



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: December 2025

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 December 2025

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 actually gets and what he believes he would have gotten had he chosen differently. Ex post beliefs are critical to that evaluation, and 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., investments, delegation, gambling, market entry, and information revelation.

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

JEL codes: C72; D91

*Battigalli: Bocconi University, IGIER & BIDSa, Italy; pierpaolo.battigalli@unibocconi.it. Dufwenberg: Purdue University, USA; University of Gothenburg, Sweden; mdufwenb@purdue.edu. Lin: Southwestern University of Finance & Economics, China; senranlin@outlook.com. We thank Carlo Andreatta, Samuele Dotta, Tore Ellingsen, Mark Machina, Silvia Meneghesso, 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.*

(Piet Hein)

Regret can be a powerful emotion, as the grook by Piet Hein exemplifies. People regret things they say (hurt someone’s 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 emotional pain. A substantial literature in psychology explores regret,¹ but few economists paid attention despite the potential of anticipated 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 and there is heterogeneous work on specific games.³ However, no comprehensive approach has been available. Until now; we develop a theory for a general class of game forms.

¹See, e.g., Zeelenberg (1999), Connolly & Butler (2006), Zeelenberg & Pieters (2007*a,b*), and references therein. Compare also Cadish’s 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).

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

³Engelbrecht-Wiggans (1989), Engelbrecht-Wiggans & Katok (2008), Filiz-Ozbay & Ozbay (2007), and Bergemann et al. (2025) study auctions; Syam, Krishnamurthy & Hess (2008) and Zou, Zhou & Jiang (2020) investigate product differentiation; Cerrone, Feri & Neary (2025) explore specific static games.

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




Figure 1 Two job-market game forms

Example 1 *The game forms in Figs. 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 $p \in (\frac{3}{4}, 1)$ 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 only 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).*

⁴ *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.*

Our analysis generates several insights with broad applied importance. A basic consideration is this: Regret involves rumination on what-might-have-been, had one chosen differently. Player i 's regret then depends on i 's post-play beliefs about how other players (including

⁴Whether (i) or (ii) is more realistic depends on circumstances (e.g., whether Bob has a friend who works "elsewhere").

Chance) would have responded to counterfactual choices. These belief-dependent pangs of regret are anticipated by i at earlier stages. It is therefore necessary to employ an expanded notion of utility relative to what is offered by traditional game theory. Rather, the modeling is based on the framework of *psychological game theory* (Geanakoplos, Pearce & Stacchetti 1989; Battigalli & Dufwenberg 2009; Battigalli, Corrao & Dufwenberg 2019). Whereas in traditional games, utility depends on players’ choices only, in “psychological games” utility depends also on certain beliefs.⁵ Example 1 illustrates: Bob’s utility following (*deny*, *take*) depends on his belief regarding his would-have-been material payoff had he chosen *reject*: these beliefs are different in cases (i) and (ii).

That basic consideration spawns multiple further insights: First, *information structure across terminal nodes* is critical to predictions, because it influences the post-play beliefs that shape regret. 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. (We postpone an illustration; Example 1 does not suffice.)

Third, we derive multiple results regarding the relevance or irrelevance of players caring about regret. In part, this line of 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 emerges 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. The intuition can be understood by drawing an analogy between how in Example 1 it matters to Bob’s preference whether or not G is revealed post-play, and how in decision problems it may matter whether or not there is post-choice state-revelation.

Fourth, we extend that line of inquiry (concerning the relevance, or irrelevance, of linear regret) to multi-player games. Our results depend on game forms’ features such as “essentially simultaneous-moves” or “observed actions” (terms defined in Section 2). Further related insights can be gleaned by noticing similarities between the insights (see the previous paragraph) in single-player games and their counterparts with multiple players. Example 1

⁵Regret 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.

serves as an illustration of how the relevance of regret depends on the game form and its information structure: (i) exhibiting irrelevance, or (ii) exhibiting relevance.

Section 2 presents our general approach. Section 3 analyzes game forms with a single personal player, 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. Section 7 compares our approach to several strands of related literature, including DRT which establishes the “third” insight flagged for above. Section 8 concludes. The Appendix contains some of the proofs.

2 Modeling regret

We consider finite sequential game forms with monetary payoffs played by players with perfect recall. Adding to a game form utility functions featuring aversion to regret, we obtain a regret game. We present the primitive and derived elements of the extensive-form representation that are essential in our analysis.

2.1 Game form

The game tree is a *finite* set of feasible sequences of action profiles closed under the weak **precedence** relation (“prefix-of”) \preceq , hence including the empty sequence \emptyset as its **root** (cf. Osborne & Rubinstein 1994, and Battigalli et al. 2020).⁶ A profile of two or more actions obtains when two or more players move simultaneously. We allow for the possibility 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. In 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.⁷ Adding to the game tree the vectors of chance probabilities and of monetary payoffs, we obtain the game form.

⁶Sequence $x = (x^1, \dots, x^k)$ weakly precedes, or is a prefix of, sequence $y = (y^1, \dots, y^\ell)$, written $x \preceq y$, if $k \leq \ell$ and $(x^1, \dots, x^k) = (y^1, \dots, y^k)$.

⁷Although 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).

Game tree Let I be the nonempty and *finite* set of personal **players**. If the game contains chance moves, we let $0 \notin I$ denote the **Chance** pseudo-player and $I_0 = I \cup \{0\}$ the player set including Chance. The rules of the game determine a finite tree $\bar{X} \subseteq \bigcup_{\ell=0}^L A^\ell$ of (partial or complete) **plays** (feasible sequences of actions profiles), where L is an upper bound on the height of the tree, $A = \bigcup_{\emptyset \neq J \subseteq I_0} \times_{j \in J} A_j$ 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 **player correspondence** $\iota : X \rightrightarrows I_0$ such that $A(x) = \times_{j \in \iota(x)} A_j(x)$, and $|A_j(x)| \geq 2$ for each $x \in X$ and $j \in \iota(x)$. In words, $\iota(x)$ is the set of players who are active and (if $|\iota(x)| \geq 2$) choose simultaneously immediately after play x . Since \bar{X} (canonically ordered by the “prefix-of” relation \preceq) is a tree, we use “play” and **node** as synonymous. For each $j \in I_0$, $X_j = \{x \in X : j \in \iota(x)\}$ is the set of nodes where j is **active**. The **information** of each personal player $i \in I$ is represented by a partition \bar{H}_i of $X_i \cup \{\emptyset\} \cup Z$ that refines $\{\{\emptyset\}, X_i \setminus \{\emptyset\}, Z\}$; That is, each personal player knows at least when the game starts and ends, and when he is active.⁸ Thus, denoting by $\bar{H}_i(y)$ the information set of i containing $y \in X_i \cup \{\emptyset\} \cup Z$, we have $\{\emptyset\} \in \bar{H}_i$, $\bar{H}_i(x) \subseteq X_i$ for each $x \in X_i$, and $\bar{H}_i(z) \subseteq Z$ for each $z \in Z$. We also let H_i and $H_i(\cdot)$ respectively denote the restrictions of \bar{H}_i and $\bar{H}_i(z)$ to the instances where i is active. Furthermore, we also introduce an information partition 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 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 $j \in I_0$, $h \in H_j$, and $x \in h$. We assume that H_0 and each \bar{H}_i satisfy the *perfect recall* property, which – for personal players – is a cognitive assumption (cf. Osborne & Rubinstein 1994, ch. 12, and Battigalli & Generoso 2024). This implies that the information sets of each player i are partially ordered by precedence.⁹ In graphical representations, distinct nodes in the same information set of the relevant player are joined by a dashed line (see Example 1).

⁸Even if player i is not active at the root, it is convenient to include his information state at the beginning of the game.

⁹Note that the ability to remember is a cognitive personal trait, not a property of the rules of the game. Our interpretation for personal players is that each $h \in H_i$ is a set of plays compatible with a *personal history* of signals received and actions played by i , which makes H_i objectively determined by the rules of the game. A sufficient condition that justifies this interpretation is that each player’s memory (a personal trait) is perfect. See Battigalli & Generoso (2024).

Game form To obtain a game form, we add the probabilities of chance moves and personal players' monetary (or material) payoffs. **Chance probabilities** are given by a *strictly positive* vector $(\pi_0(\cdot|h))_{h \in H_0} \in \times_{h \in H_0} \Delta^\circ(A_0(h))$.¹⁰ The (monetary) **payoff** of each $i \in I$ is determined by a function $m_i : Z \rightarrow \mathbb{R}$. The **game form** is the mathematical structure

$$\Gamma = \left\langle I, \bar{X}, \iota, \pi_0, H_0, (\bar{H}_i, m_i)_{i \in I} \right\rangle$$

representing the rules of the games.

To illustrate, in Example 1 there are no truly simultaneous moves: the players' map $\iota(\cdot)$ is single-valued and obvious from the figures.¹¹ 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

$$H_B(\omega, deny, take) = \{(G, deny, take), (\neg G, deny, take)\}$$

for each $\omega \in \{G, \neg G\}$. 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 $h \in H_j$. Thus, the set of strategies of $j \in I_0$ is the cross-product $S_j = \times_{h \in H_j} A_j(h)$. 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, terminal node) $z = \zeta(s, s_0)$ induced by each profile (s, s_0) .¹² 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

¹⁰For 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.

¹¹It 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., $\iota(x)$ contains at least two players for some $x \in X$).

¹²For all $\ell \in \{1, \dots, L\}$ and $(a^1, \dots, a^\ell) \in Z$, $\zeta(s_0, s) = (a^1, \dots, a^\ell)$ 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 $k \in \{1, \dots, \ell - 1\}$.

$S_{I_0}(h) = \bigcup_{x \in h} S_{I_0}(x)$. 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 $x \in h \in \bar{H}_i$. Thus, the factorization holds in games with $|I_0| = 2$.

Special assumptions about information Some of our results concern game forms with **essentially simultaneous moves**, that is, for each player $j \in I_0$, the collection of information sets where j is active contains only one element, viz. $H_j = \{h\}$, such that, for each terminal node $z \in Z$, there is a unique node $x \in h$ preceding z (i.e., a strict prefix of complete play z). This makes the path function $\zeta : S \times S_0 \rightarrow Z$ a bijection, i.e., $S \times S_0$ and Z are isomorphic, strategies coincide with actions, and the elements of Z are – essentially – profiles of actions.¹³ Furthermore, three nested special cases of information structure are worth mentioning. We list them 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 every information set $h \in \bigcup_{i \in I} \bar{H}_i$ (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.¹⁴

Since regret depends on terminal information, it is also relevant to consider associated properties. From most to least general, game form Γ features:

- **observed payoffs** if each monetary payoff function m_i is \bar{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 $i \in I$.

