The Recommendation Principle in Information Design

Bravo Working Paper # 2025-002

Zeky Murra-Anton* Giacomo Rubbini† Roberto Serrano‡

First version: February 2025; This Version: September 2025

Click here for latest version

Abstract

The recommendation principle states that, without losing generality, a sender can focus on experiments in which messages are recommendations in the receiver's action space. We study the general validity of the recommendation principle and characterize the conditions under which it holds in persuasive communication settings. Using our characterization, we construct a simple test that asserts the validity of the recommendation principle. We apply our test to examples in the literature where the principle has been assumed or proved for that specific application.

Keywords: Bayesian Persuasion, Information Design, Recommendation Principle, Recommendation Experiments, Behavioral Persuasion

JEL Codes: C72, D82, D83, D91

*Market Development, ISO New England. Contact: zmurraanton@iso-ne.com. This paper has not been funded or endorsed by ISO New England. All the views herein expressed are the authors' only.

†Department of Economics, University of Bologna. Contact: giacomo.rubbini2@unibo.it

‡Department of Economics, Brown University. Contact: roberto serrano@brown.edu

The Recommendation Principle in Information Design

Zeky Murra-Anton* Giacomo Rubbini† Roberto Serrano‡

First version: February 2025; This Version: September 2025

Click here for latest version

Abstract

The recommendation principle states that, without losing generality, a sender can focus on experiments in which messages are recommendations in the receiver's action space. We study the general validity of the recommendation principle and characterize the conditions under which it holds in persuasive communication settings. Using our characterization, we construct a simple test that asserts the validity of the recommendation principle. We apply our test to examples in the literature where the principle has been assumed or proved for that specific application.

Keywords: Bayesian Persuasion, Information Design, Recommendation Principle, Recommendation Experiments, Behavioral Persuasion

JEL Codes: C72, D82, D83, D91

*Market Development, ISO New England. Contact: zmurraanton@iso-ne.com. This paper has not been funded or endorsed by ISO New England. All the views herein expressed are the authors' only.

†Department of Economics, University of Bologna. Contact: giacomo.rubbini2@unibo.it

‡Department of Economics, Brown University. Contact: roberto_serrano@brown.edu

1. Introduction

Information plays a crucial role in countless economic environments: from financial markets to political campaigns, the ability of an informed agent to shape others' beliefs and influence their subsequent actions is a fundamental ingredient of many economic models. Kamenica and Gentzkow (2011) provide a foundational framework for studying this kind of problem, showing that an informed party (sender) can influence the choices of an uninformed opponent (receiver) by committing to an experiment whose results her opponent can observe.

Analogous to the revelation principle in mechanism design, the *recommendation principle* considerably simplifies the search for the optimal experiment by showing that no generality is lost by restricting attention to *direct recommendation experiments*— that is, experiments in which the sender recommends an action to the receiver.[1] The useful *recommendation principle* is critical to ensure the tractability and applicability of the persuasion framework of Kamenica and Gentzkow (2011).

Despite the widespread use of the recommendation principle in the literature, its general validity remains an open question. This paper fills this gap by identifying the key features of the receiver's decision-making environment that ensure it is without loss of generality to focus on direct recommendation experiments, and by providing an intuitive framework to diagnose failures of the principle.

Our model mirrors the standard persuasion framework with one sender and one receiver. The sender (he), before learning about the state of the world, commits to a communication experiment consisting of an arbitrary message space and an arbitrary family of conditional probability distributions over messages. Then, a potentially privately informed receiver (she), who knows the sender's experiment, observes the realization of one of its messages (possibly together with her private payoff type), updates her beliefs about the state of the world, and takes an action.[2]

To discuss the action choice problem of the receiver, we assume that the receiver chooses her action to maximize her *value function*, that is, her indirect utility as a function only of her action, type, posterior and, possibly, the distribution of posteriors generated by the

[1]In particular, the recommendation principle makes it possible to study the sender's payoff maximization problem with standard optimization techniques: rather than maximizing over a (potentially very complex) set of possible experiments, it is enough to focus on the much simpler problem maximizing over the set of all distributions over action recommendations.

[2]In models of Bayesian persuasion, the sender can be considered an informed party, as long as he commits to an experiment before realizing his private information. This assumption, contrasting with the literature on cheap talk, implies that the experiment choice does not contain information about the state of the world (Crawford and Sobel, 1982).

experiment chosen by the sender. We show how several setups and models in the literature can be analyzed based on such a value function.

If the receiver's value function does not depend on the posterior distribution generated by the experiment, we show that the recommendation principle holds if and only if the set of posteriors for which each action is optimal is convex. This is the case, for example, in the standard framework of Kamenica and Gentzkow (2011) in which the value function is linear in the posterior beliefs. If the set of receiver's actions is binary, we show that the recommendation principle holds whenever the value functions associated with the two actions intersect only once.

If the receiver's optimal action also depends on the posterior distribution generated by the experiment, we show that a slightly stronger version of the recommendation principle holds if and only if "pooling" posteriors inducing the same action do not cause the receiver to change her favorite action. We apply this framework to models of persuaded search (Mekonnen et al., 2025), of persuasion of an inattentive receiver of (Dall'Ara, 2024), and non-Bayesian persuasion (de Clippel and Zhang, 2022). Our analysis shows how the recommendation principle extends beyond the standard framework without a significant change in the machinery behind the results. In addition, this analysis sheds light on the aspects of a communication model that would lead to violations of the recommendation principle.

We are not the first to observe that convexity and invariance in the pooling of messages are relevant to establish a recommendation principle.[3] For example, Lipnowski and Mathevet (2018) notice that the recommendation principle may fail (in a specific sense) whenever the set of posteriors for which an action is optimal is not convex. Similarly, Doval and Skreta (2024) proves a recommendation principle for a class of persuasion problems with constraints if the sender cannot increase its payoff by pooling different messages. Our paper complements these analyses by providing a converse to their results while generalizing the applicability of the recommendation principle to a much broader class of persuasion problems, including stationary dynamic problems (Mekonnen et al., 2025), non-Bayesian updating (de Clippel and Zhang, 2022), and bounded rationality (Dall'Ara, 2024). This is possible because our results do not rely in any way on the specifics of the payoff function of the sender, focusing instead on the decision-making environment of the receiver only. The current paper also relates to the recent work on convex choice settings, for example, in which the set of types for which a given action is optimal is convex (Kartik et al., 2023; Kartik and Kleiner, 2025).

The remainder of the paper proceeds as follows. Section 2 presents the model that

[3]More broadly, our findings also relate to efforts trying to characterize the class of mechanism design problems in which the revelation principle holds (Saran, 2011; Rubbini, 2024b; Xiong, 2024) and in which incentive compatibility is a necessary condition for implementation Rubbini (2024a).

defines the general framework for information design and persuasion problems. Section 3 introduces the conditions under which the recommendation principle holds in problems where the receiver's utility from taking a given action depends on her posterior but not on the distribution of posteriors generated by the experiment chosen by the sender. Section 4 extends the model by removing this limitation. Finally, Section 5 discusses possible generalizations of our model. Proofs of results and one such generalization are relegated to Appendix B, and one additional generalization is the content of an Online Appendix.

2. Model

2.1. Setup

Denote by Ω a set of states of the world, with typical element ω, and a set of payoff types Θ for the receiver, with typical element θ.

We analyze the following game. First, without observing the state of the world, the sender commits to an experiment $\pi = (S, \sigma)$, consisting of a realization space S and a family of distributions $\sigma = \{\sigma(\cdot \mid \omega)\}_{\omega \in \Omega}$ over S, where $\sigma(s \mid \omega)$ represents the probability that the message s is sent in the state ω. For simplicity of exposition, we assume that the spaces S and Ω are finite.[4]

The receiver then observes the experiment π, a realization s, and a realization θ of her private type θ, and updates her prior belief $\mu_0 \in \Delta(\Omega)$ to posterior $\mu^\pi(s) \in \Delta(\Omega)$. As we interpret θ as the receiver's payoff type, we make the standard assumption that the receiver's prior does not depend on her payoff type θ.[5] Then the receiver chooses an action from the finite action space A. We denote a typical action by $a \in A$. As the receiver's actions may be different depending on her payoff type, we define an action plan as the function $\bar{a} : \Theta \to A$, where $\bar{a}(\theta)$ is the action that the receiver takes according to plan $\bar{a}$. We denote by A^Θ the set of all possible action plans. Finally, since the posterior beliefs of the receiver depend on the realization s, we denote by G^π the distribution of posteriors induced by experiment π.

The receiver chooses an action to maximize her *value function*, which is defined as $v : B \times A \times \Theta \to \mathbb{R}$ and $B \subseteq \Delta(\Omega)$ is a convex set of posteriors. We interpret v as a reduced form expression capturing all other factors that could affect the receiver's utility, such as other players' actions.[6]

[4]None of our results crucially hinge on these finiteness assumptions except for Theorem 1: we discuss the impact of these finiteness assumptions and how to generalize the model in paragraph 5.

[5]In the online appendix, we discuss that more general case, which requires significantly more cumbersome notation, offering similar results.

[6]Although this paper focuses on value functions for a simple application to existing models, our results do not conceptually hinge on the receiver's choices being represented by a real-valued function,

2.2. Implementation and the Recommendation Principle

We set a *receiver's environment* $\mathcal{E}$ as a tuple $\{\Omega, B, \Theta, A, v\}$. We formalize the receiver's behavior induced by an experiment.

Definition 1 (Implementation of Action Distributions)**.** *Experiment* π *implements an action distribution* $d^\pi : \Omega \times \Theta \to \Delta(A)$ *whenever:*

1. $\mu^0 = \sum_{s \in S} \sigma^\pi(s)\mu^\pi(s)$,
2. *there exists* $\alpha^\pi : S \times \Theta \to A$ *such that:*

   (a) *for all* $\theta \in \Theta$, $s \in supp(\sigma^\pi)$, *and* $a' \in A$:

$$v(\mu^\pi(s), \alpha^\pi(s, \theta), \theta) \geq v(\mu^\pi(s), a', \theta)$$

   (b) *for all* $\theta \in \Theta$, *we have:*

$$d^\pi(\omega, \theta)[a] = \sum_{s \in S} \mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}}(s)\sigma(s|\omega).$$

Requirement 1 is the standard *feasibility* or Bayesian plausibility requirement that the average posterior is equal to the prior.[7] Meanwhile, 2.a represents the *obedience* requirement that the receiver's action is optimal given her beliefs. Finally, 2.b summarizes the fact that the probability that the action a is played in state ω by type θ is given by the aggregate probability of the signal realizations that, conditional on ω, induce the action a in a receiver of type θ. We note that to implement d^π, we require that there exists *at least one* strategy α^π inducing d^π.

This assumption is in line with the model of Kamenica and Gentzkow (2011), who require that there exists at least one equilibrium of the sender-receiver game inducing the distribution over actions the sender seeks to implement.[8]

We say that *the recommendation principle holds for the receiver environment* $\mathcal{E}$ whenever, for any possible prior, any action distribution the sender can implement can also be implemented by recommending an action plan to the receiver. Formally:

---

and could be generalized by assuming the receiver's choices can be represented by a correspondence mapping $B \times \Theta$ to non-empty subsets of A.

[7]For ease of exposition, most of the discussion focuses on a Bayesian receiver. In Section 4.1, we show that our results extend naturally to a class of non-Bayesian updating rules as well.

[8]In case of multiple equilibria, Kamenica and Gentzkow (2011) focus attention on the equilibrium yielding the highest payoff to the sender. Without this assumption, restricting attention to direct experiments may lead to loss of generality: we refer the reader to the discussion in Bergemann and Morris (2019) and references therein.

Definition 2. *The recommendation principle holds for the receiver environment $\mathcal{E}$ if, for any $\mu_0 \in B$, any action distribution the sender can implement through experiment π can be implemented via the direct recommendation experiment $\pi^* = (A^\Theta, \sigma^{\pi^*})$ and strategy $\alpha^{\pi^*}(\bar{a}, \theta) = \bar{a}(\theta)$ for all $\theta \in \Theta$.*

2.3. Posterior Convexity

The key property of the value function that characterizes the set of persuasion problems that can be simplified using the recommendation principle is that the set of posteriors that support each action plan as optimal is convex. That is, a property of optimal-action posterior convexity, or posterior convexity, for short:

Definition 3 (Posterior Convexity (PC))**.** *A value function v satisfies posterior convexity whenever, if $\mu, \mu' \in B$, $\bar{a} \in A^\Theta$ and for all $a' \in A$, having that*

$$v(\mu, \bar{a}(\theta), \theta) \geq v(\mu, a', \theta) \text{ and } v(\mu', \bar{a}(\theta), \theta) \geq v(\mu', a', \theta) \text{ for all } \theta \in \Theta,$$

then for all $\lambda \in [0, 1]$, One also has that

$$v(\lambda\mu + (1-\lambda)\mu', \bar{a}(\theta), \theta) \geq v(\lambda\mu + (1-\lambda)\mu', a', \theta) \text{ for all } \theta \in \Theta.$$

In particular, PC holds whenever the value function is linear in the posterior. This is the case, for instance, in the standard setup of Kamenica and Gentzkow (2011) or in the special case of a non-Bayesian receiver who systematically distorts Bayesian posteriors (de Clippel and Zhang, 2022).

