

# Least-Squares Regret and Partially Strategic Players

Chlump Chatkupt

A thesis submitted for the degree of  
*Doctor of Philosophy*

Department of Mathematics  
London School of Economics and Political Science

August 4, 2015

---

## Declaration

I certify that the dissertation that I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is my own work.

The copyright of this dissertation rests with the author. Quotation from it is permitted provided that full acknowledgement is made. This dissertation may not be reproduced without the prior written consent of the author.

I warrant that this authorization does not to the best of my belief infringe the rights of any third party.

The question of the word count of a mathematical work makes little sense. Still, we note that the code

---

```
texcount -merge -sum Dissertation.tex
```

---

yields the output

---

```
File: Dissertation.tex
Encoding: ascii
Sum count: 34550
Words in text: 31776
Words in headers: 1
Words outside text (captions, etc.): 282
Number of headers: 1
Number of floats/tables/figures: 45
Number of math inlines: 2161
Number of math displayed: 330
```

---

## Abstract

Noncooperative game theory enjoys a vast canon of solution concepts. The predominant solution concept is *Nash equilibrium* (Nash, 1950a; Nash, 1951). Other solution concepts include generalizations and refinements of Nash equilibrium as well as alternatives to it.

Despite their successes, the established solution concepts are in some ways unsatisfactory. In particular, for many games, such as the *Centipede Game* (Rosenthal, 1981), the *p-Beauty Contest* (Moulin, 1986; Simonsen, 1988), and the notorious *Traveler's Dilemma* (Basu, 1994; Basu, 2007), many of the solution concepts yield solutions that are both unreasonable in theory and refuted by the experimental evidence. And when a solution concept manages to yield the expected or reasonable solutions for such games, it often suffers from other difficulties such as unwieldy complexity or reliance on *ad hoc* or game-specific constructions that may fail to be generalizable.

We propose a new solution concept, which we call *least-squares regret*, that yields the expected or reasonable solutions for games that have thus far proved to be problematic, such as the *Traveler's Dilemma*; that is simple; that involves no *ad hoc* or game-specific constructions and can thus be applied immediately and consistently to any arbitrary game; that exhibits nice properties; and that is grounded in human psychology. Intuitively, we suppose that a player chooses a strategy so as to minimize the divergence from perfect play overall. In particular, we suppose that a player is partially strategic and chooses a strategy so as to minimize the sum, across all partial profiles of strategies of the other players, of the squares of the regrets, where the regret of a strategy with respect to a partial profile is the difference of the best-response payoff with respect to the partial profile and the payoff from choosing the strategy with respect to the partial profile.

The aim of this work is to develop the solution concept of least-squares regret; explore its properties; assess its performance with respect to various games of interest; determine its merits and demerits, especially in relation to other solution concepts; review its weaknesses; introduce a refinement, which we call *mutual weighted least-squares regret*, that addresses some of the weaknesses; and propose some questions for further research.

## Acknowledgements

A creative work, even if it is produced independently, is never a solitary effort, but a product also of the lives that have come to define the creator. Thus, this work belongs as much to those who have played a role in my life as it does to me, and I wish here to add their names to what is necessarily a collaborative work.

I am grateful for their measureless love to my father Surachat, my mother Sansnee, and my brother Patrick, about whom not enough can be said.

I am indebted for reasons too great to express to my beloved friends: Ahmad Abu-Khazneh; John Allen; Amanda Anjum; Greg Bognar; Susanne Burri; Angela Chang; Catherine Chen; Kwan Chen; Ying Zhu Chin; Yichen Chiu; Benjamin Ferguson; Johannes Himmelreich; Jee-Hee Hong; Yuna Lee; Tomás Mayr; Emilio Mora; William Ruddick; Esha Senchaudhuri; Hamid Seyedsayamdost; Orri Stefansson; Stephanie Tang; Stuart Theobald; Derek Wan; Helen Wong; and Emaad Zaman. To countless others, the lack of mention here is evidence only of my own failing; your significance is immeasurable.

I am grateful for her boundless patience, understanding, and support to Debbie Carmody, who deserves far more recognition and appreciation than I could ever articulate.

I wish to thank the Department of Mathematics for institutional support and the London School of Economics and Political Science for the full PhD scholarship that funded my studies.

I wish to thank for additional support the Department of Philosophy, Logic and Scientific Method and, in particular, Richard Bradley, Roman Frigg, Christian List, and Bryan Roberts.

My debt to Luc Bovens, whose guidance and encouragement have never wavered, is too immense to convey, but it is always felt.

It is impossible to conclude this note without recognizing the debt that I owe to the mentor who, with his keenness and generosity and unfailing sympathy and encouragement, influenced and supported me most, my supervisor Bernhard von Stengel, whose pupil I shall always remain.

*For S. B.*  
*and*  
*B. v. S.*

---

# Contents

<b>Contents</b>	<b>6</b>
<b>1 Introduction</b>	<b>8</b>
1.1 Solution Concepts for Noncooperative Games . . . . .	9
1.2 Motivation . . . . .	13
1.3 The Concept of Regret . . . . .	16
1.4 Nonstrategic, Partially Strategic, and Fully Strategic Reasoning . .	17
1.5 Contributions of this Dissertation . . . . .	19
<b>2 Least-Squares Regret</b>	<b>23</b>
2.1 Notation and Basic Concepts . . . . .	23
2.2 Formal Definition . . . . .	25
2.3 An Example . . . . .	29
2.4 An Alternative Definition . . . . .	32
2.5 Discussion . . . . .	36
<b>3 Alternative Solution Concepts</b>	<b>41</b>
3.1 Least-Squares Regret and Maximin . . . . .	41
3.2 Least-Squares Regret and Iterated Regret Minimization . . . . .	43
<b>4 The Traveler's Dilemma</b>	<b>53</b>
4.1 Definition of the Game . . . . .	53
4.2 Standard Analyses of the Game . . . . .	54
4.3 Experimental Evidence . . . . .	56
4.4 Alternative Analyses . . . . .	56
4.5 Least-Squares Regret . . . . .	57
<b>5 Properties</b>	<b>61</b>
5.1 Existence . . . . .	62
5.2 Full Equivalence . . . . .	62

5.3	Dominated Strategies . . . . .	64
5.4	Iterative Elimination of Dominated Strategies . . . . .	67
5.5	Uniformly Dominant Strategies . . . . .	68
5.6	Strategic Equivalence . . . . .	71
5.7	Risk Dominance . . . . .	76
5.8	Convexity . . . . .	80
5.9	Uniqueness . . . . .	82
<b>6</b>	<b>Further Examples</b>	<b>85</b>
6.1	Dollar Auction . . . . .	85
6.2	Bertrand Competition . . . . .	87
6.3	Inspection Games . . . . .	89
6.4	Matching Pennies . . . . .	91
6.5	Chicken . . . . .	94
6.6	Coordination Games . . . . .	95
6.7	Battle of the Sexes . . . . .	99
6.8	The Two-Person Bargaining Problem . . . . .	101
<b>7</b>	<b>Weaknesses of Least-Squares Regret</b>	<b>107</b>
7.1	Independence of Irrelevant Alternatives . . . . .	107
7.2	Maximin and Two-Person Zero-Sum Games . . . . .	109
7.3	Framing Effects . . . . .	111
7.4	Fully Strategic Reasoning . . . . .	112
<b>8</b>	<b>Mutual Weighted Least-Squares Regret</b>	<b>116</b>
8.1	Formal Definition . . . . .	117
8.2	An Example . . . . .	118
8.3	Existence . . . . .	120
8.4	Mutual Weighted Least-Squares Regret and Nash Equilibrium . . .	123
8.5	Recursion and Convergence . . . . .	126
<b>9</b>	<b>Conclusion</b>	<b>128</b>
9.1	Further Questions . . . . .	128
	<b>Bibliography</b>	<b>131</b>

# 1

---

## Introduction

Noncooperative game theory enjoys a vast canon of solution concepts. The predominant solution concept is *Nash equilibrium* (Nash, 1950a; Nash, 1951). Other solution concepts include generalizations and refinements of Nash equilibrium as well as alternatives to it.

Despite their successes, the established solution concepts are in some ways unsatisfactory. In particular, for many games, such as the *Centipede Game* (Rosenthal, 1981), the *p-Beauty Contest* (Moulin, 1986; Simonsen, 1988), and the notorious *Traveler's Dilemma* (Basu, 1994; Basu, 2007), the last of which we address in detail in Chapter 4, many of the solution concepts yield solutions that are both unreasonable in theory and refuted by the experimental evidence. And when a solution concept manages to yield the expected or reasonable solutions for such games, it often suffers from other difficulties such as unwieldy complexity or reliance on *ad hoc* or game-specific constructions that may fail to be generalizable.

Thus, we are led naturally to ask: is it possible to develop a solution concept that yields the expected or reasonable solutions for games that have thus far proved to be problematic if not also for other games of interest; that is sufficiently simple so as to be both trivial to apply and also a plausible characterization of typical reasoning and behavior; that involves no *ad hoc* or game-specific constructions and can thus be applied immediately and consistently to any arbitrary game; that exhibits nice mathematical and conceptual properties; and that is grounded in human psychology? This dissertation seeks to develop such a solution concept.

While there may be reasons to act in accordance with one or another of the established solution concepts, we propose that there is an alternative way to reason about a game. This alternative way to reason is best understood by considering how a player might envisage his decision problem. We consider a brief informal characterization here and an illustrative example in Section 1.2.

In playing a game, a player must choose a single strategy in ignorance of the



strategies chosen by the other players. But whether or not a strategy is a best response may depend on the strategies that the other players choose. Thus, a strategy may not be perfect with respect to every course of play: while it may constitute perfect play with respect to some strategies of the other players, it may diverge from perfect play with respect to other strategies. Furthermore, a strategy may diverge more or less from perfect play depending on the extent to which it falls short.

We suppose that a player chooses a strategy so as to minimize the divergence from perfect play overall. In particular, we suppose that a player chooses a strategy so as to minimize the sum, across all partial profiles of strategies of the other players, of the squares of the regrets, where the regret of a strategy with respect to a partial profile is the difference of the best-response payoff with respect to the partial profile and the payoff from choosing the strategy with respect to the partial profile. This idea is the basis of our solution concept, which we call *least-squares regret*. This dissertation is concerned with defining, developing, defending, assessing, and refining least-squares regret.

The remainder of this chapter is concerned with explaining the need to develop an alternative solution concept and discussing informally the essential ideas behind least-squares regret. Section 1.1 considers some of the established solution concepts and their inadequacies. Section 1.2 motivates least-squares regret by considering a simple decision problem. Section 1.3 discusses the concept of regret, its origin in decision theory, its use in game theory, and evidence from experimental economics suggesting its role in human psychology. Section 1.4 considers the various degrees to which a player might be strategic. Section 1.5 outlines the contributions of this dissertation.

## 1.1 Solution Concepts for Noncooperative Games

We noted earlier that the established solution concepts are in some ways unsatisfactory, suggesting the need to develop an alternative solution concept such as least-squares regret. In this section, we consider briefly some of the established solution concepts and the issues surrounding them.

One of the oldest and most developed solution concepts is *maximin* (von Neumann, 1928; Wald, 1939; Wald, 1945; von Neumann and Morgenstern, 1947; Wald, 1950). While maximin is the standard solution concept for two-person zero-sum games, its significance when it comes to other classes of games is less certain. In particular, focusing exclusively on the minimum payoff may be too restrictive, and for many games of interest, maximin yields unsatisfactory solutions. We examine maximin in more detail, especially in relation to least-squares regret, in

Sections 3.1 and 7.2.

That Nash equilibrium is the predominant solution concept and has enjoyed considerable success is undeniable. For perspectives on Nash equilibrium and its influence, see Myerson (1999) and Holt and Roth (2004).

But as is well known, it is not without its problems. Many solution concepts have been developed in order to address these problems and others.

In one sense, Nash equilibrium is too weak: it may yield a multiplicity of solutions and fail to exclude unreasonable ones. Selection criteria and refinements of Nash equilibrium—such as the *focal-point effect* (Schelling, 1960); *perfect equilibrium* (Selten, 1975); *proper equilibrium* (Myerson, 1978); *sequential equilibrium* (Kreps and Wilson, 1982); *persistent equilibrium* (Kalai and Samet, 1984); *stable equilibrium* (Kohlberg and Mertens, 1986); and *payoff dominance* and *risk dominance* (Harsanyi and Selten, 1988)—introduce additional constraints in order to exclude unreasonable equilibria and to guide equilibrium selection in games with multiple equilibria.

In another sense, Nash equilibrium is too strong: it may exclude reasonable outcomes, sometimes on the basis of implausible assumptions. For example, it assumes very strong conditions on the players such as the condition that each player have no mistaken beliefs about the strategies of the other players, even when such expectations may be untenable. Generalizations of Nash equilibrium, such as *rationalizability* (Bernheim, 1984; Pearce, 1984), weaken the conditions assumed in order to include more reasonable outcomes and to obtain more accurate characterizations of rationality.

As characterizations of human behavior, however, these solution concepts have mixed records. For many games, such as those noted earlier, these solution concepts fail to capture the observed behavior. A salient example is the *finitely repeated Prisoners' Dilemma*, in which cooperation is ruled out by the solution concepts while being routinely confirmed by experiments (Smale, 1980; Axelrod, 1981).

One common approach to explain observed or so-called “anomalous” behavior in a problematic game is to modify the game in a slight, but instrumental, way and then to proceed as usual with some standard solution concept or another. Such an approach minimizes the departure from standard game theory. For example, Kreps, Milgrom, Roberts, and Wilson (1982) explain cooperation in the finitely repeated Prisoners' Dilemma by modeling the game as a Bayesian game that includes an irrational or “cooperative” type.

Some more recent solution concepts are readier to depart from standard game theory for the sake of better explanations of observed behavior. Many of these solution concepts—such as *quantal response equilibrium* (McKelvey and Palfrey,

1995; McKelvey and Palfrey, 1998), *level-k thinking* (Stahl and Wilson, 1994; Nagel, 1995; Stahl and Wilson, 1995; Camerer, Ho, and Chong, 2004; Crawford, Costa-Gomes, and Iriberry, 2013), and *noisy introspection* (Goeree and Holt, 1999; Goeree and Holt, 2004)—can capture the observed behavior exceedingly well. Their successes can be explained in part by their appeals to noise or bounded reasoning capacity.

Other recent solution concepts employ sophisticated constructions and procedures in order to model how people reason and to explain observed behavior. According to *minimax weighted expected regret* (Halpern and Leung, 2014), the beliefs of a player about uncertain events are modeled by specifying a set of weighted probability distributions, and a player chooses a strategy so as to minimize the maximum weighted expected regret. According to *common belief in utility proportional beliefs* (Bach and Perea, 2014), each player holds utility proportional beliefs, according to which the differences of the probabilities of the strategies of the other players are proportional to the differences of the utilities of the strategies, believes that each player holds such utility proportional beliefs, and so on.

The solution concepts described above have their advantages and may be appropriate for many games or situations. Nevertheless, as noted earlier, there are reasons to be dissatisfied with certain aspects of these solution concepts and thus reasons to seek an alternative solution concept.

As noted earlier, for many games of interest, many of the established solution concepts yield solutions that are both unreasonable in theory and refuted by the experimental evidence. The failure of the solution concepts to capture or explain the expected or reasonable behavior in these games renders the solution concepts inappropriate and less useful for characterizing reasoning and behavior. This problem is particularly pointed when, as in the case of the Traveler's Dilemma discussed in Chapter 4, the discrepancies are striking and perplexing.

The trouble with the common approach of modifying a game to capture the observed behavior is that the modifications are generally *ad hoc* or game-specific. For example, the method of Kreps, Milgrom, Roberts, and Wilson (1982) to explain cooperation in the finitely repeated Prisoners' Dilemma turns on the *pat* definition of an irrational or "cooperative" type, which is specific to the game. But the *ad hoc* or game-specific nature of such modifications means that the common approach of modifying a game cannot be applied to a game without some preparatory work, for example, defining an irrational or "cooperative" type and what is entailed. Furthermore, there is no guarantee that a certain modification, being specific to the game for which it was developed, can be generalized to other games. For example, for many games, it is not clear what an irrational or "cooperative" type would be or whether such a type would have any sensible meaning. Thus, while modifying

a game may yield the desired effect, such a method may not be appropriate in general.

A solution concept may be game-specific in other unsatisfactory ways, such as depending on particular specifications of a game that are decision-theoretically negligible. For example, it is well known that quantal response equilibrium is not scale invariant and thus depends on the particular utility functions specified in a game (Wright and Leyton-Brown, 2010). But the dependence of a solution concept on the particular specifications of a game that are decision-theoretically negligible, such as the dependence of quantal response equilibrium on the particular utility functions specified in a game, may render the solution concept inconsistent, yielding different solutions for games that are considered decision-theoretically equivalent.

Many of the more recent solution concepts involve specifying structures in advance, for example, an error structure, as in quantal response equilibrium or noisy introspection, or a typology of players, as in level- $k$  thinking, and then estimating free parameters, for example, a precision parameter or a population distribution parameter. Such an approach can be problematic. The need to make preliminary specifications, for example, the behavior of a level-0 type in level- $k$  thinking, and the *ad hoc* or game-specific nature of such specifications threaten ease of application and generalizability. Furthermore, the flexibility allowed in specifying a model, the focus on in-sample parameter estimation, and the freedom to choose the best parameter estimates mean that, without suitable restrictions, a model can often be adjusted to fit virtually any set of data, may be susceptible to overfitting, and may fail to generalize to other games. For more on these issues, see, for example, Goeree, Holt, and Palfrey (2005); Haile, Hortaçsu, and Kosenok (2008); Wright and Leyton-Brown (2010); Burchardi and Penczynski (2012); and Crawford, Costa-Gomes, and Iriberri (2013).

Finally, many of the established solution concepts are sufficiently complex so as to be both nontrivial to apply and also dubious characterizations of typical reasoning and behavior. For example, Nash equilibrium and related solution concepts involve complex computations of fixed points and assume of the players unbounded reasoning capacity. Even weaker solution concepts, such as rationalizability, may involve reasoning of the infinitely iterated or complex sort that is not likely to be plausible. Furthermore, a solution concept that is simple in the one sense may be complex in the other sense. For example, while level- $k$  thinking involves the simple and fairly plausible assumption of bounded reasoning capacity, solving a game using level- $k$  thinking involves nontrivial parameter estimations and extensive testing.

Much more can be said, of course, about these solution concepts and others.

For further discussions on various solution concepts and related issues, see, for example, Kreps (1990); Myerson (1991); Fudenberg and Tirole (1991); van Damme (1991); Osborne and Rubinstein (1994); Aumann (1997); Rubinstein (1998); Goeree and Holt (1999); Goeree and Holt (2001); van Damme (2002); Hillas and Kohlberg (2002); Govindan and Wilson (2008); Halpern (2008); Wright and Leyton-Brown (2010); Fudenberg (2010); Shubik (2012); Crawford, Costa-Gomes, and Iriberri (2013); Crawford (2013); and Camerer and Ho (2015).

Whatever might be said about the various solution concepts that have been developed, it is clear that interest in alternative solution concepts persists, and new solution concepts continue to be proposed and discussed.

For our part, it is our dissatisfaction with certain aspects of the established solution concepts that inspired our search for an alternative solution concept. What we are after is simply a solution concept that is less susceptible to the problems described above, that is, a solution concept that yields the expected or reasonable solutions for the games of interest; that involves no *ad hoc* or game-specific constructions and can be applied immediately and consistently to any arbitrary game; that involves no exogenous structures or free parameters; that is simple; that has nice properties; and that has some basis in actual human reasoning.

It is worth emphasizing that what we are after is neither a refinement nor a generalization of Nash equilibrium. And that is because, for many games of interest, the set of solutions that seem reasonable and the set of equilibria may be disjoint. The Traveler's Dilemma, discussed in Chapter 4, is a notable example of this point. Thus, what we are after is an alternative to Nash equilibrium.

## 1.2 Motivation

In this section, we explore the motivation behind least-squares regret as a way to reason about an uncertain situation, such as a game, and thus as a solution concept for noncooperative games.

Suppose that a player is faced with the following decision problem. Let  $\Theta = \{\theta_1, \theta_2, \theta_3, \theta_4, \theta_5, \theta_6\}$  be the set of states of the world, let  $X = \{x, y\}$  be the set of strategies available to the player, and let the function  $u: \Theta \times X \rightarrow \mathbf{R}$  be the payoff function for the player with the payoffs as shown in Table 1.1. Only one state will obtain. The trouble is that the player does not know which state will obtain. Suppose that the uniform distribution is applied to the set  $\Theta$  of states of the world. The question is how to reason about such a problem.

According to the traditional approach, the aim is to maximize the expected payoff (Borel, 1921; von Neumann, 1928; von Neumann and Morgenstern, 1947;

Table 1.1 Payoffs in a decision problem

X \ $\Theta$	$\theta_1$	$\theta_2$	$\theta_3$	$\theta_4$	$\theta_5$	$\theta_6$
x	1	2	3	4	5	5
y	2	3	4	5	6	0

Savage, 1954). Given the uniform distribution over the set  $\Theta$  of states of the world, the expected payoff from choosing  $x$  and the expected payoff from choosing  $y$  are identical and equal to  $10/3$ . Thus, neither strategy is better than the other provided that the aim is to maximize the expected payoff.

Nevertheless, we propose that, despite the equality of the expected payoff from choosing  $x$  and the expected payoff from choosing  $y$ , there is a sense in which  $x$  is more reasonable than  $y$ . In every state apart from  $\theta_6$ , choosing  $x$  yields a payoff that is only slightly less than the best-response payoff that is achieved by choosing  $y$ , and in state  $\theta_6$ , choosing  $x$  yields the best-response payoff, which is significantly greater than the payoff from choosing  $y$ . In every state apart from  $\theta_6$ , choosing  $y$  yields the best-response payoff, which is only slightly greater than the payoff from choosing  $x$ , and in state  $\theta_6$ , choosing  $y$  yields a payoff that is significantly less than the best-response payoff that is achieved by choosing  $x$ .

Clearly, neither strategy is perfect with respect to every state. Each strategy constitutes perfect play with respect to some state and diverges from perfect play with respect to some other state. But by choosing  $x$ , the player minimizes the divergence from perfect play overall: whatever state ultimately obtains, choosing  $x$  ends up being, if not perfect play, then very close to it. By contrast, while choosing  $y$  yields the best-response payoff in every state apart from  $\theta_6$ , if  $\theta_6$  obtains, then choosing  $y$  falls far short of perfect play.

We can make this reasoning explicit as follows. We suppose that the player computes the sum

$$\sum_{\theta \in \Theta} \left( \max_{z \in X} u(\theta, z) - u(\theta, x) \right)^2 = 5,$$

which can be thought of as a measure of the extent to which  $x$  diverges from perfect play overall, and the sum

$$\sum_{\theta \in \Theta} \left( \max_{z \in X} u(\theta, z) - u(\theta, y) \right)^2 = 25,$$

which can be thought of as a measure of the extent to which  $y$  diverges from perfect play overall. Furthermore, we suppose that the player chooses a strategy so as to minimize the divergence from perfect play overall and thus chooses  $x$ .

More generally, we suppose that a player chooses a strategy so as to minimize the sum, across all states of the world, of the squares of the regrets, where the regret of a strategy with respect to a state is the difference of the best-response payoff with respect to the state and the payoff from choosing the strategy with respect to the state.

We extend this idea to noncooperative games in the natural way. In particular, we suppose that a player views the strategies of the other players as uncertain events, with each partial profile of strategies constituting a state of the world. Furthermore, we suppose that a player chooses a strategy so as to minimize the sum, across all partial profiles of strategies of the other players, of the squares of the regrets, where the regret of a strategy with respect to a partial profile is the difference of the best-response payoff with respect to the partial profile and the payoff from choosing the strategy with respect to the partial profile—hence the name *least-squares regret*.

Thus, least-squares regret can be seen as an extension of the decision-theoretic approach to game theory advocated by Kadane and Larkey (1982) and Raiffa (1982) and explored further by Roth and Schoumaker (1983). The decision-theoretic approach to game theory involves viewing a game from the perspective of each player considered individually, assessing a probability distribution to characterize the beliefs of the player about the strategies of the other players, and then identifying the strategies of the player that maximize his expected payoff with respect to the assessed probability distribution. This approach to game theory contrasts sharply with the standard approach, championed by Harsanyi (1982) and others. The standard approach to game theory involves solving the decision problems of the players considered altogether, as if the decision problems constituted a system of equations in several unknowns.

Least-squares regret differs from the decision-theoretic approach to game theory in two ways. Least-squares regret involves assigning the uniform distribution, formalized in terms of unweighted summation, while the decision-theoretic approach involves assessing a probability distribution that need not be uniform. Furthermore, least-squares regret supposes that a player chooses a strategy so as to minimize the sum of the squares of the regrets while the decision-theoretic approach supposes that a player chooses a strategy so as to maximize the expected payoff with respect to the assessed probability distribution.

We recognize that assuming the uniform distribution may be too restrictive. We discuss this assumption in Sections 1.4, 2.2, 2.5, and 7.4. In Chapter 8, we

introduce a refinement of least-squares regret that involves assessing probability distributions induced by the strategies of the other players and replacing unweighted summation with weighted summation, where the weights are the assessed probabilities. Refining least-squares regret in this way brings it even closer to the decision-theoretic approach to game theory.

### 1.3 The Concept of Regret

Central to least-squares regret, as the name suggests and as discussed in Section 1.2, is the concept of regret. Intuitively, the regret of a strategy is a measure of the extent to which the strategy falls short of perfect play.

The concept of regret has a long history, originating in decision theory with Niehans (1948) and Savage (1951). Notably, Savage (1951) introduces the concept of regret, explicitly distinguishing it from the concept of negative income, partly in order to highlight the failure of Wald (1950) and others to distinguish the two concepts and argues that the minimax principle is less pessimistic and more reasonable when applied to regret than when applied to negative income.

Despite advocating the minimax principle as applied to regret, Savage (1951) professes that no absolute justification for the minimax principle can be given and concedes that an argument might well exist for preferring to minimize the average regret or some other aggregate measure. One such alternative is precisely least-squares regret, which can be thought of as involving the minimization of the average of the squares of the regrets. In Section 3.2, we show that focusing exclusively on the maximum regret, as the minimax principle specifies, may be too restrictive and that it may be better to minimize the sum of the squares of the regrets, as least-squares regret specifies.

While the concept of regret had since its introduction gained considerable currency in decision theory, it is only recently that an interest has grown in employing it in game theory. Many solution concepts and models now incorporate the concept of regret, largely to obtain more accurate characterizations of observed behavior. Linhart and Radner (1989) apply the concept of regret to the problem of sealed-bid bargaining with multiple variables and show that an analysis based on bidding so as to minimize the maximum regret has a number of advantages over equilibrium analysis. The approach to learning known as *learning direction theory* (Selten and Stoecker, 1986; Selten and Buchta, 1999; Selten, Abbink, and Cox, 2005) considers learning via adjustments made on the basis of regrets and not on the basis of experienced payoffs, as in standard reinforcement learning theories. Filiz-Ozbay and Ozbay (2007) explain overbidding in first-price auctions in terms of anticipation of loser regret. Renou and Schlag (2010) introduce a new



solution concept called *minimax regret equilibrium* according to which each player in a game chooses a belief about the behavior of the other players according to a minimax regret criterion and then chooses a best response. Halpern and Pass (2012) introduce a new solution concept called *iterated regret minimization*, which we examine in Sections 3.2 and 6.8, that involves iteratively eliminating all of the strategies that fail to minimize the maximum regret.

Furthermore, as studies in experimental economics show, regret appears to play a significant role in human psychology. Ritov (1996) shows that preferences can change depending on the extent to which uncertainty can be expected to be resolved and that a principal component of this effect is the anticipated experience of regret. Grosskopf, Erev, and Yechiam (2006) show that people tend to be extremely sensitive to foregone payoffs and that foregone payoff information can have significant effects on choice behavior depending on the environment.

Thus, there seems to be good reason to incorporate the concept of regret into the analysis of noncooperative games. We propose least-squares regret as one particular way to do so.

## 1.4 Nonstrategic, Partially Strategic, and Fully Strategic Reasoning

As discussed in Section 1.1, solution concepts can differ in the degree of sophistication assumed of a player, with some solution concepts supposing a more sophisticated player and other solution concepts supposing a less sophisticated player.

A player might be nonstrategic. In particular, a player might reason about the other players to no appreciable degree, form no assumptions about them, and then, without thinking and without any particular aim, choose a strategy randomly according to some probability distribution, say, the uniform distribution.

Alternatively, a player might be partially strategic. In particular, a player might reason about the other players to a limited degree, form only rudimentary assumptions about them, and then, in some principled way, choose a strategy accordingly.

Finally, a player might be fully strategic. In particular, a player might reason about the other players to an unlimited degree, form very sophisticated assumptions about them, including assumptions about how they might reason about one another, and then, in some principled way, choose a strategy accordingly.

This typology is fairly standard and fruitfully used. For example, it is central to level- $k$  thinking (Stahl and Wilson, 1994; Nagel, 1995; Stahl and Wilson, 1995;

Camerer, Ho, and Chong, 2004; Crawford, Costa-Gomes, and Iriberri, 2013). In level- $k$  thinking, a level-0 type is assumed to be nonstrategic, forming no beliefs about the other players and choosing a strategy randomly according to some probability distribution, typically the uniform distribution, and, for any level  $k$  such that  $k \geq 1$ , a level- $k$  type is assumed to be partially strategic, believing the other players to constitute a population of types from levels 0 to  $k-1$  and choosing a strategy accordingly.

Most solution concepts suppose that a player is fully strategic. But as seen in Section 1.1, such an assumption may be a dubious characterization of typical reasoning and behavior and yield a solution concept that is nontrivial to apply and fails to capture the expected or reasonable behavior, especially in games in which a player is less likely to be fully strategic. As noted earlier, many solution concepts, addressing these inadequacies, suppose instead that a player is partially strategic.

In defining least-squares regret, we likewise suppose that a player is partially strategic in order to address the inadequacies just noted. In particular, as discussed in Section 1.2, we suppose that a player treats uniformly the partial profiles of strategies of the other players and chooses a strategy so as to minimize the sum of the squares of the regrets. Equivalently, to use the language of level- $k$  thinking, we suppose that a player is a level-1 type who believes that each of the other players is a level-0 type randomizing according to the uniform distribution and chooses a strategy accordingly.

Supposing complete ignorance as to which one of a set of mutually exclusive events will obtain, it is not unjustifiable to regard them as all on a par.

One familiar justification of this position invokes the principle of maximum entropy (Jaynes, 1957a; Jaynes, 1957b). This principle asserts that when inferences are to be made on the basis of partial information, the proper distribution to use is the one with the maximum entropy subject to whatever is known since such a distribution yields the most unbiased representation of the knowledge of the state of the system under consideration. Notably, where nothing is known, the maximum-entropy distribution is the uniform distribution.

Another familiar justification invokes the principle of insufficient reason, which goes back to Bernoulli (1713) and Laplace (1825). This principle asserts that if there is no reason to judge any one of a set of mutually exclusive events to be likelier than any other, then the distribution to assign is the uniform distribution. Chernoff (1954) and Milnor (1954) provide formal justifications of the principle. Sinn (1980) shows that two of the axioms necessary for expected utility theory, the axiom of ordering and the axiom of independence, imply the principle.

Such a position seems particularly apt when it comes to characterizing a player

who is partially strategic to a very limited degree. We suppose, as is standard, that a player is completely ignorant of the strategy choices of the other players, has no past experience with the other players, and thus has no special information about which strategies the other players are choosing. Furthermore, we suppose that a player is partially strategic to a very limited degree and thus makes no inferences about how the other players might behave. Thus, it is natural to suppose that a player treats uniformly the partial profiles of strategies of the other players. Indeed, to suppose otherwise would be to ascribe to him strategic reasoning capacity or information about the strategy choices of the other players that he might not have.

Supposing a player to be partially strategic in this way has a number of advantages. Such an assumption may be a plausible characterization of typical reasoning and behavior, especially since people are in general neither nonstrategic nor fully strategic, but strategic to a very limited degree. Furthermore, as shown in Chapters 3, 4, 5, and 6, such an assumption makes for a solution concept that is mathematically and conceptually simple, exhibits nice properties, and yields the expected or reasonable solutions for a number of games of interest.

Of course, such an assumption is also an obvious point of criticism. The trouble is that a player characterized as above, while partially strategic, seems to be insufficiently strategic. In particular, even under complete ignorance, a player who is partially strategic, even to a limited degree, might be capable of reasoning fairly competently about how the other players might behave and thus judge some partial profiles of strategies of the other players to be likelier than other ones. We review these concerns in Section 2.5, and in Section 7.4, we discuss some of the problems that can arise from supposing that a player is partially strategic as characterized above. In Chapter 8, we introduce a refinement of least-squares regret that considers fully strategic players.

## 1.5 Contributions of this Dissertation

The primary contribution of this dissertation is a new solution concept for noncooperative games, one that yields the expected or reasonable solutions for games that have so far proved problematic, such as the Traveler's Dilemma, discussed in Chapter 4, and for other games of interest, such as those in Chapter 6; that is simple; that involves no *ad hoc* or game-specific constructions and can thus be applied immediately and consistently to any arbitrary game; that exhibits nice properties, such as those established in Chapter 5; and that is grounded in human psychology. The greater part of this dissertation is concerned with defining and developing least-squares regret; exploring its properties; assessing its performance

with respect to various games of interest; determining its merits and demerits, including the weaknesses discussed in Chapter 7; and introducing a refinement of least-squares regret in Chapter 8 that addresses some of the weaknesses.

Early in this chapter and in Sections 1.2 and 1.4, we presented an informal characterization of least-squares regret in order to motivate it and to fix ideas. As noted earlier, the idea behind least-squares regret is to choose a strategy so as to minimize the divergence from perfect play overall. But in order to proceed with the development and assessment of least-squares regret, it is necessary first to present a formal characterization of it. In Chapter 2, we formally define least-squares regret, study an illustrative example to show how least-squares regret is applied, consider briefly and set aside an alternative definition of least-squares regret with respect to randomized strategies, and then discuss the assumptions underlying least-squares regret and make some preliminary observations.

In certain ways, least-squares regret is not unlike some other solution concepts. Two solution concepts in particular to which least-squares regret bears a resemblance are maximin, mentioned in Section 1.1, and iterated regret minimization, mentioned in Section 1.3. Given the similarities, it is instructive to study least-squares regret in relation to each of these solution concepts and to see what sets least-squares regret apart. In Chapter 3, we compare least-squares regret with each of these solution concepts and give some reasons for preferring least-squares regret.

Early in this chapter and in Section 1.1, we mentioned as an illustration of some of the inadequacies of the standard solution concepts the notorious Traveler's Dilemma. Given its status as a puzzle with no universally accepted resolution, it is worthwhile to understand its recalcitrance. In Chapter 4, we describe and study the Traveler's Dilemma in detail. We consider the failure of the standard solution concepts to yield the reasonable solution that is supported by both intuition and the experimental evidence. After describing some alternative analyses, we show how least-squares regret resolves the puzzle.

