



Institutional Members: CEPR, NBER and Università Bocconi

WORKING PAPER SERIES

Regret in Games

Pierpaolo Battigalli, Martin Dufwenberg and Senran Lin

Working Paper n. 722

This Version: September 27, 2026

IGIER – Università Bocconi, Via Guglielmo Röntgen 1, 20136 Milano –Italy
<http://www.igier.unibocconi.it>

The opinions expressed in the working papers are those of the authors alone, and not those of the Institute, which takes non institutional policy position, nor those of CEPR, NBER or Università Bocconi.

Regret in Games*

Pierpaolo Battigalli, Martin Dufwenberg & Senran Lin

Draft of September 27, 2026

Abstract

We develop a general approach to exploring how regret influences strategic interaction and risky choice. Regret is captured by the payoff gap between what a player gets and what he believes he would have gotten had he chosen differently. Ex post beliefs are critical to that evaluation; the modeling therefore draws on tools from psychological game theory. Predictions depend in novel ways on the information structure across end-nodes and assumptions regarding the precise nature of chance moves. Regret can have a powerful impact in a variety of economic settings including, e.g., information acquisition, auctions, coordination, climate action, market entry, and delegation.

Keywords: regret, psychological games, belief-dependent motivation, rumination, information

JEL codes: C72; D91

*Battigalli: Bocconi University, IGIER & BIDS, Italy; pierpaolo.battigalli@unibocconi.it. Dufwenberg: Purdue University, USA; University of Gothenburg, Sweden; mdufwenb@purdue.edu. Lin: University of Edinburgh, United Kingdom; senranlin@outlook.com. We thank Carlo Andreatta, Samuele Dotta, Tore Ellingsen, Tangren Feng, Federico Golonia, Laura Lo Schiavo, Mark Machina, Silvia Meneghesso, William Miller, Viola Sigismondi, Nicolas Sourisseau, Bob Sugden, Jörgen Weibull, and the participants of several seminars and conferences for their comments. This project is funded by the European Union (ERC, TRAITS-GAMES, 101142844). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

1 Introduction

*Losing one glove is certainly painful,
but nothing compared to the pain,
of losing one, throwing away the other,
and finding the first one again.*

Regret can be a powerful emotion, as the above grook (by Piet Hein) exemplifies. People regret things they say (hurt feelings?), things they did (quit a job?), and how they lived their life (on their deathbed!). If regret is anticipated, behavior may be adjusted to avoid the pang. A substantial literature in psychology, as well as related neuroscience research, explores regret.¹ Yet, few economists paid attention despite the potential of regret to influence behavior and shape economic outcomes. A branch of “decision regret theory” (DRT) explores how regret affects single decision makers (e.g., risky choice and choice over menus),² there is some work exploring how regret may influence insurance and investment decisions,³ and there is heterogeneous work on specific games.⁴ However, no comprehensive game-theoretic approach has been available. Until now, we develop a theory for a general class of game forms.

An initial example helps flag how and why regret shapes behavior in novel ways:

¹For psychology, see, e.g., Zeelenberg (1999), Connolly & Butler (2006), Zeelenberg & Pieters (2007*a,b*), and references therein. Compare also Cadish’s (2001) book *Damn!: Reflections on Life’s Biggest Regrets* in which the author concludes that “regrets are universal; nearly everyone has them. Regrets transcend age, gender, race, culture, nationality, religion, language, social status, and geographic location” (p. 2). For neuroscience see, e.g., Camille, Coricelli, Sallet, Pradate-Diehl, Duhamel & Sirigu (2004), Coricelli, Critchley, Joffily, O’Doherty, Sirigu & Dolan (2005), Frydman & Camerer (2016), and Fioretti, Vostroknutov & Coricelli (2022).

²The DRT label is ours. The pioneers in this literature are Bell (1982) and Loomes & Sugden (1982), and subsequent work includes Loomes & Sugden (1987), Sugden (1993), Quiggin (1990, 1994), Sarver (2008), Bikhchandani & Segal (2011, 2014), Halpern & Pass (2012), and Krämer & Stone (2013).

³See, e.g., Shefrin & Statman (1985), Braun & Muermann (2004), Qin (2015), Frydman & Camerer (2016), Fioretti *et al.* (2022), and Armantier, Foncel & Treich (2023).

⁴Engelbrecht-Wiggans (1989), Engelbrecht-Wiggans & Katok (2008), Filiz-Ozbay & Ozbay (2007), and Bergemann, Breuer, Cramton, Hirsch, Ndiaye & Ockenfels (2025) study auctions; Syam, Krishnamurthy & Hess (2008) and Zou, Zhou & Jiang (2020) investigate product differentiation; Vravosinos (2023) and Cerrone, Feri & Neary (2025) examine regret-related motives in some specific static games.




Figure 1: Two job-market game forms

Example 1 The game forms in Figures 1(i) & (ii) describe two subtly different situations: Bob received an exploding job offer from Agency A, but requested a deadline extension. Chance (=player 0) determines the state: with probability $p \in (\frac{3}{4}, \frac{6}{7})$ it is good (G): Bob can get a better offer elsewhere. (Agency) A can grant or deny the extension. In the former case, Bob will eventually get the best job available. In the latter case, he must take or reject A's offer, without knowing the state. Game forms (i) & (ii) differ as regards whether or not, conditional on play proceeding with (deny, take), Bob is informed about the state *ex post* (as indicated by the information set connecting those terminal nodes where he is not informed).⁵ A's payoffs (on top of Bob's) reflect its preferences: early taking (payoff = 2) trumps granting the extension (= 1) which trumps rejection (= 0). For Bob, the payoffs reflect his salary, not his utility which is affected also by regret. If avoiding regret is important enough, our prediction is (extend, reject) in (i) and (deny, take) in (ii). Behavior switches completely depending on the terminal information structure! The intuition: Bob prefers to take only in (ii) where he is sheltered from the pangs of regret that in (i) would plague him were he to take and learn that G occurred. As A anticipates this, the conclusion follows.

The key feature that distinguishes our approach is that players' utilities depend on their *ex post* beliefs about others' strategies: Regret-prone persons ponder what would have happened had they chosen differently. In games, this depends on others' strategies. Instances can occur where a player cannot know another's strategy, and to compute regret the player must consult beliefs. Figure 1(ii) exemplifies this; following choices *deny + take*, Bob cannot know the (realized) strategy of Chance. His belief regarding what would have happened had he chosen *reject* is therefore different in game forms (i) and (ii), and Bob's regret is calculated differently in the two cases.

⁵Whether (i) or (ii) is more realistic depends on, e.g., whether Bob has a friend who works "elsewhere".

Lack of information about others’ strategies can have two sources: *imperfect feedback* and *counterfactual choices*. Our analysis so far illustrates imperfect feedback, manifest in Figure 1(ii) via Bob’s non-singleton information set. Example 1 can also illustrate the role of counterfactual choices, if we instead focus on player A: If A chooses *extend*, then A cannot know the counterfactual choices of Bob. To compute how much to regret the choice *extend*, A must therefore form beliefs about what would have happened had A chosen *deny*. This belief-dependence of A’s utility (following *extend*) may seem vacuous, in the sense that said belief was pretty much pinned down via our earlier analysis of Bob’s behavior. However, as we show later, in other games there will be more flexibility both in how players may form beliefs and in how this impacts their regret, utility, and behavior.

The belief-dependence of utility transcends traditional game theory (where utility depends on players’ choices only) and places our exercise in the realm of the broader mathematical framework called *psychological game theory* (Geanakoplos, Pearce & Stacchetti 1989; Battigalli & Dufwenberg 2009).⁶ Our analysis uncovers a series of novel insights:

First, *information structure across terminal nodes* is critical to predictions. Example 1 illustrates: recall how “behavior switches completely depending on the terminal information structure.” This prediction stands in stark contrast to what can happen in traditional game theory, where such information is irrelevant (and therefore typically not made explicit).

Second, we uncover novel issues regarding chance moves that are moot without regret in the picture. A ruminating player, who contemplates what could have been, will need to ask himself whether counterfactual choices by Chance would have been perfectly correlated with realized choices. The answers naturally differ from one strategic situation to another. We argue that one needs a notion of *information sets of Chance* to handle the matter, and show how this affects predictions.

Third, we derive multiple results regarding the relevance or irrelevance of regret for predictions. In part, this exploration may be read as a contribution to DRT. A basic result in DRT is that if “regret costs” are *linear*, then DRT is, in fact, equivalent to traditional expected utility theory. The classical DRT setting may be seen as a special case of our general game-theoretic approach. We show that the linear-irrelevance result does not extend outside the setting explored in DRT, even for one-player game forms. When the analysis is extended to multi-player games, our sufficient conditions for linear-irrelevance concern (lack of!) imperfect feedback and counterfactual choices.

Fourth, fifth, etc., we convey many more results that relate to, e.g., best responses, sequential equilibria, and suggested directions for applied work.

Section 2 presents our general approach, covering finite extensive-game forms with monetary payoffs played by agents with perfect recall, and emphasizing relevant properties of the interim and ex post information structure. Section 3 analyzes game forms with a single personal player,

⁶Regret joins other motivations – e.g., other emotions, reciprocity, and social image – that also are belief-dependent, in different ways. See Battigalli & Dufwenberg (2022) for a survey of psychological game theory.

plus Chance; these simple structures allow us to highlight clearly some key signature features of our theory. Section 4 analyzes strategic interaction; we consider best replies and sequential equilibria, and prove several related results regarding when and how regret influences play. Section 5 analyzes information avoidance and acquisition. Section 6 presents examples intended to inspire applications, illustrating how regret may have significant impact in a variety of economic settings and how our general results can be applied in these settings. Section 7 compares our approach to decision-theoretic studies related to regret, including DRT. Section 8 concludes. The Appendix contains all of the proofs.

2 Modeling regret

We next define, in turn, game forms, beliefs, and players' regret and utility.

2.1 Game form

Game tree Let I be the nonempty, finite set of personal **players**. Allowing for games that contain chance moves, we let $0 \notin I$ denote the **Chance** pseudo-player and define $I_0 = I \cup \{0\}$. The rules of the game determine a finite tree $\bar{X} \subseteq \bigcup_{\ell=0}^L A^\ell$ of feasible sequences of action profiles closed under the weak **precedence** relation (“prefix-of”) $\preceq$, hence including the empty sequence $\emptyset$ as its **root**;⁷ L is an upper bound on the height of the tree, and $A = \bigcup_{\emptyset \neq J \subseteq I_0} \times_{j \in J} A_j$ is the set of action profiles and A_j is the set of (potentially feasible) **actions** of j . With this, $A(x) = \{a \in A : (x, a) \in \bar{X}\}$ is the set of **feasible action** profiles given $x \in \bar{X}$, $Z = \{x \in \bar{X} : A(x) = \emptyset\}$ is the set of complete, or **terminal** plays, and $X = \bar{X} \setminus Z$ is the set of partial, or **nonterminal** plays. There is a nonempty-valued **active player correspondence** $\iota : X \rightrightarrows I_0$ such that $A(x) = \times_{j \in \iota(x)} A_j(x)$ for each $x \in X$ and $j \in \iota(x)$. To avoid trivialities and at the same time simplify the statement of some results, we assume that $|A_i(x)| \geq 2$ for each $x \in X$ and every *personal* player $i \in \iota(x) \setminus \{0\}$. As for Chance, we assume that it has at least a “pseudo-move” at the root, $0 \in \iota(\emptyset)$ (allowing $|A_0(\emptyset)| = 1$), and $|A_0(x)| \geq 2$ for every $x \in X \setminus \{\emptyset\}$ with $0 \in \iota(x)$. This allows us to interpret $A_0(\emptyset)$ – when appropriate – as a set Ω of “states of nature”, letting $|A_0(\emptyset)| = 1$ represent the case of a commonly known state. If $|\iota(x)| \geq 2$ then multiple players choose simultaneously after play x . Since $\bar{X}$ (canonically ordered by the “prefix-of” relation $\preceq$) is a tree, we use “play” and **node** as synonymous. (Terminal plays/nodes are also called “end nodes.”) For each $j \in I_0$, $X_j = \{x \in X : j \in \iota(x)\}$ is the set of nodes where j is **active**.⁸ We say that there are **no chance moves** if $X_0 = \{\emptyset\}$ and Chance has a pseudo-move at the root, that is, $|A_0(\emptyset)| = 1$.

⁷Cf. Osborne & Rubinstein (1994) and Battigalli, Leonetti & Maccheroni (2020). Sequence $x = (x^1, \dots, x^k)$ weakly precedes, or is a prefix of, sequence $y = (y^1, \dots, y^\ell)$, written $x \preceq y$, if $k \leq \ell$ and $(x^1, \dots, x^k) = (y^1, \dots, y^k)$.

⁸Of course, Chance is only “pseudo-active” at the root if $|A_0(\emptyset)| = 1$.

We allow that a player obtains information even if he is not active. In particular, to model regret it is key to represent the information that players obtain at the end of the game.⁹ The **information** of each $i \in I$ is represented by a partition $\bar{\mathcal{H}}_i$ of $X_i \cup \{\emptyset\} \cup Z$ that refines $\{\{\emptyset\}, X_i \setminus \{\emptyset\}, Z\}$; that is, each i knows at least when the game starts and ends, and when i is active.¹⁰ Thus, denoting by $\bar{H}_i(y)$ the information set of i containing $y \in X_i \cup \{\emptyset\} \cup Z$, we have $\{\emptyset\} \in \bar{\mathcal{H}}_i$, $\bar{H}_i(x) \subseteq X_i$ for each $x \in X_i$, and $\bar{H}_i(z) \subseteq Z$ for each $z \in Z$. We also let $\mathcal{H}_i$ and $H_i(\cdot)$ respectively denote the restrictions of $\bar{\mathcal{H}}_i$ and $\bar{H}_i(\cdot)$ to the instances where i is active. It is also convenient to denote by $\mathcal{Z}_i$ the **terminal information partition** of i , that is, the subcollection of $\bar{\mathcal{H}}_i$ that partitions Z . Finally, we also introduce an information partition $\mathcal{H}_0$ of X_0 for Chance (the importance of this will be explained in Section 3, and there is no reason to specify the terminal information of Chance).

Since players must know what actions are possible at each node, the feasible-actions correspondence $A_j(\cdot)$ must be $\mathcal{H}_j$ -measurable, i.e., $H_j(x) = H_j(y)$ implies $A_j(x) = A_j(y)$ for all $j \in I_0$ and $x, y \in X_j$. Thus, it makes sense to write $A_j(h) = A_j(x)$ for all $j \in I_0$, $h \in \mathcal{H}_j$, and $x \in h$. We assume that $\mathcal{H}_0$ and each $\bar{\mathcal{H}}_i$ satisfy the *perfect recall* property. This implies that the information sets of each i are partially ordered by precedence.¹¹ In graphical representations, distinct nodes in the same information set of the relevant player are joined by a dashed line (see Example 1).

Game form To obtain a **game form** Γ – the mathematical structure that describes the rules of play – we add **Chance probabilities**, a *strictly positive* vector $(\pi_0(\cdot|h))_{h \in \mathcal{H}_0} \in \times_{h \in \mathcal{H}_0} \Delta^\circ(A_0(h))$,¹² and monetary (or material) **payoffs** for each $i \in I$ determined by function $m_i : Z \rightarrow \mathbb{R}$:

$$\Gamma = \left\langle I, \bar{X}, \iota, \pi_0, \mathcal{H}_0, (\bar{\mathcal{H}}_i, m_i)_{i \in I} \right\rangle.$$

To illustrate, in Example 1 there are no truly simultaneous moves: the players' map $\iota(\cdot)$ is

⁹In many real games (e.g., chess) players may also be inactive and yet obtain information as the play unfolds, but this is not important in our analysis of regret, because we assume that the only regret that matters is felt at the end of the game. Therefore, we follow the tradition by not representing such information. Although this simplification is innocuous given how we model regret, the non-terminal information of inactive players may instead be important for other psychological motivations. See, e.g., Battigalli & Dufwenberg (2022) and Battigalli & Generoso (2024).

¹⁰Even if i is not active at the root, it is convenient to include i 's information state at the root.

¹¹Note that for $i \in I$, perfect recall is a cognitive personal trait, not a property of the rules of the game. Our interpretation is that each $h \in \bar{\mathcal{H}}_i$ is a set of plays compatible with a **personal history** of signals received and actions played by $i \in I$, which makes $\bar{\mathcal{H}}_i$ objectively determined by the rules of the game. A sufficient condition that justifies this interpretation is that i 's memory (a personal trait) is perfect. See Osborne & Rubinstein (1994, ch. 12) and Battigalli & Generoso (2024).

¹²For any finite set C , we let $\Delta(C) = \{\gamma \in \mathbb{R}_+^C : \sum_{c \in C} \gamma(c) = 1\}$ denote the set of all probability mass functions on C , and we let $\Delta^\circ(C) = \Delta(C) \cap \mathbb{R}_{++}^C$ denote its relative interior.

single-valued and obvious from the figures.¹³ The set of terminal nodes (complete plays) is

$$Z = \{(G, ext), (\neg G, ext), (G, deny, take), (\neg G, deny, take), (G, deny, rej), (\neg G, deny, rej)\},$$

and X (resp., $\bar{X}$) is made of all the predecessor/prefixes (resp., weak predecessors/prefixes) of Z , including the empty sequence (root) $\emptyset$. The nonterminal information structure satisfies, for each realization (state) $\omega \in \{G, \neg G\}$, $H_A(\omega) = \{G, \neg G\}$ and $H_B(\omega, deny) = \{(G, deny), (\neg G, deny)\}$. The terminal information structure of (i) is perfect (made of singletons for each personal player), while the terminal information structure of (ii) satisfies

$$\bar{H}_B(\omega, deny, take) = \{(G, deny, take), (\neg G, deny, take)\}$$

for each $\omega \in \Omega = A_0(\emptyset) = \{G, \neg G\}$. Chance probabilities and payoffs are obvious from the pictures.

Strategies Information-dependent behavior is described by **strategies**, i.e., for each $j \in I_0$, functions $h \mapsto s_j(h) \in A_j(h)$, with $h \in \mathcal{H}_j$. Thus, the set of strategies of $j \in I_0$ is the cross-product $S_j = \times_{h \in \mathcal{H}_j} A_j(h)$. We let $S = \times_{i \in I} S_i$ and $S_{-i} = \times_{j \in I \setminus \{i\}} S_j$ respectively denote the sets of strategy profiles of personal players and of personal co-players of $i \in I$. The **path function** $\zeta : S \times S_0 \rightarrow Z$ specifies the path (complete play, end node) $z = \zeta(s, s_0)$ induced by each profile (s, s_0) .¹⁴ With this, $S_{I_0}(x) = \{(s, s_0) : x \preceq \zeta(s, s_0)\}$ is the set of profiles inducing play/node x . It can be checked that $S_{I_0}(x) = \times_{j \in I_0} S_j(x)$, where $S_j(x)$ is the projection of $S_{I_0}(x)$ onto S_j . The set of strategy profiles reaching information set h is $S_{I_0}(h) = \bigcup_{x \in h} S_{I_0}(x)$. Unlike $S_{I_0}(x)$, if $|I_0| \geq 3$, it may be impossible to factorize $S_{I_0}(h)$ into its single-player projections. But perfect recall implies that $S_{I_0}(h) = S_i(h) \times S_{I_0 \setminus \{i\}}(h)$ and $S_i(x) = S_i(h)$ for all $x \in h \in \bar{\mathcal{H}}_i$. Thus, the factorization holds in games with $|I_0| = 2$.