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 deviators¹⁵

¹³Game trees with essentially simultaneous moves can be “simultanized”: 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).

¹⁴Traditional 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.

¹⁵For example, for each realization $\omega \in \{G, \neg G\} = S_0$, $S(H_B(\omega, \text{deny})) = \{\text{deny}\} \times S_B \times S_0$ and, in the second version, $S(H_B(\omega, \text{deny}, \text{take})) = \{\text{take}\} \times \{\text{deny}\} \times S_0$.

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 hold or not: In some decision problems of interest terminal information does not resolve all the uncertainty about payoffs. For example, the player might care about some material consequence (e.g., consequences of global warming) that shall realize only after the player is dead.

Mixed strategy of Chance While $s \in S$ is endogenously determined by the strategic analysis of the game, 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:¹⁶

$$\sigma_0(s_0) = \prod_{h \in H_0} \pi_0(s_0(h) | h) \quad (1)$$

for all $s_0 \in S_0$.

2.2 Beliefs

Players’ beliefs are central to describing both belief-dependent utilities that incorporate regret and how players analyze the game. Conditional on $h \in \bar{H}_i$, player $i \in I$ holds belief $\alpha_i(\cdot | h) \in \Delta(S_{I_0 \setminus \{i\}}(h))$ about his co-players’ strategies (information-dependent actions), and $\alpha_i = (\alpha_i(\cdot | h))_{h \in \bar{H}_i}$ 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 $h \in \bar{H}_i$ 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 $\Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}})$ denote i ’s set of consistent (first-order) beliefs.

¹⁶See Kuhn (1953), and also Section 3 where we explain why we need “information sets of Chance.”

To gain some intuition, suppose that, for all $h \in \bar{H}_i$, i obtains logically independent information about Chance and the set of personal 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, consistency condition (2) is equivalent to the requirement that α_i is obtained from σ_0 and a conditional probability system $\alpha_{i,-i}$ on $(S_{I \setminus \{i\}}, \bar{H}_i)$ 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 $(S_{I \setminus \{i\}}, \bar{H}_i)$ is a vector $(\alpha_{i,-i}(\cdot | h))_{h \in \bar{H}_i} \in \times_{h \in \bar{H}_i} \Delta(S_{I \setminus \{i\}}(h))$ that satisfies the chain rule: for all $h, h' \in \bar{H}_i$ and $s_{-i} \in S_{I \setminus \{i\}}$, if $s_{-i} \in S_{I \setminus \{i\}}(h') \subseteq S_{I \setminus \{i\}}(h)$,¹⁷

$$\alpha_{i,-i}(s_{-i} | h) = \alpha_{i,-i}(s_{-i} | h') \alpha_{i,-i}(S_{I \setminus \{i\}}(h') | h).$$

In particular, (2) implies that if two distinct information sets $h, h' \in \bar{H}_i$ give the same information about co-players because they only differ due to past actions taken by i , $S_{I_0 \setminus \{i\}}(h) = S_{I_0 \setminus \{i\}}(h')$, then 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).

The following result summarizes the previous analysis of consistency.

Remark 1 *In games with observed deviators and in one-person games against Chance, condition (3) is equivalent to the consistency condition (2).*

Proof: The equivalence between (3) and (2) in games with observed deviators follows from Proposition 3.1 in Battigalli (1996). In one-person game forms against Chance ($I_0 = \{i\} \cup \{0\}$), perfect recall implies that $S_{I_0}(h) = S_i(h) \times S_0(h)$ for all $h \in \bar{H}_i$. Hence, such game forms have observed deviators. ■

2.3 Regret and utility

At terminal information sets, players obtain information about their material payoff 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.

¹⁷That is, h is less informative than h' about personal co-players, e.g. (but not necessarily), because h precedes h' .

Definition 1 The **regret** of player $i \in I$ at $h \in \bar{H}_i$ ($h \subseteq Z$) given beliefs α_i is

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

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

Note that 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 the definition for two-player game forms with observed actions and without chance moves put forward by Battigalli & Dufwenberg (2009).

Definition 2 The **utility** of player $i \in I$ is the function $u_i : Z \times \Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}}) \rightarrow \mathbb{R}$ 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 systems of beliefs $\alpha_i \in \Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}})$, 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 a strictly increasing function. 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 (in Section 7.1) our approach to DRT, 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 f_i and θ_i into a new function. Doing so would, however, obscure the psychological interpretation. Function f_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), 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 et al. 2019 and Battigalli & Dufwenberg 2022).

A game form Γ coupled with utilities $u = (u_i)_{i \in I}$, as given by (6), comprises what we will call a **regret game** $G = (\Gamma, u)$. We say that G is a **linear regret game** if regret is linear for every player $i \in I$. We assume that 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.¹⁸




Figure 2 A three-person game form

Example 2 Consider Fig. 2, with a three-player game form. The payoffs of players 2 and 3 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. What is player 1's utility of Out? Despite that Out ends the game, the answer is not a given number. Rather, it depends on 1's beliefs about the choices of players 2 and 3. For example, if player 1 assigns probability 1 to (\mathbf{B}, \mathbf{A}) then his utility of Out equals $1 - \theta_1 \cdot (2 - 1) = 1 - \theta_1$, while if he assigns probability 1 to (\mathbf{A}, \mathbf{B}) then said utility equals 1. In general,

$$u_1(\text{Out}, \alpha_1) = 1 - \theta_1 (\max \{ \alpha_1(\mathbf{A}, \mathbf{A} | \text{Out}) + \alpha_1(\mathbf{B}, \mathbf{B} | \text{Out}) + 2\alpha_1(\mathbf{B}, \mathbf{A} | \text{Out}), 1 \} - 1),$$

where $\alpha_1(s_2, s_3 | h)$ is the subjective probability assigned to (s_2, s_3) conditional on h .

¹⁸We denote with boldface letters like \mathbf{a}_j conditional statements of the form “if h occurred, action a_j would be chosen,” where $h \in H_j$ and $a_j \in A_j(h)$; i.e., bold letters in examples denote strategies, or parts of strategies.

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:¹⁹




Figure 3 A game form where Chance is a follower

Example 3 Consider Fig. 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 1 chooses s then he would not learn what Chance would have chosen had he selected r . Player 1's belief about Chance's strategy is as follows:

$$\alpha_1(\mathbf{g}|s) = 1 - \alpha_1(\mathbf{b}|s) = \frac{3}{5}.$$
²⁰

Given this belief, player 1's linear regret for choosing s is:

$$r_1(s, \alpha_1) = \left(\frac{3}{5} \cdot 2 + \frac{2}{5} \cdot 0 \right) - 1 = \frac{6}{5} - 1 = \frac{1}{5}.$$

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

$$u_1(s, \alpha_1) = 1 - \frac{\theta_1}{5}$$

¹⁹In this section, we let i denote the only personal player in general formulas and statements. In the examples, $i = 1$.

²⁰Formally, this equation is implied by consistency of beliefs with Chance probabilities (Eq. (3) can be used in this case). We will not repeat similar observations concerning consistency in the following examples.

and the expected utility of choosing r is

$$u_1(r, \alpha_1) = \frac{3}{5} \cdot 2 + \frac{2}{5} \cdot [0 - \theta_1 \cdot (1 - 0)] = \frac{6}{5} - \frac{2}{5}\theta_1.$$

Condition (3) determines beliefs in one-person games with Chance; this explains why α_1 does not appear explicitly on the right-hand side of this expression. Comparing expected utilities one sees that player 1's preferred choice between r or s depends on whether $\theta_1 < 1$ or $\theta_1 > 1$.

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

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

Proof: In one-person games against Chance with essentially simultaneous moves, the set of terminal histories can be written as the set of strategy (action) pairs, $Z = S_i \times S_0$, letting the path function ζ be the identity on $S_i \times S_0$. Perfect feedback implies that \bar{H} is the finest partition of $\{\emptyset\} \cup (S_i \times S_0)$ and terminal information sets $h \subseteq S_i \times S_0$ are singletons $h = \{(s'_i, s_0)\}$, where s'_i is the strategy (action) chosen by i . Thus, according to (4), the linear regret of i given $h = \{(s'_i, s_0)\}$ is simply

$$r_i(\{(s'_i, s_0)\}, \alpha_i) = \max_{s_i \in S_i} m_i(s_i, s_0) - m_i(s'_i, s_0), \quad (7)$$

which is independent of the ex-post beliefs specified by belief system α_i , because there is no ex post uncertainty about the choice of Chance, s_0 . For the same reason, the psychological utility of any pair (terminal node) (s'_i, s_0) may also be expressed as belief-independent:

$$\begin{aligned} u_i(s'_i, s_0, \alpha_i) &= m_i(s'_i, s_0) - \theta_i \left(\max_{s_i \in S_i} m_i(s_i, s_0) - m_i(s'_i, s_0) \right) \\ &= (1 + \theta_i) m_i(s'_i, s_0) - \theta_i \max_{s_i \in S_i} m_i(s_i, s_0). \end{aligned}$$

Therefore, the expected utility maximization problem is

$$\max_{s'_i \in S_i} \sum_{s_0 \in S_0} \alpha_i(s_0 | \emptyset) ((1 + \theta_i) m_i(s'_i, s_0) - \theta_i M_i(s_0)),$$

where $M_i(s_0) = \max_{s_i \in S_i} m_i(s_i, s_0)$. Since $(1 + \theta_i) > 0$ and $M_i(s_0)$ does not depend on the choice variable, this is equivalent to maximizing the expected monetary payoff:

$$\arg \max_{s'_i \in S_i} \sum_{s_0 \in S_0} \alpha_i(s_0 | \emptyset) u_i(s'_i, s_0, \alpha_i) = \arg \max_{s'_i \in S_i} \sum_{s_0 \in S_0} \alpha_i(s_0 | \emptyset) m_i(s'_i, s_0). \blacksquare$$




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

Example 4 To illustrate Proposition 1, consider Fig. 4, where $\lambda \in (0, 1)$ is the probability of choice b . By (6) and (7), player 1 prefers s over r iff

$$\begin{aligned} \lambda \cdot 1 + (1 - \lambda) \cdot (1 - \theta_1(2 - 1)) &\geq \lambda \cdot (0 - \theta_1(1 - 0)) + (1 - \lambda) \cdot 2 \Leftrightarrow \\ (2\lambda - 1)(1 + \theta_1) &\geq 0 \Leftrightarrow \\ \lambda &\geq \frac{1}{2}. \end{aligned}$$

Regret is irrelevant to the prediction, as player 1 maximizes expected payoff independently of θ_1 .