Part of evaluating a solution concept involves determining the properties that it exhibits and the criteria that it satisfies. Ideally, a solution concept should exhibit nice and desirable properties and satisfy the relevant criteria. In Chapter 5, we establish some notable mathematical and conceptual properties of least-squares regret, including existence; invariance with respect to full equivalence; the relationship between least-squares regret and dominated strategies; the differences between least-squares regret and iterative elimination of dominated strategies; the relationship between least-squares regret and uniformly dominant strategies; invariance with respect to certain well-known transformations of the payoffs in a game that leave unchanged the best-response correspondences of the players;

the equivalence of least-squares regret and risk dominance when it comes to the equilibria in pure strategies of a  $2 \times 2$  game; convexity of the set of solutions of a game; and uniqueness of a solution under specific conditions.

Another part of evaluating a solution concept involves determining whether it yields the expected or reasonable solutions for the games of interest. A solution concept that failed to yield the expected or reasonable solutions would not be considered successful no matter how superior it might be in other respects. In Chapter 6, we apply least-squares regret to a number of well-known games, in particular, the Dollar Auction; Bertrand competition; inspection games; Matching Pennies; Chicken; coordination games; Battle of the Sexes; and the two-person bargaining problem. We show how, with respect to many of these games, least-squares regret yields reasonable solutions in line with intuition and the experimental evidence and outperforms standard solution concepts such as Nash equilibrium and how, with respect to some of these games, least-squares regret yields unsatisfactory solutions.

All solution concepts have their weaknesses, and least-squares regret is no different. As shown in Chapter 6, for some games, least-squares regret yields unsatisfactory solutions. Furthermore, least-squares regret has some other weaknesses not unlike those afflicting the solution concepts discussed in Sections 1.1 and 1.3. In Chapter 7, we consider some of the weaknesses of least-squares regret, including its failure to satisfy the principle of Independence of Irrelevant Alternatives, its divergence from maximin and Nash equilibrium with respect to two-person zero-sum games, its susceptibility to framing effects, and the problems that arise from supposing that a player is partially strategic in the sense discussed in Sections 1.2, 1.4, 2.2, and 2.5.

As noted throughout this dissertation, least-squares regret considers partially strategic players. Supposing a player to be partially strategic has many advantages, as the greater part of this dissertation illustrates. But as discussed in Sections 1.2, 1.4, 2.2, 2.5, and 7.4, such an assumption has its limitations. In particular, such an assumption can yield unsatisfactory solutions and, being oversimple, may be an implausible characterization of typical reasoning and behavior, which may be more sophisticated. Recognizing the importance of relaxing this assumption and addressing its limitations, we introduce in Chapter 8 a refinement of least-squares regret, which we call *mutual weighted least-squares regret*, that considers fully strategic players. This refinement may be a fruitful alternative to least-squares regret, especially for games in which a player is likely to be significantly, if not fully, strategic. Much of Chapter 8 is devoted to developing mutual weighted least-squares regret, studying an illustrative example to show how mutual weighted least-squares regret is applied, proving a general existence theorem, comparing

mutual weighted least-squares regret with Nash equilibrium, and understanding whether computation of a solution can be reduced to a recursive process.

We view this dissertation essentially as an exploration of least-squares regret and its significance. But we recognize that room remains for further development. In Chapter 9, we conclude the dissertation by proposing some questions for further research. The proposed topics include a deeper defense of least-squares regret; a refinement of least-squares regret that considers highly strategic, but not fully strategic, players; extensions of least-squares regret to other classes of games, such as games with infinite strategy sets, Bayesian games, and games in extensive form; and applications of least-squares regret to applied areas such as mechanism design.

# 2

---

## Least-Squares Regret

In this chapter, we present a formal characterization of least-squares regret. Section 2.1 presents the notation and basic concepts used throughout this dissertation. Section 2.2 formally defines least-squares regret. Section 2.3 studies an illustrative example. Section 2.4 considers briefly and sets aside an alternative definition of least-squares regret with respect to randomized strategies. Section 2.5 discusses the assumptions underlying least-squares regret and makes some preliminary observations.

### 2.1 Notation and Basic Concepts

In this section, we present standard notation and the basic concepts of noncooperative game theory. We follow for the most part the presentation of Myerson (1991) and deviate where appropriate.

A *game in strategic form* is any  $\Gamma$  of the form

$$\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N}),$$

where  $N$  is the nonempty set of players and, for any player  $i$  in  $N$ , the nonempty set  $C_i$  is the *pure-strategy set* for player  $i$ , a *pure strategy* for player  $i$  is any  $c_i$  in  $C_i$ , and the function  $u_i: \times_{j \in N} C_j \rightarrow \mathbf{R}$  is the *utility function* for player  $i$ .

A *pure-strategy profile* is any vector  $c = (c_j)_{j \in N}$  in  $\times_{j \in N} C_j$ .

For any game  $\Gamma$  in strategic form, the game  $\Gamma$  is *finite* if and only if the set  $N$  of players is finite and, for every player  $i$  in  $N$ , the pure-strategy set  $C_i$  is finite. In the interest of tractability, we restrict attention to finite games in strategic form.

For any player  $i$  in  $N$ , let  $N - i$  be the set such that

$$N - i = N \setminus \{i\},$$

and let  $C_{-i}$  be the set of partial profiles of pure strategies of the other players, that is,

$$C_{-i} = \prod_{j \in N-i} C_j.$$

For any player  $i$  in  $N$ , any partial profile  $c_{-i} = (c_j)_{j \in N-i}$  in  $C_{-i}$ , and any pure strategy  $c_i$  in  $C_i$ , let  $(c_{-i}, c_i)$  be the pure-strategy profile in  $\prod_{j \in N} C_j$  such that the  $i$ -component is  $c_i$  and all other components are as in  $c_{-i}$ .

For any finite set  $Z$ , let  $\Delta(Z)$  be the set of probability distributions over the set  $Z$ , that is,

$$\Delta(Z) = \{q: Z \rightarrow \mathbf{R} \mid \sum_{y \in Z} q(y) = 1 \text{ and } q(z) \geq 0, \quad \forall z \in Z\}.$$

For any player  $i$  in  $N$ , the set  $\Delta(C_i)$  is the *randomized-strategy set* for player  $i$ , and a *randomized strategy* for player  $i$  is any  $\sigma_i = (\sigma_i(c_i))_{c_i \in C_i}$  in  $\Delta(C_i)$ .

A *randomized-strategy profile* is any vector  $\sigma = (\sigma_j)_{j \in N}$  in  $\prod_{j \in N} \Delta(C_j)$ .

For any player  $i$  in  $N$ , any partial profile  $\sigma_{-i} = (\sigma_j)_{j \in N-i}$  in  $\prod_{j \in N-i} \Delta(C_j)$ , and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , let  $(\sigma_{-i}, \sigma_i)$  be the randomized-strategy profile such that the  $i$ -component is  $\sigma_i$  and all other components are as in  $\sigma_{-i}$ .

As is standard, we suppose that the players choose their pure strategies independently. For any randomized-strategy profile  $\sigma = (\sigma_j)_{j \in N}$  in  $\prod_{j \in N} \Delta(C_j)$  and any pure-strategy profile  $c = (c_j)_{j \in N}$  in  $\prod_{j \in N} C_j$ , the probability that  $c$  obtains in play is just

$$\prod_{j \in N} \sigma_j(c_j).$$

For any player  $i$  in  $N$ , the utility function  $u_i: \prod_{j \in N} C_j \rightarrow \mathbf{R}$  is extended to the domain  $\prod_{j \in N} \Delta(C_j)$  to yield the function  $u_i: \prod_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  such that

$$u_i(\sigma) = \sum_{c \in \prod_{j \in N} C_j} \left( \prod_{j \in N} \sigma_j(c_j) \right) u_i(c).$$

For any player  $i$  in  $N$  and any pure strategy  $c_i$  in  $C_i$ , let  $\sigma_i$  be also the randomized strategy in  $\Delta(C_i)$  such that the pure strategy  $c_i$  is assigned probability 1.

Using standard linear algebra notation, for any player  $i$  in  $N$  and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , let

$$\sigma_i = \sum_{c_i \in C_i} \sigma_i(c_i) c_i.$$



## 2.2 Formal Definition

When it comes to playing a game, the essential question is which strategy to choose. One natural way to evaluate a strategy is to determine how grave an error it would be to choose it. Intuitively, the less of an error it would be to play a strategy, the better the strategy is. Thus, the idea is to choose a strategy with the minimum degree of error. What follows is essentially a characterization of the error associated with choosing a strategy.

In playing a game, a player faces other players, each of whom is likewise choosing a strategy. How well a strategy ends up performing depends on the strategies that the other players choose and, in particular, on the outcome that obtains in play. Thus, we suppose that a player is concerned with the partial profiles of strategies of the other players that can ultimately obtain in play.

Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Consider any player  $i$  in  $N$ . The set

$$C_{-i} = \prod_{j \in N-i} C_j$$

is the set of partial profiles of pure strategies of the other players. In the end, after the other players have settled on their pure strategies, exactly one partial profile will obtain in play. The trouble for player  $i$  is that he does not know the strategy choices of the other players.

The set  $C_{-i}$  is the set of partial profiles that can ultimately obtain in play. Thus, we suppose that player  $i$  is concerned with the partial profiles in  $C_{-i}$  and chooses a pure strategy accordingly.

Now, consider any pure strategy  $c_i$  in  $C_i$  and any partial profile  $c_{-i}$  in  $C_{-i}$ . If the other players play as in  $c_{-i}$ , then the payoff to player  $i$  from choosing  $c_i$  is  $u_i(c_{-i}, c_i)$ .

But  $c_i$  may or may not be a best response to  $c_{-i}$ . Thus,  $u_i(c_{-i}, c_i)$  may or may not be the best-response payoff  $\max_{d_i \in C_i} u_i(c_{-i}, d_i)$  that could be achieved with respect to  $c_{-i}$ .

The best-response payoff  $\max_{d_i \in C_i} u_i(c_{-i}, d_i)$  is notable. Since it is impossible to achieve a payoff greater than the best-response payoff, there is no use in being concerned with payoffs beyond it. It is the maximum payoff that could be achieved with perfect play with respect to  $c_{-i}$ . Thus, it represents the ideal or target or benchmark payoff with respect to  $c_{-i}$ .

If  $c_i$  is not a best response to  $c_{-i}$ , then player  $i$  errs in choosing  $c_i$  in the very trivial sense that he is not playing as well as he possibly could have with respect to  $c_{-i}$ . By the informal notion of error here, we mean simply a divergence from perfect play.

We can quantify the extent to which a pure strategy diverges from perfect play with respect to a partial profile. For any player  $i$  in  $N$ , any partial profile  $c_{-i}$  in  $C_{-i}$ , and any pure strategy  $c_i$  in  $C_i$ , the *regret* of the pure strategy  $c_i$  with respect to the partial profile  $c_{-i}$  is

$$\max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, c_i).$$

Intuitively, this payoff difference is a measure of the extent to which  $c_i$  falls short of a best response with respect to  $c_{-i}$ . It can be interpreted as the regret from choosing  $c_i$  with respect to  $c_{-i}$ .

Just as the payoff of a pure strategy with respect to a partial profile can vary depending on the partial profile, so can the regret of a pure strategy with respect to a partial profile. Thus, it is necessary to compute the regret of a pure strategy with respect to each partial profile.

What is of interest ultimately is a measure of the regret of a pure strategy overall. But since the regret of a pure strategy with respect to a partial profile can vary depending on the partial profile and since it is impossible to know which partial profile will obtain in play, it is necessary to evaluate a pure strategy with respect to all of the partial profiles considered at once.

We suppose the squaring of regret for technical reasons and for mathematical convenience. But it is worth noting that the squaring of regret amounts to supposing that a larger regret is far more significant in an economic or psychological sense. We discuss the squaring of regret in Sections 2.5 and 3.2.

We suppose also that a player is partially strategic in the sense described in Sections 1.2 and 1.4. In particular, we suppose that a player treats uniformly the partial profiles of pure strategies of the other players. This assumption is an obvious point of criticism. We discuss it in Sections 2.5 and 7.4, and in Chapter 8, we introduce a refinement of least-squares regret that considers fully strategic players.

Furthermore, we suppose that a player computes the regret of a pure strategy by taking the sum, across all partial profiles of pure strategies of the other players, of the squares of the regrets.

For any player  $i$  in  $N$ , let  $\rho_i: C_i \rightarrow \mathbf{R}$  be the *regret function in pure strategies* for player  $i$  such that

$$\rho_i(c_i) = \sum_{c_{-i} \in C_{-i}} \left( \max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, c_i) \right)^2.$$

Intuitively, for any pure strategy  $c_i$  in  $C_i$ , the value  $\rho_i(c_i)$  is the regret from choosing  $c_i$ .

It is straightforward to extend the definitions above to incorporate randomized strategies. We proceed in the natural way.

Consider any player  $i$  in  $N$ . The set  $\times_{j \in N-i} \Delta(C_j)$  is the set of partial profiles of randomized strategies of the other players. In the end, after the other players have settled on their randomized strategies, exactly one partial profile will obtain in play. The trouble for player  $i$  is that he does not know the strategy choices of the other players.

But in a sense, the partial profiles in  $\times_{j \in N-i} \Delta(C_j)$  are not really the ones that can ultimately obtain in play. While each of the other players may choose a randomized strategy, each player will randomize and end up choosing a pure strategy. Thus, the outcome that obtains after everything is settled will be some partial profile in  $C_{-i}$ .

The set  $C_{-i}$  is the set of partial profiles that can ultimately obtain in play. Thus, we suppose that player  $i$  is concerned with the partial profiles in  $C_{-i}$  and chooses a randomized strategy accordingly.

Now, consider any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$  and any partial profile  $c_{-i}$  in  $C_{-i}$ . If the other players play as in  $c_{-i}$ , then the payoff to player  $i$  from choosing  $\sigma_i$  is  $u_i(c_{-i}, \sigma_i)$ . (Recall the definition from earlier: for any player  $i$  in  $N$  and any pure strategy  $c_i$  in  $C_i$ , let  $c_i$  be also the randomized strategy in  $\Delta(C_i)$  such that the pure strategy  $c_i$  is assigned probability 1.) Clearly,

$$u_i(c_{-i}, \sigma_i) = \sum_{c_i \in C_i} \sigma_i(c_i) u_i(c_{-i}, c_i), \quad \forall i \in N, \quad \forall \sigma_i \in \Delta(C_i), \quad \forall c_{-i} \in C_{-i}.$$

Intuitively, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$  and any partial profile  $c_{-i}$  in  $C_{-i}$ , the payoff  $u_i(c_{-i}, \sigma_i)$  is the expected payoff to player  $i$  from choosing  $\sigma_i$  with respect to  $c_{-i}$ .

But  $\sigma_i$  may or may not be a best response to  $c_{-i}$ . Thus,  $u_i(c_{-i}, \sigma_i)$  may or may not be the best-response payoff  $\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i)$  that could be achieved with respect to  $c_{-i}$ .

We note in passing that, clearly,

$$\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) = \max_{d_i \in C_i} u_i(c_{-i}, d_i), \quad \forall i \in N, \quad \forall c_{-i} \in C_{-i}.$$

Still, we prefer in this context to specify randomized strategies as opposed to pure strategies for the sake of symmetry and clarity and to make explicit, especially in what follows, that randomized strategies are being considered.

The best-response payoff  $\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i)$  is notable. Since it is impossible to achieve a payoff greater than the best-response payoff, there is no use in being concerned with payoffs beyond it. It is the maximum payoff that could be achieved with perfect play with respect to  $c_{-i}$ . Thus, it represents the ideal or target or benchmark payoff with respect to  $c_{-i}$ .

If  $\sigma_i$  is not a best response to  $c_{-i}$ , then player  $i$  errs in choosing  $\sigma_i$  in the very trivial sense that he is not playing as well as he possibly could have with respect to  $c_{-i}$ . Again, by the informal notion of error here, we mean simply a divergence from perfect play.

We can quantify the extent to which a randomized strategy diverges from perfect play with respect to a partial profile. For any player  $i$  in  $N$ , any partial profile  $c_{-i}$  in  $C_{-i}$ , and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the *regret* of the randomized strategy  $\sigma_i$  with respect to the partial profile  $c_{-i}$  is

$$\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i).$$

Intuitively, this payoff difference is a measure of the extent to which  $\sigma_i$  falls short of a best response with respect to  $c_{-i}$ . It can be interpreted as the regret from choosing  $\sigma_i$  with respect to  $c_{-i}$ .

Just as the payoff of a randomized strategy with respect to a partial profile can vary depending on the partial profile, so can the regret of a randomized strategy with respect to a partial profile. Thus, it is necessary to compute the regret of a randomized strategy with respect to each partial profile.

What is of interest ultimately is a measure of the regret of a randomized strategy overall. But since the regret of a randomized strategy with respect to a partial profile can vary depending on the partial profile and since it is impossible to know which partial profile will obtain in play, it is necessary to evaluate a randomized strategy with respect to all of the partial profiles considered at once.

We appeal to the assumptions and constructions given in the case of pure strategies and make adjustments as needed. In particular, we suppose the squaring of regret; that a player is partially strategic in the sense described in Sections 1.2 and 1.4; that a player treats uniformly the partial profiles of pure strategies of the other players; and that a player computes the regret of a randomized strategy by taking the sum, across all partial profiles of pure strategies of the other players, of the squares of the regrets. As noted earlier, we discuss these assumptions in Sections 2.5, 3.2, and 7.4, and in Chapter 8, we introduce a refinement of least-squares regret that considers fully strategic players.

For any player  $i$  in  $N$ , let  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the *regret function in randomized strategies* for player  $i$  such that

$$\rho_i(\sigma_i) = \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right)^2.$$

Intuitively, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the value  $\rho_i(\sigma_i)$  is the regret from choosing  $\sigma_i$ .

The regret of a strategy is a measure of the divergence from the best-response payoffs of the payoffs that the strategy could yield and is thus a measure of the extent to which the strategy diverges from perfect play overall. By choosing a strategy with the minimum regret, a player minimizes the divergence and gets as close as possible to perfect play overall and thus to securing the best-response payoff whatever the other players might do.

As discussed in Sections 1.2 and 1.4, we suppose that a player is partially strategic in the sense that, given his rudimentary assumptions about the other players, he responds accordingly. In particular, we suppose that a player chooses a strategy so as to minimize the divergence from the best-response payoffs.

If pure strategies are considered, then a player chooses a pure strategy so as to minimize the regret function in pure strategies. For any player  $i$  in  $N$  and any pure strategy  $c_i$  in  $C_i$ , the pure strategy  $c_i$  is a *pure least-squares regret strategy* for player  $i$  if and only if

$$\rho_i(c_i) \leq \rho_i(d_i), \quad \forall d_i \in C_i.$$

For any pure-strategy profile  $c = (c_j)_{j \in N}$  in  $\times_{j \in N} C_j$ , the pure-strategy profile  $c$  is a *least-squares regret profile in pure strategies* of  $\Gamma$  if and only if

$$\rho_i(c_i) \leq \rho_i(d_i), \quad \forall i \in N, \quad \forall d_i \in C_i.$$

Thus, in a least-squares regret profile in pure strategies, each player chooses a pure least-squares regret strategy.

If randomized strategies are considered, then a player chooses a randomized strategy so as to minimize the regret function in randomized strategies. For any player  $i$  in  $N$  and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the randomized strategy  $\sigma_i$  is a *randomized least-squares regret strategy* for player  $i$  if and only if

$$\rho_i(\sigma_i) \leq \rho_i(\tau_i), \quad \forall \tau_i \in \Delta(C_i).$$

For any randomized-strategy profile  $\sigma = (\sigma_j)_{j \in N}$  in  $\times_{j \in N} \Delta(C_j)$ , the randomized-strategy profile  $\sigma$  is a *least-squares regret profile in randomized strategies* of  $\Gamma$  if and only if

$$\rho_i(\sigma_i) \leq \rho_i(\tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i).$$

Thus, in a least-squares regret profile in randomized strategies, each player chooses a randomized least-squares regret strategy.

### 2.3 An Example

For an illustration of least-squares regret, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 2.1.

Table 2.1 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	2
	$y_1$	1	0

Consider first least-squares regret with respect to pure strategies. As discussed in Section 2.2, each player chooses a pure least-squares regret strategy.

Consider player 1. Notice that  $x_1$  is a best response if player 2 chooses  $x_2$  and falls short if player 2 chooses  $y_2$  and that  $y_1$  falls short if player 2 chooses  $x_2$  and is a best response if player 2 chooses  $y_2$ . Whatever player 1 chooses, he risks playing imperfectly depending on what player 2 chooses. Thus, player 1 chooses a pure strategy so as to minimize the regret function  $\rho_1$  in pure strategies.

Consider  $x_1$ . If player 2 chooses  $x_2$ , then the payoff to player 1 is 3 while the best-response payoff is 3, and so, the regret is  $3 - 3 = 0$ . If player 2 chooses  $y_2$ , then the payoff to player 1 is 0 while the best-response payoff is 1, and so, the regret is  $1 - 0 = 1$ . The regret of  $x_1$  is

$$\begin{aligned}\rho_1(x_1) &= \left( \max_{d_1 \in C_1} u_1(d_1, x_2) - u_1(x_1, x_2) \right)^2 + \left( \max_{d_1 \in C_1} u_1(d_1, y_2) - u_1(x_1, y_2) \right)^2 \\ &= (3 - 3)^2 + (1 - 0)^2 \\ &= 1.\end{aligned}$$

Consider  $y_1$ . If player 2 chooses  $x_2$ , then the payoff to player 1 is 0 while the best-response payoff is 3, and so, the regret is  $3 - 0 = 3$ . If player 2 chooses  $y_2$ , then the payoff to player 1 is 1 while the best-response payoff is 1, and so, the regret is  $1 - 1 = 0$ . The regret of  $y_1$  is

$$\begin{aligned}\rho_1(y_1) &= \left( \max_{d_1 \in C_1} u_1(d_1, x_2) - u_1(y_1, x_2) \right)^2 + \left( \max_{d_1 \in C_1} u_1(d_1, y_2) - u_1(y_1, y_2) \right)^2 \\ &= (3 - 0)^2 + (1 - 1)^2 \\ &= 9.\end{aligned}$$

The regret of  $x_1$  is less than the regret of  $y_1$ , and so, the unique pure least-squares regret strategy for player 1 is  $x_1$ .

Consider player 2. Notice that  $x_2$  falls short if player 1 chooses  $x_1$  and is a best response if player 1 chooses  $y_1$  and that  $y_2$  is a best response if player 1 chooses  $x_1$  and falls short if player 1 chooses  $y_1$ . Whatever player 2 chooses, he risks playing

imperfectly depending on what player 1 chooses. Thus, player 2 chooses a pure strategy so as to minimize the regret function  $\rho_2$  in pure strategies.

Consider  $x_2$ . If player 1 chooses  $x_1$ , then the payoff to player 2 is 0 while the best-response payoff is 2, and so, the regret is  $2 - 0 = 2$ . If player 1 chooses  $y_1$ , then the payoff to player 2 is 1 while the best-response payoff is 1, and so, the regret is  $1 - 1 = 0$ . The regret of  $x_2$  is

$$\begin{aligned}\rho_2(x_2) &= \left( \max_{d_2 \in C_2} u_2(x_1, d_2) - u_2(x_1, x_2) \right)^2 + \left( \max_{d_2 \in C_2} u_2(y_1, d_2) - u_2(y_1, x_2) \right)^2 \\ &= (2 - 0)^2 + (1 - 1)^2 \\ &= 4.\end{aligned}$$

Consider  $y_2$ . If player 1 chooses  $x_1$ , then the payoff to player 2 is 2 while the best-response payoff is 2, and so, the regret is  $2 - 2 = 0$ . If player 1 chooses  $y_1$ , then the payoff to player 2 is 0 while the best-response payoff is 1, and so, the regret is  $1 - 0 = 1$ . The regret of  $y_2$  is

$$\begin{aligned}\rho_2(y_2) &= \left( \max_{d_2 \in C_2} u_2(x_1, d_2) - u_2(x_1, y_2) \right)^2 + \left( \max_{d_2 \in C_2} u_2(y_1, d_2) - u_2(y_1, y_2) \right)^2 \\ &= (2 - 2)^2 + (1 - 0)^2 \\ &= 1.\end{aligned}$$

The regret of  $y_2$  is less than the regret of  $x_2$ , and so, the unique pure least-squares regret strategy for player 2 is  $y_2$ .

Thus, the unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ , which gives the payoff allocation  $(0, 2)$ .

Now, consider least-squares regret with respect to randomized strategies. As discussed in Section 2.2, each player chooses a randomized least-squares regret strategy.

Consider player 1. Notice that whatever player 1 chooses, he risks playing imperfectly depending on what player 2 chooses. Thus, player 1 chooses a randomized strategy so as to minimize the regret function  $\rho_1$  in randomized strategies.

Consider any randomized strategy  $\sigma_1$  in  $\Delta(C_1)$ . If player 2 chooses  $x_2$ , then the payoff to player 1 is  $3\sigma_1(x_1)$  while the best-response payoff is 3, and so, the regret is  $3 - 3\sigma_1(x_1)$ . If player 2 chooses  $y_2$ , then the payoff to player 1 is  $1 - \sigma_1(x_1)$  while the best-response payoff is 1, and so, the regret is  $1 - (1 - \sigma_1(x_1))$ . The regret

function  $\rho_1: \Delta(C_1) \rightarrow \mathbf{R}$  is

$$\begin{aligned}\rho_1(\sigma_1) &= \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, x_2) - u_1(\sigma_1, x_2) \right)^2 + \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, y_2) - u_1(\sigma_1, y_2) \right)^2 \\ &= (3 - 3\sigma_1(x_1))^2 + (1 - (1 - \sigma_1(x_1)))^2 \\ &= 10(\sigma_1(x_1))^2 - 18\sigma_1(x_1) + 9,\end{aligned}$$

and so,

$$\frac{d(\rho_1(\sigma_1))}{d\sigma_1(x_1)} = 20\sigma_1(x_1) - 18.$$

The regret function  $\rho_1$  is minimized at the point  $\sigma_1(x_1) = 0.9$ , and so, the unique randomized least-squares regret strategy for player 1 is  $0.9x_1 + 0.1y_1$ .

Consider player 2. Notice that whatever player 2 chooses, he risks playing imperfectly depending on what player 1 chooses. Thus, player 2 chooses a randomized strategy so as to minimize the regret function  $\rho_2$  in randomized strategies.

Consider any randomized strategy  $\sigma_2$  in  $\Delta(C_2)$ . If player 1 chooses  $x_1$ , then the payoff to player 2 is  $2(1 - \sigma_2(x_2))$  while the best-response payoff is 2, and so, the regret is  $2 - 2(1 - \sigma_2(x_2))$ . If player 1 chooses  $y_1$ , then the payoff to player 2 is  $\sigma_2(x_2)$  while the best-response payoff is 1, and so, the regret is  $1 - \sigma_2(x_2)$ . The regret function  $\rho_2: \Delta(C_2) \rightarrow \mathbf{R}$  is

$$\begin{aligned}\rho_2(\sigma_2) &= \left( \max_{\tau_2 \in \Delta(C_2)} u_2(x_1, \tau_2) - u_2(x_1, \sigma_2) \right)^2 + \left( \max_{\tau_2 \in \Delta(C_2)} u_2(y_1, \tau_2) - u_2(y_1, \sigma_2) \right)^2 \\ &= (2 - 2(1 - \sigma_2(x_2)))^2 + (1 - \sigma_2(x_2))^2 \\ &= 5(\sigma_2(x_2))^2 - 2\sigma_2(x_2) + 1,\end{aligned}$$

and so,

$$\frac{d(\rho_2(\sigma_2))}{d\sigma_2(x_2)} = 10\sigma_2(x_2) - 2.$$

The regret function  $\rho_2$  is minimized at the point  $\sigma_2(x_2) = 0.2$ , and so, the unique randomized least-squares regret strategy for player 2 is  $0.2x_2 + 0.8y_2$ .

Thus, the unique least-squares regret profile in randomized strategies is

$$(0.9x_1 + 0.1y_1, 0.2x_2 + 0.8y_2),$$

which gives the payoff allocation  $(0.62, 1.46)$ .

## 2.4 An Alternative Definition

In this section, in order to clarify and affirm the established definition of the regret function in randomized strategies given in Section 2.2, we consider briefly and set aside an alternative definition.



Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Now, recall the established definition. For any player  $i$  in  $N$ , the regret function in randomized strategies for player  $i$  is the function  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  such that

$$\begin{aligned} \rho_i(\sigma_i) &= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right)^2 \\ &= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - \sum_{c_i \in C_i} \sigma_i(c_i) u_i(c_{-i}, c_i) \right)^2. \end{aligned}$$

Intuitively, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the value  $\rho_i(\sigma_i)$  is the sum, across all partial profiles of pure strategies of the other players, of the squares of the regrets, where the regret of  $\sigma_i$  with respect to a partial profile is the difference of the best-response payoff with respect to the partial profile and the expected payoff from choosing  $\sigma_i$  with respect to the partial profile.

According to the established definition, as discussed in Section 2.2, a player chooses a randomized strategy so as to minimize the regret function in randomized strategies.

Now, consider the following alternative definition. For any player  $i$  in  $N$ , let  $\zeta_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the alternative regret function in randomized strategies for player  $i$  such that

$$\zeta_i(\sigma_i) = \sum_{c_{-i} \in C_{-i}} \sum_{c_i \in C_i} \sigma_i(c_i) \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, c_i) \right)^2.$$

Intuitively, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the value  $\zeta_i(\sigma_i)$  is the sum, across all partial profiles of pure strategies of the other players, of the expected values of the squares of the regrets, where the expected value induced by  $\sigma_i$  with respect to a partial profile is the probability-weighted average with respect to  $\sigma_i$  of the squares of all possible regrets with respect to the partial profile.

According to the alternative definition just described, a player chooses a randomized strategy so as to minimize the alternative regret function in randomized strategies.

For an illustration of some of the differences between the two definitions, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 2.2.

Consider first the alternative definition. Consider player 1 and any randomized strategy  $\sigma_1$  in  $\Delta(C_1)$ . Given  $\sigma_1$ , player 1 chooses  $x_1$  with probability  $\sigma_1(x_1)$  and  $y_1$  with probability  $\sigma_1(y_1) = 1 - \sigma_1(x_1)$ . If player 2 chooses  $x_2$ , then the regret from choosing  $x_1$  is  $3 - 3 = 0$ , and the regret from choosing  $y_1$  is  $3 - 0 = 3$ . If player 2 chooses  $y_2$ , then the regret from choosing  $x_1$  is  $1 - 0 = 1$ , and the regret from

Table 2.2 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	3	0
	$y_1$	0	1

choosing  $y_1$  is  $1 - 1 = 0$ . The alternative regret function  $\zeta_1: \Delta(C_1) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \zeta_1(\sigma_1) &= \sum_{c_2 \in C_2} \sum_{c_1 \in C_1} \sigma_1(c_1) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, c_2) - u_1(c_1, c_2) \right)^2 \\ &= \sigma_1(x_1)(3 - 3)^2 + (1 - \sigma_1(x_1))(3 - 0)^2 + \sigma_1(x_1)(1 - 0)^2 + (1 - \sigma_1(x_1))(1 - 1)^2 \\ &= 9 - 8\sigma_1(x_1). \end{aligned}$$

Thus, the unique randomized strategy that minimizes the alternative regret function  $\zeta_1$  is  $x_1$ , which assigns probability 1 to the pure strategy  $x_1$ .

Notably, for many games, the alternative definition yields solutions that involve no randomization at all, even when some randomization, which can be seen as a method of hedging, might be expected or reasonable.

Now, consider the established definition. The unique randomized least-squares regret strategy for player 1 is  $0.9x_1 + 0.1y_1$ . Thus, minimizing the regret function  $\rho_1$  involves randomizing between  $x_1$  and  $y_1$  and not playing either with probability 1.

Indeed, for many games, the established definition yields solutions that involve some randomization, especially when randomization, interpreted as a method of hedging, might be expected or reasonable.

For another illustration of some of the differences between the two definitions, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 2.3.

Table 2.3 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1	0
	$y_1$	0	1

Consider first the alternative definition. Consider player 1 and any randomized strategy  $\sigma_1$  in  $\Delta(C_1)$ . Given  $\sigma_1$ , player 1 chooses  $x_1$  with probability  $\sigma_1(x_1)$  and  $y_1$  with probability  $\sigma_1(y_1) = 1 - \sigma_1(x_1)$ . If player 2 chooses  $x_2$ , then the regret from choosing  $x_1$  is  $1 - 1 = 0$ , and the regret from choosing  $y_1$  is  $1 - 0 = 1$ . If player 2 chooses  $y_2$ , then the regret from choosing  $x_1$  is  $1 - 0 = 1$ , and the regret from choosing  $y_1$  is  $1 - 1 = 0$ . The alternative regret function  $\zeta_1: \Delta(C_1) \rightarrow \mathbf{R}$  is

$$\begin{aligned}\zeta_1(\sigma_1) &= \sum_{c_2 \in C_2} \sum_{c_1 \in C_1} \sigma_1(c_1) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, c_2) - u_1(c_1, c_2) \right)^2 \\ &= \sigma_1(x_1)(1 - 1)^2 + (1 - \sigma_1(x_1))(1 - 0)^2 + \sigma_1(x_1)(1 - 0)^2 + (1 - \sigma_1(x_1))(1 - 1)^2 \\ &= 1.\end{aligned}$$

Thus, any randomization between  $x_1$  and  $y_1$  minimizes the alternative regret function  $\zeta_1$ .

Notably, for many games, the alternative definition yields infinitely many solutions, even when uniqueness might be expected or reasonable.

Now, consider the established definition. The unique randomized least-squares regret strategy for player 1 is  $0.5x_1 + 0.5y_1$ . Thus, minimizing the regret function  $\rho_1$  involves randomizing uniformly between  $x_1$  and  $y_1$ .

Indeed, for many games, the established definition yields unique solutions, especially when uniqueness might be expected or reasonable. For two theorems that describe sufficient conditions for the uniqueness of a solution yielded by the established definition, see Section 5.9.

For further illustrations of the established definition reflecting hedging via randomization and yielding randomized or unique solutions in line with intuition and the experimental evidence, see the examples in Sections 6.3, 6.4, 6.5, 6.6, and 6.7.

The alternative definition may well have its merits. Nevertheless, we note several reasons to prefer the established definition.

Unlike the established definition, which involves convex combinations of the payoffs in a game, the alternative definition involves convex combinations of the squares of the regrets. Consequently, in contrast to the established definition, the alternative definition is conceptually somewhat unnatural; structurally unlike the definition of the regret function in pure strategies; mathematically inconvenient in certain ways; deprived of the natural and intuitive interpretations discussed in Section 2.5 and of certain of the nice properties established in Chapter 5; and, as the examples illustrate, incapable of yielding the expected or reasonable solutions for a number of games, such as those discussed in Chapter 6.

## 2.5 Discussion

As discussed in Sections 1.2, 1.4, and 2.2, least-squares regret involves a number of assumptions. In this section, we discuss these assumptions and make some preliminary observations.

One preliminary observation is that whether or not randomized strategies are considered can matter greatly. For an illustration of this point, consider again the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 2.1 in Section 2.3 and reproduced in Table 2.4.

