{q^\beta_B - q^\alpha_B}{\hat{p}}\Big[p_{N-2}(\hat{p},\alpha,1)q^\beta_A(1-\hat{p}) & \\
\qquad -q^\beta_B\hat{p}\left(1-p_{N-2}(\hat{p},\alpha,1)\right)\Big] & \text{if } \kappa < q^\alpha_A + q^\alpha_B < 1\\
\hat{\phi}(p_{N-2}(\hat{p},\alpha,1)) & \text{if } q^\alpha_A + q^\alpha_B < \kappa < 1.
\end{cases}
\end{aligned}$$

Next consider $p_{N-2}(p_{N-3},\alpha,0)$. It can be verified that $p_{N-2}(p_{N-3},\alpha,0) < x^\alpha_1$ if and only if $q^\alpha_A + q^\alpha_B > 1$. Furthermore, $p_{N-2}(p_{N-3},\alpha,0) < x^\beta_2$ if and only if

$(1-q_A^\alpha)(1-q_A^\beta) > (1-q_B^\beta)(1-q_B^\alpha)$, where the latter is equivalent to $q_A^\alpha + q_B^\alpha > \kappa$. Moreover, $p_{N-2}(p_{N-3},\alpha,0) < x_2^\alpha$ if and only if $q_A^\alpha > q_B^\alpha$. Hence, if $q_A^\alpha + q_B^\alpha > 1$ then $p_{N-2}(p_{N-3},\alpha,0) < x_1^\alpha$; if $\kappa < q_A^\alpha + q_B^\alpha < 1$ then $x_1^\alpha < p_{N-2}(p_{N-3},\alpha,0) < x_2^\beta$; and if $q_A^\alpha + q_B^\alpha < \kappa$ then $x_2^\beta < p_{N-2}(p_{N-3},\alpha,0) < x_2^\alpha$. This implies that:

$$\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,0)\right)$$
$$= \begin{cases} \hat{\phi}(p_{N-2}(\hat{p},\alpha,0)) & \text{if } q_A^\alpha + q_B^\alpha > 1 > \kappa \\ \hat{\phi}(p_{N-2}(\hat{p},\alpha,0)) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\Big[q_A^\alpha(1-\hat{p})p_{N-2}(\hat{p},\alpha,0) & \\ \qquad -\hat{p}q_B^\alpha(1-p_{N-2}(\hat{p},\alpha,0))\Big] & \text{if } \kappa < q_A^\alpha + q_B^\alpha < 1 \\ \kappa\hat{\phi}(p_{N-2}(\hat{p},\alpha,0)) & \text{if } q_A^\alpha + q_B^\alpha < \kappa < 1. \end{cases}$$

Next consider $p_{N-2}(p_{N-3},\beta,1)$. Note that $p_{N-2}(p_{N-3},\beta,1) < x_1^\alpha$ if and only if $q_A^\alpha q_A^\beta < q_B^\alpha q_B^\beta$, which is equivalent to $q_A^\alpha(\kappa - q_A^\alpha) < q_B^\alpha(\kappa - q_B^\alpha)$, and in turn, to $q_A^\alpha + q_B^\alpha > \kappa$. Furthermore, $p_{N-2}(p_{N-3},\beta,1) < x_2^\beta$ if and only if $q_B^\beta + q_A^\beta < 1$, which is equivalent to $q_B^\alpha + q_A^\alpha > 2\kappa - 1$. We therefore need to distinguish between the case of $\kappa < \frac{1}{2}$ and $\kappa > \frac{1}{2}$.