Information Interim and ex post information play a key role in our analysis. Here we describe some relevant properties of the information structure. We say that a game form has **essentially simultaneous moves** if each player effectively makes a single move, without conditioning on information concerning previous actions of other players. More formally, each personal player $i \in I$ is active at a unique information set h ($\mathcal{H}_i = \{h\}$), Chance has at most one information set at which it makes a nontrivial move, and – for all $z \in Z$ and $h \in \bigcup_{i \in I} \mathcal{H}_i$ – there is a (unique) $x \in h$ (Chance need not move only at the root). Thus, the path function $\zeta : S \times S_0 \rightarrow Z$ is a bijection,

¹³It may be argued that the moves of Chance and A are “essentially simultaneous” and one can show that our analysis is invariant to interchanging essentially simultaneous moves. See Battigalli *et al.* (2020). Some of the following examples feature truly simultaneous moves (i.e., $\iota(x) \geq 2$ for some $x \in X$).

¹⁴For all $\ell \in \{1, \dots, L\}$ and $(a^1, \dots, a^\ell) \in Z$, $\zeta(s, s_0) = (a^1, \dots, a^\ell)$ if $a^1 = (s_j(\{\emptyset\}))_{j \in \iota(\emptyset)}$ and $a^{k+1} = (s_j(H_j(a^1, \dots, a^k)))_{j \in \iota(a^1, \dots, a^k)}$ for each $k \in \{1, \dots, \ell - 1\}$.

i.e., $S \times S_0$ and Z are isomorphic, and terminal nodes are essentially action profiles.¹⁵

Next we consider extensive-form representations of simultaneous Bayesian games with monetary payoffs, that is, game forms in which Chance moves first and personal players may have imperfect and asymmetric information about its realization. A game form features **asymmetric information and essentially simultaneous moves** by personal players if only Chance moves at the root ($\iota(\emptyset) = \{0\}$), Chance has no other move ($X_0 = \{\emptyset\}$), and – for every $\omega \in \Omega := A_0(\emptyset)$ and $i \in I$ – there is a unique nonterminal information set $h \in \mathcal{H}_i$, with $A_i(h) = A_i$, containing a node following ω ; furthermore, for every $z \in Z$ following ω , there is a unique $x \in h$ following ω and preceding z .¹⁶

Three nested special cases of information structure are worth mentioning, listed from most to least general; game form Γ features:

- **observed deviators** if $S_{I_0}(h) = \times_{j \in I_0} S_j(h)$ for each $h \in \bigcup_{i \in I} \bar{\mathcal{H}}_i$ (see Battigalli, 1996);
- **observed actions** (or “almost perfect information”) if $H_j(x)$ is a singleton for all $x \in X$ and $j \in \iota(x)$, and $\bar{H}_i(z)$ is a singleton for all $z \in Z$ and $i \in I$;
- **perfect information** if there are *observed actions* and, for all $x \in X$, there is only one active player, i.e., $\iota : X \rightarrow I_0$ is a function.¹⁷

Since regret depends on *terminal* information, it is also relevant to consider associated properties. Game form Γ features:

- **observed payoffs** if each monetary payoff function m_i is $\bar{\mathcal{H}}_i$ -measurable, i.e., $\bar{H}_i(z') = \bar{H}_i(z'')$ implies $m_i(z') = m_i(z'')$ for all $z', z'' \in Z$ and $i \in I$;
- **perfect feedback** if $\bar{H}_i(z)$ is a singleton for all $z \in Z$ and $i \in I$;
- **ex post observed actions of personal players** if the game form features *asymmetric information and essentially simultaneous moves* of personal players, and, for all *distinct* terminal nodes $z', z'' \in Z$ and personal players $i \in I$ it holds that $\bar{H}_i(z') = \bar{H}_i(z'')$ *if and only if* both of the following conditions are satisfied: (a) there are *distinct* realizations $\omega', \omega'' \in \Omega = A_0(\emptyset)$, an information set $h \in \mathcal{H}_i$, and distinct plays $x', x'' \in h$ such that $(\omega') \preceq x' \prec z'$ and $(\omega'') \preceq x'' \prec z''$; (b) for every $j \in I$ and any $y', y'' \in X_j$ such that $y' \prec z', y'' \prec z''$ we have that $a_j(y' \rightarrow z') = a_j(y'' \rightarrow z'')$, where $a_j(y \rightarrow z)$ is the action of j at y leading to z .

¹⁵Game trees with essentially simultaneous moves can be “simultanzied”: In the derived game tree with “truly” simultaneous moves, $\iota(\emptyset) = I_0$ and $H_j = \{\{\emptyset\}\}$, $A_j(\emptyset) = A_j(x)$ for any $j \in I_0$ and any node $x \in h$ of the *original* game tree with “essentially” simultaneous moves. Strategy sets and path functions of the two game trees are the same up to re-labeling. See Battigalli *et al.* (2020).

¹⁶This property fails in Example 1 because $z = (G, \text{extend})$ does not cross Bob’s information set.

¹⁷Traditional extensive-form representations assume that there is always at most one active player, which eliminates the difference between observed actions and perfect information. Osborne & Rubinstein (1994, pp. 102 and 115) call “perfect information” the weaker property of observed actions.

Note that observed actions (hence, also perfect information) imply perfect feedback and observed payoffs. To illustrate, in both versions of Example 1 there are observed deviators¹⁸ and observed payoffs. The first version satisfies perfect feedback, the second does not. Neither perfect information nor observed actions holds in either version.

The observed-payoffs assumption is natural in many applications, but we regard it as a property that may or may not hold: In some situations terminal information does not resolve all the uncertainty about payoffs. A juror, e.g., may care about but never observe whether a defendant is guilty. Or, to take an example we explore (Section 6.2), someone concerned with global warming may care about consequences that are realized only after his death. Ex post observed actions of personal players holds in the extensive-form representations of simultaneous Bayesian games with monetary payoffs when players' actions are revealed ex post, but their private information is not revealed. With private values (i.e., the payoff of each $i \in I$ is measurable with respect to the private information of i), this further implies observed payoffs. If $|\Omega| = 1$, i.e., there are no chance moves, ex post observed actions of personal players is equivalent to requiring both essentially simultaneous moves and perfect feedback.

Next we consider the feedback about co-players' strategies induced by terminal information sets of player $i \in I$ and its dependence on the strategy played by i . The path function ζ and the terminal information structure of i determine a **feedback map**

$$\begin{aligned} F_i : S \times S_0 &\rightarrow \mathcal{Z}_i \\ (s, s_0) &\mapsto \bar{H}_i(\zeta(s, s_0)) \end{aligned}$$

that associates each strategy profile (s, s_0) with the corresponding terminal information set $\bar{H}_i(\zeta(s, s_0))$ of i . Now consider the section $F_{i,s_i} : S_{-i} \times S_0 \rightarrow \mathcal{Z}_i$ of this map at any strategy s_i , that is, $(s_{-i}, s_0) \mapsto F_{i,s_i}(s_{-i}, s_0) = \bar{H}_i(\zeta(s_i, s_{-i}, s_0))$. Taking the preimages of terminal information sets reachable under s_i (those in the image collection $F_{i,s_i}(S_{-i} \times S_0) \subseteq \mathcal{Z}_i$), we obtain the partition of $S_{-i} \times S_0$ that describes the information about the behavior of others that i may obtain if he plays s_i :

$$\mathcal{F}_i(s_i) := \{F_{i,s_i}^{-1}(\bar{h})\}_{\bar{h} \in F_{i,s_i}(S_{-i} \times S_0)}.$$

Realization-equivalent strategies induce the same feedback about co-players' strategies:

Remark 1 Fix any pair of strategies $(s'_i, s''_i) \in S_i \times S_i$; if $\zeta(s'_i, s_{-i}, s_0) = \zeta(s''_i, s_{-i}, s_0)$ for all $(s_{-i}, s_0) \in S_{-i} \times S_0$, then $\mathcal{F}_i(s'_i) = \mathcal{F}_i(s''_i)$.

So, when considering the informativeness of strategies, we can focus on the corresponding **reduced strategies**, i.e., their realization-equivalence classes.

¹⁸For example, for each realization $\omega \in \{G, \neg G\} = S_0$, $S(H_B(\omega, deny)) = \{deny\} \times S_B \times S_0$ and, in the second version, $S(H_B(\omega, deny, take)) = \{take\} \times \{deny\} \times S_0$.

The next property can help us identify cases where maximizing expected psychological utility is equivalent to maximizing expected monetary payoff, so that concerns about regret do not affect behavior. Player i in Γ has **own-strategy independent feedback about others** if $\mathcal{F}_i(s'_i) = \mathcal{F}_i(s''_i)$ for all $s'_i, s''_i \in S_i$. With this, game form Γ features

- **own-strategy independence of feedback about others** (OSIFO) if each personal player's feedback about others is independent of his played strategy.

OSIFO is a strong property that typically fails when a personal player i obtains interim information about other personal players' moves, as in Example 1, where *extend* prevents the ex post observation of Bob's strategy, whereas *deny* allows it.¹⁹ However, OSIFO holds in relevant economic situations with (essentially) simultaneous moves, e.g., Cournot oligopolies with a known market demand function, where firms can back out the total output of competitors from the observed market price, or public good games with a known production function, where agents can infer others' total contributions from the observed level of the public good. The following observation is relevant for our results:

Remark 2 *Own-strategy independence of feedback about others holds if the game form features (1) essentially simultaneous moves and perfect feedback, or (2) asymmetric information and essentially simultaneous moves with ex post observed actions of personal players.*

Furthermore, by Remark 1, OSIFO is invariant under transformations of the game tree that do not change (up to isomorphisms) the map from reduced strategies to end nodes, that is, transformations in the two groups of (i) interchanging (or “simultanizing”) essentially simultaneous moves, and (ii) coalescing sequential moves of the same player, together with its inverse, sequential-agents splitting; see Battigalli, Leonetti, & Maccheroni (2020).²⁰

Mixed strategy of Chance While $s \in S$ is endogenously determined by the strategic analysis of the game, any $s_0 \in S_0$ is determined at random according to the chance probability vector π_0 ; i.e., we obtain the (exogenous) “mixed strategy of Chance” $\sigma_0 \in \Delta^\circ(S_0)$ by means of the standard transformation of behavior strategies into mixed strategies under the assumption that chance moves are independent across information sets:²¹

$$\sigma_0(s_0) = \prod_{h \in \mathcal{H}_0} \pi_0(s_0(h) | h). \quad (1)$$

¹⁹It is notable that this property is also relevant for the self-confirming equilibrium (SCE) concept, for which ex post information is key. Indeed, OSIFO and observed payoffs imply that SCE is equivalent to Nash equilibrium. See Battigalli, Cerreia, Maccheroni, & Marinacci (2015, 2016); the latter reference offers a detailed analysis of information feedback.

²⁰For instance, if a game form Γ can be transformed into Γ' with essentially simultaneous moves and ex post observed actions of personal players by coalescing sequential moves by the same players, then Γ satisfies OSIFO whenever Γ' does.

²¹See Kuhn (1953), and also Section 3 where we explain why we need “information sets of Chance.”

2.2 Beliefs

Players' beliefs are central to describing both belief-dependent utilities that incorporate regret and how players analyze the game. Conditional on observing $h \in \bar{\mathcal{H}}_i$, player $i \in I$ infers that $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)$ and holds belief $\alpha_i(\cdot | S_{I_0 \setminus \{i\}}(h)) =: \alpha_i(\cdot | h) \in \Delta(S_{I_0 \setminus \{i\}}(h))$ about his co-players' strategies (information-dependent actions). With this, $\alpha_i = (\alpha_i(\cdot | h))_{h \in \bar{\mathcal{H}}_i}$ is the **system of (first-order) beliefs** of i . Clearly, such subjective beliefs must be disciplined by the rules of conditional probabilities and the knowledge of the chance probability function π_0 , or the corresponding mixed strategy σ_0 . Therefore, we impose the following **consistency** condition on α_i : there is a converging sequence of strictly positive measures on personal co-players' strategy profiles $(\sigma_{-i}^n)_{n=1}^\infty \in \Delta^\circ(S_{I \setminus \{i\}})^\mathbb{N}$ such that

$$\alpha_i(s_{-i}, s_0 | h) = \lim_{n \rightarrow \infty} \frac{\sigma_{-i}^n(s_{-i}) \sigma_0(s_0)}{\sum_{(s'_{-i}, s'_0) \in S_{I_0 \setminus \{i\}}(h)} \sigma_{-i}^n(s'_{-i}) \sigma_0(s'_0)} \quad (2)$$

for all $h \in \bar{\mathcal{H}}_i$ and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)$. This essentially expresses the assumption that player i knows σ_0 (that is, π_0) and regards the strategic play of personal co-players and Chance as stochastically independent (cf. Battigalli 1996; Kohlberg & Reny 1997). We let $\Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}})$ denote i 's set of consistent (first-order) beliefs.

To gain intuition, suppose that, for all $h \in \bar{\mathcal{H}}_i$, i obtains logically independent information about Chance and i 's other co-players, i.e., $S_{I_0 \setminus \{i\}}(h) = S_{I \setminus \{i\}}(h) \times S_0(h)$, a weakening of the observed-deviators assumption. Then, condition (2) implies that α_i be obtained from σ_0 and a conditional probability system $\alpha_{i,-i}$ on $(S_{I \setminus \{i\}}, \bar{\mathcal{H}}_i)$ so that

$$\alpha_i(s_{-i}, s_0 | h) = \alpha_{i,-i}(s_{-i} | h) \frac{\sigma_0(s_0)}{\sigma_0(S_0(h))}, \quad (3)$$

where a **conditional probability system** on $(S_{I \setminus \{i\}}, \bar{\mathcal{H}}_i)$ is $(\alpha_{i,-i}(\cdot | h))_{h \in \bar{\mathcal{H}}_i} \in \times_{h \in \bar{\mathcal{H}}_i} \Delta(S_{I \setminus \{i\}}(h))$ satisfying the *chain rule*: for all $h, h' \in \bar{\mathcal{H}}_i$ and $s_{-i} \in S_{I \setminus \{i\}}$, if $s_{-i} \in S_{I \setminus \{i\}}(h') \subseteq S_{I \setminus \{i\}}(h)$,²²

$$\begin{aligned} \alpha_{i,-i}(s_{-i} | h) & : = \alpha_{i,-i}(s_{-i} | S_{I \setminus \{i\}}(h)) = \alpha_{i,-i}(s_{-i} | S_{I \setminus \{i\}}(h')) \alpha_{i,-i}(S_{I \setminus \{i\}}(h') | S_{I \setminus \{i\}}(h)) =: \\ & \alpha_{i,-i}(s_{-i} | h') \alpha_{i,-i}(S_{I \setminus \{i\}}(h') | h). \end{aligned}$$

In particular, (2) implies that if two distinct information sets $h, h' \in \bar{\mathcal{H}}_i$ differ only because of i 's past actions and hence give the same information about co-players, then $S_{I_0 \setminus \{i\}}(h) = S_{I_0 \setminus \{i\}}(h')$ and i must hold the same beliefs conditional on both information sets, $\alpha_i(\cdot | h) = \alpha_i(\cdot | h')$. This is easier to see in the special case covered by (3).

²²That is, h is less informative than h' about i 's co-players $j \in I$, e.g. (but not necessarily), because h precedes h' .

2.3 Regret and utility

At terminal information sets, players obtain information about their material payoffs and experience (painful) regret as they consult their terminal beliefs and ruminate about what would have happened had they chosen differently. They then take into account information they have come to obtain about which choices others, including Chance, made.

Definition 1 The *regret* of player $i \in I$ at terminal information set $h \in \mathcal{Z}_i$ given beliefs α_i is

$$r_i(h, \alpha_i) := \phi_i \left(\max_{s_i \in S_i} \left\{ \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | h) m_i(\zeta(s_i, s_{-i}, s_0)) \right\} - \bar{m}_i(h, \alpha_i) \right), \quad (4)$$

where function $\phi_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, strictly increasing and satisfies $\phi_i(0) = 0$, and

$$\bar{m}_i(h, \alpha_i) := \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | h) m_i(\zeta(s'_i, s_{-i}, s_0)) \quad (5)$$

for any $s'_i \in S_i(h)$. We say that i 's regret is **linear** if ϕ_i is the identity function.

Perfect recall implies that $\bar{m}_i(h, \alpha_i)$ is well defined, because the r.h.s. of (5) does not depend on which s'_i one picks in $S_i(h)$. This follows as $S_i(h) = S_i(z)$ for every $z \in h$. If Γ features observed payoffs, then $\bar{m}_i(h, \alpha_i)$ is independent of α_i , because i knows the material payoff implied by h , (m_i is constant on h). Formula (4) subsumes as a special case Battigalli & Dufwenberg's (2009) definition for two-player game forms with observed actions and without chance moves.

Definition 2 The *utility* of player $i \in I$ is the function $u_i : Z \times \Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}}) \rightarrow \mathbb{R}$ defined by

$$u_i(z, \alpha_i) = m_i(z) - \theta_i r_i(\bar{H}_i(z), \alpha_i), \quad (6)$$

for all $z \in Z$ and consistent $\alpha_i \in \Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}})$, where $\theta_i \geq 0$ measures i 's "regret sensitivity."

Several features of u_i warrant commentary: First, we assume that u_i is linear in i 's monetary payoff. We could, alternatively, have allowed for non-linearity (e.g., to capture classical risk aversion), by replacing $m_i(z)$ with $v_i(m_i(z))$ where $v_i : [\min_{z \in Z} m_i(z), \max_{z \in Z} m_i(z)] \rightarrow \mathbb{R}$ is strictly increasing. We might then also have substituted $v_i(m_i(z))$ and $\mathbb{E}_{\alpha_i}[v_i(m_i(\cdot)) | h]$ for $m_i(z)$ and $\bar{m}_i(h, \alpha_i)$ in formula (4) of Definition 1. The reason for our choice is that we wish to focus on one behavioral idea at a time, in our case regret rather than risk aversion. We will, however, consider alternative formulations when we compare our approach to DRT (see Section 7.1, and also Section 6.2), as scholars in that tradition have made such assumptions. Loomes & Sugden (1982), e.g., refer to (a function like) v_i as a "choiceless utility."

Second, from a mathematical point of view we could, obviously, have merged ϕ_i and θ_i into a new function. Doing so would, however, obscure the psychological interpretation. Function ϕ_i

captures how pangs of regret vary with the size of the monetary loss relative to what could-have-been. Parameter θ_i captures i 's overall concern for avoiding regret; its separate presence facilitates comparative statics on how concern for regret shapes outcomes.

Third, note that u_i is a belief-dependent utility, by virtue of how α_i affects r_i in (6). Recall that α_i is a system of beliefs about other players (including Chance), and it does not contain a belief about i 's own behavior interpreted as his plan. Therefore, utility function (6) satisfies “own-plan independence” and represents dynamically consistent preferences (see the related discussion in Battigalli, Corrao, & Dufwenberg 2019 and Battigalli & Dufwenberg 2022).

A game form Γ coupled with utilities $u = (u_i)_{i \in I}$, as given by (6), comprises a **regret game** $G = (\Gamma, u)$, which we say is **linear** if regret is linear for every $i \in I$. We assume that each player i maximizes the expectation of u_i given α_i .

Equation (4) highlights that players' regret, and consequently their utilities, depend on their beliefs. The following example illustrates this relationship.²³




Figure 2: A three-person game form

Example 2 Consider Figure 2, a three-player game form. The payoffs of players j and k will not matter for the points we make, and the figure should be read as not disclosing that. Furthermore, given perfect recall, assumptions about post-play information do not matter for the following calculations. Assume linear regret. What is player i 's utility of Out? Despite that Out ends the game, the answer is not a given number. Rather, it depends on i 's beliefs about the choices of players j and k . For example, if player i assigns probability 1 to (B, A) then his utility of Out equals $1 - \theta_i \cdot (2 - 1) = 1 - \theta_i$, while if he assigns probability 1 to (A, B) then said utility equals 1. In general,

$$u_i(\text{Out}, \alpha_i) = 1 - \theta_i (\max \{ \alpha_i(\mathbf{A}, \mathbf{A} | \text{Out}) + \alpha_i(\mathbf{B}, \mathbf{B} | \text{Out}) + 2\alpha_i(\mathbf{B}, \mathbf{A} | \text{Out}), 1 \} - 1),$$

where $\alpha_i(s_j, s_k | h)$ is the subjective probability assigned to (s_j, s_k) conditional on h .