Table 2.4 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0    2	3    0
	$y_1$	1    0	0    1

As noted in Section 2.3, in this game, for player 1, the unique pure least-squares regret strategy is  $x_1$ , and the unique randomized least-squares regret strategy is  $0.9x_1 + 0.1y_1$ , and for player 2, the unique pure least-squares regret strategy is  $y_2$ , and the unique randomized least-squares regret strategy is  $0.2x_2 + 0.8y_2$ . Thus, for a player, the set of pure least-squares regret strategies and the set of randomized least-squares regret strategies need not coincide, and so, for a game, the set of least-squares regret profiles in pure strategies and the set of least-squares regret profiles in randomized strategies need not coincide.

Furthermore, as the example illustrates, the support of a randomized least-squares regret strategy may contain pure strategies that are not pure least-squares regret strategies. As just observed, the unique randomized least-squares regret strategy for player 1 assigns positive probability to  $y_1$ , which is not a pure least-squares regret strategy for player 1, and the unique randomized least-squares regret strategy for player 2 assigns positive probability to  $x_2$ , which is not a pure least-squares regret strategy for player 2.

In general, as will be evident throughout, it is important to be clear about whether or not randomized strategies are considered.

As noted in Sections 1.2 and 2.2, we suppose the squaring of regret for technical reasons and for mathematical convenience. Squaring the regret simplifies the mathematics and yields nice mathematical properties. For example, as shown in Section 5.8, one agreeable consequence of squaring the regret is that the regret function in randomized strategies is convex. Convexity of the regret function

in randomized strategies implies that computation of randomized least-squares regret strategies is a convex optimization problem, which can be solved efficiently using standard well-developed techniques. More generally, squaring the regret allows for computations that are elementary, straightforward, and easily executed using established techniques.

As mentioned in Section 2.2, the squaring of regret amounts to supposing that a larger regret is far more significant in an economic or psychological sense. In this way, the squaring of regret can be thought of as formalizing a cognitive pattern along the lines of, say, loss aversion (Kahneman and Tversky, 1979; Kahneman and Tversky, 1984; Tversky and Kahneman, 1991). In particular, what the squaring of regret formalizes is an attitude toward regret. We discuss this attitude and compare it with an alternative attitude with respect to particular games in Sections 3.2 and 6.8.

Furthermore, squaring the regret and assuming that partial profiles are treated uniformly allow for a natural and intuitive geometric interpretation of least-squares regret. Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Consider any player  $i$  in  $N$ . The best-response payoffs form a payoff vector  $(\max_{d_i \in C_i} u_i(c_{-i}, d_i))_{c_{-i} \in C_{-i}} = (\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i))_{c_{-i} \in C_{-i}}$  in  $\mathbf{R}^{|C_{-i}|}$ . Now, consider any pure strategy  $c_i$  in  $C_i$  and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ . Notice that  $c_i$  yields the payoff vector  $(u_i(c_{-i}, c_i))_{c_{-i} \in C_{-i}}$  in  $\mathbf{R}^{|C_{-i}|}$  and that  $\sigma_i$  yields the payoff vector  $(u_i(c_{-i}, \sigma_i))_{c_{-i} \in C_{-i}}$  in  $\mathbf{R}^{|C_{-i}|}$ . The regret function  $\rho_i: C_i \rightarrow \mathbf{R}$  can be interpreted as a distance function that specifies, for any pure strategy  $c_i$  in  $C_i$ , the squared Euclidean distance from the payoff vector yielded by  $c_i$  to the best-response payoff vector. The regret function  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  can be interpreted as a distance function that specifies, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the squared Euclidean distance from the payoff vector yielded by  $\sigma_i$  to the best-response payoff vector. Thus, there is a very natural and concrete sense in which minimizing regret minimizes the divergence from the best-response payoffs.

Squaring the regret and assuming that partial profiles are treated uniformly allow also for an interpretation of least-squares regret in terms of the standard ordinary least-squares method of parameter estimation. For a fuller discussion on this statistical method, see any standard econometrics text, for example, Greene (2011) or Wooldridge (2013), on which the following discussion is based.

For an illustration, consider simple linear regression. Suppose that  $y$  and  $x$  are two variables characterizing a population and that the aim is to determine how  $y$  depends on  $x$ . A *simple regression model* is any equation of the form

$$y = \beta_0 + \beta_1 x + u,$$

where the constant  $\beta_0$  is the *intercept parameter*, the constant  $\beta_1$  is the *slope param-*

eter, and the variable  $u$  is the *error term* representing all explanatory factors other than  $x$  that affect  $y$ .

The aim is to derive an estimate  $\hat{\beta}_0$  of the intercept parameter  $\beta_0$  and an estimate  $\hat{\beta}_1$  of the slope parameter  $\beta_1$  using a sample drawn from the population. Let  $N = \{1, \dots, n\}$  be the nonempty set of observations, and let

$$\{(x_i, y_i) \in \mathbf{R}^2 \mid i \in N\}$$

be a random sample drawn from the population. Since the sample is drawn from the population, which is characterized by the simple regression model, it follows that

$$y_i = \beta_0 + \beta_1 x_i + u_i, \quad \forall i \in N,$$

where, for any observation  $i$  in  $N$ , the variable  $u_i$  is the error term for observation  $i$  representing all explanatory factors other than  $x_i$  that affect  $y_i$ .

For any estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  and any observation  $i$  in  $N$ , the *fitted value*  $\hat{y}_i$  for observation  $i$  with respect to the estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  is the number such that

$$\hat{y}_i = \hat{\beta}_0 + \hat{\beta}_1 x_i,$$

and the *residual*  $\hat{u}_i$  for observation  $i$  with respect to the estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  is the difference of the actual value  $y_i$  and the fitted value  $\hat{y}_i$ , that is,

$$\hat{u}_i = y_i - \hat{y}_i = y_i - \hat{\beta}_0 - \hat{\beta}_1 x_i.$$

For any estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$ , the *sum of the squared residuals* with respect to the estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  is the sum

$$\sum_{i=1}^n \hat{u}_i^2 = \sum_{i=1}^n (y_i - \hat{\beta}_0 - \hat{\beta}_1 x_i)^2.$$

The ordinary least-squares method involves choosing the estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  so as to minimize the sum of the squared residuals. An *ordinary least-squares regression line* is any equation of the form

$$\hat{y} = \hat{\beta}_0 + \hat{\beta}_1 x,$$

where the estimates  $\hat{\beta}_0$  and  $\hat{\beta}_1$  minimize the sum of the squared residuals and the variable  $\hat{y}$  is the fitted version of  $y$ . Intuitively, the ordinary least-squares regression line is the line that best approximates the sample.

As the name suggests, least-squares regret is essentially an adaptation of the ordinary least-squares method of parameter estimation. Just as the idea behind simple linear regression is to choose estimates that define a line that

best approximates the sample, the idea behind least-squares regret is to choose a strategy that defines a payoff vector that best approximates the best-response payoff vector. In the case of simple linear regression, divergence is defined in terms of the residual, and the aim is to choose estimates so as to minimize the sum of the squares of the residuals. In the case of least-squares regret, divergence is defined in terms of regret, and the aim is to choose a strategy so as to minimize the sum of the squares of the regrets. Thus, least-squares regret can be seen as a special case of the problem of ordinary least-squares, which, as is well known, is well studied and has a very complete theory, arises in a number of areas, and can be solved very efficiently using established techniques.

While we assume the squaring of regret for the reasons just given, we recognize that there may be other operations that are equally or more reasonable. And while we consider the question of alternative formulations only briefly in Section 3.2, it would be interesting and worthwhile to consider further reasons for squaring or not squaring regret.

As discussed in Sections 1.2, 1.4, and 2.2, we suppose that a player is partially strategic. In particular, we suppose that a player treats uniformly the partial profiles of pure strategies of the other players; computes the regret of a strategy by taking the sum, across all partial profiles of pure strategies of the other players, of the squares of the regrets; and chooses a strategy so as to minimize the regret function. These assumptions imply that fully strategic reasoning and information that might determine how the other players might behave—such as the payoffs of the other players—are disregarded.

As might be expected, these assumptions confer certain advantages. They make for a solution concept that is mathematically and conceptually simple and easy to apply. They imply that players can be considered separately, allowing a game to be decomposed into independent parts and making the computation of solutions trivial. They lead to reasonable solutions for a number of games, as shown in Chapters 4 and 6. And they enable the capture of certain experimentally robust effects, as shown in Sections 6.3, 6.4, and 6.6.

Of course, we recognize that the proposed assumptions can be problematic. Such assumptions may be inappropriate or too restrictive, especially if the aim is to characterize the behavior of fully strategic players capable of reasoning about one another and responding accordingly. Also, such assumptions can lead to unsatisfactory solutions for some games, as shown in Sections 6.4 and 6.7. Furthermore, such assumptions imply that unlike in, say, a Nash equilibrium, the strategies that the players choose may differ from the ones that the players expect to be chosen. Finally, as Harsanyi (1982) argues in his criticism of the decision-theoretic approach to game theory advocated by Kadane and Larkey

(1982) and Raiffa (1982) and explored further by Roth and Schoumaker (1983), disregarding fully strategic reasoning and information that might determine how the other players might behave amounts to throwing away essential information and deprives game theory of its substance and remit.

These problems merit consideration. We discuss them more fully in Chapter 7, and in Chapter 8, we introduce a refinement of least-squares regret that addresses them.



# 3

---

## Alternative Solution Concepts

Two solution concepts to which least-squares regret bears a resemblance are maximin, mentioned in Section 1.1, and iterated regret minimization, mentioned in Section 1.3.

In this chapter, we study least-squares regret in relation to each of these solution concepts. Section 3.1 compares least-squares regret with maximin. Section 3.2 compares least-squares regret with iterated regret minimization.

### 3.1 Least-Squares Regret and Maximin

One solution concept that rivals least-squares regret is maximin (von Neumann, 1928; Wald, 1939; Wald, 1945; von Neumann and Morgenstern, 1947; Wald, 1950), which involves choosing a strategy so as to maximize the minimum payoff. For an illustration of some of the differences between least-squares regret and maximin, consider again the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 2.1 in Section 2.3 and reproduced in Table 3.1.

Table 3.1 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0    2	
	$y_1$	3    0	
		1    0	
		0    1	

Consider first maximin with respect to pure strategies. For player 1, the minimum payoff from choosing  $x_1$  is 0, and the minimum payoff from choosing  $y_1$  is 0, and so, both  $x_1$  and  $y_1$  are pure maximin strategies. For player 2, the


minimum payoff from choosing  $x_2$  is 0, and the minimum payoff from choosing  $y_2$  is 0, and so, both  $x_2$  and  $y_2$  are pure maximin strategies. Thus, neither player has a favored strategy, and all outcomes are possible.

Now, consider maximin with respect to randomized strategies. As Figure 3.1 shows, for player 1, the minimum expected payoff is maximized when

$$3\sigma_1(x_1) = 1 - \sigma_1(x_1),$$

that is, at the point  $\sigma_1(x_1) = 0.25$ , and so, the unique randomized maximin strategy is  $0.25x_1 + 0.75y_1$ . As Figure 3.2 shows, for player 2, the minimum expected payoff

Figure 3.1 Expected payoffs of player 1




is maximized when

$$2(1 - \sigma_2(x_2)) = \sigma_2(x_2),$$

that is, at the point  $\sigma_2(x_2) = 2/3$ , and so, the unique randomized maximin strategy is  $(2/3)x_2 + (1/3)y_2$ . Thus, the unique profile of randomized maximin strategies is

Figure 3.2 Expected payoffs of player 2



$$(0.25x_1 + 0.75y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(3/4, 2/3)$ . Notably, player 1 favors  $y_1$ , and player 2 favors  $x_2$ .

Now, consider least-squares regret. Recall from Section 2.3 that the unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ , which gives the payoff allocation  $(0, 2)$ , and that the unique least-squares regret profile in randomized strategies is

$$(0.9x_1 + 0.1y_1, 0.2x_2 + 0.8y_2),$$

which gives the payoff allocation  $(0.62, 1.46)$ . Notably, player 1 favors  $x_1$ , and player 2 favors  $y_2$ .

That least-squares regret and maximin should yield different solutions is unsurprising. They represent different ways to reason about a game. Least-squares regret involves assessing personal payoffs in the form of regrets and choosing a strategy so as to minimize the divergence from the best-response payoffs. Maximin involves assessing personal payoffs, but not regrets, and choosing a strategy so as to maximize the minimum payoff. Notably, the exclusive focus of maximin on the minimum payoff characterizes an attitude that can be described as pessimistic or conservative.

The discrepancy between least-squares regret and maximin is particularly significant when it comes to two-person zero-sum games, for which maximin is the standard solution concept. We compare least-squares regret with maximin with respect to two-person zero-sum games in Section 7.2.

## 3.2 Least-Squares Regret and Iterated Regret Minimization

Least-squares regret is closely related to the solution concept of *iterated regret minimization* (Halpern and Pass, 2012). In this section, we compare the two solution concepts.

As the name suggests, iterated regret minimization is based on the concept of regret. Just as with least-squares regret, the regret of a strategy with respect to a partial profile of strategies of the other players is the difference of the best-response payoff with respect to the partial profile and the payoff from choosing the strategy with respect to the partial profile.

The idea behind iterated regret minimization is simple. Given any finite game in strategic form, fix some initial set of strategy profiles, typically the set of pure-strategy profiles or the set of randomized-strategy profiles. Given the initial set of strategy profiles, each player eliminates all of the strategies available to him that fail to minimize the maximum regret. The process is repeated *ad infinitum*, each time on the set of strategy profiles that remain after the previous round of elimination.

Iterated regret minimization is defined formally as follows. We use our notation where convenient to facilitate comparison and defer to Halpern and Pass (2012) for the original formulations. For the sake of economy, we define iterated regret minimization with respect to both pure strategies and randomized strategies at once.

Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$  and any set  $S = \times_{j \in N} S_j$  such that

$$S_i \subseteq C_i, \quad \forall i \in N$$

in the case of pure strategies or

$$S_i \subseteq \Delta(C_i), \quad \forall i \in N$$

in the case of randomized strategies, let  $\text{regret}_i^S: S_i \rightarrow \mathbf{R}$  be the function such that

$$\text{regret}_i^S(s_i) = \max_{s_{-i} \in S_{-i}} \left( \max_{t_i \in S_i} u_i(s_{-i}, t_i) - u_i(s_{-i}, s_i) \right),$$

where  $S_{-i} = \times_{j \in N-i} S_j$ . Intuitively, for any strategy  $s_i$  in  $S_i$  and any set  $S = \times_{j \in N} S_j$  of strategy profiles, the value  $\text{regret}_i^S(s_i)$  is the maximum regret from choosing  $s_i$  with respect to  $S$ .

For any player  $i$  in  $N$ , let  $\mathcal{RM}_i: \mathcal{P}(\times_{j \in N} C_j) \rightarrow \mathcal{P}(C_i)$  in the case of pure strategies or  $\mathcal{RM}_i: \mathcal{P}(\times_{j \in N} \Delta(C_j)) \rightarrow \mathcal{P}(\Delta(C_i))$  in the case of randomized strategies be the function such that

$$\mathcal{RM}_i(S) = \underset{t_i \in S_i}{\text{argmin}} \text{regret}_i^S(t_i).$$

Intuitively, for any set  $S = \times_{j \in N} S_j$  of strategy profiles, the value  $\mathcal{RM}_i(S)$  is the set of strategies in  $S_i$  that minimize the maximum regret function  $\text{regret}_i^S$  with respect to  $S$ . Clearly,  $\mathcal{RM}_i(S) \subseteq S_i$ .

Now, let  $\mathcal{RM}: \mathcal{P}(\times_{j \in N} C_j) \rightarrow \mathcal{P}(\times_{j \in N} C_j)$  in the case of pure strategies or  $\mathcal{RM}: \mathcal{P}(\times_{j \in N} \Delta(C_j)) \rightarrow \mathcal{P}(\times_{j \in N} \Delta(C_j))$  in the case of randomized strategies be the function such that

$$\mathcal{RM}(S) = \times_{j \in N} \mathcal{RM}_j(S).$$

Clearly,  $\mathcal{RM}(S) \subseteq S$ .

The process of iterative elimination is defined recursively as follows. For any player  $i$  in  $N$  and any set  $S = \times_{j \in N} S_j$  such that

$$S_i \subseteq C_i, \quad \forall i \in N$$

in the case of pure strategies or

$$S_i \subseteq \Delta(C_i), \quad \forall i \in N$$

in the case of randomized strategies, let

$$\mathcal{RM}_i^1(S) = \mathcal{RM}_i(S),$$

$$\mathcal{RM}_i^{k+1}(S) = \mathcal{RM}_i(\mathcal{RM}^k(S)), \quad \forall k \in \{1, 2, 3, \dots\}, \quad \text{and}$$

$$\mathcal{RM}_i^\infty(S) = \bigcap_{k=1}^{\infty} \mathcal{RM}_i^k(S).$$

Intuitively, for any set  $S = \times_{j \in N} S_j$  of strategy profiles, the value  $\mathcal{RM}_i^\infty(S)$  is the set of strategies in  $S_i$  that survive iterated regret minimization with respect to  $S$ .

For any set  $S = \times_{j \in N} S_j$  such that

$$S_i \subseteq C_i, \quad \forall i \in N$$

in the case of pure strategies or

$$S_i \subseteq \Delta(C_i), \quad \forall i \in N$$

in the case of randomized strategies, let

$$\mathcal{RM}^1(S) = \mathcal{RM}(S),$$

$$\mathcal{RM}^{k+1}(S) = \mathcal{RM}(\mathcal{RM}^k(S)), \quad \forall k \in \{1, 2, 3, \dots\}, \quad \text{and}$$

$$\mathcal{RM}^\infty(S) = \bigcap_{k=1}^{\infty} \mathcal{RM}^k(S).$$

Intuitively, for any set  $S = \times_{j \in N} S_j$  of strategy profiles, the value  $\mathcal{RM}^\infty(S)$  is the set of strategy profiles in  $S$  that survive iterated regret minimization with respect to  $S$ .

Halpern and Pass (2012) describe sufficient conditions for convergence to a nonempty fixed point.

**THEOREM 3.2.1 (HALPERN AND PASS (2012)).** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any set  $S = \times_{j \in N} S_j$  such that*

$$S_i \subseteq C_i, \quad \forall i \in N$$

*in the case of pure strategies or*

$$S_i \subseteq \Delta(C_i), \quad \forall i \in N$$

*in the case of randomized strategies, if the set  $S$  is nonempty and closed, then  $\mathcal{RM}^\infty(S)$  is nonempty and*

$$\mathcal{RM}(\mathcal{RM}^\infty(S)) = \mathcal{RM}^\infty(S).$$

In particular, if  $S = \times_{j \in N} C_j$  or  $S = \times_{j \in N} \Delta(C_j)$ , then, by Theorem 3.2.1, the set  $\mathcal{RM}^\infty(S)$  is nonempty and

$$\mathcal{RM}(\mathcal{RM}^\infty(S)) = \mathcal{RM}^\infty(S).$$

Thus, for any finite game in strategic form, a solution is guaranteed to exist.

For an illustration of iterated regret minimization, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 3.2. For the sake of simplicity in what follows, let  $C$  be the set such that

$$C = \times_{j \in N} C_j.$$

Table 3.2 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	2
	$y_1$	1	0

Consider first iterated regret minimization with respect to pure strategies. For player 1, with respect to  $C$ , the maximum regret from choosing  $x_1$  is 1, and the maximum regret from choosing  $y_1$  is 2, and so,  $x_1$  is the unique pure strategy that survives iterated regret minimization. For player 2, with respect to  $C$ , the maximum regret from choosing  $x_2$  is 2, and the maximum regret from choosing  $y_2$  is 1, and so,  $y_2$  is the unique pure strategy that survives iterated regret minimization. Thus, the unique pure-strategy profile that survives iterated regret minimization is  $(x_1, y_2)$ , which gives the payoff allocation  $(0, 2)$ .

Now, consider iterated regret minimization with respect to randomized strategies. To simplify the computations, we appeal to the following proposition, which establishes that, at the first step, only pure-strategy partial profiles need be considered.

**PROPOSITION 3.2.1 (HALPERN AND PASS (2012)).** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Let  $S = \times_{j \in N} \Delta(C_j)$ . Then*


$$\text{regret}_i^S(\sigma_i) = \max_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right), \quad \forall i \in N, \quad \forall \sigma_i \in \Delta(C_i).$$

Now, let  $S = \times_{j \in N} \Delta(C_j)$ . Consider player 1. Notice that, by Proposition 3.2.1,

$$\text{regret}_1^S(\sigma_1) = \max\{2 - 2\sigma_1(x_1), 1 - (1 - \sigma_1(x_1))\}, \quad \forall \sigma_1 \in \Delta(C_1).$$

As Figure 3.3 shows, the maximum regret function  $regret_1^S$  is minimized at the point  $\sigma_1(x_1) = 2/3$ , and so, the unique randomized strategy that survives iterated regret minimization for player 1 is  $(2/3)x_1 + (1/3)y_1$ . Consider player 2. Notice that,

Figure 3.3 Regret of player 1




by Proposition 3.2.1,

$$regret_2^S(\sigma_2) = \max\{2 - 2(1 - \sigma_2(x_2)), 1 - \sigma_2(x_2)\}, \quad \forall \sigma_2 \in \Delta(C_2).$$

As Figure 3.4 shows, the maximum regret function  $regret_2^S$  is minimized at the point  $\sigma_2(x_2) = 1/3$ , and so, the unique randomized strategy that survives iterated regret minimization for player 2 is  $(1/3)x_2 + (2/3)y_2$ . Thus, the unique randomized-

Figure 3.4 Regret of player 2



strategy profile that survives iterated regret minimization is

$$((2/3)x_1 + (1/3)y_1, (1/3)x_2 + (2/3)y_2),$$

which gives the payoff allocation  $(2/3, 1)$ .

Both least-squares regret and iterated regret minimization involve minimizing regret in one way or another. But despite the similarities between the two solution concepts, there are two differences that deserve mentioning.

One difference concerns iterative processes. Least-squares regret involves no iterative process while iterated regret minimization involves iterative elimination.

We exclude iterative elimination to forestall complications. Iterative elimination is an added complication and may be computationally and cognitively demanding, especially in more complex games, and introduces other difficulties. For example, the reason for eliminating a strategy might turn on the maximum regret potentially induced by some partial profile of strategies of the other players that is subsequently eliminated and assumed not to occur. Such difficulties recall those afflicting iterative elimination of weakly dominated strategies (Samuelson, 1992; Mas-Colell, Whinston, and Green, 1995).

Halpern and Pass (2012) consider as a remedy lexicographic belief systems reminiscent of the lexicographic probability systems of Blume, Brandenburger, and Dekel (1991) and Brandenburger, Friedenberg, and Keisler (2008). But the remedy itself is fairly involved.

In any case, whether iterative elimination is included or not is a minor difference between least-squares regret and iterated regret minimization, and we point it out mainly for the sake of completeness. Indeed, iterative elimination could easily be added to least-squares regret or removed from iterated regret minimization. But we note that by eschewing iterative elimination, we skirt entirely the complications noted above.

The more significant difference between least-squares regret and iterated regret minimization concerns the assessment of regret. Least-squares regret involves choosing a strategy so as to minimize the sum of the squares of the regrets while iterated regret minimization involves choosing a strategy so as to minimize the maximum regret. These different imperatives amount to quite different attitudes toward regret.

It is easy to see how, ignoring iterative elimination, least-squares regret and iterated regret minimization are related. Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Consider iterated regret minimization (without iterative elimination) as applied to the initial set  $S = \times_{j \in N} C_j$  of pure-strategy profiles. For any number  $p$  such that  $p \geq 1$  and any player  $i$  in  $N$ , let  $\rho_i: C_i \rightarrow \mathbf{R}$  be the function such that

$$\rho_i(c_i) = \sum_{c_{-i} \in C_{-i}} \left( \max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, c_i) \right)^p.$$

This function with the parameter  $p$  is a natural generalization of the regret function in pure strategies defined in Section 2.2. Notice that when the function is raised to the power of  $1/p$ , the result is an application of the  $p$ -norm. Intuitively, the greater is the value of  $p$ , the more larger regrets matter.



Now, notice that

$$\lim_{p \rightarrow \infty} (\rho_i(c_i))^{1/p} = \max_{c_{-i} \in C_{-i}} \left( \max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, c_i) \right), \quad \forall c_i \in C_i.$$

In the limit as  $p \rightarrow \infty$ , only the maximum regret matters. Thus, the natural generalization of least-squares regret that involves minimizing the function  $(\rho_i(c_i))^{1/p}$  subsumes iterated regret minimization (without iterative elimination) as an extreme case that is derived by letting  $p \rightarrow \infty$ .

Least-squares regret and iterated regret minimization can yield different solutions. For certain games, studied below, least-squares regret yields more agreeable solutions. The reason is that focusing exclusively on the maximum regret, as iterated regret minimization requires, may be too restrictive.

For an example in which least-squares regret outperforms iterated regret minimization, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 3.3.

Table 3.3 Payoffs of player 1 in a game in strategic form

		2				
		$v_2$	$w_2$	$x_2$	$y_2$	$z_2$
1	$u_1$	1	1	1	0	1
	$v_1$	1	1	1	1	0
	$w_1$	1	1	1	0	0
	$x_1$	1	1	0	0	0
	$y_1$	1	0	0	0	0
	$z_1$	0	0	0	0	0

Consider first iterated regret minimization. Notice that, for player 1,

$$\text{regret}_1^C(c_1) = \max_{c_2 \in C_2} \left( \max_{d_1 \in C_1} u_1(d_1, c_2) - u_1(c_1, c_2) \right) = 1, \quad \forall c_1 \in C_1.$$

Each pure strategy in  $C_1$  yields a maximum regret of 1. Thus,

$$\mathcal{RM}_1^\infty(C) = C_1.$$

Each pure strategy in  $C_1$  survives iterated regret minimization.

But this solution seems unreasonable. The pure strategies in  $C_1$  are not all on a par. On the contrary: there seems to be a natural and reasonable ordering over them. Notice that, for player 1, the pure strategy  $z_1$  is strongly dominated (for example, by a randomization between  $u_1$  and  $v_1$ ) and weakly dominated by every other strategy; that  $y_1$  is weakly dominated by  $x_1$ ; that  $x_1$  is weakly dominated by  $w_1$ ; that each of  $w_1$  and  $x_1$  and  $y_1$  is weakly dominated by  $u_1$  and by  $v_1$ ; and that neither  $u_1$  nor  $v_1$  is dominated. Thus, the most attractive strategies are  $u_1$  and  $v_1$ , followed in turn by  $w_1$ , and then by  $x_1$ , and then by  $y_1$ , and then finally by  $z_1$ .

Iterated regret minimization is here incapable of recognizing certain strategies as more reasonable than others, and the reason is the exclusive focus on the maximum regret. For example,  $z_1$  is regarded as just as reasonable as  $u_1$ —even though  $z_1$  can never yield a positive payoff and  $u_1$  yields the best-response payoff with respect to all but one of the strategies of player 2—simply because  $z_1$  and  $u_1$  yield the same maximum regret of 1.

As the game shown in Table 3.3 illustrates, the exclusive focus on the maximum regret can be problematic. In particular, such a focus means that domination may go unrecognized, with the result that dominated strategies may be equated with undominated strategies. In general, evaluating a strategy solely in terms of the maximum regret that it might yield means that the enumeration of partial profiles of strategies of the other players to which the strategy fails to be a best response—whether it be just one partial profile or all of them—matters not at all.

By contrast, least-squares regret yields a more reasonable solution. Notice that, for player 1,

$$\begin{aligned}\rho_1(u_1) &= \rho_1(v_1) = 1, \\ \rho_1(w_1) &= 2, \\ \rho_1(x_1) &= 3, \\ \rho_1(y_1) &= 4, \quad \text{and} \\ \rho_1(z_1) &= 5.\end{aligned}$$

Thus, least-squares regret captures exactly the natural and reasonable ordering described earlier and yields as the sole solutions precisely the two undominated pure strategies  $u_1$  and  $v_1$ . In general, least-squares regret rejects dominated strategies; for more on this point, see Theorem 5.6.2 in Section 5.6 and the discussion in Section 5.3.

For another example in which least-squares regret outperforms iterated regret minimization, consider the two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form, where

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 0 \leq x \leq 1\},$$

the utility function for player 1 is

$$\begin{aligned} u_1(c_1, c_2) &= 0 - \varepsilon && \text{if } c_1 = c_2 = 0, \\ &= 1 && \text{if } c_1 = 0 \text{ and } c_2 > 0, \\ &= 1 && \text{if } c_1 > 0 \text{ and } c_2 = 0, \\ &= 0 && \text{if } c_1 > 0 \text{ and } c_2 > 0, \end{aligned}$$

and  $\varepsilon$  is any small positive number.

Consider first iterated regret minimization. Notice that, for player 1,

$$\begin{aligned} \text{regret}_1^C(0) &= \max_{c_2 \in C_2} \left( \max_{d_1 \in C_1} u_1(d_1, c_2) - u_1(0, c_2) \right) = 1 + \varepsilon \quad \text{and} \\ \text{regret}_1^C(c_1) &= \max_{c_2 \in C_2} \left( \max_{d_1 \in C_1} u_1(d_1, c_2) - u_1(c_1, c_2) \right) = 1, \quad \forall c_1 \in \{x \in \mathbf{R} \mid 0 < x \leq 1\}. \end{aligned}$$

Thus,

$$\mathcal{RM}_1^\infty(C) = C_1 \setminus \{0\}.$$

The pure strategy 0 is eliminated, and every other pure strategy survives iterated regret minimization.

But this solution seems unreasonable. With respect to almost all—that is, except for a set of measure zero—of the pure strategies of player 2, the pure strategy 0 yields the best-response payoff of 1, and every other pure strategy yields the minimum payoff of 0. Furthermore, choosing the pure strategy 0 fails to be a best response if and only if  $c_2 = 0$ , and then the regret is  $1 + \varepsilon$ , which is only marginally greater than 1. Thus, it seems that the sole reasonable solution is the pure strategy 0.

The trouble, again, is the exclusive focus on the maximum regret. The only reason that the pure strategy 0 is eliminated and every other pure strategy is preserved is that the maximum regret from choosing the pure strategy 0 is  $1 + \varepsilon$  while the maximum regret from choosing any other pure strategy is 1. But this is not an overwhelming reason to eliminate the pure strategy 0 in favor of the other pure strategies considering that the differences

$$\text{regret}_1^C(0) - \text{regret}_1^C(c_1) = \varepsilon, \quad \forall c_1 \in \{x \in \mathbf{R} \mid 0 < x \leq 1\},$$

are essentially negligible—and are all the more so the smaller is  $\varepsilon$ —and that, as noted earlier, the pure strategy 0 outperforms every other pure strategy with respect to almost all of the pure strategies of player 2.

The exclusive focus on the maximum regret seems excessively strict and pessimistic. It entails favoring one strategy over another as long as the maximum

regret from choosing the latter is greater than that from choosing the former, no matter how negligible the difference and no matter how superior the unfavored strategy might be in other respects.

By contrast, least-squares regret yields a more reasonable solution. Notice that, for player 1,

$$\begin{aligned}\rho_1(0) &= 0 \quad \text{and} \\ \rho_1(c_1) &= 1, \quad \forall c_1 \in \{x \in \mathbf{R} \mid 0 < x \leq 1\}.\end{aligned}$$

Thus, least-squares regret yields as the sole solution precisely the pure strategy 0.

For another example in which least-squares regret yields a more agreeable solution than does iterated regret minimization, see the extended discussion on the two-person bargaining problem in Section 6.8. Theorem 6.8.2 establishes that in the two-person bargaining problem, if each player has a concave utility function, then the final payoff allocation yielded by least-squares regret is at least as great as that yielded by iterated regret minimization. Intuitively, the reason is that the exclusive focus on the maximum regret required by iterated regret minimization induces one to be more conservative than one would be if one were instead to act in accordance with least-squares regret.

# 4

---

## The Traveler's Dilemma

As noted in Chapter 1, the puzzle of the *Traveler's Dilemma* (Basu, 1994; Basu, 2007) is notoriously recalcitrant. What is most vexing about the game is that all of the standard solution concepts converge upon a unique solution that is both unreasonable in theory and refuted by the experimental evidence. Interest in the game remains high, and to this date, no generally accepted resolution of the puzzle exists.

In this chapter, we study the game in detail. Section 4.1 formally defines the game. Section 4.2 considers the standard solution concepts and shows how all yield the same unsatisfactory solution. Section 4.3 discusses the experimental evidence, which refutes the standard analyses. Section 4.4 describes some alternative analyses of the game. Section 4.5 shows how least-squares regret resolves the puzzle and yields solutions in line with both intuition and the experimental evidence.

### 4.1 Definition of the Game

The Traveler's Dilemma game is defined formally as follows. An airline loses the identical belongings of two travelers and reimburses them according to the following scheme. Each traveler must submit independently and privately to the airline a quotation that can be any number between 2 and 100, representing the value of the belongings. If the numbers match, each traveler is reimbursed that amount; otherwise, the traveler quoting the lower number is reimbursed the lower number plus some amount  $\alpha$  such that  $\alpha > 1$  while the other traveler is reimbursed the lower number less  $\alpha$ . The number  $\alpha$  can be thought of as the reward for quoting the lower number and the penalty for quoting the higher number. In this game, the set of players is  $N = \{1, 2\}$ . Now, there are two versions of the game: the discrete version and the continuous version. In the discrete

version, the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{Z} \mid 2 \leq x \leq 100\}.$$

In the continuous version, the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 2 \leq x \leq 100\}.$$

In both versions, for any player  $i$  in  $N$ , the utility function is

$$\begin{aligned} u_i(c_1, c_2) &= \min\{c_1, c_2\} - \alpha && \text{if } i \notin \underset{j \in \{1,2\}}{\operatorname{argmin}} c_j, \\ &= \min\{c_1, c_2\} + \alpha && \text{if } \{i\} = \underset{j \in \{1,2\}}{\operatorname{argmin}} c_j, \\ &= c_i && \text{if } c_1 = c_2, \end{aligned}$$

where  $\alpha$  is any real number, specified in advance, such that  $\alpha > 1$ .

## 4.2 Standard Analyses of the Game

There are many different ways to solve the game. In this section, we consider the standard analyses. For simplicity in what follows, restrict analysis to the standard discrete version where  $\alpha = 2$ .

Consider first maximin. Notice that the minimum payoff from quoting 2 is 2 and that, for any quotation other than 2, the minimum payoff is 0. Thus, the strategy of quoting 2 is the unique maximin strategy for each player.

Now, consider Nash equilibrium. At first glance, it may seem reasonable to conclude that each player should quote 100—the maximum number—since each will then be reimbursed that amount. But observe that there is always the incentive for one player to undercut the other player by the slimmest margin whenever possible. Assuming that the other player quotes 100, one would do better to deviate and quote 99. But, anticipating this, the other player would then do better to deviate and quote 98. And, anticipating this, one would then do better to deviate and quote 97. Continuing thus leads to the conclusion that, in the unique equilibrium of the game, each player quotes 2—the minimum number.

Seen in this light, the Traveler's Dilemma turns out to be very much like a one-shot simultaneous-move version of the famous *Centipede Game* (Rosenthal, 1981), a game well known for embodying the paradox of rationality posed by backward induction. This similarity is no accident. Indeed, the Traveler's Dilemma was devised to illustrate that the paradox is deeper than had been thought since it can arise in a one-shot simultaneous-move game and cannot be resolved by finding fault with structures relating to the extensive form. Thus, the Traveler's

Dilemma, involving backward induction at the level of introspection and purged of supplementary elements such as play over time, can be seen as one of the purest forms of the paradox of rationality.