If $\kappa < \frac{1}{2}$ then $p_{N-2}(p_{N-3},\beta,1) < x_2^\beta$, which means that in this case, if $q_A^\alpha + q_B^\alpha > \kappa$ then $p_{N-2}(p_{N-3},\beta,1) < x_1^\alpha$, and if $q_A^\alpha + q_B^\alpha < \kappa$ then $x_1^\alpha < p_{N-2}(p_{N-3},\beta,1) < x_2^\beta$. Hence, if $\kappa < \frac{1}{2}$ then

$$\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,1)\right)$$
$$= \begin{cases} \hat{\phi}(p_{N-2}(\hat{p},\beta,1)) & \text{if } \kappa < q_A^\alpha + q_B^\alpha,\ \kappa < \frac{1}{2} \\ \hat{\phi}(p_{N-2}(\hat{p},\beta,1)) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\Big[q_A^\alpha(1-\hat{p})p_{N-2}(\hat{p},\beta,1) & \\ \qquad -\hat{p}q_B^\alpha(1-p_{N-2}(\hat{p},\beta,1))\Big] & \text{if } q_A^\alpha + q_B^\alpha < \kappa < \frac{1}{2}. \end{cases}$$

If $\kappa > \frac{1}{2}$ then, if $2\kappa - 1 < q_A^\alpha + q_B^\alpha < \kappa$ then $x_1^\alpha < p_{N-2}(p_{N-3},\beta,1) < x_2^\beta$. Moreover, $p_{N-2}(p_{N-3},\beta,1) < x_2^\alpha$ if and only if $(1-q_A^\alpha)q_A^\beta < q_B^\beta(1-q_B^\alpha)$, which always holds. Hence, we have

$$\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,1)\right)$$

$$
= \begin{cases} \hat{\phi}(p_{N-2}(\hat{p},\beta,1)) & \text{if } \frac{1}{2}<\kappa<q_A^\alpha+q_B^\alpha \\ \hat{\phi}(p_{N-2}(\hat{p},\beta,1)) + \frac{q_B^\beta-q_B^\alpha}{\hat{p}}\Big[q_A^\alpha(1-\hat{p})p_{N-2}(\hat{p},\beta,1) & \\ \qquad -\hat{p}q_B^\alpha(1-p_{N-2}(\hat{p},\beta,1))\Big] & \text{if } 2\kappa-1<q_A^\alpha+q_B^\alpha<\kappa,\ \frac{1}{2}<\kappa<1 \\ \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,1)) & \text{if } q_A^\alpha+q_B^\alpha<2\kappa-1, \frac{1}{2}<\kappa<1. \end{cases}
$$

Now consider $p_{N-2}(p_{N-3},\beta,0)$. First, $p_{N-2}(p_{N-3},\beta,0) > x_1^\alpha$ is equivalent to $(1-q_A^\beta)q_A^\alpha > q_B^\alpha(1-q_B^\beta)$, which always holds. Similarly, $p_{N-2}(p_{N-3},\beta,0) > x_2^\beta$ always holds. Next, $p_{N-2}(p_{N-3},\beta,0) < x_2^\alpha$ if and only if $(1-q_A^\beta)(1-q_A^\alpha) > (1-q_B^\alpha)(1-q_B^\beta)$, which can be shown to be equivalent to $\kappa < q_A^\alpha + q_B^\alpha$. Furthermore, $p_{N-2}(p_{N-3},\beta,0) > x_1^\beta$ if and only if $q_A^\beta(1-q_A^\beta) > q_B^\beta(1-q_B^\beta)$, which is equivalent to $q_B^\beta + q_A^\beta > 1$, and in turn, to $q_B^\alpha + q_A^\alpha < 2\kappa - 1$. We again distinguish between the two cases of $\kappa < \frac{1}{2}$ and $\kappa > \frac{1}{2}$.

If $\kappa < \frac{1}{2}$ then $q_B^\alpha + q_A^\alpha > 2\kappa - 1$ always and hence $p_{N-2}(p_{N-3},\beta,0) < x_1^\beta$. Therefore,

$$
\begin{aligned}
&\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,0)\right) \\
&\quad = \begin{cases} \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,0)) & \text{if } \kappa<q_A^\alpha+q_B^\alpha,\ \kappa<\frac{1}{2} \\ \hat{\phi}(p_{N-2}(\hat{p},\beta,0)) + \frac{q_B^\beta-q_B^\alpha}{\hat{p}}\Big[p_{N-2}(\hat{p},\beta,0)q_A^\beta(1-\hat{p}) & \\ \qquad -q_B^\beta\hat{p}(1-p_{N-2}(\hat{p},\beta,0))\Big] & \text{if } q_A^\alpha+q_B^\alpha<\kappa<\frac{1}{2}. \end{cases}
\end{aligned}
$$