²³We denote with boldface letters like $\mathbf{a}_j$ conditional statements like “if h occurred, action a_j would be chosen,” where $h \in \mathcal{H}_j$, $a_j \in A_j(h)$; i.e., bold letters in examples denote strategies, or parts of them.

We next relate (4) to feedback about others. Consider player $i \in I$ who just played strategy $\bar{s}_i$, observed $\bar{h} \in \mathcal{Z}_i$, and ruminates with the benefit of hindsight about his (expected) payoff had he played strategy s_i instead. He infers that co-players played a strategy profile in subset $S_{I_0 \setminus \{i\}}(\bar{h}) = F_{i, \bar{s}_i}^{-1}(\bar{h}) \in \mathcal{F}_i(\bar{s}_i)$ – where $\mathcal{F}_i(\bar{s}_i)$ is the ex post feedback partition defined in Section 2.1 – and holds terminal belief $\alpha_i(s_{-i}, s_0 | \bar{h}) := \alpha_i(s_{-i}, s_0 | S_{I_0 \setminus \{i\}}(\bar{h})) \in \Delta(S_{I_0 \setminus \{i\}}(\bar{h}))$. This shows that the hindsight-expected payoff of s_i , a key component of the regret term, generally depends on the played strategy $\bar{s}_i$. However, when two strategies induce the same ex post information partition of $S_{-i} \times S_0$, they yield the same hindsight-expected payoff. Let

$$\bar{M}_i(\bar{s}_i, \bar{s}_{-i}, \bar{s}_0, s_i, \alpha_i) := \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}_i, \bar{s}_{-i}, \bar{s}_0)))} \alpha_i(s_{-i}, s_0 | \bar{H}_i(\zeta(\bar{s}_i, \bar{s}_{-i}, \bar{s}_0))) m_i(\zeta(s_i, s_{-i}, s_0))$$

denote the hindsight-expected payoff from ruminating about strategy s_i , given that $(\bar{s}_i, \bar{s}_{-i}, \bar{s}_0)$ has been played.

Lemma 1 *For every consistent belief system α_i , for all $\bar{s}'_i, \bar{s}''_i \in S_i$, and $(\bar{s}_{-i}, \bar{s}_0) \in S_{I_0 \setminus \{i\}}$, if $\mathcal{F}_i(\bar{s}'_i) = \mathcal{F}_i(\bar{s}''_i)$ then $\bar{M}_i(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0, s_i, \alpha_i) = \bar{M}_i(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0, s_i, \alpha_i)$.*

If player i has OSIFO, then $\mathcal{F}_i(\bar{s}'_i) = \mathcal{F}_i(\bar{s}''_i)$ for all $\bar{s}'_i, \bar{s}''_i \in S_i$. Therefore, the hindsight-expected payoff from ruminating about s_i is the same whether $\bar{s}'_i$ or $\bar{s}''_i$ is the actually played strategy.

Corollary 1 *If the game form features own-strategy independence of i 's feedback about others, then, for every consistent belief system α_i , all $\bar{s}'_i, \bar{s}''_i \in S_i$, $(\bar{s}_{-i}, \bar{s}_0) \in S_{-i} \times S_0$, and $s_i \in S_i$, $\bar{M}_i(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0, s_i, \alpha_i) = \bar{M}_i(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0, s_i, \alpha_i)$.*

3 One-person games

The influence of regret on behavior is manifest via comparative statics on predictions as θ_i changes. One-person games provide the cleanest context for gaining several key insights into this relationship, so we start there and begin with an example:




Figure 3: A game form where Chance is a follower

Example 3 Consider Figure 3, a game form with a single personal player and Chance, which features perfect feedback (in this case, equivalent to observed payoffs). Assume linear regret and that $\lambda = \frac{2}{5}$. If player i chooses s , then he cannot learn what Chance would have chosen had he selected r . Player i 's belief about Chance's strategy is as follows:

$$\alpha_i(\mathbf{g}|s) = 1 - \alpha_i(\mathbf{b}|s) = \frac{3}{5}.^{24}$$

Given this belief, player i 's linear regret for choosing s is:

$$r_i(s, \alpha_i) = \left(\frac{3}{5} \cdot 2 + \frac{2}{5} \cdot 0 \right) - 1 = \frac{6}{5} - 1 = \frac{1}{5}.$$

According to (6), the utility of choosing s is

$$u_i(s, \alpha_i) = 1 - \frac{\theta_i}{5}$$

and the expected utility of choosing r is

$$u_i(r, \alpha_i) = \frac{3}{5} \cdot 2 + \frac{2}{5} \cdot [0 - \theta_i \cdot (1 - 0)] = \frac{6}{5} - \frac{2}{5}\theta_i.$$

Condition (3) determines beliefs in one-person games with Chance; this explains why α_i does not appear explicitly on the right-hand side of this expression. Comparing expected utilities one sees that player i 's preferred choice between r and s depends on whether $\theta_i < 1$ or $\theta_i > 1$.

Example 3 provides a first illustration that regret can matter to predictions. When and how can this happen more generally? The following “irrelevance result” – a special case of Theorem 1 in Section 4.1 – will prove useful for providing (eventually, also more general) answers:

²⁴Formally, this equation is implied by consistency of beliefs with Chance probabilities (Equation (3) can be used in this case). We will not repeat similar observations concerning consistency in the following examples.

Proposition 1 *Suppose a one-person game form Γ features essentially simultaneous moves and perfect feedback. If regret is linear, then the strategies (actions) of player $i \in I$ that maximize the expectation of u_i are the same as those that maximize the expectation of m_i .*




Figure 4: A one-person essentially simultaneous game form with perfect feedback

Example 4 *To illustrate Proposition 1, consider Figure 4, where $\lambda \in (0, 1)$ is the probability of choice b . By (6), player i prefers s over r iff*

$$\begin{aligned} \lambda \cdot 1 + (1 - \lambda) \cdot (1 - \theta_i(2 - 1)) &\geq \lambda \cdot (0 - \theta_i(1 - 0)) + (1 - \lambda) \cdot 2 \Leftrightarrow \\ (2\lambda - 1)(1 + \theta_i) &\geq 0 \Leftrightarrow \lambda \geq \frac{1}{2}. \end{aligned}$$

Regret is irrelevant to the prediction; player i maximizes expected payoff independently of θ_i .

The irrelevance of regret under the conditions of Proposition 1 concerns implied behavior, not that feelings of regret are absent when a game is played. Example 4 illustrates this, since i experiences regret following paths (b, r) and (g, s) . Relaxing the conditions of Proposition 1 reveals scenarios where predictions vary with θ_i . Example 3 demonstrates that regret can matter for games that do not feature essentially simultaneous moves. Similarly, regret can matter in the presence of imperfect feedback:




Figure 5: A one-person essentially simultaneous game form with imperfect feedback

Example 5 Consider Figure 5, which looks like Figure 4 except for the information structure over end-nodes. Assume that $\frac{1}{3} < \lambda < \frac{1}{2}$. By (6), player i prefers s over r iff

$$\begin{aligned} 1 - \theta_i[\lambda \cdot 0 + (1 - \lambda) \cdot 2 - 1] &\geq \lambda \cdot (-\theta_i) + (1 - \lambda) \cdot 2 \Leftrightarrow \\ \theta_i &\geq \frac{1 - 2\lambda}{3\lambda - 1}. \end{aligned} \quad (7)$$

Unlike in Example 4, this example does not feature perfect feedback, as the terminal information set $\bar{H}_i(b, s) = \bar{H}_i(g, s)$ is not a singleton. In the associated linear regret game, for any $\lambda \in (\frac{1}{3}, \frac{1}{2})$ player i prefers s to r if θ_i is high enough. And the higher is θ_i , the lower λ must be for i to be willing to choose r .

We next discuss five related key insights:

Belief-dependent utility First, in Examples 4 and 5, player i 's utility following choice s depends on his belief regarding what he would have gotten had he chosen r . This is the feature that places our exercise within the mathematical framework of psychological game theory. Moreover, since the involved belief differs between Example 4 and 5, this explains why we may get different predictions in the two cases (if θ_i is high enough).²⁵

Terminal information Second, in traditional game theory, the information structure across end-nodes is irrelevant (and therefore usually not specified). Our theory does not have that property; a comparison of Examples 4 and 5 makes that clear. The two game forms differ only in the

²⁵Note that the effect is analogous to the one we sketched in Example 1, for the choice *take*, helping us get a deeper insight also there. In both cases, a belief-dependent preference is responsible for regret to matter.

terminal information structure, and yet the predictions may be different (for high enough θ_i).²⁶

Timing To cover this third topic, we need a new example:




Figure 6: Chance is a follower with perfect information

Example 6 Consider Figure 6. Its game form does not feature essentially simultaneous moves. Chance is a follower with perfect information and four strategies. Each terminal node reveals only the action taken by Chance subsequent to the action chosen by player i , not the action that would have been taken following the foregone action. Assume linear regret and that $\frac{1}{3} < \lambda < \frac{1}{2}$. According to (6), player i prefers s over r if and only if (7) holds; decisions depend on θ_i .

Compare Examples 4 and 6. The involved game forms differ as regards the order of moves, but they share the property that player i chooses between s and r without information about Chance's choice. If θ_i is high enough, this subtle distinction leads to different predictions!

Information set of Chance Fourth, and related to the timing issue just illustrated, modeling Chance moves requires special care. To explain why, we need one more example:

²⁶Again, the effect is analogous to, and thus helps explain the one we sketched in Example 1.




Figure 7: Chance is a follower with imperfect information

Example 7 See Figure 7. Again, the order of moves has been flipped relative to Figure 4, but Chance is now a follower with imperfect information. The “interchange of essentially simultaneous moves” transformation applies. Assume linear regret. As in Example 4, player i prefers s over r iff $\lambda \geq \frac{1}{2}$.

Compare Examples 6 and 7. The difference concerns a form of counterfactual correlation: If player i chooses s and learns that Chance chose g , does that imply that Chance would also have chosen action g had player i chosen r ? The answer is no in Figure 6 and yes in Figure 7. Note that either figure may make economic sense: Figure 7 might reflect that the chance move corresponds to an outcome in the stock market, which is independent of whether player i invests (i.e., chooses r) or not (chooses s). Figure 6 could instead describe a situation where r corresponds to player i entering a casino and placing a bet, whereas s involves going home and flipping a coin in the kitchen! If he goes home, then what would have happened had he entered the casino player i will never know, and it is independent of his coin flip. For many theories of decision making, player i would make the same choice in Figures 6 and 7 (and for this reason information sets of Chance are usually not specified). This is not the case in our theory; the psychology of rumination and regret will be different in the two cases. In Example 7, but not in Example 6, following path (s, g) , player i would know that he would have made a lot of money had he instead chosen r . Hence, following (s, g) , player i would feel considerable regret in Example 7 and less regret in Example 6, and (as we have shown) the behavior in the two cases is (for high enough θ_i) different.

Risk aversion Fifth, the effect of increasing θ_i , in (e.g.) Example 5, may look like an increase in risk aversion. However, this is a new breed, not based on a dislike for mean-preserving spreads represented by the concavity of the utility of money. To see this, study how player i would behave in the game forms of Examples 4 and 5 if his utility were a concave transformation of the given

numbers, e.g., a “CRRA utility”: $v_i(x) = x^{\frac{1}{\gamma}}$ with $\gamma > 1$. We can here note that “the higher is risk aversion parameter γ , the lower λ must be for i to be willing to gamble.” However, *the effect would be the same in the game forms of both figures!* By contrast, with u_i given by (6) predictions differ. The issue is not concavity, but rather how pangs of regret get soothed if information about counterfactual consequences is hidden.²⁷ In Example 5, lack of information after s shelters player i from much of the regret that would plague him if he knew whether Chance had chosen b or g .²⁸

4 Strategic interaction

To model behavior in general game forms we first analyze best-reply correspondences. Indeed, most of our results hold for best replies. Then we move on and – to keep readers in their comfort zone – we use the traditional sequential equilibrium (SE) solution concept.²⁹

4.1 Best replies

We rely on a sequential version of the best-reply concept, which requires some preliminary definitions. For any player $i \in I$ and nonterminal information set $h \in \mathcal{H}_i$, let

$$\mathcal{H}_i^{|h} := \{h' \in \mathcal{H}_i | h' \not\preceq h\}$$

denote the collection of nonterminal information sets of i that do not precede h .³⁰ For any $h \in \mathcal{H}_i$ and $s_i \in S_i$, let $s_i^{|h}$ denote the h -replacement of s_i , i.e., the strategy that coincides with s_i on $\mathcal{H}_i^{|h}$ and selects i ’s action leading to h at each predecessor $h' \in \mathcal{H}_i \setminus \mathcal{H}_i^{|h}$ (if any). With this, for any consistent system of beliefs $\alpha_i \in \Delta^{\tilde{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}})$ we consider the set of (pure) sequential best replies

$$BR_i(\alpha_i) := \left\{ s_i \in S_i : \forall h \in \mathcal{H}_i, s_i^{|h} \in \arg \max_{s'_i \in S_i(h)} \mathbb{E}_{s'_i, \alpha_i}[u_i|h] \right\},$$

where

$$\mathbb{E}_{s'_i, \alpha_i}[u_i|h] := \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} u_i(\zeta(s'_i, s_{-i}, s_0), \alpha_i) \alpha_i(s_{-i}, s_0|h).$$

²⁷For related experimental evidence distinguishing regret-sensitive choice from risk attitudes, with a focus on the role of feedback, see, e.g., Zeelenberg, Beattie, van der Pligt, & de Vries (1996), in a one-stage decision problem, and Fioretti *et al.* (2022), in dynamic individual choice.

²⁸In a setting somewhat similar to ours, Krähmer & Stone (2013) argue that anticipated regret aversion explains uncertainty aversion.

²⁹Our analysis of best replies makes it clear that many of our results and insights also apply to other solution concepts, whose use is advocated in Battigalli & Dufwenberg (2009) and Battigalli, Corrao, & Dufwenberg (2019). For the classic definition of SE in traditional games see Kreps & Wilson (1982). For extensions to psychological games see Battigalli & Dufwenberg (2009) and Battigalli *et al.* (2019). Note that the differences between the definitions of SE in the latter two articles are moot for regret games.

³⁰Recall that $\{\emptyset\} \in \tilde{\mathcal{H}}_i$, but $\{\emptyset\} \in \mathcal{H}_i$ if and only if i is active at the root. In this case, $\mathcal{H}_i^{\{\emptyset\}} = \mathcal{H}_i$.

In words, s_i is a sequential best reply to α_i if, for each nonterminal information set h of i , the h -replacement of s_i (a kind of continuation strategy) maximizes i 's expected utility conditional on h . As noted above, since the utility functions satisfy own-plan independence, they represent dynamically consistent preferences. Thus, finiteness of the game form and standard dynamic programming results imply that sequential best replies always exist.³¹

Remark 3 *For all personal players $i \in I$ and consistent belief systems α_i , $BR_i(\alpha_i) \neq \emptyset$.*

Furthermore, whenever the initial belief $\alpha_i(\cdot)$ assigns strictly positive probability to each “strategic-form information set” $S_{I_0 \setminus \{i\}}(h)$ ($h \in \mathcal{H}_i$), $BR_i(\alpha_i) = \arg \max_{s_i \in S_i} \mathbb{E}_{s_i, \alpha_i}[u_i]$ coincides with the set of strategies that maximize ex ante expected utility given the initial belief $\alpha_i(\cdot)$. In the special case of one-person games, the aforementioned condition is satisfied because $\alpha_i(\cdot) = \sigma_0(\cdot)$ is the mixed strategy of chance, which is strictly positive; see (1), (2), and (3). When we need to emphasize the dependence of best replies on the sensitivity parameter θ_i in psychological utility function (6), we write $BR_i^{\theta_i}(\alpha_i)$. If $\theta_i = 0$ then $BR_i^0(\alpha_i)$ is the set of strategies that maximize material expected payoff at each $h \in \mathcal{H}_i$:

$$BR_i^0 = \left\{ s_i \in S_i : \forall h \in \mathcal{H}_i, s_i|_h \in \arg \max_{s'_i \in S_i(h)} \mathbb{E}_{s'_i, \alpha_i}[m_i|h] \right\}.$$

The analysis of best replies is similar to that of expected utility maximization in one-person games, with Chance replaced by the set of co-players. The two key differences are first that in multi-person games belief systems are endogenous, i.e., determined by the strategic analysis of the game by means of solution concepts, and second that in games with sequential moves, surprises (initially unexpected information sets) may occur because vector α_i need not be strictly positive. Thus, $BR_i(\alpha_i)$ may be a strict subset of the set of ex ante expected utility maximizers, because sequential best replies must also maximize expected utility given revised beliefs conditional on unexpected information sets. Of course, surprises cannot occur in games with essentially simultaneous moves, to which we can apply the same logic of one-person games to analyze best replies.

Proposition 1 states conditions under which regret does not matter in one-person games. Does regret “matter” in strategic settings? We address the issue by fixing a game form and exploring whether the set $BR_i^{\theta_i}(\alpha_i)$ of sequential best replies to a consistent system of beliefs changes with θ_i . We already know that the answer can be yes, via Examples 3, 5, and 6 in which the optimal choices were, trivially, sequential best replies. We will soon see more examples, in multi-player games. However, to appreciate how those work, it is useful to first provide a backdrop of circumstances where regret does not matter.

Theorem 1 *Suppose that player $i \in I$ has own-strategy independent feedback about others. For every consistent belief system α_i , if u_i features linear regret, $BR_i^{\theta_i}(\alpha_i)$ is independent of the*

³¹Intuitively, one constructs a strategy s_i by “folding-back planning” on the tree of non-terminal information sets. Consistency of α_i implies that s_i is a sequential best reply to α_i . See, e.g., Battigalli, Catonini & Manili (2023).

sensitivity parameter θ_i and coincides with the set of strategies that maximize expected material payoff at every information set.

The result applies, e.g., to finite-action versions of the Cournot and public-good settings discussed in Section 2.1, as well as sealed-bid auctions where bids are revealed ex post. Intuitively, OSIFO implies that maximum hindsight-expected payoff is independent of the played strategy (Corollary 1). Under linear regret, such maximum payoff and expected payoff given terminal beliefs enter the utility function as additive terms multiplied respectively by $-\theta_i$ and $+\theta_i$. By the law of iterated expectations (which holds by consistency of α_i), the expectation of the second term is simply θ_i times expected monetary payoff, so expected utility is given by $(1 + \theta_i)$ times expected monetary payoff plus a term that is independent of the adopted strategy.

