

Contagious Ambiguity*

Mira Frick

Ryota Iijima

Daisuke Oyama

August 20, 2026

Abstract

We study the strategic impact of ambiguity through the channel of higher-order beliefs. We show that even small amounts of prior ambiguity about game payoffs can generate arbitrarily large amounts of higher-order ambiguity. This gives rise to a novel form of contagion: vanishingly small payoff ambiguity can select “secure” actions (e.g., non-participation) as the unique equilibrium outcome even when those actions are almost dominated. We highlight two main implications. First, classical robustness results under probabilistic uncertainty break down under ambiguity when players can deviate to secure actions. Second, if a designer can introduce small amounts of payoff ambiguity into a game, this can serve as a powerful tool for unique implementation.

*Frick: Princeton University (mfrick@princeton.edu); Iijima: Princeton University (riijima@princeton.edu); Oyama: University of Tokyo (oyama@e.u-tokyo.ac.jp). For helpful comments, we thank Roberto Corrao, Drew Fudenberg, Faruk Gul, Atsushi Kajii, Peter Klibanoff, Elliot Lipnowski, Stephen Morris, Sujoy Mukerji, Pietro Ortoleva, Ludovic Renou, Tomasz Strzalecki, Takashi Ui, Mu Zhang, as well as numerous seminar and conference audiences. We also gratefully acknowledge the financial support from Sloan Research Fellowships (Frick and Iijima) and from JSPS KAKENHI Grant 23K01303 (Oyama).

1 Introduction

In many economic environments, agents face ambiguity, i.e., are unable to form a precise probabilistic belief about fundamentals. For example, investors may have ambiguous beliefs about the profitability of an investment, or consumers may face ambiguity about the quality of a new technology. Many such settings naturally feature strategic externalities (e.g., through synergies or network effects). In this case, agents' behavior may depend not only on their first-order ambiguity (i.e., ambiguous beliefs about fundamentals) but also on their higher-order ambiguity (i.e., ambiguous beliefs about other agents' beliefs).

In this paper, we study the strategic impact of ambiguity through the channel of higher-order beliefs. We show that even small amounts of prior ambiguity about game payoffs can generate arbitrarily large amounts of higher-order ambiguity. This gives rise to a novel form of contagion: vanishingly small payoff ambiguity can select “secure” actions (e.g., non-participation) as the unique equilibrium outcome even when those actions are almost dominated. We highlight two main implications. First, classical robustness results under probabilistic uncertainty break down under ambiguity when players can deviate to secure actions. Second, if a designer can introduce small amounts of payoff ambiguity into a game, this can serve as a powerful tool for unique implementation. The following example illustrates the main ideas.

1.1 Illustrative Example

Consider a complete-information game with two players $i \in \{1, 2\}$ and binary action sets $A_i = \{0, 1\}$ for each i . Each player i 's payoff function is given by

$$u_i(a) = \mathbf{1}_{a_i=a_{-i}=1} - ca_i$$

for some $c \in (0, 1)$. Thus, a player who plays $a_i = 1$ always incurs a unilateral cost of c , and receives a benefit of 1 if and only if the other player also plays $a_{-i} = 1$. Action $a_i = 0$ is *safe*, yielding a payoff of 0 regardless of the other player's action. There are two strict Nash equilibria: $a = \mathbf{1}$ where both players choose action 1, and $a = \mathbf{0}$ where both players choose action 0. This stylized game captures coordination problems that are relevant in various economic settings. For example, action $a_i = 1$ might represent participation in a joint investment or coordinated attack, while $a_i = 0$ corresponds

to non-participation in these activities.

Observe that when $c = 0$, action $a_i = 1$ is weakly dominant for each player i . Thus, one might expect that when c is sufficiently small (i.e., $a_i = 1$ is “almost dominant”), equilibrium **1** survives against small amounts of uncertainty about game payoffs. Under probabilistic uncertainty, this point was formalized by [Kajii and Morris \(1997\)](#). They considered ε -elaborations, i.e., incomplete-information games with a common prior that assigns probability at least $1 - \varepsilon$ to both players being *normal* types, who know that their payoffs coincide with the original game payoffs u_i . Their analysis implies that as long as $c < 0.5$, **1** is *robust*, i.e., every ε -elaboration with small enough ε admits an equilibrium where **1** is played with high probability.

In this paper, we instead study the implications of small amounts of ambiguity about game payoffs. In the baseline model, we represent ambiguous beliefs using the multiple-prior approach and focus on maxmin expected utility preferences ([Gilboa and Schmeidler, 1989](#)). Thus, we allow for *ambiguous* ε -elaborations, i.e., incomplete-information games with a common *set* of priors, each of which assigns probability at least $1 - \varepsilon$ to both players being normal types. We find that **1** is not robust against such elaborations, because ambiguity makes **0** *contagious*: For every $\varepsilon > 0$, there is an ambiguous ε -elaboration where **0** is played with probability 1 in the unique equilibrium. This is the case even if c is arbitrarily close to 0.

To illustrate, consider an ambiguous elaboration of the following form. Each player i has a set of types $T_i = \mathbb{N} = \{0, 1, \dots\}$. All types $t_i \geq 1$ are normal with payoffs given by u_i , but type $t_i = 0$ has a payoff function under which action 0 is strictly dominant. Players share a common set of priors $\mathcal{P} = \{P^\lambda : \lambda \in [0, 1]\} \subseteq \Delta(T_1 \times T_2)$. Here, for each $\lambda \in [0, 1]$, the prior $P^\lambda \in \Delta(T_1 \times T_2)$ is given by

$$P^\lambda(t_1, t_2) = \begin{cases} \varepsilon(1 - \varepsilon)^n \lambda & \text{if } (t_1, t_2) = (n + 1, n) \text{ for some } n \in \mathbb{N}, \\ \varepsilon(1 - \varepsilon)^n (1 - \lambda) & \text{if } (t_1, t_2) = (n, n + 1) \text{ for some } n \in \mathbb{N}, \end{cases}$$

and $P^\lambda(t_1, t_2) = 0$ otherwise. Thus, all priors P^λ assign probability $1 - \varepsilon$ to both players having payoffs u_i (i.e., to the event $(t_1, t_2) \geq (1, 1)$), so there is only a small amount of prior ambiguity about game payoffs. However, there is significant ambiguity about the relative ranking of players’ types, as priors differ in the probabilities $P^\lambda(t_1 \geq t_2) = \lambda$ they assign to player 1’s type being higher than player 2’s. This type space can be viewed as an ambiguous analog of an email-game type space à la

Rubinstein (1989), where players share the single common prior $P^{\frac{1}{2}}$.

Each player i knows her own type t_i and uses this information to form posteriors about her opponent's type. There are many belief-updating rules under ambiguity. For concreteness, assume for now that players perform full Bayesian updating (Pires, 2002): i 's set of posteriors at type t_i is given by the Bayesian updates $\{P^\lambda(\cdot | t_i) : \lambda \in [0, 1]\}$ of all the priors conditional on t_i . In particular, every normal type $t_i = n \geq 1$ believes that the opponent's type t_{-i} is either $n - 1$ or $n + 1$, and perceives the relative probability to be fully ambiguous: For each $n \geq 1$,

$$\{P^\lambda(t_{-i} = n - 1 | t_i = n) = 1 - P^\lambda(t_{-i} = n + 1 | t_i = n) : \lambda \in [0, 1]\} = [0, 1].$$

Note that all normal types know they are normal. Moreover, all types $t_i \geq 2$ also know that their opponent is normal, and thus have no ambiguity about either their own or their opponent's payoffs. Nevertheless, the fact that type $t_i = 1$ faces substantial ambiguity about her opponent's payoffs (which in turn induces substantial higher-order ambiguity for types $t_i \geq 2$) is enough to trigger the following contagion: First, each type $t_i = 1$ chooses action 0. This is because action 0 yields a safe payoff of 0, whereas the payoff to playing action 1 is $-c$ under i 's worst-case posterior that assigns probability 1 to $t_{-i} = 0$ for whom action 0 is strictly dominant. As a result, the payoff to action 1 is also $-c$ under the worst-case posterior of each type $t_i = 2$ (which assigns probability 1 to $t_{-i} = 1$), and hence $t_i = 2$ plays action 0. Inductively, all types must play action 0 in the unique equilibrium.¹

The above contagion suggests that ambiguity has dual implications for two important questions in economic theory. First, adopting a robust-prediction perspective à la Kajii and Morris (1997) and a large follow-up literature (see Section 4.1), one can ask which equilibria of the complete-information game are robust to small amounts of *exogenous* payoff uncertainty. From this perspective, the implication is that considering payoff ambiguity can greatly limit robust predictions relative to the case of probabilistic uncertainty: Even equilibria in almost dominant actions (e.g., **1** when $c \approx 0$) can fail to be robust, as a small amount of payoff ambiguity can, through its effect on higher-order beliefs, lead players to opt for safe actions. For example, even a very profitable joint venture may fail to attract investors, as vanishingly small prior ambiguity about its profitability may trigger a contagion towards non-investment.

¹Indeed, more strongly, 0 is uniquely rationalizable for all types.

Second, a more recent literature adopts a design perspective and asks how a principal can *endogenously* inject small amounts of payoff uncertainty into a game with the goal of implementing some target action profile (e.g., [Bergemann and Morris, 2019](#); [Halac, 2025](#)). From this perspective, the above suggests that ambiguity can be a useful tool. Specifically, suppose the principal seeks to implement the safe action profile $\mathbf{0}$ (e.g., prevent a coordinated attack) using bilateral contracts that pay each player i a subsidy $w_i \geq 0$ if i chooses $a_i = 0$. If the principal’s goal is partial implementation, i.e., to implement $\mathbf{0}$ as *an* equilibrium, no subsidies are needed. However, it is natural for the principal to be concerned about strategic uncertainty, i.e., that players may coordinate on an undesirable equilibrium, and hence seek to implement $\mathbf{0}$ as the *unique* equilibrium. If the principal can only offer deterministic payment schemes or introduce probabilistic uncertainty about players’ subsidies, uniquely implementing $\mathbf{0}$ is in general more costly than partial implementation.² However, suppose the principal can design *ambiguous payment schemes*, modeled as a set of distributions over subsidies (w_i). This generates higher-order ambiguity, where each type of player i knows her own realized subsidy but faces ambiguity about the other player’s subsidy, their beliefs about i ’s subsidy, etc. We will see that ambiguous payment schemes reduce the unique-implementation cost of $\mathbf{0}$ to (arbitrarily close to) zero, i.e., allow the principal to eliminate strategic uncertainty for free. The idea is to introduce vanishingly small prior ambiguity about the possibility that some players receive subsidies that make $\mathbf{0}$ dominant, and to exploit this to create a contagion towards $\mathbf{0}$ as above.

1.2 Overview of Results

Moving beyond this example, in the following we analyze general games with finitely many players $i \in I$ and finite action sets A_i for each i . [Section 2](#) models players’ higher-order ambiguity via an ambiguous type space, which specifies a countable set of types T_i for each player, along with a compact set $\mathcal{P}$ of common prior beliefs over type profiles. Each player i of type t_i forms a set of posterior beliefs $\mathcal{P}_{t_i}$ over opponent types. We allow for fairly general belief-updating rules, including full Bayesian updating (as in the example), maximum-likelihood updating ([Gilboa and Schmeidler, 1993](#)), and many other updating rules proposed in the literature.

[Section 3](#) presents the main results. Our first main result ([Theorem 1](#)) highlights

²Specifically, the unique implementation cost of $\mathbf{0}$ under deterministic (respectively, stochastic but unambiguous) payment schemes is $1 - c$ (respectively, $\max\{1 - 2c, 0\}$).

a key difference between higher-order beliefs under ambiguity vs. probabilistic uncertainty. In unambiguous type spaces with a common singleton prior $\mathcal{P} = \{P\}$, [Kajii and Morris’s \(1997\)](#) critical path theorem implies that if players have high prior confidence in some event E , this limits how much higher-order uncertainty about E there can be: That is, if $P(E) \approx 1$, then for all not too large p , the probability that players have common p -belief ([Monderer and Samet, 1989](#)) in E must also be close to one. In stark contrast, [Theorem 1](#) shows that under ambiguity, prior confidence imposes no limits on higher-order uncertainty: For arbitrarily small ε , we exhibit type spaces where even though *all* prior beliefs $P \in \mathcal{P}$ assign probability at least $1 - \varepsilon$ to an event E , common ε -belief of E fails at all types. Here, we use a natural extension of common p -belief to ambiguous type spaces ([Ganguli, 2007](#)).

Second, we build on [Theorem 1](#) to study contagion in games with small amounts of payoff ambiguity and maxmin expected utility preferences ([Section 5.1](#) considers more general ambiguity preferences). Generalizing the illustrative example, we find that ambiguity creates a new force for contagion: Suppose an action profile a^* is *secure*, i.e., each player i ’s worst-case payoff to playing action a_i^* when opponents play an arbitrary a_{-i} is better than the best payoff she can achieve by deviating to any $a_i \neq a_i^*$ when opponents stick with a_{-i}^* . Then [Theorem 2](#) shows that a^* is contagious under ambiguous elaborations. Any strict Nash equilibrium in safe actions (e.g., $\mathbf{0}$ in the example) is secure, but as we discuss, secure equilibria also arise in natural coordination games that do not feature safe actions. The contagion of secure equilibria is driven by the higher-order-belief effects of ambiguity in [Theorem 1](#): Even when all priors put probability $1 - \varepsilon$ on the original game payoffs and only probability ε on “crazy” types for whom a_i^* is strictly dominant, players may face substantial higher-order ambiguity about each other’s beliefs regarding the probability of crazy types. When a^* is secure, then as in the example, such higher-order ambiguity iteratively forces all types of all players to play a^* .

[Section 4](#) illustrates the implications of these results for the dual questions highlighted in the example. First, [Section 4.1](#) considers the robustness of equilibrium predictions to small exogenous payoff uncertainty. Under probabilistic uncertainty, the literature has identified various sufficient conditions for robustness that formalize the central idea that equilibria in “almost dominant” actions are robust. For instance, [Kajii and Morris \(1997\)](#) show that an equilibrium $\hat{a}$ is robust if it is p -dominant for some $p < 1/|I|$ (i.e., each $\hat{a}_i$ is a strict best-response to any opponent strategy that

puts probability at least p on $\hat{a}_{-i}$). In contrast, our analysis implies that these conditions do not guarantee robustness to small payoff ambiguity. For instance, as the example illustrates, even if $\hat{a}$ is p -dominant for arbitrarily small $p > 0$, it can be eliminated by small amounts of payoff ambiguity, because by Theorem 2 such payoff ambiguity may select a secure $a^* \neq \hat{a}$ as the unique equilibrium. We also discuss how to combine p -dominance and security to obtain stronger sufficient conditions that ensure robustness to small payoff ambiguity.

Second, Section 4.2 considers how a principal who can endogenously introduce small payoff ambiguity into a game can exploit this for unique implementation. Building on Theorem 2, we highlight two cases in which such ambiguity can reduce the unique implementation cost of a target action profile a^* all the way to its partial implementation cost. First, generalizing the example, Corollary 3 shows this is the case if a^* is sufficiently close to secure. Second, Corollary 4 shows that, when combined with ambiguity, a simple departure from bilateral contracts—*insurance contracts*—allows the principal to uniquely implement *any* action profile a^* at no additional cost relative to partial implementation. In contrast, in both these cases, if the principal can only introduce probabilistic payoff uncertainty, the unique implementation cost can exceed the partial implementation cost by a significant margin.

Finally, while all aforementioned results apply under general belief-updating rules, Section 3.3 shows that some updating rules generate even more extreme departures from the case of probabilistic uncertainty. Indeed, we provide a sense in which updating rules that are less conservative (i.e., induce smaller sets of posteriors) can magnify the higher-order belief effects of ambiguity. As an extreme example, consider maximum-likelihood updating, which refines full Bayesian updating by restricting each type t_i 's set of posteriors $\mathcal{P}_{t_i}$ to the Bayesian updates of only those priors under which t_i is most likely. In this case, we derive a starker failure of the critical path theorem: even a less demanding notion of common belief that we term *common ε -justifiability* fails for arbitrarily small ε . Consequently, under maximum-likelihood updating, *any* strict Nash equilibrium is contagious under ambiguous elaborations. Intermediate results hold for updating rules in between full Bayesian and maximum-likelihood updating.

1.3 Related Literature

We contribute to the literature on ambiguous beliefs in strategic environments. Many papers study the impact of ambiguous beliefs in specific classes of games, without an explicit focus on higher-order beliefs.³ Epstein and Wang (1996) and Ahn (2007) provide foundations for ambiguous type spaces based on hierarchies of ambiguous beliefs. Kajii and Ui (2005) formulate the equilibrium concept we use in our analysis: each player’s strategy is required to be worst-case optimal against opponents’ strategies given interim beliefs updated from a common set of priors.⁴ Some recent papers focus on different aspects of higher-order ambiguity. Huo, Pedroni, and Pei (2024) and Cerreia-Vioglio, Corrao, and Lanzani (2022) study how ambiguous beliefs affect the unique equilibrium in beauty-contest games à la Morris and Shin (2002). Ui (2025) studies binary-action global games with signals of ambiguous precision and full Bayesian updating. He shows that, unlike low-quality unambiguous signals, more ambiguous signal technologies increase the range of signals where a safe action is played, and that sufficiently ambiguous signal technologies can induce an action that is both safe and ex-ante risk-dominant to be played at all signals in the unique equilibrium.

We instead focus on contrasting the higher-order belief effects of small payoff ambiguity with those under probabilistic payoff uncertainty. Under probabilistic uncertainty, Rubinstein (1989) provides the seminal observation that high prior confidence in an event E need not guarantee approximate common certainty of E , and as a result, small payoff uncertainty can eliminate some equilibria of a complete-information game by uniquely selecting another equilibrium. Qualifying this contagion result, Kajii and Morris (1997) bound the amount of higher-order uncertainty that can be generated by small prior uncertainty, and use this to show that small payoff uncertainty cannot eliminate equilibria that are p -dominant for small enough p . In contrast, we show that small prior payoff ambiguity can generate arbitrarily large higher-order ambiguity, and consequently, gives rise to a more extreme form of contagion than probabilistic uncertainty: even equilibria that are arbitrarily close to dominant can be eliminated in favor of a secure equilibrium.⁵

³See, e.g., Salo and Weber (1995); Ellis (2016); Chen (2026) in the context of auctions, voting, and social learning.