Part of what makes the earlier equilibrium analysis so unreasonable is that it presumes a kind of superrationality on the part of the players. Each player must engage in a long chain of reasoning involving myriad assumptions of one another's extreme sophistication in order to reach the equilibrium solution.

The other standard solution concepts fare no better. Consider iterative elimination of weakly dominated strategies. Notice that, for each player, the strategy of quoting 100 is weakly dominated by the strategy of quoting 99 and can thus be eliminated. But then, for each player, the strategy of quoting 99 becomes weakly dominated by the strategy of quoting 98 and can thus be eliminated. Continuing thus leaves each player with the strategy of quoting 2 as the unique iteratively undominated strategy in the weak sense. But as with equilibrium analysis, this analysis presumes a kind of superrationality on the part of the players.

Iterative elimination of strongly dominated strategies is even more complicated. Notice that, for each player  $i$  in  $N$ , the strategy of quoting 100 is strongly dominated by the randomized strategy  $\sigma_i$  in  $\Delta(C_i)$  such that

$$\begin{aligned} \sigma_i(100) &= 0 \quad \text{and} \\ \sigma_i(c_i) &= \frac{\varepsilon^{100-c_i}}{\sum_{k=1}^{98} \varepsilon^k}, \quad \forall c_i \in \{x \in \mathbf{Z} \mid 2 \leq x \leq 99\}, \end{aligned}$$

where  $\varepsilon$  is any small positive number, and the strategy of quoting 100 can thus be eliminated. But then, for each player  $i$  in  $N$ , the strategy of quoting 99 becomes strongly dominated by the randomized strategy  $\sigma_i$  in  $\Delta(C_i \setminus \{100\})$  such that

$$\begin{aligned} \sigma_i(99) &= 0 \quad \text{and} \\ \sigma_i(c_i) &= \frac{\varepsilon^{99-c_i}}{\sum_{k=1}^{97} \varepsilon^k}, \quad \forall c_i \in \{x \in \mathbf{Z} \mid 2 \leq x \leq 98\}, \end{aligned}$$

where  $\varepsilon$  is any small positive number, and the strategy of quoting 99 can thus be eliminated. Continuing thus leaves each player with the strategy of quoting 2 as the unique iteratively undominated strategy in the strong sense. Just as before, a kind of superrationality on the part of the players is presumed.

Rationalizability (Bernheim, 1984; Pearce, 1984) is likewise unsatisfactory. As is well known, in two-person games, the set of rationalizable strategies for a player is just the set of strategies that survive iterative elimination of strongly dominated strategies. Thus, given the foregoing analysis, the strategy of quoting 2 is also the unique rationalizable strategy for each player. But determining the set of rationalizable strategies for a player requires an analysis just as involved as the foregoing ones.

It is natural to be dissatisfied with these solution concepts. All of them yield the strategy of quoting 2 as the unique rational strategy, defying intuition, and, moreover, all of them, with the exception of maximin, presume a kind of superrationality on the part of the players.

### 4.3 Experimental Evidence

Notwithstanding the analyses described in Section 4.2, it seems that no reasonable person would ever choose to quote 2 or to engage in the complicated reasoning required. In fact, experimental research demonstrates that people consistently reject the strategy of quoting 2 in favor of quoting a larger number and, moreover, that those who do so generally win higher payoffs. For detailed results, see, for example, Capra, Goeree, Gómez, and Holt (1999); Goeree and Holt (2001); Rubinstein (2006); Cabrera, Capra, and Gómez (2007); Rubinstein (2007); Chakravarty, Dechenaux, and Roy (2010); and Basu, Becchetti, and Stanca (2011).

Even professional game theorists fail to play in accordance with standard game theory. An experiment conducted on members of the Game Theory Society reveals deviations from standard game theory just as marked as in other experiments (Becker, Carter, and Naeve, 2005). In the experiment, members were asked to submit a strategy, pure or randomized, for the one-shot Traveler's Dilemma. Fifty-one entries were received. Of the forty-five specifying a pure strategy, ten specified the maximum quotation of 100; thirty-one specified a quotation of 96 or greater; thirty-eight specified a quotation of 90 or greater; and only three specified the prescribed quotation of 2. Just as in the other experiments noted above, there is a salient concentration at the top of the scale—the very opposite of what standard game theory suggests. Furthermore, the pure strategy that turns out to do best against the average strategy is the strategy of quoting 97, with its expected payoff of 85.09, and worst of all is the prescribed strategy of quoting 2, with its expected payoff of 3.92. These results thus appear to confirm the suspicion noted earlier: even knowing all that standard game theory has to say, there is something rational about quoting a larger number. The confirmation lies in the payoff received.

Another study shows that, in general, the greater is the value of  $\alpha$ , the lower is the number that is quoted, as intuition suggests (Capra, Goeree, Gómez, and Holt, 1999). Importantly, however, none of the standard solution concepts is sensitive to the value of  $\alpha$ .

### 4.4 Alternative Analyses

Numerous alternative analyses have been proposed. Capra, Goeree, Gómez, and



Holt (1999) and Anderson, Goree, and Holt (2002) show that logit equilibrium models can fit the experimental data very well. Becker, Carter, and Naeve (2005) adapt the approach of Kreps, Milgrom, Roberts, and Wilson (1982) used to explain cooperation in the finitely repeated Prisoners' Dilemma and model the Traveler's Dilemma as a Bayesian game that includes an irrational or "cooperative" type. Rubinstein (2006) and Rubinstein (2007) analyze response times and suggest that the distribution of observed strategy choices may be due at least in part to variations in cognitive effort or sophistication. Cabrera, Capra, and Gómez (2007) develop a noisy introspection model in which the reasoning process consists in iteratively computing responses, with some error, until some specified stopping rule is satisfied. Chakravarty, Dechenaux, and Roy (2010) consider the ability of pre-play communication to induce coordination and show that while ill-defined communication does little to foster cooperation, precise communication can lead to higher numbers being chosen. Basu, Becchetti, and Stanca (2011) extend the results of Capra, Goeree, Gómez, and Holt (1999) and show the dominance of one's own bonus-penalty amount, document heterogeneity of player types, find evidence of inadequacy in strategic thinking, and show that strategy choice and treatment effects are largely explained by risk aversion. Bavly (2012) considers the effect of introducing uncertainty about the range of available strategies. Halpern and Pass (2012) show how iterated regret minimization yields solutions consistent with intuition and the experimental evidence. Baghestanian (2014) develops a level- $k$  model with heterogeneous types to explain the observed data.

## 4.5 Least-Squares Regret

In this section, we show how least-squares regret yields for the Traveler's Dilemma solutions in line with intuition and the experimental evidence, at once resolving the puzzle and outperforming all of the standard solution concepts.

But before continuing, it is instructive to consider first the question of whether people, in fact, reason in terms of regret in a game such as the Traveler's Dilemma. To answer this question, notice that there is a general intuition that quoting 2, as standard game theory prescribes, is a poor strategy all things considered and that one would do better to quote a higher number. Quoting 2 is a best response if and only if the other player quotes either 2 or 3. But in any other circumstance, one would do better—and potentially substantially better—to quote a higher number. Furthermore, by quoting 2, one limits one's maximum payoff to 4, a meager amount considering the range of potential payoffs. Notice also that the greater is the quotation of the other player, the more one loses out by quoting 2. These observations suggest that regret may be an important strategic consideration,

especially in a game such as the Traveler's Dilemma.

Furthermore, as the studies of Ritov (1996) and Grosskopf, Erev, and Yechiam (2006) described in Section 1.3 illustrate, there is evidence in experimental economics that suggests that regret plays a significant role in decision-making. There is thus reason to think that regret might well play a role also in reasoning about the Traveler's Dilemma.

It is easy to see how least-squares regret resolves the Traveler's Dilemma puzzle. For concreteness, consider the standard discrete version of the game where the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{Z} \mid 2 \leq x \leq 100\}$$

and  $\alpha = 2$ .

For an illustration, consider the strategy of quoting 2. Notice that, for any player  $i$  in  $N$ , the regret from quoting 2 is

$$\begin{aligned} \rho_i(2) &= \sum_{c_{-i}=2}^{100} \left( \max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, 2) \right)^2 \\ &= (2-2)^2 + (4-4)^2 + \dots + (101-4)^2 \\ &= 308945. \end{aligned}$$


The regret from quoting 2 is high. And the reasons are precisely those given earlier. As the computation indicates, quoting 2 is a best response if and only if the other player quotes either 2 or 3, but in any other circumstance, one would do better to quote a higher number, and, moreover, the greater is the quotation of the other player, the more one loses out by quoting 2.

For another illustration, consider the strategy of quoting 100. Notice that, for any player  $i$  in  $N$ , the regret from quoting 100 is


$$\begin{aligned} \rho_i(100) &= \sum_{c_{-i}=2}^{100} \left( \max_{d_i \in C_i} u_i(c_{-i}, d_i) - u_i(c_{-i}, 100) \right)^2 \\ &= (2-0)^2 + (4-1)^2 + \dots + (101-100)^2 \\ &= 878. \end{aligned}$$

The regret from quoting 100 is low. And the reason is what intuition suggests. Notably, while in no circumstance would it be a best response to quote 100, doing so guarantees in every conceivable circumstance a payoff very close to the best-response payoff.

In general, quoting a low number is associated with high regret while quoting a high number is associated with low regret, as intuition suggests and as Figure 4.1 shows.

Figure 4.1 Regret of player  $i$  in the Traveler's Dilemma

In fact, for any player  $i$  in  $N$ , the pure strategies that minimize the regret function  $\rho_i$  are the strategy of quoting 96 and the strategy of quoting 97, as Figure 4.2 shows.

Figure 4.2 Regret of player  $i$  in the Traveler's Dilemma

Thus, unlike the standard solution concepts, least-squares regret yields solutions in line with the experimental evidence and with the intuition that quoting a high number is a better strategy all things considered than quoting a low number. Notably, the solutions yielded by least-squares regret are consistent with the experimental result discussed earlier that the strategy of quoting 97 is the pure strategy that does best against the average strategy (Becker, Carter, and Naeve, 2005).

Furthermore, it is easy to see that least-squares regret captures the expected sensitivity to the value of  $\alpha$ . To see this, consider the continuous version of the game where the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 2 \leq x \leq 100\}.$$

Consider any player  $i$  in  $N$  and any pure strategy  $c_i$  in  $C_i$ . If the other player chooses some pure strategy  $c_{-i}$  in  $C_{-i}$  that is less than  $c_i$ , then the payoff to player

$i$  from choosing  $c_i$  is  $c_{-i} - \alpha$  while the best-response payoff is  $c_{-i} + \alpha$  (since the best response is to undercut the other player by the slimmest margin). If the other player chooses some pure strategy  $c_{-i}$  in  $C_{-i}$  that is greater than  $c_i$ , then the payoff to player  $i$  from choosing  $c_i$  is  $c_i + \alpha$  while the best-response payoff is again  $c_{-i} + \alpha$ . Thus, for any player  $i$  in  $N$ ,

$$\begin{aligned}\rho_i(c_i) &= \int_2^{c_i} ((c_{-i} + \alpha) - (c_{-i} - \alpha))^2 dc_{-i} + \int_{c_i}^{100} ((c_{-i} + \alpha) - (c_i + \alpha))^2 dc_{-i} \\ &= \int_2^{c_i} 4\alpha^2 dc_{-i} + \int_{c_i}^{100} (c_{-i} - c_i)^2 dc_{-i} \\ &= (100 - c_i)^3/3 + 4\alpha^2 c_i - 8\alpha^2,\end{aligned}$$

and so,

$$\frac{d(\rho_i(c_i))}{dc_i} = 4\alpha^2 - (100 - c_i)^2.$$

Setting the derivative equal to 0 and solving for  $c_i$  yields

$$c_i = \frac{200 \pm \sqrt{(-200)^2 - 4(100^2 - 4\alpha^2)}}{2} = 100 \pm 2\alpha.$$

The regret function  $\rho_i$  is minimized at the smaller critical point, which is also the only critical point in the domain

$$C_i = \{x \in \mathbf{R} \mid 2 \leq x \leq 100\}.$$

Thus, the unique pure least-squares regret strategy for player  $i$  is  $c_i = 100 - 2\alpha$ . This strategy is a strictly decreasing function of  $\alpha$ , and so, the greater is the value of  $\alpha$ , the lower is the number that is yielded by least-squares regret.

Thus, again, unlike the standard solution concepts, least-squares regret yields solutions in line with the experimental evidence and with the intuition that strategy choice is sensitive to the value of  $\alpha$ . Notably, the sensitivity captured by least-squares regret is consistent with the experimental result discussed earlier that the greater is the value of  $\alpha$ , the lower is the number that is quoted (Capra, Goeree, Gómez, and Holt, 1999).

# 5

---

## Properties

In appraising a solution concept and assessing its performance, numerous questions arise, such as whether the solution concept satisfies a general existence theorem; whether, for any game, the solution concept yields as solutions all intuitively reasonable outcomes and rules out all intuitively unreasonable ones; whether, for any game, the set of solutions yielded by the solution concept is invariant when the game is transformed in a way that is considered irrelevant; whether the fundamental logic of the solution concept is intuitive, credible, and compelling as a characterization of typical reasoning and behavior.

In this chapter, we establish some notable mathematical and conceptual properties of least-squares regret. Section 5.1 presents a general existence theorem. Section 5.2 presents a theorem that establishes that least-squares regret is invariant with respect to full equivalence. Section 5.3 examines the relationship between least-squares regret and dominated strategies. Section 5.4 shows that least-squares regret and iterative elimination of dominated strategies can yield different solutions. Section 5.5 examines the relationship between least-squares regret and uniformly dominant strategies. Section 5.6 presents two theorems that establish that least-squares regret is invariant with respect to certain well-known transformations of the payoffs in a game that leave unchanged the best-response correspondences of the players. Section 5.7 presents a theorem that establishes that when it comes to the equilibria in pure strategies of a  $2 \times 2$  game in strategic form, least-squares regret is equivalent to risk dominance. Section 5.8 presents a theorem that establishes that, for any finite game in strategic form, least-squares regret yields a convex set of solutions. Section 5.9 presents two theorems that describe sufficient conditions for the uniqueness of a solution yielded by least-squares regret.

## 5.1 Existence

The following general existence theorem establishes that every finite game in strategic form has at least one least-squares regret profile in pure strategies and at least one least-squares regret profile in randomized strategies. Thus, for any finite game in strategic form, a solution is guaranteed to exist.

**THEOREM 5.1.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Then there exist some least-squares regret profile in pure strategies and some least-squares regret profile in randomized strategies.*

*Proof.* Consider any player  $i$  in  $N$ . Consider first pure strategies. Since the set  $C_i$  is finite, the regret function  $\rho_i: C_i \rightarrow \mathbf{R}$  has a minimum. And so, the set

$$\operatorname{argmin}_{d_i \in C_i} \rho_i(d_i)$$

is nonempty.

Thus, the set

$$\prod_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i)$$

is nonempty since it is the Cartesian product of nonempty sets.

Now, consider randomized strategies. Since the set  $\Delta(C_i)$  is a nonempty compact set and since the regret function  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  is continuous,  $\rho_i$  has a minimum. And so, the set

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is nonempty.

Thus, the set

$$\prod_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is nonempty since it is the Cartesian product of nonempty sets. ■

## 5.2 Full Equivalence

A utility function is simply a mathematical representation of the preferences of an individual and is unique up to strictly increasing affine transformation. Thus, replacing any number of the utility functions in a game with decision-theoretically equivalent ones leaves unchanged the underlying preference structure of the game. Since the original game and the transformed game represent the same fundamental situation, they must be considered decision-theoretically equivalent.

Myerson (1991) introduces the concept of *full equivalence* of finite games in strategic form to characterize the equivalence just described. Full equivalence might more illuminatingly be called *cardinal equivalence*.

Full equivalence of finite games in strategic form is defined formally as follows. For any finite games  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  and  $\hat{\Gamma} = (N, (C_i)_{i \in N}, (\hat{u}_i)_{i \in N})$  in strategic form, the games  $\Gamma$  and  $\hat{\Gamma}$  are *fully equivalent* if and only if, for every player  $i$  in  $N$ , there exist real numbers  $A_i$  and  $B_i$  such that  $A_i > 0$  and

$$\hat{u}_i(c_{-i}, c_i) = A_i u_i(c_{-i}, c_i) + B_i, \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i.$$

As a matter of consistency, a solution concept should not yield different solutions for fully equivalent games. Any solution concept that fails on this score must be seen as being faulty in a significant way. For example, as noted in Section 1.1, quantal response equilibrium (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998) is not scale invariant and can thus be inconsistent, yielding different solutions for fully equivalent games (Wright and Leyton-Brown, 2010).

The following theorem establishes that least-squares regret is invariant with respect to full equivalence in the sense that fully equivalent games have the same solutions.

**THEOREM 5.2.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  and  $\hat{\Gamma} = (N, (C_i)_{i \in N}, (\hat{u}_i)_{i \in N})$  be any finite games in strategic form such that  $\Gamma$  and  $\hat{\Gamma}$  are fully equivalent. For any player  $i$  in  $N$ , let  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\Gamma$ , let  $\hat{\rho}_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\hat{\Gamma}$ , let  $\rho_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\Gamma$ , and let  $\hat{\rho}_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\hat{\Gamma}$ . Then*

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i)$$

and

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

*Proof.* Consider any player  $i$  in  $N$ . By assumption, the games  $\Gamma$  and  $\hat{\Gamma}$  are fully equivalent, and so, there exist real numbers  $A_i$  and  $B_i$  such that  $A_i > 0$  and

$$\hat{u}_i(c_{-i}, c_i) = A_i u_i(c_{-i}, c_i) + B_i, \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i.$$

Clearly, it follows that

$$\hat{u}_i(c_{-i}, \sigma_i) = A_i u_i(c_{-i}, \sigma_i) + B_i, \quad \forall c_{-i} \in C_{-i}, \quad \forall \sigma_i \in \Delta(C_i).$$

Thus,

$$\begin{aligned}
& \hat{\rho}_i(\sigma_i) \\
&= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} \hat{u}_i(c_{-i}, \tau_i) - \hat{u}_i(c_{-i}, \sigma_i) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\
&= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} (A_i u_i(c_{-i}, \tau_i) + B_i) - (A_i u_i(c_{-i}, \sigma_i) + B_i) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\
&= A_i^2 \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\
&= A_i^2 \rho_i(\sigma_i), \quad \forall \sigma_i \in \Delta(C_i).
\end{aligned}$$

Since  $A_i > 0$ , the regret functions  $\rho_i$  and  $\hat{\rho}_i$  differ by a strictly increasing linear transformation. And so,

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

From the above, it follows that

$$\hat{\rho}_i(c_i) = A_i^2 \rho_i(c_i), \quad \forall c_i \in C_i.$$

Since  $A_i > 0$ , the regret functions  $\rho_i$  and  $\hat{\rho}_i$  differ by a strictly increasing linear transformation. And so,

$$\operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i). \quad \blacksquare$$

### 5.3 Dominated Strategies

It is instructive to examine the relationship between least-squares regret and dominated strategies. Intuitively, if a pure strategy is weakly or strongly dominated for a player, then, for any partial profile of strategies of the other players, the regret of the dominated strategy with respect to the partial profile is at least as great as the regret of the dominating strategy with respect to the partial profile, and there



exists some partial profile such that the regret of the dominated strategy with respect to the partial profile is strictly greater than the regret of the dominating strategy with respect to the partial profile.

This intuition leads immediately to two conclusions. To begin, a pure strategy that is weakly or strongly dominated for a player by some randomized strategy cannot, when assigned probability 1, minimize the regret function in randomized strategies. Furthermore, a pure strategy that is weakly or strongly dominated for a player by some other pure strategy cannot minimize the regret function in pure strategies. The following proposition establishes these facts.

**PROPOSITION 5.3.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$ , any pure strategies  $c_i$  and  $\hat{c}_i$  in  $C_i$ , and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , if  $c_i$  is weakly or strongly dominated for player  $i$  by  $\sigma_i$ , then*

$$\rho_i(\sigma_i) < \rho_i(c_i),$$

and if  $c_i$  is weakly or strongly dominated for player  $i$  by  $\hat{c}_i$ , then

$$\rho_i(\hat{c}_i) < \rho_i(c_i).$$

*Proof.* Consider any player  $i$  in  $N$ . Let  $c_i$  be any pure strategy in  $C_i$ , and let  $\sigma_i$  be any randomized strategy in  $\Delta(C_i)$ . Suppose that  $c_i$  is weakly or strongly dominated for player  $i$  by  $\sigma_i$ . Then

$$\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \leq \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i},$$

and there exists some partial profile  $\hat{c}_{-i}$  in  $C_{-i}$  such that

$$\max_{\tau_i \in \Delta(C_i)} u_i(\hat{c}_{-i}, \tau_i) - u_i(\hat{c}_{-i}, \sigma_i) < \max_{\tau_i \in \Delta(C_i)} u_i(\hat{c}_{-i}, \tau_i) - u_i(\hat{c}_{-i}, c_i).$$

Thus,

$$\rho_i(\sigma_i) < \rho_i(c_i).$$

Let  $c_i$  and  $\hat{c}_i$  be any pure strategies in  $C_i$ . Suppose that  $c_i$  is weakly or strongly dominated for player  $i$  by  $\hat{c}_i$ . From the above, it follows that

$$\rho_i(\hat{c}_i) < \rho_i(c_i). \quad \blacksquare$$

Whether or not randomized strategies are considered can matter substantially. For an illustration of this point, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.1.

Consider the pure strategy  $z_1$ . Notice that it is for player 1 weakly dominated, for example, by the randomized strategy  $(2/3)x_1 + (1/3)y_1$ , and strongly dominated,

Table 5.1 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	3
	$y_1$	3	0
	$z_1$	1	1

for example, by the randomized strategy  $0.5x_1 + 0.5y_1$ . Thus, by Proposition 5.3.1, the pure strategy  $z_1$  cannot, when assigned probability 1, be a randomized least-squares regret strategy for player 1. But notice that if only pure strategies are considered, then  $z_1$  is neither weakly nor strongly dominated for player 1 and, moreover, since

$$\rho_1(z_1) = 8 < 9 = \rho_1(x_1) = \rho_1(y_1),$$

it is, in fact, the unique pure least-squares regret strategy for player 1.

It is clear that the support of a randomized least-squares regret strategy cannot contain a weakly or strongly dominated strategy. The following proposition establishes this fact.

**PROPOSITION 5.3.2.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$ , any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , and any pure strategy  $c_i$  in  $C_i$ , if  $\sigma_i$  is a randomized least-squares regret strategy for player  $i$  and  $c_i$  is weakly or strongly dominated for player  $i$ , then  $\sigma_i(c_i) = 0$ .*

*Proof.* Consider any player  $i$  in  $N$ . Let  $\sigma_i$  be any randomized strategy in  $\Delta(C_i)$ , and let  $c_i$  be any pure strategy in  $C_i$ . Suppose that  $\sigma_i$  is a randomized least-squares regret strategy for player  $i$  and that  $c_i$  is weakly or strongly dominated for player  $i$  and that  $\sigma_i(c_i) > 0$ . Since  $c_i$  is weakly or strongly dominated for player  $i$ , there exists some randomized strategy  $\xi_i$  in  $\Delta(C_i)$  such that

$$u_i(c_{-i}, \xi_i) \geq u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i},$$

and there exists some partial profile  $\hat{c}_{-i}$  in  $C_{-i}$  such that

$$u_i(\hat{c}_{-i}, \xi_i) > u_i(\hat{c}_{-i}, c_i).$$

Now, consider the randomized strategy  $\hat{\sigma}_i$  in  $\Delta(C_i)$  that is exactly like  $\sigma_i$  except that  $c_i$  is assigned probability 0 and  $\xi_i$  is assigned probability  $\sigma_i(c_i)$ . That is,  $\hat{\sigma}_i$  is

the randomized strategy in  $\Delta(C_i)$  such that

$$\begin{aligned}\hat{\sigma}_i(c_i) &= 0 \quad \text{and} \\ \hat{\sigma}_i(d_i) &= \sigma_i(d_i) + \sigma_i(c_i)\xi_i(d_i), \quad \forall d_i \in C_i \setminus \{c_i\}.\end{aligned}$$

Clearly,

$$\max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \hat{\sigma}_i) \leq \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i), \quad \forall c_{-i} \in C_{-i},$$

and there exists some partial profile  $\bar{c}_{-i}$  in  $C_{-i}$  such that

$$\max_{\tau_i \in \Delta(C_i)} u_i(\bar{c}_{-i}, \tau_i) - u_i(\bar{c}_{-i}, \hat{\sigma}_i) < \max_{\tau_i \in \Delta(C_i)} u_i(\bar{c}_{-i}, \tau_i) - u_i(\bar{c}_{-i}, \sigma_i).$$

Thus,

$$\rho_i(\hat{\sigma}_i) < \rho_i(\sigma_i).$$

But this is a contradiction. Thus,  $\sigma_i(c_i) = 0$ . ■

Finally, it is instructive to consider the question of iterative elimination of dominated strategies. Consider again the Traveler's Dilemma discussed in Chapter 4. Recall from Section 4.2 that, for each player, the strategy of quoting 96 and the strategy of quoting 97 each end up being eliminated following both iterative elimination of weakly dominated strategies and iterative elimination of strongly dominated strategies. Furthermore, recall from Section 4.5 that, for each player, the pure least-squares regret strategies are the strategy of quoting 96 and the strategy of quoting 97.

This discrepancy does not contradict Proposition 5.3.1, which ignores iterative elimination of dominated strategies. Notice that the strategy of quoting 96 and the strategy of quoting 97 become weakly or strongly dominated for a player only after several rounds of elimination. Neither strategy is weakly or strongly dominated prior to any elimination, and so, it is not contradictory that each is a pure least-squares regret strategy. By contrast, the strategy of quoting 100 is weakly and strongly dominated prior to any elimination and is also not a pure least-squares regret strategy, precisely as Proposition 5.3.1 establishes.

The point is that whether or not iterative elimination of dominated strategies is considered can matter greatly. In general, as discussed in Section 5.4, least-squares regret and iterative elimination of dominated strategies can yield different solutions.

## 5.4 Iterative Elimination of Dominated Strategies

Least-squares regret and iterative elimination of dominated strategies can yield different solutions. For an illustration of this point, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.2.

Table 5.2 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0 3	1 1
	$y_1$	0 1	1 2

This game is dominance solvable. Notice that  $x_2$  is strongly dominated for player 2 by  $y_2$  and can thus be eliminated. But once  $x_2$  is eliminated,  $x_1$  becomes strongly dominated for player 1 by  $y_1$  and can thus be eliminated. Thus, the unique solution yielded by iterative elimination of dominated strategies is  $(y_1, y_2)$ , which gives the payoff allocation  $(2, 1)$ .

Least-squares regret yields a different solution. The unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ , which gives the payoff allocation  $(1, 1)$ .

The discrepancy observed in this example is a consequence of supposing that a player is partially strategic in the sense discussed in Sections 1.2, 1.4, 2.2, and 2.5. In Section 7.4, we consider more fully the problems, such as the discrepancy above, that can arise from making such an assumption.

## 5.5 Uniformly Dominant Strategies

Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$  and any pure strategy  $c_i$  in  $C_i$ , the pure strategy  $c_i$  is *uniformly dominant* for player  $i$  if and only if

$$u_i(c_{-i}, c_i) \geq u_i(c_{-i}, d_i), \quad \forall d_i \in C_i, \quad \forall c_{-i} \in C_{-i}.$$

Intuitively, a uniformly dominant strategy for a player is at least as good as every other pure strategy with respect to each partial profile of strategies of the other players. Such a strategy is an obvious choice for playing the game. Notice that a player may have more than one uniformly dominant strategy.

We note in passing that while the term *weakly dominant* is sometimes used to mean what we call *uniformly dominant*, we prefer our term since the idea of a pure strategy being weakly dominant for a player typically involves at least one strict inequality with respect to some or another partial profile of strategies of the other players. With the term *uniformly dominant*, there is no such ambiguity and thus no risk of confusion.

The following proposition establishes that if a player has at least one uniformly dominant strategy, then every uniformly dominant strategy is a pure least-squares

regret strategy and, when assigned probability 1, is a randomized least-squares regret strategy, and every pure least-squares regret strategy is a uniformly dominant strategy. Thus, the following proposition formalizes the intuition that a uniformly dominant strategy is an obvious choice for playing a game.

**PROPOSITION 5.5.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$ , if there exists some pure strategy  $\bar{c}_i$  in  $C_i$  such that  $\bar{c}_i$  is uniformly dominant for player  $i$ , then, for any pure strategy  $c_i$  in  $C_i$ , if  $c_i$  is uniformly dominant for player  $i$ , then  $c_i$  is a pure least-squares regret strategy for player  $i$  and, when assigned probability 1, is a randomized least-squares regret strategy for player  $i$ , and if  $c_i$  is a pure least-squares regret strategy for player  $i$ , then  $c_i$  is uniformly dominant for player  $i$ .*

*Proof.* Consider any player  $i$  in  $N$ . Suppose that there exists some pure strategy  $\bar{c}_i$  in  $C_i$  such that  $\bar{c}_i$  is uniformly dominant for player  $i$ . Then  $\rho_i(\bar{c}_i) = 0$ .

Consider any pure strategy  $c_i$  in  $C_i$ . If  $c_i$  is uniformly dominant for player  $i$ , then  $\rho_i(c_i) = 0$ , and so,  $c_i$  is a pure least-squares regret strategy for player  $i$  and, when assigned probability 1, is a randomized least-squares regret strategy for player  $i$ . If  $c_i$  is not uniformly dominant for player  $i$ , then  $\rho_i(c_i) > 0$ , and so, it is not a pure least-squares regret strategy for player  $i$ . ■

The following proposition concerns randomized strategies and establishes that if a player has at least one uniformly dominant strategy, then the support of any randomized least-squares regret strategy can contain only uniformly dominant strategies. Intuitively, this proposition establishes that if one has at least one uniformly dominant strategy, then one can randomize however one pleases as long as one randomizes exclusively over uniformly dominant strategies.

**PROPOSITION 5.5.2.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$ , if there exists some pure strategy  $\bar{c}_i$  in  $C_i$  such that  $\bar{c}_i$  is uniformly dominant for player  $i$ , then, for any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$  and any pure strategy  $c_i$  in  $C_i$ , if  $\sigma_i$  is a randomized least-squares regret strategy for player  $i$  and  $\sigma_i(c_i) > 0$ , then  $c_i$  is uniformly dominant for player  $i$ .*

*Proof.* Consider any player  $i$  in  $N$ . Suppose that there exists some pure strategy  $\bar{c}_i$  in  $C_i$  such that  $\bar{c}_i$  is uniformly dominant for player  $i$ . Then  $\rho_i(\bar{c}_i) = 0$ .

Consider any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$  and any pure strategy  $c_i$  in  $C_i$ . Suppose that  $\sigma_i$  is a randomized least-squares regret strategy for player  $i$  and that  $\sigma_i(c_i) > 0$  and that  $c_i$  is not uniformly dominant for player  $i$ . Then  $\rho_i(\sigma_i) > 0$ . But then

$$\sigma_i \notin \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i).$$

But this is a contradiction. Thus,  $c_i$  is uniformly dominant for player  $i$ . ■

For an illustration of these two propositions, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form known as the *Prisoners' Dilemma* shown in Table 5.3.

Table 5.3 Prisoners' Dilemma game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	2    3	
	$y_1$	0    1	

It is easy to see that Proposition 5.5.1 holds. Consider any player  $i$  in  $N$ . Clearly,  $y_i$  is for player  $i$  both the unique uniformly dominant strategy and the unique pure least-squares regret strategy. Thus, for any player  $i$  in  $N$ , the set of uniformly dominant strategies for player  $i$  and the set of pure least-squares regret strategies for player  $i$  are identical. Furthermore, since  $\rho_i(y_i) = 0$ , it follows that  $y_i$ , when assigned probability 1, is also a randomized least-squares regret strategy for player  $i$ .

It is likewise easy to see that Proposition 5.5.2 holds. Consider any player  $i$  in  $N$ . The regret function  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_i(\sigma_i) &= \left( \max_{\tau_i \in \Delta(C_i)} u_i(x_{-i}, \tau_i) - u_i(x_{-i}, \sigma_i) \right)^2 + \left( \max_{\tau_i \in \Delta(C_i)} u_i(y_{-i}, \tau_i) - u_i(y_{-i}, \sigma_i) \right)^2 \\ &= (3 - (2\sigma_i(x_i) + 3(1 - \sigma_i(x_i))))^2 + (1 - (1 - \sigma_i(x_i)))^2 \\ &= 2(\sigma_i(x_i))^2. \end{aligned}$$

The regret function  $\rho_i$  is minimized at the point  $\sigma_i(x_i) = 0$ , and so, the unique randomized least-squares regret strategy for player  $i$  is  $y_i$ , which assigns probability 1 to the pure strategy  $y_i$ , the unique uniformly dominant strategy for player  $i$ .

In the absence of uniformly dominant strategies, randomizing over several pure strategies may be better than playing any particular pure strategy with probability 1. For an illustration of this point, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.4.

In this game, the unique randomized least-squares regret strategy for player 1 is  $0.5x_1 + 0.5y_1$ . Thus, minimizing the regret function  $\rho_1$  involves randomizing between  $x_1$  and  $y_1$  and not playing either with probability 1.

But this is not to say that in the absence of uniformly dominant strategies, randomizing over several pure strategies need be better than playing any particular

Table 5.4 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	3	0
	$y_1$	1	2

pure strategy with probability 1. Indeed, in the absence of uniformly dominant strategies, it may be best to play some particular pure strategy for sure. For an illustration of this point, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.5.

Table 5.5 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	2	-9
	$y_1$	1	0
	$z_1$	0	1

Consider player 1. The regret function  $\rho_1 : \Delta(C_1) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_1(\sigma_1) &= \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, x_2) - u_1(\sigma_1, x_2) \right)^2 + \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, y_2) - u_1(\sigma_1, y_2) \right)^2 \\ &= (2 - (2\sigma_1(x_1) + \sigma_1(y_1)))^2 + (1 - (1 - 10\sigma_1(x_1) - \sigma_1(y_1)))^2 \\ &= 104(\sigma_1(x_1))^2 + 2(\sigma_1(y_1))^2 + 24\sigma_1(x_1)\sigma_1(y_1) - 8\sigma_1(x_1) - 4\sigma_1(y_1) + 4. \end{aligned}$$

It is straightforward to verify that the regret function  $\rho_1$  is minimized when

$$\sigma_1(x_1) = 0 \quad \text{and} \quad \sigma_1(y_1) = 1.$$

Thus, the unique randomized least-squares regret strategy for player 1 is  $y_1$ , which assigns probability 1 to the pure strategy  $y_1$ .

## 5.6 Strategic Equivalence

In Section 5.2, we considered the concept of full equivalence of finite games in strategic form (Myerson, 1991). But other concepts of equivalence have been

proposed. Of particular significance are certain well-known transformations of the payoffs in a game—most notably, those transformations that characterize the concept of *strategic equivalence* of finite games in strategic form—that leave unchanged the best-response correspondences of the players. In this section, we show that least-squares regret is invariant with respect to such transformations.