Theorem 1 and Remark 2 – which gives sufficient conditions for OSIFO – yield:

Proposition 2 *Consider a game form with either (1) essentially simultaneous moves and perfect feedback, or (2) asymmetric information and essentially simultaneous moves with ex post observed actions of personal players. For every player $i \in I$ and consistent belief system α_i , if u_i features linear regret, $BR_i^{\theta_i}(\alpha_i)$ is independent of the sensitivity parameter θ_i and coincides with the set of strategies/actions that maximize expected material payoff.*

Another form of θ_i -irrelevance arises in games *without Chance* under the general regret function (4) when a player is “certain” about others’ strategies whenever he moves. Formally, i ’s belief system α_i features **deterministic beliefs when active** if, for each nonterminal $h \in \mathcal{H}_i$, i ’s belief at h assigns probability 1 to a specific pure strategy profile of the other personal players:

Condition 1 (Deterministic Beliefs When Active) *Fix $i \in I$; for every $h \in \mathcal{H}_i$, there is some $s_{-i}^* \in S_{I \setminus \{i\}}(h)$ such that $\alpha_i(s_{-i}^*|h) = 1$.*

Proposition 3 *Consider a game form without chance moves. Fix any $i \in I$ and a consistent belief system α_i that satisfies deterministic beliefs when active; the set of sequential best replies $BR_i^{\theta_i}(\alpha_i)$ is independent of θ_i and coincides with the set of strategies that maximize expected material payoff at each information set $h \in \mathcal{H}_i$.*

The intuition for this result is relatively simple. Consider player i who has to move at h . By consistency of α_i , a kind of law of iterated expectations holds: if i is certain of the co-players’ strategies s_{-i}^* conditional on h , he expects to be certain of s_{-i}^* conditional on every future information set he deems reachable from h as he considers his continuation strategies, including the terminal information sets. This implies that, when certain, player i expects to feel no regret at the end of the game if and only if he plans to carry out an expected material payoff maximizing strategy.

It is worth noting that this invariance result can be extended to a regret game with Chance, where Chance takes an action only at the root ($X_0 = \{\emptyset\}$) and where player i observes Chance’s choice before moving.

4.2 Sequential equilibrium

As with traditional games, sometimes pure strategy equilibria do not exist; thus, we consider equilibria in randomized strategies. For any $h \in \mathcal{H}_i$, let $\pi_i(\cdot|h) \in \Delta(A_i(h))$ denote a probability mass function on the feasible actions of i given h . A **behavior strategy** for i is a probability vector $\pi_i = (\pi_i(\cdot|h))_{h \in \mathcal{H}_i}$. We do not assume that players actually randomize, but if $\pi = (\pi_i)_{i \in I}$ is a sequential equilibrium strategy profile then we may interpret π_i in terms of the common first-order beliefs held by i 's opponents about i .³² We say that π is **strictly randomized** if $\pi_i(a_i|h) > 0$ for all i , $h \in \mathcal{H}_i$, and $a_i \in A_i(h)$.

We consider profiles of behavior strategies and belief systems $(\pi_i, \alpha_i)_{i \in I}$ (a.k.a. “assessments”) that satisfy a strong consistency condition.³³ For any player $i \in I$, pure strategy $s_i \in S_i$, and behavior strategy π_i , let

$$\mathbb{P}_{\pi_i}(s_i) = \prod_{h \in \mathcal{H}_i} \pi_i(s_i(h)|h)$$

interpreted as the probability assigned to s_i by π_i , where $s_i(h)$ is the choice s_i makes at h (cf. Kuhn 1953). For $i = 0$, we get $P_{\pi_0}(s_0) = \sigma_0(s_0)$, where σ_0 is the (fully) mixed strategy of Chance defined in (1) of Section 2.1.

Definition 3 A profile $(\pi_i, \alpha_i)_{i \in I}$ of behavior strategies and systems of beliefs (assessment) is **fully consistent** if there is a converging sequence $(\pi_i^k)_{i \in I} \rightarrow (\pi_i)_{i \in I}$ of strictly randomized behavior strategy profiles such that, for all $i \in I$, $h \in \bar{\mathcal{H}}_i$, and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)$,

$$\alpha_i(s_{-i}, s_0|h) = \lim_{k \rightarrow \infty} \frac{\sigma_0(s_0) \prod_{j \in I \setminus \{i\}} \mathbb{P}_{\pi_j^k}(s_j)}{\sum_{(s'_{-i}, s'_0) \in S_{I_0 \setminus \{i\}}(h)} \sigma_0(s'_0) \prod_{j \in I \setminus \{i\}} \mathbb{P}_{\pi_j^k}(s'_j)}. \quad (8)$$

Full consistency requires that α_i be coherent with the behavior strategies of i 's co-players; furthermore, it strengthens consistency, requiring that beliefs satisfy a strong form of independence across co-players (cf. Battigalli 1996; Kohlberg & Reny 1997). In particular, in games with *observed deviators*, where factorization $S_{I_0 \setminus \{i\}}(h) = \times_{j \in I_0 \setminus \{i\}} S_j(h)$ holds, we have

$$\alpha_i(s_{-i}, s_0|h) = \frac{\sigma_0(s_0)}{\sum_{s'_0 \in S_0(h)} \sigma_0(s'_0)} \prod_{j \in I \setminus \{i\}} \lim_{k \rightarrow \infty} \frac{\mathbb{P}_{\pi_j^k}(s_j)}{\sum_{s'_j \in S_j(h)} \mathbb{P}_{\pi_j^k}(s'_j)} \quad (9)$$

for all $i \in I$, $h \in \bar{\mathcal{H}}_i$, and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)$. In games with *observed actions*, where $\mathcal{H}_i \cong X_i$ for each $i \in I$, full consistency is characterized by the following condition:

$$\alpha_i(s_{-i}, s_0|x) = \prod_{j \in I_0 \setminus \{i\}} \mathbb{P}_{\pi_j}(s_j|x) = \prod_{j \in I_0 \setminus \{i\}} \prod_{x' \in X_j^x} \pi_j(s_j(x')|x'). \quad (10)$$

³²An alternative interpretation is that π_i represents a (possibly) non-deterministic plan of i and there are common first-order beliefs on i 's behavior. See Battigalli *et al.* (2019).

³³Equivalent to that of Kreps & Wilson (1982), provided that one also covers terminal information sets.

for all $i \in I$, $x \in X_i$, and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(x)$.³⁴ Finally, in games with *essentially simultaneous moves*, where players get no information before moving and strategies coincide with actions, full consistency reduces to the following independence condition:

$$\alpha_i(a_{-i}, a_0) = \prod_{j \in I_0 \setminus \{i\}} \pi_j(a_j)$$

for all $i \in I$, and $(a_{-i}, a_0) \in A_{I_0 \setminus \{i\}}$.

Definition 4 An assessment $(\pi_i, \alpha_i)_{i \in I}$ is a **sequential equilibrium (SE)** if it is fully consistent and satisfies **sequential rationality**: for all $i \in I$, $h \in \mathcal{H}_i$, $s_i \in S_i$,

$$\mathbb{P}_{\pi_i}(s_i) > 0 \Rightarrow s_i \in BR_i(\alpha_i).$$

It can be checked that sequential rationality is equivalent to requiring, for each $h \in \mathcal{H}_i$,

$$\mathbb{P}_{\pi_i}(s_i^h | h) > 0 \Rightarrow s_i^h \in \arg \max_{s'_i \in S_i(h)} \mathbb{E}_{s'_i, \alpha_i}[u_i | h],$$

where

$$\mathbb{P}_{\pi_i}(s_i^h | h) = \prod_{h' \in \mathcal{H}_i^h} \pi_i(s_i(h') | h')$$

is the probability assigned by π_i to the h -replacement of s_i conditional on h . By full consistency, future behavior is expected to comply with the equilibrium (behavior) strategies independently of whether deviations have occurred, because such deviations are interpreted as mistakes in carrying out co-players' plans, and mistakes are not expected to occur in the future, just as in Selten's "trembling-hand" perfect equilibrium (cf. Kreps & Wilson 1982, Proposition 6).

Proposition 4 Any regret game G , defined as in Section 2, admits an SE.

Indeed, one can check that the "trembling-hand" argument of Battigalli & Dufwenberg (2009, Theorem 9) can be extended from continuous psychological games based on finite game forms with observed actions to all continuous psychological games based on finite game forms with perfect recall. Since G is based on a finite game form with perfect recall and psychological utilities (6) are continuous, G must have an SE.

By inspection of Definition 4, Theorem 1 implies that the equilibrium set is linear-regret-invariant when OSIFO holds, and in particular (by Proposition 2) in essentially simultaneous game forms with perfect feedback and in asymmetric-information and essentially simultaneous game forms with observed actions of personal players:

³⁴In analogy with the definition of $\mathcal{H}_j^h$, $X_j^x := \{x' \in X_j : x' \not\prec x\}$ is the set of decision nodes of j in the subtree with root x .

Corollary 2 Consider a game form Γ with either (1) essentially simultaneous moves and perfect feedback, or (2) asymmetric information and essentially simultaneous moves with ex post observed actions of personal players. Let G^θ be the associated linear regret game ($\theta = (\theta_i)_{i \in I}$). The set of SEs in G^θ is invariant w.r.t. θ . Moreover, each i behaves as if he maximized expected material payoff.

Each of conditions (1) and (2) of Corollary 2 (and of Proposition 2) cannot in general be relaxed. For condition (1), this was already shown in Examples 3 and 6, where Chance does not move simultaneously with player i . The next example expands on this point in a strategic setting without Chance, where conditions (1) and (2) coincide:




Figure 8: A parameterized three-person game form

Example 8 Consider Figure 8, an extension of Figure 2, which includes players j and k 's "Matching-Pennies" payoffs in the right sub-game and – for player i – mirrors Figure 2 if $x = 0$. Assume linear regret, and that each player has perfect information across end-nodes (i.e., perfect feedback); thus, relative to condition (1) of Corollary 2, only the requirement of essentially simultaneous moves is relaxed. Assume that $x \in (0, 1)$ and that $\theta_i \geq \theta_j = \theta_k = 0$. In any SE, j and k assign probability $\frac{1}{2}$ to each choice of the Matching-Pennies sub-game. To see what i does, note that his utility from Out equals $1 - \theta_i \cdot \frac{x}{4}$ while his utility from In equals $1 + \frac{x}{4} - \frac{1}{4} \cdot \theta_i \cdot 1$, so i prefers Out to In iff $\theta_i \geq \frac{x}{1-x}$. I.e., i 's equilibrium choice is not θ_i -invariant.

We next show that perfect feedback cannot be relaxed also in a strategic setting without Chance:




Figure 9: A two-person essentially simultaneous game form with imperfect feedback

Example 9 Consider the game form in Figure 9 and assume linear regret. If information across end-nodes were perfect then the situation would be covered by Corollary 2; we would have θ_i -invariance. This is not the case in Figure 9 as $\{(U, L), (D, L)\}$ is a doubleton information set for j . Assume that $\theta_j \geq \theta_i = 0$. It can be checked that there is a unique SE in which j assigns probability $\frac{1}{2}$ to each of his choices (indifference condition for i), while i chooses U with θ_j -dependent probability $\frac{1+2\cdot\theta_j}{2+3\cdot\theta_j}$ (indifference condition for j).

Our analysis of best replies yields another form of regret invariance of SE. Formally, a sequential equilibrium $(\pi_i, \alpha_i)_{i \in I}$ is said to be with deterministic beliefs when active if, for every player i , belief system α_i satisfies the deterministic-beliefs Condition 1. With this, by inspection of Definition 4, Proposition 3 implies the following:

Corollary 3 Consider a game form Γ without chance moves. Let G^θ be a regret game with $\theta = (\theta_i)_{i \in I}$, where no additional restrictions are imposed on $(\phi_i)_{i \in I}$. The set of SEs of G^θ with deterministic beliefs when active is invariant with respect to θ . Consequently, each player behaves as if he maximized expected material payoff.

The result also holds in games where Chance takes an action only at the root ($X_0 = \{\emptyset\}$), which is observed by all players before moving.

Condition 1 imposes a constraint on beliefs, distinct from the standard concept of “pure strategy,” which allows for off-path non-deterministic beliefs about unobserved past moves.³⁵ Formally,

³⁵Conversely, an SE with deterministic beliefs when active is not necessarily an SE in pure strategies, as there may be one player who randomizes at the root and whose initial action is observed.

a sequential equilibrium $(\pi_i, \alpha_i)_{i \in I}$ is in pure strategies if, for every personal player $i \in I$, there exists a unique strategy s_i^* that fully describes player i 's behavior at every information set:

Condition 2 (Pure Strategies) *For every $i \in I$ there is some $s_i^* \in S_i$ such that, for all $h \in \mathcal{H}_i$, $\pi_i(s_i^*(h)|h) = 1$.*

While the previous invariance result does not hold in general for the set of SEs in pure strategies (see Example 10 below), we establish that it does hold when restricting attention to games with observed actions. Given characterization (10) (full consistency in games with observed actions) we obtain the following invariance result.

Proposition 5 *Consider a game form Γ without chance moves and with observed actions. Let G^θ be a regret game with $\theta = (\theta_i)_{i \in I}$, where no additional restrictions are imposed on $(\phi_i)_{i \in I}$. The set of SEs in pure strategies in G^θ is invariant with respect to θ . Consequently, each player behaves as if he maximized expected material payoff.*

Examples 3 and 8 show, respectively, that the absence of Chance and the restriction to pure strategies are critical to Proposition 5. We now demonstrate that the statement does not hold in a game with imperfectly observed actions.




Figure 10: A game form with imperfectly observed actions

Example 10 *Consider the game form in Figure 10 with $x \in (0, 1)$ and $y > 2$. Assume linear regret and fix $\theta_i = 0$. We construct a belief system (λ, λ_c) that supports the assessment $(a, c, \lambda, \lambda_c)$ as a sequential equilibrium (SE) for sufficiently large θ_j but not for small θ_j . This demonstrates that the set of SEs in pure strategies is not invariant with respect to θ . By consistency, we have $\lambda = \lambda_c$, allowing us to express pure-strategy consistent assessments as triples (s_i, s_j, λ) . Since $y > 2$, c is conditionally dominated by a 50/50 mixture on b and d when $\theta_j = 0$. Thus, the only pure-strategy SE assessments are of the form $(\ell, b, \lambda = 1)$*

and $(r, d, \lambda = 0)$, and the path (a) is never an SE outcome—neither in pure nor mixed strategies—when θ_j is small. The parameterized utilities for j are given by:

$$\begin{aligned} u_j((\ell, b), \lambda) &= y - \theta_j \cdot 0 &= u_j((r, d), \lambda), \\ u_j((\ell, c), \lambda) &= 1 - \theta_j \cdot [\max\{\lambda y, (1 - \lambda)y\} - 1] &= u_j((r, c), \lambda), \\ u_j((\ell, d), \lambda) &= 0 - \theta_j \cdot y &= u_j((r, b), \lambda). \end{aligned}$$

Now, consider $\lambda = \frac{1}{2}$. The expected utilities of player j 's strategies, conditional on h (i.e., conditional on observing that player i did not choose a), are:

$$\begin{aligned} \bar{u}_{j,h}\left(b, \lambda = \frac{1}{2}\right) &= \frac{1}{2}y - \frac{1}{2}\theta_j y = \bar{u}_{j,h}\left(d, \lambda = \frac{1}{2}\right), \\ \bar{u}_{j,h}\left(c, \lambda = \frac{1}{2}\right) &= 1 - \theta_j \left(\frac{1}{2}y - 1\right). \end{aligned}$$

The pair of parameters (θ_j, y) makes $(a, c, \lambda = \frac{1}{2})$ an SE if and only if:

$$\begin{aligned} 1 - \theta_j \left(\frac{1}{2}y - 1\right) &\geq \frac{1}{2}y - \frac{1}{2}\theta_j y \\ \theta_j &\geq \frac{1}{2}y - 1 > 0. \end{aligned}$$

This game form features observed deviators; hence this property is insufficient for the invariance result of Corollary 3; the stronger observed-actions property is tight.

Outside the limited class of cases where regret does not matter, comparative statics on $\theta = (\theta_i)_{i \in I}$ may have dramatic effects. In Section 6, we exhibit this for a variety of economic settings. It is natural to wonder whether, perhaps, if all θ_i 's move in a common direction, the set of SEs would change monotonically (with respect to inclusion). It only takes a brief reflection to realize that the equilibrium set does not always change monotonically. Example 3 provides a counterexample. The shift, on increasing θ_i , from SE risky to SE safe implies that the solution set neither expands nor contracts. For small values of $\theta = (\theta_i)_{i \in I}$, however, we have the following limit result where slightly higher sensitivity to regret leads to (weakly) sharper predictions:

Proposition 6 *Fix a game form. Let G^θ be the associated regret game $(\theta = (\theta_i)_{i \in I})$. Let $SE(G^\theta)$ be the associated set of SEs. Let G^0 denote the game where $\theta_i = 0$ for all $i \in I$. With this, $\lim_{n \rightarrow \infty} SE(G^{\theta^n}) \subseteq SE(G^0)$ for every sequence $(\theta^n)_{n \in \mathbb{N}}$ such that $\lim_{n \rightarrow \infty} \max_{i \in I} \theta_i^n = 0$. Moreover, the inclusion may be strict.*

Intuitively, the result follows from the upper-hemicontinuity of the equilibrium correspondence $\theta \mapsto SE(G^\theta)$, see Lemma 3 in the Appendix.³⁶ Sharper predictions for small regret sensitivity

³⁶Given that the SE correspondence is compact-valued, we consider the Hausdorff metric topology over compact subsets of $\times_{i \in I, h \in \mathcal{H}_i} \Delta(A_i(h))$.

come from the lack of lower-hemicontinuity at 0. To see that the inclusion may be strict, consider Example 8 and assume $x = 0$. Then, the path in which player i chooses In is an SE outcome of G^0 . (Indeed, every behavior strategy profile in which players j and k randomize uniformly is part of an SE assessment of G^0 , since uniform randomization by j and k renders i indifferent.) However, for any sequence θ^n with $\theta_i^n > 0$, no path that includes In can be an SE outcome of G^{θ^n} .

5 Information avoidance

We now analyze how anticipated regret affects different forms of information acquisition, pushing in the direction of information avoidance. First we show that strategies that yield more ex post information have higher hindsight-expected payoffs. Other things being equal, this lowers expected utility by increasing anticipated regret. We then show by example that informative signals that are materially valuable may be avoided because their acquisition would entail higher anticipated regret.

Recall from Section 2.1 that $\mathcal{F}_i(s_i)$ is the ex post information partition of $S_{-i} \times S_0$ representing the feedback that player i can obtain about co-players' behavior if he plays strategy s_i . We compare the informativeness of strategies by comparing the corresponding partitions. By Remark 1, the same informativeness partial order holds for the corresponding reduced strategies.

Definition 5 Fix $s'_i, s''_i \in S_i$; s'_i is *at least as ex post informative as* s''_i if $\mathcal{F}_i(s'_i)$ is finer than $\mathcal{F}_i(s''_i)$, i.e., every cell in $\mathcal{F}_i(s''_i)$ is a union of cells in $\mathcal{F}_i(s'_i)$.

For example, suppose player i is active at the root and can terminate the game immediately by choosing $\text{stop} \in A_i(\emptyset)$. Every strategy s'_i with $s'_i(\emptyset) \neq \text{stop}$ is at least as ex post informative as any s''_i with $s''_i(\emptyset) = \text{stop}$, because the latter induces the trivial partition $\mathcal{F}_i(s''_i) = \{S_{-i} \times S_0\}$, which is necessarily (weakly) coarser than $\mathcal{F}_i(s'_i)$. For any *terminal* information set $h \in \mathcal{Z}_i$ and belief system α_i , let $M_i(h, \alpha_i)$ denote the hindsight-maximum expected payoff conditional on h given α_i featured in the psychological utility formula (6).

Lemma 2 Fix a consistent belief system α_i and two strategies $s'_i, s''_i \in S_i$. If s'_i is at least as ex post informative as s''_i , then for every nonterminal information set $h \in \mathcal{H}_i$ such that $s'_i, s''_i \in S_i(h)$,

$$\begin{aligned} & \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 \mid h) M_i(\bar{H}_i(\zeta(s'_i, s_{-i}, s_0)), \alpha_i) \\ & \geq \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 \mid h) M_i(\bar{H}_i(\zeta(s''_i, s_{-i}, s_0)), \alpha_i). \end{aligned}$$

Together, Lemmata 1 and 2 highlight how ex post information enters regret: terminal information determines which co-players' strategy profiles remain possible when foregone alternatives are evaluated with hindsight. When two strategies induce the same ex post information partition

of $S_{-i} \times S_0$, Lemma 1 establishes pointwise equivalence of the corresponding hindsight-expected payoffs and, consequently, equality of the expected maximum hindsight payoffs conditional on h . Lemma 2 instead compares these expected maximum hindsight payoffs when one induced partition is weakly coarser than the other.