In some cases, it is easier to check the following (stronger) condition:

Definition 4 (Strong Posterior Convexity (SPC))**.** *A value function v satisfies Strong Posterior Convexity whenever, if $a, a' \in A$ and $\mu, \mu' \in B$ are such that*

$$v(\mu, a, \theta) \geq v(\mu, a', \theta) \text{ and } v(\mu', a, \theta) \geq v(\mu', a', \theta)$$

then for all $\lambda \in [0, 1]$,

$$v(\lambda\mu + (1-\lambda)\mu', a, \theta) \geq v(\lambda\mu + (1-\lambda)\mu', a', \theta)$$

In fact, the PC and SPC coincide whenever the set of possible receiver types is a singleton. More generally, it is easy to see that SPC implies PC. Suppose that there exist μ and μ' such that $\bar{a}(\theta)$ is optimal for the type θ at the posteriors μ and μ'. Although SPC implies that $\bar{a}(\theta)$ is also optimal for the type θ in any linear combination of μ and μ', PC implies the same only if $\bar{a}(\theta')$ is optimal for θ' in any posteriors μ and μ', *for all* $\theta' \neq \theta$.

To illustrate PC, consider a simple set-up with two states ($\Omega = \{\omega, \omega'\}$) and two actions ($A = \{a, a'\}$) in which the planner is interested. As the state space is binary, the receiver's beliefs are captured by the probability $p \in [0, 1]$ that the state is ω. As we do not discuss types at this moment, let us omit θ as an argument for v.

For a simple example of a PC value function, we can turn to the standard Bayesian Persuasion model in Kamenica and Gentzkow (2011). The value function for this problem can be written as:

$$v(p, a) = pu(a, \omega) + (1 - p)u(a, \omega'),$$

where p is the probability of state ω. It is immediate to notice that PC holds: for all $p, p' \in [0, 1]$ with $v(p, a) \geq v(p, a')$ and $v(p', a) \geq v(p', a')$, the fact that v is affine in the posterior entails that we have $v(\lambda p + (1 - \lambda)p', a) \geq v(\lambda p + (1 - \lambda)p', a')$ for all $\lambda \in [0, 1]$.

It is also instructive to inspect the plot of the two value functions. Figure 1 shows graphically how the set of posteriors such that action a is optimal is convex, as is the set of posteriors at which action a' is optimal. It is sometimes useful to instead consider the function $\psi(p) = v(p, a) - v(p, a')$: a will be preferred to a' at all posteriors p such that $\psi(p) \geq 0$, while a' will be preferred to a for all posteriors p such that $\psi(p) \geq 0$.

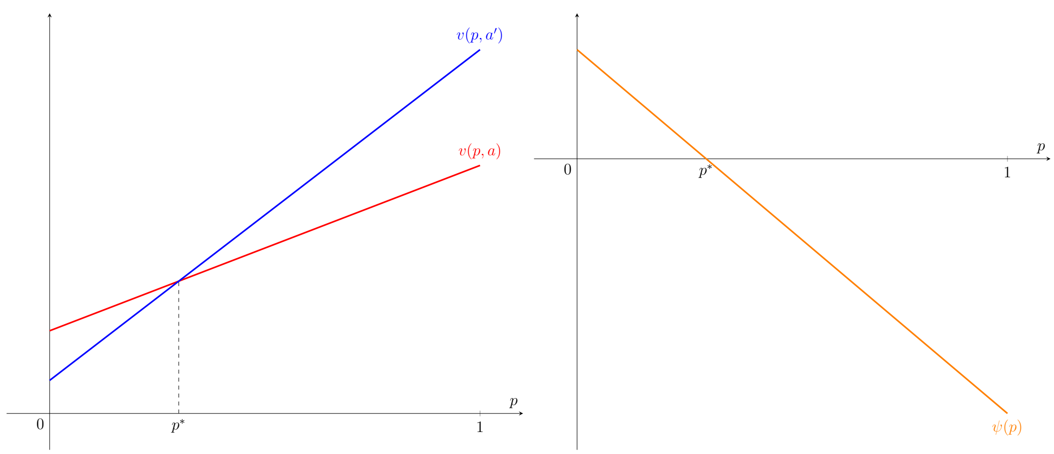

Figure 1: In the standard Bayesian Persuasion problem (Kamenica and Gentzkow, 2011), value functions are linear in the posterior and cross only once. This entails the function $\phi(\mu) = v(\mu, a) - v(\mu, a')$ crosses the horizontal axis only once.

This convexity means that a (a') is optimal for all posteriors that are a linear combination of posteriors such that the action a (a') is optimal. Therefore, for the standard Bayesian persuasion problem, the recommendation principle holds since the associated value function v is PC. That is, because the sets of posteriors supporting each action as optimal are convex.

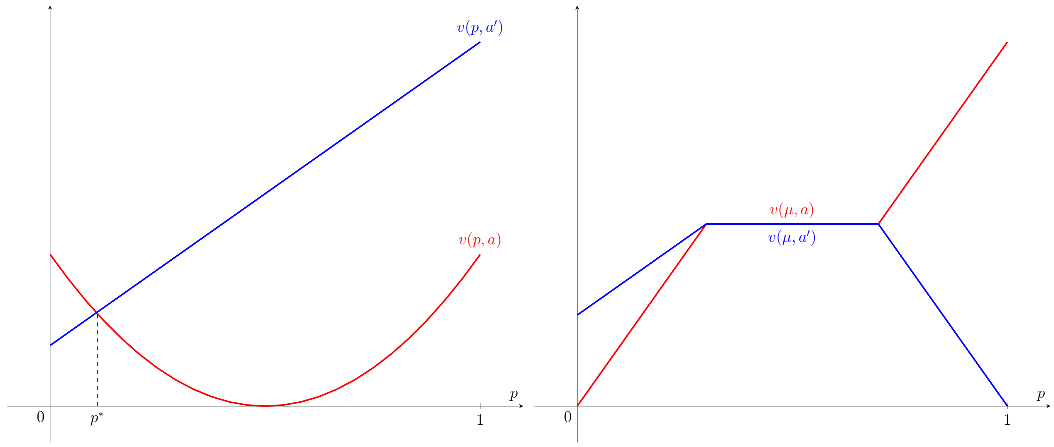

Figure 2: The recommendation principle may hold even if the value function is not linear in the posterior (left). The recommendation principle holds whenever the locus of points at which two continuous value functions intersect is convex (right).

Although v is linear in the posterior in the model of Kamenica and Gentzkow (2011), this is not necessary for the recommendation principle to hold. For example, consider the following v:

$$v(p,a) = pu(a,\omega) + (1-p)u(a,\omega'),$$

$$v(p,a') = \frac{1}{4} - p(1-p).$$

As the two functions cross only once, the sets of posteriors supporting each action are convex: this is a case in which the recommendation principle holds even if utilities are posterior dependent, since the payoff of the second action depends only on the degree of certainty, i.e., how certain she is about the state. Moreover, v is not linear in the posterior for action a': however, as long as the two curves cross only once in the interval $[0,1]$, the recommendation principle holds, as the sets of posteriors such that a is optimal and a' is optimal are convex.

For the two-action case, the single crossing of $v(\cdot,a)$ and $v(\cdot,a')$ is sufficient for the recommendation principle to hold. However, it is not necessary: for the recommendation principle to hold, it suffices that the two curves cross only in a convex set of points as in the example presented in the right panel of Figure 2.

3. Optimal Action Independent of Posterior Distribution

3.1. Main Results for the Independent Case

In this section, we study a persuasion model with a receiver whose choices only depend on the Bayesian posterior and her type.[9] Our main result in the section is that PC characterizes the class of value functions v and, therefore, of receiver's environments where the recommendation principle holds. Formally:

Theorem 1. *Suppose the receiver's choices only depend on the Bayesian posterior and her payoff type. Then, the recommendation principle holds if and only if v satisfies PC.*

As argued above, PC amounts to the convexity of the set of posteriors that makes each action plan $\bar{a}$ optimal.

This result is related to an observation in Lipnowski and Mathevet (2018), who remark that the recommendation principle may fail in their setup whenever the convexity of the set of posteriors making a given action optimal fails.[10] In fact, our result shows that the recommendation principle can hold even when the receiver prefers finer information. For example, let $|\Omega| = 2$, $A = \{0, 1\}$, and $v(a, p) = u(a, p) = ap^2 + (1 - a)c$. Even if v is convex in p for all $a \in \{0, 1\}$, it satisfies SPC and the recommendation principle holds.

Theorem 1 immediately leads to the following sufficient condition for the recommendation principle:

Corollary 1. *Let $v^*(\cdot, \theta) = \max_{a' \neq a} v(\cdot, a, \theta)$. If v is continuous in the posterior and the locus of $\mu \in B$ such that $v(\cdot, a, \theta) = v^*(\cdot, \theta)$ is a hyperplane for all $a \in A$ and $\theta \in \Theta$, the recommendation principle holds.*

Corollary 1 is a generalized single crossing condition: in particular, if $|A| = |\Omega| = 2$, the recommendation principle holds whenever, for all $\theta \in \Theta$, $v(\cdot, a, \theta)$ and $v(\cdot, a', \theta)$ intersect at most at one $\mu \in B$. While Corollary 1 is a sufficient condition for PC to hold, it is by no means necessary. The right panel of Figure 2 is a case in point: rather than for a single posterior, in that case $v(\cdot, a, \theta)$ and $v(\cdot, a', \theta)$ intersect for an interval of posteriors in B.

Corollary 1 is particularly useful in setups in which the locus of the points at which v and v^* intersect is "thin", in the sense that it has zero measure. This is the case, for instance, if small perturbations of the posteriors lead to a strict preference for one action or

[9] This also allows us to capture some forms of non-Bayesian updating, as discussed in de Clippel and Zhang (2022).

[10] Their definition of recommendation principle, however, is different from ours. In their setup, sender and receiver share the same payoff function, and they say the RP holds whenever direct recommendations maximize the payoff of both players.

another, i.e., whenever for all $\theta \in \Theta$, $\mu \in B$ and $\epsilon > 0$, there exists a μ' in the ϵ-ball around μ such that either $v(\mu', a, \theta) > v^*(\mu', \theta)$ or $v(\mu', a, \theta) < v^*(\mu', \theta)$. If $|\Omega| = 2$ and v and v^* are differentiable in p, this is the case whenever the derivatives of $\frac{\partial v(p,a,\theta)}{\partial p}$ is either strictly higher ors strictly lower than $\frac{\partial v(p,\theta)}{\partial p}$ for all $p \in [0, 1]$ and $a \in A$.

If v and v^* intersect in such a "thin" locus of points, we moreover have a partial converse to Corollary 1: whenever the recommendation principle holds, the set of posteriors at which any two actions $a, a' \in A$ are optimal must be a hyperplane. This follows from the fact that, if both a, a' are optimal at posteriors $\mu, \mu' \in B$, then both must be optimal at all posteriors $\lambda\mu + (1 - \lambda)\mu'$ on the line segment between μ and μ'.

3.2. Examples and Applications

Standard Bayesian Persuasion We can capture in our framework the standard setting in Kamenica and Gentzkow (2011) if we assume that the value function v of the receiver is equal to her expected utility given the action and the posterior considered:

$$v(\mu, a, \theta) = \sum_{\omega \in \Omega} \mu[\omega] u(\omega, a, \theta).$$

It is immediate to see that as v is an affine function of the posterior, given any type θ and any action a:

$$v(\lambda\mu + (1-\lambda)\mu', a, \theta) = \sum_{\omega \in \Omega} (\lambda\mu[\omega] + (1-\lambda)\mu'[\omega]) u(\omega, a, \theta) = \lambda v(\mu, a, \theta) + (1-\lambda) v(\mu', a, \theta).$$

Consequently, for all a, a' and μ, μ', if $v(\mu, a, \theta) \geq v(\mu, a', \theta)$ and $v(\mu', a, \theta) \geq v(\mu', a', \theta)$, then $v(\lambda\mu + (1-\lambda)\mu', a, \theta) \geq v(\lambda\mu + (1-\lambda)\mu', a')$. In turn, v satisfies SPC and the recommendation principle holds.

Non-Bayesian Expected Utility de Clippel and Zhang (2022) prove that the recommendation principle holds for non-Bayesian persuasion models where the receiver's posteriors are determined from passing Bayesian posteriors through a *distortion function* D_θ and each distortion function maps each convex combination of Bayesian posteriors μ, μ' to a convex combination of the distorted posteriors $D_\theta(\mu)$ and $D_\theta(\mu')$.[11] We verify that this property implies that v satisfies the PC property, immediately leading to the recommenda-

[11] de Clippel and Zhang (2022) also allow the distortion function to depend on the receiver's prior: we discuss this extension in Section 4.1.

tion principle. This observation allows us to make use of our framework whenever we can still express the receiver's beliefs in terms of the Bayesian posterior, even if she is not a Bayesian updater. de Clippel and Zhang (2022) show that this class of updating rules is quite large, and it includes many models of non-Bayesian updating discussed in previous literature.