Strategic equivalence of finite games in strategic form is defined formally as follows. For any finite games  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  and  $\hat{\Gamma} = (N, (C_i)_{i \in N}, (\hat{u}_i)_{i \in N})$  in strategic form, the games  $\Gamma$  and  $\hat{\Gamma}$  are *strategically equivalent* if and only if, for every player  $i$  in  $N$ , there exist some real number  $A_i$  and some function  $B_i: \times_{j \in N-i} C_j \rightarrow \mathbf{R}$  such that  $A_i > 0$  and

$$\hat{u}_i(c_{-i}, c_i) = A_i u_i(c_{-i}, c_i) + B_i(c_{-i}), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i.$$

Such transformations are particularly notable. For example, Moulin and Vial (1978) show that, for any two-person game, such transformations are the only transformations that leave unchanged the best-response correspondences of the players (setting aside the latitude to make a dominated strategy arbitrarily better or worse as long as it remains dominated). Furthermore, Moulin and Vial (1978) introduce the concept of a two-person *strategically zero-sum game*, which is defined as any game that is strategically equivalent to some two-person zero-sum game, and show that, with respect to a number of correlation schemes, including the scheme characterized by the concept of a correlated strategy (Aumann, 1974; Aumann, 1987), strategically zero-sum games are precisely those two-person games whose completely randomized equilibria cannot be improved upon.

As a matter of consistency, a solution concept should not yield different solutions for strategically equivalent games. Any solution concept that fails on this score must be seen as being faulty in a significant way. For example, as noted in Sections 1.1 and 5.2, quantal response equilibrium (McKelvey and Palfrey, 1995; McKelvey and Palfrey, 1998) is not scale invariant and can thus be inconsistent, yielding different solutions for strategically equivalent games (Wright and Leyton-Brown, 2010).

The following theorem establishes that least-squares regret is invariant with respect to strategic equivalence in the sense that strategically equivalent games have the same solutions.

**THEOREM 5.6.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  and  $\hat{\Gamma} = (N, (C_i)_{i \in N}, (\hat{u}_i)_{i \in N})$  be any finite games in strategic form such that  $\Gamma$  and  $\hat{\Gamma}$  are strategically equivalent. For any player  $i$  in  $N$ , let  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\Gamma$ , let  $\hat{\rho}_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\hat{\Gamma}$ , let  $\rho_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\Gamma$ ,*



and let  $\hat{\rho}_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\hat{\Gamma}$ . Then

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i)$$

and

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

*Proof.* Consider any player  $i$  in  $N$ . By assumption, the games  $\Gamma$  and  $\hat{\Gamma}$  are strategically equivalent, and so, there exist some real number  $A_i$  and some function  $B_i: \times_{j \in N-i} C_j \rightarrow \mathbf{R}$  such that  $A_i > 0$  and

$$\hat{u}_i(c_{-i}, c_i) = A_i u_i(c_{-i}, c_i) + B_i(c_{-i}), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i.$$

Clearly, it follows that

$$\hat{u}_i(c_{-i}, \sigma_i) = A_i u_i(c_{-i}, \sigma_i) + B_i(c_{-i}), \quad \forall c_{-i} \in C_{-i}, \quad \forall \sigma_i \in \Delta(C_i).$$

Thus,

$$\begin{aligned} & \hat{\rho}_i(\sigma_i) \\ &= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} \hat{u}_i(c_{-i}, \tau_i) - \hat{u}_i(c_{-i}, \sigma_i) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\ &= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} (A_i u_i(c_{-i}, \tau_i) + B_i(c_{-i})) - (A_i u_i(c_{-i}, \sigma_i) + B_i(c_{-i})) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\ &= A_i^2 \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right)^2, \quad \forall \sigma_i \in \Delta(C_i) \\ &= A_i^2 \rho_i(\sigma_i), \quad \forall \sigma_i \in \Delta(C_i). \end{aligned}$$

Since  $A_i > 0$ , the regret functions  $\rho_i$  and  $\hat{\rho}_i$  differ by a strictly increasing linear transformation. And so,

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

From the above, it follows that

$$\hat{\rho}_i(c_i) = A_i^2 \rho_i(c_i), \quad \forall c_i \in C_i.$$

Since  $A_i > 0$ , the regret functions  $\rho_i$  and  $\hat{\rho}_i$  differ by a strictly increasing linear transformation. And so,

$$\operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i). \quad \blacksquare$$

In some ways, a weakly or strongly dominated strategy for a player is strategically irrelevant and negligible. Thus, it seems reasonable to suppose that, as long as a weakly or strongly dominated strategy remains so, it can be made arbitrarily better or worse without affecting the behavior of the player. The following theorem establishes that least-squares regret is invariant with respect to such transformations of weakly or strongly dominated strategies.

**THEOREM 5.6.2.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  and  $\hat{\Gamma} = (N, (C_i)_{i \in N}, (\hat{u}_i)_{i \in N})$  be any finite games in strategic form. For any player  $i$  in  $N$ , let  $D_i$  be the set of weakly or strongly dominated strategies for player  $i$  in  $\Gamma$ , let  $\hat{D}_i$  be the set of weakly or strongly dominated strategies for player  $i$  in  $\hat{\Gamma}$ , let  $\rho_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\Gamma$ , let  $\hat{\rho}_i: \Delta(C_i) \rightarrow \mathbf{R}$  be the regret function in randomized strategies for player  $i$  in  $\hat{\Gamma}$ , let  $\rho_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\Gamma$ , and let  $\hat{\rho}_i: C_i \rightarrow \mathbf{R}$  be the regret function in pure strategies for player  $i$  in  $\hat{\Gamma}$ . If, for every player  $i$  in  $N$ ,*

$$D_i = \hat{D}_i$$

and

$$\hat{u}_i(c_{-i}, c_i) = u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i \setminus D_i,$$

then

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i)$$

and

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

*Proof.* Consider any player  $i$  in  $N$ . Consider first randomized strategies. Suppose that

$$D_i = \hat{D}_i = \emptyset.$$

Since, by assumption,

$$\hat{u}_i(c_{-i}, c_i) = u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i \setminus D_i,$$

it follows that  $\Gamma$  and  $\hat{\Gamma}$  are identical. Thus, it is trivially true that

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

Now, suppose that

$$D_i = \hat{D}_i \neq \emptyset.$$

By Proposition 5.3.2 in Section 5.3, the support of a randomized least-squares regret strategy cannot contain a weakly or strongly dominated strategy. Thus, only the undominated strategies in  $C_i$  can be included in the support of a randomized least-squares regret strategy. But, by assumption,

$$\hat{u}_i(c_{-i}, c_i) = u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i \setminus D_i.$$

And so,

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i) = \bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \hat{\rho}_i(\tau_i).$$

Now, consider pure strategies. Suppose that

$$D_i = \hat{D}_i = \emptyset.$$

Since, by assumption,

$$\hat{u}_i(c_{-i}, c_i) = u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i \setminus D_i,$$

it follows that  $\Gamma$  and  $\hat{\Gamma}$  are identical. Thus, it is trivially true that

$$\operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

Now, suppose that

$$D_i = \hat{D}_i \neq \emptyset.$$

By Proposition 5.3.1 in Section 5.3, a pure strategy that is weakly or strongly dominated for a player by some other pure strategy cannot minimize the regret

function in pure strategies. Thus, only an undominated strategy in  $C_i$  can be a pure least-squares regret strategy. But, by assumption,

$$\hat{u}_i(c_{-i}, c_i) = u_i(c_{-i}, c_i), \quad \forall c_{-i} \in C_{-i}, \quad \forall c_i \in C_i \setminus D_i,$$

And so,

$$\operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i).$$

Thus,

$$\bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \rho_i(d_i) = \bigtimes_{i \in N} \operatorname{argmin}_{d_i \in C_i} \hat{\rho}_i(d_i). \quad \blacksquare$$

## 5.7 Risk Dominance

When it comes to the equilibria in pure strategies of a  $2 \times 2$  game in strategic form, there is a strong relationship between least-squares regret and the concept of risk dominance proposed by Harsanyi and Selten (1988). In this section, we examine this relationship more closely.

Particularly significant here is the well-known fact that when it comes to the equilibria in pure strategies of a  $2 \times 2$  game in strategic form, risk dominance is equivalent to  $1/2$ -dominance (Harsanyi and Selten, 1988; Morris, Rob, and Shin, 1995; Fudenberg and Levine, 1998). This equivalence is defined formally as follows.

Let  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  be any finite two-person game in strategic form such that, for each player  $i$  in  $N$ , it is the case that  $|C_i| = 2$ . For any pure-strategy profile  $(c_1, c_2)$  in  $C_1 \times C_2$  such that  $(c_1, c_2)$  is an equilibrium of  $\Gamma$ , the pure-strategy profile  $(c_1, c_2)$  is risk dominant if and only if

$$\sum_{c_{-i} \in C_{-i}} 0.5 u_i(c_{-i}, c_i) \geq \sum_{c_{-i} \in C_{-i}} 0.5 u_i(c_{-i}, d_i), \quad \forall i \in N, \quad \forall d_i \in C_i$$

or, more simply,

$$\sum_{c_{-i} \in C_{-i}} u_i(c_{-i}, c_i) \geq \sum_{c_{-i} \in C_{-i}} u_i(c_{-i}, d_i), \quad \forall i \in N, \quad \forall d_i \in C_i.$$

Intuitively, for any pure-strategy profile  $(c_1, c_2)$  in  $C_1 \times C_2$  such that  $(c_1, c_2)$  is an equilibrium, the pure-strategy profile  $(c_1, c_2)$  is risk dominant if and only if, for each player  $i$  in  $N$ , the pure strategy  $c_i$  is a best response to the uniform randomization between the two pure strategies of the other player.

The following theorem establishes that when it comes to the equilibria in pure strategies of a  $2 \times 2$  game in strategic form, least-squares regret is equivalent to risk dominance.

**THEOREM 5.7.1.** *Let  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  be any finite two-person game in strategic form such that, for each player  $i$  in  $N$ , it is the case that  $|C_i| = 2$ . Then, for any pure-strategy profile  $(c_1, c_2)$  in  $C_1 \times C_2$  such that  $(c_1, c_2)$  is an equilibrium of  $\Gamma$ , the pure-strategy profile  $(c_1, c_2)$  is risk dominant if and only if  $(c_1, c_2)$  is a least-squares regret profile in pure strategies.*

*Proof.* For concreteness, for any player  $i$  in  $N$ , let  $C_i = \{x_i, y_i\}$ . Consider any player  $i$  in  $N$ . The game  $\Gamma$  can be shown as in Table 5.6.

Table 5.6 Payoffs of player  $i$  in a game in strategic form

	$-i$	$x_{-i}$	$y_{-i}$
$i$		$x_i$	$y_i$
$x_i$		$u_i(x_{-i}, x_i)$	$u_i(y_{-i}, x_i)$
$y_i$		$u_i(x_{-i}, y_i)$	$u_i(y_{-i}, y_i)$

Let  $(x_{-i}, x_i)$  be an equilibrium of  $\Gamma$ . Notice that there may be more than one equilibrium. Since  $(x_{-i}, x_i)$  is an equilibrium,  $u_i(x_{-i}, x_i)$  is the best-response payoff with respect to  $x_{-i}$ .

Suppose that the equilibrium  $(x_{-i}, x_i)$  is risk dominant. There are two cases to consider depending on whether  $u_i(y_{-i}, x_i)$  or  $u_i(y_{-i}, y_i)$  is the best-response payoff with respect to  $y_{-i}$ .

Suppose first that  $u_i(y_{-i}, x_i)$  is the best-response payoff with respect to  $y_{-i}$ . Clearly,

$$u_i(x_{-i}, x_i) + u_i(y_{-i}, x_i) \geq u_i(x_{-i}, y_i) + u_i(y_{-i}, y_i)$$

and  $0 = \rho_i(x_i) \leq \rho_i(y_i)$ . Indeed, this conclusion follows immediately also from the fact that  $x_i$  is uniformly dominant for player  $i$  and Proposition 5.5.1 in Section 5.5.

Now, suppose that  $u_i(y_{-i}, y_i)$  is the best-response payoff with respect to  $y_{-i}$ . Then

$$\begin{aligned} \rho_i(x_i) &= (u_i(y_{-i}, y_i) - u_i(y_{-i}, x_i))^2 \quad \text{and} \\ \rho_i(y_i) &= (u_i(x_{-i}, x_i) - u_i(x_{-i}, y_i))^2. \end{aligned}$$

Since  $(x_{-i}, x_i)$  is risk dominant,

$$u_i(x_{-i}, x_i) + u_i(y_{-i}, x_i) \geq u_i(x_{-i}, y_i) + u_i(y_{-i}, y_i).$$

Thus,

$$u_i(x_{-i}, x_i) - u_i(x_{-i}, y_i) \geq u_i(y_{-i}, y_i) - u_i(y_{-i}, x_i) \geq 0,$$

and so,  $\rho_i(x_i) \leq \rho_i(y_i)$ .

Thus, the equilibrium  $(x_{-i}, x_i)$  is a least-squares regret profile in pure strategies.

Now, suppose that the equilibrium  $(x_{-i}, x_i)$  is a least-squares regret profile in pure strategies. There are two cases to consider depending on whether  $u_i(y_{-i}, x_i)$  or  $u_i(y_{-i}, y_i)$  is the best-response payoff with respect to  $y_{-i}$ .

Suppose first that  $u_i(y_{-i}, x_i)$  is the best-response payoff with respect to  $y_{-i}$ . Clearly,  $0 = \rho_i(x_i) \leq \rho_i(y_i)$  and

$$u_i(x_{-i}, x_i) + u_i(y_{-i}, x_i) \geq u_i(x_{-i}, y_i) + u_i(y_{-i}, y_i).$$

Now, suppose that  $u_i(y_{-i}, y_i)$  is the best-response payoff with respect to  $y_{-i}$ . Then

$$\begin{aligned} \rho_i(x_i) &= (u_i(y_{-i}, y_i) - u_i(y_{-i}, x_i))^2 \quad \text{and} \\ \rho_i(y_i) &= (u_i(x_{-i}, x_i) - u_i(x_{-i}, y_i))^2. \end{aligned}$$

Since  $(x_{-i}, x_i)$  is a least-squares regret profile in pure strategies,  $\rho_i(x_i) \leq \rho_i(y_i)$ . Thus,

$$u_i(x_{-i}, x_i) - u_i(x_{-i}, y_i) \geq u_i(y_{-i}, y_i) - u_i(y_{-i}, x_i) \geq 0,$$

and so,

$$u_i(x_{-i}, x_i) + u_i(y_{-i}, x_i) \geq u_i(x_{-i}, y_i) + u_i(y_{-i}, y_i).$$

Thus, the equilibrium  $(x_{-i}, x_i)$  is risk dominant. ■

For an illustration of Theorem 5.7.1 and the equivalence that it describes, consider the  $2 \times 2$  game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.7.

Table 5.7 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	2 3	1 0
	$y_1$	0 2	2 2

This game has three equilibria:  $(x_1, x_2)$ , which gives the payoff allocation  $(3, 2)$ ;  $(y_1, y_2)$ , which gives the payoff allocation  $(2, 2)$ ; and

$$((2/3)x_1 + (1/3)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(2, \frac{4}{3})$ . Notably, the equilibrium  $(x_1, x_2)$  is payoff dominant, and the equilibrium  $(y_1, y_2)$  is risk dominant.

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is precisely the risk-dominant equilibrium  $(y_1, y_2)$ .

Least-squares regret with respect to randomized strategies yields a similar solution. The unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, 0.2x_2 + 0.8y_2),$$

which gives the payoff allocation  $(1.72, 1.52)$ .

This solution is quite reasonable. In particular, it comes quite close to being the risk-dominant equilibrium. Still, this solution is somewhat unsatisfactory in that it is worse for each player than the risk-dominant equilibrium that is achieved by considering only pure strategies. But this is to be expected: hedging via randomization comes at a price.

Of course, how well a least-squares regret profile in randomized strategies fares relative to a risk-dominant equilibrium in pure strategies depends on the other payoffs that can be achieved. For an illustration of this point, consider the  $2 \times 2$  game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 5.8, which differs from the game shown in Table 5.7 only in that  $u_1(x_1, x_2)$  is increased from 3 to 10 and  $u_1(y_1, x_2)$  is increased from 2 to 9.

Table 5.8 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	2 10	1 0
	$y_1$	0 9	2 2

In this game, as Theorem 5.7.1 establishes, the unique least-squares regret profile in pure strategies is again precisely the risk-dominant equilibrium  $(y_1, y_2)$ , which gives the payoff allocation  $(2, 2)$ , and the unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, 0.2x_2 + 0.8y_2),$$

which gives the payoff allocation  $(3.12, 1.52)$ . Notably, this solution in randomized strategies is better for player 1 than the risk-dominant equilibrium that is achieved by considering only pure strategies.

## 5.8 Convexity

As discussed in Sections 2.2, 2.5, and 3.2, we suppose the squaring of regret for technical reasons and for mathematical convenience. As noted earlier, one agreeable consequence of squaring the regret is that the regret function in randomized strategies is convex. Thus, minimization of the regret function in randomized strategies defines a convex set of randomized least-squares regret strategies. It follows that, for any finite game in strategic form, the set of least-squares regret profiles in randomized strategies is convex. The following theorem establishes these facts.

**THEOREM 5.8.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Then the set*

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

*is convex.*

*Proof.* Consider any player  $i$  in  $N$ . Let  $\sigma_i$  and  $\hat{\sigma}_i$  be any randomized strategies in  $\Delta(C_i)$ . Let  $\lambda$  be any real number such that  $0 \leq \lambda \leq 1$ , and let  $\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i$  be the randomized strategy in  $\Delta(C_i)$  such that

$$(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i)(c_i) = \lambda\sigma_i(c_i) + (1 - \lambda)\hat{\sigma}_i(c_i), \quad \forall c_i \in C_i.$$

Then

$$\rho_i(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) = \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) \right)^2.$$

But notice that

$$u_i(c_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) = \lambda u_i(c_{-i}, \sigma_i) + (1 - \lambda) u_i(c_{-i}, \hat{\sigma}_i).$$

And so,

$$\begin{aligned} & \rho_i(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) \\ &= \sum_{c_{-i} \in C_{-i}} \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \hat{\sigma}_i) + \lambda(u_i(c_{-i}, \hat{\sigma}_i) - u_i(c_{-i}, \sigma_i)) \right)^2. \end{aligned}$$

Now, consider the function  $f_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$f_i(\lambda) = \rho_i(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i).$$

Notice that

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} = 2 \sum_{c_{-i} \in C_{-i}} (u_i(c_{-i}, \hat{\sigma}_i) - u_i(c_{-i}, \sigma_i))^2.$$



Clearly,

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} \geq 0, \quad \forall \lambda \in [0, 1],$$

and so, the function  $f_i$  is convex.

Now, notice that

$$f_i(0) = \rho_i(0\sigma_i + (1-0)\hat{\sigma}_i) = \rho_i(\hat{\sigma}_i)$$

and

$$f_i(1) = \rho_i(1\sigma_i + (1-1)\hat{\sigma}_i) = \rho_i(\sigma_i).$$

Observe that the line connecting the points  $(0, \rho_i(\hat{\sigma}_i))$  and  $(1, \rho_i(\sigma_i))$  in  $\mathbf{R}^2$  is just the function  $g_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$g_i(\lambda) = \lambda\rho_i(\sigma_i) + (1-\lambda)\rho_i(\hat{\sigma}_i).$$

But since the function  $f_i$  is convex,

$$f_i(\lambda) = \rho_i(\lambda\sigma_i + (1-\lambda)\hat{\sigma}_i) \leq g_i(\lambda) = \lambda\rho_i(\sigma_i) + (1-\lambda)\rho_i(\hat{\sigma}_i), \quad \forall \lambda \in [0, 1].$$

Thus, the regret function  $\rho_i$  is convex. And so, the set

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is convex.

Thus, the set

$$\prod_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is convex since it is the Cartesian product of convex sets. ■

Convexity is appealing for several reasons. Convexity of the regret function in randomized strategies implies that computation of randomized least-squares regret strategies is a convex optimization problem, which can be solved efficiently using standard well-developed techniques. Furthermore, convexity of the regret function in randomized strategies implies that the player has a convex set of randomized least-squares regret strategies, each of which is considered reasonable, and can randomize over such strategies in any manner.

## 5.9 Uniqueness

Theorem 5.8.1 in Section 5.8 establishes that, for any finite game in strategic form, the set of least-squares regret profiles in randomized strategies is convex. But this means that a finite game in strategic form may have infinitely many solutions. The following theorem describes a sufficient condition for the uniqueness of a least-squares regret profile in randomized strategies.

**THEOREM 5.9.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. If, for every player  $i$  in  $N$  and for any randomized strategies  $\sigma_i$  and  $\hat{\sigma}_i$  in  $\Delta(C_i)$  such that  $\sigma_i \neq \hat{\sigma}_i$ , there exists some partial profile  $c_{-i}$  in  $C_{-i}$  such that*

$$u_i(c_{-i}, \sigma_i) \neq u_i(c_{-i}, \hat{\sigma}_i),$$

then the set

$$\bigtimes_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is a singleton.

*Proof.* Consider any player  $i$  in  $N$ . Let  $\sigma_i$  and  $\hat{\sigma}_i$  be any randomized strategies in  $\Delta(C_i)$  such that  $\sigma_i \neq \hat{\sigma}_i$ . Let  $\lambda$  be any real number such that  $0 \leq \lambda \leq 1$ , and let  $\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i$  be the randomized strategy in  $\Delta(C_i)$  such that

$$(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i)(c_i) = \lambda\sigma_i(c_i) + (1 - \lambda)\hat{\sigma}_i(c_i), \quad \forall c_i \in C_i.$$

Just as in the proof of Theorem 5.8.1 in Section 5.8, consider the function  $f_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$f_i(\lambda) = \rho_i(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i).$$

Recall that

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} = 2 \sum_{c_{-i} \in C_{-i}} (u_i(c_{-i}, \hat{\sigma}_i) - u_i(c_{-i}, \sigma_i))^2.$$

By assumption, there exists some partial profile  $c_{-i}$  in  $C_{-i}$  such that

$$u_i(c_{-i}, \sigma_i) \neq u_i(c_{-i}, \hat{\sigma}_i).$$

Thus,

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} > 0, \quad \forall \lambda \in [0, 1],$$

and so, the function  $f_i$  is strictly convex.

Now, notice that

$$f_i(0) = \rho_i(0\sigma_i + (1 - 0)\hat{\sigma}_i) = \rho_i(\hat{\sigma}_i)$$

and

$$f_i(1) = \rho_i(1\sigma_i + (1-1)\hat{\sigma}_i) = \rho_i(\sigma_i).$$

Observe that the line connecting the points  $(0, \rho_i(\hat{\sigma}_i))$  and  $(1, \rho_i(\sigma_i))$  in  $\mathbf{R}^2$  is just the function  $g_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$g_i(\lambda) = \lambda\rho_i(\sigma_i) + (1-\lambda)\rho_i(\hat{\sigma}_i).$$

But since the function  $f_i$  is strictly convex,

$$f_i(\lambda) = \rho_i(\lambda\sigma_i + (1-\lambda)\hat{\sigma}_i) < g_i(\lambda) = \lambda\rho_i(\sigma_i) + (1-\lambda)\rho_i(\hat{\sigma}_i), \quad \forall \lambda \in (0, 1).$$

Thus, the regret function  $\rho_i$  is strictly convex. And so, the set

$$\underset{\tau_i \in \Delta(C_i)}{\operatorname{argmin}} \rho_i(\tau_i)$$

is a singleton since, by Theorem 5.1.1 in Section 5.1, it is nonempty and since, as just shown, the regret function  $\rho_i$  is strictly convex.

Thus, the set

$$\prod_{i \in N} \underset{\tau_i \in \Delta(C_i)}{\operatorname{argmin}} \rho_i(\tau_i)$$

is a singleton since it is the Cartesian product of singleton sets. ■

In the case of a small finite game in strategic form in which each player has at most two pure strategies, the sufficient condition for uniqueness is far simpler.

**THEOREM 5.9.2.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form such that, for every player  $i$  in  $N$ , it is the case that  $|C_i| \leq 2$ . If, for every player  $i$  in  $N$ , it is the case that  $|C_i| = 1$  or, for any pure strategies  $c_i$  and  $\hat{c}_i$  in  $C_i$  such that  $c_i \neq \hat{c}_i$ , there exists some partial profile  $c_{-i}$  in  $C_{-i}$  such that*

$$u_i(c_{-i}, c_i) \neq u_i(c_{-i}, \hat{c}_i),$$

then the set

$$\prod_{i \in N} \underset{\tau_i \in \Delta(C_i)}{\operatorname{argmin}} \rho_i(\tau_i)$$

is a singleton.

*Proof.* Consider any player  $i$  in  $N$ . Suppose that  $|C_i| = 1$ . Then, clearly,  $|\Delta(C_i)| = 1$ . And so, the set

$$\underset{\tau_i \in \Delta(C_i)}{\operatorname{argmin}} \rho_i(\tau_i)$$

is a singleton.

Now, suppose that  $|C_i| = 2$ . Let  $c_i$  and  $\hat{c}_i$  be any pure strategies in  $C_i$  such that  $c_i \neq \hat{c}_i$ . By assumption, there exists some partial profile  $c_{-i}$  in  $C_{-i}$  such that

$$u_i(c_{-i}, c_i) \neq u_i(c_{-i}, \hat{c}_i).$$

Let  $\sigma_i$  and  $\hat{\sigma}_i$  be any randomized strategies in  $\Delta(C_i)$  such that  $\sigma_i \neq \hat{\sigma}_i$ . Then

$$\begin{aligned} u_i(c_{-i}, \sigma_i) &= \sigma_i(c_i)u_i(c_{-i}, c_i) + \sigma_i(\hat{c}_i)u_i(c_{-i}, \hat{c}_i) \\ &= \sigma_i(c_i)u_i(c_{-i}, c_i) + (1 - \sigma_i(c_i))u_i(c_{-i}, \hat{c}_i) \\ &= \sigma_i(c_i)(u_i(c_{-i}, c_i) - u_i(c_{-i}, \hat{c}_i)) + u_i(c_{-i}, \hat{c}_i) \end{aligned}$$

and

$$\begin{aligned} u_i(c_{-i}, \hat{\sigma}_i) &= \hat{\sigma}_i(c_i)u_i(c_{-i}, c_i) + \hat{\sigma}_i(\hat{c}_i)u_i(c_{-i}, \hat{c}_i) \\ &= \hat{\sigma}_i(c_i)u_i(c_{-i}, c_i) + (1 - \hat{\sigma}_i(c_i))u_i(c_{-i}, \hat{c}_i) \\ &= \hat{\sigma}_i(c_i)(u_i(c_{-i}, c_i) - u_i(c_{-i}, \hat{c}_i)) + u_i(c_{-i}, \hat{c}_i). \end{aligned}$$

But

$$\sigma_i(c_i) \neq \hat{\sigma}_i(c_i)$$

and

$$u_i(c_{-i}, c_i) - u_i(c_{-i}, \hat{c}_i) \neq 0.$$

And so,

$$\sigma_i(c_i)(u_i(c_{-i}, c_i) - u_i(c_{-i}, \hat{c}_i)) + u_i(c_{-i}, \hat{c}_i) \neq \hat{\sigma}_i(c_i)(u_i(c_{-i}, c_i) - u_i(c_{-i}, \hat{c}_i)) + u_i(c_{-i}, \hat{c}_i).$$

Thus, there exists some partial profile  $c_{-i}$  in  $C_{-i}$  such that

$$u_i(c_{-i}, \sigma_i) \neq u_i(c_{-i}, \hat{\sigma}_i).$$

And so, by Theorem 5.9.1, the set

$$\operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is a singleton.

Thus, the set

$$\prod_{i \in N} \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\tau_i)$$

is a singleton since it is the Cartesian product of singleton sets. ■

Intuition suggests that, for any generic finite game in strategic form, the set of least-squares regret profiles in randomized strategies is a singleton. It would be interesting to verify this conjecture.

# 6

---

## Further Examples

One significant question that arises when appraising a solution concept is whether or not it yields the expected or reasonable solutions for the games to which it is applied and, in particular, for games beyond those that inspired its development.

In this chapter, we move beyond the Traveler's Dilemma and show that least-squares regret yields reasonable solutions for many other well-known games—and, for some, solutions that are even more reasonable than those yielded by standard solution concepts—and unsatisfactory solutions for some. Section 6.1 studies the Dollar Auction game and shows how least-squares regret yields a reasonable solution and outperforms Nash equilibrium. Section 6.2 studies Bertrand competition and shows how least-squares regret yields a reasonable solution in line with the experimental evidence and outperforms Nash equilibrium. Section 6.3 studies inspection games and shows how least-squares regret, unlike Nash equilibrium, captures certain intuitive and experimentally robust effects. Section 6.4 studies Matching Pennies games and evaluates the successes and failures of least-squares regret and Nash equilibrium. Section 6.5 studies the Chicken game and shows how least-squares regret yields reasonable solutions and outperforms Nash equilibrium. Section 6.6 studies coordination games, for which least-squares regret yields reasonable solutions. Section 6.7 studies the Battle of the Sexes game, for which least-squares regret yields unsatisfactory solutions. Section 6.8 studies the two-person bargaining problem and presents two new theorems.

### 6.1 Dollar Auction

Consider the *Dollar Auction* game (Raiffa, 1982). In this game, a single dollar is up for auction. There are two risk-neutral players, and each must privately and independently submit a bid that can be any real number between 0 and 1. The high bidder wins the auction and pays the amount of his bid. The low bidder loses

the auction and pays nothing. In the event of a tie, each player has a probability of 0.5 of winning the auction and paying the amount of his bid and a probability of 0.5 of losing the auction and paying nothing. Thus, in this game, the set of players is  $N = \{1, 2\}$ , the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 0 \leq x \leq 1\},$$

and, for any player  $i$  in  $N$ , the utility function is

$$\begin{aligned} u_i(c_1, c_2) &= 0 && \text{if } i \notin \operatorname{argmax}_{j \in \{1,2\}} c_j, \\ &= 1 - c_i && \text{if } \{i\} = \operatorname{argmax}_{j \in \{1,2\}} c_j, \\ &= (1 - c_i)/2 && \text{if } c_1 = c_2. \end{aligned}$$

In the unique equilibrium of this game, each player bids 1 for sure and receives a net payoff of 0. For any other profile of bids, there is always an incentive for one player to outbid the other player by the slimmest margin. In this way, the Dollar Auction game models the dynamics of bidding up the price of an item in an auction.

But this solution seems unreasonable. It seems that no player would ever bid 1. Indeed, for each player, bidding 1 is a weakly dominated strategy—one that is, in fact, weakly dominated by every other strategy—since it is the only bid that cannot yield a positive payoff and no bid can yield a negative payoff. Thus, it is difficult to recommend playing the equilibrium strategy of bidding 1 for sure. Any other bid would be better. If there is any chance at all of the other player, for whatever reason, bidding strictly less than 1, it would be rational for one likewise to bid some amount strictly less than 1. Here, the equilibrium strategy seems wrongheaded.

By contrast, least-squares regret yields a more reasonable solution. Consider any player  $i$  in  $N$  and any bid  $c_i$  in  $C_i$ . If the other player chooses some bid  $c_{-i}$  in  $C_{-i}$  that is less than  $c_i$ , then the payoff to player  $i$  from choosing  $c_i$  is  $1 - c_i$  while the best-response payoff is  $1 - c_{-i}$  (since the best response is to outbid the other player by the slimmest margin). If the other player chooses some bid  $c_{-i}$  in  $C_{-i}$  that is greater than  $c_i$ , then the payoff to player  $i$  from choosing  $c_i$  is 0 while the best-response payoff is again  $1 - c_{-i}$ . Thus, for any player  $i$  in  $N$ ,

$$\begin{aligned} \rho_i(c_i) &= \int_0^{c_i} ((1 - c_{-i}) - (1 - c_i))^2 dc_{-i} + \int_{c_i}^1 ((1 - c_{-i}) - 0)^2 dc_{-i} \\ &= \int_0^{c_i} (c_i - c_{-i})^2 dc_{-i} + \int_{c_i}^1 (1 - c_{-i})^2 dc_{-i} \\ &= c_i^2 - c_i + 1/3, \end{aligned}$$

and so,

$$\frac{d(\rho_i(c_i))}{dc_i} = 2c_i - 1.$$

Setting the derivative equal to 0 and solving for  $c_i$  yields the unique pure least-squares regret strategy  $c_i = 0.5$ . Such a bid, being strictly less than 1, is more reasonable than the equilibrium strategy: it always does at least as well as the equilibrium strategy, and with respect to half of the pure-strategy set of the other player, it does strictly better. Such a bid makes intuitive sense also: it balances the need to bid high enough to win the auction and the need to bid low enough to avoid a costly victory.

Of course, the unique least-squares regret profile in pure strategies  $(0.5, 0.5)$ , while reasonable, cannot be an equilibrium since each player, should he expect the other player to choose 0.5, would prefer to outbid the other player by the slimmest margin. As discussed in Sections 1.2, 1.4, 2.2, and 2.5, least-squares regret disregards fully strategic reasoning, and this disregard is an obvious point of criticism. We discuss this issue in Section 7.4 and address it in Chapter 8.

## 6.2 Bertrand Competition

Consider *Bertrand competition*, which models the price-setting behaviors of competitors in a market (Bertrand, 1883). In this game, there are two risk-neutral players competing in a market to sell 100 units of some homogeneous commodity, and each must privately and independently set a price to publicize that can be any real number between 0 and 200. Assume that the players have no costs and that consumers choose on the basis of price alone. The player setting the lower price wins the entire market and sells all 100 units at his chosen price. The player setting the higher price loses the entire market and sells nothing. If the two players set the same price, they split the market evenly. Thus, in this game, the set of players is  $N = \{1, 2\}$ , the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 0 \leq x \leq 200\},$$

and, for any player  $i$  in  $N$ , the utility function is

$$\begin{aligned} u_i(c_1, c_2) &= 0 && \text{if } i \notin \operatorname{argmin}_{j \in \{1, 2\}} c_j, \\ &= 100c_i && \text{if } \{i\} = \operatorname{argmin}_{j \in \{1, 2\}} c_j, \\ &= 50c_i && \text{if } c_1 = c_2. \end{aligned}$$

One may notice several striking similarities between Bertrand competition and the Dollar Auction game in Section 6.1. This is not surprising. The two games are

more or less structurally equivalent. More precisely, Bertrand competition, as it has been described here, can be seen as a kind of reverse auction.

In the unique equilibrium of this game, each player sets a price of 0 for sure and receives a net payoff of 0. For any other profile of prices, there is always an incentive for one player to undercut the other player by the slimmest margin. In this way, Bertrand competition models the dynamics of price wars or convergence to a market equilibrium.

But this solution seems unreasonable. It seems that no player would ever set a price of 0. Indeed, for each player, setting a price of 0 is a weakly dominated strategy—one that is, in fact, weakly dominated by every other strategy—since it is the only price that cannot yield a positive payoff and no price can yield a negative payoff. Thus, it is difficult to recommend playing the equilibrium of strategy of setting a price of 0 for sure. Any other price would be better. If there is any chance at all of the other player, for whatever reason, setting a price strictly greater than 0, it would be rational for one likewise to set a price strictly greater than 0. Here, the equilibrium strategy seems wrongheaded.