If $\kappa > \frac{1}{2}$ then

$$
\begin{aligned}
&\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,0)\right) \\
&\quad = \begin{cases} \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,0)) & \text{if } \frac{1}{2}<\kappa<q_A^\alpha+q_B^\alpha \\ \hat{\phi}(p_{N-2}(\hat{p},\beta,0)) + \frac{q_B^\beta-q_B^\alpha}{\hat{p}}\Big[p_{N-2}(\hat{p},\beta,0)q_A^\beta(1-\hat{p}) & \\ \qquad -q_B^\beta\hat{p}(1-p_{N-2}(\hat{p},\beta,0))\Big] & \text{if } 2\kappa-1<q_A^\alpha+q_B^\alpha<\kappa,\ \kappa>\frac{1}{2} \\ \hat{\phi}(p_{N-2}(\hat{p},\beta,0)) & \text{if } q_A^\alpha+q_B^\alpha<2\kappa-1,\ \kappa>\frac{1}{2}. \end{cases}
\end{aligned}
$$

We can now calculate the expectations in (13). We distinguish between the different cases considered above. Suppose first that $q_A^\alpha + q_B^\alpha > 1$. Then the following must hold: $\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,1)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\alpha,1))$, $\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,0)\right) = \hat{\phi}(p_{N-2}(\hat{p},\alpha,0))$, $\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,1)\right) = \hat{\phi}(p_{N-2}(\hat{p},\beta,1))$, and $\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta,0)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,0))$. Replacing $p_{N-2}(\hat{p},\sigma,1)$ and $p_{N-2}(\hat{p},\sigma,0)$ using (14), and recalling that $\hat{\phi}(p) = (q_B^\beta - q_B^\alpha)(p - \hat{p})/\hat{p}$, it can be shown that

$$\begin{aligned}\mathbb{E}\left[\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha)\right)\right] &= \left(\hat{p}q_A^\alpha + (1-\hat{p})q_B^\alpha\right)\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,1)\right)\\ &\quad + \left(\hat{p}(1-q_A^\alpha) + (1-\hat{p})(1-q_B^\alpha)\right)\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,0)\right)\\ &= (\kappa-1)(q_B^\beta - q_B^\alpha)(1-\hat{p})(q_A^\alpha - q_B^\alpha),\end{aligned}$$

and similarly

$$\mathbb{E}\left[\left|\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha)\right|\right)\right] = (\kappa+1)(q_B^\beta - q_B^\alpha)(1-\hat{p})(q_A^\alpha - q_B^\alpha).$$

Furthermore,

$$\mathbb{E}\left[\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta)\right)\right] = (q_B^\beta - q_B^\alpha)(1-\hat{p})(q_A^\alpha - q_B^\alpha)(\kappa-1)$$

and

$$\mathbb{E}\left[\left|\phi_{N-1}^W\left(p_{N-2}(\hat{p},\beta)\right)\right|\right] = (q_B^\beta - q_B^\alpha)(1-\hat{p})(q_A^\alpha - q_B^\alpha)(\kappa+1).$$

Hence, from (13), when $\kappa < 1$ and $q_A^\alpha + q_B^\alpha > 1$,

$$\phi_{N-2}^W(\hat{p}) = (q_B^\beta - q_B^\alpha)(1-\hat{p})(q_A^\alpha - q_B^\alpha)(\kappa-1).$$