Proposition 7 *Fix a nonterminal information set $h \in \mathcal{H}_i$, consider strategies $s'_i, s''_i \in S_i(h)$ and a consistent belief system α_i such that $\mathbb{E}_{\alpha_i, s'_i}[m_i|h] \leq \mathbb{E}_{\alpha_i, s''_i}[m_i|h]$ (resp. $\mathbb{E}_{\alpha_i, s'_i}[m_i|h] < \mathbb{E}_{\alpha_i, s''_i}[m_i|h]$). If s'_i is at least as ex post informative as s''_i , then, under linear regret, $\mathbb{E}_{\alpha_i, s'_i}[u_i|h] \leq \mathbb{E}_{\alpha_i, s''_i}[u_i|h]$ (resp. $\mathbb{E}_{\alpha_i, s'_i}[u_i|h] < \mathbb{E}_{\alpha_i, s''_i}[u_i|h]$).*

Intuitively, this follows from the fact that finer ex post information allows a more fine-tuned evaluation of hindsight-maximized expected material payoffs, increasing anticipated regret. We next illustrate this for the interesting case in which s'_i and s''_i yield the same expected payoff, but the more ex post informative strategy has strictly lower expected utility.³⁷




Figure 11: A one-player game form with Chance

Example 11 *In Figure 11, player i chooses among saving money, trading a stock, and placing a bet in a casino. Chance chooses a state $\omega = (\omega^s, \omega^b)$ in $\{w, l\} \times \{w, l\}$, where the first (second) coordinate represents whether the stock (the bet) yields a gain. Both activities lose with probability $\lambda \in (0, 1)$. Player i always observes the stock outcome regardless of his choice, but he observes the bet outcome only if he bets. Trading a stock and placing a bet yield the same expected material payoff. However, stock-trading is less ex post informative than betting. Under linear regret, the anticipated regret from trading is $\lambda(\max\{2(1 - \lambda), 1\} - 0) = \max\{2(1 - \lambda)\lambda, \lambda\}$ and the anticipated regret from betting is $2(1 - \lambda)\lambda + \lambda^2$. For any $\lambda \in (0, 1)$, trading a stock yields strictly lower anticipated regret than placing a bet and therefore strictly higher expected utility.*

The information-avoidance result extends beyond linear regret under stronger assumptions about the material-payoff implications of the strategies being compared.

³⁷Section 6.1 applies Proposition 7 to auctions.

The next example shows that informative and materially valuable signals may be avoided, implying that ex post less informative strategies may be strictly preferred even if they yield strictly lower expected material payoffs.




Figure 12: Information acquisition game form

Example 12 In the game form of Figure 12, player i can choose to bet or not. Betting yields a payoff of 1 if successful and -1 if unsuccessful, while not betting yields 0. The signal — H or L — indicates whether the probability of success (w) is high (0.6) or low (0.3). Before deciding whether to bet, the player can acquire (a) information about ω costlessly. Assume linear regret. Suppose $\theta_i \in (0.5, 1)$, a range of interest because θ_i is high enough to affect information acquisition, but not so high as to overturn the choice the player would make once informed. In the associated regret game, without information the player will not bet: the ex ante winning probability is $(0.6 + 0.3)/2 = 0.45 < 0.5$, yielding an expected utility of

$$0.45 + 0.55 \cdot (-1 - \theta_i \cdot 1) < 0.$$

With information, his second choice depends on the realized signal. If it is high, $\omega = H$, he bets, since

$$0.6 + 0.4 \cdot (-1 - \theta_i \cdot 1) > 0 - \theta_i \cdot (0.6 - 0.4),$$

provided that $\theta_i < 1$. If it is low, $\omega = L$, he will not bet, since

$$0.3 + 0.7 \cdot (-1 - \theta_i \cdot 1) < 0.$$

Thus, acquiring information improves the decision with a high signal realization. Indeed, acquiring the information also leads to a net gain in the expected material payoffs. Nevertheless, player i will not acquire information. Once informed, he bets if the signal is high, which raises the level of anticipated regret. Since

$\theta_i > 0.5$ implies

$$0 > \frac{1}{2} \cdot 0.6 + 0.4 \cdot (-1 - \theta_i \cdot 1) + \frac{1}{2} \cdot 0,$$

the player strictly prefers not acquiring information and not betting (reduced strategy $r.n$) to acquiring information and betting iff the signal is high (reduced strategy $a.b^H.n^L$). The latter (reduced) strategy yields strictly higher expected material payoff while the former is strictly less ex post informative ($\mathcal{F}_i(r.n) = \{S_0\}$, $\mathcal{F}_i(a.b^H.n^L) = \{\{H.w\}, \{H.l\}, \{L.w, L.l\}\}$) implying strictly lower anticipated regret.

6 Applications

In an effort to inspire future applied theoretical and empirical research, we now explore a few illustrative examples. The settings are simple, yet rich enough to have meaningful economic interpretations. We highlight various intriguing effects that arise if the players are concerned with regret, which may also be indicative of the relevance of regret in richer related settings.

6.1 Auctions

A large literature (e.g., Engelbrecht-Wiggans 1989, Engelbrecht-Wiggans & Katok 2008, Filiz-Ozbay & Ozbay 2007, and Bergemann, Breuer, Cramton, Hirsch, Ndiaye & Ockenfels 2025) has argued that anticipated regret helps explain bidding in auctions. A key idea/finding is that if information about losing bids is not disclosed, then bidders compete harder because they worry less about the “winner regret” that might otherwise accompany the ex post insight that they might have won at lower cost. Hiding losing bids means adjusting a game form’s terminal information structure. Under traditional game-theoretic assumptions, doing so does not impact players’ incentives. From this point of view, the idea/finding is puzzling. However, scholars have appealed to intuition when winner regret was simply assumed away in game forms with undisclosed losing bids. By contrast, as we have shown in many examples, in our theory terminal information structure impacts players’ incentives without any ad hoc switch of utility assumption. Here we sketch an explanation, based on our previous findings, of why anticipated regret avoidance makes bidding more aggressive when losing bids are hidden.³⁸

To fit our finite game framework, we consider sufficiently fine grids of valuations and bids. Furthermore, we assume *linear regret*. Finally, we also assume, for the sake of simplicity, that there are two bidders—i.e., $I = \{1, 2\}$ —where bidder 1 wins the object when they tie, and that regret sensitivities are commonly known.³⁹ Using our previous analysis, we derive a result about bidder 1’s best replies to *given beliefs* $\hat{\alpha}_1$ under different information feedback regimes. A salient example of the beliefs $\hat{\alpha}_1$ is what is implied by the equilibrium bidding functions under expected payoff maximization, i.e., with $\theta_1 = \theta_2 = 0$.

³⁸A similar argument shows that hiding winning bids makes bidding less aggressive.

³⁹We can show that our argument applies more generally.

We compare two information-feedback regimes: **full revelation**, in which bids are perfectly observed ex post, and **hiding losing bids**, in which the winner does not observe the losing bid, while the loser observes the winning bid. We consider bidder 1 with valuation v_1 and regret sensitivity θ_1 under each information-feedback regime. A similar argument applies to bidder 2. Since there are *private values*, only the competitor's distribution of bids matters for expected utility maximization. Fix a distribution of the competitor's bids $\hat{\alpha}_{1,2} \in \Delta(A_2)$. To streamline the analysis, we assume that bidder 1 has a unique best reply to $\hat{\alpha}_{1,2}$ under each information-feedback regime.⁴⁰ In particular, $\hat{\alpha}_{1,2}$ may be the standard equilibrium distribution.⁴¹

Proposition 8 *For each valuation v_1 , the best reply to $\hat{\alpha}_{1,2}$ under hiding losing bids is (weakly) higher than under full revelation.*

To see this, note that a first-price auction corresponds to a game form with asymmetric information and simultaneous moves. Under full revelation, the actions of personal players are observed ex post. Therefore, the assumption of own-strategy independence of feedback about others holds (Remark 2) and the best reply for each valuation is the bid that maximizes expected monetary payoff independently of the regret sensitivity parameter (Theorem 1). Keeping beliefs fixed, let us compare with the hiding losing bids regime. The key consideration is that, *if losing bids are hidden, lower bids are ex post more informative than higher bids*. Intuitively, since losers observe the winning bid while winners only learn that the competitor's bid was lower than theirs, a lower bid yields more losing cases, in which the competitor's bid is observed, and reveals more than a higher bid when both bids win.⁴² Thus, by Proposition 7, bids strictly below the expected payoff-maximizing one are strictly less attractive because they are more ex post informative; with this, the best reply must be (weakly) higher than under the full revelation.⁴³

Building on this result, we can show that *if* best replies to the competitor's distribution of bids are increasing with respect to first-order stochastic dominance,⁴⁴ then *anticipated-regret avoidance makes equilibrium bidding more aggressive when losing bids are hidden than under full revelation*.

This auction example also suggests a broader implication. Auctions are special cases of mechanism design problems. If our theory can help explain auctions, it may shed light on mechanism design issues more generally. In particular, we have in mind aspects of information design. By judicious adjustment of regret-prone players' ex post information, a planner's ability to implement may be enhanced.

⁴⁰It is common to assume uniqueness of the best reply to the standard equilibrium distribution.

⁴¹That is, the pushforward of the distribution of valuations under the competitor's equilibrium bidding function.

⁴²See Lemma 4 in the Appendix, which formally establishes that, under the hiding-losing-bids regime, lower bids yield a finer ex post information partition.

⁴³In the uniform-valuation model, we can show under reasonable assumptions that, for every $\theta_1 > 0$, if the grid of bids is sufficiently fine the inequality is strict.

⁴⁴It is sufficient that this monotonicity condition hold for best replies with $\theta_1 = 0$.

6.2 Climate action and payoff observability

Climate scientists argue that global warming is anthropogenic in nature and that carbon taxation could solve the problem. Skeptics dispute that account. Many folks are in between, unsure about the physics or policy impact. So, who will support climate action? Environmental economists van der Ploeg & Rezai (2019) (vdP&R) tackle that question from a variety of decision-theoretic angles. Applying our linear-regret model to their game (Figure 13(i)), some striking conclusions emerge:




Figure 13: Three climate-action game forms

Player i , who is unsure whether the scientists or the skeptics are right, must “price carbon” (PC) or “do nothing” (DN). The chance move in Figure 13(i) reflects i ’s state of mind; he assigns probability λ to the scientists being right.⁴⁵ VdP&R assume that $0 < c < a < b$: If Player i chooses DN and the skeptics are right, then i ’s (normalized) payoff is 0 ; c , the “costs of unnecessarily distorting energy decisions” (vdP&R, p. 73), is small relative to b (which reflects a catastrophe); finally, $a < b$ as PC reflects the best policy in a warming world while $a > c$ “since pricing carbon would not mitigate all emissions from now onward and human emissions do not damage welfare if climate deniers are right” (ibid).⁴⁶

⁴⁵In other words, the initial chance move should be interpreted as an unknown state of nature, and λ is a subjective probability.

⁴⁶VdP&R make additional assumptions that render PC a fairly obvious choice, notably that λ is “large” and that $b - a > c$. We entertain more possibilities on the grounds that: (i) If skeptics (with $\lambda = 0$) exist then presumably also quasi-skeptics with λ “small or medium” exist. (ii) The game form in Figure 13(i) abstracts from how climate action is in reality a multi-player game played across the world. An individual may then perceive significant opportunity costs of choice PC, if he believes that countries other than his will take enough climate action to thwart the crisis regardless of his choice. This would induce an incentive to “free-ride:” as a shorthand for such reasoning, a might be higher than vdp&r suggest so that, possibly, $b - a < c$. All of this noted, we add that these matters are not very important to our analysis because (as we explain below) we focus on a form of population-based marginal impact

If Player i maximizes expected payoffs (or expected utilities, if payoff equals utility) then he prefers PC to DN if $\lambda > c/(b - a + c)$, and vice versa. If others are like him, except that their (homegrown, choiceless; cf. text after (6)) parameter values (λ, a, b, c) differ, we can divide the population into those who choose (and presumably vote for) PC or those who favor DN. We now ask how the size of these groups changes if voters are affected by (linear) regret, as we model it.

The game form in Figure 13(i) is covered by Propositions 1 & 2, so incorporating regret changes no conclusion for that game form.

However, arguably Figure 13(i) is not the “right” game form in which to conduct the analysis. It features observable payoffs. While that is a natural assumption in many economic contexts, this need not be the case here, for two distinct reasons. First, Player i may not be able to distinguish the terminal nodes where he gets $-c$ and $-a$; if so, this moves us from Figure 13(i) to 13(ii). The justification is that these game forms are both highly simplified models. In the real world, even if the argument that vDP&R offer for why $c \neq a$ is valid, lots of realistic background noise is not incorporated into Figs 13(i)&(ii). For example, the temperature on earth is probably subject to random shocks that shroud the clarity with which it can be determined which of the two worldviews is more accurate. Translated to Figure 13(ii), the outcomes where Player i gets $-a$ or $-c$ may appear indistinguishable as they happen, in the sense that they would not reveal Chance’s initial choice. At each node, global warming was avoided; the reasons for this (i.e., the choices of Chance and Player i) differ, but because of unmodeled background noise, i cannot tell.

If the game form in Figure 13(ii) is relevant, then the impact of regret is felt. Namely, the range of parameters (λ, a, b, c) for which Player i will choose PC expands (with θ_i). To verify this, note that the logic of the solution of Player i ’s decision problem in Figure 13(ii) resembles that in Example 5, and the conclusions are analogous. All in all, the more regret-prone that decision makers are, the higher the number of them who will choose PC.

Second, and perhaps more dramatically, Player i may not be able to distinguish the terminal nodes where he gets $-b$ and 0. At first glance this may appear implausible, since b is presumably much larger than 0. However, the catastrophic effects of climate change may take time to realize. While Player i may (deeply) care about the difference between $-b$ and 0, he may well be *dead* when Chance’s choice is revealed! If so, and if the logic for why he also cannot distinguish the outcomes where he gets $-a$ and $-c$ holds, then we move from Figure 13(ii) to Figure 13(iii). It is straightforward to now verify that, again, regret will not influence the analysis. That is, even if Player i cares about the welfare of his offspring just as much as his own, he cannot experience regret after he dies, and thus regret has no effect. His offspring also cannot experience regret, since they did not make the choice.

Let us qualify that further: Is it really likely that player i will be dead once the outcome associated with $-b$ is revealed? This may depend on how much longer i has to live when DN is chosen, and his related beliefs may affect his choice. It may be that for (relatively more) young of changing θ_i rather than the frequency of Players i who might choose PC.

people the relevant game form is Figure 13(ii), while for (relatively more) older people it is Figure 13(iii). We would then expect regret to influence decisions differently for different generations.

6.3 Coordination through feedback

Cerrone, Feri & Neary (2025) propose and experimentally test a novel game form in which two players' choices are connected through feedback. In the game form, shown in Figure 14, each player chooses between s , which yields a fixed payoff of i , and ℓ , which yields either $R > 1$ with probability p or 0. If at least one player chooses ℓ , the outcome of that option is revealed ex post to both players.




Figure 14: Two-player gamble game form

This innovative setting creates an environment where one player's choice affects the other through feedback, but only when players care about feedback, for example, when players are regret averse. For such players, the likelihood of experiencing regret from a foregone ℓ increases when the co-player also chooses ℓ , since the chance to learn the co-player's strategy is higher. Thus, players prefer to make the same choices: both- s or both- ℓ . Cerrone *et al.* analyze this game form by assuming that a player who receives no feedback about a foregone choice does not experience regret, and predict two types of equilibria, one in which both choose ℓ , and another in which both choose s . Thus, two regret-averse players may coordinate on either action.

Assuming *linear regret*, our model predicts a similar outcome as theirs. First note that, when $pR = 1$, the expected payoffs from choosing ℓ and s are the same. Moreover, for each player, ℓ is a strategy that is at least as ex post informative as s in the sense of Definition 5. Therefore, Proposition 7 applies, implying that each player weakly prefers s to ℓ . Intuitively, whenever this preference is strict, R must exceed $\frac{1}{p}$ to make the two strategies indifferent. The following

calculation confirms this intuition.

In general, consider the associated regret game with $\theta_i = \theta_j = \theta > 0$. Let $q_i = \alpha_i(\ell)$ denote the probability that player i assigns to j choosing ℓ . Player i prefers ℓ if

$$pR - (1-p)\theta \geq 1 - pq_i\theta(R-1) - (1-q_i) \cdot \theta \max\{pR-1, 0\}$$

$$\Leftrightarrow R \geq \frac{1}{p} \left[\frac{1 + \theta + \theta(1-p)(1-q_i)}{1 + \theta} \right].$$

As q_i increases, the anticipated regret from choosing s rises, making s less attractive. Accordingly, the threshold value of R required for choosing ℓ decreases in q_i . If $R \in \left[\frac{1}{p}, \frac{1}{p} \frac{1+\theta+\theta(1-p)}{1+\theta} \right]$, then there are two pure (sequential) equilibrium strategy profiles: (s, s) and (ℓ, ℓ) .

6.4 Market entry

Compare the game forms in Figures 15(i) & (ii). In each case, two firms, A and B, decide whether or not to enter an emerging market facing “entry cost” $c \in (0.5, 1)$. The revenue of a firm that stays out equals 0, and if one or both enter(s) the market then the entrant(s) get a revenue of 1. The game forms differ only as regards what happens if both A and B choose *In*. In (i), the firms split the market, getting revenue $\frac{1}{2}$ each; this case resembles what is seen, e.g., in the wireless carrier industry. In (ii), a single winner is instead selected at random (by a 50/50 chance move): this resembles what happens in, e.g., the medical industry, when A and B engage in a patent race.




Figure 15: Two market-entry game forms

Suppose the firms’ CEOs are regret-neutral: $\theta_A = \theta_B = 0$. Each of the associated games has a unique completely mixed SE such that each firm will enter with a probability of $2(1 - c)$.

If instead the CEOs are regret-averse, i.e., $\theta_A, \theta_B > 0$, and both u_A and u_B feature linear regret, then the regret game corresponding to each game form still has a unique completely mixed SE. However, the strategies involved are different. In (i), which involves essentially simultaneous moves

as well as perfect feedback, we know that regret invariance must hold (by Corollary 2): each firm will enter with probability $2(1 - c)$. By contrast, in (ii), for each firm $i \in \{A, B\}$, let $j \neq i$ denote the other firm. Firm i will enter with probability

$$q_i := \frac{(1 - c)(1 + \theta_j)}{\theta_j + (1 - c \cdot \theta_j)/2} < 2(1 - c).$$

Note that q_i is decreasing in θ_j .

Scholars in IO have taken an interest in how (behavioral) CEOs' characteristics may influence market entry; see, e.g., Goldfarb & Xiao's (2011) "cognitive hierarchy" analysis that permits heterogeneity as regards "strategic sophistication." Our example suggests that regret *may* provide another route, and our preceding results provide guidance as to why this is the case. Namely, by Corollary 2, regret cannot matter in the regret games associated with the game form in (i).

6.5 Delegation

Our final topic concerns strategic delegation, and our analysis suggests that regret can influence outcomes in this arena. Player A, now the captain of a soccer team, is fouled and awarded a penalty kick. A can take the kick himself, or delegate it to his teammate B. The interaction also involves the opposing side's goalkeeper (G). See Figure 16. First, A decides whether to DIY or to delegate the kick to B. Then, whoever is the kicker, and G, make their choices (kick & dive) – left or right – simultaneously, with the (material) payoffs as indicated in the figure.