⁴See, e.g., Billot, Chateauneuf, Gilboa, and Tallon (2000); Kajii and Ui (2009); Cerreia-Vioglio, Corrao, and Lanzani (2022) for characterizations of the common-prior assumption under ambiguity.

⁵It is worth contrasting this finding with Weinstein and Yildiz (2007), who instead adopt an interim perspective on robustness. Applied to the complete-information game in the example (rep-

In highlighting the role of secure actions (e.g., non-participation) as a source of contagion in games, we relate to an extensive literature (going back to [Dow and Werlang, 1992](#)) that proposes ambiguity as a source of market non-participation (see, e.g., [Billot, Mukerji, and Tallon, 2020](#), for a survey). While this literature mostly focuses on non-negligible payoff ambiguity and abstracts away from strategic externalities, our analysis suggests that if one accounts for strategic externalities, even small payoff ambiguity can have a large impact through the channel of higher-order beliefs.

Our analysis complements [Oyama and Tercieux \(2010, 2012\)](#), who consider elaborations with probabilistic but non-common priors and show that any strict Nash equilibrium is contagious with respect to such elaborations.⁶ We highlight that ambiguity creates an alternative force for contagion that applies even when the common-prior assumption is maintained. However, reflecting the difference in mechanism, under ambiguity, the extent of contagion depends on the conservativeness of players' updating rules and the security of actions, neither of which is relevant under (common or non-common) probabilistic priors.⁷ While players in our main model share a common set of priors, Section 5.1 shows that our results extend to settings where players share a single non-probabilistic or second-order prior; thus, the departures from probabilistic uncertainty we highlight do not rely on ambiguity being modeled via multiple priors. By maintaining the common-prior assumption, our results also apply to settings where, starting from a complete-information benchmark, a designer induces an ambiguous type space through contracts or information, as in the unique-implementation problem in Section 4.2. By contrast, in line with the agreeing-to-disagree logic ([Aumann, 1976](#)), inducing a non-common-prior type space would require the designer to manipulate agents' prior beliefs themselves, rather than

resented by a type t^{CK} in the universal type space), their analysis implies that, even if $c \approx 0$ so that $\mathbf{1}$ is almost dominant, one can slightly perturb (in the product topology) t^{CK} to some unambiguous type t^* for which $\mathbf{0}$ is uniquely rationalizable. However, [Kajii and Morris's \(1997\)](#) analysis implies that in any type space whose common prior assigns high enough probability to the original game payoffs, such types t^* must have very small prior probability. In contrast, we show that one can find ambiguous common-prior elaborations where $\mathbf{0}$ is uniquely rationalizable for *all* types. [Yokota \(2024\)](#) extends [Weinstein and Yildiz's \(2007\)](#) results to perturbations of ambiguous types.

⁶See also [Kajii and Morris \(1997\)](#) (pp. 1296-7), who note that the critical path theorem breaks down with non-common probabilistic priors.

⁷As Section 3.3 discusses, this difference relates to the novel feature under ambiguity that one must distinguish higher-order uncertainty in the sense of common belief vs. common justifiability. Under general updating rules, the critical path theorem fails with respect to common belief but can hold with respect to common justifiability, and failures of common belief (resp. justifiability) are needed to generate contagion of secure (resp. arbitrary strict Nash) equilibria.

merely provide them with different information.

The unique-implementation application relates to a recent literature on the endogenous design of ambiguity. Several papers study how a designer can benefit from inducing ambiguity about agents’ own payoffs via ambiguous information, contracts, or mechanisms (e.g., di Tillio, Kos, and Messner, 2016; Kellner and Le Quement, 2018; Beauchêne, Li, and Li, 2019; Tang and Zhang, 2021; Dütting, Feldman, Peretz, and Samuelson, 2024; Cheng, Klibanoff, Mukerji, and Renou, 2024). In a multi-agent mechanism design setting, Bose and Renou (2014) show that inducing first-order ambiguity (about other agents’ payoff types) is a useful tool for partial implementation. In contrast, in our contracting setting, we highlight the channel of higher-order ambiguity (about other agents’ beliefs) and its usefulness for unique implementation.

Finally, how to model belief updating under ambiguity is a key methodological question (e.g., Machina and Siniscalchi, 2014). The literature has proposed and axiomatized many updating rules (e.g., Pires, 2002; Gilboa and Schmeidler, 1993; Hanany and Klibanoff, 2007; Epstein and Schneider, 2007; Gul and Pesendorfer, 2021; Cheng, 2022). Some recent work compares different updating rules, for example, in terms of their performance in single-agent learning problems (Frick, Iijima, and Ishii, 2025) or their evolutionary stability (Sadowski and Sarver, 2026). Our analysis in this paper accommodates a large class of updating rules and sheds light on their implications for higher-order beliefs. In particular, we highlight that less conservative alternatives to full Bayesian updating rules, such as maximum-likelihood updating and its variants, can magnify the higher-order belief effects of prior ambiguity, leading to more drastic strategic implications. Another well-known issue is that belief updating under ambiguity can give rise to dynamic inconsistency (e.g., Siniscalchi, 2009). While dynamic inconsistency is an important force in many applications of ambiguity to strategic settings, Section 5.3 shows that our main insights do not hinge on this channel by considering a dynamically consistent equilibrium formulation à la Hanany, Klibanoff, and Mukerji (2020).

2 Setting

2.1 Type Space and Belief Updating

Fix a finite set I of at least two agents. Agents’ higher-order ambiguity is captured by an (*ambiguous*) *type space* ($T = \prod_{i \in I} T_i, \mathcal{P}$), which specifies a countable set T_i of

types for each agent $i \in I$ and a compact set $\mathcal{P} \subseteq \Delta(T)$ of common prior distributions over type profiles.⁸ We assume that $P(\{t_i\} \times T_{-i}) > 0$ for all $t_i \in T_i$ and $P \in \mathcal{P}$, where $T_{-i} = \prod_{j \neq i} T_j$. The type space is unambiguous if $\mathcal{P} = \{P\}$ is a singleton.

Agents know their own type, so the ambiguity of agent i of type t_i is captured by a set of posterior beliefs $\mathcal{P}_{t_i} \subseteq \Delta(T_{-i})$ over opponent types. The literature has proposed many belief-updating rules under ambiguity. Perhaps the most commonly studied rule is **full Bayesian updating**: For each prior $P \in \mathcal{P}$, type t_i forms the Bayesian update $P(\cdot \mid t_i) \in \Delta(T_{-i})$ (i.e., $P(t_{-i} \mid t_i) = \frac{P(t_i, t_{-i})}{P(\{t_i\} \times T_{-i})}$ for all $t_{-i} \in T_{-i}$). Her set of posteriors is then

$$\mathcal{P}_{t_i}^{\text{FB}} = \{P(\cdot \mid t_i) : P \in \mathcal{P}\}.$$

This updating rule is quite conservative, as the agent does not use the information that her own type is t_i to eliminate any priors.⁹ Motivated by this, the literature has developed less conservative refinements of full Bayesian updating. For instance, under **maximum-likelihood updating**,

$$\mathcal{P}_{t_i}^{\text{ML}} = \{P(\cdot \mid t_i) : P \in \operatorname{argmax}_{P' \in \mathcal{P}} P'(\{t_i\} \times T_{-i})\}.$$

That is, after observing t_i , the agent restricts attention to the priors $P \in \mathcal{P}$ that predicted t_i with the highest probability and only forms the Bayesian updates of these priors.

To capture these and many other commonly studied updating rules, we allow for any updating rule that satisfies the following assumption:

Assumption 1. *For every $(T, \mathcal{P})$, $i \in I$, and $t_i \in T_i$, the set of posteriors $\mathcal{P}_{t_i}$ is non-empty, closed, and satisfies:*

(i). $\mathcal{P}_{t_i} \subseteq \mathcal{P}_{t_i}^{\text{FB}}$,

(ii). for any $P, P' \in \mathcal{P}$ with $\operatorname{marg}_{T_i} P = \operatorname{marg}_{T_i} P'$,

$$P(\cdot \mid t_i) \in \mathcal{P}_{t_i} \iff P'(\cdot \mid t_i) \in \mathcal{P}_{t_i}.$$

⁸For any topological space X , $\Delta(X)$ denotes the set of Borel probability measures on X , endowed with the induced topology of weak convergence. We endow T_i with the discrete topology.

⁹See, e.g., [Seidenfeld and Wasserman \(1993\)](#) and [Shishkin and Ortoleva \(2023\)](#) for conceptual and experimental critiques of full Bayesian updating.

Condition (i) requires agents' updating rules to (weakly) refine full Bayesian updating. Condition (ii) imposes an own-information invariance requirement on this refinement: If P and P' have the same marginals over T_i , then they generate the same distribution over agent i 's private information, so i 's realized type t_i cannot serve as a basis for eliminating one prior but not the other.¹⁰ Assumption 1 nests full Bayesian and maximum-likelihood updating, as well as natural hybrids whose sets of posteriors are in between $\mathcal{P}_{t_i}^{\text{FB}}$ and $\mathcal{P}_{t_i}^{\text{ML}}$ (e.g., Epstein and Schneider, 2007; Cheng, 2022). It also nests some updating rules, such as proxy updating (Gul and Pesendorfer, 2021), whose sets of posteriors need not be supersets of $\mathcal{P}_{t_i}^{\text{ML}}$.¹¹ We allow updating rules to vary across agents and types.

2.2 Common Belief

We will be interested in the impact of higher-order ambiguity on strategic interactions. Monderer and Samet (1989) introduced common p -belief as a natural quantification of the degree of higher-order uncertainty in strategic settings. We will make use of the following extension to ambiguous type spaces due to Ganguli (2007).¹² Fix a type space $(T, \mathcal{P})$ and a belief-updating rule. Consider any $p \in [0, 1)$ and product event $E = \prod_{i \in I} E_i \subseteq T$. Define the event that E is *individually p -believed* by

$$B^p(E) = \prod_{i \in I} B_i^p(E), \quad \text{where} \quad B_i^p(E) = \{t_i \in E_i : P_i(E_{-i}) \geq p \ \forall P_i \in \mathcal{P}_{t_i}\}. \quad (1)$$

Define the event that E is *commonly p -believed* by

$$C^p(E) = \bigcap_{k \in \mathbb{N}} (B^p)^k(E),$$

where $(B^p)^0(E) = E$ and $(B^p)^k(E) = B^p((B^p)^{k-1}(E))$ for all $k \geq 1$. Section 3.3 also introduces a novel relaxation of common p -belief that will play an important role in contrasting the strategic impact of different belief-updating rules.

¹⁰All our results remain valid if condition (ii) is relaxed to the requirement that $\mathcal{P}_{t_i}$ contain a nonempty subset $\tilde{\mathcal{P}}_{t_i}$ satisfying condition (ii).

¹¹Formally, proxy updating posits a proxy belief $\hat{P}_i \in \mathcal{P}$ for each $i \in I$, and the set of posteriors for each t_i is given by $\mathcal{P}_{t_i} = \{P(\cdot | t_i) : P \in \mathcal{P} \text{ and } \text{marg}_{T_i} P = \text{marg}_{T_i} \hat{P}_i\}$.

¹²Ganguli (2007) used this extension of common p -belief to generalize the agreeing-to-disagree result and the no-trade theorem to settings with ambiguity.

2.3 Ambiguous Incomplete-Information Games

We model strategic interactions under ambiguity by an ambiguous *incomplete-information game*: This consists of a type space $(T, \mathcal{P})$, along with a finite action set A_i and bounded payoff function $v_i : A \times T \rightarrow \mathbb{R}$ for each agent $i \in I$, where $A = \prod_{j \in I} A_j$. As usual, each $v_i(\cdot, t)$ is extended linearly to mixed action profiles in $\prod_{j \in I} \Delta(A_j)$. A strategy for agent i is a mapping $\sigma_i : T_i \rightarrow \Delta(A_i)$. We assume maxmin expected utility preferences à la Gilboa and Schmeidler (1989) (Section 5.1 considers more general ambiguity preferences). That is, given any updating rule satisfying Assumption 1, type t_i 's payoff to playing action a_i against opponent strategies $\sigma_{-i} = (\sigma_j)_{j \neq i}$ is¹³

$$V_i(a_i, \sigma_{-i}, t_i) = \min_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}).$$

We extend i 's payoffs linearly to mixed actions $\alpha_i \in \Delta(A_i)$, i.e.,

$$V_i(\alpha_i, \sigma_{-i}, t_i) = \sum_{a_i \in A_i} \alpha_i(a_i) V_i(a_i, \sigma_{-i}, t_i).$$

This implicitly assumes that a malevolent Nature chooses $P_i \in \mathcal{P}_{t_i}$ after i 's pure action has been realized, so that i cannot hedge against ambiguity by mixing. However, as Section 5.2 discusses, our main insights go through even if mixing is given hedging value by reversing this timing. An *equilibrium* is a strategy profile $\sigma^* = (\sigma_i^*)_{i \in I}$ such that for each $i \in I$ and $t_i \in T_i$,

$$\sigma_i^*(t_i) \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} V_i(\alpha_i, \sigma_{-i}^*, t_i).$$

Fixed-point arguments as in Kajii and Ui (2005) imply that an equilibrium exists. While this equilibrium formulation imposes interim optimality at each type t_i , Section 5.3 considers an alternative formulation based on ex-ante optimality.

¹³The minimum over $P_i \in \mathcal{P}_{t_i}$ is well-defined as $\mathcal{P}_{t_i}$ is compact: By Assumption 1, $\mathcal{P}_{t_i}$ is a closed subset of $\mathcal{P}_{t_i}^{\text{FB}}$ (which is compact by compactness of $\mathcal{P}$).

3 Main Results

3.1 Higher-Order Ambiguity vs. Probabilistic Uncertainty

Our first main result highlights a key difference between higher-order beliefs under ambiguity vs. probabilistic uncertainty. Since this result is purely epistemic, it does not rely on any of the assumptions in Section 2.3 about agents' payoffs and ambiguity preferences in an underlying game.

In unambiguous type spaces $(T, \{P\})$, a key insight due to [Kajii and Morris \(1997\)](#) is that if agents have high prior confidence in an event E , this limits how much higher-order uncertainty about E there can be. Indeed, while [Rubinstein \(1989\)](#) highlighted that even if $P(E)$ is close to 1, common p -belief of E can fail for large p , [Kajii and Morris's \(1997\)](#) critical path theorem shows that there is a lower bound on the probability of common p -belief at moderate p :

Critical Path Theorem ([Kajii and Morris, 1997](#)). *For every unambiguous type space $(T, \{P\})$, product event E , and $p \in (0, 1/|I|)$,*

$$P(C^p(E)) \geq 1 - (1 - P(E)) \frac{1 - p}{1 - |I|p}.$$

In particular, the critical path theorem implies that if $P(E) \approx 1$, then $P(C^p(E)) \approx 1$ for any $p < 1/|I|$. In contrast, our first main result shows a stark failure of the critical path theorem under ambiguity:

Theorem 1. *Fix any $\varepsilon \in (0, 1)$. There is a type space $(T, \mathcal{P})$ and product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \mathcal{P}$ but $C^\varepsilon(E) = \emptyset$ under all updating rules satisfying Assumption 1.*

That is, under ambiguity, prior confidence imposes no limits on higher-order uncertainty: Even if an event E has high probability under every prior belief $P \in \mathcal{P}$, common ε -belief of E can fail globally even for arbitrarily small ε . Observe that the event E in Theorem 1 is unambiguous, as all priors $P \in \mathcal{P}$ assign the same probability to E . What drives the failure of the critical path theorem is that there is prior ambiguity about the probabilities of certain subevents of E .

More specifically, the proof of Theorem 1 (Appendix B.1) adapts the construction in the illustrative example in Section 1.1 to general updating rules and more than two agents. As in the example, we make use of type spaces with ordered types t_i and

two key features: (i) there is only small ambiguity about the marginal distributions of individual agents’ types—indeed, the proof uses a type space where $\text{marg}_{T_i} P$ is the same for all $P \in \mathcal{P}$, so marginals are fully unambiguous; but (ii) there is significant ambiguity about the correlation across agents’ types, in particular, their relative rankings.¹⁴ The key insight behind Theorem 1 is that these two features can give rise to events E (e.g., the event that all agents’ types are above some cutoff, as in the example) that have (unambiguous) high prior probability but nevertheless are subject to arbitrarily large higher-order uncertainty. In contrast, by the critical path theorem, this is impossible if the uncertainty about the relative rankings of types is probabilistic, as in the seminal email-game type space à la Rubinstein (1989).

As prior work has highlighted (e.g., in the literatures on global games and information design), type spaces with rank uncertainty are relevant in a variety of economic contexts, from currency attacks, where speculators are uncertain about how their private signals about fundamentals compare with those of others, to organizations, where employees may be uncertain about relative pay rankings. While such rank uncertainty is typically modeled as probabilistic, Theorem 1 suggests that the higher-order belief implications can be very different if it is instead ambiguous.

It is worth noting that the type space $(T, \mathcal{P})$ and event E in Theorem 1 can be chosen uniformly across all updating rules satisfying Assumption 1. Moreover, as we discuss in Section 5.1, the fact that Theorem 1 uses a multiple-prior formulation of ambiguity is not important. Analogs of Theorem 1 can be derived when ambiguity is instead modeled via a common single second-order or non-additive prior.

3.2 Secure Actions and Contagion

Building on Theorem 1, we now explore how ambiguity magnifies the impact of small amounts of payoff uncertainty in games.

To capture games with small amounts of payoff uncertainty, we fix a complete-information game $G = (A, u)$, where there is common certainty that i ’s payoff function is $u_i : A \rightarrow \mathbb{R}$. We then consider *(ambiguous) ε -elaborations* of G , i.e., incomplete-

¹⁴In contrast to the full ambiguity about relative rankings in the example, for any given $\varepsilon \in (0, 1)$, the prior probabilities that each agent’s type is higher than her opponents’ can be bounded away from zero and one. Moreover, unlike in the example, T_i and $\mathcal{P}$ in the proof of Theorem 1 are finite.

information games $(T, \mathcal{P}, (A_i, v_i)_{i \in I})$ as in Section 2.3 such that

$$P(T^*) \geq 1 - \varepsilon, \forall P \in \mathcal{P}.$$

Here, $T^* = \prod_{i \in I} T_i^*$, where

$$T_i^* = \{t_i \in T_i : v_i(\cdot, (t_i, t_{-i})) = u_i(\cdot) \forall t_{-i} \in T_{-i}\}$$

denotes the set of i 's **normal types**, i.e., types whose payoff functions coincide with the complete-information payoffs u_i . When $|\mathcal{P}| = 1$, this corresponds to Kajii and Morris's (1997) notion of unambiguous ε -elaborations.