Next, consider the case $\kappa < q_A^\alpha + q_B^\alpha < 1$. Then

$$\begin{aligned}&\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,1)\right)\\ &= \hat{\phi}(p_{N-2}(\hat{p},\alpha,1)) + \frac{q_B^\beta - q_B^\alpha}{\hat{p}}\left[p_{N-2}(\hat{p},\alpha,1)q_A^\beta(1-\hat{p}) - q_B^\beta\hat{p}\left(1 - p_{N-2}(\hat{p},\alpha,1)\right)\right.\\ &= \frac{(q_B^\beta - q_B^\alpha)}{\hat{p}q_A^\alpha + (1-\hat{p})q_B^\alpha}\left[q_A^\alpha + q_A^\alpha q_A^\beta(1-\hat{p}) + \hat{p}q_A^\alpha q_B^\beta - (1+q_B^\beta)\left(\hat{p}q_A^\alpha + (1-\hat{p})q_B^\alpha\right)\right]\end{aligned}$$

and

$$\begin{aligned}&\phi_{N-1}^W\left(p_{N-2}(\hat{p},\alpha,0)\right)\\ &= (q_B^\beta - q_B^\alpha)\left[\frac{(1-q_A^\alpha) + q_A^\alpha(1-\hat{p})(1-q_A^\alpha) + q_B^\alpha\hat{p}(1-q_A^\alpha)}{\hat{p}(1-q_A^\alpha) + (1-\hat{p})(1-q_B^\alpha)} - 1 - q_B^\alpha\right],\end{aligned}$$

and, furthermore, as in the previous case, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,1)\right) = \hat{\phi}(p_{N-2}(\hat{p},\beta,1))$ and $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,0)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,0))$. In particular, as in the previous case,

$$\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(\kappa - 1)$$

and

$$\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right|\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(\kappa + 1).$$

Turning to $\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right]$ and $\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right|\right]$, it can easily be verified that

$$\begin{aligned}\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right] &= \left(\hat{p}q^\alpha_A + (1-\hat{p})q^\alpha_B\right)\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,1)\right)\\ &\quad + \left(\hat{p}(1-q^\alpha_A) + (1-\hat{p})(1-q^\alpha_B)\right)\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,0)\right)\\ &= (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[\kappa + 1 - 2(q^\alpha_A + q^\alpha_B)\right],\end{aligned}$$

and

$$\begin{aligned}\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right|\right] &= \left(\hat{p}q^\alpha_A + (1-\hat{p})q^\alpha_B\right)\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,1)\right)\right|\\ &\quad + \left(\hat{p}(1-q^\alpha_A) + (1-\hat{p})(1-q^\alpha_B)\right)\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,0)\right)\right|\\ &= (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1+\kappa).\end{aligned}$$

We can now derive $\phi^W_{N-2}(\hat{p})$ for this case:

$$\phi^W_{N-2}(\hat{p}) = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[\kappa - (q^\alpha_A + q^\alpha_B)\right].$$

Now consider the case $2\kappa - 1 < q^\alpha_A + q^\alpha_B < \kappa$. In this case, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,1)\right) = \hat{\phi}(p_{N-2}(\hat{p},\alpha,1))$, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,0)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\alpha,0))$, and we can now calculate $\phi^W_{N-1}(p_{N-2}(\hat{p},\beta,1))$ and $\phi^W_{N-1}(p_{N-2}(\hat{p},\beta,0))$ to obtain

$$\begin{aligned}&\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,1)\right)\\ &= \hat{\phi}(p_{N-2}(\hat{p},\beta,1)) + \frac{q^\beta_B - q^\alpha_B}{\hat{p}}\left[q^\alpha_A(1-\hat{p})p_{N-2}(\hat{p},\beta,1) - \hat{p}q^\alpha_B(1-p_{N-2}(\hat{p},\beta,1))\right]\\ &= \frac{q^\beta_B - q^\alpha_B}{\hat{p}q^\beta_A + (1-\hat{p})q^\beta_B}(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[\kappa - 1 - (q^\alpha_A + q^\alpha_B)\right]\end{aligned}$$

and

$$\begin{aligned}
&\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,0)\right) \\
&= \hat{\phi}(p_{N-2}(\hat{p},\beta,0)) + \frac{q^\beta_B - q^\alpha_B}{\hat{p}}\Big[p_{N-2}(\hat{p},\beta,0)q^\beta_A(1-\hat{p}) - q^\beta_B\hat{p}(1-p_{N-2}(\hat{p},\beta,0))\Big] \\
&= \frac{q^\beta_B - q^\alpha_B}{\hat{p}(1-q^\beta_A) + (1-\hat{p})(1-q^\beta_B)}(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[2\kappa - (q^\alpha_B + q^\alpha_A)\right].
\end{aligned}$$