Note that the irrelevance of regret under the conditions of Proposition 1 concern implied behavior, and not that feelings of regret are absent when a game is played. Example 4 illustrates this, since player 1 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 for games that do not feature essentially simultaneous moves. Similarly, regret can matter in the presence of imperfect feedback, as we show next:




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

Example 5 Consider Fig. 5, which looks like Fig. 4 except for the information structure over end-nodes. Assume that $0 < \lambda < \frac{1}{2}$. By (6), player 1 prefers s over r iff

$$\begin{aligned} 1 - \theta_1[\lambda \cdot 0 + (1 - \lambda) \cdot 2 - 1] &\geq \lambda \cdot (-\theta_1) + (1 - \lambda) \cdot 2 \Leftrightarrow \\ \theta_1 &\geq \frac{1 - 2\lambda}{3\lambda - 1}. \end{aligned} \quad (8)$$

Unlike in Example 4, the game form in this example does not feature perfect feedback, as the terminal information set $\bar{H}_1(b, s) = \bar{H}_1(g, s)$ is not a singleton. In the associated linear regret game, for any $\lambda > \frac{1}{3}$ player 1 prefers s to r if θ_1 is high enough. And the higher is θ_1 , the lower λ must be for player 1 to be willing to choose r .

We next discuss five related key insights:

Belief-dependent utility First, in Examples 4 and 5, player 1'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_1 is high enough).²¹

²¹Note also that the effect is analogous to the one we sketched in Example 1, for the choice *take*, helping us get a deeper insight also into that example. In both cases, a belief-dependent preference is responsible for regret to matter.

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, and a comparison of Examples 4 and 5 makes that clear. The game forms used in these two examples differ only in the terminal information structure, and yet the predictions may be different (for high enough θ_1).²²

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




Figure 6 Chance is a follower with perfect information

Example 6 Consider Fig. 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 1, not the action that would have been taken following the forgone action. Assume that $0 < \lambda < \frac{1}{2}$. According to (6), player 1 prefers s over r if and only if (8) holds; decisions depend on θ_1 .

Compare Examples 4 and 6. The involved game forms differ as regards the order of moves, but they share the property that player 1 chooses between s or r without information about Chance's choice. If θ_1 is high enough, this subtle distinction leads to different predictions!

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

²² Again, 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 Fig. 7. Again, the order of moves has been flipped relative to Fig. 4, but this time Chance is a follower with imperfect information. The “interchange of essentially simultaneous moves” transformation applies.²³ As in Example 4, player 1 prefers s over r iff $\lambda \geq \frac{1}{2}$.

Now compare Examples 6 and 7. The difference concerns a form of counterfactual correlation: If player 1 chooses s and learns that Chance chose g , does that imply that Chance would also have chosen the action on the right g had player 1 chosen r ? The answer is no in Fig. 6 and yes in Fig. 7. Note that either figure may make economic sense: Fig. 7 might reflect that the chance move corresponds to an outcome in the stock market, which is independent of whether player 1 invests (i.e., chooses r) or not (chooses s). Fig. 6 could instead describe a situation where r corresponds to player 1 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 1 will never know and it is independent of his coin flip. For many theories of decision making, player 1 would make the same choice in Figs. 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 1 would know that he would have made a lot of money had he instead chosen r . Hence, following (s, g) , player 1 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_1) different.

²³Such transformations are defined and analyzed, for the class of game forms we consider, in Battigalli et al. (2020). See the references therein on the literature on transformations of games.

Risk aversion Fifth, the effect of increasing θ_1 , 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 1 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”: $v_1(x) = x^{\frac{1}{\gamma}}$ with $\gamma > 1$. We can here note that “the higher is risk aversion parameter γ , the lower λ must be for player 1 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. In Example 5, lack of information after the safe action shelters player 1 from much of the regret that would plague him if he knew whether Chance had chosen s or r .

4 Strategic interaction

To model behavior in general game forms (with $|I| \geq 1$) we start by analyzing 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.²⁴ Our 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 et al. (2019).

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 $h \in H_i$, let

$$H_i^{|h} := \{h' \in H_i | h' \not\prec h\}$$

denote the collection of nonterminal information sets of i that do not precede h .²⁵ For any $h \in H_i$ 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 $H_i^{|h}$ and selects i 's action leading to h at each predecessor $h' \in H_i \setminus H_i^{|h}$ (if any).²⁶ With this, for any consistent system of beliefs $\alpha_i \in \Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}})$ we consider the set of (pure)

²⁴For 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 mute for regret games.

²⁵Recall that $\{\emptyset\} \in \bar{H}_i$, but $\{\emptyset\} \in H_i$ if and only if i is active at the root. In this case, $H_i^{\{\emptyset\}} = H_i$.

²⁶Such actions are well defined by perfect recall.

sequential best replies

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

where

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

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 psychological utility functions considered here satisfy own-plan independence, they represent dynamically consistent preferences. Finiteness of the game form and standard dynamic programming results imply that sequential best replies always exist:²⁷

Remark 2 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” $S_{I_0 \setminus \{i\}}(h)$ ($h \in H_i$), $BR_i(\alpha_i) = \arg \max_{s_i \in S_i} \mathbb{E}_{s_i, \alpha_i}[u_i]$ coincides with the set of strategies that maximize the 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 $h \in H_i$:

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

The analysis of best replies is similar to the analysis of expected utility maximization in one-person games, with Chance replaced by the whole 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,

²⁷Intuitively, one can construct a strategy s_i by “folding-back planning” on the tree of non-terminal information sets H_i . Consistency of α_i implies that such s_i is a sequential best reply to α_i . See, e.g., Battigalli *et al.* (2023).

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 this question 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 the (single-player) Examples 3, 5, and 6. The optimal choices of those decision problems 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.

Proposition 2 *Consider a game form Γ with (i) essentially simultaneous moves, and (ii) perfect feedback. Then, 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.*

Proof: Since the game has essentially simultaneous moves, strategies coincide with actions, and

$$BR_i(\alpha_i) = \arg \max_{a_i \in A_i(\emptyset)} \mathbb{E}_{a_i, \alpha_i}[u_i].$$

Such maximizers are the same as in the fictitious one-person game where i is the only personal player and Chance plays each action profile (a_{-i}, a_0) with probability $\alpha_i(a_{-i}, a_0)$. By Proposition 1, the set of maximizers is independent of θ_i . The thesis follows. ■

Another form of regret-irrelevance arises in games *without Chance* under the general regret function (4) when a player is “certain” about others’ strategies whenever he has to move. Formally, player i ’s belief system α_i features **deterministic beliefs when active** if, for each nonterminal information set $h \in H_i$, player 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 $h \in H_i$, there is some $s_{-i}^* \in S_{I \setminus \{i\}}(h)$ such that $\alpha_{i, -i}(s_{-i}^* | h) = 1$.*

Proposition 3 *Consider a game form Γ without chance moves (i.e., $X_0 = \emptyset$). Fix any $i \in I$ and a consistent belief system α_i that satisfies deterministic beliefs when active; then, the set of sequential best replies $BR_i^{\theta_i}(\alpha_i)$ is independent of the sensitivity parameter θ_i and coincides with the set of strategies that maximize expected material payoff at each information set $h \in H_i$.*

The formal proof of this result is in the Appendix. The intuition 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.

Finally, 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.²⁸

Lemma 1 *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 $(\hat{\theta}_i, \hat{\alpha}_i, \hat{s}_i) \in \mathbb{R}_+ \times \Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}}) \times S_i$ and every sequence $(\theta_i^n, \alpha_i^n, s_i^n)_{n \in \mathbb{N}} \in (\mathbb{R}_+ \times \Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}}) \times S_i)^{\mathbb{N}}$ 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 , $\hat{s}_i \in BR_i^{\hat{\theta}_i}(\hat{\alpha}_i)$.*

4.2 Sequential equilibrium

As with traditional games, sometimes pure strategy equilibria do not exist; thus, we consider equilibria in randomized strategies. For any $h \in H_i$, let $\pi_i(\cdot|h) \in \Delta(A_i(h))$ denote a probability distribution on the feasible actions of i . A **behavior strategy** for i is a probability vector $\pi_i = (\pi_i(\cdot|h))_{h \in H_i}$. 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 .²⁹ We say that π is **strictly randomized** if $\pi_i(a_i|h) > 0$ for all $i, h \in H_i$ 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.³⁰ For any player $i \in I$, pure strategy $s_i \in S_i$, and behavior strategy π_i , let

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

²⁸We endow the set $\Delta^{\bar{H}_i}(S_{I_0 \setminus \{i\}}) \subseteq \mathbb{R}^{\bar{H}_i \times S_{I_0 \setminus \{i\}}}$ of consistent systems of beliefs with the topology induced by the Euclidean metric of $\mathbb{R}^{\bar{H}_i \times S_{I_0 \setminus \{i\}}}$, and S_i with the discrete topology. For completeness, we include a proof in the Appendix.

²⁹An 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).

³⁰Equivalent to the one put forward by Kreps & Wilson (1982), provided that we also consider terminal information sets.

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 $i \in I$, $h \in \bar{H}_i$, and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)$,

$$\alpha_i(s_{-i}, s_0 | h) = \lim_{k \rightarrow \infty} \frac{\sigma_0(s_0) \prod_{j \in I \setminus \{i\}} \mathbb{P}_{\pi_j^k}(s_j)}{\sum_{(s'_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \sigma_0(s'_0) \prod_{j \in I \setminus \{i\}} \mathbb{P}_{\pi_j^k}(s'_j)}. \quad (9)$$

Full consistency requires the belief system of each player to be coherent with the behavior strategies of the co-players; furthermore, it strengthens consistency by 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

$$\alpha_i(s_{-i}, s_0 | h) = \frac{\sigma_0(s_0)}{\sum_{s'_0 \in S_0(h)} \sigma_0(s'_0)} \prod_{j \in I \setminus \{i\}} \lim_{k \rightarrow \infty} \frac{\mathbb{P}_{\pi_j^k}(s_j)}{\sum_{s'_j \in S_j(h)} \mathbb{P}_{\pi_j^k}(s'_j)} \quad (10)$$

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

$$\alpha_i(s_{-i}, s_0 | x) = \prod_{j \in I_0 \setminus \{i\}} \mathbb{P}_{\pi_j}(s_j | x) = \prod_{j \in I_0 \setminus \{i\}} \prod_{x' \in X_j^{|x|}} \pi_j(s_j(x') | x'). \quad (11)$$

for all $i \in I$, $x \in X_i$, and $(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(x)$. 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 $i \in I$, $h \in H_i$, $s_i \in S_i$,

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

It can be checked that the sequential rationality condition is equivalent to requiring that, for each $h \in H_i$,

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

where

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

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 been observed, because such deviations are interpreted as mere mistakes in carrying out co-players' plans, and such mistakes are not expected to occur in the future, just as in Selten's (1975) "trembling-hand" perfect equilibrium (cf. Kreps & Wilson 1982, Proposition 6).