Figure 16: Delegation vs. DIY game form

Assume linear regret. Solving for SE, it is clear that each player will kick/dive left or right with equal probability, independently of any regret sensitivities. Now consider A's initial choice, whether or not to delegate. If $\theta_A = 0$, then A is clearly indifferent. If $\theta_A > 0$, then A will strictly

prefer to delegate as this minimizes his expected regret. This is consistent with Proposition 6: as regret sensitivities go to zero, the regret-sensitive SE set need not converge to the full set of no-regret SEs. Note that A can end up regretting his choice in two ways: He may delegate, see B fail to score, and then experience regret in the amount of $\frac{1}{2}$, which is the would-have-been expected score from choosing to DIY. Alternatively, A may take the kick himself, miss, and then experience regret in the amount of 1, as he would have been able to score had he attempted the kick in the opposite direction. While A cannot avoid the pangs of regret that may come with his delegation decision, he still prefers to delegate because this choice shuts down the regret associated with his own kicking.

The example implicitly presumes that A and B are equally proficient at penalty shooting. The fact that, in SE, A strictly prefers to delegate makes it clear that one could develop an extension where A is more skilled than B, yet still delegates the task to B.

7 Comparison with Decision Theory

In this section we compare our game-theoretic analysis of regret-averse players with the decision theoretic analysis of regret. The two approaches have different goals and features, but they overlap in one-person games, which makes an explicit comparison useful. At a high level, we take a functional-form representation as given and explore its implications in game forms with *general* information structures, while decision theory derives representations from axioms on preferences over spaces of conceivable alternatives typically assuming *special* information structures. The axiomatic approach yields transparent, modular representations of assumptions about behavior and, in principle, allows the identification of an agent's regret trait from the observation of choices in controlled conditions. To the best of our knowledge, the extant decision theory concerning regret assumes that all the uncertainty is resolved ex post. We argued that this is too restrictive, and even impossible in genuinely sequential games, where counterfactual choices cannot be observed ex post.

In Section 7.1 we consider early decision regret theory (DRT) and show how we adapt one of its functional form representations to allow for imperfect ex post resolution of uncertainty. We consider the menu-choice approach to regret in Section 7.2, and minimax regret in Section 7.3.

7.1 Early decision regret theory

Since our approach subsumes games with a single personal player plus Chance as a special case, it is natural to wonder about its connection with the classic DRT. In this section we explore this issue. First, recall what the DRT pioneers did: Bell (1982) and Loomes & Sugden (1982) focused

on pairwise choice,⁴⁷ while extensions to larger choice sets were proposed subsequently (see also Sugden 1993). We relate to the following version, proposed by Quiggin (1994):⁴⁸

Let Ω , $C \subseteq \mathbb{R}$, and $A \subseteq C^\Omega$ be finite, non-empty sets of states, consequences (in the form of monetary payments), and Savage acts that a decision maker (DM) can choose. Let $\pi(\omega) > 0$ be the exogenously given probability of $\omega \in \Omega$. Let $V : C \rightarrow \mathbb{R}$ describe DM's "choiceless utility" (Loomes & Sugden's terminology) of consequences. After DM chooses $a \in A$, state $\omega \in \Omega$ is revealed and DM (being regretful) ruminates on what-could-have-been. His regret-adjusted utility $U : \Omega \times A \rightarrow \mathbb{R}$, of which he maximizes the expectation, is

$$U(a, \omega) = V(a(\omega)) - f \left(\max_{a' \in A} V(a'(\omega)) - V(a(\omega)) \right), \quad (11)$$

where function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, strictly increasing and satisfies $f(0) = 0$.

To account for non-linearities in choiceless utility, we extend our original model to encompass the DRT framework. Specifically, we replace $m_i(z)$ and $\bar{m}_i(h, \alpha_i) = \mathbb{E}_{\alpha_i} [m_i|h]$ in (6) and (4) with $v_i(m_i(z))$ and $\mathbb{E}_{\alpha_i} [v_i \circ m_i|h]$, where $v_i : [\min_{z \in Z} m_i(z), \max_{z \in Z} m_i(z)] \rightarrow \mathbb{R}$ is a strictly increasing function that may be non-linear.⁴⁹ Consequently, i 's extended utility function becomes:

$$\tilde{u}_i(z, \alpha_i) = (v_i \circ m_i)(z) - \theta_i \tilde{r}_i(\bar{H}_i(z), \alpha_i); \quad (12)$$

where the extended regret function $\tilde{r}_i$ is:

$$\tilde{r}_i(h, \alpha_i) := \phi_i \left(\max_{s_i \in S_i} \left\{ \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0|h) (v_i \circ m_i)(\zeta(s_i, s_{-i}, s_0)) \right\} - \mathbb{E}_{\alpha_i} [v_i \circ m_i|h] \right), \quad (13)$$

where function $\phi_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, strictly increasing and satisfies $\phi_i(0) = 0$.⁵⁰

Given this extended model, for an arbitrary decision model in the DRT framework (Ω, C, π, A, U) , there always exists an extended regret game $(\Gamma, \tilde{u}_i)$ with a personal player i and Chance that can recast this decision problem.

To construct such a regret game, consider a game form Γ in which Chance first selects a state $\omega \in S_0 := \Omega$ according to the probability distribution $\pi(\cdot) \in \Delta(\Omega)$, i.e., $\pi_0(\omega|\emptyset) = \pi(\omega)$. Player i (DM) then chooses $s_i \in S_i$ without observing ω , the game ends afterwards. Furthermore, let i be able to observe ω at the end of the game; the game form features perfect feedback. For simplicity,

⁴⁷Related axiomatic work by Bikhchandani & Segal (2011, 2014) studies when regret-based preferences are transitive in the correlated-lottery and independent-lottery cases, respectively.

⁴⁸Quiggin shows that his approach is implied by "irrelevance of statewise dominated alternatives," a quite intuitive axiom. See also Loomes & Sugden (1987).

⁴⁹Perfect recall implies that $\mathbb{E}_{\alpha_i} [v_i \circ m_i|h]$ is independent of the strategy played by i . Hence, it is determined by $\alpha_i(\cdot|h) \in \Delta(S_{I_0 \setminus \{i\}}(h))$.

⁵⁰An alternative approach, adopted in some papers (e.g., Bikhchandani & Segal, 2014), assumes that regret arises from comparing the realized choiceless utility with the certainty equivalent of the foregone lottery.

define the path function $\zeta : S_i \times S_0 \rightarrow Z$ as the identity on $Z = S_i \times S_0$.

To establish the relationship between Savage acts and strategies, consider any act $a \in A$ in the decision problem. Construct a strategy s_i and monetary payoff function $m_i(\cdot)$ such that, for all $\omega \in \Omega$, $m_i(s_i, \omega) = a(\omega)$. Since i chooses s_i without observing ω , each strategy s_i defines a section of the monetary payoff function, $m_i(\cdot, s_i) : \Omega \rightarrow C$, which corresponds to act a . According to this construction, S_i is isomorphic to A , and the set of end nodes can be written as $Z = A \times \Omega$.

Given Γ , we identify a specific $\tilde{u}_i$ that takes the same value as U at each end node (ω, a) :

Remark 4 *Given the DRT utility function $U(a, \omega) = V(a(\omega)) - f\left(\max_{a' \in A} V(a'(\omega)) - V(a(\omega))\right)$, if the extended utility function $\tilde{u}_i$ satisfies:*

(i) *at each $(\omega, a) \in Z$, $(v_i \circ m_i)(\omega, a) = V(a(\omega))$, and*

(ii) *for all $x \in \mathbb{R}_+$, $\theta_i \cdot \phi_i(x) = f(x)$,*

then for each $(\omega, a) \in Z$ and each belief α_i ,

$$\tilde{u}_i((\omega, a), \alpha_i) = U(a, \omega).$$

A basic DRT-insight is that if f is linear (e.g., $f(x) = k \cdot x$ with $k > 0$), then regret has no impact on choice: an act $a \in A$ maximizes the expectation of U iff it maximizes the expectation of V .⁵¹

Building on our earlier alignment result, we now restate this property within our extended model. We continue to say that regret is *linear* if ϕ_i is the identity function on $\mathbb{R}_+$:

Corollary 4 *Suppose that $(\Gamma, \tilde{u}_i)$ features essentially simultaneous moves and perfect feedback. If regret is linear, then the strategies (actions) of player i that maximize the expectation of $\tilde{u}_i$ are the same as those that maximize the expectation of $v_i \circ m_i$.*

In $(\Gamma, \tilde{u}_i)$, a strategy s_i maximizes the expectation of $\tilde{u}_i(\cdot)$ iff it maximizes the expectation of $(v_i \circ m_i)(\cdot)$. Although the linear irrelevance result provides valuable insights, its application is limited to settings with simultaneous chance moves and perfect feedback. As demonstrated in Section 3, relaxing these conditions reveals that regret matters even when it is linear.

7.2 Regret with menu choice

Building on the work of Dekel, Lipman, & Rustichini (2001) on preference for flexibility, Sarver (2008) employs an axiomatic approach to model preferences over menus. He characterizes the decision maker's preference over menus with a utility representation as if the decision maker chooses

⁵¹See, e.g., Loomes & Sugden (1982, p. 810). For this reason, DRT-scholars have focused on the implications of f being concave or convex.

a lottery from a menu and subsequently experiences regret when the chosen lottery turns out to be ex post suboptimal. Let $p, q \in \Delta$ denote lotteries in a simplex, and let $A \in \mathcal{A} \subseteq 2^\Delta$ denote a conceivable menu. Following Dekel *et al.* (2001), the state space describing uncertainty is represented as a set $\mathcal{U}$ of affine utility functions defined on the convex set Δ of lotteries. The consideration of lotteries is essential to derive a linear structure on the collection $\mathcal{A}$ of conceivable menus (closed subsets of Δ) and obtain an essentially unique representation from axioms about *menu preferences*, but it is not crucial for the comparison with our model. Sarver’s utility representation, which characterizes the decision maker’s preference over menus, is given by

$$U(A) = \max_{p \in A} \int_{\mathcal{U}} \left[u(p) - K \cdot \left[\max_{q \in A} u(q) - u(p) \right] \right] \mu(du), \quad (14)$$

where μ is a subjective probability measure on the state space $\mathcal{U}$, $u : \Delta \rightarrow \mathbb{R}$ is an (affine) utility of lotteries, and $K \geq 0$. The representation yields a standard subjective expected utility formula for the *commitment preference*, i.e., the preference over the subcollection of singleton menus: $\{p\} \succeq \{q\}$ if and only if $\int_{\mathcal{U}} u(p) \mu(du) \geq \int_{\mathcal{U}} u(q) \mu(du)$. Indeed, regret arises not from the choice among menus, but from the choice of lotteries within a given menu A . Furthermore, regret is experienced when uncertainty is resolved (i.e., the true u is observed) and before the realization of the risky lottery.

Since regret enters linearly in U and utility-state u is observed ex post, the chosen lottery p within a menu A must maximize expected utility $\int_{\mathcal{U}} u(p) \mu(du)$ (cf. Proposition 1); thus, regret only affects the value of menus. Sarver’s central axiom is that a decision maker weakly prefers a smaller menu to an extended menu that includes an additional lottery that would not be chosen under the commitment preference: if $p \in A$ and $\{p\} \succeq \{q\}$, then $A \succeq A \cup \{q\}$. Note that, without regret ($K = 0$), the antecedent of the foregoing implication yields an equivalence: $A \sim A \cup \{q\}$. The axiom instead allows for strict preference $A \succ A \cup \{q\}$, reflecting that even an unchosen lottery can generate greater ex post regret in some utility states, thereby amplifying anticipated regret and reducing the attractiveness of the extended menu.

To compare with our framework, we represent the menu-choice problem using a class of game forms that mirrors the essence of Sarver’s setup. We dispense with lotteries and assume instead that choices in menus yield state-dependent monetary payoffs. There is a set of states Ω and a set of “basic actions” B that can be chosen by player i . An action-state pair $(b, \omega) \in B \times \Omega$ yields monetary payoff $m_{b, \omega}$. Chance/nature first selects a state $\omega \in \Omega$ with probability $\sigma_0(\omega)$. Player i then chooses a nonempty menu $M \subseteq B$ in a feasible set $\mathcal{M} \subseteq 2^B \setminus \{\emptyset\}$ without observing ω .⁵² We assume that all the singletons $\{b\}$ ($b \in B$) are contained in $\mathcal{M}$. If the chosen menu contains multiple actions, he continues to choose one of them with no information about ω (if the menu is a singleton, he is forced to choose the only action in the menu). The game form features *perfect feedback*, thus ω is observed ex post. The set of strategies is $S_i = \mathcal{M} \times (\times_{M \in \mathcal{M}} M)$ and the set

⁵²Thus, according to our general notation, $m_i(\omega, M, b) = m_{b, \omega}$ for every menu M containing b .

of **reduced strategies** is $\{(M, b) \in \mathcal{M} \times B : b \in M\}$. Since our analysis of decision problems is invariant to coalescing sequential moves, we can focus on reduced strategies. We relate preference $\succsim$ over menus to expected material payoffs and to the expected psychological utility implied by (reduced) strategies.




Figure 17: Menu-choice game form

Example 13 In Figure 17, we illustrate an instance of such a game form. Chance first selects $\omega \in \{\omega_1, \omega_2\}$. Player i subsequently chooses a menu from $\mathcal{M} = \{\{a\}, \{b\}, \{a, b\}, \{c\}\}$. If the menu $M = \{a, b\}$ is chosen, player i then chooses between a and b . At each terminal node, player i recalls his choice and learns ω ; thus, he earns the payoff he obtains and, when the chosen menu is not a singleton, the payoff he could have obtained from the other choice.

Note that in this setup own-strategy independence of feedback about others (OSIFO) is trivially satisfied: player i observes realization ω independently of the adopted (reduced) strategy (see Section 2.1). Therefore, under *linear regret*, preferences satisfy expected payoff maximization (see Theorem 1): $M \succsim M'$ if and only if

$$\max_{a \in M} \sum_{\omega \in \Omega} \pi_0(\omega) m_{a,\omega} \geq \max_{b \in M'} \sum_{\omega \in \Omega} \pi_0(\omega) m_{b,\omega}.$$

It follows that Sarver's central axiom holds as an indifference: if $a \in M$ and $\{a\} \succsim \{b\}$, then $M \sim M \cup \{b\}$. The reason is that player i can feel regret also for an ex post suboptimal choice of menu as illustrated by the following elementary example.

Example 14 Consider the game form of Example 13 and suppose that $m_{a1} = m_{a2} = 1$, $m_{b1} = 0$, $m_{b2} = 2$, $m_{c1} = m_{c2} = 0$, and $\frac{1}{2} < p < 1$. Both $(\{a\}, a)$ and $(\{a, b\}, a)$ are ex ante optimal, $\{a\} \succ \{b\}$, and $\{a\} \sim \{a, b\}$. Both $(\{b\}, b)$ and $(\{a, b\}, b)$ are ex post optimal after ω_2 is observed. Whether $(\{a\}, a)$

or $(\{a, b\}, a)$ is chosen, the regret after observing ω_2 is the same, i.e., $2 - 1 = 1$. In words, after ω_2 is observed, $\{a\}$ is an ex post suboptimal choice of menu and $(\{a, b\}, a)$ is an equally ex post suboptimal choice of basic action within menu $\{a, b\}$.

We note that the central axiom of Sarver may also hold as a strict inequality in our setup if we adopt a different assumption about terminal information that is consistent with Sarver’s analysis: player i observes ex post only the payoffs of actions in the chosen menu. This is a form of imperfect feedback. Formally,

$$\bar{H}_i(\omega, M, b) = \{(\omega', M', b') : (M', b') = (M, b), \forall a \in M, m_{a, \omega'} = m_{a, \omega}\}$$

for all (ω, M, b) .⁵³ While this different assumption about the information structure does not eliminate regret for the choice of menu, it nonetheless introduces a benefit of choosing smaller menus: OSIFO does not hold anymore, small menus reveal less about the state and, by Proposition 7 (information avoidance), may be more attractive when the same basic action is chosen afterward. To illustrate, modifying the previous example according to this assumption, $(\{a\}, a)$ does not reveal the state while $(\{a, b\}, a)$ does. Since these two reduced strategies yield the same expected payoff, the former is preferred because it entails less anticipated regret.

A recent extension of Sarver’s framework is provided by Wang (2022), who introduces an additional decision stage by allowing the decision maker to choose among collections of menus and studies information acquisition. In his setting, the decision maker first selects a collection of menus together with a Blackwell experiment, then chooses a menu, after which information is realized and a lottery is selected within the menu. Unlike Sarver, regret in Wang’s main representation arises from comparisons across menus within the chosen collection, rather than from the choice of the collection itself or from the choice of lotteries.⁵⁴ Within this setting, Wang studies the trade-off between the acquisition of instrumentally valuable information and information avoidance driven by regret over the menu choice. This is related to our analysis in Section 5. However, our definition of regret is not restricted to a particular stage of choice, and information avoidance arises as an implication of the framework rather than being imposed axiomatically.

7.3 Minimax Regret

Another relevant concept is *minimax regret*, originally proposed by Savage (1951), and later axiomatized (Hayashi 2008; Stoye 2011). Minimax regret has been extended to strategic interaction as a solution concept (Renou & Schlag 2010, 2011; Halpern & Pass 2012), and serves as a normative criterion in robust mechanism design (Bergemann & Schlag 2008, 2011).

⁵³In Sarver’s formalism, after choosing menu A and any lottery in A , the decision maker only observes the restriction $u|_A$ of utility-state u . This yields the same valuation of menus as in Sarver’s analysis.

⁵⁴He also provides an extension that allows for an additive regret concern from the choice of lottery.

Although it incorporates the notion of regret, the minimax regret principle does not aim to capture the actual experience of regret. Rather, it serves as a pragmatic decision rule, particularly suited to environments where a probability distribution over states of nature is unavailable.

Following this principle, the DM chooses a mixed action to minimize the worst-case regret—the greatest possible regret induced by any state of nature. Using our notation, the minimax regret principle for player i is formally expressed as:

$$\min_{\sigma_i \in \Delta(S_i)} \max_{s_0 \in S_0} \left\{ \sum_{s_i \in S_i} \sigma_i(s_i) \left[\max_{s'_i \in S_i} m_i(s'_i, s_0) - m_i(s_i, s_0) \right] \right\}. \quad (15)$$

Central to this principle is the focus on worst-case regret, which depends on the action being evaluated. As shown in Equation (15), the evaluation is independent of both the objective chance probabilities π_0 and the player’s subjective conditional beliefs α_i . This allows the DM to apply the principle even under complete ignorance about the probabilities of different states. In contrast, our model assumes players maximize their expected utility, which depends on π_0 (which shapes α_i) and on their terminal beliefs $(\alpha_i(\cdot|\bar{h}))_{\bar{h} \in Z_i}$.

Similar to DRT, the minimax regret principle also assumes that the DM can observe the state of nature (strategy of Chance) *ex post*, allowing him to evaluate regret as defined in Equation (15). Accordingly, the setup can be represented as a single-player game with Chance and essentially simultaneous moves and perfect feedback, such as the game form shown in Example 4. In that example, player i , who is assumed to have linear regret, prefers action s over r iff the probability of state b , λ , is at least $\frac{1}{2}$. Under the minimax regret principle, however, letting $p := \sigma_i(s)$, Equation (15) becomes $\min_{p \in [0,1]} \max\{1 - p, p\}$. The unique solution is $p^* = 1/2$, regardless of the value of λ .

In sum, while the minimax regret principle shares the use of regret as a key component, its normative, belief-independent nature contrasts sharply with our model.