Remark 1. In an ε -elaboration, payoff uncertainty is small in two senses: First, players have high confidence that payoff functions are given by the profile $u = (u_i)_{i \in I}$, as all priors assign probability at least $1 - \varepsilon$ to this event. Second, even though we allow for ambiguous type spaces, the first point also implies that there is only *small prior ambiguity about game payoffs*: Indeed, for each $P \in \mathcal{P}$, consider the induced probability measure $\mu_P \in \Delta(\mathbb{R}^{A \times I})$ over game payoffs.¹⁵ Then the fact that $\mu_P(\{u\}) \geq 1 - \varepsilon$ for all $P \in \mathcal{P}$ implies that the total variation distance $d_{TV}(\mu_P, \mu_{P'})$ is at most ε for any $P, P' \in \mathcal{P}$. In contrast, it is important for our results that ε -elaborations do not limit the distance between priors in $\mathcal{P}$ when viewed as distributions over the entire type space; that is, there can be *large prior ambiguity about payoff-irrelevant events*.¹⁶ ▲

To highlight the additional strategic implications of small payoff ambiguity relative to the case of small probabilistic uncertainty, we focus on the classic question of contagion (Rubinstein, 1989). Call a Nash equilibrium a^* of G **contagious** if, for every $\varepsilon > 0$, there is an ε -elaboration of G whose unique equilibrium σ^* satisfies $\sigma_i^*(a_i^* | t_i) = 1$ for all $i \in I$, $t_i \in T_i$. That is, even if game G admits multiple equilibria, introducing arbitrarily small payoff uncertainty can select a^* as the unique equilibrium.¹⁷ As we will see in Section 4, understanding how ambiguity affects contagion has implications for two dual questions: First, which equilibrium predictions are

¹⁵That is, $\mu_P(B) = P(\{t \in T : (v_i(\cdot, t))_{i \in I} \in B\})$ for all measurable $B \subseteq \mathbb{R}^{A \times I}$.

¹⁶Indeed, Appendix E.2 formalizes a sense in which, if $(T, \mathcal{P})$ features small prior ambiguity about all events in T , then it satisfies an approximate version of the critical path theorem, and as a result, $(T, \mathcal{P})$ does not support the contagion in Theorem 2 below.

¹⁷Our results remain valid even if we strengthen the requirement that a^* is played with probability 1 in the unique equilibrium of the elaboration to the requirement that a^* is uniquely obtained via iterated elimination of dominated strategies.

robust when a cautious analyst perceives small exogenous ambiguity about players' payoffs (Section 4.1); second, which target actions a principal can uniquely implement when he can endogenously design such payoff ambiguity (Section 4.2).

Our second main result shows that ambiguity introduces a new force for contagion—security:

Definition 1. *Action profile a^* in G is **secure** if for every player i ,*

$$\min_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i}) > \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*). \quad (2)$$

That is, each player i 's worst-case payoff when she plays action a_i^* and other agents play an arbitrary a_{-i} is better than the best payoff she can achieve by deviating to any $a_i \neq a_i^*$ when other players stick with a_{-i}^* . Thus, security strengthens the strict Nash equilibrium requirement that $u_i(a_i^*, a_{-i}^*) > \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*)$ by introducing worst-case reasoning about opponents' behavior into the left-hand side, i.e., into each player i 's calculation of her payoffs to playing a_i^* . Many important games admit secure equilibria. One example are strict Nash equilibria a^* where each action a_i^* is **safe**, i.e., i 's payoff $u_i(a_i^*, a_{-i})$ is constant across opponent actions a_{-i} . Safe actions are natural in economic settings where there is an option of non-participation, as in the introductory example. However, secure equilibria also arise in important coordination games that do not feature safe actions:

Example 1 (Security without safety). *Suppose $A = X^I$ and $u_i(a) = b(a)\mathbf{1}_{\{a_i=a_j \forall j\}} - c(a_i)$, where $b(a), c(a_i) \geq 0$ for all $a \in A$. That is, players receive a benefit if they all play the same action (e.g., adopt the same technology), while each player unilaterally incurs some cost for each action choice. Then coordinating on the least costly action is secure, i.e., playing a^* with $a_i^* = x^*$ for all i where $\{x^*\} = \operatorname{argmin}_{x \in X} c(x)$. $\blacktriangle$*

Theorem 2. *Fix any complete-information game G . Under any updating rule satisfying Assumption 1, any secure action profile a^* in G is contagious.*

Theorem 2 implies that small payoff ambiguity has more extreme strategic implications than small probabilistic payoff uncertainty. To understand the gap, consider the introductory example, where $A_i = \{0, 1\}$ and $u_i(a) = \mathbf{1}_{a_i=a_{-i}=1} - ca_i$ for some $c \in (0, 1)$. There, $\mathbf{0}$ is contagious with respect to probabilistic ε -elaborations if and only if $c \geq \frac{1}{2}$, i.e., if and only if the other strict Nash equilibrium $\mathbf{1}$ is not

risk-dominant. In contrast, since action 0 is safe, Theorem 2 implies that under ambiguous elaborations, $\mathbf{0}$ is contagious for *all* $c > 0$ no matter how small, i.e., even if action 0 is almost dominated by 1. Likewise, while in general games G no if-and-only-if characterization of contagion under probabilistic uncertainty is known, Kajii and Morris’s (1997) analysis implies that a necessary condition for a^* to be contagious under probabilistic uncertainty is that no other equilibrium $\hat{a}$ is *p -dominant* for some $p < 1/|I|$ (i.e., $\hat{a}_i$ is a strict best-response under u_i to any opponent strategy that assigns probability at least p to $\hat{a}_{-i}$); see Section 4.1 for more discussion.¹⁸ In contrast, under ambiguity, Theorem 2 implies that a secure a^* is contagious even if another equilibrium $\hat{a}$ is p -dominant for arbitrarily small $p > 0$. Thus, vanishingly small payoff ambiguity can select a secure a^* over an equilibrium in almost dominant strategies.

The following example illustrates how the security of a^* allows one to exploit the failure of the critical path theorem in Theorem 1 to construct elaborations where a^* is the unique equilibrium. The proof of Theorem 2 (Appendix B.2) adapts the ideas in the example to general games with more than two players.

Example 2 (Illustration of Theorem 2). *Consider the introductory example. To illustrate why $a^* = \mathbf{0}$ is contagious under general updating rules, consider ε -elaborations where $a_i^* = 0$ is dominant for non-normal types. Then, in any equilibrium σ , $T^*(\mathbf{1}) := \{t : \sigma(\mathbf{1}|t) > 0\} \subseteq T^*$. Moreover, since $a_i^* = 0$ is safe, a type t_i is willing to choose action 1 in equilibrium only if **all posteriors** of t_i assign probability at least $p = c$ to $T^*(\mathbf{1})$; otherwise, the maxmin expected utility of playing action 1 is less than the payoff 0 to choosing 0.¹⁹ Thus, $T^*(\mathbf{1}) \subseteq B^p(T^*(\mathbf{1}))$ (i.e., $T^*(\mathbf{1})$ is p -evident), which (by standard arguments) implies $T^*(\mathbf{1}) \subseteq C^p(T^*)$. But by Theorem 1, the elaboration can be constructed in such a way that $C^p(T^*) = \emptyset$. Thus, $\sigma(\mathbf{0} | t) = 1$ for all t . $\blacktriangle$*

3.3 Role of Belief Updating

While Theorems 1–2 apply under general updating rules, some updating rules give rise to results that are even more extreme departures from probabilistic uncertainty.

¹⁸For generic binary-action supermodular games, Oyama and Takahashi (2020) provide an if-and-only-if condition for contagion under probabilistic uncertainty in terms of monotone potential maximization.

¹⁹More generally, if a^* is secure, a type t_i is willing to choose some action $a_i \neq a_i^*$ in equilibrium only if all of t_i ’s posteriors assign sufficiently high probability to opponents deviating from a_{-i}^* .

Indeed, we now illustrate that updating rules that are less conservative (i.e., induce smaller sets of posteriors) can magnify the contagious effects of ambiguity. For expositional simplicity, we focus on contrasting the predictions under the most conservative updating rule satisfying Assumption 1, full Bayesian updating, with those under the refinement of maximum-likelihood updating. Appendix E.1 provides intermediate results for a class of updating rules that are in between these two updating rules.

To contrast the higher-order belief effects of ambiguity across different updating rules, we introduce the following relaxation of p -belief. Define the event that E is *individually p -justifiable* by

$$\overline{B}^p(E) = \prod_{i \in I} \overline{B}_i^p(E), \quad \text{where } \overline{B}_i^p(E) = \{t_i \in E_i : \exists P_i \in \mathcal{P}_{t_i} \text{ s.t. } P_i(E_{-i}) \geq p\}.$$

Define the event that E is *commonly p -justifiable* by

$$\overline{C}^p(E) = \bigcap_{k \in \mathbb{N}} (\overline{B}^p)^k(E),$$

where $(\overline{B}^p)^0(E) = E$ and $(\overline{B}^p)^k(E) = \overline{B}^p((\overline{B}^p)^{k-1}(E))$ for all $k \geq 1$. Note that $\overline{B}_i^p(E)$ replaces the universal quantifier over posteriors $P_i \in \mathcal{P}_{t_i}$ in the p -belief operator $B_i^p(E)$ with an existential quantifier. In unambiguous type spaces, $B^p(E) = \overline{B}^p(E)$ and $C^p(E) = \overline{C}^p(E)$; more generally, $B^p(E) \subseteq \overline{B}^p(E)$, and hence $C^p(E) \subseteq \overline{C}^p(E)$ by standard monotonicity arguments. Thus, common p -justifiability is in general a weaker requirement than common p -belief.

Proposition 1.

1. Under full Bayesian updating, the critical path theorem holds under common justifiability.²⁰ For every type space $(T, \mathcal{P})$, product event E , $P \in \mathcal{P}$, and $p \in (0, 1/|I|)$, we have $P(\overline{C}^p(E)) \geq 1 - (1 - P(E))^{\frac{1-p}{1-|I|p}}$.
2. Under maximum-likelihood updating, Theorem 1 extends to common justifiability: For any $\varepsilon \in (0, 1)$, there is a type space $(T, \mathcal{P})$ and a product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \mathcal{P}$ but $\overline{C}^\varepsilon(E) = \emptyset$.

Thus, under full Bayesian updating, high prior confidence in an event E can coexist with failures of common ε -belief of E for arbitrarily small ε , but it at least guaran-

²⁰We use “under full Bayesian updating” to mean that for every $(T, \mathcal{P})$, $i \in I$, and $t_i \in T_i$, we have $\mathcal{P}_{t_i} = \mathcal{P}_{t_i}^{\text{FB}}$, with the analogous meaning for “under maximum-likelihood updating”.

tees common p -justifiability of E for moderate p . In contrast, maximum-likelihood updating exacerbates the higher-order belief effects of ambiguity: No amount of prior confidence in E is enough to even rule out failures of common ε -justifiability for arbitrarily small ε .

Example 3 (Illustration of Proposition 1). *Consider the introductory example. Under maximum-likelihood updating, every normal type $t_1 \geq 1$ of player 1 has the unique posterior $P^1(\cdot \mid t_1)$, as $P^1(\{t_1\} \times T_2) > P^\lambda(\{t_1\} \times T_2)$ for any $\lambda \neq 1$. Similarly, for every normal type $t_2 \geq 1$ of player 2, $\mathcal{P}_{t_2}^{\text{ML}} = \{P^0(\cdot \mid t_2)\}$.²¹ Thus, each normal type t_i 's unique posterior assigns probability one to the event that $t_{-i} = t_i - 1$. As a result, the set of normal type profiles $E = \{1, 2, \dots\}^2$ is not individually ε -justifiable at $t_i = 1$, and inductively, E is not commonly ε -justifiable at any type profile. $\blacktriangle$*

Under maximum-likelihood updating, the stronger failure of the critical path theorem in Proposition 1 also gives rise to an extreme form of contagion that does not rely on security. (For the intermediate class of updating rules in Appendix E.1, we instead provide a sufficient condition for contagion that retains a weaker form of security.)

Proposition 2. *Fix any complete-information game G . Under maximum-likelihood updating, any strict Nash equilibrium a^* of G is contagious.*

The following example illustrates the connection between failures of common justifiability and the contagion result for general strict Nash equilibria. Beyond the question of contagion, the contrast between Examples 2 and 4 (and the proofs of Theorem 2 and Proposition 2) highlights a broader methodological point: Under ambiguity, different belief operators— p -belief vs. p -justifiability—are needed to capture players' coordination incentives for secure vs. non-secure actions.

Example 4 (Illustration of Proposition 2). *Consider the non-secure action profile 1 in the introductory example. To see why this is contagious under maximum-likelihood updating, consider ε -elaborations where $a_i^* = 1$ is dominant for non-normal types, so $T^*(\mathbf{0}) := \{t : \sigma(\mathbf{0} \mid t) > 0\} \subseteq T^*$ in any equilibrium σ . Now, t_i is willing to choose action 0 in equilibrium only if **some posterior** of t_i assigns probability at least $p = 1 - c$ to $T^*(\mathbf{0})$; otherwise, the maxmin expected utility of playing action 1 is greater than the payoff 0 to choosing 0. Thus, in contrast with the requirement*

²¹However, crazy types $t_1 = 0$ (resp. $t_2 = 0$) use P^0 (resp. P^1) to form posteriors.

that $T^*(\mathbf{1}) \subseteq B^p(T^*(\mathbf{1}))$ in Example 2, we now need $T^*(\mathbf{0}) \subseteq \overline{B}^p(T^*(\mathbf{0}))$, and hence $T^*(\mathbf{0}) \subseteq \overline{C}^p(T^*)$. But under maximum-likelihood updating, Proposition 1 shows we can construct an elaboration where even $\overline{C}^p(T^*) = \emptyset$. Then, $\sigma(\mathbf{1} \mid t) = 1$ for all t . $\blacktriangle$

4 Implications

We now highlight the dual implications of our main results for robust equilibrium predictions and unique implementation in games. Throughout this section, for each type space $(T, \mathcal{P})$, we fix some arbitrary updating rules $(\mathcal{P}_{t_i})$ satisfying Assumption 1.

4.1 Equilibrium (Non-)Robustness

Consider the question of equilibrium robustness studied by Kajii and Morris (1997) and a large follow-up literature.²² An analyst views a complete-information game $G = (A, u)$ as only an approximation of the true environment, allowing for the possibility that players face small exogenous payoff uncertainty. Crucially, the analyst is cautious with respect to such payoff uncertainty and takes a robust approach to equilibrium predictions: for an equilibrium of G to be robust, he requires that it remain a valid prediction under *every* ε -elaboration with sufficiently small ε .

Formally, extending Kajii and Morris's (1997) definition to allow for payoff ambiguity, call an equilibrium $\hat{a}$ of G **robust** if for every $\gamma < 1$, there is $\bar{\varepsilon} > 0$ such that every ε -elaboration of G with $\varepsilon \leq \bar{\varepsilon}$ admits an equilibrium σ with $\sum_{t \in T} P(t) \sigma(\hat{a} \mid t) \geq \gamma$ for all $P \in \mathcal{P}$; that is, in all close enough elaborations of G , $\hat{a}$ is played with high probability in some equilibrium.²³ For $\hat{a}$ to be robust, it is necessary that no other equilibrium a^* of G is contagious. Thus, Theorem 2 immediately implies the following:

Corollary 1. *Fix any complete-information game G . An equilibrium $\hat{a}$ of G is not robust if there is some secure action profile $a^* \neq \hat{a}$.*

The literature has developed various sufficient conditions for an equilibrium $\hat{a}$ to be robust to probabilistic payoff uncertainty. In particular, Kajii and Morris (1997)

²²See, e.g., Ui (2001); Morris and Ui (2005); Oyama and Tercieux (2009); Chassang and Takahashi (2011); Oyama and Takahashi (2020); Pei and Strulovici (2025), and the survey by Kajii and Morris (2020a,b). More broadly, this question also underlies the large literature on robust mechanism design (for a survey, see Bergemann and Morris, 2013).

²³Note that since we are fixing some updating rules $(\mathcal{P}_{t_i})$ for each type space $(T, \mathcal{P})$, this robustness notion is weaker than one that would require robustness under all possible updating rules satisfying Assumption 1. This makes the implications we state below for *non*-robustness stronger.

show that this is the case if $\hat{a}$ is p -dominant for $p < 1/|I|$, which reduces to the requirement that $\hat{a}$ is risk-dominant in (symmetric) 2×2 coordination games; follow-up work has introduced other sufficient conditions that also reduce to risk-dominance in 2×2 coordination games. The key implication of Corollary 1 is that none of these generalizations of risk-dominance is sufficient to guarantee robustness to payoff ambiguity: As we saw in the introductory 2×2 game, even if $\hat{a}$ is p -dominant for $p \approx 0$, this does not rule out the presence of other secure equilibria.

At the same time, our analysis also yields sufficient conditions for robustness to payoff ambiguity by appropriately combining p -dominance with security. Call $\hat{a} \in A$ *p -secure* if for all i ,

$$\min_{a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i}) > \max_{a_i \neq \hat{a}_i} u_i(a_i, \alpha_{-i}), \quad \forall \alpha_{-i} \in \Delta(A_{-i}) \text{ with } \alpha_{-i}(\hat{a}_{-i}) \geq p. \quad (3)$$

Note that replacing the left-hand side of (3) with $u_i(\hat{a}_i, \alpha_{-i})$ yields p -dominance. Thus, analogous to how security strengthens strict Nash, p -security strengthens p -dominance by incorporating worst-case reasoning about opponents' behavior into each i 's calculation of her payoffs to playing $\hat{a}_i$. Likewise, (3) strengthens the security requirement $\min_{a_{-i} \in A_{-i}} u_i(\hat{a}_i, a_{-i}) > \max_{a_i \neq \hat{a}_i} u_i(a_i, \hat{a}_{-i})$ by incorporating p -dominance reasoning into the right-hand side. Any p -dominant $\hat{a}$ where $\hat{a}_i$ is safe for all i is p -secure.