Calculating the expectations $\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\sigma)\right)\right]$, it can be shown that

$$\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1-\kappa)$$

and

$$\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right|\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1+\kappa),$$

and, similarly,

$$\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[3\kappa - 1 - 2(q^\alpha_A + q^\alpha_B)\right]$$

and

$$\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right|\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(\kappa+1).$$

We can now derive $\phi^W_{N-2}(\hat{p})$ for this case:

$$\phi^W_{N-2}(\hat{p}) = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)\left[\kappa - (q^\alpha_A + q^\alpha_B)\right].$$

Note that this is the same as in the case $\kappa < q^\alpha_A + q^\alpha_B < 1$.

Finally, consider the case $q^\alpha_A + q^\alpha_B < 2\kappa - 1, \kappa > \frac{1}{2}$. Here, as in the previous case, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,1)\right) = \hat{\phi}(p_{N-2}(\hat{p},\alpha,1))$, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha,0)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\alpha,0))$, and furthermore

$$\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1-\kappa)$$

and

$$\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\alpha)\right)\right|\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1+\kappa).$$

Moreover, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,1)\right) = \kappa\hat{\phi}(p_{N-2}(\hat{p},\beta,1))$ and, furthermore, $\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta,0)\right) = \hat{\phi}(p_{N-2}(\hat{p},\beta,0))$. It can then be shown that

$$\mathbb{E}\left[\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1-\kappa)$$

and similarly,

$$\mathbb{E}\left[\left|\phi^W_{N-1}\left(p_{N-2}(\hat{p},\beta)\right)\right|\right] = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1+\kappa).$$

We can now derive $\phi^W_{N-2}(\hat{p})$ for this case:

$$\phi^W_{N-2}(\hat{p}) = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B)(1-\kappa),$$

We therefore conclude that, for the case of $\kappa < 1$,

$$\phi^W_{N-2}(\hat{p}) = (q^\beta_B - q^\alpha_B)(1-\hat{p})(q^\alpha_A - q^\alpha_B) \times \begin{cases} \kappa - 1 & \text{if } q^\alpha_A + q^\alpha_B > 1 \\ \kappa - (q^\alpha_A + q^\alpha_B) & \text{if } 2\kappa - 1 < q^\alpha_A + q^\alpha_B < 1 \\ 1 - \kappa & \text{if } q^\alpha_A + q^\alpha_B < 2\kappa - 1 \text{ and } \kappa > \frac{1}{2}. \end{cases}$$

In particular, when $\kappa < 1$, it is only possible that $\phi^W_{N-2}(\hat{p}) = 0$ when $\kappa - (q^\alpha_A + q^\alpha_B) = 0$. But since by definition $\kappa = q^\alpha_A + q^\beta_A$, this means $q^\beta_A = q^\alpha_B$ and $q^\alpha_A = q^\beta_B$. In other words, it is only possible that $\phi^W_{N-2}(\hat{p}) = 0$ under a symmetric information structure.

It can be verified that the same property holds for $\kappa > 1$ (that case is analogous and hence omitted). Therefore, it is only possible that $\phi^W_{N-2}(\hat{p}) = 0$ when $\kappa = (q^\alpha_A + q^\alpha_B)$. The result then follows from Lemma 5. ■

References

Banerjee, Abhijit V. 1992. "A simple model of herd behavior." *The Quarterly Journal of Economics*, 107(3): 797–817.

Bergemann, Dirk, and Ulrich Hege. 1998. "Venture capital financing, moral hazard, and learning." *Journal of Banking & Finance*, 22(6-8): 703–735.