We can capture the setup of de Clippel and Zhang (2022) in our framework by considering the value function

$$v(\mu, a, \theta) = \sum_{\omega \in \Omega} D_\theta(\mu(\omega))u(\omega, a, \theta).$$

As the value function of the receiver depends on her type θ, we also allow for settings in which the receiver's updating rule is correlated with her private payoff type θ.

de Clippel and Zhang (2022) show the recommendation principle holds whenever the distortion function is an affine function of the posterior. For all types θ, let D_θ be such that for all $\lambda \in [0, 1]$ and posteriors μ, μ' there exists $\gamma \in [0, 1]$ such that:

$$D_\theta(\lambda\mu + (1 - \lambda)\mu') = \gamma D_\theta(\mu) + (1 - \gamma)D_{\mu_0,\theta}(\mu').$$

Therefore, we can rewrite:

$$v(\lambda\mu + (1 - \lambda)\mu', a, \theta) = \sum_{\omega \in \Omega} D_\theta(\lambda\mu[\omega] + (1 - \lambda)\mu'[\omega])u(\omega, a, \theta).$$

Substituting for $\lambda\mu[\omega] + (1 - \lambda)\mu'[\omega]$:

$$\sum_{\omega \in \Omega} D_\theta(\lambda\mu[\omega] + (1 - \lambda)\mu'[\omega])u(\omega, a, \theta) = \sum_{\omega \in \Omega} [\gamma D_\theta(\mu[\omega]) + (1 - \gamma)D_\theta(\mu'[\omega])]u(\omega, a, \theta).$$

This entails for all $a \in A$:

$$v(\lambda\mu + (1 - \lambda)\mu', a, \theta) = \gamma v(\mu, a, \theta) + (1 - \gamma)v(\mu', a, \theta).$$

Therefore, whenever $v(\mu, a, \theta) \geq v(\mu, a', \theta)$ and $v(\mu', a, \theta) \geq v(\mu', a', \theta)$:

$$\begin{aligned} v(\lambda\mu + (1-\lambda)\mu', a, \theta) &= \gamma v(\mu, a, \theta) + (1-\gamma) v(\mu', a, \theta) \\ &\geq \gamma v(\mu, a', \theta) + (1-\gamma) v(\mu', a', \theta) \\ &= v(\lambda\mu + (1-\lambda)\mu', a', \theta) \end{aligned}$$

It follows that v satisfies SPC and, by Theorem 1, that the recommendation principle holds.

3.3. Information Design without Commitment

While this paper focuses on setups in which the sender commits to an information structure, PC is also helpful in analyzing models in which the sender has limited or no commitment power. We can interpret PC as the requirement that a receiver playing a given action after observing two different messages also follows the corresponding direct recommendation of playing that action. PC is, therefore, a condition about the receiver's choices: as long as the receiver's choices do not depend on the set of information structures the sender can pick, the presence or absence of commitment power on the part of the sender might not matter to establish whether the receiver will follow or not an action recommendation.

If the interaction between sender and receiver is analyzed through the lens of Perfect Bayesian Equilibrium (PBE henceforth) or Sequential Equilibrium (SE henceforth), this feature is enough to yield a recommendation principle analogue. In line with the related literature, assume that the space of messages S is fixed and that $A^\Theta \subset S$.

Theorem 2. *If PC holds and there exists a PBE (SE) e of the game without commitment, then there exists a PBE (SE) e' in which the sender uses the direct recommendation strategy σ^{π^*} with supp$(\sigma^{\pi^*}) \subseteq A^\Theta$ and the receiver plays the action that $\bar{a}$ recommends for her type.*

For both PBE and SE, the argument behind Theorem 2 is based on supporting equilibrium e' through the same off-path beliefs that supported equilibrium e. Importantly, this result does not mean that we can derive a recommendation principle for other refinements of PBE. In equilibria with more stringent off-path belief restrictions, it may be unfeasible to support a direct recommendation equilibrium.

4. Optimal Action Dependent on Posterior Distribution

Our results extend to models in which the optimal action for the receiver depends not only on her posterior, but also on the distribution over posteriors the experiment induces. By doing this, we include now models in which the receiver may be taking her action ex ante, i.e., before observing the realization of a given message. One example is models where the receiver can pay a cost to access a signal about the state of the world (e.g., Dall'Ara (2024); Mekonnen et al. (2025)). Another example—for some non-Bayesian updating rules—is the case of a receiver that does not update beliefs based on the Bayes rule (de Clippel and Zhang, 2022).

Let π generate the posterior distribution G, inducing the action distribution d^π. The first step is to determine the direct recommendation experiment associated with π^*. As π^* "pools" together all the messages that induce the same action plan, one may think that the distribution H over the posteriors induced by π^* will simply "pool" together all the posteriors that induce the same action plan. That is, one may think H will assign to the "average" posterior inducing $\bar{a}$ the aggregate probability of posteriors inducing $\bar{a}$ under G. We first establish that this intuition is generally incorrect.

Example 1. *Suppose that the receiver has no private information and that $\Omega = A = \{x, y\}$ with $\mu_0[x] = p_0 = \frac{1}{2}$. Let π be such that $S = \{x, y, z\}$, $\sigma(y|y) = 1$, and $\sigma(x|x) = \sigma(z|x) = \frac{1}{2}$. Denoting $\mu^\pi(\cdot)[x]$ as $p(x)$, this implies $p(y) = 0$ and $p(x) = p(z) = 1$. Suppose now v is such that $v(1, x) = v(1, y)$ and $v(0, y) > v(0, x)$. If the receiver plays y when she observes messages y or z, and x otherwise, π implements the action distribution d^π so that y is played for sure when the state is y, and x and y are played with equal probability when the state is x. Notice that π induces a posterior distribution that assigns probability $\frac{1}{2}$ to posterior 1 and probability $\frac{1}{2}$ to posterior 0. Therefore, "pooling" all the posteriors inducing actions x and y would generate a distribution over posteriors that is not properly defined, since it would assign probability 1 to posterior $\frac{1}{2}$ and probability $\frac{1}{2}$ to posterior 1. Consider now the associated direct recommendation experiment $\pi^* = (A, \sigma^{\pi^*})$, which recommends x and y with equal probability when the state is x, and recommends y with probability one when the state is y. The associated distribution of posteriors assigns probability $\frac{1}{4}$ to posterior 1 and probability $\frac{3}{4}$ to posterior $\frac{1}{3}$.*

Notice that the issue discussed in the example arises as we "pool together" posteriors 0 and 1 because they induce the same optimal action y, but we do not account for the fact that the receiver plays x half of the time when the posterior is 1. In other words, we fail to account for the fact that two experiments inducing the same distribution over posteriors may induce different distributions over actions whenever more than one action plan is optimal at each posterior.

To simplify exposition, in this section, we address this issue by assuming v is such that there exists a unique action plan $\bar{a}$ that maximizes v for each posterior μ, type θ, and posterior distribution $\bar{a}$. This is the case, for instance, if we assume that the potential receiver's indifferences are broken in favor of the sender's preferred action, as is often assumed in the literature (Bergemann and Morris, 2019). Appendix B, in the appendix, contains a more general version of our result that does not rely on this assumption.

Let $\mathcal{G} \subseteq \Delta(\Delta(\Omega))$ denote the class of distributions G such that $\text{supp}(G) \subseteq B$. For any $G \in \mathcal{G}$ and $\bar{a} \in A^{\Theta}$, denote as $B_{\bar{a}}^{G}$ the set of posteriors for which action plan $\bar{a}$ is optimal for the receiver:[12]

$$B_{\bar{a}}^{G} = \{\mu \in B : \bar{a}(\theta) \in \arg\max_{a' \in A} v(G, \mu, a', \theta) \text{ for all } \theta \in \Theta\}$$

We can now define the average posterior inducing $\bar{a}$ as:

$$\bar{\mu}_{\bar{a}}^{G} = \sum_{\mu \in B_{\bar{a}}^{G}} \frac{G(\mu)}{G(B_{\bar{a}}^{G})} \mu,$$

where $G(B_{\bar{a}}^{G}) = \sum_{\mu \in B_{\bar{a}}^{G}} G(\mu)$ is the total probability the original experiment induces a posterior μ such that $\bar{a}$ is optimal. We can then write the distribution over posterior H^G induced by the direct recommendation experiment as:

$$H^G(\mu) = \sum_{\bar{a} \in A^{\Theta}} \mathbf{1}_{\{\mu = \bar{\mu}_{\bar{a}}^{G}\}} G(B_{\bar{a}}^{G}).$$

That is, the probability H^G assigns to posterior μ is equal to the probability of any recommendation $\bar{a}$ that induces posterior μ in the direct recommendation experiment.

We are now ready to define the main condition for this section:

Definition 5 (Convex Posterior Pooling (CPP))**.** *We say a function v satisfies convex posterior pooling whenever for all $G \in \mathcal{G}$, $\bar{a} \in A^{\Theta}$, $a' \in A$, and $\theta \in \Theta$:*

$$v(H^G, \bar{\mu}_{\bar{a}}^{G}, \bar{a}(\theta), \theta) \geq v(H^G, \bar{\mu}_{\bar{a}}^{G}, a', \theta).$$

In words, CPP requires that any action plan that is optimal in the posterior μ and distribution G is still optimal for each possible pooled posterior $\bar{\mu}_{\bar{a}}^{G}$ and pooled distribution over the posteriors H^G. Notice also that CPP implies PC: It is enough to consider a posterior distribution G such that $G(\mu) = \lambda$ and $G(\mu') = 1 - \lambda$ for $\mu, \mu' \in B_{\bar{a}}$.

We can then generalize Theorem 1 as follows:

Theorem 3. *Suppose the receiver's optimal action also depends on the posterior distributions generated by the experiment. Then, the recommendation principle holds if and only if v satisfies CPP.*

[12] Recall we define $\mu : \Theta \to \Delta(\Omega)$ to capture the possibility that the receiver's posteriors may also depend on her type.

CPP holds, in particular, whenever $v(G, \cdot, \cdot, \cdot)$ satisfies PC for all G with $\text{supp}(G) \subseteq B$ and v depends on G only through the expectation of a function of the posterior μ, which is affine over the set of posteriors that induce each action plan $\bar{a}$. This is the case, for example, if v depends on G only through the ex ante expected state. Section 4.1 shows this property is satisfied by some models considered in the literature (Mekonnen et al., 2025; Dall'Ara, 2024; de Clippel and Zhang, 2022), as well as by novel examples that have not yet been explored, such as models with reference-dependent preferences or moral wiggle-room (see the next subsection).

Formally, denote by $\phi = \{\phi_k\}_{k \in K}$ any (possibly infinite) collection of functions, where $\phi : B \times A \times \Theta \to \mathbb{R}$. We say ϕ is *posterior-linear* over set $\tilde{B} \subseteq \Delta(\Omega)$ whenever for all $\mu, \mu' \in \tilde{B}$ and $\lambda \in [0, 1]$:

$$\lambda \phi(\mu, \cdot, \cdot) + (1 - \lambda)\phi(\mu', \cdot, \cdot) = \phi(\lambda \mu + (1 - \lambda)\mu', \cdot, \cdot).$$

That is, ϕ is posterior-linear if every function ϕ_k in the collection is linear in the posteriors in set $\tilde{B}$.

Denoting $\sum_{\mu \in \Delta(\Omega)} G(\mu)\phi(\mu, a, \theta)$ as $\phi^G(a, \theta)$, we then have the following corollary:

Corollary 2. *Suppose that $v(G, \cdot, \cdot, \cdot)$ satisfies PC for all $G \in \mathcal{G}$ and that there exists a collection of functions ϕ such that for all $\mu \in B$, $a \in A$, and $\theta \in \Theta$:*

$$v(G, \mu, a, \theta) = \hat{v}(\phi^G(a, \theta), \mu, a, \theta).$$

If ϕ is posterior-linear over $B^G_{\bar{a}}$ for all $\bar{a} \in A^\Theta$, the recommendation principle holds.

If v depends on G only through the expected value of collection ϕ, the posterior linearity of ϕ over the set $B^G_{\bar{a}}$ of posteriors that make $\bar{a}$ optimal ensures that "pooling" those posteriors together by using a direct recommendation experiment do not affect the value of v. Together with the fact that v satisfies PC for all G, this implies that the way v ranks the different actions available to the receiver is also unaffected. Although posterior linearity is sufficient to ensure that PCC holds and is satisfied in many of the examples presented below, it is not a necessary condition for PCC to hold if pooling posteriors affects the value of v but not the ranking it induces on the receiver's actions.

4.1. Examples and Applications

Persuaded Search The value function approach developed above is also suitable to capture dynamic persuasion applications that are *history independent*, in the sense that the payoff of the action the receiver takes after observing a message only depends on her current posterior and action.

An example is the sequential search problem of Mekonnen et al. (2025), where an agent (receiver) cannot observe the quality of the good sampled in each period, but can acquire information from a principal (sender). The paper's analysis focuses on stationary equilibria, i.e., in which the contract offered by the sender and the behavior of the receiver do not depend on the game's history, both on and off the equilibrium path.

Mekonnen et al. (2025) focus their attention on binary ("pass/fail") messages, and then show that this leads to no loss of generality, since all stationary equilibria of the game are payoff-equivalent for the players. Indeed, in a binary action framework, a pass-fail test can be interpreted as a recommendation experiment if messages are relabeled accordingly. In Mekonnen et al. (2025), the receiver will stop the search if and only if the sampled good passes the test.

By applying Theorem 3, we can confirm that it is without loss of generality to focus on such binary experiments, as the value function associated with the problem (receiver's expected utility in each period) satisfies CPP. Besides being significantly simpler, our approach also allows us to show that this result does not hinge on the choice of the sender's utility function, provided such a function only depends on the receiver's action.

Mekonnen et al. (2025) define a contract as a pair $(p_t, \pi_t)_{t\geq 0}$ of (possibly period-dependent) price and experiment, which the receiver can decide whether to accept or reject. Given the focus on stationary equilibria, we will restrict attention to contracts in which the receiver pays the same price and accesses the same experiment each period. In each period, the receiver can stop searching and buy the good ($a = s$) or continue searching in the next period ($a = c$).

We will now show that if the sender can induce some distribution of actions via contract (p, π), he can induce the same distribution over actions via contract (p, π^d), where π^d is a direct recommendation experiment. Denote as $m(\mu)$ the expected value of the good sampled in the current period given posterior μ:

$$m(\mu) = \int_{\Omega} \omega \, d\mu(\omega).$$

The value of buying the good for the receiver is then equal to her utility from stopping the search and buying the good:[13]

$$v(G, s, \mu) = v(s, \mu) = m(\mu).$$

[13]As the receiver holds no private information, we omit θ from the arguments of v.