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

Proof: 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, Proposition 2 implies that the equilibrium set is linear-regret-invariant in essentially simultaneous game forms with perfect feedback:

Corollary 1 *Consider a game form Γ with (i) essentially simultaneous moves, and (ii) perfect feedback. Let G^θ be the associated linear regret game where $\theta = (\theta_i)_{i \in I}$. The set of SEs in G^θ is invariant with respect to θ . Moreover, each player behaves as if he maximized expected material payoff.*

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

moves:




Figure 8 A parameterized three-person game form

Example 8 Consider Fig. 8, an extension of Fig. 2, which includes players 2 and 3's "Matching-Pennies" payoffs in the right sub-game and – for player 1 – mirrors the game form in Fig. 2 when $x = 0$. Assume linear regret, and that each player has perfect information across end-nodes (i.e., perfect feedback); thus, relative to Proposition 2, only the simultaneous-move condition (i) is relaxed. Assume that $x \in (0, 1)$ and that $\theta_1 \geq \theta_2 = \theta_3 = 0$. In any SE, 2 and 3 assign probability $\frac{1}{2}$ to each choice of the Matching-Pennies subgame. To see what 1 does, note that his utility from Out equals $1 - \theta_1 \cdot \frac{x}{4}$ while his utility from In equals $1 + \frac{x}{4} - \frac{1}{4} \cdot \theta_1 \cdot 1$, so 1 prefers Out to In iff $\theta_1 \geq \frac{x}{1-x}$. That is, 1's equilibrium choice is not invariant w.r.t. θ_1 .

The next example shows that condition (ii) of Corollary 1 (perfect feedback) cannot be relaxed.




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

Example 9 Consider the game form in Fig. 9 and assume linear regret. If information across end-nodes were perfect then the situation would be covered by Corollary 1, and we would have regret-sensitivity invariance. However, this is not the case in Fig. 9 where $\{(U, L), (D, L)\}$ is a doubleton information set for player 2. Assume that $\theta_2 \geq \theta_1 = 0$. It can be checked that there is unique SE in which 2 assigns probability $\frac{1}{2}$ to each of his choices (indifference condition for 1), while 1 chooses U with θ_2 -dependent probability $\frac{1+2\theta_2}{2+3\theta_2}$ (indifference condition for 2).

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 2 Consider a game form Γ without chance moves (i.e., $X_0 = \emptyset$). Let G^θ be a regret game with $\theta = (\theta_i)_{i \in I}$, where no additional restrictions are imposed on $(f_i)_{i \in I}$. 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.³¹

³¹Conversely, 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.

Formally, 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 $h \in H_i$, $\pi_i(s_i^*(h)|h) = 1$.*

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 (11) (full consistency in games with observed actions) we obtain the following invariance result.

Proposition 5 *Consider a game form Γ without chance moves ($X_0 = \emptyset$) 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 $(f_i)_{i \in I}$. 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.*

Proof: By the observed-actions assumption, for each player $i \in I$ we can replace information sets in H_i with decision nodes in X_i . Let s_i^* denote the equilibrium pure strategy of i as per Condition 2. By characterization (11) of full consistency in games with observed actions, it follows that, for all personal co-players $j \in I \setminus \{i\}$ and $x \in X_j$, we have $\alpha_{j,i}(s_i^*|x) = 1$. Since there are no chance moves, Condition 1 is satisfied, that is, any sequential equilibrium in pure strategies satisfies deterministic beliefs. Thus, the result follows from Proposition 3. ■

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 action.




Figure 10 A game form with imperfectly observed actions

Example 10 Consider the game form in Fig. 10 with $x \in (0, 1)$ and $y > 2$. Assume linear regret and fix $\theta_1 = 0$. 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_2 but not for small θ_2 . 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_1, s_2, λ) . Since $y > 2$, c is conditionally dominated by a 50/50 mixture on b and d when $\theta_2 = 0$. 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 in mixed strategies—when θ_2 is zero or very small. The parameterized utilities of the terminal nodes for player 2 are given by:

$$\begin{aligned} u_2((\ell, b), \lambda) &= y - \theta_2 \cdot 0 &= u_2((r, d), \lambda), \\ u_2((\ell, c), \lambda) &= 1 - \theta_2 \cdot [\max\{\lambda y, (1 - \lambda)y\} - 1] &= u_2((r, c), \lambda), \\ u_2((\ell, d), \lambda) &= 0 - \theta_2 \cdot y &= u_2((r, b), \lambda). \end{aligned}$$

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

$$\begin{aligned} \bar{u}_{2,h}\left(b, \lambda = \frac{1}{2}\right) &= \frac{1}{2}y - \frac{1}{2}\theta_2 y = \bar{u}_{2,h}\left(d, \lambda = \frac{1}{2}\right), \\ \bar{u}_{2,h}\left(c, \lambda = \frac{1}{2}\right) &= 1 - \theta_2\left(\frac{1}{2}y - 1\right). \end{aligned}$$

The pair of parameters (θ_2, y) makes $(a, c, \lambda = \frac{1}{2})$ an SE if and only if:

$$\begin{aligned} 1 - \theta_2 \left(\frac{1}{2}y - 1 \right) &\geq \frac{1}{2}y - \frac{1}{2}\theta_2 y \\ \theta_2 &\geq \frac{1}{2}y - 1 > 0. \end{aligned}$$

Note that this game form features observed deviators. Thus, this example shows that observed deviators are insufficient for the invariance result of Corollary 2; 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). Such a suspicion might be strengthened by recent results of Weinstein (2016) and Battigalli et al. (2015, 2016). These authors show that if one fixes a game form with monetary outcomes and then explores the effect of increased risk aversion, and (in Battigalli et al.'s case) also ambiguity aversion, then the prediction set (e.g., various forms of equilibria, or rationalizable choices) expands. However, it only takes a brief reflection to realize that, with regret, those monotonicity results have no analog. Example 3 provides a counterexample. The shift, on increasing θ_1 , 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 and let G^θ be the associated regret game where $\theta = (\theta_i)_{i \in I}$. Let $SE(G^\theta)$ be the associated set of SEs. Let G^0 denote the game where $\theta_i = 0$ for all i , and let $\bar{\theta} = \max_{i \in I} \theta_i$. With this, $\lim_{\bar{\theta} \rightarrow 0} SE(G^{\bar{\theta}}) \subseteq SE(G^0)$. Moreover, the inclusion may be strict.*

Intuitively, the result follows from the upper-hemicontinuity of the equilibrium correspondence $\theta \mapsto SE(G^\theta)$, which follows from Lemma 1. Sharper predictions for small regret aversion come from the lack of lower-hemicontinuity at 0. We prove the general weak inclusion in the Appendix. To see the the inclusion may be strict, consider Example 8 and assume $x = 0$. Then, the path in which player 1 chooses In is an SE outcome of G^0 . (Indeed, every behavior strategy profile in which players 2 and 3 randomize uniformly is part of an SE assessment of G^0 , since uniform randomization by players 2 and 3 renders player 1 indifferent.) However, for any sequence θ^n with $\theta_1^n > 0$, no path that includes In can be an SE outcome of G^{θ^n} .³²

³²A perspective one may take on Proposition 6 is that regret could serve as a refinement device. Think

5 Information avoidance

In this section we analyze how anticipated regret avoidance affects different forms of information acquisition, pushing in the direction of information avoidance. First we show that strategies that yield more information *ex post* — other things being equal — have lower expected psychological utility because they entail higher anticipated regret. Next we show by example that informative signals that are materially valuable may be avoided because their acquisition would entail higher anticipated regret.

We begin by defining a relation for comparing strategies based on the information they reveal about co-players' strategy profiles at the end of the game. For each $z \in Z$, $S_{I_0 \setminus \{i\}}(\bar{H}_i(z))$ denotes the set of co-players' strategy profiles that i cannot rule out given his terminal information at z . With this, the overall information that i can infer about $S_{I_0 \setminus \{i\}}$ by playing s_i is represented by collection

$$\mathcal{P}_i(s_i) := \{E_{I_0 \setminus \{i\}} \in 2^{S_{I_0 \setminus \{i\}}} : \exists z \in Z_i(s_i), E_{I_0 \setminus \{i\}} = S_{I_0 \setminus \{i\}}(\bar{H}_i(z))\}.$$

It is worth noting that any two realization-equivalent strategies yield the same *ex post* information:

Remark 3 *Fix any pair of strategies $s'_i, s''_i \in S_i$; if $\zeta(s'_i, s_{-i}) = \zeta(s''_i, s_{-i})$ for all $s_{-i} \in S_{-i}$, then $\mathcal{P}_i(s'_i) = \mathcal{P}_i(s''_i)$.*

This implies that, when comparing the informativeness of strategies, we can focus on the corresponding **reduced strategies**, i.e., their realization-equivalence classes.³³ Furthermore, we prove in the Appendix that, for each s_i , $\mathcal{P}_i(s_i)$ is a *partition* of $S_{I_0 \setminus \{i\}}$.³⁴ Therefore, we can compare the informativeness of two strategies by means of the canonical partial order over partitions:

Definition 5 *Fix $s'_i, s''_i \in S_i$, we say that s'_i is **as least as ex post informative as** s''_i if $\mathcal{P}_i(s'_i)$ is finer than $\mathcal{P}_i(s''_i)$, that is, for every $E''_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s''_i)$, there exists $E'_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s'_i)$ such that $E'_{I_0 \setminus \{i\}} \subseteq E''_{I_0 \setminus \{i\}}$.*

of the comparison sometimes made between Nash equilibrium (NE) and perfect equilibrium (PE). NE does not consider the possibility of mistakes; PE involves perturbed games where players err, with a focus on the limit as errors vanish. PE turns out to refine NE. The parallel in our case would be a reference to vanishingly small concerns for regret rather than mistakes. The construction refines the set of SEs relative to when regret concerns are absent to start with. We find this perspective intriguing as the refinement-use of a slight degree of a motivational concern seems unusual (but see Wu & Jiang (1962) for something remotely in this spirit). It is, however, conceptually contrived, because if regret is an important emotion, then the main case of interest should involve $\theta_i > 0$ rather than the case where $\theta_i \rightarrow 0$.