Under full Bayesian updating, Proposition 1 can be used to show that p -security for $p < 1/|I|$ is sufficient for robustness. At the same time, since less conservative updating rules can exacerbate contagion, the corresponding conditions for robustness are in general more demanding. Indeed, under the extreme case of maximum-likelihood updating, Proposition 2 implies that no equilibrium is robust when there are multiple strict Nash equilibria, so no degree of p -security with $p > 0$ ensures robustness. More generally, Appendix E.1 considers a parametrized class of updating rules in between full Bayesian and maximum-likelihood: Now, p -secure equilibria for $p < \bar{p}$ are robust, where $\bar{p}$ is smaller the less conservative the updating rule.

Corollary 2.

1. *Under full Bayesian updating, any p -secure equilibrium $\hat{a}$ of G with $p < 1/|I|$ is robust.*
2. *Under maximum-likelihood updating, no game G with multiple strict Nash equilibria admits a robust equilibrium.*

4.2 Unique Implementation in Games

Consider a classic problem of contracting with externalities (for a survey, see [Halac, 2025](#)). Given some complete-information game $G = (A, u)$ describing players' interdependent incentives, a principal seeks to design payments to each player to implement some target action profile $a^* \in A$ as an equilibrium; crucially, to avoid coordination on another, less desirable equilibrium, the principal requires the equilibrium to be unique. Formally, the principal can offer each agent i a bilateral contract that pays $w_i \geq 0$ if and only if i plays action a_i^* , and the principal's goal is to choose the cheapest payment scheme that induces a^* as the unique equilibrium.

Earlier work studied this unique implementation problem when the principal is restricted to deterministic payments w_i (e.g., [Segal, 2003](#)). More recently, [Halac, Lipnowski, and Rappoport \(2021\)](#) highlighted that the principal can do better by introducing probabilistic higher-order uncertainty via stochastic payments w_i .²⁴ Moving beyond this, we allow the principal to design payment schemes that involve higher-order ambiguity, in the sense that each agent i knows her own payment but faces ambiguity about others' payments, their beliefs about i 's payment, etc. Such ambiguity is natural in many contracting problems with externalities; for example, in organizations managers may provide employees with only vague information about coworkers' wages, and in the context of technology adoption, firms may offer personalized subsidies to new adopters, creating ambiguity among consumers about each other's offers. Based on the previous analysis, we show that ambiguous payments can lead to further significant reductions in the principal's unique implementation cost.

Principal's problem. The principal chooses an *ambiguous payment scheme* $\mathcal{C} = (T, \mathcal{P}, w)$: Here, $(T, \mathcal{P})$ is an ambiguous type space, and $w = (w_i)_{i \in I}$ specifies a bounded payment $w_i : T_i \rightarrow \mathbb{R}_+$ for each agent i as a function of her type t_i , which is paid if and only if i plays a_i^* . We identify each payment scheme $\mathcal{C} = (T, \mathcal{P}, w)$ with the induced ambiguous incomplete-information game with payoff functions

$$v_i(a, t) = u_i(a) + w_i(t_i) \mathbf{1}_{a_i = a_i^*}.$$

Denote by $\Sigma(\mathcal{C})$ the set of equilibria of $\mathcal{C}$. Slightly abusing notation, we denote by a^* the strategy profile σ of $\mathcal{C}$ such that $\sigma(a^* | t) = 1$ for all $t \in T$.

²⁴We focus on contractible actions (as in [Segal, 2003](#); [Sakovics and Steiner, 2012](#)), but our analysis can be extended to settings with moral hazard (as in [Halac, Lipnowski, and Rappoport, 2021](#)).

The principal seeks to choose the cheapest payment scheme $\mathcal{C}$ that implements a^* as the unique equilibrium of $\mathcal{C}$: The *unique implementation cost* W^{a^*} is given by

$$W^{a^*} = \inf_{\mathcal{C}=(T,\mathcal{P},w)} \min_{P \in \mathcal{P}} \sum_{t \in T} P(t) \sum_{i \in I} w_i(t_i) \quad \text{subject to} \quad (4)$$

$$\{a^*\} = \Sigma(\mathcal{C}), \quad (\text{unique implementation})$$

$$\sum_{t \in T} P(t) \sum_{i \in I} w_i(t_i) = \sum_{t \in T} P'(t) \sum_{i \in I} w_i(t_i), \quad \forall P, P' \in \mathcal{P}. \quad (\text{consistency})$$

Deterministic payment schemes correspond to $|T| = |\mathcal{P}| = 1$; more generally, unambiguous payment schemes correspond to $|T| \geq 1$ but $|\mathcal{P}| = 1$, i.e., allow for probabilistic higher-order uncertainty. Let $W_{\text{prob}}^{a^*}$ denote the unique implementation cost when the principal is restricted to unambiguous payment schemes.

In allowing for ambiguous payment schemes, we impose a consistency condition as in di Tillio, Kos, and Messner (2016) and Dütting, Feldman, Peretz, and Samuelson (2024), which requires the principal to be indifferent among all distributions $P \in \mathcal{P}$. The interpretation is that this is necessary for the principal to be able to credibly introduce ambiguity about which of several payment schemes he is using; otherwise, agents might be able to narrow down the set of possible schemes by reasoning about the principal's incentives.

The implementation cost in (4) considers a $P \in \mathcal{P}$ that minimizes the principal's expected payment. An interpretation is that the principal himself knows which $P \in \mathcal{P}$ he uses, but keeps the choice of $P \in \mathcal{P}$ ambiguous for agents (e.g., by using vague language or by conditioning the choice of P on an event, such as the composition of an Ellsberg urn, whose realization is known to the principal but ambiguous to agents). However, by the consistency condition, (4) could equivalently consider $P \in \mathcal{P}$ that maximize the principal's expected payment, capturing a principal who commits to also not learning the realized $P \in \mathcal{P}$ (e.g., by delegating its choice to an intermediary) and has maxmin expected utility preferences.

Unique vs. partial implementation. To illustrate how ambiguous payment schemes can facilitate unique implementation, we contrast the unique implementation cost W^{a^*} with the *partial implementation cost* $W_{\text{partial}}^{a^*}$, i.e., the value of (4) when the unique implementation condition $\{a^*\} = \Sigma(\mathcal{C})$ is relaxed to the condition that $a^* \in \Sigma(\mathcal{C})$. For partial implementation, there is no benefit to using ambiguous payment schemes: $W_{\text{partial}}^{a^*}$ is achieved by a deterministic payment scheme with payments

$w_{i,\text{partial}}^* = \max_{a_i \in A_i} u_i(a_i, a_{-i}^*) - u_i(a^*)$ for each i (see Lemma C.1).

Under partial implementation, the principal accepts some strategic uncertainty about which equilibrium of $\mathcal{C}$ will be played. Consequently, partial implementation is in general strictly less costly than unique implementation. However, based on Theorem 2, we exhibit two settings in which ambiguous payment schemes bring the unique implementation cost all the way down to the partial implementation cost, i.e., allow the principal to fully eliminate strategic uncertainty for free.

Unique implementation of secure actions. First, consider target actions a^* that are sufficiently close to secure:

Corollary 3. *We have $W^{a^*} = W_{\text{partial}}^{a^*}$ if a^* satisfies*

$$\min_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i}) + w_{i,\text{partial}}^* \geq \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*) \quad \text{for all } i \in I. \quad (5)$$

Condition (5) requires that in the modified complete-information game where each agent receives the partial implementation subsidy $w_{i,\text{partial}}^*$ for playing a_i^* , a^* satisfies (a weak inequality version of) security. Clearly, (5) holds if a^* is secure in the original game G (in which case $w_{i,\text{partial}}^* = 0$ as a^* is already an equilibrium of G); (5) also holds if each a_i^* is safe, regardless of whether or not a^* is an equilibrium of G .

Corollary 3 highlights that for actions that satisfy condition (5), ambiguous payment schemes greatly facilitate unique implementation relative to unambiguous payment schemes. Indeed, even for secure actions, unique implementation using unambiguous payment schemes can be arbitrarily more costly than partial implementation:

Example 5 (Ambiguous vs. unambiguous payment schemes). *Consider an instance of the technology adoption game in Example 1: $A_i = \{0, 1\}$ and $u_i(a_i, a_{-i}) = b(a_i)\mathbf{1}_{a_i=a_{-i}} - ca_i$ for all $i = 1, 2$, where $b(1) > b(0) > c > 0$. Since $\mathbf{0}$ is secure, Corollary 3 implies that, under ambiguous payment schemes, $W^{\mathbf{0}} = W_{\text{partial}}^{\mathbf{0}} = 0$ for all b and c . In contrast, under unambiguous payment schemes, Appendix C.4 shows that $W_{\text{prob}}^{\mathbf{0}} = \max\{b(1) - b(0) - 2c, 0\}$. In particular, $W_{\text{prob}}^{\mathbf{0}} \rightarrow \infty$ as $b(1) - b(0) \rightarrow \infty$. $\blacktriangle$*

The proof of Corollary 3 (Appendix C.3) builds on Theorem 2 by designing an ambiguous payment scheme that, given condition (5), creates contagion toward a^* : All normal types receive a payment $w_{i,\text{partial}}^* + \delta$ with $\delta \approx 0$, whereas non-normal types receive a w_i that is large enough to make a_i^* a dominant action. Here, analogous to Theorems 1–2, players know their own payments but face ambiguity about their

type's ranking relative to other players, which generates higher-order ambiguity about other players' (beliefs about) payments. Relative to Theorem 2, additional care is required to ensure that the consistency condition is satisfied. However, the amount of payoff ambiguity introduced by the principal is again small, in that the probability of normal types is arbitrarily close to 1 under all $P \in \mathcal{P}$. Just as in Theorems 1–2, the same ambiguous payment scheme can be used for all updating rules satisfying Assumption 1, so that the principal does not need to know players' updating rules.²⁵

Unique implementation via insurance contracts. Second, we exhibit a limited and natural departure from bilateral contracts—insurance contracts—that, when combined with ambiguity, reduces the unique implementation cost of *any* target action to its partial implementation cost.

An insurance contract reimburses each player i for any loss caused by other players' deviations from the target action profile a^* provided i plays a_i^* , and thus makes a_i^* a safe action.²⁶ Formally, it specifies payments $w_i \geq 0$ for each i such that if i plays a_i^* , she receives a total subsidy of $u_i(a^*) - u_i(a_i^*, a_{-i}) + w_i$ for any opponent action profile a_{-i} , and hence a constant payoff of $u_i(a^*) + w_i$ regardless of a_{-i} ; if i does not play a_i^* , i receives no subsidy.²⁷ Allowing for higher-order ambiguity about (w_i) , an ***ambiguous insurance contract*** specifies a type space $(T, \mathcal{P})$ and bounded subsidies $w_i : T_i \rightarrow \mathbb{R}_+$ that induce the incomplete-information game with payoffs

$$v_i(a, t) = \begin{cases} u_i(a^*) + w_i(t_i) & \text{if } a_i = a_i^* \\ u_i(a) & \text{otherwise.} \end{cases}$$

Denote by $\bar{W}^{a^*}$ the unique implementation cost of a^* when in (4) the principal optimizes over ambiguous insurance contracts, and by $\bar{W}_{\text{partial}}^{a^*}$ the corresponding partial implementation cost. These problems reduce to the original bilateral contracting setting with modified game payoffs $\bar{u}_i$ given by $\bar{u}_i(a) = u_i(a^*)$ if $a_i = a_i^*$ and $\bar{u}_i(a) = u_i(a)$ otherwise. Since each agent's target action a_i^* is safe under this modified game, Corollary 3 immediately implies the following:

²⁵That said, if the principal knows players' updating rules, then $W^{a^*} = W_{\text{partial}}^{a^*}$ may hold for a larger class of target actions (e.g., for all a^* under maximum-likelihood updating, by Proposition 2).

²⁶This idea bears a broad resemblance to some schemes used in practice. For example, FDIC deposit insurance discourages bank runs by making non-withdrawal safe for insured depositors.

²⁷This formulation allows for negative off-path subsidies. However, the same result continues to hold if subsidies are replaced with $\max\{u_i(a^*) - u_i(a_i^*, a_{-i}) + w_i, 0\}$, ensuring limited liability.

Corollary 4. *For every $a^* \in A$, we have $\overline{W}^{a^*} = \overline{W}_{\text{partial}}^{a^*} = W_{\text{partial}}^{a^*}$.*

Thus, when combined with (even small amounts of) payoff ambiguity, insurance contracts are a highly effective tool for unique implementation. The combination with ambiguity is key, as the unique implementation cost $\overline{W}_{\text{prob}}^{a^*}$ under unambiguous insurance contracts in general strictly exceeds $W_{\text{partial}}^{a^*}$. For instance, in Example 5, $\overline{W}_{\text{prob}}^1 = \max\{2c - 2b(1) + b(0), 0\}$ (see Appendix C.4), which can be arbitrarily more costly than $W_{\text{partial}}^1 = 0$.²⁸

5 Discussion

5.1 More General Ambiguity Representations

Beyond maxmin expected utility, our main insights extend to more general classes of ambiguity-sensitive preferences. First, to capture players with a mix of ambiguity-averse and ambiguity-seeking tendencies, Online Appendix F.1 considers α -maxmin expected utility. We provide a condition for contagion that reduces to security under maxmin expected utility. Under the opposite extreme of maxmax preferences, the condition implies that a strict Nash equilibrium a^* is contagious if every action $a_i \neq a_i^*$ is safe. Thus, if agents are ambiguity-seeking, the presence of safe actions serves as a source of contagion for *non-safe* equilibria. However, as in the maxmin case, this again means that almost-dominated actions can be contagious.

Second, while we have followed the canonical approach of representing ambiguity by a set of priors, our main results extend to other standard representations of ambiguity based on a single second-order prior or non-additive prior. In particular, in Online Appendix F.2–F.3, we adapt ambiguous type spaces and elaborations to agents with smooth preferences (Klibanoff, Marinacci, and Mukerji, 2005) with a common second-order prior $\mu \in \Delta(\Delta(T))$, as well as to agents with Choquet expected utility preferences (Schmeidler, 1989) and a common non-additive Choquet capacity ν over

²⁸As the literature notes, fully general non-bilateral contracts can equalize unique- and partial-implementation costs even with deterministic payments, but there are various difficulties with using such contracts (e.g., Segal, 2003; Halac, 2025). Insurance contracts mitigate some of these concerns: (i) while this broader class may condition i 's payment on the details of opponents' actions, which may be very costly to verify, insurance contracts only condition on i 's action and payoff loss $u_i(a^*) - u_i(a_i^*, a_{-i})$; (ii) while the aforementioned deterministic contracts make each a_i^* strictly dominant via large payments if opponents deviate, insurance contracts merely make a_i^* safe, which requires less knowledge of the full payoff matrix and makes them less vulnerable to side contracts among agents.

T . We derive analogs of the failure of the critical path theorem in Theorem 1. We also show that secure action profiles a^* remain contagious under both these models.

5.2 Hedging against Ambiguity through Mixing

The formulation of incomplete-information games in Section 2.3 rules out hedging against ambiguity through mixing: When player i mixes, Nature selects a worst-case posterior $P_i \in \mathcal{P}_{t_i}$ after i 's pure action has been realized. An alternative formulation reverses this order, thereby giving mixing hedging value. That is, Nature selects a posterior after observing i 's mixed action but before the corresponding pure action is realized, so that i 's interim payoff from a mixed action $\alpha_i \in \Delta(A_i)$ is

$$\tilde{V}_i(\alpha_i, \sigma_{-i}, t_i) = \min_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) v_i(\alpha_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}).$$

In general, it is well known that giving mixing hedging value can substantially change equilibrium predictions.²⁹ However, our main insights go through under this alternative formulation. In particular, every strict Nash equilibrium in safe actions remains contagious, and the implications in Section 4 continue to hold whenever security is strengthened to safety. The reason is that the contagion argument for equilibria in safe actions a_i^* does not require Nature to choose different posteriors against different pure deviations from a_i^* . Rather, at each inductive step, the construction yields a single posterior $P_i \in \mathcal{P}_{t_i}$ for the type t_i under consideration under which *every* pure $a_i \neq a_i^*$ has expected payoff strictly below the constant payoff from the safe action a_i^* . As a result, this also implies that t_i 's worst-case expected payoff of every mixed action $\alpha_i \neq \delta_{a_i^*}$ is strictly lower than the payoff from a_i^* . Hence, $\delta_{a_i^*}$ is the unique surviving action for type t_i , allowing the inductive argument to proceed as before.³⁰

²⁹For example, in Dütting, Feldman, Peretz, and Samuelson (2024) the advantage of ambiguous contracts over unambiguous contracts disappears when the agent can hedge by mixing. See also Bose and Renou (2014); di Tillio, Kos, and Messner (2016).

³⁰However, this argument need not extend to secure a^* where a_i^* is not safe. The construction still yields a single posterior $P_i \in \mathcal{P}_{t_i}$ under which every pure $a_i \neq a_i^*$ has an expected payoff strictly below the *worst-case* payoff of a_i^* . But when a_i^* is not safe, this does not rule out that some mixed $\alpha_i \neq \delta_{a_i^*}$ with $a_i^* \in \text{supp } \alpha_i$ yields a better worst-case payoff than a_i^* .

5.3 Dynamically Consistent Agents

As is well known, belief-updating under ambiguity gives rise to dynamic inconsistency: A strategy that maximizes a player's interim worst-case payoff based on her set of posteriors $\mathcal{P}_{t_i}$ at each type need not in general maximize her ex-ante worst-case payoff based on the set of priors $\mathcal{P}$. However, dynamic inconsistency is not what drives the departures from probabilistic uncertainty that we have highlighted.

To shut down the dynamic inconsistency channel, we replace the equilibrium notion in Section 2.3 based on interim optimality with one based on ex-ante optimality.³¹ Define i 's ex-ante payoff to playing pure strategy $s_i : T_i \rightarrow A_i$ against opponents' mixed strategy profile σ_{-i} by

$$V_i(s_i, \sigma_{-i}) = \min_{P \in \mathcal{P}} \sum_{t \in T} P(t) v_i(s_i(t_i), \sigma_{-i}(t_{-i}), t_i, t_{-i}),$$

and extend V_i linearly to mixed strategies σ_i of player i .³² An *ex-ante equilibrium* is a strategy profile $\sigma^* = (\sigma_i^*)_{i \in I}$ such that, for each $i \in I$, $\sigma_i^* \in \operatorname{argmax}_{\sigma_i} V_i(\sigma_i, \sigma_{-i}^*)$. Thus, each player's contingent plan σ_i^* is a worst-case best-response against σ_{-i}^* based on the set of priors $\mathcal{P}$. Analogously to Section 3.2, for any complete-information game G , we call $a^* \in A$ *ex-ante contagious* if for every $\varepsilon > 0$, G admits an ε -elaboration whose unique ex-ante equilibrium satisfies $\sigma_i^*(a_i^*|t_i) = 1$ for all $i \in I$, $t_i \in T_i$.