By contrast, least-squares regret yields a more reasonable solution. Consider any player  $i$  in  $N$  and any price  $c_i$  in  $C_i$ . If the other player chooses some price  $c_{-i}$  in  $C_{-i}$  that is less than  $c_i$ , then the payoff to player  $i$  from choosing  $c_i$  is 0 while the best-response payoff is  $100c_{-i}$  (since the best response is to undercut the other player by the slimmest margin). If the other player chooses some price  $c_{-i}$  in  $C_{-i}$  that is greater than  $c_i$ , then the payoff to player  $i$  from choosing  $c_i$  is  $100c_i$  while the best-response payoff is again  $100c_{-i}$ . Thus, for any player  $i$  in  $N$ ,

$$\begin{aligned} \rho_i(c_i) &= \int_0^{c_i} (100c_{-i} - 0)^2 dc_{-i} + \int_{c_i}^{200} (100c_{-i} - 100c_i)^2 dc_{-i} \\ &= 100^2 \int_0^{c_i} c_{-i}^2 dc_{-i} + 100^2 \int_{c_i}^{200} (c_{-i} - c_i)^2 dc_{-i} \\ &= 100^2(200c_i^2 - 200^2c_i + 200^3/3), \end{aligned}$$

and so,

$$\frac{d(\rho_i(c_i))}{dc_i} = 100^2(400c_i - 200^2).$$

Setting the derivative equal to 0 and solving for  $c_i$  yields the unique pure least-squares regret strategy  $c_i = 100$ . Such a price, being strictly greater than 0, is more reasonable than the equilibrium strategy: it always does at least as well as the equilibrium strategy, and with respect to half of the pure-strategy set of the other player, it does strictly better. Such a price makes intuitive sense also: it balances the need to set a price low enough to capture the market and the need to set a price high enough to avoid costly underpricing. Notably, the solution yielded



by least-squares regret agrees with the experimental evidence, which shows that people consistently choose prices well above that specified by the equilibrium strategy (Dufwenberg and Gneezy, 2000).

Of course, the unique least-squares regret profile in pure strategies (100, 100), while reasonable, cannot be an equilibrium since each player, should he expect the other player to choose 100, would prefer to undercut the other player by the slimmest margin. As discussed in Sections 1.2, 1.4, 2.2, and 2.5, least-squares regret disregards fully strategic reasoning, and this disregard is an obvious point of criticism. We discuss this issue in Section 7.4 and address it in Chapter 8.

### 6.3 Inspection Games

Consider the finite two-person inspection game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 6.1 (Fudenberg and Tirole, 1991; von Stengel, 2011).

Table 6.1 An inspection game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	4
	$y_1$	-1	-2

This game models the dynamics of compliance inspection. In this game, player 1 is the inspector, and player 2 is the inspectee. Player 1 has two pure strategies: the “abstain” strategy  $x_1$  and the “inspect” strategy  $y_1$ . Player 2 has two pure strategies: the “comply” strategy  $x_2$  and the “cheat” strategy  $y_2$ . The aim of player 1 is to ensure that player 2 complies with the regulations and to catch any violations, but inspections are costly. Player 2 has an incentive to cheat, but getting caught is very costly.

The unique equilibrium of this game is

$$((5/7)x_1 + (2/7)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(-4/3, 0)$ . Notably, player 1 favors the “abstain” strategy  $x_1$ , and player 2 favors the “comply” strategy  $x_2$ .

Least-squares regret yields a quite different solution. The unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, (25/29)x_2 + (4/29)y_2),$$

which gives the payoff allocation  $(-148/145, -144/145)$  or roughly  $(-1.02, -0.99)$ . Notably, player 1 favors the “inspect” strategy  $y_1$ , and player 2 favors the “comply” strategy  $x_2$ .

Now, consider the two-person inspection game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  shown in Table 6.2, which differs from the game shown in Table 6.1 only in that  $u_2(y_1, y_2)$  is decreased from  $-10$  to  $-20$ . This change models an increase in the cost of getting caught.

Table 6.2 An inspection game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	4
	$y_1$	0	-20
		-1	-2

The unique equilibrium of this game is

$$((5/6)x_1 + (1/6)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(-4/3, 0)$ .

In equilibrium, each player randomizes so as to render the other player indifferent between his strategies, and so, the behavior of a player depends not on his own payoffs, but on the payoffs of the other player. Thus, increasing the cost of getting caught reduces the probability of the “inspect” strategy  $y_1$  being chosen while having no effect on the behavior of player 2.

But this is peculiar. The cost of getting caught is greater in this game than in the game shown in Table 6.1. Thus, one would expect the probability of the “cheat” strategy  $y_2$  being chosen to be lower in this game than in the game shown in Table 6.1. More generally, one would expect a player to be sensitive to his own payoffs in the natural way.

Such expectations are not only reasonable; they are well supported. Indeed, *own-payoff effects*—and the failure of Nash equilibrium to capture them—are well recognized and experimentally robust (Ochs, 1995; McKelvey, Palfrey, and Weber, 2000; Goeree and Holt, 2001; Goeree, Holt, and Palfrey, 2003). Furthermore, with respect to inspection games of the sort considered here, own-payoff effects are known to play a role in determining behavior (Nosenzo, Offerman, Sefton, and van der Veen, 2014).

Now, consider least-squares regret. The unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, (25/26)x_2 + (1/26)y_2),$$

which gives the payoff allocation  $(-56/65, -38/65)$  or roughly  $(-0.86, -0.58)$ .

According to least-squares regret, the behavior of a player depends on his own payoffs. Thus, increasing the cost of getting caught reduces the probability of the “cheat” strategy  $y_2$  being chosen, as expected, while also having no effect on the behavior of player 1.

Thus, least-squares regret is notable here for yielding solutions in line with intuition and the experimental evidence, outperforming Nash equilibrium, and capturing the experimentally robust own-payoff effects described earlier.

## 6.4 Matching Pennies

Consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form known as *Matching Pennies* shown in Table 6.3.

Table 6.3 Matching Pennies game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	40	80
	$y_1$	80	40

The unique equilibrium of this game is

$$(0.5x_1 + 0.5y_1, 0.5x_2 + 0.5y_2),$$

which gives the payoff allocation  $(60, 60)$ .

As might be expected, the experimental evidence confirms this solution (Goeree and Holt, 2001).

Least-squares regret agrees with both Nash equilibrium and the experimental evidence. For each player  $i$  in  $N$ , the unique randomized least-squares regret strategy for player  $i$  is  $0.5x_i + 0.5y_i$ .

Now, consider the asymmetric Matching Pennies game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  shown in Table 6.4, which differs from the game shown in Table 6.3 only in that  $u_1(x_1, x_2)$  is increased from 80 to 320.

The unique equilibrium of this game is

$$(0.5x_1 + 0.5y_1, (1/8)x_2 + (7/8)y_2),$$

which gives the payoff allocation  $(75, 60)$ .

Table 6.4 An asymmetric Matching Pennies game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	40 320	80 40
	$y_1$	80 40	40 80

In equilibrium, each player randomizes so as to render the other player indifferent between his strategies, and so, the behavior of a player depends not on his own payoffs, but on the payoffs of the other player. Thus, increasing  $u_1(x_1, x_2)$  from 80 to 320 reduces the probability of  $x_2$  being chosen while having no effect on the behavior of player 1.

But this is peculiar. The payoff  $u_1(x_1, x_2)$  is greater in this game than in the game shown in Table 6.3. Thus, one would expect the probability of  $x_1$  being chosen to be greater in this game than in the game shown in Table 6.3. More generally, one would expect a player to be sensitive to his own payoffs in the natural way. In fact, the experimental evidence confirms these intuitions: in this game,  $x_1$  is chosen with probability 0.96 (Goeree and Holt, 2001).

Now, consider least-squares regret. The unique least-squares regret profile in randomized strategies is

$$(0.98x_1 + 0.02y_1, 0.5x_2 + 0.5y_2),$$

which gives the payoff allocation  $(\frac{888}{5}, 60)$  or  $(177.6, 60)$ .

According to least-squares regret, the behavior of a player depends on his own payoffs. Thus, increasing  $u_1(x_1, x_2)$  from 80 to 320 increases the probability of  $x_1$  being chosen, as expected, while also having no effect on the behavior of player 2. Notably, the unique randomized least-squares regret strategy for player 1 involves choosing  $x_1$  with probability 0.98, which very nearly matches the observed probability of 0.96.

Now, consider the asymmetric Matching Pennies game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  shown in Table 6.5, which differs from the game shown in Table 6.3 only in that  $u_1(x_1, x_2)$  is decreased from 80 to 44.

The unique equilibrium of this game is

$$(0.5x_1 + 0.5y_1, (\frac{10}{11})x_2 + (\frac{1}{11})y_2),$$

which gives the payoff allocation  $(\frac{480}{11}, 60)$  or roughly  $(43.64, 60)$ .

Again, in equilibrium, each player randomizes so as to render the other player indifferent between his strategies, and so, the behavior of a player depends not on

Table 6.5 An asymmetric Matching Pennies game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	40 44	80 40
	$y_1$	80 40	40 80

his own payoffs, but on the payoffs of the other player. Thus, decreasing  $u_1(x_1, x_2)$  from 80 to 44 increases the probability of  $x_2$  being chosen while having no effect on the behavior of player 1.

But again, this is peculiar. The payoff  $u_1(x_1, x_2)$  is lower in this game than in the game shown in Table 6.3. Thus, one would expect the probability of  $x_1$  being chosen to be lower in this game than in the game shown in Table 6.3. And again, more generally, one would expect a player to be sensitive to his own payoffs in the natural way. In fact, the experimental evidence confirms these intuitions: in this game,  $x_1$  is chosen with probability 0.08 (Goeree and Holt, 2001).

Now, consider least-squares regret. The unique least-squares regret profile in randomized strategies is

$$((1/101)x_1 + (100/101)y_1, 0.5x_2 + 0.5y_2),$$

which gives the payoff allocation  $(^{6042/101}, 60)$  or roughly  $(59.82, 60)$ .

Again, according to least-squares regret, the behavior of a player depends on his own payoffs. Thus, decreasing  $u_1(x_1, x_2)$  from 80 to 44 decreases the probability of  $x_1$  being chosen, as expected, while also having no effect on the behavior of player 2. Notably, the unique randomized least-squares regret strategy for player 1 involves choosing  $x_1$  with probability  $1/101$ , which matches fairly closely the observed probability of 0.08.

Thus, least-squares regret is notable here for yielding solutions in line with intuition and the experimental evidence, outperforming Nash equilibrium, and capturing the experimentally robust own-payoff effects (Ochs, 1995; McKelvey, Palfrey, and Weber, 2000; Goeree and Holt, 2001; Goeree, Holt, and Palfrey, 2003).

But this is not to say that least-squares regret yields solutions that are wholly consistent with the experimental evidence. Notice that in all three games above, the unique randomized least-squares regret strategy for player 2 is  $0.5x_2 + 0.5y_2$ . But the experimental evidence shows  $y_2$  being chosen with probability 0.84 in the game shown in Table 6.4 and  $x_2$  being chosen with probability 0.8 in the game shown in Table 6.5 (Goeree and Holt, 2001). Player 2 thus appears to anticipate

the strategy choices of player 1 and to respond accordingly in a manner consistent with Nash equilibrium.

Least-squares regret fails to capture the observed behavior simply because, as discussed in Sections 1.2, 1.4, 2.2, and 2.5, it disregards fully strategic reasoning. This disregard is an obvious point of criticism, and we discuss it in Section 7.4. Interesting questions remain, however, for example, why the behavior of player 1 is determined more by own-payoff effects than by considerations of what player 2 might do and how considerations of the behavior of other players can be incorporated into least-squares regret. We propose one remedy in Chapter 8.

## 6.5 Chicken

Consider the finite two-person anticoordination game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form known as *Chicken* shown in Table 6.6.

Table 6.6 Chicken game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	4	6
	$y_1$	1	-3

In this game, each player  $i$  in  $N$  has two pure strategies: the “cautious” strategy  $x_i$  and the “bold” strategy  $y_i$ . Each player would most prefer to be bold himself and the other to be cautious, but for each player to be bold would be catastrophic. The best symmetric outcome occurs when each player is cautious; importantly, this outcome is also efficient.

This game has three equilibria:  $(y_1, x_2)$ , which gives the payoff allocation  $(6, 1)$ ;  $(x_1, y_2)$ , which gives the payoff allocation  $(1, 6)$ ; and

$$((2/3)x_1 + (1/3)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(3, 3)$ .

Notice that the best symmetric outcome  $(x_1, x_2)$ , which gives the payoff allocation  $(4, 4)$ , is, in fact, also a remarkably good outcome. For each player, this outcome is only slightly worse than his most preferred outcome, notably better than the equilibrium in which he is cautious and the other bold, better than the unique symmetric equilibrium, and considerably better than the catastrophic

outcome in which each player is bold. And yet, this outcome, though both good and efficient, is unachievable by players who choose only equilibrium strategies.

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is precisely the best symmetric outcome  $(x_1, x_2)$ . Thus, least-squares regret with respect to pure strategies yields as the unique solution precisely the good and efficient outcome that is unachievable by players who choose only equilibrium strategies.

Least-squares regret with respect to randomized strategies yields a similar solution. The unique least-squares regret profile in randomized strategies is

$$(0.8x_1 + 0.2y_1, 0.8x_2 + 0.2y_2),$$

which gives the payoff allocation  $(3.56, 3.56)$ .

This solution is quite reasonable. In particular, it comes quite close to being the best symmetric outcome and is better for each player than the qualitatively similar symmetric equilibrium described above. Still, this solution is somewhat unsatisfactory in that it is worse for each player than the best symmetric outcome that is achieved by considering only pure strategies. But this is to be expected: hedging via randomization comes at a price.

## 6.6 Coordination Games

Consider the finite two-person coordination game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 6.7.

Table 6.7 A coordination game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	3	0
	$y_1$	0	1

This game has three equilibria:  $(x_1, x_2)$ , which gives the payoff allocation  $(3, 3)$ ;  $(y_1, y_2)$ , which gives the payoff allocation  $(1, 1)$ ; and

$$(0.25x_1 + 0.75y_1, 0.25x_2 + 0.75y_2),$$

which gives the payoff allocation  $(0.75, 0.75)$ . Notably, in the unique equilibrium in randomized strategies, each player  $i$  in  $N$  favors  $y_i$ , making the outcome  $(y_1, y_2)$  quite likely.

But these are peculiar results. It seems that the sole reasonable outcome is  $(x_1, x_2)$ , which is clearly superior to the rest. Thus, one would expect the outcome  $(x_1, x_2)$  to be the unique solution. And while it is possible to generate it as the unique solution using standard solution concepts, doing so requires an appeal to other principles such as the focal-point effect (Schelling, 1960) or payoff dominance (Harsanyi and Selten, 1988).

Furthermore, the efficient outcome  $(x_1, x_2)$  is available and clearly better for each player. Thus, one would expect each player  $i$  in  $N$  to favor  $x_i$ .

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is precisely the efficient outcome  $(x_1, x_2)$ . Thus, least-squares regret with respect to pure strategies yields as the unique solution precisely the sole reasonable outcome  $(x_1, x_2)$ , all without needing to appeal to other principles.

Least-squares regret with respect to randomized strategies yields a similar solution. The unique least-squares regret profile in randomized strategies is

$$(0.9x_1 + 0.1y_1, 0.9x_2 + 0.1y_2),$$

which gives the payoff allocation  $(2.44, 2.44)$ .

While this solution is not the efficient outcome  $(x_1, x_2)$ , it is quite reasonable. Each player  $i$  in  $N$  favors  $x_i$ , as expected. Furthermore, since each player strongly favors enacting the efficient outcome  $(x_1, x_2)$ , making it very likely, the resulting outcome is only marginally worse. Finally, choosing with positive probability not to enact the efficient outcome can be seen as a sensible hedge to handle the possibility that the other player chooses not to enact the efficient outcome; the payoff reduction is simply the price of hedging via randomization.

Thus, least-squares regret is notable here for yielding solutions in line with intuition, all without needing to appeal to other principles, and outperforming Nash equilibrium.

Now, consider the finite two-person coordination game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 6.8, which differs from the game shown in Table 6.7 only in that, for each player  $i$  in  $N$ , the payoff  $u_i(x_1, x_2)$  is increased from 3 to 10.

This game has three equilibria:  $(x_1, x_2)$ , which gives the payoff allocation  $(10, 10)$ ;  $(y_1, y_2)$ , which gives the payoff allocation  $(1, 1)$ ; and

$$((1/11)x_1 + (10/11)y_1, (1/11)x_2 + (10/11)y_2),$$

which gives the payoff allocation  $(10/11, 10/11)$  or roughly  $(0.91, 0.91)$ . Notably, in the unique equilibrium in randomized strategies, each player  $i$  in  $N$  strongly favors  $y_i$ , making the outcome  $(y_1, y_2)$  very likely.

Furthermore, in equilibrium, each player randomizes so as to render the other player indifferent between his strategies, and so, the behavior of a player depends



Table 6.8 A coordination game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	10	0
	$y_1$	0	1

not on his own payoffs, but on the payoffs of the other player. Thus, for each player  $i$  in  $N$ , increasing  $u_i(x_1, x_2)$  from 3 to 10 reduces the probability of  $x_i$  being chosen.

But again, these are peculiar results. It seems that the sole reasonable outcome is  $(x_1, x_2)$ , which is clearly superior to the rest. In fact, it is even more outstanding in this game than in the game shown in Table 6.7. Thus, again, one would expect the outcome  $(x_1, x_2)$  to be the unique solution. But just as before, to generate it as the unique solution using standard solution concepts requires an appeal to other principles such as the focal-point effect (Schelling, 1960) or payoff dominance (Harsanyi and Selten, 1988). And, perhaps disappointingly, this fact holds no matter how outstanding the outcome  $(x_1, x_2)$  might be.

Furthermore, the efficient outcome  $(x_1, x_2)$  is available and clearly better for each player. In fact, the extent of its superiority for each player is even greater in this game than in the game shown in Table 6.7. Thus, again, one would expect each player  $i$  in  $N$  to favor  $x_i$ .

Finally, for each player  $i$  in  $N$ , the payoff  $u_i(x_1, x_2)$  is greater in this game than in the game shown in Table 6.7. Thus, one would expect, for each player  $i$  in  $N$ , the probability of  $x_i$  being chosen to be greater in this game than in the game shown in Table 6.7. More generally, one would expect a player to be sensitive to his own payoffs in the natural way.

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is precisely the efficient outcome  $(x_1, x_2)$ . Thus, least-squares regret with respect to pure strategies yields as the unique solution precisely the sole reasonable outcome, all without needing to appeal to other principles.

Least-squares regret with respect to randomized strategies yields a similar solution. The unique least-squares regret profile in randomized strategies is

$$((^{100/101}x_1 + (^{1/101}y_1), (^{100/101}x_2 + (^{1/101}y_2)),$$

which gives the payoff allocation  $(^{100001/10201}, ^{100001/10201})$  or roughly  $(9.80, 9.80)$ .

While this solution is not the efficient outcome  $(x_1, x_2)$ , it is quite reasonable. Each player  $i$  in  $N$  favors  $x_i$ , as expected. Furthermore, since each player very

strongly favors enacting the efficient outcome  $(x_1, x_2)$ , making it almost inevitable, the resulting outcome is only marginally worse. Also, as noted earlier, choosing with positive probability not to enact the efficient outcome can be seen as a sensible hedge to handle the possibility that the other player chooses not to enact the efficient outcome; the payoff reduction is simply the price of hedging via randomization. Finally, for each player  $i$  in  $N$ , increasing  $u_i(x_1, x_2)$  from 3 to 10 increases the probability of  $x_i$  being chosen, as expected.

Thus, least-squares regret is notable here for yielding solutions in line with intuition, all without needing to appeal to other principles, outperforming Nash equilibrium, and capturing the experimentally robust own-payoff effects (Ochs, 1995; McKelvey, Palfrey, and Weber, 2000; Goeree and Holt, 2001; Goeree, Holt, and Palfrey, 2003).

Needless to say, least-squares regret cannot guarantee efficiency. Consider the finite two-person coordination game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 6.9.

Table 6.9 A coordination game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	3      2.5	3      0
	$y_1$	0      1	2.5      1

The unique least-squares regret profile in pure strategies is the inefficient outcome  $(y_1, y_2)$ , which gives the payoff allocation  $(1, 1)$ , and the unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, 0.2x_2 + 0.8y_2),$$

which gives the payoff allocation  $(1.16, 1.16)$ .

These solutions seem unsatisfactory. The efficient outcome  $(x_1, x_2)$  is available and clearly better for each player. Thus, one would expect each player  $i$  in  $N$  to favor  $x_i$ .

Still, it seems reasonable for each player  $i$  in  $N$  to favor  $y_i$ . If the other player chooses  $x_{-i}$ , then  $y_i$  yields a payoff that is only marginally less than the best-response payoff that is achieved by choosing  $x_i$ , and if the other player chooses  $y_{-i}$ , then  $y_i$  yields the best-response payoff, which is notably greater than the payoff from choosing  $x_i$ . Thus, it seems reasonable, all things considered, to favor  $y_i$ . Indeed, this is the very idea behind least-squares regret. In a sense, the inefficient

solutions yielded by least-squares regret are reasonable insofar as the logic just described is.

These solutions are reasonable for yet another reason: the inefficient outcome  $(y_1, y_2)$  is a risk-dominant equilibrium. Indeed, as Theorem 5.7.1 in Section 5.7 establishes, when it comes to the equilibria in pure strategies of a  $2 \times 2$  game in strategic form, least-squares regret is equivalent to risk dominance. Thus, the convergence here of least-squares regret and risk dominance is precisely what one would expect.

## 6.7 Battle of the Sexes

Consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form known as *Battle of the Sexes* shown in Table 6.10.

Table 6.10 Battle of the Sexes game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1    0	3    0
	$y_1$	0    3	0    1

This game has three equilibria:  $(x_1, x_2)$ , which gives the payoff allocation  $(3, 1)$ ;  $(y_1, y_2)$ , which gives the payoff allocation  $(1, 3)$ ; and

$$(0.75x_1 + 0.25y_1, 0.25x_2 + 0.75y_2),$$

which gives the payoff allocation  $(0.75, 0.75)$ . The first equilibrium is the most preferred outcome of player 1, and the second equilibrium is the most preferred outcome of player 2, and so, the players prefer different outcomes. In the third equilibrium, the players act in a random and uncoordinated manner, but each favors enacting his most preferred outcome. This third equilibrium is notably inefficient.

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is the conflict outcome  $(x_1, y_2)$ , which gives the payoff allocation  $(0, 0)$ , and the unique least-squares regret profile in randomized strategies is

$$(0.9x_1 + 0.1y_1, 0.1x_2 + 0.9y_2),$$

which gives the payoff allocation  $(0.36, 0.36)$ .

These solutions are admittedly unsatisfactory. But it is worth noting that such outcomes may well obtain in play, especially in the absence of any coordination, learning, or other mechanisms.

Least-squares regret can sometimes yield surprisingly gratifying solutions. Consider the modified finite Battle of the Sexes game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 6.11, which differs from the game shown in Table 6.10 only in that, for each player  $i$  in  $N$ , the payoff  $u_i(y_1, x_2)$  is increased from 0 to 2.5.

Table 6.11 Modified Battle of the Sexes game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1 3	0 0
	$y_1$	2.5 2.5	3 1

This game has three equilibria:  $(x_1, x_2)$ , which gives the payoff allocation  $(3, 1)$ ;  $(y_1, y_2)$ , which gives the payoff allocation  $(1, 3)$ ; and

$$((1/3)x_1 + (2/3)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(2, 2)$ .

Notice that this third equilibrium is worse for each player than the nonequilibrium compromise outcome  $(y_1, x_2)$ , which gives the payoff allocation  $(2.5, 2.5)$  and is only slightly worse for each player than his most preferred outcome. In fact, the compromise outcome  $(y_1, x_2)$  is a remarkably good outcome, and yet, it is unachievable by players who choose only equilibrium strategies.

Least-squares regret yields a more satisfying solution. The unique least-squares regret profile in pure strategies is precisely the compromise outcome  $(y_1, x_2)$ . Thus, least-squares regret with respect to pure strategies yields as the unique solution the remarkably good outcome that is unachievable by players who choose only equilibrium strategies.

Least-squares regret with respect to randomized strategies yields a similar solution. The unique least-squares regret profile in randomized strategies is

$$(0.2x_1 + 0.8y_1, 0.8x_2 + 0.2y_2),$$

which gives the payoff allocation  $(2.24, 2.24)$ .

This solution is quite reasonable. In particular, it comes close to being the compromise outcome and is better for each player than the qualitatively similar

third equilibrium described above. Still, this solution is somewhat unsatisfactory in that it is worse for each player than the compromise outcome that is achieved by considering only pure strategies. But this is to be expected: hedging via randomization comes at a price.

## 6.8 The Two-Person Bargaining Problem

Consider the *two-person bargaining problem* (Nash, 1950b; Nash, 1953). In this game, two players must divide some good between themselves, and each must propose the proportion of the good that he himself will take. If the proportions sum to no more than 1, each gets his demand; otherwise, each gets nothing. Thus, in this game, the set of players is  $N = \{1, 2\}$ , the pure-strategy sets are

$$C_1 = C_2 = \{x \in \mathbf{R} \mid 0 \leq x \leq 1\},$$

and, for any player  $i$  in  $N$ , the utility function is

$$\begin{aligned} u_i(c_1, c_2) &= 0 && \text{if } c_1 + c_2 > 1, \\ &= f_i(c_i) && \text{if } c_1 + c_2 \leq 1, \end{aligned}$$

where the function  $f_i: C_i \rightarrow \mathbf{R}$  is increasing and twice differentiable. Without loss of generality, for any player  $i$  in  $N$ , let the function  $f_i$  be normalized so that

$$f_i(0) = 0 \quad \text{and} \quad f_i(1) = 1.$$

For any player  $i$  in  $N$ , let the function  $F_i: C_i \rightarrow \mathbf{R}$  be any antiderivative of  $f_i$ .

For any allocation  $(c_1, c_2)$  in  $C_1 \times C_2$ , the allocation  $(c_1, c_2)$  is *feasible* if and only if  $c_1 + c_2 \leq 1$ .

It is easy to see that least-squares regret here yields reasonable solutions. Consider any player  $i$  in  $N$  and any demand  $c_i$  in  $C_i$ . If the other player chooses some demand  $c_{-i}$  in  $C_{-i}$  that is less than  $1 - c_i$ , then the allocation is feasible, but inefficient, and the payoff to player  $i$  from choosing  $c_i$  is  $f_i(c_i)$  while the best-response payoff is  $f_i(1 - c_{-i})$  (since the best response is to choose  $1 - c_{-i}$ ). If the other player chooses some demand  $c_{-i}$  in  $C_{-i}$  that is greater than  $1 - c_i$ , then the allocation is not feasible, and the payoff to player  $i$  from choosing  $c_i$  is 0 while the best-response payoff is again  $f_i(1 - c_{-i})$ . Thus, for any player  $i$  in  $N$ ,

$$\begin{aligned} \rho_i(c_i) &= \int_0^{1-c_i} (f_i(1 - c_{-i}) - f_i(c_i))^2 dc_{-i} + \int_{1-c_i}^1 (f_i(1 - c_{-i}) - 0)^2 dc_{-i} \\ &= \int_0^{1-c_i} (f_i(c_i))^2 - 2f_i(c_i)f_i(1 - c_{-i}) dc_{-i} + \int_0^1 (f_i(1 - c_{-i}))^2 dc_{-i} \\ &= (f_i(c_i))^2(1 - c_i) - 2f_i(c_i)(F_i(1) - F_i(c_i)) + \int_0^1 (f_i(1 - c_{-i}))^2 dc_{-i}, \end{aligned}$$

and so,

$$\frac{d(\rho_i(c_i))}{dc_i} = (f_i(c_i))^2 - 2 \frac{d(f_i(c_i))}{dc_i} [(F_i(1) - F_i(c_i)) - f_i(c_i)(1 - c_i)]$$

and

$$\begin{aligned} \frac{d^2(\rho_i(c_i))}{dc_i^2} &= 2f_i(c_i) \frac{d(f_i(c_i))}{dc_i} + 2 \left( \frac{d(f_i(c_i))}{dc_i} \right)^2 (1 - c_i) \\ &\quad - 2 \frac{d^2(f_i(c_i))}{dc_i^2} [(F_i(1) - F_i(c_i)) - f_i(c_i)(1 - c_i)]. \end{aligned}$$

In the standard formulation of the two-person bargaining problem, each player is risk neutral. It is easy to see that if, for each player  $i$  in  $N$ , the function  $f_i$  is linear, then least-squares regret yields precisely the Nash bargaining solution. Consider any player  $i$  in  $N$ . If the function  $f_i$  is linear, then

$$\begin{aligned} \frac{d(\rho_i(c_i))}{dc_i} &= c_i^2 - 2[(1/2 - c_i^2/2) - c_i(1 - c_i)] \\ &= 2c_i - 1. \end{aligned}$$

Setting the derivative equal to 0 and solving for  $c_i$  yields the unique pure least-squares regret strategy  $c_i = 0.5$ . Thus, the final allocation is  $(0.5, 0.5)$ , which is precisely the Nash bargaining solution.

What is notable about least-squares regret is that it generates precisely the Nash bargaining solution without requiring the axioms described in Nash (1950b) and Nash (1953). Furthermore, it is trivial to apply, intuitive, and parsimonious.

It is natural to wonder whether, by acting in accordance with least-squares regret, the players will end up with a feasible allocation. The following theorem establishes that if, for each player  $i$  in  $N$ , the function  $f_i$  is concave, then the final allocation is feasible.

**THEOREM 6.8.1.** *Consider any two-person bargaining problem as defined above. If, for each player  $i$  in  $N$ , the function  $f_i$  is concave, then, for any allocation  $(c_1, c_2)$  in  $C_1 \times C_2$ , if*

$$\rho_1(c_1) \leq \rho_1(d_1), \quad \forall d_1 \in C_1$$

and

$$\rho_2(c_2) \leq \rho_2(d_2), \quad \forall d_2 \in C_2,$$

then  $(c_1, c_2)$  is feasible.

*Proof.* Consider any player  $i$  in  $N$ . Since the function  $f_i$  is nonnegative, increasing, and concave,

$$\begin{aligned} \frac{d^2(\rho_i(c_i))}{dc_i^2} &= 2f_i(c_i) \frac{d(f_i(c_i))}{dc_i} + 2 \left( \frac{d(f_i(c_i))}{dc_i} \right)^2 (1 - c_i) \\ &\quad - 2 \frac{d^2(f_i(c_i))}{dc_i^2} [(F_i(1) - F_i(c_i)) - f_i(c_i)(1 - c_i)] \geq 0, \quad \forall c_i \in C_i. \end{aligned}$$

Thus, the regret function  $\rho_i$  is convex.

Consider the point  $c_i = 0.5$ . Notice that every concave function must lie somewhere between the following two extreme cases. In the first extreme case (in which the function  $f_i$  is linear),

$$f_i(0.5) = 0.5, \quad \frac{d(f_i(0.5))}{dc_i} = 1, \quad F_i(1) = 0.5, \quad \text{and} \quad F_i(0.5) = 1/8,$$

and so,

$$\frac{d(\rho_i(0.5))}{dc_i} = 0.$$

In the second extreme case,

$$f_i(0.5) = 1 \quad \text{and} \quad \frac{d(f_i(0.5))}{dc_i} = 0,$$

and so,

$$\frac{d(\rho_i(0.5))}{dc_i} = 1.$$

Thus, for any concave function  $f_i$ ,

$$0 \leq \frac{d(\rho_i(0.5))}{dc_i} \leq 1.$$

The point to note here is that the first derivative of the regret function  $\rho_i$  at the point  $c_i = 0.5$  is nonnegative. If

$$\frac{d(\rho_i(0.5))}{dc_i} = 0,$$

which holds if and only if the function  $f_i$  is linear, then the point  $c_i = 0.5$  is the unique point that minimizes the regret function  $\rho_i$ . But if

$$0 < \frac{d(\rho_i(0.5))}{dc_i} \leq 1,$$

so that the first derivative of the regret function  $\rho_i$  at the point  $c_i = 0.5$  is strictly positive, then any point that minimizes the regret function  $\rho_i$  must be strictly less than 0.5 since  $\rho_i$  is convex.

And so, for any point  $c_i$  in  $C_i$ , if

$$\rho_i(c_i) \leq \rho_i(d_i), \quad \forall d_i \in C_i,$$

then  $c_i \leq 0.5$ . Thus, for any allocation  $(c_1, c_2)$  in  $C_1 \times C_2$ , if

$$\rho_1(c_1) \leq \rho_1(d_1), \quad \forall d_1 \in C_1$$

and

$$\rho_2(c_2) \leq \rho_2(d_2), \quad \forall d_2 \in C_2,$$

then  $c_1 + c_2 \leq 1$ , and so,  $(c_1, c_2)$  is feasible. ■

Intuitively, a risk-averse player is conservative with his demands and unwilling to risk demanding an amount that might lead to an allocation that is not feasible. Risk aversion on the part of each player thus guarantees that the players will never end up outside of the set of feasible allocations.

It is instructive to compare least-squares regret with iterated regret minimization (Halpern and Pass, 2012) with respect to the two-person bargaining problem. The following theorem establishes that if, for each player  $i$  in  $N$ , the function  $f_i$  is concave, then the final payoff allocation generated by least-squares regret is at least as great as that generated by iterated regret minimization. Thus, the two-person bargaining problem is another example in which least-squares regret outperforms iterated regret minimization.

**THEOREM 6.8.2.** *Consider any two-person bargaining problem as defined above. For each player  $i$  in  $N$ , let the function  $f_i$  be concave. Let  $(\bar{c}_1, \bar{c}_2)$  be any allocation in  $C_1 \times C_2$  such that*

$$\rho_1(\bar{c}_1) \leq \rho_1(d_1), \quad \forall d_1 \in C_1$$

and

$$\rho_2(\bar{c}_2) \leq \rho_2(d_2), \quad \forall d_2 \in C_2.$$

Let  $(\hat{c}_1, \hat{c}_2)$  be any allocation in  $C_1 \times C_2$  such that

$$\hat{c}_1 \in \mathcal{RM}_1^\infty(C_1 \times C_2)$$

and

$$\hat{c}_2 \in \mathcal{RM}_2^\infty(C_1 \times C_2).$$

Then

$$(u_1(\bar{c}_1, \bar{c}_2), u_2(\bar{c}_1, \bar{c}_2)) \geq (u_1(\hat{c}_1, \hat{c}_2), u_2(\hat{c}_1, \hat{c}_2)).$$

*Proof.* Consider any player  $i$  in  $N$ . Since the function  $f_i$  is nonnegative, increasing, and concave,

$$\begin{aligned} \frac{d^2(\rho_i(c_i))}{dc_i^2} &= 2f_i(c_i) \frac{d(f_i(c_i))}{dc_i} + 2 \left( \frac{d(f_i(c_i))}{dc_i} \right)^2 (1 - c_i) \\ &\quad - 2 \frac{d^2(f_i(c_i))}{dc_i^2} [(F_i(1) - F_i(c_i)) - f_i(c_i)(1 - c_i)] \geq 0, \quad \forall c_i \in C_i. \end{aligned}$$

Thus, the regret function  $\rho_i$  is convex.