³³See the Theorem 1 in Kuhn (1953) and the Appendix.

³⁴See Lemma 3.

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 ex post informative as any s''_i with $s''_i(\emptyset) = \text{stop}$. Indeed, s''_i induces the trivial partition $S_{I_0 \setminus \{i\}}$, $\mathcal{P}_i(s''_i) = \{S_{I_0 \setminus \{i\}}\}$, which is necessarily (weakly) coarser than $\mathcal{P}_i(s'_i)$.

Proposition 7 *Fix a nonterminal information set $h \in H_i$, consider strategies $s'_i, s''_i \in S_i$ 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, the result follows from the fact that finer ex post information allows a more fine-tuned evaluation of hindsight-maximized expected material payoffs. This increases anticipated regret. The following example illustrates this result in the interesting case in which s'_i and s''_i yield the same expected material payoff, but the more ex post informative strategy has strictly lower expected psychological utility.




Figure 11 A one-player game form with chance

Example 11 *In Fig. 11, player 1 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 coordinate represents whether stock trading yields a gain and the second whether betting yields a gain. Both stock trading and betting lose with probability $\lambda \in (0, 1)$. Player 1 always observes the stock outcome, regardless of his action, whereas he observes the bet outcome only if he bets.*

In this game form, trading a stock and placing a bet yield the same expected material payoff. However, stock-trading is less ex post informative than betting. The anticipated regret from trading a stock is

$$\lambda(\max\{2(1 - \lambda), 1\} - 0) = \max\{2\lambda - 2\lambda^2, \lambda\},$$

and the anticipated regret from placing a bet is

$$2(1 - \lambda)\lambda + \lambda^2.$$

For any $\lambda \in (0, 1)$, trading a stock yields strictly higher anticipated regret than placing a bet and therefore strictly lower expected utility.

We can show that the information-avoidance result extends beyond linear regret under stronger assumptions about the material-payoff implications of strategies being compared. The following example shows that informative and materially valuable signals may be avoided, which implies 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 Fig. 12, the player can choose to bet or not. Betting yields a payoff of 1 if successful and -1 if unsuccessful, while not betting always 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. Suppose $\theta_1 \in (0.5, 1)$, a range of interest because sensitivity is strong enough to affect information acquisition, but not so strong 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

$$0.45 + 0.55 \cdot (-1 - \theta_1 \cdot 1) < 0.$$

With information, his subsequent decision depends on the realized signal. If it is high, $\omega = H$, he will bet, since

$$0.6 + 0.4 \cdot (-1 - \theta_1 \cdot 1) > 0 - \theta_1 \cdot (0.6 - 0.4),$$

provided that $\theta_1 < 1$. If it is low, $\omega = L$, he will not bet, since

$$0.3 + 0.7 \cdot (-1 - \theta_1 \cdot 1) < 0.$$

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 1 will not acquire information. Once informed, he bets if the signal is high, which raises the level of anticipated regret. Since $\theta_1 > 0.5$ implies

$$0 > \frac{1}{2} \cdot 0.6 + 0.4 \cdot (-1 - \theta_1 \cdot 1) + \frac{1}{2} \cdot 0,$$

the player strictly prefers avoiding information acquisition without betting (reduced strategy $r.n$) to acquiring information and betting iff the signal is high (reduced strategy $a.b^H.n^L$). Indeed, the latter (reduced) strategy yields a strictly higher expected material payoff, the former is strictly less *ex post* informative ($\mathcal{P}_i(r.n) = \{S_0\}$ and $\mathcal{P}_i(a.b^H.n^L) = \{\{H.w\}, \{H.l\}, \{L.w, L.l\}\}$) and implies a strictly lower anticipated regret.

6 Suggestions for 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 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. We apply the linear-regret version of our model to their game form (Fig. 13(i)). Some novel and striking conclusions emerge:




Figure 13 Climate-action game forms

Player 1, who is unsure whether the scientists or the skeptics are right, must “price carbon” (PC) or “do nothing” (DN). The chance move in Fig. 13(i) reflects 1’s state of mind; he assigns probability λ to the scientists being right. VdP&R assume that $0 < c < a < b$: If Player 1 chooses DN and the skeptics are right, then 1’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).³⁵

³⁵ VdP&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 Fig. 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,

If Player 1 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) 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 Fig. 13(i) is covered by Propositions 1 & 2, so incorporating regret changes no conclusion for that game form.

However, arguably Fig. 13(i) is not the “right” game form in which to conduct the analysis. Fig. 13(i) 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 1 may not be able to distinguish the terminal nodes where he gets $-c$ and $-a$; if so, this moves us from Fig. 13(i) to Fig. 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 Fig. 13(ii), the outcomes where Player 1 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 1) differ, but because of unmodeled background noise, 1 cannot tell.³⁶

If the game form in Fig. 13(ii) is relevant, then the impact of regret is felt. Namely, the range of parameters (λ, a, b, c) for which Player 1 will choose PC expands (with θ_1). To verify this, note that the logic of the solution of Player 1’s decision problem in Fig. 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 1 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 1 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 Fig. 13(ii) to Fig. 13(iii). It is straightforward to now verify that, again, regret will not influence the analysis. That is, even if Player 1 cares about the welfare of his offspring just

$b - a < c$. 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 of changing θ_1 rather than the frequency of Players 1 who might choose PC.

³⁶A similarly motivated example of non-observable payoffs appears in Dufwenberg & Nordblom’s (2022, Fig. 1, p. 9) analysis of tax evasion.

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 observation further: Is it really likely that Player 1 will be dead once the outcome associated with $-b$ is revealed? This may depend on how much longer Player 1 has to live when DN is chosen, and his related beliefs may affect his choice. It may be that for (relatively more) young people the relevant game form is Fig. 13(ii), while for (relatively more) older people it is Fig. 13(iii). We would then expect regret to influence decisions differently depending on which generation is considered.

6.2 Market entry

Compare the two game forms in Figs. 14(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 firms enters the market then the entrant(s) get access to 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 14 Two market-entry game forms

Suppose the CEOs of both firms are regret-neutral, $\theta_A = \theta_B = 0$. For each game form, the associated game has a unique mixed equilibrium such that each firm will enter with a probability of $2(1 - c)$.

If instead the CEOs are regret-averse, $\theta_A, \theta_B > 0$, then the regret game corresponding to each game form still has a unique 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 1): each firm will enter with probability $2(1 - c)$. By contrast, in (ii), firm $i \in \{A, B\}$ will enter with probability

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

where θ_j is the *other* firm's regret sensitivity. Note that q_i is decreasing in θ_j .

Scholars in IO have taken an interest in how (behavioral) CEO's 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 some guidance as regards whether this is the case. Namely, we know by Corollary 1 that regret can not matter in the regret games associated with the game form in (i).

6.3 Delegation

Our next example again concerns two similar game forms, but this time neither is covered by any of our regret-irrelevance results. Nevertheless, regret subtly influences play differently in the two cases. The 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 Fig. 15. 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) – L or R – simultaneously, with the (material) payoffs as indicated in the figure.




Figure 15 Delegation vs. DIY game form

Solving for SE, it is clear that each player will shoot left or right with equal probability, independently of any regret sensitivities. Now study 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 regrets. 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 scoring of 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 is indicative that one could develop an extension where A is more skilled than B yet still delegates the task to B .

6.4 Auctions and mechanism design

A large literature has argued that anticipated regret helps explain bidding in auctions. See, e.g., Engelbrecht-Wiggans (1989), Engelbrecht-Wiggans & Katok (2008), Filiz-Ozbay & Ozbay (2007), and Bergemann et al. (2025). 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 lesser 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. We propose that it may be fruitful to use our theory to explore auctions. While we shall not do so in this paper, we offer two related comments:

First, our theory will not allow that winner regret is simply assumed away when switching from game forms with to those without disclosed losing bids. Rather, for purposes of computing regret, winners would consult their ex post beliefs. Yet, the effect may be similar.

Second, auctions may be viewed as 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.

7 Comparison with related work

7.1 Classical 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, while extensions to larger choice sets were proposed subsequently. We relate to the following version, proposed by Quiggin (1994):³⁷

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

$$U(a, \omega) = V(a(\omega)) - f\left(\max_{a' \in A} V(a'(\omega)) - V(a(\omega))\right), \quad (12)$$

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 and then show that this extension encompasses the DRT framework described above. Specif-

³⁷Quiggin shows that his approach is implied by "irrelevance of statewise dominated alternatives," a quite intuitive axiom. See also Loomes & Sugden (1987).

ically, we replace $m_i(z)$ and $\bar{m}_i(z)$ in (6) and (4) with $v_i(m_i(z))$ and $v_i(\bar{m}_i(z))$, 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. Consequently, i 's extended utility function becomes:

$$\tilde{u}_i(z, \alpha_i) = (v_i \circ m_i)(z) - \theta_i \cdot \tilde{r}_i(\cdot, \cdot); \quad (13)$$

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

$$\tilde{r}_i(h, \alpha_i) := f_i \left(\max_{s_i \in S_i} \left\{ \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | h) (v_i \circ m_i)(\zeta(s_i, s_{-i}, s_0)) \right\} - v_i(\bar{m}_i(h, \alpha_i)) \right). \quad (14)$$

Function $f_i : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is continuous, strictly increasing and satisfies $f_i(0) = 0$.³⁸

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 (Γ, \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, 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 an arbitrary act $a \in A$ in the decision problem. Construct a strategy s_i and a 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 the 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 now identify a specific \tilde{u}_1 that takes the same value as U at each end node (ω, a) :

Proposition 8 *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 $x \in \mathbb{R}_+$, $\theta_i \cdot f_i(x) = f(x)$,*

³⁸ An alternative approach, adopted in some decision theory papers on regret (e.g., Bikhchandani & Segal, 2014), assumes that regret arises from comparing the realized choiceless utility with the certainty equivalent (CE) of the forgone lottery.

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

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

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

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

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

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

where the second equality follows from Condition (ii). ■

A basic DRT-insight—compare, e.g., Loomes & Sugden (p. 810)—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 .³⁹

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

Corollary 3 *Suppose that (Γ, \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 $v_i \circ m_i$.*

Proof: This result follows in light of Proposition 1, relying crucially on the linearity of regret, i.e., f_1 being the identity function. The argument proceeds analogously, with m_i replaced by $(v_i \circ m_i)$. ■

In (Γ, \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.

³⁹For this reason, DRT-scholars have focused on the implications of f being concave or convex.