Under the ex-ante formulation, ambiguity continues to magnify contagion relative to probabilistic uncertainty. Indeed, this formulation leads to an extreme contagion result, paralleling that under interim optimality and maximum-likelihood updating:

Proposition 3. *Any strict Nash equilibrium a^* of G is ex-ante contagious.*

Thus, there is a sense in which players' incentives can be more sensitive to small payoff ambiguity under ex-ante than under interim optimality. One novel channel under ex-ante optimality is that players' worst-case beliefs, and hence incentives, are sensitive to the cardinal specification of utilities in an elaboration. For example, if

³¹One interpretation of this solution concept is that players follow the dynamically consistent updating rule in Hanany and Klibanoff (2007). Posteriors under this updating rule are subsets of $\mathcal{P}_{t_i}^{\text{FB}}$, but this rule violates consequentialism, in that $\mathcal{P}_{t_i}$ can depend on opponent strategies and game payoffs. In contrast, our interim-optimal analysis imposed consequentialism, because we fixed players' updating rules across opponent strategies and games.

³²We embed mixed strategies $\sigma_i \in (\Delta(A_i))^{T_i}$ into mixtures $\Delta(A_i^{T_i})$ over pure strategies in the standard way.

utility functions take the form $v_i(\cdot, t_i, t_{-i}) = u_i(\cdot) + b_i(t_i)$ for some $b_i : T_i \rightarrow \mathbb{R}$, then $b_i(\cdot)$ affects i 's incentives under ex-ante optimization, but not under interim optimization.

6 Concluding Remarks

We conclude by discussing some directions for future research. First, whereas the unique-implementation application in Section 4.2 introduced higher-order ambiguity into complete-information games via ambiguous payment schemes, it is also worthwhile to consider information design settings where players share a common prior about a payoff-relevant fundamental θ and a designer can introduce higher-order ambiguity via ambiguous signals about θ . Several papers study the design of probabilistic signal structures that uniquely implement a target action profile via a process of contagion (e.g., [Moriya and Yamashita, 2020](#); [Hoshino, 2022](#); [Morris, Oyama, and Takahashi, 2024](#)).³³ Our analysis suggests that allowing for ambiguous signals may further facilitate unique implementation by expanding the class of target action profiles for which contagion can be induced.

Second, to illustrate the contrast with probabilistic uncertainty, we identified simple and economically interpretable sufficient conditions for contagion and robustness under ambiguous elaborations. However, it would be valuable to obtain sharper characterizations. As noted in Section 3.2, even under probabilistic elaborations, if-and-only-if conditions for contagion and robustness are not known for general games. A promising direction could be to analyze structured classes of games, such as binary-action supermodular games, for which [Oyama and Takahashi \(2020\)](#) obtain sharp characterizations under probabilistic uncertainty. Under ambiguity, such characterizations would additionally depend on the class of updating rules.

³³Contagion is also used in information design under worst-case equilibrium selection (e.g., [Li, Song, and Zhao, 2023](#); [Inostroza and Pavan, 2025](#)).

Appendix

A Preliminaries

A.1 Generalized Operators

We extend the ambiguous p -belief and p -justifiability operators from Sections 2.2 and 3.3 to more general $\mathbf{f}$ -operators (Morris and Shin, 2007; Oyama and Takahashi, 2020).

For each agent i , consider a function $f_i : 2^{I \setminus \{i\}} \rightarrow \mathbb{R}$ such that

$$f_i(\emptyset) < 0 \quad \text{and} \quad f_i(T) \leq f_i(S) \text{ for all } T \subseteq S \subseteq I \setminus \{i\}. \quad (6)$$

Let $\mathbf{f} = (f_i)_{i \in I}$. Throughout the appendix, we treat product events $E \subseteq T$ as profiles $E = (E_i)_{i \in I}$ of their projections onto each coordinate i . As such, the inclusion $(E_i)_{i \in I} \subseteq (E'_i)_{i \in I}$ means that $E_i \subseteq E'_i$ for each $i \in I$, the intersection $(E_i)_{i \in I} \cap (E'_i)_{i \in I}$ equals $(E_i \cap E'_i)_{i \in I}$, and $(E_i)_{i \in I} = \emptyset$ if and only if $E_i = \emptyset$ for all $i \in I$.

Given any type space $(T, \mathcal{P})$, any belief-updating rule satisfying Assumption 1, and any product event $E = (E_j)_{j \in I}$, define the *individual $\mathbf{f}$ -belief* and *common $\mathbf{f}$ -belief* events by

$$\begin{aligned} B_i^{f_i}(E) &= \left\{ t_i \in E_i : \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) f_i(\{j \neq i : t_j \in E_j\}) \geq 0 \quad \forall P_i \in \mathcal{P}_{t_i} \right\}, \\ B^{\mathbf{f}}(E) &= (B_i^{f_i}(E))_{i \in I}, \quad (B^{\mathbf{f}})^0(E) = E, \quad (B^{\mathbf{f}})^k(E) = (B_i^{f_i}((B^{\mathbf{f}})^{k-1}(E)))_{i \in I} \quad \forall k \geq 1, \\ C_i^{\mathbf{f}}(E) &= \bigcap_{k \in \mathbb{N}} (B^{\mathbf{f}})^k_i(E), \quad C^{\mathbf{f}}(E) = (C_i^{\mathbf{f}}(E))_{i \in I}. \end{aligned}$$

Define the *individual $\mathbf{f}$ -justifiability* and *common $\mathbf{f}$ -justifiability* events by

$$\begin{aligned} \overline{B}_i^{f_i}(E) &= \left\{ t_i \in E_i : \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) f_i(\{j \neq i : t_j \in E_j\}) \geq 0 \quad \text{for some } P_i \in \mathcal{P}_{t_i} \right\}, \\ \overline{B}^{\mathbf{f}}(E) &= (\overline{B}_i^{f_i}(E))_{i \in I}, \quad (\overline{B}^{\mathbf{f}})^0(E) = E, \quad (\overline{B}^{\mathbf{f}})^k(E) = (\overline{B}_i^{f_i}((\overline{B}^{\mathbf{f}})^{k-1}(E)))_{i \in I} \quad \forall k \geq 1, \\ \overline{C}_i^{\mathbf{f}}(E) &= \bigcap_{k \in \mathbb{N}} (\overline{B}^{\mathbf{f}})^k_i(E), \quad \overline{C}^{\mathbf{f}}(E) = (\overline{C}_i^{\mathbf{f}}(E))_{i \in I}. \end{aligned}$$

If each f_i is defined by $f_i(S) = \mathbf{1}_{\{S=I \setminus \{i\}\}} - p$ for all $S \subseteq I \setminus \{i\}$, these definitions reduce

to the individual and common p -belief and p -justifiability definitions in Sections 2.2 and 3.3.

The following lemma summarizes some basic properties of $\mathbf{f}$ -belief and $\mathbf{f}$ -justifiability that we will use throughout the appendix without explicit reference. We omit the proof, which follows standard arguments. Call a product event E **lower $\mathbf{f}$ -evident** (resp., **upper $\mathbf{f}$ -evident**) if $E_i \subseteq B_i^{f_i}(E)$ (resp., $E_i \subseteq \overline{B}_i^{f_i}(E)$) for all $i \in I$.

- Lemma A.1.**
1. If $E \subseteq E'$, then $B_i^{f_i}(E) \subseteq B_i^{f_i}(E')$ and $\overline{B}_i^{f_i}(E) \subseteq \overline{B}_i^{f_i}(E')$.
 2. If $(E^n)_{n \in \mathbb{N}}$ is a sequence of events that is weakly decreasing with respect to set inclusion, then $B_i^{f_i}(\bigcap_{n \in \mathbb{N}} E^n) = \bigcap_{n \in \mathbb{N}} B_i^{f_i}(E^n)$ and $\overline{B}_i^{f_i}(\bigcap_{n \in \mathbb{N}} E^n) = \bigcap_{n \in \mathbb{N}} \overline{B}_i^{f_i}(E^n)$.
 3. $C^{\mathbf{f}}(E)$ (resp. $\overline{C}^{\mathbf{f}}(E)$) is the largest lower (resp. upper) $\mathbf{f}$ -evident event contained in E .

A.2 Strategic Role of Common Belief and Justifiability

The following lemma extends to settings with ambiguity the well-known result that a p -dominant equilibrium of a complete-information game G can be played in an equilibrium of an elaboration of G whenever there is common p -belief in normal types. It also establishes an analogous relationship between p -secure equilibria (as defined in Section 4.1) and the weaker notion of common p -justifiability.

Lemma A.2. Fix any $p \in (0, 1)$.

1. Suppose $\hat{a}$ is a p -dominant equilibrium of G . Then for any ε -elaboration of G , there exists an equilibrium σ with $\sigma(\hat{a}|t) = 1$ at all $t \in C^p(T^*)$.
2. Suppose $\hat{a}$ is a p -secure equilibrium of G . Then for any ε -elaboration of G , there exists an equilibrium σ with $\sigma(\hat{a}|t) = 1$ for all $t \in \overline{C}^p(T^*)$.

Proof. For the first part, let $\hat{\Sigma}_i$ be the set of i 's strategies such that $\sigma_i(\hat{a}_i|t_i) = 1$ for all $t_i \in C_i^p(T^*)$. Since $C^p(T^*)$ is p -evident (in the sense of lower $\mathbf{f}$ -evidence defined above), i 's best response against any strategy profile in $\hat{\Sigma}_{-i}$ is in $\hat{\Sigma}_i$. Thus, by applying Kakutani's fixed-point theorem to the best-response correspondences restricted to strategy profiles with $\sigma_i \in \hat{\Sigma}_i$ for all i , we have an equilibrium $\hat{\sigma}$ such that $\hat{\sigma}(\hat{a}|t) = 1$ for all $t \in C^p(T^*)$. The proof of the second part follows an analogous argument, by considering the best-response correspondences restricted to strategy profiles where $\sigma(\hat{a}|t) = 1$ for all $t \in \overline{C}^p(T^*)$. $\square$

B Proofs for Section 3

B.1 Proof of Theorem 1

We prove the following generalization of Theorem 1 to **f**-operators.

Theorem B.1. *Fix any $\varepsilon \in (0, 1)$. There is a type space $(T, \mathcal{P})$ and product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \mathcal{P}$ but $C_i^{\mathbf{f}}(E) = \emptyset$ for all functions $\mathbf{f} = (f_i)_{i \in I}$ satisfying (6), all $i \in I$, and all updating rules satisfying Assumption 1.*

Proof. Pick $N \in \mathbb{N}$ large enough that $\frac{2}{N+1} \leq \varepsilon$. Let $T_i = \{0, 1, \dots, N\}$ for each $i \in I$. Let $\mathcal{P} = \{P^i : i \in I\}$, where each P^i is defined by $P^i(t) = \frac{1}{N+1}$ for any $t \in T$ such that $t_j \equiv t_i - 1$ for all $j \neq i$, where $\equiv$ denotes equality modulo $N + 1$.³⁴ Observe that $\text{marg}_{T_i} P^j$ is the same (viz., uniform on $\{0, 1, \dots, N\}$) for all j . Thus, by Assumption 1, $\mathcal{P}_{t_i} = \mathcal{P}_{t_i}^{\text{FB}}$ for every t_i .

Consider the product event $E = (E_i)_{i \in I}$ with $E_i = \{1, 2, \dots, N\}$ for each i . Then for all $i \in I$, $P^i(E) = \frac{N-1}{N+1} \geq 1 - \varepsilon$. At the same time, for each $i \in I$ and $n \in \{0, 1, \dots, N-1\}$, we have

$$P^i(\{t_j = n \forall j \neq i\} \mid t_i = n+1) = 1. \quad (7)$$

Take any functions $\mathbf{f} = (f_i)_{i \in I}$ satisfying (6). To show that $C_i^{\mathbf{f}}(E) = \emptyset$ for each $i \in I$, it suffices to prove inductively that for all $n \in \mathbb{N}$,

$$(B^{\mathbf{f}})_i^n(E) \subseteq \{n+1, \dots, N\} \text{ for each } i,$$

where we set $\{n+1, \dots, N\} = \emptyset$ for all $n \geq N$. The claim is immediate for $n = 0$, as $(B^{\mathbf{f}})^0(E) = E$. Suppose the claim holds for some n with $0 \leq n \leq N-1$. Consider any i . Then for $t_i = n+1$, (7) and the inductive hypothesis imply that

$$\sum_{t_{-i} \in T_{-i}} P^i(t_{-i} \mid t_i) f_i\left(\left\{j \neq i : t_j \in (B^{\mathbf{f}})_j^n(E)\right\}\right) = f_i(\emptyset) < 0.$$

Hence, $B_i^{f_i}((B^{\mathbf{f}})^n(E)) \subseteq \{n+2, \dots, N\}$. Thus, $(B^{\mathbf{f}})_i^{n+1}(E) \subseteq \{n+2, \dots, N\}$. This establishes the claim for any $n \leq N$ and, in particular, shows that $(B^{\mathbf{f}})_i^N(E) = \emptyset$ for all i . Then we also have $(B^{\mathbf{f}})_i^n(E) = \emptyset$ for all i and $n > N$. $\square$

³⁴Thus, P^i assigns equal probability to the $N+1$ type profiles $(t_i = N, t_{-i} = \mathbf{N}-1)$, $(t_i = N-1, t_{-i} = \mathbf{N}-2)$, $\dots$, $(t_i = 1, t_{-i} = \mathbf{0})$, $(t_i = 0, t_{-i} = \mathbf{N})$ and probability zero to all other t .

B.2 Proof of Theorem 2

Since a^* is secure, there is $\gamma < 1$ such that for all i , $a_i \neq a_i^*$, and $Q_i \in \Delta(A_{-i})$ with $Q_i(a_{-i}^*) \geq \gamma$, we have

$$\min_{a_{-i} \in A_{-i}} u_i(a_i^*, a_{-i}) > \sum_{a_{-i} \in A_{-i}} Q_i(a_{-i}) u_i(a_i, a_{-i}). \quad (8)$$

For each i , define f_i by

$$f_i(S) = \begin{cases} \gamma - 1 & \text{if } S = \emptyset \\ \gamma & \text{if } S \neq \emptyset. \end{cases}$$

Consider any $\varepsilon > 0$. By Theorem B.1, there is a type space $(T, \mathcal{P})$ and a product event E such that $P(E) \geq 1 - \varepsilon$ for all $P \in \mathcal{P}$ but $C_i^f(E) = \emptyset$ for each $i \in I$. Define a bounded payoff function $v_i : A \times T \rightarrow \mathbb{R}$ for each i such that $v_i(\cdot, (t_i, t_{-i})) = u_i(\cdot)$ for all $t_i \in E_i$ and $t_{-i} \in T_{-i}$ and such that a_i^* is strictly dominant for all $t_i \notin E_i$. Then $(T, \mathcal{P}, v)$ is an ε -elaboration of u .

Consider any equilibrium σ of this elaboration. Let $\hat{T}_i = \{t_i \in T_i : \sigma_i(a_i^* | t_i) < 1\}$ for each $i \in I$, and let $\hat{T} = (\hat{T}_i)_{i \in I}$. Note that $\sigma_i(a_i^* | t_i) = 1$ for each $t_i \notin E_i$, and thus $\hat{T} \subseteq E$. Below we show that $\hat{T}_i = \emptyset$ for each $i \in I$, so that $\sigma(a^* | t) = 1$ for all $t \in T$.

For each $i \in I$ and $t_i \in \hat{T}_i$, there is $a_i \neq a_i^*$ with $\min_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i}} P_i(t_{-i}) u_i(a_i, \sigma_{-i}(t_{-i})) \geq \min_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i}} P_i(t_{-i}) u_i(a_i^*, \sigma_{-i}(t_{-i}))$. By (8), this implies $P_i((T_j \setminus \hat{T}_j)_{j \neq i}) \leq \gamma$ for all $P_i \in \mathcal{P}_{t_i}$, and hence $t_i \in B_i^{f_i}(\hat{T})$. Thus, $\hat{T}$ is lower $\mathbf{f}$ -evident. Since $\hat{T} \subseteq E$, we have $\hat{T} \subseteq C^f(E)$. Then $\hat{T}_i = \emptyset$ for each $i \in I$, as $C_i^f(E) = \emptyset$ by construction. $\square$

B.3 Proof of Proposition 1

B.3.1 First part

Take any $p < \frac{1}{|I|}$, type space $(T, \mathcal{P})$, product event E , and $P \in \mathcal{P}$. Let $\hat{B}^p(\cdot)$ and $\hat{C}^p(\cdot)$ be the p -belief operator and common p -belief operator under the unambiguous type space $(T, \{P\})$. Then $\hat{B}^p(\cdot) \subseteq \overline{B}^p(\cdot)$. By monotonicity and continuity of $\hat{B}^p(\cdot)$ and $\overline{B}^p(\cdot)$ (see Lemma A.1), standard arguments imply $\hat{C}^p(\cdot) \subseteq \overline{C}^p(\cdot)$. Thus, by Kajii and Morris's (1997) critical path theorem, we have $P(\overline{C}^p(E)) \geq P(\hat{C}^p(E)) \geq 1 - (1 - P(E))^{\frac{1-p}{1-|I|p}}$, as claimed.

B.3.2 Second part

We prove the following more general version based on $\mathbf{f}$ -operators.