8 Concluding remarks

People regret a lot of things. It is illuminating to conduct Google searches or to ask LLMs to get a feeling for that. For example, when we entered the search terms “regret death bed” the following statements popped up: “I wish I’d had the courage to live a life true to myself, not the life others expected of me.” ... “I wish I hadn’t worked so hard.” ... “I wish I’d had the courage to express my feelings.” ... “I wish I had stayed in touch with my friends.” ... “I wish I had let myself be happier.” We feel that a good way to wrap up our paper is to describe them. Doing so highlights several aspects that we feel are relevant when assessing our work and useful for inspiring further work on related topics.

The quotes suggest that regret can matter importantly for human happiness. Nevertheless,

economists have been slow to study how regret shapes decisions and life outcomes. To do that, one needs a generally applicable model, and, to the best of our knowledge, ours is the first comprehensive framework offered for regret in games. We hope that our contribution represents an important step in allowing economists to explore whether and how regret shapes economic outcomes.

Comprehensive as our framework may be, the examples we have scrutinized, to a degree, may seem to not address themes as grand as those signaled by the quotes (re, e.g., courage, expression of feelings, work, friends, happiness, etc.). While we have not only developed a basic framework but also applied it to explore several important economic situations, we invite further applied work using the tools we have developed. Our mathematical framework also leaves room for extensions. In particular, extending the analysis to multi-period games would allow regret to occur before the game ends (cf. Battigalli *et al.* 2019, Section 4).

There is a sense, however, in which the above quotes may be misleading as regards when and how regret matters. They by and large illustrate *painful* episodes when regret actually *occurred*. We wish to emphasize that similar episodes often shape outcomes by being *avoided*. That is, a particular regret episode (conditional on a possible course of action) is anticipated but ultimately leads a player to behave in a way that renders that particular regret episode *counterfactual* (even though regret need not disappear altogether from the realized path).⁵⁵

Perhaps there is a more general lesson in this. Think of all the world outcomes where people are not suffering much from regret, e.g., instances where someone studied hard, broke up a relationship, or became a professional golfer, and in the end felt that they had by and large lived their life well. These observations do not constitute counter-examples to the idea that regret is important to human decision-making, as anticipated counterfactual regret episodes may still have played a key role in shaping past behavior and overall outcomes. We close our paper with some lyrics which we read as a powerful and plausible reflection of this possibility. Take it from Paul Anka & Frank Sinatra:

Regrets, I've had a few
But then again, too few to mention
I did what I had to do
And saw it through without exception

⁵⁵It is interesting to compare our observations here to what can be said of other emotions, such as guilt, disappointment, or anger, which similarly may be rendered counterfactual because they carry negative valence or otherwise undesirable consequences. Battigalli & Dufwenberg (2022, p. 843) make a related observation concerning a difference between the natural focus of economists and that of psychologists. For economists, counterfactual emotions are obviously important if they influence behavior and who gets what. Psychologists, by contrast, tend to focus on the impact of emotional experiences that actually occur.

9 Appendix

Proof of Remark 2 Part (1) obviously holds by inspection of the relevant definitions. We prove part (2). Suppose that Γ features asymmetric information with essentially simultaneous moves and ex post observed actions of personal players. By essentially simultaneous moves of personal players, for each player i , the interim information partition $\mathcal{H}_i$ is equivalent to a corresponding “type” partition $\mathcal{T}_i$ of $\Omega = A_0(\emptyset)$ according to a bijection η_i between $\mathcal{T}_i$ and $\mathcal{H}_i$: for all $\omega', \omega'' \in \Omega$, $t_i, t'_i, t''_i \in \mathcal{T}_i$, we have $\omega', \omega'' \in t_i \in \mathcal{T}_i$ if and only if there is one and only one information set $h = \eta_i(t_i) \in \mathcal{H}_i$ with two nodes (histories) $x', x'' \in h$ following ω' and ω'' respectively, whereas if ω' and ω'' belong to distinct cells t'_i and t''_i of $\mathcal{T}_i$ then any two successors $x', x'' \in X_i$ of ω' and ω'' belong to distinct information sets $\eta_i(t'_i) = h'$ and $\eta_i(t''_i) = h''$, that is, $(\omega') \prec x' \in h'$, $(\omega'') \prec x'' \in h''$, with $h' \cap h'' = \emptyset$. Therefore, we can regard strategies as maps from $\mathcal{T}_i$ to A_i .⁵⁶ For each $j \in I$, let $T_j(\omega)$ denote the cell of $\mathcal{T}_j$ containing ω . Given that actions of personal players (but not the private information of others) are observed ex post, $\mathcal{Z}_i$ is isomorphic to $A_i \times A_{-i} \times \mathcal{T}_i$, and $F_i(s, \omega)$ is isomorphic to $\times_{j \in I} \{s_j(T_j(\omega))\} \times T_i(\omega)$ for every $(s, \omega) \in S \times \Omega$, where $(s_j(T_j(\omega)))_{j \in I}$ is the action profile selected by strategy profile s at ω . With this, for every $s_i : \mathcal{T}_i \rightarrow A_i$,

$$\mathcal{F}_i(s_i) \cong \{(s_{-i}, \omega) \in S_{-i} \times \Omega : T_i(\omega) = t_i, \forall j \in I \setminus \{i\}, s_j(T_j(\omega)) = a_j\}_{(a_{-i}, t_i) \in A_{-i} \times \mathcal{T}_i}$$

(read $\cong$ as “equal up to relabeling”), where the r.h.s. is independent of s_i . ■

Proof of Lemma 1 By definition, $(\bar{s}_{-i}, \bar{s}_0)$ is possible under both terminal information sets $\bar{H}_i(\zeta(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0))$ and $\bar{H}_i(\zeta(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0))$, that is,

$$(\bar{s}_{-i}, \bar{s}_0) \in S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0))) \cap S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0))).$$

Since $S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0))) \in \mathcal{F}_i(\bar{s}'_i)$, $S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0))) \in \mathcal{F}_i(\bar{s}''_i)$, and $\mathcal{F}_i(\bar{s}'_i) = \mathcal{F}_i(\bar{s}''_i)$, it follows that

$$S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}'_i, \bar{s}_{-i}, \bar{s}_0))) = S_{I_0 \setminus \{i\}}(\bar{H}_i(\zeta(\bar{s}''_i, \bar{s}_{-i}, \bar{s}_0)))$$

as both sets must be the same cell of the same ex post information partition of $S_{-i} \times S_0$. This implies the result. ■

Proof of Theorem 1 We prove that, for any strategy $\bar{s}_i \in S_i(h)$ played at h and following information sets,

$$\mathbb{E}_{\bar{s}_i, \alpha_i}[u_i|h] = (1 + \theta_i) \mathbb{E}_{\bar{s}_i, \alpha_i}[m_i|h] - \theta_i M_i(h, \alpha_i), \quad (16)$$

⁵⁶When moves are truly simultaneous – that is, $\iota(\omega) = I$ for every $\omega \in \Omega$ – this is literally true: for every $i \in I$, $\mathcal{T}_i = \mathcal{H}_i$ and $S_i = A_i^{\mathcal{T}_i}$.

where $M_i(h, \alpha_i)$ denotes the expectation conditional on h of the maximum hindsight-expected payoff, and is independent of $\bar{s}_i$. This implies the result.

The expectation conditional on h of the ex post expected payoff when $\bar{s}_i$ is played is

$$\mathbb{E}_{\bar{s}_i, \alpha_i} [\bar{m}_i(\alpha_i, \bar{H}_i(\zeta(\bar{s}_i, \cdot))) | h].$$

By consistency of α_i , the law of iterated expectations holds and the expectation conditional on h of the expectation conditional on terminal information of m_i coincides with the expectation conditional on h of m_i :

$$\textbf{Claim 1} \quad \mathbb{E}_{\bar{s}_i, \alpha_i} [\bar{m}_i(\bar{H}_i(\zeta(\bar{s}_i, \cdot)), \alpha_i) | h] = \mathbb{E}_{\bar{s}_i, \alpha_i} [\mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | \bar{H}_i(\zeta(\bar{s}_i, \cdot))] | h] = \mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | h].$$

Proof of Claim 1 The first equality holds by definition of $\bar{m}_i(\bar{h}, \alpha_i)$. We derive the second equality, explicitly writing in α_i the conditioning events corresponding to information sets, that is, for any $\hat{h} \in \bar{\mathcal{H}}_i$, we write $\alpha_i(\cdot | S_{I_0 \setminus \{i\}}(\hat{h}))$ instead of the abbreviation $\alpha_i(\cdot | \hat{h})$. Also, let

$$\mathcal{F}_i(\bar{s}_i, h) := \{E_{I_0 \setminus \{i\}} \in \mathcal{F}_i(\bar{s}_i) : E_{I_0 \setminus \{i\}} \subseteq S_{I_0 \setminus \{i\}}(h)\}.$$

By perfect recall, $\mathcal{F}_i(\bar{s}_i, h)$ is a partition of $S_{I_0 \setminus \{i\}}(h)$. With this,

$$\begin{aligned} & \mathbb{E}_{\bar{s}_i, \alpha_i} [\mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | \bar{H}_i(\zeta(\bar{s}_i, \cdot))] | h] = \\ & \sum_{E_{I_0 \setminus \{i\}} \in \mathcal{F}_i(\bar{s}_i, h)} \alpha_i(E_{I_0 \setminus \{i\}} | S_{I_0 \setminus \{i\}}(h)) \sum_{(s_{-i}, s_0) \in E_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0 | E_{I_0 \setminus \{i\}}) m_i(\zeta(\bar{s}_i, s_{-i}, s_0)) = \\ & \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | E_{I_0 \setminus \{i\}}) \alpha_i(E_{I_0 \setminus \{i\}} | S_{I_0 \setminus \{i\}}(h)) m_i(\zeta(\bar{s}_i, s_{-i}, s_0)) = \\ & \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | S_{I_0 \setminus \{i\}}(h)) m_i(\zeta(\bar{s}_i, s_{-i}, s_0)) = \mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | h], \end{aligned}$$

where the first and second equalities hold because $\mathcal{F}_i(\bar{s}_i, h)$ is a partition of $S_{I_0 \setminus \{i\}}(h)$, the third follows from the chain rule (by consistency of α_i), and the last equality holds by definition. $\square$

Corollary 1 implies that, by OSIFO, the maximum hindsight-expected payoff in (4) is independent of the played strategy. Thus, by linearity of regret and Claim 1, the expectation of (6) conditional on h given $\bar{s}_i$ and α_i is

$$\begin{aligned} \mathbb{E}_{\bar{s}_i, \alpha_i} [u_i | h] &= \mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | h] - \theta_i \mathbb{E}_{\bar{s}_i, \alpha_i} \left[\max_{s_i \in S_i} \mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | \bar{H}_i(\zeta(\bar{s}_i, \cdot))] | h \right] + \theta_i \mathbb{E}_{\bar{s}_i, \alpha_i} [\mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | \bar{H}_i(\zeta(\bar{s}_i, \cdot))] | h] \\ &= (1 + \theta_i) \mathbb{E}_{\bar{s}_i, \alpha_i} [m_i | h] - \theta_i M_i(h, \alpha_i). \quad \blacksquare \end{aligned}$$

Proof of Proposition 3 Since there are no chance moves, S_0 is a singleton, and we ease notation by eliminating the symbol s_0 in this proof. Let α_i satisfy the deterministic-beliefs Condition 1:

for any given nonterminal information set $h \in \mathcal{H}_i$, there exists a strategy profile $s_{-i}^* \in S_{-i}(h)$ such that $\alpha_i(s_{-i}^*|h) = 1$. Given $\alpha_i(s_{-i}^*|h) = 1$, player i 's expected utility simplifies to:

$$\begin{aligned}\mathbb{E}_{s_i, \alpha_i}[u_i|h] &= \sum_{s_{-i} \in S_{-i}(h)} \alpha_i(s_{-i}|h) \cdot u_i(\zeta(s_i^{|h|}, s_{-i}), \alpha_i) \\ &= u_i(\zeta(s_i^{|h|}, s_{-i}^*), \alpha_i) \\ &= m_i(\zeta(s_i^{|h|}, s_{-i}^*)) - \theta_i \cdot \phi_i \left(\max_{s'_i \in S_i} \mathbb{E}_{s'_i, \alpha_i}[m_i|\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))] - m_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right).\end{aligned}\tag{17}$$

Applying the *chain rule*, we derive the following condition for any $s_i \in S_i$ and any terminal information set $\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))$ resulting from the strategy profile $(s_i^{|h|}, s_{-i}^*)$:

$$\begin{aligned}\alpha_i(s_{-i}^*|h) &= \alpha_i \left(s_{-i}^* | \bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right) \cdot \alpha_i \left(S_{I \setminus \{i\}} \left(\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right) | h \right) \\ &= \alpha_i \left(s_{-i}^* | \bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right) \cdot 1.\end{aligned}$$

Since the leftmost term equals 1, it follows that $\alpha_i \left(s_{-i}^* | \bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right) = 1$. Thus, for any $s_i^{|h|} \in S_i(h)$, which leads to $\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))$, the expected material payoff of considering $s'_i \in S_i$ at this terminal information set satisfies:

$$\mathbb{E}_{s'_i, \alpha_i}[m_i|\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))] = m_i(\zeta(s'_i, s_{-i}^*)).\tag{18}$$

Substituting this into Equation (17) yields:

$$\mathbb{E}_{s_i, \alpha_i}[u_i|h] = m_i(\zeta(s_i^{|h|}, s_{-i}^*)) - \theta_i \cdot \phi_i \left(\max_{s'_i \in S_i} \{m_i(\zeta(s'_i, s_{-i}^*))\} - m_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right).$$

Since the term $\max_{s'_i \in S_i} \{m_i(\zeta(s'_i, s_{-i}^*))\}$ does not depend on $s_i^{|h|} \in S_i(h)$ and the expression

$$\phi_i \left(\max_{s'_i \in S_i} \{m_i(\zeta(s'_i, s_{-i}^*))\} - m_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right)$$

is (strictly) decreasing in $m_i(\zeta(s_i^{|h|}, s_{-i}^*))$, it follows that for any $s_i, \tilde{s}_i \in S_i$,

$$\begin{aligned}\mathbb{E}_{s_i, \alpha_i}[u_i|h] &\geq \mathbb{E}_{\tilde{s}_i, \alpha_i}[u_i|h] \\ \Leftrightarrow m_i(\zeta(s_i^{|h|}, s_{-i}^*)) &\geq m_i(\zeta(\tilde{s}_i^{|h|}, s_{-i}^*)) \\ \Leftrightarrow \mathbb{E}_{s_i, \alpha_i}[m_i|h] &\geq \mathbb{E}_{\tilde{s}_i, \alpha_i}[m_i|h].\end{aligned}$$

Therefore, s_i maximizes expected utility if and only if it maximizes expected material payoffs. Consequently, the set of maximizers is independent of θ_i . The result follows. ■

Proof of Proposition 5 By the observed-actions assumption, for each $i \in I$ we can replace information sets in $\mathcal{H}_i$ with decision nodes in X_i . Let s_i^* denote the equilibrium pure strategy of i as per Condition 2. By characterization (10) of full consistency in games with observed actions, it follows that, for all $j \in I \setminus \{i\}$ and $x \in X_j$, we have $\alpha_{j,i}((s_i^*)^{|x|} | x) = 1$, where $(s_i^*)^{|x|}$ denotes the strategy that coincides with s_i^* at all nodes not preceding x and selects i 's actions leading to x at each of i 's decision nodes preceding x . Since there are no chance moves, Condition 1 is satisfied with the adjusted strategies forming the required opponents' profile. That is, any sequential equilibrium in pure strategies satisfies deterministic beliefs. Thus, the result follows from Proposition 3. ■

Proof of Lemma 2 To ease notation, we consider the ex ante comparison at the root information set $h = \{\emptyset\}$ (the proof for other information sets is analogous).

For any $\hat{s}_i \in S_i$, since $M_i(\bar{H}_i(\zeta(\hat{s}_i, s_{-i}, s_0)), \alpha_i)$ depends on (s_{-i}, s_0) only through the induced terminal information set $\bar{h} = \bar{H}_i(\zeta(\hat{s}_i, s_{-i}, s_0))$, we have

$$\begin{aligned} & \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0) M_i(\bar{H}_i(\zeta(\hat{s}_i, s_{-i}, s_0)), \alpha_i) \\ &= \sum_{\bar{h} \in \mathcal{Z}_i(\hat{s}_i)} \alpha_i(S_{I_0 \setminus \{i\}}(\bar{h})) M_i(\bar{h}, \alpha_i) \\ &= \sum_{E_{I_0 \setminus \{i\}} \in \mathcal{F}_i(\hat{s}_i)} \alpha_i(E_{I_0 \setminus \{i\}}) M_i(E_{I_0 \setminus \{i\}}, \alpha_i), \end{aligned}$$

where $\mathcal{Z}_i(\hat{s}_i)$ is the collection of terminal information sets compatible with $\hat{s}_i$, $\alpha_i(S_{I_0 \setminus \{i\}}(\bar{h}))$ is the ex ante probability of reaching $\bar{h} \in \mathcal{Z}_i(\hat{s}_i)$, and — with a slight abuse of notation — we write $M_i(E_{I_0 \setminus \{i\}}, \alpha_i) = M_i(S_{I_0 \setminus \{i\}}(\bar{h}), \alpha_i) = M_i(\bar{h}, \alpha_i)$ for every terminal information set $\bar{h}$ and the corresponding set $E_{I_0 \setminus \{i\}} = S_{I_0 \setminus \{i\}}(\bar{h})$.

Since s'_i is at least as ex post informative as s''_i , each cell $E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)$ is the union of cells $E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})$, where

$$\mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}}) := \{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i) : E'_{I_0 \setminus \{i\}} \subseteq E''_{I_0 \setminus \{i\}}\}$$

is the partition of $E''_{I_0 \setminus \{i\}}$ induced by s'_i . Thus,

$$\begin{aligned} & \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i)} \alpha_i(E'_{I_0 \setminus \{i\}}) M_i(E'_{I_0 \setminus \{i\}}, \alpha_i) \\ &= \sum_{E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)} \sum_{E'_i \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) M_i(E'_{I_0 \setminus \{i\}}, \alpha_i), \end{aligned}$$

and we must show

$$\begin{aligned}
& \sum_{E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)} \alpha_i(E''_{I_0 \setminus \{i\}}) M_i(E''_{I_0 \setminus \{i\}}, \alpha_i) \\
& \leq \sum_{E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)} \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) M_i(E'_{I_0 \setminus \{i\}}, \alpha_i). \tag{19}
\end{aligned}$$

For each $E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)$, using the chain rule and the fact that $\mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})$ is a partition of $E''_{I_0 \setminus \{i\}}$ we obtain

$$\begin{aligned}
& \alpha_i(E''_{I_0 \setminus \{i\}}) M_i(E''_{I_0 \setminus \{i\}}, \alpha_i) \\
& = \alpha_i(E''_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} \sum_{(s_{-i}, s_0) \in E''_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0 | E''_{I_0 \setminus \{i\}}) m_i(\zeta(s_i, s_{-i}, s_0)) \\
& = \max_{s_i \in S_i} \sum_{(s_{-i}, s_0) \in E''_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0 | E''_{I_0 \setminus \{i\}}) \alpha_i(E''_{I_0 \setminus \{i\}}) m_i(\zeta(s_i, s_{-i}, s_0)) \\
& = \max_{s_i \in S_i} \sum_{(s_{-i}, s_0) \in E''_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0) m_i(\zeta(s_i, s_{-i}, s_0)) \\
& = \max_{s_i \in S_i} \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) \sum_{(s_{-i}, s_0) \in E'_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0 | E'_{I_0 \setminus \{i\}}) m_i(\zeta(s_i, s_{-i}, s_0)) \\
& \leq \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} \sum_{(s_{-i}, s_0) \in E'_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0 | E'_{I_0 \setminus \{i\}}) m_i(\zeta(s_i, s_{-i}, s_0)) \\
& = \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) M_i(E'_{I_0 \setminus \{i\}}, \alpha_i),
\end{aligned}$$

where the inequality holds because the maximum of a weighted sum is at most the weighted sum of maxima. Taking the summation of both sides over $E''_{I_0 \setminus \{i\}} \in \mathcal{F}_i(s''_i)$ we obtain (19). ■

The continuity of u_i with respect to terminal beliefs as well as θ_i implies a quite standard upper-hemicontinuity result that is useful for equilibrium analysis.⁵⁷

Lemma 3 *Fix any $i \in I$, the sequential best reply correspondence $(\theta_i, \alpha_i) \mapsto BR_i^{\theta_i}(\alpha_i)$ is upper-hemicontinuous: for every $(\hat{\theta}_i, \hat{\alpha}_i, \hat{s}_i) \in \mathbb{R}_+ \times \Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}}) \times S_i$ and every sequence $(\theta_i^n, \alpha_i^n, s_i^n)_{n \in \mathbb{N}} \in (\mathbb{R}_+ \times \Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}}) \times S_i)^{\mathbb{N}}$ converging to $(\hat{\theta}_i, \hat{\alpha}_i, \hat{s}_i)$ such that $s_i^n \in BR_i^{\theta_i^n}(\alpha_i^n)$ for all n , $\hat{s}_i \in BR_i^{\hat{\theta}_i}(\hat{\alpha}_i)$.*

⁵⁷Endow the set $\Delta^{\bar{\mathcal{H}}_i}(S_{I_0 \setminus \{i\}}) \subseteq \mathbb{R}^{\bar{\mathcal{H}}_i \times S_{I_0 \setminus \{i\}}}$ of consistent systems of beliefs with the topology induced by the Euclidean metric of $\mathbb{R}^{\bar{\mathcal{H}}_i \times S_{I_0 \setminus \{i\}}}$, and S_i with the discrete topology.