Meanwhile, the value of continuing the search solves:

$$v(G,c,\mu) = \int_{\Delta(\Omega)} \max\{\delta v(G,c,\mu'), m(\mu')\}\, dG(\mu') = \int_{B_s^G} m(\mu')\, dG(\mu') + \int_{B_c^G} \delta v(G,c,\mu')\, dG(\mu'),$$

where B_s^G is the set of posteriors $\mu' \in \Delta(\Omega)$ such that stopping the search and buying the good is optimal ($m(\mu') \geq \delta v(G,c,\mu')$) and B_c^G is the set of posteriors such that continuing the search is optimal ($m(\mu') < \delta v(G,c,\mu')$).[14]

As the value of continuing the search does not depend on the posterior μ but only on its distribution G, we can rewrite:

$$v(G,c) = \frac{1}{1-\delta G(B_c^G)} \int_{B_s^G} m(\mu')\, dG(\mu').$$

We can now easily check that v satisfies CPP. Notice that v satisfies PC for all $G \in \mathcal{G}$, as for any given G both $v(\mu, s)$ and $v(G,c)$ are linear in μ. In addition, note that $v(G,c)$ depends on G only through the averages of $\phi_1 = m$ and $\phi_2 = \mathbf{1}_{\{\mu \in B_c^G\}}(\mu)$, which are affine over both B_c^G and B_s^G. Corollary 2 then establishes that v satisfies CPP.

This result implies that it is without loss of generality to focus on experiments that directly recommend to the receiver whether she should continue or stop searching: modulo a relabeling of the messages, this is equivalent to restricting attention to the class of binary pass/fail experiments.

As any distribution of actions that can be induced by an experiment can be induced by a pass/fail recommendation, it follows immediately that it is enough to restrict attention to contracts in which the sender offers such an experiment. This avoids the need for a lengthy proof to show that all stationary equilibria of the game are payoff-equivalent, as is the case in Mekonnen et al. (2025). Furthermore, our result establishes that, when utility for the receiver has this form, it is enough to focus on a sender that only offers binary experiments even when the sender has a different utility function than the one considered in Mekonnen et al. (2025), as long as his utility function only depends on the action the receiver takes. This means, for example, that focusing on binary experiments is enough no matter how patient the sender is or whether the sender is consistent from an intertemporal perspective.

Persuasion of an Inattentive receiver Dall'Ara (2024) considers a setup in which

[14] As $m(\mu) = 0$ for $\mu \in B_s^G \cap B_c^G$, double-counting elements of $B_s^G \cap B_c^G$ is immaterial.

a sender persuades a receiver to take a costly action and paying attention to the sender's message is costly for the receiver. Both the cost of taking the action and the cost of paying attention to the sender's message are private information of the receiver.

We can capture this setup within our framework by letting $A = \{0, 1\}$ and $\theta = (\beta, c)$, where $\beta \in \mathbb{R}$ is a parameter that captures the effort cost of paying attention and c is the material cost the receiver bears from taking action $a = 1$. Denoting again the expected value of ω given μ as $m(\mu)$, we can write the utility of the receiver given a belief μ and an effort level e as:

$$a(m(\mu) - c) - \beta k(e),$$

where parameters $\beta \in \mathbb{R}$ and $k : [0, 1] \to \mathbb{R}$ capture the cost of paying attention to the sender's message and k is known to the sender. Before sending the message, the receiver chooses an attention level $e^*(G, (\beta, c))$ to maximize:

$$e^* \int_{\Delta(\Omega)} \max\{m(\mu) - c, 0\}\, dG(\mu') + (1 - e)m(\mu_0) - \beta k(e^*).$$

We can capture this setup by letting the receiver's value function be

$$v(G, \mu, a, (\beta, c)) = a(m(\mu) - c) - \beta k(e^*(G, (\beta, c))).$$

Notice that $e^*(G, \theta)$ depends on G only through the value of information G provides, which can be rewritten as:

$$\int_{\Delta(\Omega)} \max\{m(\mu) - c, 0\}\, dG(\mu') = \int_{B_1(\beta,c)} (m(\mu) - c)\, dG(\mu') + \int_{B_0(\beta,c)} 0\, dG(\mu') =$$
$$\int_{B_1(\beta,c)} (m(\mu) - c)\, dG(\mu'),$$

where $B_1^G(\beta, c)$ denotes the set of posteriors such that $m(\mu) \geq c$ and $B_0^G(\beta, c)$ denotes the set of posteriors such that $m(\mu) < c$. As before, we can see that $m(\mu) - c$ is an affine function of μ, so its expectation is not affected by pooling posteriors that all belong to either $B_1^G(\beta, c)$ or to $B_0^G(\beta, c)$. By Corollary 2, it is then without loss of generality to focus on binary type-dependent action recommendations regardless of the utility function of the sender, as long as it depends only on the action taken by the receiver. It would also be possible to extend this result to an arbitrary action space A, provided that the utility of each action a is affine in the posterior.

Notice also that the functional form of the cost function and the fact it does not depend on G are crucial to the result: for instance, if k depended on G through a non-affine function of the posterior, pooling posteriors may affect the action taken by the receiver by affecting her optimal level of effort e^*.

Expectations-based Reference Dependence Our framework can also capture setups in which the receiver's preferences may directly depend on the distribution of posteriors. We present an example based on the intuition that losses with respect to a reference point loom larger than equivalent gains.[15] Building on the intuition that the reference point is determined by endogenous expectations (Kőszegi and Rabin, 2006), we assume that the receiver forms a reference point based on the distribution of posteriors in the experiment.

For simplicity, we focus on a simple model where $A = \{0, 1\}$ and $\Omega \subseteq \mathbb{R}$. The receiver's utility in each state is reference-dependent, that is,

$$u(a, \omega, G) = \begin{cases} a\omega + \chi(a\omega - r(G^\pi)) & \text{if } a\omega \geq r(G), \\ a\omega + \chi\lambda(a\omega - r(G^\pi)) & \text{if } a\omega < r(G), \end{cases}$$

where $\chi > 0$ captures how much the receiver values the gain-loss component of utility, $\lambda > 1$ is the loss aversion coefficient, and $r : \Delta(\Delta(\Omega)) \to \mathbb{R}$ represents the reference point, defined as the ex-ante expected value of taking action $a = 1$:

$$r(G) = \int_{\Delta(\Omega)} m(\mu)\, dG(\mu).$$

We can capture this feature of the receiver in our framework by assuming that the receiver evaluates actions according to the ex ante expected value of u:

$$v(G, \mu, a) = \int_{\Omega} u(a, \omega, G)\, d\mu(\omega).$$

As m is affine, $r(G)$ is not affected by pooling posteriors that induce the same action: as v is affine in the posterior and therefore satisfies PC for all $G \in \mathcal{G}$, Corollary 2 ensures v satisfies CPP.

Notice that a similar argument would hold if we replaced m with any affine function ϕ of the posterior. Although ϕ being affine is sufficient to ensure that v satisfies CPP, it is not a necessary condition. As an example, consider $\phi(\mu) = \max\{0, m(\mu)\}$:

$$r(G) = \int_{\Delta(\Omega)} \max\{0, m(\mu)\}\, dG(\mu).$$

[15] We refer the reader to O'Donoghue and Sprenger (2018) for an overview of the literature about reference-dependent preferences.

As $a = 1$ maximizes v if and only if $m(\mu) \geq 0$ and the maximum function is piecewise linear, $r(G)$ is unaffected by pooling posteriors that are both nonnegative or nonpositive.

Shortlisting Actions Following on the intuition that a decision-maker may be overwhelmed by the amount of choices at her disposal (see, for instance, Lleras et al. 2017), we consider the case of a receiver who can pay limited attention to the menu of actions she has available. When learning of the distribution of posteriors generated by the experiments, the receiver "shortlists" the actions that grant the top-k ex ante expected payoffs, where $k < |A|$. Then, when the message from the experiments is realized, she picks the shortlisted action granting her the highest expected payoff given the message she observed.

For the sake of simplicity, assume that the receiver's expected utility $u(a, \mu)$ is bounded below by 0. We can then capture this setup by assuming that the receiver's value function takes the value $v(G, \mu, a) = u(a, \mu)$ if a is on the shortlist $\kappa(G)$ and $v(G, \mu, a) = 0$ otherwise.

It is easy to check that v satisfies CPP. Suppose a maximizes $v(G, \mu, \cdot)$ and $v(G, \mu', \cdot)$, so that $a \in \kappa(G)$. As ex ante payoffs are linear in the posterior, by Corollary 2, we have $\kappa(G) = \kappa(H^G)$. This immediately implies that a maximizes $v(H^G, \mu, \cdot)$ and $v(H^G, \mu', \cdot)$, and that v satisfies CPP.

Non-Bayesian Expected Utility (continued) The non-Bayesian framework of the previous section can be extended to accommodate distortion functions $D_{\mu_0,\theta}$ that depend on the prior: it is enough to notice that the distribution over Bayesian posteriors generated by any experiment π must be such that the average posterior coincides with the prior μ_0. We can then capture the setup of de Clippel and Zhang (2022) by setting

$$v(G, \mu, a, \theta) = v(\mu_0, \mu, a, \theta) = \sum_{\omega \in \Omega} D_{\mu_0,\theta}(\mu(\omega))u(\omega, a, \theta).$$

Suppose now that $D_{\mu_0,\theta}$ be is such that for all $\lambda \in [0, 1]$, $\theta \in \Theta$, posteriors μ, μ', and prior μ_0, there exists $\gamma \in [0, 1]$ such that:

$$D_{\mu_0,\theta}(\lambda\mu + (1 - \lambda)\mu') = \gamma D_{\mu_0,\theta}(\mu) + (1 - \gamma)D_{\mu_0,\theta}(\mu').$$

By a similar argument to that in the previous section, this entails that for all $a \in A$:

$$v(\mu_0, \lambda\mu + (1 - \lambda)\mu', a, \theta) = \gamma v(\mu_0, \mu, a) + (1 - \gamma)v(\mu_0, \mu', a, \theta).$$

As $\gamma \in [0, 1]$, $v(\mu_0, \mu, a, \theta) \geq v(\mu_0, \mu, a', \theta)$ and $v(\mu_0, \mu', a, \theta) \geq v(\mu_0, \mu', a', \theta)$, it follows that $v(\mu_0, \cdot, \cdot, \cdot)$ satisfies SPC for all $\mu_0 \in B$. As the function $\phi(\mu, \cdot, \cdot) = \mu$ is posterior-linear for all posteriors in B, Corollary 2 then implies that the recommendation principle holds. Notice that the discussion can also be easily extended to cases in which the receiver's value function depends on G both through the distortion function and directly, as would be the case if we reconsidered the two previous examples with a non-Bayesian receiver.

Moral Wiggle Room The concept of *moral wiggle room* (Dana et al., 2007) has been proposed to explain why subjects exhibit less prosocial behavior in experiments in which they have "an excuse" to do so. For instance, a CEO may prioritize shareholders' interests when the impact of her actions on their welfare is clear, but act more self-interestedly when it is uncertain whether her decisions serve her own interests or those of the shareholders.

To capture this intuition with a simple example, assume that a CEO can either invest in a project or not ($A = \{0, 1\}$), and the value of the project for shareholders is positive or negative ($\Omega = \{-1, 1\}$). The CEO gets a payoff of $\beta < 1$ from investing (regardless of the quality of the project), plus a payoff $m(\mu)$ from the expected value of the project. Before investing, the CEO can privately observe a report on the quality of the project, modeled as an experiment that generates a posterior distribution G. The CEO incurs an *image cost* $\psi \in (\beta, 1)$ from investing in the project if she has no "excuse" she can use to justify investing, i.e., if there exists no posterior in the support of G such that $m(\mu) \geq 0$. We can interpret ψ as a cost associated with shame, social image, or loss of reputation that the CEO incurs when she cannot justify her decision to invest by claiming that she received a message that the project was not profitable.

We can capture this setup using the following function v:

$$v(G, \mu, a) = a(\beta + m(\mu) - \psi(G)),$$

where $\psi(G) = \psi$ if $m(\mu) < 0$ for all μ supported by G and $\psi(G) = 0$ otherwise.

This value function violates CPP even if the set of posteriors that support each action as optimal is convex for every distribution of posteriors G. Consider the distribution of posteriors G assigning probability $\frac{1}{2}$ to posterior $\mu = 0$, $\frac{\beta}{4\psi}$ to posterior $\mu' = \frac{1-\beta}{2}$, and $\frac{1}{2} - \frac{\beta}{4\psi}$ to posterior $\mu'' = \frac{1-\beta+\psi}{2}$. As $\beta < 1$, the action $a = 0$ is optimal for the receiver when the posterior is $\mu = 0$, while the posterior $\mu' = \frac{1-\beta}{2}$ is high enough to make the receiver indifferent between both actions when she faces no image cost ψ, but not so high that $a = 1$ is still optimal when image costs are present. The posterior $\mu'' = \frac{1-\beta+\psi}{2}$ is instead high enough to ensure that $a = 1$ is optimal even when there are image costs.

Notice that $m(\bar{\mu}) < 0$, where $\bar{\mu} = \frac{\beta}{4\psi}\mu' + (\frac{1}{2} - \frac{\beta}{4\psi})\mu''$. Therefore, pooling the posteriors μ' and μ'' means that the CEO loses any excuse to justify investing in the project, as all posteriors in the support of the new distribution H^G imply a negative expected value of the project. It follows that $\psi(H^G) = \psi$: given that $m(\bar{\mu})$ is not high enough to ensure that $a = 1$ is optimal when image costs are present, v does not satisfy CPP.