Proposition B.1. *Consider maximum-likelihood updating. For all $\varepsilon \in (0, 1)$, there is a type space $(T, \mathcal{P})$ and a product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \mathcal{P}$ but $\overline{C}_i^{\mathbf{f}}(E) = \emptyset$ for all functions $\mathbf{f} = (f_i)_{i \in I}$ satisfying (6) and all $i \in I$.*

Proof. Fix any $\varepsilon \in (0, 1)$. Let $T = (T_i)_{i \in I}$, with $T_i = \mathbb{N}$ for all i . Choose $r \in (1 - \varepsilon, 1)$, and set

$$c = \frac{(1 - \varepsilon)(1 - r)}{r}, \quad d = \varepsilon - c > 0.$$

Let $\mathcal{P} = \{P^i : i \in I\}$, where each P^i is defined by

$$P^i(t) = \begin{cases} cr^n & \text{if } \exists n \in \mathbb{N} \text{ such that } t = \mathbf{n} + \mathbf{1}_i \\ d & \text{if } t = \mathbf{0} \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathbf{1}_i \in \mathbb{R}^I$ denotes the i th standard basis vector. Consider the product event $E = \{1, 2, \dots\}^I$. Then for all $i \in I$,

$$P^i(E) = P^i(\{\mathbf{n} + \mathbf{1}_i : n > 0\}) = 1 - \varepsilon.$$

At the same time, for each $n \in \mathbb{N}$ and $t_i = n + 1$, we have $\mathcal{P}_{t_i}^{\text{ML}} = \{P^i(\cdot \mid t_i)\}$. This is because for all $j \neq i$,

$$P^i(t_i) = cr^n > cr^{n+1} = P^j(t_i).$$

Moreover, $P^i(\cdot \mid t_i)$ satisfies (7).

Take any functions $\mathbf{f} = (f_i)_{i \in I}$ satisfying (6). To show that $\overline{C}_i^{\mathbf{f}}(E) = \emptyset$ for each $i \in I$, it suffices to prove inductively that for all $n \in \mathbb{N}$, $\left(\overline{B}^{\mathbf{f}}\right)_i^n(E) \subseteq \{n+1, n+2, \dots\}$ for each i . The claim is immediate for $n = 0$, as $\left(\overline{B}^{\mathbf{f}}\right)^0(E) = E$. Suppose the claim holds for some $n \geq 0$. Consider any i . Then for $t_i = n + 1$, (7) and the inductive hypothesis imply that

$$\sum_{t_{-i} \in T_{-i}} P^i(t_{-i} \mid t_i) f_i \left(\left\{ j \neq i : t_j \in \left(\overline{B}^{\mathbf{f}}\right)_j^n(E) \right\} \right) = f_i(\emptyset) < 0.$$

Hence, $\overline{B}_i^{f_i} \left(\left(\overline{B}^f \right)^n (E) \right) \subseteq \{n+2, n+3, \dots\}$, as required. $\square$

B.4 Proof of Proposition 2

Consider any strict Nash equilibrium a^* of G . Then we can pick $\gamma < 1$ such that for all i , $a_i \neq a_i^*$, and $Q_i \in \Delta(A_{-i})$ with $Q_i(a_{-i}^*) \geq \gamma$, we have

$$\sum_{a_{-i} \in A_{-i}} Q_i(a_{-i}) u_i(a_i^*, a_{-i}) > \sum_{a_{-i} \in A_{-i}} Q_i(a_{-i}) u_i(a_i, a_{-i}). \quad (9)$$

Define $\mathbf{f}$ as in the proof of Theorem 2. Construct a type space $(T, \mathcal{P})$ and product event E as in the proof of Proposition B.1 and use this to construct an ε -elaboration $(T, \mathcal{P}, v)$ as in the proof of Theorem 2. Let $\hat{T}_i = \{t_i : \sigma_i(a_i^* | t_i) < 1\}$ for each $i \in I$, and let $\hat{T} = (\hat{T}_i)_{i \in I}$. For each $i \in I$ and $t_i \in \hat{T}_i$, there exist $a_i \neq a_i^*$ and $P_i \in \mathcal{P}_{t_i}$ with $\sum_{t_{-i}} P_i(t_{-i}) u_i(a_i, \sigma_{-i}(t_{-i})) \geq \sum_{t_{-i}} P_i(t_{-i}) u_i(a_i^*, \sigma_{-i}(t_{-i}))$. By (9), this implies $P_i((T_j \setminus \hat{T}_j)_{j \neq i}) \leq \gamma$, and hence $t_i \in \overline{B}_i^{f_i}(\hat{T})$. Thus, $\hat{T}$ is upper $\mathbf{f}$ -evident. Since $\hat{T} \subseteq E$, we have $\hat{T} \subseteq \overline{C}^{\mathbf{f}}(E)$. Then $\hat{T}_i = \emptyset$ for each $i \in I$, as $\overline{C}_i^{\mathbf{f}}(E) = \emptyset$ by construction. Hence, $\sigma(a^* | t) = 1$ for all $t \in T$. $\square$

C Proofs for Section 4

C.1 Proof of Corollary 2

The first part follows from the second part of Lemma A.2 and the first part of Proposition 1. The second part follows from Proposition 2. $\square$

C.2 Partial Implementation

The following lemma shows that to partially implement a^* , the principal can without loss of optimality focus on deterministic payment schemes.

Lemma C.1. *For every $a^* \in A$, $W_{\text{partial}}^{a^*} = \sum_{i \in I} w_{i,\text{partial}}^*$, where $(w_{i,\text{partial}}^*)$ is the deterministic payment scheme given by $w_{i,\text{partial}}^* = \max_{a_i \in A_i} u_i(a_i, a_{-i}^*) - u_i(a_i^*)$ for all i .*

Proof. To see that $W_{\text{partial}}^{a^*} \leq \sum_{i \in I} w_{i,\text{partial}}^*$, note that the deterministic payment scheme where each agent i receives $w_{i,\text{partial}}^*$ for playing a_i^* induces a^* as an equilibrium.

To show that $W_{\text{partial}}^{a^*} \geq \sum_{i \in I} w_{i,\text{partial}}^*$, take any payment scheme $(T, \mathcal{P}, w)$ that partially implements a^* . Then for each type t_i , we must have $u_i(a^*) + w_i(t_i) \geq \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*)$, i.e., $w_i(t_i) \geq \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*) - u_i(a^*)$. Since $w_i(\cdot)$ is non-negative, we also have $w_i(t_i) \geq \max_{a_i \in A_i} u_i(a_i, a_{-i}^*) - u_i(a^*) = w_{i,\text{partial}}^*$. Thus, for every $P \in \mathcal{P}$, we have $\sum_{t \in T} P(t) \sum_{i \in I} w_i(t_i) \geq \sum_{i \in I} w_{i,\text{partial}}^*$. $\square$

C.3 Proof of Corollary 3

Fix any target action profile $a^* \in A$ satisfying condition (5). Since $W_{\text{partial}}^{a^*} \leq W^{a^*}$ by definition, it suffices to show that for any $\kappa > 0$, there is $\mathcal{C}$ satisfying ([unique implementation](#)) and ([consistency](#)) such that $\min_{P \in \mathcal{P}} \sum_{t \in T} P(t) \sum_{i \in I} w_i(t_i) \leq W_{\text{partial}}^{a^*} + \kappa$.

Let $w_i^* = \max_{a_i \in A_i} u_i(a_i, a_{-i}^*) - u_i(a^*) + \frac{\kappa}{2|I|}$. Then a^* is a strict Nash equilibrium of the game with modified utility functions $\tilde{u}_i(a) = u_i(a) + w_i^* \mathbf{1}_{a_i = a_i^*}$, and in this game a^* is secure by (5). Choose $K > 0$ large enough that $u_i(a_i^*, a_{-i}) + K > u_i(a_i, a_{-i})$ for all $i \in I$, $a_i \neq a_i^*$, and $a_{-i} \in A_{-i}$.

For any $\varepsilon \leq \frac{\kappa}{2|I|K}$, let $(T, \mathcal{P})$ and E be the type space and event used in the proof of Theorem B.1. The construction in that proof implies that $\sum_{i \in I} P(((T_i \setminus E_i), T_{-i})) = \sum_{i \in I} P'(((T_i \setminus E_i), T_{-i})) \leq |I|\varepsilon$ for all $P, P' \in \mathcal{P}$. Let $\mathcal{C} = (T, \mathcal{P}, w)$, where for each $i \in I$,

$$w_i(t_i) = \begin{cases} w_i^* & \text{for } t_i \in E_i \\ w_i^* + K & \text{for } t_i \notin E_i. \end{cases} \quad (10)$$

Note that the incomplete-information game induced by $\mathcal{C}$ is an ε -elaboration of $(A, \tilde{u})$ and that by the choice of K , a_i^* is strictly dominant for each $t_i \notin E_i$. Moreover, the proof of Theorem 2 shows that $\mathcal{C}$ has a^* as its unique equilibrium.

Since $\sum_{i \in I} w_i^* = W_{\text{partial}}^{a^*} + \kappa/2$, the total expected payment under each $P \in \mathcal{P}$ is

$$\sum_{t \in T} P(t) \sum_{i \in I} w_i(t_i) = \sum_{i \in I} (w_i^* + P(((T_i \setminus E_i), T_{-i}))K) \leq W_{\text{partial}}^{a^*} + \frac{\kappa}{2} + |I|\varepsilon K \leq W_{\text{partial}}^{a^*} + \kappa.$$

Moreover, by construction of $(T, \mathcal{P}, E)$, the total expected payment is the same for all $P \in \mathcal{P}$, so ([consistency](#)) holds. This proves that for any $\kappa > 0$, $W^{a^*} \leq W_{\text{partial}}^{a^*} + \kappa$. $\square$

C.4 Unambiguous Implementation

To determine the unique implementation cost under unambiguous payment schemes in Example 5, we show a more general result for binary-action supermodular games: Here, for each player $i \in I$, we have $A_i = \{0, 1\}$, and i 's payoff difference between action 1 vs. action 0, i.e., $d_i(S) = u_i(1, \mathbf{1}_S, \mathbf{0}_{I \setminus (S \cup \{i\})}) - u_i(0, \mathbf{1}_S, \mathbf{0}_{I \setminus (S \cup \{i\})})$, is non-decreasing with respect to set inclusion in the set S of opponents playing action 1.

The following result follows from arguments analogous to those in Morris, Oyama, and Takahashi (2022), so we relegate the proof to Online Appendix G. Let Π be the set of permutations of I . For $\gamma \in \Pi$, let $S_{-i}(\gamma)$ be the set of players that precede i in γ .

Proposition C.1. *In any binary-action supermodular game, we have*

$$W_{\text{prob}}^1 = \min_{\substack{z \in \mathbb{R}_+^I \\ \rho \in \Delta(\Pi)}} \left\{ \sum_{i \in I} z_i : z_i + \sum_{\gamma \in \Pi} \rho(\gamma) d_i(S_{-i}(\gamma)) \geq 0 \text{ for every } i \in I \right\}. \quad (11)$$

The inequality condition in (11) requires that, when each player i receives the deterministic payment z_i for playing action 1, action profile $\mathbf{1}$ satisfies a sequential obedience condition; in symmetric two-player games (as in Example 5), this condition is equivalent to $\mathbf{1}$ being weakly risk-dominant (see Oyama and Takahashi, 2020).

Applied to the game in Example 5 (and swapping the labels for actions 0 and 1), Proposition C.1 implies that $W_{\text{prob}}^0 = \max\{b(1) - b(0) - 2c, 0\}$, as making $\mathbf{0}$ weakly risk-dominant requires $z_i \geq 0.5(b(1) - b(0)) - c$ for each i . Likewise, to show that under unambiguous insurance contracts $\bar{W}_{\text{prob}}^1 = \max\{2c - 2b(1) + b(0), 0\}$, one can apply Proposition C.1 to the modified game where action 1 is safe with payoff $b(1) - c$.

D Proofs for Section 5

D.1 Proof of Proposition 3

Fix any strict Nash equilibrium a^* of G , and any $\varepsilon > 0$. We construct an ε -elaboration based on the type space used in the second part of Proposition 1, where $T_i = \mathbb{N}$ for all i and $\mathcal{P} = \{P^i : i \in I\}$ has the property that $P^i(t_i = 0) < P^j(t_i = 0)$ for all $i \neq j$.

Let $T_i^* = \{t_i \in T_i : t_i \geq 1\}$ be the set of i 's normal types. We specify the payoff

of the non-normal type $t_i = 0$ as

$$v_i(a, (0, t_{-i})) = \mathbf{1}_{a_i=a_i^*} + M_i,$$

where $M_i > 0$ is large enough that for all $j \neq i$,

$$P^i(t_i = 0)(M_i + 1) + (1 - P^i(t_i = 0)) \max_{a \in A} u_i(a) < P^j(t_i = 0)M_i + (1 - P^j(t_i = 0)) \min_{a \in A} u_i(a).$$

This inequality ensures that player i 's ex-ante worst-case payoff is always evaluated based on the prior P^i , i.e., for any strategy profile σ ,

$$\min_{P \in \mathcal{P}} \sum_{t \in T} P(t) v_i(\sigma(t), t) = \sum_{t \in T} P^i(t) v_i(\sigma(t), t).$$

Thus, against any opponent strategies σ_{-i} , i 's ex-ante best response

$$\sigma_i^* \in \operatorname{argmax}_{\sigma_i} \min_{P \in \mathcal{P}} \sum_{t \in T} P(t) v_i(\sigma_i(t_i), \sigma_{-i}(t_{-i}), t_i, t_{-i})$$

satisfies interim optimality for each $t_i \in T_i$ with respect to the singleton posterior belief $\{P^i(\cdot | t_i)\}$, i.e.,

$$\sigma_i^*(t_i) \in \operatorname{argmax}_{\alpha_i \in \Delta(A_i)} \sum_{a_i} \alpha_i(a_i) \sum_{t_{-i}} P^i(t_{-i} | t_i) v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}).$$

By construction, a_i^* is a strictly dominant action for $t_i = 0$. Then an analogous argument as in Proposition 2 implies that the unique ex-ante equilibrium σ^* satisfies $\sigma^*(a^* | t) = 1$ for all $t \in T$. $\square$

E Omitted Details

E.1 Role of Conservativeness in Belief Updating

To shed more light on the role of agents' belief updating, we extend the comparison of full Bayesian and maximum-likelihood updating in Section 3.3 to an intermediate class of updating rules. Under β -*mixture updating*, for all type spaces $(T, \mathcal{P})$, $i \in I$, and $t_i \in T_i$, we have

$$\mathcal{P}_{t_i} = \beta \operatorname{co} \mathcal{P}_{t_i}^{\text{FB}} + (1 - \beta) \operatorname{co} \mathcal{P}_{t_i}^{\text{ML}}. \quad (12)$$

Thus, t_i 's set of posteriors contracts the closed convex hull of the full-Bayesian posterior set towards that of the maximum-likelihood posterior set. Parameter β captures the extent of conservativeness, in the sense that the sets of posteriors $\mathcal{P}_{t_i}$ are increasing in β with respect to set inclusion. Convexification does not affect maximin payoffs or the belief and justifiability operators, so $\beta = 1$ and $\beta = 0$ are behaviorally equivalent to full Bayesian and maximum-likelihood updating, respectively.³⁵

The following proposition generalizes Proposition 1. The first part provides a weaker version of the critical path theorem with respect to common justifiability, which becomes stronger when β is higher, i.e., updating is more conservative. The second part exhibits failures of the critical path theorem, which are more extreme under smaller β , i.e., less conservative updating rules:

Proposition E.1. *Under β -mixture updating, the following holds:*

1. *For every type space $(T, \mathcal{P})$, product event E , $P \in \mathcal{P}$, and $p < \frac{1}{|I|}$, we have*

$$P(\overline{C}^{\beta p}(E)) \geq 1 - (1 - P(E)) \frac{1 - p}{1 - |I|p}.$$

2. *For any $\varepsilon \in (0, 1)$, there is a type space $(T, \mathcal{P})$ and product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \mathcal{P}$ but $\overline{C}^p(E) = \emptyset$ for all $p > \beta$.*

Proof. For the first part, let $\hat{B}^p(\cdot)$ and $\hat{C}^p(\cdot)$ be the p -belief operator and common p -belief operator under the unambiguous type space $(T, \{P\})$. For each $t_i \in T_i$ and product event $D \subseteq T$ such that $P(D_{-i}|t_i) \geq p$, we have $Q_i(D_{-i}) \geq \beta p$ for some $Q_i \in \mathcal{P}_{t_i}$. Thus, $\hat{B}^p(\cdot) \subseteq \overline{B}^{\beta p}(\cdot)$. Then the desired conclusion follows by the same argument as in the proof of the first part of Proposition 1. The proof of the second part is analogous to the second part of Proposition 1. $\square$

Proposition E.1 implies the following corollary:

Corollary E.1. *Under β -mixture updating, for any complete-information game G :*

1. *If $a^* \in A$ satisfies $\beta \min_{a_{-i}} u_i(a_i^*, a_{-i}) + (1 - \beta)u_i(a^*) > \max_{a_i \neq a_i^*} u_i(a_i, a_{-i}^*)$ for every $i \in I$, then a^* is contagious.*

³⁵While this updating rule can violate Assumption 1, this violation does not play a role in the analysis. Epstein and Schneider (2007) and Cheng (2022) propose similar hybrids of full Bayesian and maximum-likelihood updating, for which similar results can be obtained.

2. If $\hat{a}$ is a p -secure equilibrium of G with $p < \beta/|I|$, then $\hat{a}$ is robust.

Proof. The argument for the first part is analogous to the second part of Proposition 2. The second part follows from the first part of Proposition E.1 and Lemma A.2. $\square$

The condition in the first part requires that the payoff to the equilibrium action a_i^* evaluated at the worst-case belief in $\beta\Delta(A_{-i}) + (1 - \beta)\{\delta_{a_{-i}^*}\}$ is better than the payoff to any deviation a_i evaluated at the equilibrium belief $\delta_{a_{-i}^*}$. When $\beta = 1$, this reduces to a^* being secure as in Theorem 2; when $\beta = 0$, this reduces to a^* being a strict Nash equilibrium as in Proposition 2.

Extending Corollary 2, the second part shows that a p -secure $\hat{a}$ is robust when p is sufficiently small, where the requirement on p is more demanding the smaller is β .

E.2 Details for Remark 1

As discussed in Remark 1, it is important for our results that while ε -elaborations only feature small prior ambiguity about payoffs, they allow for significant ambiguity about payoff-irrelevant events. Elaborating on Remark 1, we show that if $(T, \mathcal{P})$ features small prior ambiguity about *all* events in T , then it satisfies an approximate version of the critical path theorem, and as a result, $(T, \mathcal{P})$ does not support the contagion in Theorem 2.