Bergemann, Dirk, and Ulrich Hege. 2005. "The financing of innovation: Learning and stopping." *RAND Journal of Economics*, 719–752.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A theory of fads, fashion, custom, and cultural change as informational cascades." *Journal of political Economy*, 100(5): 992–1026.

Bikhchandani, Sushil, David Hirshleifer, Omer Tamuz, and Ivo Welch. 2024. "Information cascades and social learning." *Journal of Economic Literature*, 62(3): 1040–1093.

Binns, Reuben. 2018. "Fairness in machine learning: Lessons from political philosophy." 149–159, PMLR.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender systems as mechanisms for social learning." *The Quarterly Journal of Economics*, 133(2): 871–925.

Cowgill, Bo, Fabrizio Dell'Acqua, Samuel Deng, Daniel Hsu, Nakul Verma, and Augustin Chaintreau. 2020. "Biased programmers? Or biased data? A field experiment in operationalizing AI ethics." 679–681.

Fageria, Nand Kumar. 2016. *The use of nutrients in crop plants.* CRC press.

Floridi, Luciano, Josh Cowls, Monica Beltrametti, Raja Chatila, Patrice Chazerand, Virginia Dignum, Christoph Luetge, Robert Madelin, Ugo Pagallo, Francesca Rossi, et al. 2018. "AI4People—an ethical framework for a good AI society: opportunities, risks, principles, and recommendations." *Minds and machines*, 28: 689–707.

Glazer, Jacob, Ilan Kremer, and Motty Perry. 2021. "The wisdom of the crowd when acquiring information is costly." *Management Science*, 67(10): 6443–6456.

Guo, Yingni. 2016. "Dynamic delegation of experimentation." *American Economic Review*, 106(8): 1969–2008.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. "Optimal contracts for experimentation." *The Review of Economic Studies*, 83(3): 1040–1091.

Hörner, Johannes, and Larry Samuelson. 2013. "Incentives for experimenting agents." *The RAND Journal of Economics*, 44(4): 632–663.

Kremer, Ilan, Yishay Mansour, and Motty Perry. 2014. "Implementing the "wisdom of the crowd"." *Journal of Political Economy*, 122(5): 988–1012.

Ladha, Jagdish K, Himanshu Pathak, Timothy J Krupnik, Johan Six, and Chris van Kessel. 2005. "Efficiency of fertilizer nitrogen in cereal production: retrospects and prospects." *Advances in agronomy*, 87: 85–156.

Luca, Michael. 2014. "Were OkCupid's and Facebook's experiments unethical." *Harvard Business Review Blog Network. http://blogs. hbr. org/2014/07/were-okcupids-and-facebooks-experiments-unethical (visited October 19, 2014).*

Mega, Jessica L, Sandra L Close, Stephen D Wiviott, Lei Shen, Richard D Hockett, John T Brandt, Joseph R Walker, Elliott M Antman, William Macias, Eugene Braunwald, et al. 2009. "Cytochrome p-450 polymorphisms and response to clopidogrel." *New England journal of medicine*, 360(4): 354–362.

Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wachter, and Luciano Floridi. 2016. "The ethics of algorithms: Mapping the debate." *Big Data & Society*, 3(2): 2053951716679679.

Mok, Tony S, Yi-Long Wu, Myung-Ju Ahn, Marina C Garassino, Hye R Kim, Suresh S Ramalingam, Frances A Shepherd, Yong He, Hiroaki Akamatsu, Willemijn SME Theelen, et al. 2017. "Osimertinib or platinum–pemetrexed in EGFR T790M–positive lung cancer." *New England Journal of Medicine*, 376(7): 629–640.

Smith, Lones, and Peter Sørensen. 2000. "Pathological outcomes of observational learning." *Econometrica*, 68(2): 371–398.

Welch, Ivo. 1992. "Sequential sales, learning, and cascades." *The Journal of Finance*, 47(2): 695–732.

Wolitzky, Alexander. 2018. "Learning from Others' Outcomes." *American Economic Review*, 108(10): 2763–2801.