5. Discussion and Generalizations

Unique Optimal Action Plan In Section 4, we assumed that, for each posterior, only one action plan maximizes v in order to simplify the discussion. A similar result to Theorem 3 also obtains in setups in which this assumption does not hold. The main difficulty in this case is that we need a way to keep track of the action plan the receiver could potentially choose whenever more than one is optimal at some posterior $\mu \in B$. We do so through a function $\lambda_G(\mu) \in \Delta(A^\Theta)$ that assigns positive probability only to action plans that are optimal at posterior μ and distribution G. In other words, this function tells us, for each posterior, what distribution over action plans the receiver would choose. Despite the more complex notation, the argument and characterization are very similar to the ones of Theorem 3 and just require adapting the definitions of the average posterior $\bar{\mu}_{\bar{a}}^G$ inducing $\bar{a}$ and of the distribution over posteriors H^G generated by the direct recommendation experiment accordingly. Loosely speaking, the generalized characterization requires CPP to hold for every possible function λ_G of the kind described above. We discuss such generalization in full detail in Appendix B, in the appendix.

Type-dependent Priors In the Online Appendix, we further generalize the model by allowing the receiver's priors to depend on the receiver's private type. We show that a distribution over type-dependent posteriors can be generated by an experiment if and only if the product between the probability of a posterior and the posterior-prior ratio is independent of the receiver's private type. For each distribution in this class, we obtain a characterization similar to Theorem 3, with a few adjustments that account for the fact that now every type may have a different prior. If we interpret different types of the same receiver as different receivers altogether, this same machinery can serve as a starting point to discuss whether the recommendation principle holds in setups with multiple receivers who observe the same public message. In the same way, it can help to discuss setups with multiple senders, as we can interpret the message a receiver observes from a sender as her "private type" from the perspective of other senders. While extremely interesting, these applications and extensions fall beyond the scope of the current paper, and we leave them open for future research.

Finiteness assumptions In our model, we assume S, A, and Ω are finite to avoid the more cumbersome measure-theoretic notation and the addition of technical assumptions. Our most general result (Theorem 3) does not rely in any significant way on the finiteness assumption, which could be relaxed by assuming continuity of v in a and compactness of A to ensure there is at least one optimal action for each posterior, or by slightly adapting the definition of function λ_G (in the Appendix) to ensure $\lambda_G(\mu)[\bar{a}] = 0$ for all $\bar{a} \in A^\Theta$ whenever there is no $\bar{a} \in A^\Theta$ such that for all $\theta \in \Theta$ and $a' \in A$ we have $v(G, \mu, \bar{a}(\theta), \theta) \geq v(G, \mu, a', \theta)$. Theorem 1 relies instead on the finiteness of S to simplify the presentation of CP, making the meaning of the condition clearer by presenting it in terms of a convex combination of two posteriors. The proof of Theorem 1 would still go through without any finiteness assumption by replacing Lemma 1 with a non-finite version of PC by requiring that, whenever an action plan is optimal for a (possibly infinite) set of posteriors, then it is still optimal for any (possibly infinite) linear combination of posteriors in that set.

Corollary 2 requires finiteness of A to ensure that for all posteriors μ and distribution G, there exists an action plan $\bar{a} \in A^\Theta$ such that $v(G, \mu, \bar{a}(\theta), \theta) \geq v(G, \mu, a', \theta)$ for all $a' \in A$. In this case, the finiteness of A can again be replaced with the milder assumption that v is continuous in a and that A is compact.

References

Bergemann, D. and Morris, S. . Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, March 2019. doi: 10.1257/jel.20181489. URL https://www.aeaweb.org/articles?id=10.1257/jel.20181489.

Crawford, V. P. and Sobel, J. . Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

Dall'Ara, P. . Persuading an inattentive and privately informed receiver, 2024.

Dana, J. , Weber, R. A. , and Kuang, J. X. . Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory*, 33(1):67–80, 2007. ISSN 09382259, 14320479.

de Clippel, G. and Zhang, X. . Non-bayesian persuasion. *Journal of Political Economy*, 130 (10):2594–2642, 2022.

Doval, L. and Skreta, V. . Constrained information design. *Mathematics of Operations Research*, 49(1):78–106, 2024.

Kamenica, E. and Gentzkow, M. . Bayesian persuasion. *American Economic Review*, 101 (6):2590–2615, 2011.

Kartik, N. and Kleiner, A. . Convex choice. *Working Paper*, 2025.

Kartik, N. , Lee, S. , and Rappoport, D. . Single-crossing differences in convex environments. *The Review of Economic Studies*, 91(5):2981–3012, 10 2023.

Kőszegi, B. and Rabin, M. . A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165, 11 2006. ISSN 0033-5533.

Lipnowski, E. and Mathevet, L. . Disclosure to a psychological audience. *American Economic Journal: Microeconomics*, 10(4):67–93, November 2018.

Lleras, J. S. , Masatlioglu, Y. , Nakajima, D. , and Ozbay, E. Y. . When more is less: Limited consideration. *Journal of Economic Theory*, 170:70–85, 2017.

Mekonnen, T. , Murra-Anton, Z. , and Pakzad-Hurson, B. . Persuaded search. *Journal of Political Economy*, 2025.

O'Donoghue, T. and Sprenger, C. . Reference-dependent preferences. In Bernheim, B. D. , DellaVigna, S. , and Laibson, D. , editors, *Handbook of Behavioral Economics - Foundations and Applications 1*, volume 1, pages 1–77. North-Holland, 2018.

Rubbini, G. . Mechanism design without rational expectations, 2024a.

Rubbini, G. . The revelation principle without rational expectations, 2024b.

Saran, R. . Menu-dependent preferences and revelation principle. *Journal of Economic Theory*, 146(4):1712–1720, 2011. ISSN 0022-0531.

Xiong, S. . Mechanism design with sequential-move games: Revelation principle, 2024.

A. Mathematical Appendix

Proof of Theorem 1. Suppose v satisfies PC and that π implements the distribution $d^\pi : \Omega \times \Theta \to \Delta(A)$. As π implements d^π, there exists $\alpha^\pi : S \times \Theta \to A$ that maximizes v for each message s and type θ of the receiver, and that is such that for all $s \in S$ and $\theta \in \Theta$:

$$d(\omega,\theta)[a] = \sum_{s\in S} \mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}}(s)\sigma(s|\omega)$$

For any $\bar{a} : \Theta \to A$, let $\eta(\bar{a})$ denote the set of all messages s that induce the receiver to play according to $\bar{a} : \Theta \to A$ when her strategy is $\bar{a}$. That is:

$$\eta(\bar{a}) = \{s \in S : \alpha^\pi(s) = \bar{a}\}$$

Notice that, for all $\theta \in \Theta$, the posterior after observing type-dependent recommendation $\bar{a}$ such that $\eta(\bar{a}) \neq \emptyset$ is the average posterior after observing each of the messages s that would induce the receiver to play $\bar{a}$:[16]

$$\mu^{\pi^*}(\bar{a}) = \frac{\sum_{s\in\eta(\bar{a})} \sigma(s)\mu^\pi(s)}{\sum_{s\in\eta(\bar{a})} \sigma(s)}$$

This implies the posterior after observing type-dependent action recommendation $\bar{a}$ is a convex combination of the posteriors the receiver would have after receiving a message $s \in \eta(\bar{a})$. Consider now experiment $\pi^* = (A^\Theta, \sigma^{\pi^*})$, where for all $\omega \in \Omega$ and $\bar{a} \in A^\Theta$:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sum_{s\in\eta(\bar{a})} \sigma(s|\omega)$$

Notice that, as π induces a Bayes-feasible posterior distribution, so does π^*:

$$\mu_0 = \sum_{s\in S} \sigma(s)\mu^\pi(s) = \sum_{\bar{a}\in A^\Theta} \sum_{s\in\eta(\bar{a})} \sigma(s)\mu^\pi(s) = \sum_{\bar{a}\in A^\Theta} \sigma^{\pi^*}(\bar{a})\mu^{\pi^*}(\bar{a})$$

Let $\alpha^{\pi^*} : A^\Theta \times \Theta \to A$ be such that $\alpha^{\pi^*}(\bar{a},\theta) = \bar{a}(\theta)$ for all $\theta \in \Theta$. Notice α^{π^*} generates the same distribution over actions as α^π given communication strategy σ^{π^*}:

$$d^*(\omega,\theta)[a] = \sum_{\bar{a}\in A^\Theta} \mathbf{1}_{\{\alpha^{\pi^*}(\bar{a},\theta)=a\}}(\bar{a})\sigma^{\pi^*}(\bar{a}|\omega)$$

[16]Note that μ^{π^*} and μ^π are vectors indexed by ω.

$$
\begin{aligned}
&= \sum_{\bar a \in A^\Theta} \sum_{s \in \eta(\bar a)} \mathbf{1}_{\{\alpha^{\pi^*}(\bar a,\theta)=a\}}(\bar a)\sigma(s|\omega) \\
&= \sum_{\bar a \in A^\Theta} \sum_{s \in \eta(\bar a)} \mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s)\sigma(s|\omega) \\
&= \sum_{s \in \cup_{\bar a}\eta(\bar a)} \mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s)\sigma(s|\omega) \\
&= \sum_{s \in \cup_{\bar a}\eta(\bar a)} \mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s)\sigma(s|\omega) + \sum_{s \notin \cup_{\bar a}\eta(\bar a)} \mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s)\sigma(s|\omega) \\
&= \sum_{s \in S} \mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s)\sigma(s|\omega)
\end{aligned}
$$

Where the third equality follows from the fact that for all $s \in \eta(\bar a)$ we have that $\alpha^\pi(s,\cdot) = \bar a = \alpha^{\pi^*}(\bar a,\cdot)$ and the fifth equality follows from the fact that if $s \notin \eta(\bar a)$ for all $\bar a \in A^\Theta$, then $\mathbf{1}_{\{\alpha^{\pi}(s,\theta)=a\}}(s) = 0$.

It remains to prove α^{π^*} satisfies the obedience constraint, i.e. that for all $\theta \in \Theta$, $\bar a \in \text{supp}(\sigma^{\pi^*})$ and $a' \in A$:

$$
v(\mu^{\pi^*}(\bar a), \bar a(\theta), \theta) \geq v(\mu^{\pi^*}(\bar a), a', \theta)
$$

Notice that for any $\bar a \in \text{supp}(\sigma^{\pi^*})$, $s \in \eta(\bar a)$ entails that for all $\theta \in \Theta$ and $a' \in A$:

$$
v(\mu^{\pi}(s), \bar a(\theta), \theta) \geq v(\mu^{\pi}(s), a', \theta)
$$

We then invoke the following result, proven separately:

Lemma 1. *PC holds if and only if for all $\{\mu_k, \lambda_k\}_{k=1}^K$ with $\sum_{k\leq K}\lambda_k = 1$, $K < \infty$, $\bar a \in A^\Theta$ and $a' \in A$:*

$$
v(\mu_k, \bar a(\theta), \theta) \geq v(\mu_k, a', \theta) \text{ for all } k \leq K \text{ and } \theta \in \Theta \implies
$$
$$
v(\sum_{k\leq K} \lambda_k \mu_k, \bar a(\theta), \theta) \geq v(\sum_{k\leq K} \lambda_k \mu_k, a', \theta) \text{ for all } \theta \in \Theta
$$

It follows by PC, Lemma 1, and $|S| < \infty$ that for any $\theta \in \Theta$ and $a' \in A$:

$$
v(\mu^{\pi^*}(\bar a), \bar a(\theta), \theta) = v(\sum_{s\in\eta(\bar a)} \sigma(s)\mu^\pi(s), \bar a(\theta), \theta) \geq v(\sum_{s\in\eta(\bar a)} \sigma(s)\mu^\pi(s), a', \theta) = v(\mu^{\pi^*}(\bar a), a', \theta)
$$

Concluding this direction of the proof.

As for the converse, suppose the recommendation principle holds. Take any $\lambda \in [0,1]$, $\bar a \in A^\Theta$ and posteriors $\mu, \mu' \in B$ such that $v(\mu, \bar a(\theta), \theta) \geq v(\mu, a', \theta)$, $v(\mu', \bar a(\theta), \theta) \geq v(\mu', a', \theta)$ for all $\theta \in \Theta$ and $a' \in A$.

Consider prior $\mu_0 = \lambda\mu + (1-\lambda)\mu'$, and experiment $\pi = (S, \sigma)$ where $S = \{s, s'\}$ and σ is such that $\sigma(s|\omega) = \lambda\frac{\mu[\omega]}{\mu_0[\omega]}$ and $\sigma(s'|\omega) = (1-\lambda)\frac{\mu'[\omega]}{\mu_0[\omega]}$. Notice $\mu_0 \in B$ and $\mu, \mu' \in B$ and B is convex. Moreover, by construction, $\mu^\pi(s) = \mu$, $\mu^\pi(s') = \mu'$, $\sigma(s) = \lambda$ and $\sigma(s') = 1-\lambda$.

Consider now α^π such that $\alpha^\pi(s,\theta) = \alpha^\pi(s',\theta) = \bar{a}(\theta)$ for all $\theta \in \Theta$. Notice α^π maximizes v for each type θ when the posterior is either μ or μ' by construction and by the premise of PC.