7.2 Sarver's model of menu choice

Sarver employs an axiomatic approach to model preferences over menus. He characterizes the player's preference with a utility representation as if the player chooses a lottery from a menu and subsequently experiences regret if this lottery turns out to be ex post suboptimal. To describe Sarver's model precisely, let p and q represents lotteries available to the player, and let $A \in \mathcal{A}$ denote a menu that containing at least one lottery. Additionally, μ represents a probability distribution over a set of possible ex post utility functions, \mathcal{U} . Sarver's utility representation, which characterizes the player's *preference over menus*, is given by:

$$U(A) = \max_{p \in A} \int_{\mathcal{U}} \left[u(p) - K \cdot [\max_{q \in A} u(q) - u(p)] \right] \mu(du), \quad (15)$$

where $K \geq 0$. Central to his model is the concept that regret stems not from the choice among menus but from the choice of lotteries within a menu, A . This represents a departure from the decision regret theory (Loomes & Sugden 1982; Quiggin 1994). Sarver justifies the unique perspective of his work by suggesting that the player may be uncertain about his future tastes when choosing a menu. Based on this assumption, a player exhibits a weak preference for a smaller menu over an extended one, with the latter comprising the entirety of the smaller menu plus an additional suboptimal lottery. The underlying logic is that such a suboptimal lottery, despite not being chosen, can potentially lead to greater regret in certain states. This possibility amplifies the player's anticipated regret when faced with the extended menu, thereby diminishing its attractiveness.

To investigate the menu choice problem through a game-theoretic approach, we now describe the problem using a specific class of game forms that mirrors the setup posited by Sarver (2008). This game form involves a Chance player, 0, who initially chooses a state of nature from the set Ω . Player i (DM), not knowing the specific state $s_0 \in \Omega$ chosen, then selects a menu, which could contain a single lottery or multiple lotteries, with his information set being Ω itself. If the chosen menu contains multiple lotteries, Player i then makes a choice among these lotteries. Finally, the state of nature is realized.




Figure 16 Menu-choice game form

Example 13 In Fig. 16,⁴⁰ we illustrate an instance of such a game form. Player 0, first chooses a state of nature, s_0 , from $S_0 = \{\omega_1, \omega_2\}$. Player 1 subsequently chooses an action given his information set $H_1(\omega_1) = H_1(\omega_2) = \{\omega_1, \omega_2\}$. The set of available actions (i.e., menus) is $A_1(\{\omega_1, \omega_2\}) = \{\{a\}, \{b\}, \{a, b\}, \{c\}\}$. When $\{a, b\}$ is chosen, player 1 continues to choose between the actions a and b within that menu. At each terminal node, player 1 receives a material payoff.

By analogy with Corollary 1, when the regret term in our model is linear, a player's choice of options within a menu aligns with maximizing expected material payoffs, suggesting that regret has no impact on behavior. This results in indifference between a smaller menu and an extended menu that includes an additional suboptimal lottery, diverging from Sarver's hypothesis of a weak preference. This divergence stems from our different assumption that, in contrast to Sarver's model, regret in our framework can also be generated by the menu selection itself. By subtly adjusting our model, it becomes effectively equivalent to Sarver's framework.

Recall that, as detailed in Eq. (4) we capture regret as the material payoff gap between what a player actually receives and what he believes he could have received had he made a different choice. In this subsection, we now tweak it by instead positing that the player ex post compares his actual payoff to the potential payoff had he chosen a strategy within a restricted subset, $\mathcal{L}_i(s_i) \subseteq S_i$, which depends on his chosen strategy, s_i . This modification

⁴⁰We maintain the convention that in specific examples of games with one personal player i (DM), we have $i = 1$.

incorporates the player’s limited attention while ruminating. Thus, the modified regret term is given by:

$$\tilde{r}_i(h, \alpha_i) := R \left(\max_{s_i \in \mathcal{L}_i(s_i)} \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(h)} \alpha_i(s_{-i}, s_0 | h) m_i(\zeta(s_i, s_{-i}, s_0)) - \bar{m}_i(h, \alpha_i) \right). \quad (16)$$

With this modified regret function, the choice of s_i not only directly determines the expected utility through the potential outcome but also indirectly influences the regret magnitude via the set $\mathcal{L}_i(s_i)$, allowing the player to partially regulate his feelings of regret.

We now shift our focus back to the previously discussed game forms, in which player i is the unique individual player (decision maker). While the function $\mathcal{L}_i : S_i \rightarrow 2^{S_i}$ can take a wide range of forms, we specify that $\mathcal{L}_i(s_i)$ represents the menu chosen by player i .

Assumption 1 *It is assumed that $\mathcal{L}_i(s_i) = S_i(s_i(\Omega))$, where $S_i(s_i(\Omega))$ is the set of strategies consistent with the initial action, namely the menu, selected by player i ’s strategy, s_i , upon arriving at the information set Ω .*

With Assumption 1 in place, we derive the modified utility function for player i as follows:

$$\tilde{u}_i(z, \alpha_i) = m_i(z) - \theta_i \tilde{r}_i(H_i(z), \alpha_i).$$

By applying this framework to the game forms representing the menu choice problems, one can easily verify that the expected utility, $\mathbb{E}_{\alpha_i}[\tilde{u}_i]$ is essentially equivalent to Eq. (15).

7.3 Other Work

This section will discuss three studies: two by Bikhchandani and Segal (2011, 2014) and one by Cerrone, Feri, and Neary (2025).

7.3.1 Bikhchandani & Segal (2011)

The study conducted by Bikhchandani & Segal (hereafter BS) in 2011 focuses on a class of preferences consistent with Loomes & Sugden’s (1987) generalized regret theory. This approach concentrates on preferences over pairwise lotteries, which distinguishes itself from Quiggin’s (1994) study that addresses the choice among multiple options. Notably, BS assume a perfect correlation between lotteries—where, in our terminology, Chance is active in a single information set—and that the player will learn ex post the outcomes of both lotteries, including those not chosen, an information structure that we term “perfect feedback”.

The key result of their study is that, within the class of regret-based preferences, only the expected utility model features transitivity.

However, the focus of BS’s study and ours diverge. BS investigate a unique class of utility functions that features the transitivity within the domain of Loomes-Sugden regret-based preferences. By contrast, our model examines how players with specific utility functions behave in given game forms.

7.3.2 Bikhchandani & Segal (2014)

Bikhchandani & Segal (2014) focus on the class of regret-based preferences, which they term as “distribution-regret-based preferences.” This class of preferences applies to scenarios where players face statistically independent lotteries. In such cases, independence implies that the realized outcome of one lottery provides no information about the outcome of another unchosen lottery. Consequently, a player’s regret depends on the comparison between the realized outcome and the ex ante distribution of the alternative lottery. Our game form in Example 6, which features imperfect feedback, also reflects this setting.

As in their 2011 paper, BS (2014) investigate the conditions under which the transitivity property is preserved. They highlight two conditions: “betweenness” and a novel concept they call “consistency.” Consistency refers to the condition that indifference curves maintain the same shape. BS show that if either condition is satisfied, the preference can be classified as a member of the distribution-regret-based family. Notably, the consistency condition is satisfied by both CARA and CRRA preferences, making distribution regret compatible with a wide range of non-expected utility models.

The focus of BS’s 2014 paper and ours is different as well. BS (2014) examine a specific class of utility functions that preserve transitivity within distribution-regret-based preferences. Our model, by contrast, examines how players behave in specific game forms with given utility functions.

7.3.3 Cerrone, Feri & Neary (2025)

Cerrone, Feri, and Neary investigate anticipated regret in strategic interaction. Unlike our theory, which develops a general model with sequential moves, they study a class of simultaneous-move games and test relevant hypotheses by experiments. In their setup, each player chooses between s , which yields a fixed payoff of 1, 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 all players.




Figure 17 Two-player gamble game form

This innovative setting creates an environment where one player's choice affects the other through anticipated regret. For regret-averse players, the likelihood of experiencing regret from a foregone ℓ option rises when the co-player also chooses ℓ . Thus, players prefer to make the same choices: both- s or both- ℓ . However, their model of regret differs from ours. In line with Zeelenberg et al. (1996), they assume that a player who receives no feedback about a foregone choice does not experience regret. Despite this difference, our model predicts a similar outcome in their specific game form.

Using our model, consider the associated linear regret game with $\theta_1 = \theta_2 = \theta > 0$. Let $q_i := \alpha_i(\ell)$ denote the probability that player i assigns to his player j choosing ℓ . Player i prefers to choose ℓ if

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

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

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) . Through this analysis, our model yields predictions consistent with those of Cerrone et al. (2025).

7.3.4 Minimax Regret

Another relevant concept is *minimax regret*, originally proposed by Savage (1951), and later axiomatized (Hayashi 2008; Stoye 2011). Minmax 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).

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 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:

$$\min_{\sigma_i \in \Sigma_i} \sum_{s_i \in S_i} \sigma_i(s_i) \cdot \left[\max_{s_0 \in S_0} \left\{ \max_{s'_i \in S_i} m_i(s'_i, s_0) - m_i(s_i, s_0) \right\} \right]. \quad (17)$$

Central to this principle is the focus on worst-case regret, which depends on the action being evaluated. As shown in Eq. (17), 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 $(\alpha_i(\cdot|\bar{h}))_{\bar{h} \subseteq Z, \bar{h} \in \bar{H}_i}$.

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 Eq. (17). Accordingly, the setup can be represented as a single-player game against Chance with essentially simultaneous moves and perfect feedback, such as the game form shown in Example 4. In that example, player 1, who is assumed to have linear regret, prefers action s over r iff the probability of state b , λ , is at least $\frac{1}{2}$. However, under the minimax regret principle, the worst-case regrets associated with both s and r are equal to 1, so any mixed strategy minimizes worst-case regret and thus satisfies the principle — 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 inspirational 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.” Reading these quotes, while keeping in mind the exercise we have conducted in this paper, several reflections come to mind. 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. 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.