Proof of Lemma 3 Fix $(\hat{\theta}_i, \hat{\alpha}_i, \hat{s}_i)$ and a sequence $(\theta_i^n, \alpha_i^n, s_i^n)_{n \in \mathbb{N}} \xrightarrow{n \rightarrow \infty} (\hat{\theta}_i, \hat{\alpha}_i, \hat{s}_i)$ such that $s_i^n \in BR_i^{\theta_i^n}(\alpha_i^n)$ for all n . Then, for all $h \in \mathcal{H}_i$, $s'_i \in S_i(h)$, and n ,

$$\sum_{s_{-i,0} \in S_{I_0 \setminus \{i\}}} \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \left[u_i^{\theta_i^n} \left(\zeta \left(s_i^{n|h}, s_{-i}, s_0 \right), \alpha_i^n \right) - u_i^{\theta_i^n} \left(\zeta \left(s'_i, s_{-i}, s_0 \right), \alpha_i^n \right) \right] \alpha_i^n(s_{-i}, s_0 | h) \geq 0,$$

where $u_i^{\theta_i^n}$ and $s_i^{n|h}$ respectively denote the utility function with sensitivity parameter θ_i^n and the h -replacement of s_i^n . Since the utility of terminal nodes is continuous in terminal beliefs and in θ_i , as $n \rightarrow \infty$, for all $h \in \mathcal{H}_i$, $s'_i \in S_i(h)$, we obtain

$$\sum_{s_{-i,0} \in S_{I_0 \setminus \{i\}}} \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \left[u_i^{\hat{\theta}_i} \left(\zeta \left(\hat{s}_i^h, s_{-i}, s_0 \right), \hat{\alpha}_i \right) - u_i^{\hat{\theta}_i} \left(\zeta \left(s'_i, s_{-i}, s_0 \right), \hat{\alpha}_i \right) \right] \hat{\alpha}_i(s_{-i}, s_0 | h) \geq 0,$$

that is, $\hat{s}_i \in BR_i^{\hat{\theta}_i}(\hat{\alpha}_i)$. ■

Proof of Proposition 6 Let $G^{\theta_n} \rightarrow G^0$ be a convergent sequence of regret games, and let $(\pi^n, \alpha^n)_{n \in \mathbb{N}}$ be a sequence of SEs of games G^{θ_n} , with $(\pi^n, \alpha^n) \rightarrow (\pi, \alpha)$. We want to show that (π, α) is an SE of G^0 . This implies the thesis. Sequential rationality of (π, α) follows from Lemma 3.

Since, for each n , (π^n, α^n) is an SE of G^{θ_n} , by full consistency, there exists a sequence of fully mixed behavior strategy profiles $\pi^{k_n} \rightarrow \pi^n$ (with the corresponding induced beliefs α^{k_n}) s.t. $(\pi^{k_n}, \alpha^{k_n}) \rightarrow (\pi^n, \alpha^n)$, as $k_n \rightarrow \infty$. We show that (π, α) satisfies full consistency: This follows from the fact that the set of fully consistent assessments is closed. To see this more explicitly, for each n , choose k_n large enough so that π^{k_n} and α^{k_n} are within distance $\frac{1}{n}$ of π^n and α^n , respectively. Since $\pi^n \rightarrow \pi$ and $\alpha^n \rightarrow \alpha$, we have $\pi^{k_n} \rightarrow \pi$ and $\alpha^{k_n} \rightarrow \alpha$. Thus, the graph of $SE(G^\theta)$ is closed. Because $SE(G^\theta)$ is also compact-valued, the closed graph theorem implies that SE is upper-hemicontinuous.

Moreover, we proved in the main text that the inclusion may be strict. ■

Proof of Proposition 7 We again consider the ex ante comparison at $h = \{\emptyset\}$; the general case is analogous.

As in the proof of Theorem 1, under linear regret the expected utility of playing any strategy $\bar{s}_i \in S_i$ given belief system α_i is

$$\mathbb{E}_{\bar{s}_i, \alpha_i}[u_i] = (1 + \theta_i) \mathbb{E}_{\bar{s}_i, \alpha_i}[m_i] - \theta_i \mathbb{E}_{\bar{s}_i, \alpha_i} \left[M_i \left(\bar{H}_i(\zeta(\bar{s}_i, \cdot)), \alpha_i \right) \right].$$

Since $\mathbb{E}_{\alpha_i, s'_i}[m_i] \leq \mathbb{E}_{\alpha_i, s''_i}[m_i]$ (resp. $<$), to establish $\mathbb{E}_{\alpha_i, s'_i}[u_i] \leq \mathbb{E}_{\alpha_i, s''_i}[u_i]$ (resp. $<$), it suffices to

show that

$$\mathbb{E}_{\alpha_i, s_i''} [M_i (\bar{H}_i (\zeta (s_i'', \cdot)), \alpha_i)] \leq \mathbb{E}_{\alpha_i, s_i'} [M_i (\bar{H}_i (\zeta (s_i', \cdot)), \alpha_i)],$$

which follows from inequality (19) in the proof of Lemma 2.

■

Lemma 4 *In the first-price auction model with private values in Section 6.1, under hiding losing bids, lower bids yield a finer ex post information partition.*

Proof of Lemma 4. Consider the hiding losing bids information structure, while *the loser observes the winning bid*. Consider two bids $b_1' < b_1''$. Taking into account that bidder 1 wins if there is a tie, the ex post information partitions induced by b_1' and b_1'' are

$$\begin{aligned} \mathcal{F}_1^{\text{HL}}(b_1') &= \left\{ \{0, \dots, b_1'\}, \{\{b_2\}\}_{b_2 > b_1'} \right\}, \\ \mathcal{F}_1^{\text{HL}}(b_1'') &= \left\{ \{0, \dots, b_1''\}, \{\{b_2\}\}_{b_2 > b_1''} \right\}. \end{aligned}$$

These partitions feature a collection of singletons above the given bid of 1, and a residual cell of bids weakly below the bid of 1. Since $b_1' < b_1''$ and the loser observes the competitor's bid ex post, $\mathcal{F}_1^{\text{HL}}(b_1')$ refines $\mathcal{F}_1^{\text{HL}}(b_1'')$. Indeed, letting δ denote the minimal discrete increment,

$$\{0, \dots, b_1'\} \subset \{0, \dots, b_1''\},$$

$$\forall b_2 \in \{b_1' + \delta, \dots, b_1''\}, \{b_2\} \subset \{0, \dots, b_1''\}$$

and every bid $b_2 > b_1''$ is observed under both bids of 1. Thus, *the lower bid is more informative under the hiding losing bids information feedback regime.*⁵⁸ ■

Proof of Remark 4 By construction, $(\Gamma, \tilde{u}_i)$ features essentially simultaneous moves and perfect feedback. Thus, i 's regret function at an arbitrary terminal information set $h = \{(a, \omega)\}$ is given by:

$$\begin{aligned} \tilde{r}_i(h, \alpha_i) &= \phi_i \left(\max_{s_i' \in S_i} (v_i \circ m_i)(s_i', \omega) - (v_i \circ m_i)(a, \omega) \right), \\ &= \phi_i \left(\max_{a' \in A} V(a'(\omega)) - V(a(\omega)) \right), \end{aligned}$$

⁵⁸ An analogous result holds for bidder 2.

where the second equality follows from condition (i). Therefore,

$$\begin{aligned}
\tilde{u}_i((a, \omega), \alpha_i) &= (v_i \circ m_i)(a, \omega) - \theta_i \cdot \tilde{r}_i((a, \omega), \alpha_i) \\
&= V(a(\omega)) - f\left(\max_{a' \in A} V(a'(\omega)) - V(a(\omega))\right) \\
&= U(a, \omega),
\end{aligned}$$

where the second equality follows from condition (i) and the final equality follows from condition (ii). ■

Proof of Corollary 4 This result follows from Proposition 1, relying crucially on the linearity of regret, i.e., ϕ_i being the identity function. The argument proceeds analogously, with m_i replaced by $(v_i \circ m_i)$. ■

References

- [1] ARMANTIER, O., J. FONCEL, AND N. TREICH (2023): “Insurance and Portfolio Decisions: Two Sides of the Same Coin?” *Journal of Financial Economics*, 148, 201–219.
- [2] BATTIGALLI, P. (1996): “Strategic Independence and Perfect Bayesian Equilibria,” *Journal of Economic Theory*, 70, 201–234.
- [3] BATTIGALLI, P., AND M. DUFWENBERG (2009): “Dynamic Psychological Games,” *Journal of Economic Theory*, 144, 1–35.
- [4] BATTIGALLI, P., AND M. DUFWENBERG (2022): “Belief-Dependent Motivations and Psychological Game Theory,” *Journal of Economic Literature*, 60, 833–882.
- [5] BATTIGALLI, P., AND N. GENEROSO (2024): “Information Flows and Memory in Games,” *Games & Economic Behavior*, 145, 356–376.
- [6] BATTIGALLI, P., E. CATONINI, AND J. MANILI (2023): “Belief Change, Rationality, and Strategic Reasoning in Sequential Games,” *Games & Economic Behavior*, 142, 527–551.
- [7] BATTIGALLI, P., R. CORRAO, AND M. DUFWENBERG (2019): “Incorporating Belief-Dependent Motivation in Games,” *Journal of Economic Behavior & Organization*, 167, 185–218.
- [8] BATTIGALLI, P., P. LEONETTI, AND F. MACCHERONI (2020): “Behavioral Equivalence of Extensive Game Structures,” *Games & Economic Behavior*, 121, 533–547.

- [9] BATTIGALLI, P., S. CERREIA, F. MACCHERONI, AND M. MARINACCI (2015): “Self-Confirming Equilibrium and Model Uncertainty,” *American Economic Review*, 105, 646–677.
- [10] BATTIGALLI, P., S. CERREIA, F. MACCHERONI, AND M. MARINACCI (2016): “Analysis of Information Feedback and Selfconfirming Equilibrium,” *Journal of Mathematical Economics*, 66, 40–51.
- [11] BELL, D. (1982): “Regret in Decision Making under Uncertainty,” *Operations Research*, 30, 961–981.
- [12] BERGEMANN, D., AND K. SCHLAG (2008): “Pricing Without Priors,” *Journal of the European Economic Association*, 6, 560–569.
- [13] BERGEMANN, D., AND K. SCHLAG (2011): “Robust Monopoly Pricing,” *Journal of Economic Theory*, 146, 2527–2543.
- [14] BERGEMANN, D., K. BREUER, P. CRAMTON, J. HIRSCH, Y. NDIAYE, AND A. OCKENFELS (2025): “Soft-Floor Auctions: Harnessing Regret to Improve Efficiency and Revenue,” Cowles Foundation Discussion Paper 2438.
- [15] BIKHCHANDANI, S., AND U. SEGAL (2011): “Transitive Regret,” *Theoretical Economics*, 6, 95–108.
- [16] BIKHCHANDANI, S., AND U. SEGAL (2014): “Transitive Regret over Statistically Independent Lotteries,” *Journal of Economic Theory*, 152, 237–248.
- [17] BRAUN, M., AND A. MUERMANN (2004): “The Impact of Regret on the Demand for Insurance,” *Journal of Risk and Insurance*, 71, 737–767.
- [18] CADISH, B. (2001): “Damn!: Reflections on Life’s Biggest Regrets”, *Kansas City: Andrews McMeel Publishing*, 2.
- [19] CAMILLE, N., G. CORICELLI, J. SALLET, P. PRADAT-DIEHL, J.-R. DUHAMEL, AND A. SIRIGU (2004): “The Involvement of the Orbitofrontal Cortex in the Experience of Regret,” *Science*, 304, 1167–1170.
- [20] CERRONE, C., F. FERI, AND P. NEARY (2025): “Ignorance is Bliss: A Game of Regret,” mimeo.
- [21] CONNOLLY, T., AND D. BUTLER (2006): “Regret in Economic and Psychological Theories of Choice,” *Journal of Behavioral Decision Making*, 19, 139–154.
- [22] CORICELLI, G., H. D. CRITCHLEY, M. JOFFILY, J. P. O’DOHERTY, A. SIRIGU, AND R. J. DOLAN (2005): “Regret and Its Avoidance: A Neuroimaging Study of Choice Behavior,” *Nature Neuroscience*, 8, 1255–1262.

- [23] DEKEL, E., B. LIPMAN, AND A. RUSTICHINI (2001): “Representing Preferences with a Unique Subjective State Space,” *Econometrica*, 69, 891–934.
- [24] ENGELBRECHT-WIGGANS, R. (1989): “The Effect of Regret on Optimal Bidding in Auctions,” *Management Science*, 35, 685–92.
- [25] ENGELBRECHT-WIGGANS, R., AND E. KATOK (2008): “Regret and Feedback Information in First-Price Sealed-Bid Auctions,” *Management Science* 54, 808–819.
- [26] FILIZ-OZBAY, E., AND E. OZBAY (2007): “Auctions with Anticipated Regret: Theory and Experiment,” *American Economic Review*, 97, 1407–1418.
- [27] FIORETTI, M., A. VOSTROKNUTOV, AND G. CORICELLI (2022): “Dynamic Regret Avoidance,” *American Economic Journal: Microeconomics*, 14, 70–93.
- [28] FRYDMAN, C., AND C. CAMERER (2016): “Neural Evidence of Regret and Its Implications for Investor Behavior,” *The Review of Financial Studies*, 29, 3108–3139.
- [29] GEANAKOPOLOS, J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games & Economic Behavior*, 1, 60–80.
- [30] GOLDFARB, A., AND M. XIAO (2011): “Who Thinks about the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets,” *American Economic Review*, 101, 3130–3161.
- [31] HALPERN, J., AND R. PASS (2012): “Iterated Regret Minimization: A New Solution Concept,” *Games & Economic Behavior*, 74, 194–207.
- [32] HAYASHI, T. (2008): “Regret Aversion and Opportunity Dependence,” *Journal of Economic Theory*, 139, 242–268.
- [33] KOHLBERG, E., AND P. RENY (1997): “Independence on Relative Probability Spaces and Consistent Assessments in Game Trees,” *Journal of Economic Theory*, 75, 280–313.
- [34] KRÄHMER, D., AND R. STONE (2013): “Anticipated Regret as an Explanation of Uncertainty Aversion,” *Economic Theory*, 52, 709–728.
- [35] KREPS, D., AND R. WILSON (1982): “Sequential Equilibrium,” *Econometrica*, 50, 863–94.
- [36] KUHN, H. (1953): “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games II*, ed. by H. Kuhn & A. Tucker. Princeton University Press, 193–216.
- [37] LOOMES, G., AND R. SUGDEN (1982): “Regret Theory: An Alternative Theory of Rational Choice under Uncertainty,” *Economic Journal*, 92, 805–824.

- [38] LOOMES, G., AND R. SUGDEN (1987): “Some Implications of a More General Form of Regret Theory,” *Journal of Economic Theory*, 41, 270–287.
- [39] OSBORNE, M., AND A. RUBINSTEIN (1994): *A Course in Game Theory*, MIT Press.
- [40] QIN, J. (2015): “A Model of Regret, Investor Behavior, and Market Turbulence,” *Journal of Economic Theory*, 160, 150–174.
- [41] QUIGGIN, J. (1990): “Stochastic Dominance in Regret Theory,” *Review of Economic Studies*, 57, 503–11.
- [42] QUIGGIN, J. (1994): “Regret Theory with General Choice Sets,” *Journal of Risk and Uncertainty*, 8, 153–65.
- [43] RENOU, L., AND K. H. SCHLAG (2010): “Minimax Regret and Strategic Uncertainty,” *Journal of Economic Theory*, 145, 264–286.
- [44] RENOU, L., AND K. H. SCHLAG (2011): “Implementation in Minimax Regret Equilibrium,” *Games & Economic Behavior*, 71, 527–533.
- [45] SARVER, T. (2008): “Anticipating Regret: Why Fewer Options May Be Better,” *Econometrica*, 76, 263–305.
- [46] SAVAGE, L. (1951): “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, 46, 55–67.
- [47] SHEFRIN, H., AND M. STATMAN (1985): “The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence,” *Journal of Finance*, 40, 777–790.
- [48] STOYE, J. (2011): “Axioms for Minimax Regret Choice Correspondences,” *Journal of Economic Theory*, 146, 2226–2251.
- [49] SUGDEN, R. (1993): “An Axiomatic Foundation for Regret Theory,” *Journal of Economic Theory*, 60, 159–180.
- [50] SYAM, N., P. KRISHNAMURTHY, AND J. HESS (2008): “That’s what I thought I wanted? Miswanting and Regret for a Standard Good in a Mass-Customized World,” *Marketing Science*, 27, 379–397.
- [51] VAN DER PLOEG, F., AND A. REZAI (2019): “The Agnostic’s Response to Climate Deniers: Price carbon!” *European Economic Review*, 111, 70–84.
- [52] VRAVOSINOS, O. (2023): “Regret, Blame, and Division of Responsibility in Games,” SSRN Working Paper 4602748.

- [53] WANG, Z. (2022): “Regret and Information Avoidance,” SSRN Working Paper 4269560.
- [54] ZEELLENBERG, M., J. BEATTIE, J. VAN DER PLIGT, AND N. DE VRIES (1996): “Consequences of Regret Aversion: Effects of Expected Feedback on Risky Decision Making,” *Organizational Behavior & Human Decision Processes* 65, 148–158.
- [55] ZEELLENBERG, M. (1999) “Anticipated Regret, Expected Feedback and Behavioral Decision Making,” *Journal of Behavioral Decision Making*, 12, 93–106.
- [56] ZEELLENBERG, M., AND R. PIETERS (2007a): “A Theory of Regret Regulation 1.0,” *Journal of Consumer Psychology*, 17, 3–18.
- [57] ZEELLENBERG, M., AND R. PIETERS (2007b): “A Theory of Regret Regulation 1.1,” *Journal of Consumer Psychology*, 17, 29–35.
- [58] ZOU, T., B. ZHOU, AND B. JIANG (2020): “Product-Line Design in the Presence of Consumers’ Anticipated Regret.” *Management Science*, 66, 5665–5682.