Therefore, α^π implements d^π such that for all $\omega \in \Omega$, $\theta \in \Theta$ and $a \in A$:

$$d^\pi(\omega,\theta)[a] = \sigma(s|\omega)\mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}} + \sigma(s'|\omega)\mathbf{1}_{\{\alpha^\pi(s',\theta)=a\}}$$

As the recommendation principle holds for any prior distribution (including μ_0^π), d^π can also be implemented by the direct recommendation experiment $\pi^* = (A^\Theta, \sigma^{\pi^*})$ and $\alpha^{\pi^*} : A^\Theta \times \Theta \to A$ such that $\alpha^{\pi^*}(\bar{a}, \cdot) = \bar{a}$ and $d^\pi = d^*$. Therefore, for all $\omega \in \Omega$ and $\theta \in \Theta$:

$$d^\pi(\omega,\theta)[a] = d^*(\omega,\theta)[a]$$

Substituting, this entails that for all $\omega \in \Omega$ and $\theta \in \Theta$:

$$\sigma(s|\omega)\mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}} + \sigma(s'|\omega)\mathbf{1}_{\{\alpha^\pi(s',\theta)=a\}} = \sum_{\bar{a}\in A^\Theta} \mathbf{1}_{\{\alpha^{\pi^*}(\bar{a},\theta)=a\}}\sigma^{\pi^*}(\bar{a}|\omega)$$

As $\bar{a}$ is the unique action plan in A^Θ satisfying the equality for all $\theta \in \Theta$, it follows that for all $\omega \in \Omega$:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sigma(s|\omega) + \sigma(s'|\omega)$$

As α^{π^*} maximizes v for each posterior and type of the receiver, we have for all $\theta \in \Theta$:

$$v(\mu^{\pi^*}(\bar{a}), \bar{a}(\theta), \theta) \geq v(\mu^{\pi^*}(\bar{a}), a', \theta)$$

Notice moreover that:

$$\mu^{\pi^*}(\bar{a}) = \sum_{s''\in\eta(\bar{a})} \sigma(s'')\mu^\pi(s'') = \sigma(s)\mu^\pi(s) + \sigma(s')\mu^\pi(s') = \lambda\mu + (1-\lambda)\mu'$$

Therefore, for all $\theta \in \Theta$:

$$v(\lambda\mu + (1-\lambda)\mu', \bar{a}(\theta), \theta) \geq v(\lambda\mu + (1-\lambda)\mu', a', \theta)$$

This concludes the proof. □

Proof of Lemma 1. The if direction is trivial to prove. The proof of the only if direction relies on a simple induction argument. The claim is clearly true for $K = 1$. Let $\bar{a} \in A^\Theta$ and suppose $v(\mu_k, a(\theta), \theta) \geq v(\mu_k, a', \theta)$ for all $k \leq N+1$ and $\theta \in \Theta$, and that the implication holds for $K = N$. Pick an arbitrary collection $\{\lambda_k\}_{k=1}^{N+1}$. Notice that as the implication holds for $K = N$, we have $v(\sum_{k \leq K} \lambda'_k \mu_k, \bar{a}(\theta), \theta) \geq v(\sum_{k \leq K} \lambda'_k \mu_k, a', \theta)$, where $\lambda'_k = \frac{\lambda_k}{1-\lambda_{N+1}}$. Then, by PC:

$$v(\mu_{N+1}, \bar{a}(\theta), \theta) \geq v(\mu_{N+1}, a', \theta), v(\sum_{k \leq K} \lambda'_k \mu_k, \bar{a}(\theta)) \geq v(\sum_{k \leq K} \lambda'_k \mu_k, a') \text{ for all } \theta \in \Theta$$

Implies:

$$v(\sum_{k \leq K} \lambda_k \mu_k, \bar{a}(\theta)) \geq v(\sum_{k \leq K} \lambda_k \mu_k, a') \text{ for all } \theta \in \Theta$$

As our choice of $\{\lambda_k\}_{k=1}^{N+1}$ was arbitrary, this concludes the proof. □

Proof of Corollary 1. Suppose v is continuous in the posterior and the intersection of $v(\cdot, a, \theta)$ and $v^* = \max_{a' \neq a} v(\cdot, a, \theta)$ is a hyperplane. For each $a \in A$, denote the locus of points of intersection as B^*_a. We then show that the set of posteriors supporting a as optimal is convex at state θ. Notice B^*_a partitions B/B^*_a in two sets, which we denote as $B_a = \{\mu \in B : \sum_{\omega \in \Omega} b_\omega \mu[\omega] > c\}$ and $B'_a = \{\mu \in B : \sum_{\omega \in \Omega} b_\omega \mu[\omega] < c\}$. If $v(\mu, a, \theta) \geq v^*(\mu, \theta)$ for all $\mu \in B$, a is optimal over the whole set B, which is convex by assumption. Conversely, if $v(\mu, a, \theta) \leq v^*(\mu, \theta)$ for all posteriors in B, a is optimal only for posteriors in set B^*_a, which is convex. As v is continuous in its first argument and both B_a and B'_a are convex (and thus connected), the last case we need to consider is $v(\mu, a, \theta) > v^*(\mu, \theta)$ for all $\mu \in B_a$ and $v(\mu, a, \theta) < v^*(\mu, \theta)$ for all $\mu \in B'_a$ (the reasoning is analogous in case $v(\mu, a, \theta) < v^*(\mu, \theta)$ for all $\mu \in B_a$ and $v(\mu, a, \theta) > v^*(\mu, \theta)$ for all $\mu \in B'_a$). As $B_a \cup B^*_a = \{\mu \in B : \sum_{\omega \in \Omega} b_\omega \mu[\omega] \geq c\}$ is convex and our initial choice of a was arbitrary, it follows the set of posteriors for which action a is optimal is convex for all $\theta \in \Theta$ and, therefore, that SPC holds. Applying Theorem 1 concludes the proof. □

Proof of Theorem 2. Suppose PC holds and that there exists a PBE e in which the sender picks a communication strategy σ. Denote as $\mu_e(s)$ the receiver's posterior when she observes realization s, as $a_e(s, \theta)$ the action she takes in e, and as $\eta(\bar{a}) \subseteq S$ for $\bar{a} \in A^\Theta$ the set of all messages inducing the receiver to play according to the type-dependent recommendation $\bar{a}$.

Then, we can construct a PBE e' in which the sender chooses the direct recommendation strategy $\sigma^{\pi^*}(\bar{a}) = \sum_{s \in \eta(\bar{a})} \sigma(\bar{a})$ and the receiver obeys the action that $\bar{a}$ recommends for her type. To do so, it is enough to exploit the freedom PBE leaves in the choice of off-path beliefs: for any off-path realization the receiver could observe, we will assume her beliefs in

e' are the same as those she holds in e. That is, for all $s \notin \text{supp}(\sigma^{\pi^*})$ and $\theta \in \Theta$:

$$\mu_{e'}(s) = \mu_e(s)$$

By PC and the same argument as Theorem 1, the receiver will obey any action recommendation on path. Moreover, as off-path beliefs are the same in both e and e', off-path behavior for the receiver can be the same as in equilibrium e, i.e., for all θ:

$$\alpha^{\pi^*}(s,\theta) = \alpha^{\pi}(s,\theta)$$

As off-path receiver behavior is the same in both e and e', any deviation σ' from σ^{π^*} can yield the sender a payoff no higher than the one granted by σ. As σ and σ^{π^*} implement the same action distribution, it follows that σ' is not a profitable deviation from e for the sender. Therefore, e' is a PBE.

This logic can be extended to sequential equilibrium. That is, if e is a sequential equilibrium, e' is a sequential equilibrium as well. Indeed, if e is a sequential equilibrium, there exists a completely mixed sequence $\sigma^n \to \sigma$ such that the corresponding posteriors μ^n converge to μ_e. Consider the same PBE e' as above and a sequence $\sigma^{\pi^* n}[s|\omega]$ that puts weight $\frac{1}{n}$ on the sequence $\sigma^n[s|\omega]$ supporting e and weight $(1-\frac{1}{n})$ on:

- 0 if message s is off-path in e'
- $\sum_{s' \in \eta(\bar{a})} \sigma^n[s'|\omega]$ if message s is on path in e'

Formally:

$$\sigma^{\pi^* n}[s|\omega] = \frac{1}{n}\sigma^n[s|\omega] + (1-\frac{1}{n}) \sum_{\bar{a} \in A^\Theta} \mathbf{1}_{\{s=\bar{a}\}} \sigma^{\pi^*}[\bar{a}|\omega]$$

receiver's posteriors then are:

$$\mu^n_{e'}(\omega|s) = \mu_0(\omega) \frac{\frac{1}{n}\sigma^n[s|\omega] + (1-\frac{1}{n})\sum_{\bar{a} \in A^\Theta} \mathbf{1}_{\{s=\bar{a}\}} \sigma^{\pi^*}[\bar{a}|\omega]}{\frac{1}{n}\sigma^n[s] + (1-\frac{1}{n})\sum_{\bar{a} \in A^\Theta} \mathbf{1}_{\{s=\bar{a}\}} \sigma^{\pi^*}[\bar{a}]}$$

Off-path, the receiver observes $s \notin \text{supp}(\sigma^{\pi^*})$. Therefore, her beliefs are:

$$\mu^n_{e'}(\omega|s) = \mu_0(\omega) \frac{\frac{1}{n}\sigma^n[s|\omega]}{\frac{1}{n}\sigma^n[s]}$$

As e is a sequential equilibrium, $\mu_0(\omega)\dfrac{\sigma^n[s|\omega]}{\sigma^n[s]}$ converges to $\mu_e(\omega|s,\theta) = \mu_{e'}(\omega|s,\theta)$.

On path, the receiver observes instead $\bar{a} \in \text{supp}(\sigma^{\pi^*})$ and has beliefs:

$$\mu^n_{e'}(\omega|\bar{a}) = \mu_0(\omega)\frac{\frac{1}{n}\sigma^n[\bar{a}|\omega] + (1-\frac{1}{n})\sigma^{\pi^*}[\bar{a}|\omega]}{\frac{1}{n}\sigma^n[\bar{a}] + (1-\frac{1}{n})\sigma^{\pi^*}[\bar{a}]}$$

As $\bar{a}$ is on path, the limit of the denominator of the equality above is non-zero. Therefore, the limit of $\mu^n_{e'}$ is just the ratio of the limit of the numerator and the limit of the denominator, and thus:

$$\mu^n_{e'}(\omega|\bar{a}) \to \mu_0(\omega)\frac{\sigma^{\pi^*}[\bar{a}|\omega]}{\sigma^{\pi^*}[\bar{a}]} = \mu_{e'}(\omega|\bar{a})$$

□

Proof of Theorem 3. Suppose v satisfies CPP and that π implements the distribution $d^\pi : \Omega \times \Theta \to \Delta(A)$. Denote as G the distribution over posteriors induced by π, and recall it averages to the prior μ_0 due to Bayes feasibility.

As π implements d^π, there exists $\alpha^\pi : S \times \Theta \to A$ that maximizes v for each message s and type θ of the receiver when the distribution of posteriors is G, and that is such that for all $s \in S$ and $\theta \in \Theta$:

$$d(\omega,\theta)[a] = \sum_{s\in S} \mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}}\sigma(s|\omega)$$

For any $\bar{a} : \Theta \to A$, let $\eta(\bar{a}) = \{s \in S : \alpha^\pi(s) = \bar{a}\}$ denote the set of all messages s that induce the receiver to play according to $\bar{a}$. Consider experiment $\pi^* = (A^\Theta, \sigma^{\pi^*})$, where for all $\omega \in \Omega$ and $\bar{a} \in A^\Theta$:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sum_{s\in\eta(\bar{a})} \sigma(s|\omega)$$

For all $\theta \in \Theta$ and $\bar{a}$ such that $\eta(\bar{a}) \neq \emptyset$ we have:

$$\mu^{\pi^*}(\bar{a}) = \frac{\sum_{s\in\eta(\bar{a})}\sigma(s)\mu^\pi(s)}{\sum_{s\in\eta(\bar{a})}\sigma(s)}$$

Moreover, π^* induces distribution over posteriors H^G such that:

$$\begin{aligned} H^G(\mu) &= \sum_{\bar{a}\in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}}\sigma^{\pi^*}(\bar{a}) \\ &= \sum_{\bar{a}\in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}}\left(\sum_{s\in\eta(\bar{a})}\sigma(s)\right) \\ &= \sum_{\bar{a}\in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}}G(B^G_{\bar{a}}) \end{aligned}$$

$$= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\bar{\mu}_{\bar{a}}^G = \mu\}} G(B_{\bar{a}}^G)$$

Where the last step follows from the fact that, for all $\bar{a}$ such that $\eta(\bar{a}) \neq \emptyset$:

$$\begin{aligned}
\mu^{\pi^*}(\bar{a}) &= \frac{\sum_{s \in \eta(\bar{a})} \sigma(s) \mu^\pi(s)}{\sum_{s \in \eta(\bar{a})} \sigma(s)} \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s) = \mu\}} \sigma(s) \mu \right)}{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s) = \mu\}} \sigma(s) \right)} \\
&= \frac{\sum_{\mu \in B_{\bar{a}}^G} G(\mu) \mu}{\sum_{\mu \in B_{\bar{a}}^G} G(\mu)} \\
&= \frac{\sum_{\mu \in B_{\bar{a}}^G} G(\mu) \mu}{G(B_{\bar{a}}^G)} \\
&= \bar{\mu}_{\bar{a}}^G
\end{aligned}$$

Where the third equality follows from the fact that $s \in \eta(\bar{a})$ if and only if $\mu^\pi(s) = \mu$ for $\mu \in B_{\bar{a}}^G$. As α^π maximizes v for all types θ, we have that for all $\theta \in \Theta$, $s \in \eta(\bar{a})$, and $a' \in A$:

$$v(G, \mu^\pi(s), \alpha^\pi(s, \theta), \theta) \geq v(G, \mu^\pi(s), a', \theta)$$

So that $\mu^\pi(s) \in B_{\bar{a}}^G$ for all $s \in \eta(\bar{a})$. Moreover, as $B_{\bar{a}} \subseteq \text{supp}(G)$, for all $\mu \in B_{\bar{a}}$ there exists $s \in S$ such that $\mu^\pi(s) = \mu$. By CPP and $\bar{\mu}_{\bar{a}}^G = \mu^{\pi^*}(\bar{a}, \cdot)$ for all $\bar{a} \in A^\Theta$, this implies that for all $\theta \in \Theta$ and $a' \in A$:

$$v(H^G, \mu^{\pi^*}(\bar{a}), \bar{a}(\theta), \theta) = v(H^G, \mu^{\pi^*}(\bar{a}), \alpha^\pi(s, \theta), \theta) \geq v(H^G, \mu^{\pi^*}(\bar{a}), a', \theta)$$

Consider $\alpha^{\pi^*} : A^\Theta \times \Theta \to A$ is such that $\alpha^{\pi^*}(\bar{a}, \theta) = \bar{a}(\theta)$ and for all $\theta \in \Theta$. As α^{π^*} still induces action distribution d^π given communication strategy σ^{π^*} and the obedience constraint is satisfied, we conclude the proof.