Let  $\hat{c}_i$  be any point in  $C_i$  such that

$$\hat{c}_i \in \mathcal{RM}_i^\infty(C_1 \times C_2).$$



Notice that  $\hat{c}_i = f_i^{-1}[\{0.5\}]$ .

Now, consider the point  $c_i = f_i^{-1}[\{0.5\}]$ . Notice that since the function  $f_i$  is concave,

$$\frac{d(f_i(f_i^{-1}[\{0.5\}]))}{dc_i} \geq 1$$

and

$$(F_i(1) - F_i(f_i^{-1}[\{0.5\}])) - f_i(f_i^{-1}[\{0.5\}]) (1 - f_i^{-1}[\{0.5\}]) \geq 1/8.$$

Thus, for any concave function  $f_i$ ,

$$\frac{d(\rho_i(f_i^{-1}[\{0.5\}]))}{dc_i} \leq 0.$$

If

$$\frac{d(\rho_i(f_i^{-1}[\{0.5\}]))}{dc_i} = 0,$$

which holds if and only if the function  $f_i$  is linear, then the point  $c_i = f_i^{-1}[\{0.5\}]$  is the unique point that minimizes the regret function  $\rho_i$ . But if

$$\frac{d(\rho_i(f_i^{-1}[\{0.5\}]))}{dc_i} < 0,$$

then any point that minimizes the regret function  $\rho_i$  must be strictly greater than  $f_i^{-1}[\{0.5\}]$  since  $\rho_i$  is convex.

And so, for any point  $c_i$  in  $C_i$ , if

$$\rho_i(c_i) \leq \rho_i(d_i), \quad \forall d_i \in C_i,$$

then  $c_i \geq f_i^{-1}[\{0.5\}]$ .

Let  $\bar{c}_i$  be any point in  $C_i$  such that

$$\rho_i(\bar{c}_i) \leq \rho_i(d_i), \quad \forall d_i \in C_i.$$

By Theorem 6.8.1 above, the allocation  $(\bar{c}_1, \bar{c}_2)$  is feasible, and so, the allocation  $(\hat{c}_1, \hat{c}_2)$  is likewise feasible. Thus,

$$u_i(\bar{c}_1, \bar{c}_2) = f_i(\bar{c}_i) \quad \text{and} \quad u_i(\hat{c}_1, \hat{c}_2) = f_i(\hat{c}_i).$$

And so, since  $\bar{c}_i \geq \hat{c}_i$  and the function  $f_i$  is increasing,

$$u_i(\bar{c}_1, \bar{c}_2) = f_i(\bar{c}_i) \geq f_i(\hat{c}_i) = u_i(\hat{c}_1, \hat{c}_2).$$

Thus,

$$(u_1(\bar{c}_1, \bar{c}_2), u_2(\bar{c}_1, \bar{c}_2)) \geq (u_1(\hat{c}_1, \hat{c}_2), u_2(\hat{c}_1, \hat{c}_2)). \quad \blacksquare$$

Intuitively, given risk aversion on the part of each player, choosing any pure least-squares regret strategy yields a payoff that is at least as great as that generated by choosing any pure strategy that survives iterated regret minimization. Iterated regret minimization is, relative to least-squares regret, characterized by a greater degree of conservatism and pessimism. This is unsurprising since, as discussed in Section 3.2, iterated regret minimization (without iterative elimination) can be seen as an extreme case of a natural generalization of least-squares regret.

# 7

---

## Weaknesses of Least-Squares Regret

In this chapter, we consider some of the weaknesses of least-squares regret. Section 7.1 discusses the failure of least-squares regret to satisfy the principle of Independence of Irrelevant Alternatives. Section 7.2 compares least-squares regret with maximin with respect to two-person zero-sum games. Section 7.3 examines the sensitivity of least-squares regret to framing. Section 7.4 considers some of the problems that can arise from defining least-squares regret in a way that disregards fully strategic reasoning.

### 7.1 Independence of Irrelevant Alternatives

The principle of *Independence of Irrelevant Alternatives* appears throughout game theory and economics and is defined formally as follows. Let  $X$  be any nonempty set of alternatives. The principle of Independence of Irrelevant Alternatives asserts that, for any set  $S \subseteq X$  of alternatives and any alternatives  $x$  and  $y$  in  $S$ , if the alternative  $x$  is preferred to the alternative  $y$  with respect to the set  $S$ , then, for any alternative  $z$  in  $X$  such that  $z \notin S$ , the alternative  $x$  is preferred to the alternative  $y$  with respect to the set  $S \cup \{z\}$ . Intuitively, the principle asserts that the ordering over two alternatives depends only on those alternatives and not on any others, which are deemed irrelevant to the comparison at hand.

As might be expected, least-squares regret fails to satisfy Independence of Irrelevant Alternatives. For an illustration of this point, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.1.

Notice that, for player 1,

$$\rho_1(x_1) = 1 \quad \text{and} \quad \rho_1(y_1) = 4.$$

Thus, in this game,  $x_1$  is preferred to  $y_1$ .

Table 7.1 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1	4
	$y_1$	2	2

Now, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.2, which differs from the game shown in Table 7.1 only in the addition of the pure strategy  $z_1$  and the payoffs that it yields.

Table 7.2 Payoffs of player 1 in a game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1	4
	$y_1$	2	2
	$z_1$	4	1

Notice that, for player 1,

$$\rho_1(x_1) = 9, \quad \rho_1(y_1) = 8, \quad \text{and} \quad \rho_1(z_1) = 9.$$

Thus, in this game,  $y_1$  is preferred to  $x_1$ .

Adding  $z_1$ —which, moreover, is ultimately not recommended since  $\rho_1(z_1) = 9$ —changes the ordering over  $x_1$  and  $y_1$ . Thus, Independence of Irrelevant Alternatives is violated.

But Independence of Irrelevant Alternatives is a notoriously controversial principle. It states that the ordering over two alternatives should be independent of the context in which they are presented. In practice, however, choice is often context-sensitive, and such context effects are experimentally robust; see, for example, Simonson and Tversky (1992) and Tversky and Simonson (1993). Thus, it is not clear that the failure of least-squares regret to satisfy the principle of Independence of Irrelevant Alternatives should be considered a real or significant weakness. Indeed, it seems natural to suppose, as the foregoing example illustrates, that the reasonableness of a strategy should depend on the alternatives that are

available. Intuitively, it is important to consider the entire context when making a decision, and should that context change, a re-evaluation would be advisable.

## 7.2 Maximin and Two-Person Zero-Sum Games

The standard solution concept for two-person zero-sum games is maximin (von Neumann, 1928; Wald, 1939; Wald, 1945; von Neumann and Morgenstern, 1947; Wald, 1950), which, when randomization is allowed, also coincides exactly with Nash equilibrium. Least-squares regret represents a significant departure from both of these solution concepts.

For an illustration of some of the differences, consider the finite two-person zero-sum game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.3.

Table 7.3 A two-person zero-sum game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	-4	0
	$y_1$	-1	-2
		1	2

Consider first maximin with respect to pure strategies. For player 1, the minimum payoff from choosing  $x_1$  is 0, and the minimum payoff from choosing  $y_1$  is 1, and so,  $y_1$  is the unique pure maximin strategy. For player 2, the minimum payoff from choosing  $x_2$  is -4, and the minimum payoff from choosing  $y_2$  is -2, and so,  $y_2$  is the unique pure maximin strategy. Thus, the unique profile of pure maximin strategies is  $(y_1, y_2)$ , which gives the payoff allocation  $(2, -2)$ .

Now, consider maximin with respect to randomized strategies. As Figure 7.1 shows, for player 1, the minimum expected payoff is maximized when

$$4\sigma_1(x_1) + (1 - \sigma_1(x_1)) = 2(1 - \sigma_1(x_1)),$$

that is, at the point  $\sigma_1(x_1) = 0.2$ , and so, the unique randomized maximin strategy is  $0.2x_1 + 0.8y_1$ . As Figure 7.2 shows, for player 2, the minimum expected payoff is maximized when

$$-4\sigma_2(x_2) = -\sigma_2(x_2) - 2(1 - \sigma_2(x_2)),$$

that is, at the point  $\sigma_2(x_2) = 0.4$ , and so, the unique randomized maximin strategy is  $0.4x_2 + 0.6y_2$ . Thus, the unique profile of randomized maximin strategies is

$$(0.2x_1 + 0.8y_1, 0.4x_2 + 0.6y_2),$$

Figure 7.1 Expected payoffs of player 1





Figure 7.2 Expected payoffs of player 2



which gives the payoff allocation  $(1.6, -1.6)$ . Notably, player 1 favors  $y_1$ , and player 2 favors  $y_2$ .

Now, consider least-squares regret. The unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ , which gives the payoff allocation  $(0, 0)$ , and the unique least-squares regret profile in randomized strategies is

$$\left(\left(\frac{9}{13}\right)x_1 + \left(\frac{4}{13}\right)y_1, \left(\frac{1}{17}\right)x_2 + \left(\frac{16}{17}\right)y_2\right),$$

which gives the payoff allocation  $(\frac{168}{221}, -\frac{168}{221})$  or roughly  $(0.76, -0.76)$ . Notably, player 1 favors  $x_1$ , and player 2 favors  $y_2$ .

Thus, when it comes to two-person zero-sum games, least-squares regret can yield solutions that are different from those yielded by maximin and thus also by Nash equilibrium. This discrepancy is unfortunate in light of the standard arguments for the latter two solution concepts with respect to two-person zero-sum games. From the perspective of maximin, by playing a maximin strategy, one maximizes the minimum payoff and thus assures oneself of a certain minimum payoff no matter what the other player might do. Furthermore, supposing that the other player plays his maximin strategy, one cannot do better by deviating since a maximin strategy is also an equilibrium strategy. Alternatively, from the perspective of Nash equilibrium, by playing an equilibrium strategy, one also maximizes the minimum payoff. These are rather compelling considerations.

Least-squares regret, maximin, and Nash equilibrium represent different ways to reason about a game. Least-squares regret involves partially strategic reasoning and assessing personal payoffs in the form of regrets and choosing a strategy so as to minimize the divergence from the best-response payoffs. Maximin involves assessing personal payoffs, but not regrets, and choosing a strategy so as to maximize the minimum payoff. Nash equilibrium involves fully strategic reasoning and choosing a strategy so as to maximize the expected payoff with respect to the strategy of the other player.

Which way to reason is the most reasonable depends on a number of considerations, for example, whether partially strategic or fully strategic reasoning is more appropriate; whether regrets or payoffs are more significant; and whether it is most important to minimize the divergence from the best-response payoffs, to maximize the minimum payoff, or to maximize the expected payoff with respect to the strategy of the other player.

### 7.3 Framing Effects

It has long been known that how a decision problem is presented can affect what choice is made (Kahneman and Tversky, 1979; Tversky and Kahneman, 1981; Levin, Schneider, and Gaeth, 1998). Unsurprisingly, framing effects can arise also in games, raising questions about how to take such effects into account; see, for example, Eliaz and Rubinstein (2011); Dufwenberg, Gächter, and Hennig-Schmidt (2011); Ellingsen, Johannesson, Mollerstrom, and Munkhammar (2012), and Dreber, Ellingsen, Johannesson, and Rand (2013).

In this section, we show that least-squares regret is susceptible to certain framing effects. Consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.4.

Table 7.4 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	1
	$y_1$	0	1
		10	0

Notice that, for player 1,

$$\rho_1(x_1) = 100 \quad \text{and} \quad \rho_1(y_1) = 1.$$

Thus, the unique pure least-squares regret strategy for player 1 is  $y_1$ .

Since  $u_1(y_1, x_2)$  is so great and  $u_1(x_1, x_2)$  is so low by comparison, the regret that  $y_1$  could induce is significant. Thus, given how the game is presented,  $y_1$  appears to be extremely attractive from the perspective of least-squares regret and has a dominating effect.

This effect is akin to the distortionary effect of an extreme value on the mean of a set of numbers. Just as an extreme value can pull the mean of a set of numbers in its direction, a strategy that could induce significant regret can pull least-squares regret in its direction.

But this can be problematic, as the game above shows. Notice that  $x_2$  is strongly dominated for player 2 by  $y_2$ . Given this fact,  $y_1$  is, in reality, not as attractive as least-squares regret would suggest: the significant payoff that it could yield—and thus the significant regret that it could induce—turns on player 2 acting imperfectly. Thus,  $y_1$  can be seen as a naïve strategy whose sensibility depends on imperfect or irrational play. It can be likened to a naïve chess tactic that sets up a rudimentary trap and would yield a significant gain should the opponent fall prey to the trap. But if player 2 can be expected to play sensibly, choosing  $y_2$  and not  $x_2$ , then  $y_1$  is unattractive and should thus be seen as negligible.

Thus, how a game is framed can matter a great deal for least-squares regret. This sensitivity to framing is to be expected given the basic idea behind least-squares regret. As discussed in Sections 1.2, 1.3, 2.2, and 2.5, the regret of a strategy with respect to a partial profile of strategies of the other players can be seen as a measure of error. The greater is the degree to which a strategy falls short of a best response, the graver a mistake it is to choose it. Accepting this idea means accepting that the magnitude of regret is significant and that it is important to be sensitive to it. Thus, least-squares regret should depend on the magnitudes of the regrets in one way or another. But this means also that the distortionary framing effects described above are inescapable and must be accepted as the price of defining the reasonableness of a strategy in terms of regret.

Still, we readily admit that there is something unreasonable about assigning undue significance to regret that could conceivably be considered negligible, say, because the partial profile that would generate it can be discounted. We consider this weakness in Section 7.4, and in Chapter 8, we introduce a refinement of least-squares regret that addresses this weakness and others.

## 7.4 Fully Strategic Reasoning

As discussed in Sections 1.2, 1.4, 2.2, and 2.5, least-squares regret considers partially strategic players and disregards fully strategic reasoning. In this section,



we consider some of the problems that can arise from defining least-squares regret in this way.

Consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.5.

Table 7.5 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1	0
	$y_1$	2	1

The unique least-squares regret profile in pure strategies is  $(x_1, x_2)$ , which gives the payoff allocation  $(2, 1)$ , and the unique least-squares regret profile in randomized strategies is

$$(0.8x_1 + 0.2y_1, x_2),$$

which gives the payoff allocation  $(1.6, 1)$ .

Now, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 7.6, which differs from the game shown in Table 7.5 only in the payoffs of player 2.

Table 7.6 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	0	1
	$y_1$	2	1

The unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ , which gives the payoff allocation  $(1, 1)$ , and the unique least-squares regret profile in randomized strategies is

$$(0.8x_1 + 0.2y_1, y_2),$$

which gives the payoff allocation  $(1.2, 1)$ .

Notice that the pure and randomized least-squares regret strategies for player 1 in the one game are identical to the pure and randomized least-squares regret

strategies for player 1 in the other game. Indeed, they must be identical since the utility functions for player 1 in the two games are identical. Thus, as the two games illustrate, the payoffs of player 2 are irrelevant to player 1 and can be varied arbitrarily without affecting the behavior of player 1. In general, according to least-squares regret, the behavior of a player is completely independent of the payoffs of the other players.

But this independence is too strong and perhaps unrealistic. Just as the behavior of a player depends on his own payoffs, the payoffs of the other players can be expected to influence how those players behave. And if a player can reason about how the other players might behave and to respond accordingly, it is reasonable to expect his behavior to vary depending on the payoffs of the other players. But least-squares regret, which disregards fully strategic reasoning, is unable to capture this variation and this intuition.

Furthermore, the most natural or reasonable way to solve a game may involve recognizing fully strategic reasoning. For example, the game shown in Table 7.6 can be solved by recognizing that player 2 can be expected to play  $y_2$  for sure (since  $x_2$  is strongly dominated for player 2 by  $y_2$ ) and that player 1, considering this, can be expected to play  $y_1$  for sure. In general, fully strategic reasoning may be indispensable for solving a game.

But since least-squares regret disregards fully strategic reasoning, it can fail to capture the most natural or reasonable way to solve a game and may yield an unreasonable solution. For example, the disregard of fully strategic reasoning explains why least-squares regret fails to yield for the game shown in Table 7.6 the solution  $(y_1, y_2)$  just described.

Finally, least-squares regret can yield solutions that would be unreasonable or unstable if they were to be anticipated by fully strategic players. For an illustration of this point, recall that, in the game shown in Table 7.6, the unique least-squares regret profile in pure strategies is  $(x_1, y_2)$ . But if the players are fully strategic, then this profile is unreasonable and unstable. In particular, it is not an equilibrium since player 1, should he expect player 2 to choose  $y_2$ , would prefer to deviate and choose  $y_1$ . Furthermore, it is unreasonable and unstable in the sense that player 1, should he expect player 2 to choose  $y_2$ , could reduce his regret to 0 by deviating and choosing the best response  $y_1$ . Thus, while least-squares regret may be adept at characterizing the behavior of partially strategic players, it may be less adept at characterizing the behavior of fully strategic players.

It is worth noting that these problems are not unique to least-squares regret. Indeed, it is easy to see that all of these problems plague iterated regret minimization also. Halpern and Pass (2012) consider a number of treatments, including introducing lexicographic belief systems reminiscent of the lexicographic proba-

bility systems of Blume, Brandenburger, and Dekel (1991) and Brandenburger, Friedenber, and Keisler (2008) and restricting analysis to proper subsets of the set of partial profiles of strategies of the other players.

It may be important to refine least-squares regret to incorporate fully strategic reasoning. One approach might involve, following Halpern and Pass (2012), restricting analysis to proper subsets of the set of partial profiles, for example, those resulting from the sets of undominated strategies. An alternative approach might involve assigning weights to partial profiles; we consider this particular refinement of least-squares regret in Chapter 8. The question of how best to incorporate fully strategic reasoning remains to be determined.

# 8

---

## Mutual Weighted Least-Squares Regret

As discussed in Sections 1.2, 1.4, 2.2, 2.5, and 7.4, least-squares regret considers partially strategic players and is defined accordingly. While the assumption that a player is partially strategic may be convenient and a realistic characterization of typical reasoning and behavior, it may also be too restrictive, as information about the behavior and beliefs of the other players is discarded. Furthermore, as discussed in Section 2.5, least-squares regret assumes no mutuality condition, with the result that in a least-squares regret profile, the strategies that the players choose may differ from the ones that the players expect to be chosen.

In this chapter, we introduce a refinement of least-squares regret, which we call *mutual weighted least-squares regret*, that addresses these concerns. The idea is to modify the regret function in randomized strategies defined in Section 2.2 by assigning probability weights to the partial profiles of strategies of the other players, where the distribution is just the probability distribution induced by the randomized strategies of the other players, and then to introduce a mutuality condition that requires that the randomized strategies that the players choose be precisely the ones that the players expect to be chosen.

Thus, mutual weighted least-squares regret considers fully strategic players capable of reasoning about one another. Section 8.1 formally defines mutual weighted least-squares regret. Section 8.2 studies an illustrative example. Section 8.3 presents a general existence theorem. Section 8.4 examines the relationship between mutual weighted least-squares regret and Nash equilibrium. Section 8.5 examines whether recursively updating the probability distributions by iteratively minimizing the respective weighted regret functions of the players yields convergence to a fixed point.

## 8.1 Formal Definition

Mutual weighted least-squares regret is defined formally as follows. Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. For any player  $i$  in  $N$ , the set

$$C_{-i} = \prod_{j \in N-i} C_j$$

is the set of partial profiles that can ultimately obtain in play. But the partial profiles in  $C_{-i}$  need not be all on a par; some may be likelier than others depending on how the other players choose their strategies. More precisely, the randomized strategies of the other players induce a probability distribution over the set  $C_{-i}$ . For any player  $i$  in  $N$ , any partial profile  $\sigma_{-i} = (\sigma_j)_{j \in N-i}$  in  $\prod_{j \in N-i} \Delta(C_j)$ , and any partial profile  $c_{-i}$  in  $C_{-i}$ , the probability that  $c_{-i}$  obtains in play is just

$$\prod_{j \in N-i} \sigma_j(c_j)$$

since the players choose their pure strategies independently.

The induced probabilities can be used as weights in the following way. For any player  $i$  in  $N$ , let  $\rho_i: \prod_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  be the *weighted regret function in randomized strategies* for player  $i$  such that

$$\rho_i(\sigma_{-i}, \sigma_i) = \sum_{c_{-i} \in C_{-i}} \left( \prod_{j \in N-i} \sigma_j(c_j) \right) \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \sigma_i) \right)^2.$$

This weighted regret function differs from the regret function in randomized strategies defined in Section 2.2 only in that each squared regret term is multiplied by its corresponding weight. Intuitively, for any randomized-strategy profile  $\sigma = (\sigma_j)_{j \in N}$  in  $\prod_{j \in N} \Delta(C_j)$ , the value  $\rho_i(\sigma_{-i}, \sigma_i)$  is the weighted regret from choosing  $\sigma_i$  with respect to  $\sigma_{-i}$ .

As with least-squares regret as defined in Section 2.2, we suppose that a player chooses a strategy so as to minimize the divergence from the best-response payoffs. But here, unlike with least-squares regret, the minimum is computed with respect to a partial profile, which induces the weights to be used in the computation. For any player  $i$  in  $N$ , any partial profile  $\sigma_{-i}$  in  $\prod_{j \in N-i} \Delta(C_j)$ , and any randomized strategy  $\sigma_i$  in  $\Delta(C_i)$ , the randomized strategy  $\sigma_i$  is a *randomized weighted least-squares regret strategy* for player  $i$  with respect to  $\sigma_{-i}$  if and only if

$$\rho_i(\sigma_{-i}, \sigma_i) \leq \rho_i(\sigma_{-i}, \tau_i), \quad \forall \tau_i \in \Delta(C_i).$$

Intuitively, player  $i$  expects the other players to choose their strategies according to  $\sigma_{-i}$ , computes the probability distribution induced by  $\sigma_{-i}$  over the set  $C_{-i}$ , and

then chooses a randomized strategy so as to minimize the weighted regret function  $\rho_i$  with respect to  $\sigma_{-i}$ . Notice that the randomized strategies that the players end up choosing need not coincide with the ones that the players expect to be chosen.

Now, we introduce a mutuality condition. In particular, we suppose that a player chooses a randomized weighted least-squares regret strategy and, moreover, has no mistaken beliefs about the strategies of the other players. For any randomized-strategy profile  $\sigma = (\sigma_j)_{j \in N}$  in  $\prod_{j \in N} \Delta(C_j)$ , the randomized-strategy profile  $\sigma$  is a *mutual weighted least-squares regret profile in randomized strategies* of  $\Gamma$  if and only if

$$\rho_i(\sigma_{-i}, \sigma_i) \leq \rho_i(\sigma_{-i}, \tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i).$$

Intuitively, in a mutual weighted least-squares regret profile in randomized strategies, the randomized strategies that minimize the respective weighted regret functions of the players are precisely the ones that the players expect to be chosen. Notably, a mutual weighted least-squares regret profile in randomized strategies is stable in the sense that no player could reduce his weighted regret by deviating given the randomized strategies of the other players.

## 8.2 An Example

For an illustration of mutual weighted least-squares regret, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 8.1.

Table 8.1 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1 2	0 3
	$y_1$	0 3	1 1

For any randomized-strategy profile  $(\sigma_1, \sigma_2)$  in  $\Delta(C_1) \times \Delta(C_2)$ , the randomized-strategy profile  $(\sigma_1, \sigma_2)$  is a mutual weighted least-squares regret profile in randomized strategies of the game shown in Table 8.1 if and only if  $\sigma_1$  is a randomized weighted least-squares regret strategy for player 1 with respect to  $\sigma_2$  and  $\sigma_2$  is a randomized weighted least-squares regret strategy for player 2 with respect to  $\sigma_1$ .

Consider player 1. Notice that whatever player 1 chooses, he risks playing imperfectly depending on what player 2 chooses. Thus, player 1 chooses a randomized strategy so as to minimize the weighted regret function  $\rho_1$  with respect to the randomized strategy of player 2.

Consider any randomized strategy  $\sigma_1$  in  $\Delta(C_1)$  and any randomized strategy  $\sigma_2$  in  $\Delta(C_2)$ . If player 2 ends up choosing  $x_2$ , which obtains with probability  $\sigma_2(x_2)$ , then the payoff to player 1 is  $2\sigma_1(x_1) + 3(1 - \sigma_1(x_1))$  while the best-response payoff is 3, and so, the regret is  $3 - (2\sigma_1(x_1) + 3(1 - \sigma_1(x_1)))$ . If player 2 ends up choosing  $y_2$ , which obtains with probability  $\sigma_2(y_2) = 1 - \sigma_2(x_2)$ , then the payoff to player 1 is  $3\sigma_1(x_1) + (1 - \sigma_1(x_1))$  while the best-response payoff is 3, and so, the regret is  $3 - (3\sigma_1(x_1) + (1 - \sigma_1(x_1)))$ . The weighted regret function  $\rho_1: \prod_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_1(\sigma_1, \sigma_2) &= \sigma_2(x_2) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, x_2) - u_1(\sigma_1, x_2) \right)^2 + \sigma_2(y_2) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, y_2) - u_1(\sigma_1, y_2) \right)^2 \\ &= \sigma_2(x_2)(3 - (2\sigma_1(x_1) + 3(1 - \sigma_1(x_1))))^2 + (1 - \sigma_2(x_2))(3 - (3\sigma_1(x_1) + (1 - \sigma_1(x_1))))^2, \end{aligned}$$

and so,

$$\frac{\partial(\rho_1(\sigma_1, \sigma_2))}{\partial\sigma_1(x_1)} = 8\sigma_1(x_1) + 8\sigma_2(x_2) - 6\sigma_1(x_1)\sigma_2(x_2) - 8.$$

Consider player 2. Notice that whatever player 2 chooses, he risks playing imperfectly depending on what player 1 chooses. Thus, player 2 chooses a randomized strategy so as to minimize the weighted regret function  $\rho_2$  with respect to the randomized strategy of player 1.

Consider any randomized strategy  $\sigma_2$  in  $\Delta(C_2)$  and any randomized strategy  $\sigma_1$  in  $\Delta(C_1)$ . If player 1 ends up choosing  $x_1$ , which obtains with probability  $\sigma_1(x_1)$ , then the payoff to player 2 is  $\sigma_2(x_2)$  while the best-response payoff is 1, and so, the regret is  $1 - \sigma_2(x_2)$ . If player 1 ends up choosing  $y_1$ , which obtains with probability  $\sigma_1(y_1) = 1 - \sigma_1(x_1)$ , then the payoff to player 2 is  $1 - \sigma_2(x_2)$  while the best-response payoff is 1, and so, the regret is  $1 - (1 - \sigma_2(x_2))$ . The weighted regret function  $\rho_2: \prod_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_2(\sigma_1, \sigma_2) &= \sigma_1(x_1) \left( \max_{\tau_2 \in \Delta(C_2)} u_2(x_1, \tau_2) - u_2(x_1, \sigma_2) \right)^2 + \sigma_1(y_1) \left( \max_{\tau_2 \in \Delta(C_2)} u_2(y_1, \tau_2) - u_2(y_1, \sigma_2) \right)^2 \\ &= \sigma_1(x_1)(1 - \sigma_2(x_2))^2 + (1 - \sigma_1(x_1))(1 - (1 - \sigma_2(x_2)))^2, \end{aligned}$$

and so,

$$\frac{\partial(\rho_2(\sigma_1, \sigma_2))}{\partial\sigma_2(x_2)} = 2\sigma_2(x_2) - 2\sigma_1(x_1).$$

Setting each of the partial derivatives above equal to 0 and solving the resulting system of equations yield

$$\sigma_1(x_1) = \sigma_2(x_2) = 2/3.$$

Thus, the unique mutual weighted least-squares regret profile in randomized strategies is

$$((2/3)x_1 + (1/3)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(7/3, 5/9)$ .

### 8.3 Existence

In this section, we show that mutual weighted least-squares regret satisfies a general existence theorem. The proof of the theorem relies on the following fixed-point theorem due to Kakutani (1941).

**KAKUTANI FIXED-POINT THEOREM (KAKUTANI (1941)).** *Let  $S$  be any nonempty, convex, bounded, and closed subset of a finite-dimensional vector space. Let  $F: S \rightarrow S$  be any upper-hemicontinuous point-to-set correspondence such that, for every  $x$  in  $S$ , the set  $F(x)$  is a nonempty convex subset of  $S$ . Then there exists some  $\bar{x}$  in  $S$  such that  $\bar{x} \in F(\bar{x})$ .*

The following general existence theorem establishes that every finite game in strategic form has at least one mutual weighted least-squares regret profile in randomized strategies. Thus, for any finite game in strategic form, a solution is guaranteed to exist.

**THEOREM 8.3.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Then there exists some mutual weighted least-squares regret profile in randomized strategies.*

*Proof.* The set  $\times_{j \in N} \Delta(C_j)$  of randomized-strategy profiles is a nonempty, convex, bounded, and closed subset of a finite-dimensional vector space.

For any player  $i$  in  $N$ , let  $R_i: \times_{j \in N-i} \Delta(C_j) \rightarrow \Delta(C_i)$  be the point-to-set correspondence such that

$$R_i(\sigma_{-i}) = \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(\sigma_{-i}, \tau_i).$$

Intuitively, for any partial profile  $\sigma_{-i}$  in  $\times_{j \in N-i} \Delta(C_j)$ , the set  $R_i(\sigma_{-i})$  is the set of randomized weighted least-squares regret strategies for player  $i$  with respect to  $\sigma_{-i}$ .

Consider any player  $i$  in  $N$ . Let  $\sigma_{-i}$  be any partial profile in  $\times_{j \in N-i} \Delta(C_j)$ . Since the set  $\Delta(C_i)$  is a nonempty compact set and since the weighted regret function  $\rho_i: \times_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  is continuous,  $\rho_i$  has a minimum with respect to  $\sigma_{-i}$ . Thus, for any partial profile  $\sigma_{-i}$  in  $\times_{j \in N-i} \Delta(C_j)$ , the set  $R_i(\sigma_{-i})$  is nonempty.



Now, let  $\sigma_{-i}$  be any partial profile in  $\times_{j \in N-i} \Delta(C_j)$ , and let  $\sigma_i$  and  $\hat{\sigma}_i$  be any randomized strategies in  $\Delta(C_i)$ . Let  $\lambda$  be any real number such that  $0 \leq \lambda \leq 1$ , and let  $\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i$  be the randomized strategy in  $\Delta(C_i)$  such that

$$(\lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i)(c_i) = \lambda\sigma_i(c_i) + (1 - \lambda)\hat{\sigma}_i(c_i), \quad \forall c_i \in C_i.$$

Then

$$\rho_i(\sigma_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) = \sum_{c_{-i} \in C_{-i}} \left( \prod_{j \in N-i} \sigma_j(c_j) \right) \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) \right)^2.$$

But notice that

$$u_i(c_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) = \lambda u_i(c_{-i}, \sigma_i) + (1 - \lambda)u_i(c_{-i}, \hat{\sigma}_i).$$

And so,

$$\rho_i(\sigma_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i) = \sum_{c_{-i} \in C_{-i}} \left( \prod_{j \in N-i} \sigma_j(c_j) \right) \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \hat{\sigma}_i) + \lambda(u_i(c_{-i}, \hat{\sigma}_i) - u_i(c_{-i}, \sigma_i)) \right)^2.$$

Now, consider the function  $f_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$f_i(\lambda) = \rho_i(\sigma_{-i}, \lambda\sigma_i + (1 - \lambda)\hat{\sigma}_i).$$

Notice that

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} = 2 \sum_{c_{-i} \in C_{-i}} \left( \prod_{j \in N-i} \sigma_j(c_j) \right) (u_i(c_{-i}, \hat{\sigma}_i) - u_i(c_{-i}, \sigma_i))^2.$$

Clearly,

$$\frac{d^2(f_i(\lambda))}{d\lambda^2} \geq 0, \quad \forall \lambda \in [0, 1],$$

and so, the function  $f_i$  is convex.

Now, notice that

$$f_i(0) = \rho_i(\sigma_{-i}, 0\sigma_i + (1 - 0)\hat{\sigma}_i) = \rho_i(\sigma_{-i}, \hat{\sigma}_i)$$

and

$$f_i(1) = \rho_i(\sigma_{-i}, 1\sigma_i + (1 - 1)\hat{\sigma}_i) = \rho_i(\sigma_{-i}, \sigma_i).$$

Observe that the line connecting the points  $(0, \rho_i(\sigma_{-i}, \hat{\sigma}_i))$  and  $(1, \rho_i(\sigma_{-i}, \sigma_i))$  in  $\mathbf{R}^2$  is just the function  $g_i: [0, 1] \rightarrow \mathbf{R}$  such that

$$g_i(\lambda) = \lambda\rho_i(\sigma_{-i}, \sigma_i) + (1 - \lambda)\rho_i(\sigma_{-i}, \hat{\sigma}_i).$$

But since the function  $f_i$  is convex,

$$f_i(\lambda) = \rho_i(\sigma_{-i}, \lambda\sigma_i + (1-\lambda)\hat{\sigma}_i) \leq g_i(\lambda) = \lambda\rho_i(\sigma_{-i}, \sigma_i) + (1-\lambda)\rho_i(\sigma_{-i}, \hat{\sigma}_i), \quad \forall \lambda \in [0, 1].$$

Thus, for any partial profile  $\sigma_{-i}$  in  $\times_{j \in N-i} \Delta(C_j)$ , the weighted regret function  $\rho_i$  is convex. And so, for any partial profile  $\sigma_{-i}$  in  $\times_{j \in N-i} \Delta(C_j)$ , the set  $R_i(\sigma_{-i})$  is convex.

Let  $R: \times_{j \in N} \Delta(C_j) \rightarrow \times_{j \in N} \Delta(C_j)$  be the point-to-set correspondence such that

$$R(\sigma) = \times_{i \in N} R_i(\sigma_{-i}).$$

To understand this point-to-set correspondence, consider any randomized-strategy profiles  $\sigma$  and  $\tau$  in  $\times_{j \in N} \Delta(C_j)$ . Then  $\tau \in R(\sigma)$  if and only if

$$\tau_i \in R_i(\sigma_{-i}), \quad \forall i \in N.$$

For any randomized-strategy profile  $\sigma$  in  $\times_{j \in N} \Delta(C_j)$ , the set  $R(\sigma)$  is a nonempty convex subset of  $\times_{j \in N} \Delta(C_j)$  since it is the Cartesian product of nonempty convex sets.