One natural way to limit the amount of ambiguity about all events in T is to bound the likelihood ratios $\frac{P(t)}{P'(t)}$ across different $P, P' \in \mathcal{P}$. When $\text{diam } \mathcal{P} := \sup_{P, P' \in \mathcal{P}, t \in T} \frac{P(t)}{P'(t)}$ is small, we obtain the following approximate version of the critical-path theorem:³⁶

Proposition E.2. *Consider any updating rule satisfying Assumption 1. Take any $r \geq 1$ and $p < \frac{1}{|I|r}$. Then for any type space $(T, \mathcal{P})$ such that $\text{diam } \mathcal{P} \leq \sqrt{r}$ and any product event E ,*

$$P(C^p(E)) \geq 1 - \frac{1 - rp}{1 - r|I|p}(1 - P(E)), \quad \forall P \in \mathcal{P}.$$

Proof. Fix any type space $(T, \mathcal{P})$ such that $\text{diam } \mathcal{P} \leq \sqrt{r}$ and any $P^0 \in \mathcal{P}$. Let $B^{p,0}$ and $C^{p,0}$ denote the individual and common belief operators for the unambiguous

³⁶Oyama and Tercieux (2010) derive an analogous result in the context of non-common probabilistic priors. We use the conventions $0/0 := 1$ and $x/0 := +\infty$ for every $x > 0$.

type space $(T, \{P^0\})$. We have $B^{rp,0}(E) \subseteq B^p(E)$ and hence $C^{rp,0}(E) \subseteq C^p(E)$ for any product event $E \subseteq T$: indeed, if $t_i \in B_i^{rp,0}(E)$, i.e., $t_i \in E_i$ and $P^0(\{t_i\} \times E_{-i}) \geq P^0(t_i) \times rp$, then $t_i \in E_i$ and $P'(\{t_i\} \times E_{-i}) \geq P'(t_i) \times p$ for all $P' \in \mathcal{P}$, which ensures $t_i \in B_i^p(E)$ by Assumption 1-(i).

Since $rp < 1/|I|$ by assumption, Kajii and Morris's (1997) critical path theorem applied to $(T, \{P^0\})$ implies that for any product event $E \subseteq T$,

$$P^0(C^{rp,0}(E)) \geq 1 - \frac{1 - rp}{1 - r|I|p}(1 - P^0(E)),$$

which proves the claim as $P^0(C^p(E)) \geq P^0(C^{rp,0}(E))$. □

As a corollary of Proposition E.2 and Lemma A.2, one also recovers a version of the robustness result in Kajii and Morris (1997) by bounding payoff-irrelevant ambiguity:

Corollary E.2. *Consider any updating rule satisfying Assumption 1. Suppose a^* is a p -dominant equilibrium of G with $p < \frac{1}{|I|r}$ for some $r \geq 1$. Then a^* is robust against all elaborations with $\text{diam } \mathcal{P} \leq \sqrt{r}$.*

References

- AHN, D. S. (2007): "Hierarchies of ambiguous beliefs," *Journal of Economic Theory*, 136(1), 286–301.
- AUMANN, R. J. (1976): "Agreeing to disagree," *The Annals of Statistics*, 4(6), 1236–1239.
- BEAUCHÊNE, D., J. LI, AND M. LI (2019): "Ambiguous persuasion," *Journal of Economic Theory*, 179, 312–365.
- BERGEMANN, D., AND S. MORRIS (2013): "An Introduction to Robust Mechanism Design," *Foundations and Trends® in Microeconomics*, 8(3), 169–230.
- (2019): "Information design: A unified perspective," *Journal of Economic Literature*, 57(1), 44–95.
- BILLOT, A., A. CHATEAUNEUF, I. GILBOA, AND J.-M. TALLON (2000): "Sharing beliefs: between agreeing and disagreeing," *Econometrica*, 68(3), 685–694.
- BILLOT, A., S. MUKERJI, AND J.-M. TALLON (2020): "Market allocations under ambiguity: A survey," *Revue économique*, 71(2), 267–282.
- BOSE, S., AND L. RENOU (2014): "Mechanism design with ambiguous communication devices," *Econometrica*, 82(5), 1853–1872.
- CERREIA-VIOGLIO, S., R. CORRAO, AND G. LANZANI (2022): "(Un-) Common Preferences, Ambiguity, and Coordination," working paper.

- CHASSANG, S., AND S. TAKAHASHI (2011): “Robustness to incomplete information in repeated games,” *Theoretical Economics*, 6(1), 49–93.
- CHEN, J. Y. (2026): “Sequential Learning under Informational Ambiguity,” *American Economic Review*, 116(1), 209–245.
- CHENG, X. (2022): “Relative maximum likelihood updating of ambiguous beliefs,” *Journal of Mathematical Economics*, 99, 102587.
- CHENG, X., P. KLIBANOFF, S. MUKERJI, AND L. RENOU (2024): “Persuasion with ambiguous communication,” in *Proceedings of the 25th ACM Conference on Economics and Computation*, pp. 1201–1202.
- DI TILLO, A., N. KOS, AND M. MESSNER (2016): “The design of ambiguous mechanisms,” *Review of Economic Studies*, 84(1), 237–276.
- DOW, J., AND S. R. WERLANG (1992): “Uncertainty aversion, risk aversion, and the optimal choice of portfolio,” *Econometrica*, pp. 197–204.
- DÜTTING, P., M. FELDMAN, D. PERETZ, AND L. SAMUELSON (2024): “Ambiguous contracts,” *Econometrica*, 92(6), 1967–1992.
- EICHBERGER, J., S. GRANT, AND D. KELSEY (2007): “Updating choquet beliefs,” *Journal of Mathematical Economics*, 43(7-8), 888–899.
- ELLIS, A. (2016): “Condorcet meets Ellsberg,” *Theoretical Economics*, 11(3), 865–895.
- EPSTEIN, L. G., AND M. SCHNEIDER (2007): “Learning under ambiguity,” *Review of Economic Studies*, 74(4), 1275–1303.
- EPSTEIN, L. G., AND T. WANG (1996): “‘Beliefs about beliefs’ without probabilities,” *Econometrica*, 64(6), 1343–1373.
- FRICK, M., R. IJIMA, AND Y. ISHII (2025): “Efficient Learning from Ambiguous Information,” working paper.
- GANGULI, J. V. (2007): “Common p-belief and uncertainty,” in *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 155–159.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18(2), 141–153.
- (1993): “Updating ambiguous beliefs,” *Journal of Economic Theory*, 59(1), 33–49.
- GUL, F., AND W. PESENDORFER (2021): “Evaluating ambiguous random variables from Choquet to maxmin expected utility,” *Journal of Economic Theory*, 192, 105129.
- HALAC, M. (2025): “Contracting for coordination,” *Journal of the European Economic Association*, 23(3), 815–844.
- HALAC, M., E. LIPNOWSKI, AND D. RAPPOPORT (2021): “Rank uncertainty in organizations,” *American Economic Review*, 111(3), 757–786.
- HANANY, E., AND P. KLIBANOFF (2007): “Updating preferences with multiple priors,” *Theoretical Economics*, 2(3), 261–298.
- HANANY, E., P. KLIBANOFF, AND S. MUKERJI (2020): “Incomplete information games with ambiguity averse players,” *American Economic Journal: Microeconomics*, 12(2), 135–187.

- HORIE, M. (2013): “Reexamination on updating choquet beliefs,” *Journal of Mathematical Economics*, 49(6), 467–470.
- HOSHINO, T. (2022): “Multi-Agent Persuasion: Leveraging Strategic Uncertainty,” *International Economic Review*, 63(2), 755–776.
- HUO, Z., M. PEDRONI, AND G. PEI (2024): “Bias and sensitivity under ambiguity,” *American Economic Review*, 114(12), 4091–4133.
- INOSTROZA, N., AND A. PAVAN (2025): “Adversarial coordination and public information design,” *Theoretical Economics*, 20(2), 763–813.
- KAJII, A., AND S. MORRIS (1997): “The robustness of equilibria to incomplete information,” *Econometrica*, pp. 1283–1309.
- (2020a): “Notes on “Refinements and higher-order beliefs,”” *The Japanese Economic Review*, 71(1), 35–41.
- (2020b): “Refinements and higher-order beliefs: a survey,” *The Japanese Economic Review*, 71(1), 7–34.
- KAJII, A., AND T. UI (2005): “Incomplete information games with multiple priors,” *Japanese Economic Review*, 56(3), 332–351.
- (2009): “Interim efficient allocations under uncertainty,” *Journal of Economic Theory*, 144(1), 337–353.
- KELLNER, C., AND M. T. LE QUEMENT (2018): “Endogenous ambiguity in cheap talk,” *Journal of Economic Theory*, 173, 1–17.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A smooth model of decision making under ambiguity,” *Econometrica*, 73(6), 1849–1892.
- (2009): “Recursive smooth ambiguity preferences,” *Journal of Economic Theory*, 144(3), 930–976.
- LI, F., Y. SONG, AND M. ZHAO (2023): “Global manipulation by local obfuscation,” *Journal of Economic Theory*, 207, 105575.
- MACHINA, M. J., AND M. SINISCALCHI (2014): “Ambiguity and ambiguity aversion,” in *Handbook of the Economics of Risk and Uncertainty*, vol. 1. Elsevier.
- MONDERER, D., AND D. SAMET (1989): “Approximating common knowledge with common beliefs,” *Games and Economic Behavior*, 1(2), 170–190.
- MORIYA, F., AND T. YAMASHITA (2020): “Asymmetric-information allocation to avoid coordination failure,” *Journal of Economics & Management Strategy*, 29(1), 173–186.
- MORRIS, S., D. OYAMA, AND S. TAKAHASHI (2022): “On the joint design of information and transfers,” working paper.
- (2024): “Implementation via Information Design in Binary-Action Supermodular Games,” *Econometrica*, 92(3), 775–813.
- MORRIS, S., AND H. S. SHIN (2002): “Social value of public information,” *American Economic Review*, 92(5), 1521–1534.
- (2007): “Common belief foundations of global games,” working paper.

- MORRIS, S., AND T. UI (2005): “Generalized potentials and robust sets of equilibria,” *Journal of Economic Theory*, 124(1), 45–78.
- OYAMA, D., AND S. TAKAHASHI (2020): “Generalized Belief Operator and Robustness in Binary-Action Supermodular Games,” *Econometrica*, 88(2), 693–726.
- OYAMA, D., AND O. TERCIEUX (2009): “Iterated potential and robustness of equilibria,” *Journal of Economic Theory*, 144(4), 1726–1769.
- (2010): “Robust equilibria under non-common priors,” *Journal of Economic Theory*, 145(2), 752–784.
- (2012): “On the strategic impact of an event under non-common priors,” *Games and Economic Behavior*, 74(1), 321–331.
- PEI, H., AND B. STRULOVICI (2025): “Robust implementation with costly information,” *Review of Economic Studies*, 92(1), 476–505.
- PIRES, C. P. (2002): “A rule for updating ambiguous beliefs,” *Theory and Decision*, 53(2), 137–152.
- RUBINSTEIN, A. (1989): “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge’,” *The American Economic Review*, pp. 385–391.
- SADOWSKI, P., AND T. SARVER (2026): “An Evolutionary Perspective on Updating Risk and Ambiguity Preferences,” *American Economic Review: Insights*, 8(2), 214–32.
- SAKOVICS, J., AND J. STEINER (2012): “Who matters in coordination problems?,” *American Economic Review*, 102(7), 3439–3461.
- SALO, A., AND M. WEBER (1995): “Ambiguity aversion in first-price sealed-bid auctions,” *Journal of Risk and Uncertainty*, 11, 123–137.
- SCHMEIDLER, D. (1989): “Subjective probability and expected utility without additivity,” *Econometrica*, pp. 571–587.
- SEGAL, I. (2003): “Coordination and discrimination in contracting with externalities: Divide and conquer?,” *Journal of Economic Theory*, 113(2), 147–181.
- SEIDENFELD, T., AND L. WASSERMAN (1993): “Dilation for sets of probabilities,” *The Annals of Statistics*, 21(3), 1139–1154.
- SHISHKIN, D., AND P. ORTOLEVA (2023): “Ambiguous information and dilation: An experiment,” *Journal of Economic Theory*, 208, 105610.
- SINISCALCHI, M. (2009): “Two out of three ain’t bad: A comment on “The ambiguity aversion literature: A critical assessment”,” *Economics & Philosophy*, 25(3), 335–356.
- TANG, R., AND M. ZHANG (2021): “Maxmin implementation,” *Journal of Economic Theory*, 194, 105250.
- UI, T. (2001): “Robust equilibria of potential games,” *Econometrica*, 69(5), 1373–1380.
- (2025): “Strategic ambiguity in global games,” *Games and Economic Behavior*, 149, 65–81.
- WEINSTEIN, J., AND M. YILDIZ (2007): “A structure theorem for rationalizability with application to robust predictions of refinements,” *Econometrica*, 75(2), 365–400.
- YOKOTA, Y. (2024): “Rationalizability in Regular Preference Form Games: Incomplete Information and Higher Order Uncertainty,” working paper.

Online Appendix to “Contagious Ambiguity”

Mira Frick, Ryota Iijima, and Daisuke Oyama

F More General Ambiguity Preferences

F.1 α -Maxmin

We generalize the main setting by considering α -maxmin expected utility, which allows for a mix of ambiguity-averse and ambiguity-seeking tendencies. Specifically, fix some $\alpha_i \in [0, 1]$ for each $i \in I$. Under any updating rule satisfying Assumption 1, define type t_i ’s interim payoff to playing action a_i against opponent strategy σ_{-i} by

$$\begin{aligned} V_i(a_i, \sigma_{-i}, t_i) = & \alpha_i \min_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}) \\ & + (1 - \alpha_i) \max_{P_i \in \mathcal{P}_{t_i}} \sum_{t_{-i} \in T_{-i}} P_i(t_{-i}) v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}). \end{aligned}$$

The following result generalizes Theorem 2. We omit the proof as the argument is analogous.

Proposition F.1. *Consider players who have α -maxmin preferences. Then action profile $a^* \in A$ is contagious if for each $i \in I$ and $a_i \neq a_i^*$,*

$$\alpha_i \min_{a_{-i}} u_i(a_i^*, a_{-i}) + (1 - \alpha_i) u_i(a^*) > \alpha_i u_i(a_i, a_{-i}^*) + (1 - \alpha_i) \max_{a_{-i}} u_i(a_i, a_{-i}). \quad (13)$$

Condition (13) reduces to security under maxmin expected utility ($\alpha_i = 1$ for every i). In contrast, for maxmax expected utility ($\alpha_i = 0$ for every i), the inequality takes the form $u_i(a^*) > \max_{a_{-i}, a_i \neq a_i^*} u_i(a_i, a_{-i})$. The latter condition is satisfied if, for example, a^* is a strict Nash equilibrium and every action $a_i \neq a_i^*$ is safe. Thus, if agents are ambiguity-seeking, the presence of safe actions serves as a source of contagion for non-safe equilibria. However, as in the maxmin case, this means that almost dominated actions can be contagious. For instance, consider the illustrative example from Section 1.1, where for each $i = 1, 2$, $A_i = \{0, 1\}$ and $u_i(a) = \mathbf{1}_{\{a_i = a_{-i} = 1\}} - ca_i$ for some $c \in (0, 1)$. Action profile $\mathbf{1}$ is contagious under maxmax expected utility. However, $\mathbf{0}$ is p -dominant for any $p > 1 - c$, and hence almost dominant when $c \approx 1$.

F.2 Second-order Prior

In this section, we formalize ambiguity using a second-order prior, following the smooth ambiguity model of [Klibanoff, Marinacci, and Mukerji \(2005\)](#). A *KMM*

type space (T, μ, ϕ) consists of a countable set $T = \prod_{i \in I} T_i$ of type profiles, a common second-order prior $\mu \in \Delta(\Delta(T))$, and a tuple $\phi = (\phi_i)_{i \in I}$ of aggregators. Here, μ captures uncertainty about first-order priors $P \in \Delta(T)$ in a Bayesian manner. Each $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ is a monotone transformation that captures player i 's attitude toward uncertainty over first-order priors. The analysis below does not rely on the possibility that the ϕ_i differ across players.

At type t_i , player i updates the prior $\mu \in \Delta(\Delta(T))$ to $\mu_{t_i} \in \Delta(\Delta(T_{-i}))$. As in [Klibanoff, Marinacci, and Mukerji \(2009\)](#), one common updating rule is to define μ_{t_i} as the pushforward measure of $\tilde{\mu}_{t_i} \in \Delta(\Delta(T))$ with respect to the mapping $P \mapsto P(\cdot | t_i)$, where $\tilde{\mu}_{t_i}$ is constructed by

$$d\tilde{\mu}_{t_i}(P) = \frac{P(t_i)}{\int_{\Delta(T)} P'(t_i) d\mu(P')} d\mu(P).$$

Here, Bayes' rule is applied to each first-order prior P , as well as the second-order prior μ .³⁷ We assume that $\int_{\Delta(T)} P'(t_i) d\mu(P')$ is positive for each t_i , so that the above construction is well-defined.

Define the p -belief operator by

$$B_i^p(E) = \left\{ t_i \in E_i : \int_{\Delta(T_{-i})} \phi_i(P_i(E_{-i})) d\mu_{t_i}(P_i) \geq \phi_i(p) \right\}$$

for each $i \in I$, $p \in (0, 1)$, and product event E , and define $B^p(E)$ and $C^p(E)$ analogously to the main setting.³⁸

The following result provides an analog of Theorem 1 for KMM type spaces. In particular, common ε -belief in an event E can fail even when the ex-ante probability of E is arbitrarily high, and even when the event itself is unambiguous.

Proposition F.2. *For every $\varepsilon \in (0, 1)$, there is a KMM type space (T, μ, ϕ) and a product event E such that $P(E) = P'(E) \geq 1 - \varepsilon$ for all $P, P' \in \text{supp}(\mu)$ but $C^\varepsilon(E) = \emptyset$.*

An incomplete-information game under smooth ambiguity preferences is specified by a KMM type space (T, μ, ϕ) , along with a finite set A_i of actions and a bounded payoff function $v_i : A \times T \rightarrow \mathbb{R}$ for each i . Player i 's interim payoff at t_i is given by

$$V_i(a_i, \sigma_{-i}, t_i) = \int \phi_i \left(\sum_{t_{-i}} P_i(t_{-i}) v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}) \right) d\mu_{t_i}(P_i),$$

³⁷Our results remain valid under an alternative updating procedure that applies Bayes' rule only to first-order priors, i.e., $\mu_{t_i}(\{P(\cdot | t_i) : P \in \mathcal{P}\}) = \mu(\mathcal{P})$ for each t_i and measurable $\mathcal{P} \subseteq \Delta(T)$.