As for the converse statement, suppose the recommendation principle holds. Take any G with $\text{supp}(G) \subseteq B$ and $\bar{a} \in A^\Theta$. As the recommendation principle must hold for all priors in B, let us consider prior $\mu_0 = \sum_{\mu \in \text{supp}(G)} \mu G(\mu)$. As B is convex and $\text{supp}(G) \subseteq B$, $\mu_0 \in B$. Consider now experiment $\pi = (\text{supp}(G), \sigma)$ such that for all $\omega \in \Omega$ with $\mu_0[\omega] > 0$:

$$\sigma(\mu|\omega) = \frac{\mu[\omega]}{\mu_0[\omega]} G(\mu)$$

While we let $\sigma(\mu|\omega)$ be any arbitrary lottery over $\Delta(S)$ whenever $\mu_0[\omega] = 0$.

It is immediate to notice $\sigma(\cdot|\omega) \in \Delta(S)$ whenever $\mu_o[\omega] > 0$ as:

$$\sum_{\mu \in \text{supp}(G)} \sigma(\mu|\omega) = \frac{1}{\mu_0[\omega]} \left(\sum_{\mu \in \text{supp}(G)} G(\mu)\mu[\omega] \right) = 1$$

Where the first equality follows from the fact that the prior equals the average posterior. Moreover, applying Bayes' rule yields that $\mu^\pi(\mu') = \mu$ if and only if $\mu = \mu'$. This implies π induces posterior distribution G as:

$$\sigma(\mu) = \sum_{\omega \in \Omega} \sigma(\mu|\omega)\mu_0[\omega] = G(\mu)$$

Denote now as $\eta(\bar{a})$ the set of all messages inducing action plan $\bar{a}$, i.e.:

$$\eta(\bar{a}) = \{s \in S : \mu^\pi(s) \in B^G_{\bar{a}}\}$$

As the recommendation principle holds, we can implement the action distribution d^π induced by π and α^π via a direct recommendation experiment π^*. Notice that, by an argument similar to the one in the proof of Theorem 3, π^* pools together all messages in S that induce the same action:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sum_{s \in \eta(\bar{a})} \sigma(s|\omega)$$

This implies:

$$\sigma^{\pi^*}(\bar{a}) = \sum_{\omega \in \Omega} \sum_{s \in \eta(\bar{a})} \sigma(s|\omega) = \sum_{s \in \eta(\bar{a})} \sigma(s) = G(B^G_{\bar{a}})$$

Therefore, π^* induces posterior distribution:

$$\begin{aligned} H^G(\mu) &= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a}) = \mu\}} \sigma^{\pi^*}(\bar{a}) \\ &= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a}) = \mu\}} G(B^G_{\bar{a}}) \\ &= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\bar{\mu}^G_{\bar{a}} = \mu\}} G(B^G_{\bar{a}}) \end{aligned}$$

Where the last inequality follows from the fact the posterior the receiver holds after observing message $\bar{a}$ is, for all $\omega \in \Omega$ and $\theta \in \Theta$:

$$\begin{aligned}
\mu^{\pi^*}(\bar{a})[\omega] &= \frac{\sigma^{\pi^*}(\bar{a}|\omega)}{\sigma^{\pi^*}(\bar{a})}\mu_0[\omega] \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s \in S} \mathbf{1}_{\{\mu^{\pi}(s)=\mu\}}\sigma(s|\omega)\right)}{G(B^G_{\bar{a}})}\mu_0[\omega] \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s' \in S'} \mathbf{1}_{\{\mu^{\pi'}(s')=\mu\}}\sigma(s'|\omega)\mu_0[\omega]\right)}{G(B^G_{\bar{a}})} \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s' \in S'} \mathbf{1}_{\{\mu^{\pi'}(s')=\mu\}}\sigma(s)\mu[\omega]\right)}{G(B^G_{\bar{a}})} \\
&= \frac{\sum_{\mu \in B^G_{\bar{a}}} G(\mu)\mu[\omega]}{G(B^G_{\bar{a}})} \\
&= \bar{\mu}^G_{\bar{a}}[\omega]
\end{aligned}$$

Where the fourth equality follows again from the fact $\sigma(s'|\omega)\mu_0[\omega] = \sigma(s')\mu[\omega]$ for s' such that $\mu^{\pi'}(s) = \mu$ and that $G(\mu) = \sum_{s' \in S'} \mathbf{1}_{\{\mu^{\pi'}(s')=\mu\}}\sigma(s)$. As α^{π^*} such that $\alpha^{\pi^*}(\bar{a}, \cdot) = \bar{a}$, we then have that for all $a' \in A$ and $\theta \in \Theta$:

$$v(H^G, \bar{\mu}^G_{\bar{a}}, \bar{a}(\theta), \theta) = v(H^G, \mu^{\pi^*}(\bar{a}), \bar{a}(\theta), \theta) \geq v(H^G, \mu^{\pi^*}(\bar{a}), a', \theta) = v(H^G, \bar{\mu}^G_{\bar{a}}, a', \theta)$$

This concludes the proof. □

Proof of Corollary 2. By Theorem 3, we need only to check v is CPP. Notice that for all $a \in A$, $\theta \in \Theta$, and $\bar{a} \in A^\Theta$, we have by assumption that:

$$\begin{aligned}
\phi^G(a, \theta) &= \sum_{\mu \in \Delta(\Omega)} G(\mu)\phi(\mu, a, \theta) \\
&= \sum_{\bar{a} \in A^\Theta} \sum_{\mu \in B^G_{\bar{a}}} G(\mu)\phi(\mu, a, \theta)
\end{aligned}$$

Where the second equality follows from the fact that, as A is finite, for each μ and θ there exists $\bar{a}$ maximizing v.[17] Using the fact ϕ is posterior-linear over $B^G_{\bar{a}}$ for all $\bar{a}$, we can then rewrite:

$$\phi^G(a, \theta) = \sum_{\bar{a} \in A^\Theta} \sum_{\mu \in B^G_{\bar{a}}} G(\mu)\phi(\mu, a, \theta)$$

[17]Notice that for all $\theta \in \Theta$ and $\mu \in B$, finiteness of A implies there exists $\bar{a}(\theta)$ such that $v(G, \mu, \bar{a}(\theta), \theta) \geq v(G, \mu, a', \theta)$ for all $a' \in A$.

$$
\begin{aligned}
&= \sum_{\bar{a}\in A^\Theta} G(B_{\bar{a}}^G)\phi(\bar{\mu}_{\bar{a}}^G(\theta), a, \theta) \\
&= \sum_{\bar{a}\in A^\Theta} \left(\sum_{\mu\in\Delta(\Omega)} \mathbf{1}_{\{\mu=\bar{\mu}_{\bar{a}}^G\}} G(B_{\bar{a}}^G)\phi(\mu, a, \theta) \right) \\
&= \sum_{\mu\in\Delta(\Omega)} \left(\sum_{\bar{a}\in A^\Theta} \mathbf{1}_{\{\mu=\bar{\mu}_{\bar{a}}^G\}} G(B_{\bar{a}}^G) \right) \phi(\mu, a, \theta) \\
&= \sum_{\mu\in\Delta(\Omega)} H^G(\mu)\phi(\mu, a, \theta) \\
&= \phi^{H^G}(a, \theta)
\end{aligned}
$$

Where the second and third equality follow from the definition of $\bar{\mu}_{\bar{a}}^G$. Therefore, for all $\bar{a}$, a', and θ:

$$
v(H^G, \mu_{\bar{a}}^G, \bar{a}(\theta), \theta) = v(G, \mu_{\bar{a}}^G, \bar{a}(\theta), \theta) \geq v(G, \mu_{\bar{a}}^G, a', \theta) = v(H^G, \mu_{\bar{a}}^G, a', \theta)
$$

Where the inequality follows from the fact v is PC. This concludes the proof. □

B. Generalization of Section 4

In this section, we generalize the results of Section 4 by dropping the assumption that for each $G \in \Delta(\Omega)$ and $\mu \in \text{supp}(G) \subseteq B$ there exists a unique action plan $\bar{a}$ such that $v(G, \mu, \bar{a}(\theta), \theta) \geq v(G, \mu, a', \theta)$ for all $\theta \in \Theta$ and $a' \in A$.

For any G and $\bar{a} \in A^\Theta$, let $B_{\bar{a}}^G$ be the set of posterior for which action plan $\bar{a}$ is optimal for the receiver:

$$
B_{\bar{a}}^G = \{\mu \in B : \bar{a}(\theta) \in \arg\max_{a'\in A} v(G, \mu, a', \theta) \text{ for all } \theta \in \Theta\}
$$

We then denote as $\lambda_G : \text{supp}(G) \to \Delta(A^\theta)$ any function such that $\bar{a} \in \text{supp}(\lambda_G(\mu))$ implies $\mu \in B_{\bar{a}}$. In other words, λ_G is any function mapping the support of G to a probability distribution over action plans such that any $\bar{a} \in \text{supp}(\lambda_G(\mu))$ maximizes v at μ. In some sense, we can interpret $\lambda^G(\mu)$ as telling us how the total probability $G(\mu)$ of posterior μ should be "partitioned" among all action plans that maximize v at μ. Note that if the receiver's preferences are strict at each μ, $\lambda_G(\mu)[\bar{a}] = 1$ for exactly one $\bar{a} \in A^\Theta$.

The average posterior inducing $\bar{a}$ (according to λ_G) is defined as:

$$\bar{\mu}_{\bar{a}}^{\lambda_G} = \sum_{\mu \in \text{supp}(G)} \left(\frac{\lambda_G(\mu)[\bar{a}]G(\mu)}{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)} \right) \mu$$

If we interpret $\lambda_G(\mu)[\bar{a}]G(\mu)$ as the probability $\bar{a}$ is played and that the posterior is μ, we can interpret $\bar{\mu}_{\bar{a}}^{\lambda_G}$ as the posterior the receiver would form if she observed recommendation $\bar{a}$ in the direct mechanism.[18]

The distribution H^{λ_G} induced by the direct mechanism is:

$$H^{\lambda_G}(\mu) = \sum_{\bar{a} \in A^\Theta} \mathbf{1}_{\{\mu = \bar{\mu'}_{\bar{a}}^{\lambda_G}\}} \left(\sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}]G(\mu') \right)$$

Where $\sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}]G(\mu')$ is the total probability $\bar{a}$ is played under posterior distribution G, and we sum over all possible $\bar{a}$ action plans that lead to the same average posterior.

We then say a function v satisfies *convex posterior pooling* (CPP) whenever for all $G \in \mathcal{G}$, and $\bar{a}$, $a' \in A$, θ and λ_G:

$$v(H^{\lambda_G}, \bar{\mu}_{\bar{a}}^{\lambda_G}, \bar{a}(\theta), \theta) \geq v(H^{\lambda_G}, \bar{\mu}_{\bar{a}}^{\lambda_G}, a', \theta)$$

Theorem 4. *Suppose the receiver's optimal action depends on her posterior, her payoff type, and the posterior distributions generated by the experiment. Then, the recommendation principle holds if and only if v satisfies CPP.*

We can also generalize the argument for Corollary 2. Denote as $\phi = \{\phi_k\}_{k \in K}$ any (possibly infinite) collection of functions with $\phi : B \times A \times \Theta \to \mathbb{R}$. We say ϕ is *posterior-linear* over set $\tilde{B} \subseteq \Delta(\Omega)$ whenever for all $\mu, \mu' \in \tilde{B}$ and $\lambda \in [0, 1]$:

$$\lambda\phi(\mu, \cdot, \cdot) + (1 - \lambda)\phi(\mu', \cdot, \cdot) = \phi(\lambda\mu + (1 - \lambda)\mu', \cdot, \cdot)$$

Denoting $\sum_{\mu \in \Delta(\Omega)} G(\mu)\phi(\mu, a, \theta)$ as $\phi^G(a, \theta)$, we have:

Corollary 3. *Suppose that $v(G, \cdot, \cdot, \cdot)$ is PC for all $G \in \mathcal{G}$ and that there exists a collection of functions ϕ such that for all $\mu \in B$, $a \in A$, and $\theta \in \Theta$:*

$$v(G, \mu, a, \theta) = \hat{v}(\phi^G(a, \theta), \mu, a, \theta)$$

[18]Notice in this case we can interpret $\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)$ as the total probability $\bar{a}$ is played under posterior distribution G.

If ϕ is posterior-linear over $B^G_{\bar{a}}$ for all $\bar{a} \in A^\Theta$, the recommendation principle holds.

B.1. Proofs

Proof of Theorem 4. Suppose v satisfies CPP and that π implements the distribution $d^\pi : \Omega \times \Theta \to \Delta(A)$. Denote as G the distribution over posterior induced by π, and recall it averages to the prior μ_0 due to Bayes feasibility.