Now, let  $(\sigma^k)_{k=1}^\infty$  and  $(\tau^k)_{k=1}^\infty$  be any convergent sequences, and suppose that

$$\sigma^k \in \times_{j \in N} \Delta(C_j), \quad \forall k \in \{1, 2, 3, \dots\},$$

$$\tau^k \in R(\sigma^k), \quad \forall k \in \{1, 2, 3, \dots\},$$

$$\bar{\sigma} = \lim_{k \rightarrow \infty} \sigma^k, \quad \text{and}$$

$$\bar{\tau} = \lim_{k \rightarrow \infty} \tau^k.$$

These conditions imply that

$$\rho_i(\sigma_{-i}^k, \tau_i^k) \leq \rho_i(\sigma_{-i}^k, \xi_i), \quad \forall i \in N, \quad \forall \xi_i \in \Delta(C_i), \quad \forall k \in \{1, 2, 3, \dots\}.$$

By continuity of the weighted regret function  $\rho_i$ , this in turn implies that

$$\rho_i(\bar{\sigma}_{-i}, \bar{\tau}_i) \leq \rho_i(\bar{\sigma}_{-i}, \xi_i), \quad \forall i \in N, \quad \forall \xi_i \in \Delta(C_i).$$

So,

$$\bar{\tau}_i \in R_i(\bar{\sigma}_{-i}), \quad \forall i \in N,$$

and so,  $\bar{\tau} \in R(\bar{\sigma})$ . Thus, the correspondence  $R: \times_{j \in N} \Delta(C_j) \rightarrow \times_{j \in N} \Delta(C_j)$  is upper-hemicontinuous.

Thus, by the Kakutani fixed-point theorem (Kakutani, 1941), there exists some randomized-strategy profile  $\sigma$  in  $\times_{j \in N} \Delta(C_j)$  such that  $\sigma \in R(\sigma)$ . That is,  $\sigma$  is a mutual weighted least-squares regret profile in randomized strategies. ■

## 8.4 Mutual Weighted Least-Squares Regret and Nash Equilibrium

It is natural to ask whether there is any connection between mutual weighted least-squares regret and Nash equilibrium. The following theorem answers this question simply. While the theorem is straightforward to state and to prove, we make it explicit in order to emphasize the effect of considering randomized strategies.

**THEOREM 8.4.1.** *Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Then, for any pure-strategy profile  $c$  in  $\times_{j \in N} C_j$ , the pure-strategy profile  $c$  is an equilibrium of  $\Gamma$  if and only if  $c$  is a mutual weighted least-squares regret profile in randomized strategies.*

*Proof.* Let  $c$  be any pure-strategy profile in  $\times_{j \in N} C_j$ . Suppose that  $c$  is an equilibrium. Then

$$u_i(c_{-i}, c_i) \geq u_i(c_{-i}, d_i), \quad \forall i \in N, \quad \forall d_i \in C_i.$$

And so,

$$u_i(c_{-i}, c_i) = \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i), \quad \forall i \in N.$$

Clearly,

$$\rho_i(c_{-i}, c_i) = \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, c_i) \right)^2 = 0, \quad \forall i \in N.$$

Thus,

$$\rho_i(c_{-i}, c_i) \leq \rho_i(c_{-i}, \tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i).$$

That is,  $c$  is a mutual weighted least-squares regret profile in randomized strategies.

Now, let  $c$  be any pure-strategy profile in  $\times_{j \in N} C_j$ . Suppose that  $c$  is a mutual weighted least-squares regret profile in randomized strategies. Then

$$\rho_i(c_{-i}, c_i) \leq \rho_i(c_{-i}, \tau_i), \quad \forall i \in N, \quad \forall \tau_i \in \Delta(C_i).$$

Notice that

$$\rho_i(c_{-i}, \xi_i) = \left( \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i) - u_i(c_{-i}, \xi_i) \right)^2, \quad \forall i \in N, \quad \forall \xi_i \in \Delta(C_i).$$

Since

$$c_i \in \operatorname{argmin}_{\tau_i \in \Delta(C_i)} \rho_i(c_{-i}, \tau_i) \quad \text{and} \quad \min_{\tau_i \in \Delta(C_i)} \rho_i(c_{-i}, \tau_i) = 0, \quad \forall i \in N,$$

it follows that

$$u_i(c_{-i}, c_i) = \max_{\tau_i \in \Delta(C_i)} u_i(c_{-i}, \tau_i), \quad \forall i \in N.$$

Thus,

$$u_i(c_{-i}, c_i) \geq u_i(c_{-i}, d_i), \quad \forall i \in N, \quad \forall d_i \in C_i.$$

That is,  $c$  is an equilibrium. ■

For an illustration of Theorem 8.4.1, consider the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 8.2.

Table 8.2 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	2 3	1 0
	$y_1$	1 0	2 2

This game has two equilibria in pure strategies:  $(x_1, x_2)$ , which gives the payoff allocation  $(3, 2)$ , and  $(y_1, y_2)$ , which gives the payoff allocation  $(2, 2)$ .

It is straightforward to verify that these two equilibria are also the only two mutual weighted least-squares regret profiles in randomized strategies of this game. The weighted regret function  $\rho_1: \prod_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_1(\sigma_1, \sigma_2) &= \sigma_2(x_2) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, x_2) - u_1(\sigma_1, x_2) \right)^2 + \sigma_2(y_2) \left( \max_{\tau_1 \in \Delta(C_1)} u_1(\tau_1, y_2) - u_1(\sigma_1, y_2) \right)^2 \\ &= \sigma_2(x_2)(3 - 3\sigma_1(x_1))^2 + (1 - \sigma_2(x_2))(2 - 2(1 - \sigma_1(x_1)))^2, \end{aligned}$$

and so,

$$\frac{\partial(\rho_1(\sigma_1, \sigma_2))}{\partial \sigma_1(x_1)} = 5\sigma_1(x_1)\sigma_2(x_2) + 4\sigma_1(x_1) - 9\sigma_2(x_2).$$

The weighted regret function  $\rho_2: \times_{j \in N} \Delta(C_j) \rightarrow \mathbf{R}$  is

$$\begin{aligned} \rho_2(\sigma_1, \sigma_2) &= \sigma_1(x_1) \left( \max_{\tau_2 \in \Delta(C_2)} u_2(x_1, \tau_2) - u_2(x_1, \sigma_2) \right)^2 + \sigma_1(y_1) \left( \max_{\tau_2 \in \Delta(C_2)} u_2(y_1, \tau_2) - u_2(y_1, \sigma_2) \right)^2 \\ &= \sigma_1(x_1) (2 - (2\sigma_2(x_2) + (1 - \sigma_2(x_2))))^2 + (1 - \sigma_1(x_1)) (2 - (\sigma_2(x_2) + 2(1 - \sigma_2(x_2))))^2, \end{aligned}$$

and so,

$$\frac{\partial(\rho_2(\sigma_1, \sigma_2))}{\partial \sigma_2(x_2)} = \sigma_2(x_2) - \sigma_1(x_1).$$

Setting each of the partial derivatives above equal to 0 and solving the resulting system of equations yield

$$\sigma_1(x_1) = 1 \quad \text{and} \quad \sigma_2(x_2) = 1$$

and

$$\sigma_1(x_1) = 0 \quad \text{and} \quad \sigma_2(x_2) = 0,$$

which are precisely the equilibria  $(x_1, x_2)$  and  $(y_1, y_2)$ , respectively.

When randomized strategies are considered, the set of mutual weighted least-squares regret profiles in randomized strategies and the set of equilibria need not coincide. Notice that the game shown in Table 8.2 has one equilibrium in randomized strategies, namely,

$$(0.5x_1 + 0.5y_1, 0.4x_2 + 0.6y_2),$$

which gives the payoff allocation  $(1.2, 1.5)$ . But, as just shown, the game has just two mutual weighted least-squares regret profiles in randomized strategies, and in each such profile, each player plays some pure strategy with probability 1.

For another illustration of the foregoing point, consider again the finite two-person game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  in strategic form shown in Table 8.1 in Section 8.2 and reproduced in Table 8.3.

The unique equilibrium of this game is

$$(0.5x_1 + 0.5y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(7/3, 0.5)$ , and, as noted in Section 8.2, the unique mutual weighted least-squares regret profile in randomized strategies is

$$((2/3)x_1 + (1/3)y_1, (2/3)x_2 + (1/3)y_2),$$

which gives the payoff allocation  $(7/3, 5/9)$ .

Table 8.3 A game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1    0	2    3
	$y_1$	0    1	3    1

Thus, the set of mutual weighted least-squares regret profiles in randomized strategies and the set of equilibria are disjoint. Notably, while the payoff to player 1 under mutual weighted least-squares regret is the same as that under equilibrium, the payoff to player 2 under mutual weighted least-squares regret is strictly greater than that under equilibrium.

That mutual weighted least-squares regret and Nash equilibrium should diverge when randomized strategies are considered is unsurprising. They represent different ways to reason about a game. Mutual weighted least-squares regret involves choosing a strategy so as to minimize the divergence from the best-response payoffs. Nash equilibrium involves choosing a strategy so as to maximize the expected payoff.

### 8.5 Recursion and Convergence

As discussed in Section 8.3, Theorem 8.3.1 establishes that every finite game in strategic form has at least one mutual weighted least-squares regret profile in randomized strategies. One natural question that arises is the following. Let  $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$  be any finite game in strategic form. Let  $(\sigma^k)_{k=1}^\infty$  be any sequence defined recursively thus:

$$\sigma^1 \in \prod_{j \in N} \Delta(C_j) \quad \text{and}$$

$$\sigma^{k+1} \in R(\sigma^k), \quad \forall k \in \{1, 2, 3, \dots\}.$$

The question is whether, for any sequence  $(\sigma^k)_{k=1}^\infty$  defined recursively as above, there exists some randomized-strategy profile  $\bar{\sigma}$  in  $\prod_{j \in N} \Delta(C_j)$  such that

$$\lim_{k \rightarrow \infty} \sigma^k = \bar{\sigma} \quad \text{and} \quad \bar{\sigma} \in R(\bar{\sigma}).$$

Intuitively, the question is whether recursively updating the probability distributions by iteratively minimizing the respective weighted regret functions of the players yields convergence to a mutual weighted least-squares regret profile in randomized strategies.

As might be expected, the answer is no. For an illustration of this point, consider again the Battles of the Sexes game  $\Gamma = (\{1, 2\}, C_1, C_2, u_1, u_2)$  shown in Table 6.10 in Section 6.7 and reproduced in Table 8.4.

Table 8.4 Battle of the Sexes game in strategic form

		2	
		$x_2$	$y_2$
1	$x_1$	1 3	0 0
	$y_1$	0 0	3 1

Consider the sequence  $(\sigma^k)_{k=1}^\infty$ , recursively defined as above, such that

$$\sigma^1 = (0.5x_1 + 0.5y_1, 0.5x_2 + 0.5y_2).$$

It is straightforward to verify that

$$\begin{aligned} \sigma^k &= (0.5x_1 + 0.5y_1, 0.5x_2 + 0.5y_2) && \text{if } k \text{ is odd,} \\ &= (0.9x_1 + 0.1y_1, 0.1x_2 + 0.9y_2) && \text{if } k \text{ is even.} \end{aligned}$$

That is, the sequence  $(\sigma^k)_{k=1}^\infty$  oscillates between

$$(0.5x_1 + 0.5y_1, 0.5x_2 + 0.5y_2)$$

and

$$(0.9x_1 + 0.1y_1, 0.1x_2 + 0.9y_2)$$

and so does not converge. Thus, even though, by Theorem 8.3.1 in Section 8.3, a mutual weighted least-squares regret profile in randomized strategies is guaranteed to exist, it cannot be attained via the recursive process defined above if

$$\sigma^1 = (0.5x_1 + 0.5y_1, 0.5x_2 + 0.5y_2),$$

as specified.

# 9

---

## Conclusion

As we have seen with games such as the Traveler's Dilemma in Chapter 4 and those in Chapter 6, the established solution concepts can sometimes yield solutions that seem in some ways unsatisfactory. The need to develop solution concepts that reflect more accurately and effectively how people actually reason about and play games has long been recognized, and new solution concepts continue to be developed.

The question of how people actually reason about and play games is particularly significant and not simply a matter of description or practice since it is impossible to get far in game theory without attending to it. And this is because the question of how one ought to act in a given situation depends significantly on how others can be expected actually to act. Thus, the project of developing a good normative theory is inseparable from that of developing a good descriptive theory.

We have endeavored to develop a new solution concept that seeks to provide an intuitive characterization of reasonable or observed behavior in a wide range of games, including some that have proved problematic for standard game theory. The essence of this dissertation consists in showing that our proposed solution concept of least-squares regret satisfies the relevant criteria, exhibits the desired properties, and yields the expected or reasonable solutions. Of course, work remains to be done, and Section 9.1 proposes some questions for further research.

### 9.1 Further Questions

While we have endeavored to develop and to defend least-squares regret, we view this dissertation essentially as an exploration of a new approach to solving noncooperative games. But a number of questions remain to be answered.

One important question to be answered more fully is whether people, in fact, act in accordance with least-squares regret in particular or with some form or



another of regret minimization in general.

One way to answer this question would be to develop further games and experiments and to test specifically for reasoning based on regret. The results relating to the Traveler's Dilemma in Chapter 4 and to the various games in Chapter 6 as well as the experimental results in psychology noted in Section 1.3 are promising and suggest that regret may figure importantly in reasoning about a game. Furthermore, the results relating to the games proposed in Section 3.2, games that could form the basis of further experiments, suggest that reasoning about a game is characterized more by something like least-squares regret than by some minimax regret approach such as iterated regret minimization.

Another way to determine whether people, in fact, act in accordance with least-squares regret would be to axiomatize least-squares regret in the manner of Stoye (2011) in the case of minimax regret or of von Neumann and Morgenstern (1947) in the case of expected utility maximization. An analysis of the axioms underlying least-squares regret might then help us to determine how intuitive, credible, and compelling least-squares regret might be as a characterization of typical reasoning and behavior.

As discussed in Sections 1.2, 1.4, 2.2, 2.5, and 7.4, least-squares regret considers partially strategic players. In Chapter 8, we introduced mutual weighted least-squares regret, which, as noted, considers fully strategic players. It would be worthwhile and very natural to explore the intermediate case of *weighted least-squares regret*, which can be understood as least-squares regret with the admission of nonuniform probability distributions or, equivalently, as mutual weighted least-squares regret without the mutuality condition. Intuitively, the idea is that a player might be significantly strategic and devise fairly reasonable beliefs about the strategies of the other players without being fully strategic and having perfect beliefs. The aim, then, would be to model the formation of the probability distributions to be used as weights, perhaps by appealing to least-squares regret or mutual weighted least-squares regret as benchmark models.

In the interest of tractability, we have restricted attention to finite games in strategic form and, where tractable, to simple games with infinite strategy sets. It would be worthwhile to extend least-squares regret to other classes of games.

Least-squares regret could be extended to Bayesian games. Perhaps the most straightforward approach would be simply to consider, for any Bayesian game, least-squares regret with respect to each type of each player, but there may be alternative approaches that are more flexible or fruitful.

Least-squares regret could be extended also to games in extensive form. While a game in extensive form can be reduced in the usual way to a game in strategic form and thereby made amenable to least-squares regret, there is good reason to

develop a version of least-squares regret that can be applied directly to games in extensive form. To begin, the reduction of a game in extensive form to a game in strategic form may involve the loss of significant information so that a full analysis of the initial game in extensive form may differ in important ways from that of the resulting game in strategic form. Furthermore, extending least-squares regret to games in extensive form may involve considering important and interesting intricacies and subtleties that do not arise in games in strategic form and can thus lead to a better understanding of the nature of regret and regret minimization more generally. Least-squares regret applied directly to a game in extensive form may yield solutions and insights very different from those that emerge when it is applied to the corresponding game in strategic form.

One further area to which least-squares regret can be extended is mechanism design. Just as a number of important mechanisms have been developed on the basis of Nash equilibrium, there remains the prospect of developing alternative mechanisms on the basis of least-squares regret instead. Particularly attractive areas include auction theory and bargaining theory given their practical significance and the usefulness of least-squares regret, as illustrated in Sections 6.1, 6.2, and 6.8, in reasoning about auctions and auction-like games and bargaining games, situations in which considerations of regret appear to be especially forceful.

While this dissertation is essentially an exploration of an alternative approach to solving games and while numerous questions remain, it is our hope that least-squares regret can be developed further and become a valuable supplement to the canon of solution concepts for noncooperative games.

---

# Bibliography

- Anderson, Simon P., Jacob K. Goree, and Charles A. Holt (2002). "The Logit Equilibrium: A Perspective on Intuitive Behavioral Anomalies". *Southern Economic Journal* (69)1: 21–47.
- Aumann, Robert J. (1974). "Subjectivity and Correlation in Randomized Strategies". *Journal of Mathematical Economics* (1)1: 67–96.
- (1987). "Correlated Equilibria as an Expression of Bayesian Rationality". *Econometrica* (55)1: 1–18.
- (1997). "Rationality and Bounded Rationality". *Games and Economic Behavior* (21)1–2: 2–14.
- Axelrod, Robert (1981). "The Emergence of Cooperation Among Egoists". *American Political Science Review* (75)2: 306–318.
- Bach, Christian W. and Andrés Perea (2014). "Utility Proportional Beliefs". *International Journal of Game Theory* (43)4: 881–902.
- Baghestanian, Sascha (2014). "Unraveling the Traveler's Dilemma Puzzle. A Level-*k* Approach". Goethe Universität Frankfurt am Main, Working Paper.
- Basu, Kaushik (1994). "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory". *American Economic Review* (84)2: 391–395.
- (2007). "The Traveler's Dilemma". *Scientific American* (296) 90–95.
- Basu, Kaushik, Leonardo Becchetti, and Luca Stanca (2011). "Experiments with the Traveler's Dilemma: Welfare, Strategic Choice and Implicit Collusion". *Social Choice and Welfare* (37)4: 575–595.
- Bavly, Gilad (2012). "Uncertainty in the Traveler's Dilemma". Center for the Study of Rationality, Hebrew University of Jerusalem, Discussion Paper No. 595.
- Becker, Tilman, Michael Carter, and Jörg Naeve (2005). "Experts Playing the Traveler's Dilemma". Universität Hohenheim, Discussion Paper Nr. 252/2005.
- Bernheim, B. Douglas (1984). "Rationalizable Strategic Behavior". *Econometrica* (52)4: 1007–1028.

- Bernoulli, Jakob (1713). *Ars conjectandi, opus posthumum. Accedit Tractatus de seriebus infinitis, et epistola gallicé scripta de ludo pilae reticularis*. Basel, Switzerland: Thurneysen Brothers.
- Bertrand, Joseph Louis François (1883). “Théorie des Richesses: revue de *Théories mathématiques de la richesse sociale* par Léon Walras et *Recherches sur les principes mathématiques de la théorie des richesses* par Augustin Cournot”. *Journal des Savants* (67) 499–508.
- Blume, Lawrence E., Adam M. Brandenburger, and Eddie Dekel (1991). “Lexicographic Probabilities and Choice Under Uncertainty”. *Econometrica* (59)1: 61–79.
- Borel, Félix Édouard Justin Émile (1921). “La théorie du jeu et les équations intégrales à noyau symétrique gauche”. *Comptes Rendus Hebdomadaires des Séances de l’Académie des Sciences* (173) 1304–1308.
- Brandenburger, Adam M., Amanda Friedenberg, and H. Jerome Keisler (2008). “Admissibility in Games”. *Econometrica* (76)2: 307–352.
- Burchardi, Konrad B. and Stefan P. Penczynski (2012). “Out of Your Mind: Eliciting Individual Reasoning in One Shot Games”. Working Paper.
- Cabrera, Susana, C. Mónica Capra, and Rosario Gómez (2007). “Behavior in One-Shot Traveler’s Dilemma Games: Model and Experiments with Advice”. *Spanish Economic Review* (9)2: 129–152.
- Camerer, Colin F. and Teck-Hua Ho (2015). “Behavioral Game Theory Experiments and Modeling”. In: *Handbook of Game Theory with Economic Applications*. Ed. by H. Peyton Young and Shmuel Zamir. New York, NY: Elsevier B.V., 517–573.
- Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong (2004). “A Cognitive Hierarchy Model of Games”. *Quarterly Journal of Economics* (119)3: 861–898.
- Capra, C. Mónica, Jacob K. Goeree, Rosario Gómez, and Charles A. Holt (1999). “Anomalous Behavior in a Traveler’s Dilemma?” *American Economic Review* (89)3: 678–690.
- Chakravarty, Sujoy, Emmanuel Dechenaux, and Jaideep Roy (2010). “Ill-Defined versus Precise Pre-Play Communication in the Traveler’s Dilemma”. *Southern Economic Journal* (77)2: 351–368.
- Chernoff, Herman (1954). “Rational Selection of Decision Functions”. *Econometrica* (22)4: 422–443.
- Crawford, Vincent P. (2013). “Boundedly Rational versus Optimization-Based Models of Strategic Thinking and Learning in Games”. *Journal of Economic Literature* (51)2: 512–527.
- Crawford, Vincent P., Miguel A. Costa-Gomes, and Nagore Iriberry (2013). “Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications”. *Journal of Economic Literature* (51)1: 5–62.

- Dreber, Anna, Tore Ellingsen, Magnus Johannesson, and David G. Rand (2013). "Do People Care About Social Context? Framing Effects in Dictator Games". *Experimental Economics* (16)3: 349–371.
- Dufwenberg, Martin, Simon Gächter, and Heike Hennig-Schmidt (2011). "The Framing of Games and the Psychology of Play". *Games and Economic Behavior* (73)2: 459–478.
- Dufwenberg, Martin and Uri Gneezy (2000). "Price Competition and Market Concentration: An Experimental Study". *International Journal of Industrial Organization* (18)1: 7–22.
- Eliasz, Kfir and Ariel Rubinstein (2011). "Edgar Allan Poe's Riddle: Framing Effects in Repeated Matching Pennies Games". *Games and Economic Behavior* (71)1: 88–99.
- Ellingsen, Tore, Magnus Johannesson, Johanna Mollerstrom, and Sara Munkhammar (2012). "Social Framing Effects: Preferences or Beliefs?" *Games and Economic Behavior* (76)1: 117–130.
- Filiz-Ozbay, Emel and Erkut Y. Ozbay (2007). "Auctions with Anticipated Regret: Theory and Experiment". *American Economic Review* (97)4: 1407–1418.
- Fudenberg, Drew (2010). "Predictive Game Theory". American Economic Association, Ten Years and Beyond: Economists Answer NSF's Call for Long-Term Research Agendas.
- Fudenberg, Drew and David Levine (1998). *The Theory of Learning in Games*. Cambridge, MA: The MIT Press.
- Fudenberg, Drew and Jean Tirole (1991). *Game Theory*. Cambridge, MA: The MIT Press.
- Goeree, Jacob K. and Charles A. Holt (1999). "Stochastic Game Theory: For Playing Games, Not Just for Doing Theory". *Proceedings of the National Academy of Sciences* (96)19: 10564–10567.
- (2001). "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions". *American Economic Review* (91)5: 1402–1422.
- (2004). "A Model of Noisy Introspection". *Games and Economic Behavior* (46)2: 365–382.
- Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2003). "Risk Averse Behavior in Generalized Matching Pennies Games". *Games and Economic Behavior* (45)1: 97–113.
- (2005). "Regular Quantal Response Equilibrium". *Experimental Economics* (8)4: 347–367.
- Govindan, Srihari and Robert B. Wilson (2008). "Refinements of Nash Equilibrium". In: *The New Palgrave Dictionary of Economics*. Ed. by Steven N. Durlauf and Lawrence E. Blume. New York, NY: Palgrave MacMillan.

- Greene, William H. (2011). *Econometric Analysis*. Upper Saddle River, NJ: Prentice Hall.
- Grosskopf, Brit, Ido Erev, and Eldad Yechiam (2006). “Forgone with the Wind: Indirect Payoff Information and its Implications for Choice”. *International Journal of Game Theory* (34)2: 285–302.
- Haile, Philip A., Ali Hortaçsu, and Grigory Kosenok (2008). “On the Empirical Content of Quantal Response Equilibrium”. *American Economic Review* (98)1: 180–200.
- Halpern, Joseph Y. (2008). “Beyond Nash Equilibrium: Solution Concepts for the 21st Century”. In: *Proceedings of the Twenty-Seventh ACM Symposium on Principles of Distributed Computing*. New York, NY: ACM, 1–10.
- Halpern, Joseph Y. and Samantha Leung (2014). “Weighted Sets of Probabilities and Minimax Weighted Expected Regret: A New Approach for Representing Uncertainty and Making Decisions”. *Theory and Decision*, 1–36.
- Halpern, Joseph Y. and Rafael Pass (2012). “Iterated Regret Minimization: A New Solution Concept”. *Games and Economic Behavior* (74)1: 184–207.
- Harsanyi, John C. (1982). “Subjective Probability and the Theory of Games: Comments on Kadane and Larkey’s Paper”. *Management Science* (28)2: 120–124.
- Harsanyi, John C. and Reinhard Selten (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: The MIT Press.
- Hillas, John B. and Elon Kohlberg (2002). “Foundations of Strategic Equilibrium”. In: *Handbook of Game Theory with Economic Applications*. Ed. by Robert J. Aumann and Sergiu Hart. New York, NY: Elsevier B.V., 1597–1663.
- Holt, Charles A. and Alvin E. Roth (2004). “The Nash Equilibrium: A Perspective”. *Proceedings of the National Academy of Sciences* (101)12: 3999–4002.
- Jaynes, Edwin T. (1957a). “Information Theory and Statistical Mechanics”. *Physical Review* (106)4: 620–630.
- (1957b). “Information Theory and Statistical Mechanics. II”. *Physical Review* (108)2: 171–190.
- Kadane, Joseph B. and Patrick D. Larkey (1982). “Subjective Probability and the Theory of Games”. *Management Science* (28)2: 113–120.
- Kahneman, Daniel and Amos N. Tversky (1979). “Prospect Theory: An Analysis of Decision under Risk”. *Econometrica* (47)2: 263–292.
- (1984). “Choices, Values, and Frames”. *American Psychologist* (39)4: 341–350.
- Kakutani, Shizuo (1941). “A Generalization of Brouwer’s Fixed Point Theorem”. *Duke Mathematical Journal* (8)3: 457–459.
- Kalai, Ehud and Dov Samet (1984). “Persistent Equilibria in Strategic Games”. *International Journal of Game Theory* (13)3: 129–144.

- Kohlberg, Elon and Jean-François Mertens (1986). "On the Strategic Stability of Equilibria". *Econometrica* (54)5: 1003–1037.
- Kreps, David M. (1990). *Game Theory and Economic Modelling*. New York, NY: Oxford University Press.
- Kreps, David M., Paul Milgrom, John Roberts, and Robert B. Wilson (1982). "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma". *Journal of Economic Theory* (27)2: 245–252.
- Kreps, David M. and Robert B. Wilson (1982). "Sequential Equilibria". *Econometrica* (50)4: 863–894.
- Laplace, Pierre-Simon (1825). *Essai philosophique sur les probabilités*. Paris, France: Bachelier.
- Levin, Irwin P., Sandra L. Schneider, and Gary J. Gaeth (1998). "All Frames Are Not Created Equal: A Typology and Critical Analysis of Framing Effects". *Organizational Behavior and Human Decision Processes* (76)2: 149–188.
- Linhart, Peter B. and Roy Radner (1989). "Minimax-Regret Strategies for Bargaining Over Several Variables". *Journal of Economic Theory* (48)1: 152–178.
- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory*. New York, NY: Oxford University Press.
- McKelvey, Richard D. and Thomas R. Palfrey (1995). "Quantal Response Equilibria for Normal Form Games". *Games and Economic Behavior* (10)1: 6–38.
- (1998). "Quantal Response Equilibria for Extensive Form Games". *Experimental Economics* (1)1: 9–41.
- McKelvey, Richard D., Thomas R. Palfrey, and Roberto A. Weber (2000). "The Effects of Payoff Magnitude and Heterogeneity on Behavior in  $2 \times 2$  Games with Unique Mixed Strategy Equilibria". *Journal of Economic Behavior and Organization* (42)4: 523–548.
- Milnor, John W. (1954). "Games Against Nature". In: *Decision Processes*. Ed. by Robert M. Thrall, Clyde H. Coombs, and Robert L. Davis. New York NY: John Wiley & Sons, Inc., 49–59.
- Morris, Stephen E., Rafael Rob, and Hyun Song Shin (1995). " $p$ -Dominance and Belief Potential". *Econometrica* (63)1: 145–157.
- Moulin, Hervé (1986). *Game Theory for the Social Sciences*. New York, NY: New York University Press.
- Moulin, Hervé and Jean-Philippe Vial (1978). "Strategically Zero-Sum Games: The Class of Games Whose Completely Mixed Equilibria Cannot Be Improved Upon". *International Journal of Game Theory* (7)3–4: 201–221.
- Myerson, Roger B. (1978). "Refinements of the Nash Equilibrium Concept". *International Journal of Game Theory* (7)2: 73–80.

- Myerson, Roger B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- (1999). “Nash Equilibrium and the History of Economic Theory”. *Journal of Economic Literature* (37)3: 1067–1082.
- Nagel, Rosemarie (1995). “Unraveling in Guessing Games: An Experimental Study”. *American Economic Review* (85)5: 1313–1326.
- Nash, John F. (1950a). “Equilibrium Points in  $n$ -Person Games”. *Proceedings of the National Academy of Sciences* (36)1: 48–49.
- (1950b). “The Bargaining Problem”. *Econometrica* (18)2: 155–162.
- (1951). “Noncooperative Games”. *Annals of Mathematics* (54)2: 286–295.
- (1953). “Two-Person Cooperative Games”. *Econometrica* (21)1: 128–140.
- Niehans, Jürg (1948). “Zur Preisbildung bei Ungewissen Erwartungen”. *Schweizerische Zeitschrift für Volkswirtschaft und Statistik* (84)5: 433–456.
- Nosenzo, Daniele, Theo Offerman, Martin Sefton, and Ailko van der Veen (2014). “Encouraging Compliance: Bonuses Versus Fines in Inspection Games”. *Journal of Law, Economics, and Organization* (30)3: 623–648.
- Ochs, Jack (1995). “Games with Unique Mixed Strategy Equilibria: An Experimental Study”. *Games and Economic Behavior* (10)1: 202–217.
- Osborne, Martin J. and Ariel Rubinstein (1994). *A Course in Game Theory*. Cambridge, MA: The MIT Press.
- Pearce, David G. (1984). “Rationalizable Strategic Behavior and the Problem of Perfection”. *Econometrica* (52)4: 1029–1050.
- Raiffa, Howard (1982). *The Art and Science of Negotiation*. Cambridge, MA: Harvard University Press.
- Renou, Ludovic and Karl H. Schlag (2010). “Minimax Regret and Strategic Uncertainty”. *Journal of Economic Theory* (145)1: 264–286.
- Ritov, Ilana (1996). “Probability of Regret: Anticipation of Uncertainty Resolution in Choice”. *Organizational Behavior and Human Decision Processes* (66)2: 228–236.
- Rosenthal, Robert W. (1981). “Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox”. *Journal of Economic Theory* (25)1: 92–100.
- Roth, Alvin E. and Francoise Schoumaker (1983). “Subjective Probability and the Theory of Games: Some Further Comments”. *Management Science* (29)11: 1137–1140.
- Rubinstein, Ariel (1998). *Modeling Bounded Rationality*. Cambridge, MA: The MIT Press.
- (2006). “Dilemmas of an Economic Theorist”. *Econometrica* (74)4: 865–883.
- (2007). “Instinctive and Cognitive Reasoning: A Study of Response Times”. *The Economic Journal* (117)523: 1243–1259.



- Samuelson, Larry (1992). "Dominated Strategies and Common Knowledge". *Games and Economic Behavior* (4)2: 284–313.
- Savage, Leonard J. (1951). "The Theory of Statistical Decision". *Journal of the American Statistical Association* (46)253: 55–67.
- (1954). *The Foundations of Statistics*. New York, NY: John Wiley & Sons, Inc.
- Schelling, Thomas C. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Selten, Reinhard (1975). "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games". *International Journal of Game Theory* (4)1: 25–55.
- Selten, Reinhard, Klaus Abbink, and Ricarda Cox (2005). "Learning Direction Theory and the Winner's Curse". *Experimental Economics* (8)1: 5–20.
- Selten, Reinhard and Joachim Buchta (1999). "Experimental Sealed Bid First Price Auctions with Directly Observed Bid Functions". In: *Games and Human Behavior: Essays in Honor of Amnon Rapoport*. Ed. by David V. Budescu, Ido Erev, and Rami Zwick. Mahwah, NJ: Lawrence Erlbaum Associates, 79–104.
- Selten, Reinhard and Rolf Stoecker (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames". *Journal of Economic Behavior and Organization* (7)1: 47–70.
- Shubik, Martin (2012). "The Present and Future of Game Theory". *Singapore Economic Review* (57)1: 1–14.
- Simonsen, Mario H. (1988). "Rational Expectations, Game Theory, and Inflationary Inertia". In: *The Economy as an Evolving Complex System*. Ed. by Philip W. Anderson, Kenneth J. Arrow, and David Pines. Redwood City, CA: Addison-Wesley, 205–241.
- Simonson, Itamar and Amos N. Tversky (1992). "Choice in Context: Tradeoff Contrast and Extremeness Aversion". *Journal of Marketing Research* (29)3: 281–295.
- Sinn, Hans-Werner (1980). "A Rehabilitation of the Principle of Insufficient Reason". *Quarterly Journal of Economics* (94)3: 493–506.
- Smale, Stephen (1980). "The Prisoner's Dilemma and Dynamical Systems Associated to Non-Cooperative Games". *Econometrica* (48)7: 1617–1634.
- Stahl, Dale O. and Paul W. Wilson (1994). "Experimental Evidence on Players' Models of Other Players". *Journal of Economic Behavior and Organization* (25)3: 309–327.
- (1995). "On Players' Models of Other Players: Theory and Experimental Evidence". *Games and Economic Behavior* (10)1: 218–254.
- Stoye, Jörg (2011). "Axioms for Minimax Regret Choice Correspondences". *Journal of Economic Theory* (146)6: 2226–2251.

- Tversky, Amos N. and Daniel Kahneman (1981). "The Framing of Decisions and the Psychology of Choice". *Science* (211)4481: 453–458.
- (1991). "Loss Aversion in Riskless Choice: A Reference-Dependent Model". *Quarterly Journal of Economics* (106)4: 1039–1061.
- Tversky, Amos N. and Itamar Simonson (1993). "Context-Dependent Preferences". *Management Science* (39)10: 1179–1189.
- van Damme, Eric (1991). *Stability and Perfection of Nash Equilibria*. New York, NY: Springer-Verlag.
- (2002). "Strategic Equilibrium". In: *Handbook of Game Theory with Economic Applications*. Ed. by Robert J. Aumann and Sergiu Hart. New York, NY: Elsevier B.V., 1521–1596.
- von Neumann, John (1928). "Zur Theorie der Gesellschaftsspiele". *Mathematische Annalen* (100)1: 295–320.
- von Neumann, John and Oskar Morgenstern (1947). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- von Stengel, Bernhard (2011). *Game Theory Basics*. Unpublished manuscript. London School of Economics and Political Science.
- Wald, Abraham (1939). "Contributions to the Theory of Statistical Estimation and Testing Hypotheses". *Annals of Mathematics* (10)4: 299–326.
- (1945). "Statistical Decision Functions Which Minimize the Maximum Risk". *Annals of Mathematics* (46)2: 265–280.
- (1950). *Statistical Decision Functions*. New York, NY: John Wiley & Sons, Inc.
- Wooldridge, Jeffrey M. (2013). *Introductory Econometrics: A Modern Approach*. Mason, OH: Cengage Learning.
- Wright, James R. and Kevin Leyton-Brown (2010). "Beyond Equilibrium: Predicting Human Behavior in Normal-Form Games". In: *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*. AAAI, 901–907.