³⁸The belief operator naturally depends on ϕ_i . As in the unambiguous case, $B_i^p(E)$ corresponds to the set of player i 's types in E_i who prefer to receive a bet that pays utility 1 on E_{-i} and utility 0 otherwise, rather than a lottery that pays utility 1 with probability p and utility 0 with the remaining probability. This is a standard method for eliciting an agent's subjective belief about E_{-i} .

which is linearly extended to mixed actions $\alpha_i \in \Delta(A_i)$ of player i . Equilibrium is defined analogously to the main model.

To address the question of contagion, fix a complete-information game G . An ε -elaboration is an incomplete-information game under smooth ambiguity preferences where $P(\prod_{i \in I} T_i^*) \geq 1 - \varepsilon$ for all $P \in \text{supp}(\mu)$ and T_i^* is the set of i 's normal types. The definition of a contagious action profile a^* is analogous to the main model. The following result is an analog of Theorem 2 for smooth ambiguity preferences:

Proposition F.3. *Consider players who have smooth ambiguity preferences. If an action profile a^* in G is secure, then it is contagious.*

While Proposition F.3 allows the aggregators (ϕ_i) to depend on the ε -elaboration, this dependence is not needed for the contagion argument. That is, we can fix aggregators ϕ_i uniformly across values of ε , provided that the ϕ_i are sufficiently concave. For example, under the widely used CARA specification $\phi_i(v) = -\exp[-\alpha_i v]$, the argument goes through as long as each α_i is sufficiently large.

F.3 Non-additive Prior

In this section, we formalize ambiguity using a non-additive prior, following the Choquet expected utility model of [Schmeidler \(1989\)](#). We consider a *Choquet type space* (T, ν) , where T is a countable set of type profiles and ν is a *capacity*, i.e., a mapping $\nu : 2^T \rightarrow [0, 1]$ that is monotone with respect to set inclusion and satisfies $\nu(\emptyset) = 0$ and $\nu(T) = 1$. The dual capacity is defined by $\bar{\nu}(E) := 1 - \nu(T \setminus E)$ for each $E \subseteq T$. An event E is called unambiguous if $\nu(E) = \bar{\nu}(E)$, i.e., if the capacity is additive over E and its complement $T \setminus E$.

Let ν_{t_i} denote the updated capacity on T_{-i} conditional on t_i . There are two standard updating rules in the literature. Under *Dempster-Shafer updating*, the conditional capacity is given by

$$\nu_{t_i}^{\text{DS}}(E_{-i}) = \frac{\bar{\nu}(\{t_i\} \times T_{-i}) - \bar{\nu}(\{t_i\} \times (T_{-i} \setminus E_{-i}))}{\bar{\nu}(\{t_i\} \times T_{-i})}$$

for each $E_{-i} \subseteq T_{-i}$. Under *generalized Bayesian updating*,

$$\nu_{t_i}^{\text{GB}}(E_{-i}) = \frac{\nu(\{t_i\} \times E_{-i})}{\nu(\{t_i\} \times E_{-i}) + \bar{\nu}(\{t_i\} \times (T_{-i} \setminus E_{-i}))}$$

for each $E_{-i} \subseteq T_{-i}$.³⁹ We assume the denominators above are positive so that the updating rules are well-defined, and each player follows either of these updating rules.

³⁹The latter is also called the Dempster-Fagin-Halpern rule. See [Gilboa and Schmeidler \(1993\)](#); [Eichberger, Grant, and Kelsey \(2007\)](#); [Horie \(2013\)](#) for axiomatic foundations of these two rules.

Define p -belief operators by

$$B_i^p(E) = \{t_i \in E_i : \nu_{t_i}(E_{-i}) \geq p\}$$

for each $i \in I$, $p \in (0, 1)$, and product event E , and define $B^p(E)$ and $C^p(E)$ analogously to the main setting.

The following result provides an analog of Theorem 1 for Choquet type spaces. In particular, common ε -belief in an event E can fail even when the ex-ante probability of E is arbitrarily high, and even when the event itself is unambiguous.

Proposition F.4. *For every $\varepsilon \in (0, 1)$, there is a Choquet type space (T, ν) and a product event E such that $\nu(E) = \bar{\nu}(E) \geq 1 - \varepsilon$ but $C^\varepsilon(E) = \emptyset$.*

An incomplete-information game under Choquet expected utility is given by a Choquet type space (T, ν) , along with a finite set A_i of actions and a bounded payoff function $v_i : A \times T \rightarrow \mathbb{R}$ for each i . Player i 's interim payoff at t_i is given by

$$\begin{aligned} V_i(a_i, \sigma_{-i}, t_i) &= \int_0^\infty \nu_{t_i}(\{t_{-i} : v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}) \geq x\}) dx \\ &\quad + \int_{-\infty}^0 (\nu_{t_i}(\{t_{-i} : v_i(a_i, \sigma_{-i}(t_{-i}), t_i, t_{-i}) \geq x\}) - 1) dx, \end{aligned}$$

which is linearly extended to mixed actions $\alpha_i \in \Delta(A_i)$ of player i . Equilibrium is defined analogously to the main model.

To address the question of contagion, fix a complete-information game G . An ε -elaboration is an incomplete-information game under Choquet expected utility with $\nu(\prod_{i \in I} T_i^*) \geq 1 - \varepsilon$ and $\bar{\nu}(\prod_{i \in I} T_i^*) \geq 1 - \varepsilon$, where T_i^* denotes the set of player i 's normal types. The definition of a contagious action profile a^* is analogous to the main model. The following result is an analog of Theorem 2 for Choquet expected utility:

Proposition F.5. *Consider players with Choquet expected utility preferences. If an action profile a^* in G is secure, then it is contagious.*

F.4 Proofs

F.4.1 Proofs for Appendix F.2

Proof of Proposition F.2. Take any $\varepsilon \in (0, 1)$. Choose ϕ_i for each i such that

$$\phi_i(\varepsilon) > \frac{1}{|I|} \phi_i(0) + \frac{|I| - 1}{|I|} \phi_i(1).$$

For each $i \in I$, take $T_i = \{0, 1, \dots, N\}$ and P^i as in the proof of Theorem 1. Let $\mu(\{P^i\}) = \frac{1}{|I|}$ for each $i \in I$. Let $E = \{1, 2, \dots, N\}^I$. Observe that $P^i(E) = P^j(E) \geq 1 - \varepsilon$ for all i, j .

By construction, $P^i(\{t_i\} \times T_{-i}) = P^j(\{t_i\} \times T_{-i})$ for all i, j . Thus, for each $t_i > 0$, μ_{t_i} assigns probability $\frac{1}{|I|}$ on $P^i(\cdot | t_i)$, where $P^i(t_j = t_i - 1 \forall j \neq i | t_i) = 1$. Hence, for each $t_i = n$,

$$\int_{\Delta(T_{-i})} \phi_i(P_i(\{t_{-i} \geq \mathbf{n}\})) d\mu_{t_i}(P_i) \leq \frac{1}{|I|} \phi_i(0) + \frac{|I| - 1}{|I|} \phi_i(1).$$

Given the choice of each ϕ_i , an inductive argument shows that, for all $k \in \mathbb{N}$, we have $(B^\varepsilon)^k(E) \subseteq \{k + 1, k + 2, \dots, N\}^I$. Therefore, $C^\varepsilon(E) = \emptyset$. $\square$

Proof of Proposition F.3. Since $\min_{a_{-i}} u_i(a_i^*, a_{-i}) > u_i(a_i, a_{-i}^*)$ for each $i \in I$ and $a_i \neq a_i^*$, there exist functions $(\phi_i)_{i \in I}$ such that

$$\frac{1}{|I|} \phi_i(u_i(a_i^*)) + \frac{|I| - 1}{|I|} \phi_i(\min_{a_{-i}} u_i(a_i^*, a_{-i})) > \frac{1}{|I|} \phi_i(u_i(a_i, a_{-i}^*)) + \frac{|I| - 1}{|I|} \phi_i(\max_{a_{-i}} u_i(a_i, a_{-i})) \quad (14)$$

for all $i \in I$, $a_i \neq a_i^*$.

Fix any $\varepsilon > 0$. We take the KMM type space as in the proof of Proposition F.2. For each i , $T_i^* = \{1, 2, \dots, N\}$ is the set of normal types, while a_i^* is dominant for type $t_i = 0$. Note that $P^i(\prod_k T_k^*) = P^j(\prod_k T_k^*) \geq 1 - \varepsilon$ for every i, j .

By construction, $P^i(t_i) = P^j(t_i)$ for all i, j . Thus, for each $t_i > 0$, μ_{t_i} assigns probability $\frac{1}{|I|}$ on $P^i(\cdot | t_i)$, where $P^i(t_j = t_i - 1 \forall j \neq i | t_i) = 1$. Combined with (14), an inductive argument shows that a^* is played at every type profile in the unique equilibrium. $\square$

F.4.2 Proofs for Appendix F.3

Proof of Proposition F.4. Fix any $\varepsilon \in (0, 1)$. Take the type space $(T, \mathcal{P})$ with multiple priors as constructed in the proof of the second part of Proposition 1. Then define a Choquet type space (T, ν) by setting $\nu(E) = \min_{P \in \mathcal{P}} P(E)$ (and hence $\bar{\nu}(E) = \max_{P \in \mathcal{P}} P(E)$) for each $E \subseteq T$.

Consider any i with type $t_i = n > 0$. If i follows Dempster-Shafer updating⁴⁰,

$$\begin{aligned} \nu_{t_i}(\{t_{-i} \not\prec \mathbf{n}\}) &= \frac{\max_{P \in \mathcal{P}} P(\{t_i = n\}) - \max_{P \in \mathcal{P}} P(\{t_i = n \text{ and } t_{-i} < \mathbf{n}\})}{\max_{P \in \mathcal{P}} P(\{t_i = n\})} \\ &= \frac{P^i(\{t_i = n\}) - P^i(\{t_i = n \text{ and } t_{-i} < \mathbf{n}\})}{P^i(\{t_i = n\})} = 0. \end{aligned}$$

If i follows generalized Bayesian updating,

$$\nu_{t_i}(\{t_{-i} \not\prec \mathbf{n}\}) = \frac{\min_{P \in \mathcal{P}} P(\{t_i = n \text{ and } t_{-i} \not\prec \mathbf{n}\})}{\min_{P \in \mathcal{P}} P(\{t_i = n \text{ and } t_{-i} \not\prec \mathbf{n}\}) + \max_{P \in \mathcal{P}} P(\{t_i = n \text{ and } t_{-i} < \mathbf{n}\})}$$

⁴⁰Here “ $<$ ” denotes component-wise strict inequality.

$$= \frac{P^i(\{t_i = n \text{ and } t_{-i} \not\prec \mathbf{n}\})}{P^i(\{t_i = n \text{ and } t_{-i} \not\prec \mathbf{n}\}) + P^i(\{t_i = n \text{ and } t_{-i} < \mathbf{n}\})} = 0.$$

Consider the product event $E = \{1, 2, \dots\}^I$. Then $\nu(E) = \bar{\nu}(E) \geq 1 - \varepsilon$. Based on the above calculations, one can verify inductively that for all $k \in \mathbb{N}$, we have $(B^\varepsilon)^k(E) = \{k+1, k+2, \dots\}^I$. Therefore $C^\varepsilon(E) = \emptyset$. $\square$

Proof of Proposition F.5. Fix any $\varepsilon \in (0, 1)$. Take an ε -elaboration in which the Choquet type space (T, ν) is as in the proof of Proposition F.4. For each i , let $T_i^* = \{1, 2, \dots\}$ and a_i^* be strictly dominant for type $t_i = 0$. For type $t_i = 1$, the interim payoff of a_i^* is at least $\min_{a_{-i}} u_i(a_i^*, a_{-i})$, while the interim payoff of any $a_i \neq a_i^*$ is at most

$$u_i(a_i, a_{-i}^*) + \nu_{t_i}(\{t_{-i} \not\prec \mathbf{1}\})(\max_{a_{-i}} u_i(a_i, a_{-i}) - u_i(a_i, a_{-i}^*)) = u_i(a_i, a_{-i}^*),$$

where the equality follows from the calculations in the proof of Proposition F.4 and the inequality uses the fact that type $t_j = 0$ chooses a_j^* for each $j \neq i$. Therefore, security implies that a_i^* is chosen by type $t_i = 1$. Repeating the same argument for type $t_i = n$, using $\nu_{t_i}(\{t_{-i} \not\prec \mathbf{n}\}) = 0$ and induction on n , shows that every type plays a_i^* in the unique equilibrium. $\square$

G Proof of Proposition C.1

Lower bound: Fix any unambiguous scheme $\mathcal{C} = ((T_i)_{i \in I}, \{P\}, (w_i)_{i \in I})$ that uniquely implements **1**. We apply the sequential best-response partition used in Step 1 of Proposition 4 of Morris, Oyama, and Takahashi (2022). There is a partition $T = \bigcup_{\gamma \in \Pi} T(\gamma)$ such that, for every $i \in I$ and every type t_i with positive marginal probability,

$$\sum_{\gamma \in \Pi} \sum_{\substack{t_{-i}: \\ (t_i, t_{-i}) \in T(\gamma)}} P(t_{-i} \mid t_i) [d_i(S_{-i}(\gamma)) + w_i(t_i)] > 0. \quad (15)$$

Define

$$\rho_i(\gamma \mid t_i) = \sum_{\substack{t_{-i}: \\ (t_i, t_{-i}) \in T(\gamma)}} P(t_{-i} \mid t_i), \quad \rho(\gamma) = \sum_{t \in T(\gamma)} P(t).$$

For every player i ,

$$\rho(\gamma) = \sum_{t_i \in T_i} P(\{t_i\} \times T_{-i}) \rho_i(\gamma \mid t_i), \quad \gamma \in \Pi. \quad (16)$$

Using the additive form of the payment, (15) becomes

$$w_i(t_i) + \sum_{\gamma \in \Pi} \rho_i(\gamma \mid t_i) d_i(S_{-i}(\gamma)) > 0.$$

By multiplying this inequality by $P(\{t_i\} \times T_{-i})$, summing over t_i , and using (16), we obtain

$$\bar{w}_i + \sum_{\gamma \in \Pi} \rho(\gamma) d_i(S_{-i}(\gamma)) > 0,$$

where $\bar{w}_i = \sum_{t_i \in T_i} P(\{t_i\} \times T_{-i}) w_i(t_i)$.

Thus $(\bar{w}, \rho)$ is feasible for the weak inequalities in (11). Since the expected cost of $\mathcal{C}$ is $\sum_i \bar{w}_i$, every uniquely implementing scheme has cost at least the right-hand side of (11).

Upper bound: Let (z^*, ρ^*) solve (11). Fix $\delta > 0$ and set $b_i^\delta = z_i^* + \delta$. Then

$$\sum_{\gamma \in \Pi} \rho^*(\gamma) [d_i(S_{-i}(\gamma)) + b_i^\delta] \geq \delta \quad \text{for every } i \in I. \quad (17)$$

Choose a finite number $\hat{b}_i > b_i^\delta$ such that

$$d_i(\emptyset) + \hat{b}_i > 0. \quad (18)$$

By supermodularity, action 1 is then strictly dominant for a type receiving $\hat{b}_i$.

We use the rank construction from Step 2 of Proposition 4 of [Morris, Oyama, and Takahashi \(2022\)](#). Let $T_i = \{1, 2, \dots\}$, and, for each $\gamma \in \Pi$, let $r_i(\gamma) \in \{1, \dots, n\}$ be the rank of player i in γ , where $n := |I|$. For $m \in \mathbb{N}$, define

$$t_i^{m, \gamma} = m + r_i(\gamma).$$

For $\eta \in (0, 1)$, define the common prior by

$$P_\eta(t^{m, \gamma}) = \eta(1 - \eta)^m \rho^*(\gamma), \quad m \in \mathbb{N}, \quad \gamma \in \Pi,$$

with zero probability on all other type profiles. Types of zero marginal probability may be deleted from T_i . Set

$$w_i^\eta(\tau) = \begin{cases} \hat{b}_i, & \tau \leq n - 1, \\ b_i^\delta, & \tau \geq n. \end{cases}$$

By (17), for all sufficiently small $\eta > 0$,

$$\sum_{\gamma \in \Pi} (1 - \eta)^{n - r_i(\gamma)} \rho^*(\gamma) [d_i(S_{-i}(\gamma)) + b_i^\delta] > 0 \quad \text{for every } i \in I. \quad (19)$$

Every type $\tau \leq n - 1$ strictly prefers action 1 regardless of opponents' actions, by (18). Now take $\tau \geq n$ and suppose that all types below τ have already been shown to have action 1 as their unique rationalizable action. Conditional on $t_i = \tau$, the opponents who precede i in γ have types below τ . Even if every remaining opponent chooses action 0, supermodularity implies that type τ 's conditional expected payoff gain from action 1 is at least

$$\frac{\sum_{\gamma \in \Pi} (1 - \eta)^{n-r_i(\gamma)} \rho^*(\gamma) [d_i(S_{-i}(\gamma)) + b_i^\delta]}{\sum_{\gamma \in \Pi} (1 - \eta)^{n-r_i(\gamma)} \rho^*(\gamma)} > 0. \quad (20)$$

Induction on τ therefore eliminates action 0 for every type. Thus action 1 is uniquely rationalizable for every type and, in particular, **1** is the unique equilibrium.

It remains to bound the cost. The probability that all players' types are at least n is $(1 - \eta)^{n-1}$. On this event, each player i receives b_i^δ ; off this event, total payment is at most $\sum_i \widehat{b}_i$. Hence the expected cost is no more than

$$(1 - \eta)^{n-1} \sum_{i \in I} b_i^\delta + (1 - (1 - \eta)^{n-1}) \sum_{i \in I} \widehat{b}_i.$$

Letting $\eta \downarrow 0$ through values satisfying (19) gives $W_{\text{prob}}^1 \leq \sum_{i \in I} b_i^\delta$. Since $\sum_{i \in I} b_i^\delta$ equals the right-hand side of (11) plus $n\delta$, and $\delta > 0$ is arbitrary, this proves that the right-hand side of (11) upper-bounds W_{prob}^1 . $\square$