There is a sense, however, in which the above quotes may be misleading as regards when and how regret matters. The quotes, by and large, illustrate instances when regret actually *occurred*, and they were painful instances at that. We wish to emphasize that painful emotional pangs often shape outcomes by being *avoided*. That is, a pang of regret conditional on a possible course of action is anticipated but ultimately leads a player to behave in a way that renders the pang *counterfactual*! Our modeling effort makes this clear, and our examples have amply exemplified the issue.⁴¹

Perhaps there is a more general lesson in this. Think of all the world outcomes where people are not experiencing regret, for example, instances where someone studied hard, broke up a relationship, or became a professional golfer, and in the end did *not* on the death bed

⁴¹What we noted about pangs of regret – namely, that they are counterfactual yet influence outcomes – can likewise be said of other emotions, such as guilt, disappointment, or anger, which also carry negative valence or otherwise undesirable consequences. Battigalli & Dufwenberg (2022, p. 843) discuss the issue and also observe that it “marks a difference, to a degree, between what is the natural focus of economists and psychologists: For economists it is obvious that a counterfactual emotional experience is important, if it influences behavior and who gets what. Psychologists’ discussions, by contrast, tend to focus on the impact of guilt when it actually occurs.”

ruminate about related regrets because the person felt that they had by and large lived a life well-lived. We propose that these observations do not constitute counter-examples to the idea that regret is important to human decision-making. Anticipated counterfactual regret may still have played a key role in shaping the past behavior and the overall outcomes. We hope that our model will be a useful tool for fostering appreciation and understanding of that general subject matter.

9 Appendix

Proof of Proposition 3 Since there are no chance moves, we ease notation eliminating symbol s_0 in this proof. Let α_i satisfy the deterministic-beliefs Condition 1: for any given nonterminal information set $h \in H_i$, 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:

$$\begin{aligned}\mathbb{E}_{s_i, \alpha_i}[u_i|h] &= \sum_{s_{-i} \in S_{-i}(h)} \alpha_i(s_{-i}|h) \cdot u_i(\zeta(s_i^{|h|}, s_{-i}), \alpha_i) \\ &= u_i(\zeta(s_i^{|h|}, s_{-i}^*), \alpha_i) \\ &= m_i(\zeta(s_i^{|h|}, s_{-i}^*)) - \theta_i \cdot f_i \left(\max_{s'_i \in S_i} \mathbb{E}_{s'_i, \alpha_i}[m_i|\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))] - m_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right).\end{aligned}\tag{18}$$

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 $(s_i^{|h|}, s_{-i}^*)$:

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

Since the leftmost term equals 1, it follows that $\alpha_{i,-i} \left(s_{-i}^* | \bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*)) \right) = 1$. 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:

$$\mathbb{E}_{s'_i, \alpha_i}[m_i|\bar{H}_i(\zeta(s_i^{|h|}, s_{-i}^*))] = m_i(\zeta(s'_i, s_{-i}^*)).\tag{19}$$

Substituting this into Equation (18) yields:

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

Since 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 function $f_i(\cdot)$ is (strictly) decreasing in the value of $m_i(\zeta(s_i^{|h}, s_{-i}^*))$, it follows that for any two strategies $s_i, \tilde{s}_i \in S_i$,

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

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 Lemma 1. 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 $h \in H_i$, $s'_i \in S_i(h)$, 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 $h \in H_i$, $s'_i \in S_i(h)$, we obtain

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

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 . Sequential rationality of (π, α) follows from Lemma 1.

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}) such that $(\pi^{k_n}, \alpha^{k_n}) \rightarrow (\pi^n, \alpha^n)$, as $k_n \rightarrow \infty$. Now 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 and α , 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 $SE(\theta)$ is closed. Because $SE(\theta)$ 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: First we need some preliminary definitions and results. For any strategy $s_i \in S_i$, let

$$Z_i(s_i) := \{z \in Z : \exists (s_{-i}, s_0) \in S_{-i} \times S_0, z = \zeta(s_i, s_{-i}, s_0)\}$$

denote the set of terminal nodes compatible with strategy s_i . Perfect recall implies that each set $Z_i(s_i)$ is the union of a subcollection of the terminal information partition of i . To state this result, we need a preliminary definition: for any strategy $s_i \in S_i$, let

$$\bar{\mathcal{H}}_i(s_i) = \{h \in \bar{H}_i : \exists x \in h, \exists s_{I_0 \setminus \{i\}} \in S_{I_0 \setminus \{i\}}, x \preceq \zeta(s_i, s_{I_0 \setminus \{i\}})\}$$

denote the collection of information sets of player i that can occur if s_i is played (one can show that $\bar{\mathcal{H}}_i(s_i) = \{h \in \bar{H}_i : s_i \in S_i(h)\}$).

Lemma 2 *For every $i \in I$, $s_i \in S_i$, and information set $\bar{h} \in \bar{\mathcal{H}}_i(s_i)$ with $\bar{h} \subseteq Z$, it holds that $\bar{h} \subseteq Z_i(s_i)$. Therefore, for every $i \in I$ and $s_i \in S_i$,*

$$Z_i(s_i) = \bigcup_{z \in Z(s_i)} \bar{H}_i(z).$$

Proof of Lemma 2: Fix $i \in I$, $s_i \in S_i$, $\bar{h} \in \bar{\mathcal{H}}_i(s_i)$ with $\bar{h} \subseteq Z$, and $z \in \bar{h}$ arbitrarily. We must prove that $z \in Z_i(s_i)$. Let

$$\text{Exp}_i(z) = \{\{\emptyset\}, a_i^1(z), \dots, h_i^k(z), a_i^{k+1}(z), \dots, \bar{H}_i(z)\},$$

where $\bar{H}_i(z) = \bar{h}$ (because $z \in \bar{h} \in \bar{H}_i$), denote the **experience** of i on play z , that is, the sequence of information sets observed and actions taken by i on path z (cf. Osborne & Rubinstein 1994, Ch. 11). Since $\bar{h} \in \bar{\mathcal{H}}_i(s_i)$ (that is, s_i allows \bar{h}), perfect recall implies that $a_i^1(z) = s_i(\{\emptyset\})$ and $s_i(h_i^k(z)) = a_i^{k+1}$ for every k . Therefore, z is allowed by s_i , that is, $z \in Z_i(s_i)$. ■

Recall that

$$\mathcal{P}_i(s_i) := \{E_{I_0 \setminus \{i\}} \in 2^{S_{I_0 \setminus \{i\}}} : \exists z \in Z_i(s_i), E_{I_0 \setminus \{i\}} = S_{I_0 \setminus \{i\}}(\bar{H}_i(z))\}.$$

Lemma 3 *For every $i \in I$ and $s_i \in S_i$, collection $\mathcal{P}_i(s_i)$ is a partition of $S_{I_0 \setminus \{i\}}$.*

Proof of Lemma 3: Let

$$\begin{aligned} \phi_i : \quad S_i \times S_{-i} \times S_0 &\longrightarrow 2^Z, \\ (s_i, s_{-i}, s_0) &\longmapsto \bar{H}_i(\zeta(s_i, s_{-i}, s_0)) \end{aligned}$$

denote the (strategic-form) **feedback function** of i , that is, the function that associates each strategy profile (s_i, s_{-i}, s_0) with the terminal information set of i containing the induced terminal history $\zeta(s_i, s_{-i}, s_0)$. For any $s_i \in S_i$, the section at s_i of the feedback function is

$$\begin{aligned} \phi_{i,s_i} : S_{I_0 \setminus \{i\}} &\longrightarrow 2^{Z(s_i)}, \\ (s_{-i}, s_0) &\longmapsto \bar{H}_i(\zeta(s_i, s_{-i}, s_0)). \end{aligned}$$

With this,

$$\mathcal{P}_i(s_i) = \bigcup_{z \in Z_i(s_i)} \phi_{i,s_i}^{-1}(\bar{H}_i(\zeta(s_i, s_{-i}, s_0))).$$

Since, by Lemma 2, the images of ϕ_{i,s_i} form a partition of $Z_i(s_i)$, $\mathcal{P}_i(s_i)$ must be a partition of $S_{I_0 \setminus \{i\}}$. ■

We are now ready to prove Proposition 7. 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 terminal information set $\bar{h} \subseteq Z$ of i and any α_i , let $M_i^{\alpha_i}(s_i, \bar{h})$ denote the (ex post) expected material payoff from evaluating strategy $s_i \in S_i$ at \bar{h} :

$$M_i^{\alpha_i}(s_i, \bar{h}) := \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}(\bar{h})} \alpha_i(s_{-i}, s_0 | \bar{h}) m_i(\zeta(s_i, s_{-i}, s_0)). \quad (20)$$

Under *linear* regret, player i 's utility at z is

$$u_i(z, \alpha_i) = (1 + \theta_i) m_i(z) - \theta_i \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(z)).$$

Thus, for any $\hat{s}_i \in S_i$,

$$\begin{aligned} \mathbb{E}_{\alpha_i, \hat{s}_i}[u_i] &= \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0) u_i(\zeta(\hat{s}_i, s_{-i}, s_0), \alpha_i) \\ &= \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0) \left[(1 + \theta_i) m_i(\zeta(\hat{s}_i, s_{-i}, s_0)) - \theta_i \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{h}) \right] \\ &= (1 + \theta_i) \mathbb{E}_{\alpha_i, \hat{s}_i}[m_i] - \theta_i \mathbb{E}_{\alpha_i, \hat{s}_i} \left[\max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(\hat{s}_i, \cdot))) \right]. \end{aligned}$$

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

$$\mathbb{E}_{\alpha_i, s''_i} \left[\max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(s''_i, \cdot))) \right] \leq \mathbb{E}_{\alpha_i, s'_i} \left[\max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(s'_i, \cdot))) \right].$$

For any $\hat{s}_i \in S_i$, since $M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(\hat{s}_i, s_{-i}, s_0)))$ 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

$$\begin{aligned}
& \mathbb{E}_{\alpha_i, \hat{s}_i} \left[\max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(\hat{s}_i, \cdot))) \right] \\
&= \sum_{(s_{-i}, s_0) \in S_{I_0 \setminus \{i\}}} \alpha_i(s_{-i}, s_0) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{H}_i(\zeta(\hat{s}_i, s_{-i}, s_0))) \\
&= \sum_{\bar{h} \in \bar{\mathcal{H}}_i(\hat{s}_i), \bar{h} \subseteq Z} \alpha_i(S_{I_0 \setminus \{i\}}(\bar{h})) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, \bar{h}) \\
&= \sum_{E_{I_0 \setminus \{i\}} \in \mathcal{P}_i(\hat{s}_i)} \alpha_i(E_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, E_{I_0 \setminus \{i\}}),
\end{aligned}$$

where $\alpha_i(S_{I_0 \setminus \{i\}}(\bar{h}))$ is the ex ante probability of reaching terminal information set $\bar{h} \in \bar{\mathcal{H}}_i(\hat{s}_i)$, and — with a slight abuse of notation — we write $M_i^{\alpha_i}(s_i, E_{I_0 \setminus \{i\}}) = M_i^{\alpha_i}(s_i, S_{I_0 \setminus \{i\}}(\bar{h})) = M_i^{\alpha_i}(s_i, \bar{h})$ 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 $E''_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s'_i)$ is the union of cells $E'_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s'_i, E''_{I_0 \setminus \{i\}})$, where