As π implements d^π, there exists $\alpha^\pi : S \times \Theta \to A$ that maximizes v for each message s and type θ of the receiver when the distribution of posteriors is G, and that is such that for all $s \in S$ and $\theta \in \Theta$:

$$d(\omega, \theta)[a] = \sum_{s \in S} \mathbf{1}_{\{\alpha^\pi(s,\theta)=a\}}(s)\sigma(s|\omega)$$

For any $\bar{a} : \Theta \to A$, let $\eta(\bar{a})$ denote the set of all messages s that induce the receiver to play according to $\bar{a}$. Let $\eta(\bar{a}) = \{s \in S : \alpha^\pi(s, \cdot) = \bar{a}\}$ and consider now experiment $\pi^* = (A^\Theta, \sigma^{\pi^*})$, where for all $\omega \in \Omega$ and $\bar{a} \in A^\Theta$:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sum_{s \in \eta(\bar{a})} \sigma(s|\omega)$$

Notice that for all $\theta \in \Theta$ and $\bar{a}$ such that $\eta(\bar{a}) \neq \emptyset$:

$$\mu^{\pi^*}(\bar{a}) = \frac{\sum_{s \in \eta(\bar{a})} \sigma(s)\mu^\pi(s)}{\sum_{s \in \eta(\bar{a})} \sigma(s)}$$

Moreover, π^* induces distribution over posteriors H^{λ_G} such that:

$$\begin{aligned}
H^{\lambda_G}(\mu) &= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}} \sigma^{\pi^*}(\bar{a}) \\
&= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}} \left(\sum_{s \in \eta(\bar{a})} \sigma(s) \right) \\
&= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}} \left(\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu) \right)
\end{aligned}$$

Where for all $\mu \in \text{supp}(G)$ and $\bar{a} \in A^\Theta$:

$$\lambda_G(\mu)[\bar{a}] = \frac{1}{G(\mu)} \sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s)$$

Notice the last equality above follows then from the fact that:

$$\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}] G(\mu) = \sum_{\mu \in \text{supp}(G)} \sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s) = \sum_{s \in \eta(\bar{a})} \sigma(s)$$

We can then rewrite, for all $\bar{a}$ such that $\eta(\bar{a}) \neq \emptyset$ and $\theta \in \Theta$:

$$\begin{aligned}
\mu^{\pi^*}(\bar{a}) &= \frac{\sum_{s \in \eta(\bar{a})} \sigma(s) \mu^\pi(s)}{\sum_{s \in \eta(\bar{a})} \sigma(s)} \\
&= \frac{\sum_{s \in \eta(\bar{a})} \left(\sum_{\mu \in \text{supp}(G)} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s) \mu \right)}{\sum_{s \in \eta(\bar{a})} \left(\sum_{\mu \in \text{supp}(G)} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s) \right)} \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s) \right) \mu}{\sum_{\mu \in \text{supp}(G)} \left(\sum_{s \in \eta(\bar{a})} \mathbf{1}_{\{\mu^\pi(s)=\mu\}} \sigma(s) \right)} \\
&= \frac{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}] G(\mu) \mu}{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}] G(\mu)} \\
&= \bar{\mu}_{\bar{a}}^{\lambda_G}
\end{aligned}$$

Where the fourth equality follows from the definition of λ_G. As α^π maximizes v for all types θ, we have that for all $\theta \in \Theta$, $s \in \eta(\bar{a})$, and $a' \in A$:

$$v(G, \mu^\pi(s, \theta), \alpha^\pi(s, \theta), \theta) \geq v(G, \mu^\pi(s, \theta), a', \theta)$$

So that $\mu^\pi(s) \in B_{\bar{a}}^G$ for all $s \in \eta(\bar{a})$. Moreover, as $B_{\bar{a}} \subseteq \text{supp}(G)$, for all $\mu \in B_{\bar{a}}$ there exists $s \in S$ such that $\mu^\pi(s) = \mu$. By CPP this implies that for all $\theta \in \Theta$ and $a' \in A$:

$$v(H^{\lambda_G}, \mu^{\pi^*}(\bar{a}), \bar{a}(\theta), \theta) = v(H^{\lambda_G}, \mu^{\pi^*}(\bar{a}), \alpha^\pi(s, \theta), \theta) \geq v(H^{\lambda_G}, \mu^{\pi^*}(\bar{a}), a', \theta)$$

Consider $\alpha^{\pi^*} : A^\Theta \times \Theta \to A$ is such that $\alpha^{\pi^*}(\bar{a}, \theta) = \bar{a}(\theta)$ and for all $\theta \in \Theta$. As α^{π^*} still induces action distribution d^π given communication strategy σ^{π^*} and the obedience constraint is satisfied, we conclude the proof.

As for the converse statement, suppose the recommendation principle holds. Take any G, $\lambda_G : \text{supp}(G) \to \Delta(A^\Theta)$ and $\bar{a} \in A^\Theta$, and prior $\mu_0 = \sum_{\mu \in \text{supp}(G)} \mu G(\mu) \in B$. Consider now experiment $\pi = (\text{supp}(G) \times A^\Theta, \sigma)$ such that for all $(\mu, \bar{a}) \in (\text{supp}(G) \times A^\Theta)$ and $\omega \in \Omega$ with $\mu_0[\omega] > 0$:

$$\sigma((\mu, \bar{a})|\omega) = \lambda(\mu)[\bar{a}]\frac{\mu[\omega]G(\mu)}{\mu_0[\omega]}$$

For ω with $\mu_0[\omega] = 0$, let instead $\sigma(\cdot|\omega)$ be any distribution over messages $\Delta(S)$.

We first show $\sigma(\cdot|\omega) \in \Delta(S)$ for any ω such that $\mu_0[\omega] > 0$. Notice that:

$$\begin{aligned}
\sum_{\mu \in \text{supp}(G)} \sum_{\bar{a} \in A^\Theta} \sigma((\mu, \bar{a})|\omega) &= \sum_{\mu \in \text{supp}(G)} \sum_{\bar{a} \in A^\Theta} \left(\lambda(\mu)[\bar{a}]\frac{\mu[\omega]G(\mu)}{\mu_0[\omega]}\right) \\
&= \sum_{\mu \in \text{supp}(G)} \left(\sum_{\bar{a} \in A^\Theta} \lambda(\mu)[\bar{a}]\right) \frac{\mu[\omega]G(\mu)}{\mu_0[\omega]} \\
&= \sum_{\mu \in \text{supp}(G)} \frac{\mu[\omega]G(\mu)}{\mu_0[\omega]} \\
&= \frac{1}{\mu_0[\omega]} \left(\sum_{\mu \in \text{supp}(G)} G(\mu)\mu[\omega]\right) \\
&= \frac{1}{\mu_0[\omega]}\mu_0[\omega] \\
&= 1
\end{aligned}$$

A similar argument delivers $\sigma(\cdot|\omega) \in \Delta(S)$ for $\omega \in \Omega$ with $\mu_0[\omega] = 0$ for all $\theta \in \Theta$. Moreover, it is easy to see that $\mu^\pi((\mu, \bar{a})) = \mu$ and that π induces the same distribution over posteriors as G as for all $\mu \in \text{supp}(G)$:

$$\sum_{\bar{a} \in A^\Theta} \sigma(\mu, \bar{a}) = \sum_{\bar{a} \in A^\Theta} \lambda(\mu)[\bar{a}]G(\mu) = G(\mu)$$

Consider now α^π such that $\alpha^\pi((\mu, \bar{a}), \theta) = \bar{a}(\theta)$. As the recommendation principle holds, we can implement the action distribution d^π induced by π and α^π via a direct recommendation experiment π^*. By an argument similar to the one in the proof of Theorem 3, π^* pools together all messages in S that induce the same action:

$$\sigma^{\pi^*}(\bar{a}|\omega) = \sum_{\mu \in \text{supp}(G)} \sigma((\mu, \bar{a})|\omega)$$

This implies:

$$\sigma^{\pi^*}(\bar{a}) = \sum_{\mu \in \text{supp}(G)} \sigma((\mu, \bar{a}))$$

$$= \sum_{\mu \in \text{supp}(G)} \left(\sum_{\omega \in \Omega} \sigma((\mu, \bar{a})|\omega)\mu_0[\omega] \right)$$

$$= \sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)$$

Therefore, π^* induces the posterior distribution:

$$H^{\lambda_G}(\mu) = \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}} \sigma^{\pi^*}(\bar{a})$$

$$= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\mu^{\pi^*}(\bar{a})=\mu\}} \left(\sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}]G(\mu') \right)$$

$$= \sum_{\bar{a} \in A^\theta} \mathbf{1}_{\{\bar{\mu}_{\bar{a}}^{\lambda_G}=\mu\}} \left(\sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}]G(\mu') \right)$$

Where the last inequality follows from the fact that the posterior the receiver holds after observing message $\bar{a}$ is, for all $\omega \in \Omega$:

$$\mu^{\pi^*}(\bar{a})[\omega] = \frac{\sigma^{\pi^*}(\bar{a}|\omega)}{\sigma^{\pi^*}(\bar{a})}\mu_0[\omega]$$

$$= \frac{\sum_{\mu \in \text{supp}(G)} \sigma((\mu, \bar{a})|\omega)\mu_0[\omega]}{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)}$$

$$= \frac{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)\mu[\omega]}{\sum_{\mu \in \text{supp}(G)} \lambda_G(\mu)[\bar{a}]G(\mu)}$$

$$= \bar{\mu}_{\bar{a}}^{\lambda_G}(\theta)[\omega]$$

Where the fourth equality follows again from the fact $\sigma((\mu, \bar{a})|\omega)\mu_0[\omega] = \mu[\omega]\sigma((\mu, \bar{a})) = \mu[\omega]\lambda_G(\mu)[\bar{a}]G(\mu)$ for s such that $\mu^\pi(s) = \mu$ and that $\sum_{s \in s} \mathbf{1}_{\{\mu^\pi(s)=\mu\}}\sigma(s) = G(\mu)$ for all $\mu \in \text{supp}(G)$.

As α^{π^*} such that $\alpha^{\pi^*}(\bar{a}, \cdot) = \bar{a}$, we then have that for all $a' \in A$ and $\theta \in \Theta$:

$$v(H^{\lambda_G}, \bar{\mu}_{\bar{a}}^{\lambda_G}(\theta), \bar{a}(\theta), \theta) = v(H^{\lambda_G}, \mu^{\pi^*}(\bar{a}), \bar{a}(\theta), \theta) \geq v(H^{\lambda_G}, \mu^{\pi^*}(\bar{a}), a', \theta) = v(H^{\lambda_G}, \bar{\mu}_{\bar{a}}^{\lambda_G}(\theta), a', \theta)$$

This concludes the proof. □

Proof of Corollary 3. By Theorem 4, we need only to check v is CPP. Notice that for all

$a \in A$, $\theta \in \Theta$ and λ_G:

$$\begin{aligned}
\phi^G(a,\theta) &= \sum_{\mu \in \Delta(\Omega)} G(\mu)\phi(\mu,a,\theta) \\
&= \sum_{\mu \in \Delta(\Omega)} \left(\sum_{\bar{a} \in A^\Theta} \lambda_G(\mu)[\bar{a}] \right) G(\mu)\phi(\mu,a,\theta) \\
&= \sum_{\bar{a} \in A^\Theta} \sum_{\mu \in \Delta(\Omega)} \lambda_G(\mu)[\bar{a}] G(\mu)\phi(\mu,a,\theta) \\
&= \sum_{\bar{a} \in A^\Theta} \sum_{\mu \in B^G_{\bar{a}}} \lambda_G(\mu)[\bar{a}] G(\mu)\phi(\mu,a,\theta)
\end{aligned}$$

Where the last equality follows from the fact $\lambda_G(\mu)[\bar{a}] = 0$ for all $\mu \notin B^G_{\bar{a}}$. Using the fact ϕ is posterior-linear over $B^G_{\bar{a}}$ for all $\bar{a}$, we can then rewrite:

$$\begin{aligned}
\phi^G(a,\theta) &= \sum_{\bar{a} \in A^\Theta} \left(\sum_{\mu \in B^G_{\bar{a}}} \lambda_G(\mu)[\bar{a}] G(\mu)\phi(\mu,a,\theta) \right) \\
&= \sum_{\bar{a} \in A^\Theta} \left(\sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}] G(\mu') \right) \phi(\bar{\mu}^{\lambda_G}_{\bar{a}}(\theta),a,\theta) \\
&= \sum_{\mu \in \Delta(\Omega)} \left(\sum_{\bar{a} \in A^\Theta} \mathbf{1}_{\{\mu = \bar{\mu}^{\lambda_G}_{\bar{a}}\}} \sum_{\mu' \in \text{supp}(G)} \lambda_G(\mu')[\bar{a}] G(\mu') \right) \phi(\mu(\theta),a,\theta) \\
&= \sum_{\mu \in \Delta(\Omega)} H^{\lambda_G}(\mu)\phi(\mu(\theta),a,\theta) \\
&= \phi^{H^{\lambda_G}}(a,\theta)
\end{aligned}$$

Where the second equality follows from the fact ϕ is posterior-linear. Therefore, for all $\bar{a}$, a', θ, and λ_G:

$$v(H^{\lambda_G}, \mu^{\lambda_G}_{\bar{a}}(\theta), \bar{a}(\theta), \theta) = v(G, \mu^{\lambda_G}_{\bar{a}}(\theta), \bar{a}(\theta), \theta) \geq v(G, \mu^{\lambda_G}_{\bar{a}}(\theta), a', \theta) = v(H^{\lambda_G}, \mu^{\lambda_G}_{\bar{a}}(\theta), a', \theta)$$

Where the inequality follows from the fact v is PC. This concludes the proof. □