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

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

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

and we must show

$$\begin{aligned}
& \sum_{E''_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s''_i)} \alpha_i(E''_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, E''_{I_0 \setminus \{i\}}) \\
&\leq \sum_{E''_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s''_i)} \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, E'_{I_0 \setminus \{i\}}).
\end{aligned}$$

It is sufficient to establish the inequality cell by cell. That is, we want to show that, for every $E''_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s''_i)$, the following inequality holds:

$$\alpha_i(E''_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, E''_{I_0 \setminus \{i\}}) \leq \sum_{E'_{I_0 \setminus \{i\}} \in \mathcal{P}_i(s'_i, E''_{I_0 \setminus \{i\}})} \alpha_i(E'_{I_0 \setminus \{i\}}) \max_{s_i \in S_i} M_i^{\alpha_i}(s_i, E'_{I_0 \setminus \{i\}}). \quad (21)$$

Consider first the left-hand side of (21), which can be expressed as

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

where the second equality follows from the chain rule:

$$\alpha_i(s_{-i}, s_0) = \alpha_i(s_{-i}, s_0 | E''_{I_0 \setminus \{i\}}) \alpha_i(E''_{I_0 \setminus \{i\}}) \quad \text{for } (s_{-i}, s_0) \in E''_{I_0 \setminus \{i\}}.$$

Similarly, the right-hand side of (21) can be expressed as

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

Since the maximum of a weighted sum is at most the weighted sum of maxima, the inequality follows:

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

Hence, inequality (21) follows. ■

Proof of Remark 3 By Theorem 1 in Kuhn (1953), two strategies s'_i and s''_i are realization-equivalent if and only if they are behaviorally equivalent, that is, $\bar{\mathcal{H}}_i(s'_i) = \bar{\mathcal{H}}_i(s''_i)$ and $s'_i(h) = s''_i(h)$ for all $h \in H_i \cap \bar{\mathcal{H}}_i(s'_i)$. The statement of Lemma 2 and proof of Lemma 3 make it clear that the ex post informativeness of strategies is invariant to behavioral equivalences; thus, it is invariant to realization-equivalences. ■

References

- [1] BATTIGALLI, P. (1996): “Strategic Independence and Perfect Bayesian Equilibria,” *Journal of Economic Theory*, 70, 201-234.
- [2] BATTIGALLI, P., AND M. DUFWENBERG (2009): “Dynamic Psychological Games,” *Journal of Economic Theory*, 144, 1-35.
- [3] BATTIGALLI, P., AND M. DUFWENBERG (2022): “Belief-Dependent Motivations and Psychological Game Theory,” *Journal of Economic Literature*, 60, 833-882.
- [4] BATTIGALLI, P., AND N. GENEROSO (2024): “Information Flows and Memory in Games,” *Games and Economic Behavior*, 145, 356-376.
- [5] BATTIGALLI, PIERPAOLO, EMILIANO CATONINI, AND JULIEN MANILI (2023): “Belief Change, Rationality, and Strategic Reasoning in Sequential Games,” *Games and Economic Behavior*, 142, 527-551.
- [6] BATTIGALLI P., R. CORRAO, AND M. DUFWENBERG (2019): “Incorporating Belief-Dependent Motivation in Games,” *Journal of Economic Behavior & Organization*, 167, 185-218.
- [7] BATTIGALLI, P., P. LEONETTI, AND F. MACCHERONI (2020): “Behavioral Equivalence of Extensive Game Structures,” *Games and Economic Behavior*, 121, 533-547.
- [8] BATTIGALLI, P., S. CERREIA, F. MACCHERONI, AND M. MARINACCI (2015): “Self-Confirming Equilibrium and Model Uncertainty,” *American Economic Review*, 105, 646-677.
- [9] BATTIGALLI, P., S. CERREIA-VIOGLIO, F. MACCHERONI, AND M. MARINACCI (2016): “A Note on Comparative Ambiguity Aversion and Justifiability,” *Econometrica*, 84, 1903-1916.
- [10] BELL, D. (1982): “Regret in Decision Making under Uncertainty,” *Operations Research*, 30, 961-981.
- [11] BERGEMANN, D., AND K. SCHLAG (2008): “Pricing Without Priors,” *Journal of the European Economic Association*, 6, 560-569.
- [12] BERGEMANN, D., AND K. SCHLAG (2011): “Robust Monopoly Pricing,” *Journal of Economic Theory*, 146, 2527-2543.

- [13] 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.
- [14] BIKHCHANDANI, S., AND U. SEGAL (2011): “Transitive Regret,” *Theoretical Economics*, 6, 95-108.
- [15] BIKHCHANDANI, S., AND U. SEGAL (2014): “Transitive Regret over Statistically Independent Lotteries,” *Journal of Economic Theory*, 152, 237-248.
- [16] CADISH, B. (2001): “Damn!: Reflections on Life’s Biggest Regrets”, *Kansas City: Andrews McMeel Publishing*, 2.
- [17] CERRONE, C., F. FERI, AND P. NEARY (2025): “Ignorance is Bliss: A Game of Regret,” mimeo.
- [18] CONNOLLY, T., AND D. BUTLER (2006): “Regret in Economic and Psychological Theories of Choice,” *Journal of Behavioral Decision Making*, 19, 139-154.
- [19] DUFWENBERG, M., AND K. NORDBLOM (2022): “Tax Evasion with a Conscience,” *Journal of Public Economic Theory*, 24, 5-29.
- [20] ENGELBRECHT-WIGGANS, R. (1989): “The Effect of Regret on Optimal Bidding in Auctions,” *Management Science*, 35, 685-92.
- [21] ENGELBRECHT-WIGGANS, R., AND E. KATOK (2008): “Regret and Feedback Information in First-Price Sealed-Bid Auctions,” *Management Science* 54, 808-819.
- [22] FILIZ-OZBAY, E., AND E. OZBAY (2007): “Auctions with Anticipated Regret: Theory and Experiment,” *American Economic Review*, 97, 1407-1418.
- [23] GEANAKOPOLOS, J., D. PEARCE, AND E. STACCHETTI (1989): “Psychological Games and Sequential Rationality,” *Games and Economic Behavior*, 1, 60-80.
- [24] 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.
- [25] HALPERN, J., AND R. PASS (2012): “Iterated Regret Minimization: A New Solution Concept,” *Games and Economic Behavior*, 74, 194-207.
- [26] HAYASHI, T. (2008): “Regret Aversion and Opportunity Dependence,” *Journal of Economic Theory*, 139, 242–268.

- [27] KOHLBERG, E., AND P. RENY (1997): “Independence on Relative Probability Spaces and Consistent Assessments in Game Trees,” *Journal of Economic Theory*, 75, 280-313.
- [28] KREPS, D., AND R. WILSON (1982): “Sequential Equilibrium,” *Econometrica*, 50, 863-94.
- [29] KUHN, H.W. (1953): “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games II*, ed. by H.W. Kuhn and A.W. Tucker. Princeton: Princeton University Press, 193-216.
- [30] LOOMES, G., AND R. SUGDEN (1982): “Regret Theory: An Alternative Theory of Rational Choice under Uncertainty,” *Economic Journal*, 92, 805-824.
- [31] LOOMES, G., AND R. SUGDEN (1987): “Some Implications of a More General Form of Regret Theory,” *Journal of Economic Theory*, 41, 270-287.
- [32] OSBORNE M., AND A. RUBINSTEIN (1994): *A Course in Game Theory*, MIT Press.
- [33] QUIGGIN, J. (1990): “Stochastic Dominance in Regret Theory,” *Review of Economic Studies*, 57, 503-11.
- [34] QUIGGIN, J. (1994): “Regret Theory with General Choice Sets,” *Journal of Risk and Uncertainty*, 8, 153-65.
- [35] RENOU, L., AND K. H. SCHLAG (2010): “Minimax Regret and Strategic Uncertainty,” *Journal of Economic Theory*, 145, 264–286.
- [36] RENOU, L., AND K. H. SCHLAG (2011): “Implementation in Minimax Regret Equilibrium,” *Games and Economic Behavior*, 71, 527–533.
- [37] SARVER, T. (2008): “Anticipating Regret: Why Fewer Options May Be Better,” *Econometrica*, 76, 263-305.
- [38] SAVAGE, L. (1951): “The Theory of Statistical Decision,” *Journal of the American Statistical Association*, 46, 55-67.
- [39] SELTEN, R. (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory* 4, 25–55.
- [40] STOYE, J. (2011): “Axioms for Minimax Regret Choice Correspondences,” *Journal of Economic Theory*, 146, 2226–2251.

- [41] SYAM, N., P. KRISHNAMURTHY, AND J.D. HESS (2008): “That’s what I thought I wanted? Miswanting and Regret for a Standard Good in a Mass-Customized World.” *Marketing Science*, 27, 2008, 379–397.
- [42] VAN DER PLOEG, F., AND A. REZAI (2019): “The Agnostic’s Response to Climate Deniers: Price carbon!” *European Economic Review*, 111, 70-84.
- [43] WEINSTEIN, J. (2016): “The Effect of Risk Attitudes on Strategic Behavior,” *Econometrica*, 84, 1881-1902.
- [44] WU, W., AND J. JIANG (1962): “Essential Equilibrium Points of n-person Non-cooperative Games”, *Scientia Sinica*, 11, 1307-1322.
- [45] ZEELENBERG, M., J. BEATTIE, J. VAN DER PLIGT, AND N. K. DE VRIES (1996): “Consequences of Regret Aversion: Effects of Expected Feedback on Risky Decision Making,” *Organizational Behavior and Human Decision Processes* 65, 148-158.
- [46] ZEELENBERG, M. (1999) “Anticipated Regret, Expected Feedback and Behavioral Decision Making,” *Journal of Behavioral Decision Making*, 12, 93-106.
- [47] ZEELENBERG, M., AND R. PIETERS (2007a): “A Theory of Regret Regulation 1.0,” *Journal of Consumer Psychology*, 17, 3-18.
- [48] ZEELENBERG, M., AND R. PIETERS (2007b): “A Theory of Regret Regulation 1.1,” *Journal of Consumer Psychology*, 17, 29-35.
- [49] ZOU, T., B. ZHOU, AND B. JIANG (2020): “Product-Line Design in the Presence of Consumers’ Anticipated Regret.” *Management Science*, 66, 5665–5682.