

University of St Andrews

School of Economics & Finance Discussion Papers

Complexity and Bounded Rationality in Individual Decision Problems

Theodoros M. Diasakos

School of Economics and Finance Discussion Paper No. 1314

JEL Classification:
Keywords:

School of Economics & Finance
Discussion Paper Series
issn 0962-4031
http://ideas.repec.org/s/san/wpecon.html
info: econ@st-andrews.ac.uk

Complexity and Bounded Rationality in Individual Decision Problems

Theodoros M. Diasakos*

October 7, 2013

Abstract

I develop a model of endogenous bounded rationality due to search costs, arising implicitly from the problem's complexity. The decision maker is not required to know the entire structure of the problem when making choices but can think ahead, through costly search, to reveal more of it. However, the costs of search are not assumed exogenously; they are inferred from revealed preferences through her choices. Thus, bounded rationality and its extent emerge endogenously: as problems become simpler or as the benefits of deeper search become larger relative to its costs, the choices more closely resemble those of a rational agent. For a fixed decision problem, the costs of search will vary across agents. For a given decision maker, they will vary across problems. The model explains, therefore, why the disparity, between observed choices and those prescribed under rationality, varies across agents and problems. It also suggests, under reasonable assumptions, an identifying prediction: a relation between the benefits of deeper search and the depth of the search. As long as calibration of the search costs is possible, this can be tested on any agent-problem pair. My approach provides a common framework for depicting the underlying limitations that force departures from rationality in different and unrelated decision-making situations. Specifically, I show that it is consistent with violations of timing independence in temporal framing problems, dynamic inconsistency and diversification bias in sequential versus simultaneous choice problems, and with plausible but contrasting risk attitudes across small- and large-stakes gambles.

*School of Economics & Finance, University of St Andrews, Scotland (E-mail: td31@st-andrews.ac.uk). I am grateful to Chris Shannon for invaluable advice on this and earlier versions. Discussions with David Ahn, Bob Anderson, Paolo Ghirardato, Shachar Kariv, Botond Koszegi, Matthew Rabin, Jacob Sagi, and Adam Szeidl were of significant benefit. I also received helpful comments from Art Boman, Zack Grossman, Kostis Hatzitaskos, Kristof Madarasz, George Petropoulos, Florence Neymotin, and participants in the Xth Spring Meeting of Young Economists, the 2005 Conference of the Southwestern Economic Association, the 3rd and 5th Conferences on Research in Economic Theory and Econometrics, the VIII and XIII SAET conferences, and the 2008 RUD conference. Needless to say, all errors are mine. This research was supported by a Fellowship from the State Scholarship Foundation of Greece.

1 Introduction

Most economic models assume rational decision-making which requires, at a minimum, that the agent understands the entire decision problem, including all possible plans of action and the consequences of each, has well-defined preferences over final outcomes, and chooses optimally according to these preferences. For all but the simplest problems, however, being able to make choices in this way amounts to having access to extraordinary cognitive and computational abilities and resources. It is not surprising, therefore, that, in a variety of models, the predictions arising from rationality are systematically violated across a spectrum of experimental and actual decision makers.[1] Theories of bounded rationality attempt to address these discrepancies by accounting for agents with limited abilities to formulate or solve complex problems, or to process information, as first advocated by Simon [71], [72], [73] in his pioneering work.

Bounded rationality has been incorporated into game theory (Abreu and Rubinstein [1], Piccione and Rubinstein [56], Osborne and Rubinstein [54], Rosenthal [62], Camerer *et al.* [9]), bargaining (Sabourian [65]), auctions (Crawford and Iriberri [16]), macroeconomics (Sargent [67]), and contracting (Anderlini and Felli [3], Lee and Sabourian [48]), to name but a few areas and papers, in models that take the cognitive, computational, or information-processing limitations as fixed and exogenous. Yet, a feature of bounded rationality in practice is that the exact nature and extent of such limitations vary across both decision makers and decision problems. More importantly, the variation can depend systematically on the decision problem itself. This suggests that it is the interaction between the limited abilities of the agent and the precise structure of the problem that determines the extent of bounded rationality and, hence, the resulting choices. Hence, our understanding of the effects of bounded rationality on decision-making remains incomplete until we can relate the heterogeneity of abilities to the heterogeneity of decision problems. Ultimately, the level of bounded rationality should itself be determined within the model, explaining (i) why a given agent departs from rational decision-making in some problems but not in others, and (ii) why, in a given decision problem, different agents exhibit different degrees of bounded rationality. Ideally, the variation should be connected to observable features, yielding concrete implications which could, in principle, be tested empirically.

My work takes a step towards these desiderata for general sequential decision problems of finite length. The decision maker is not required to know the entire structure of the problem when making choices. She can think ahead, through search that may be costly, to reveal more details of the problem. The novelty of my approach lies in constructing a theoretical framework for the complexity costs of search to be inferred from revealed preferences through observed choices. Specifically, I consider an agent who may have limited foresight when making a choice. I model

[1] The extant literature on violations of rational decision-making is vast. For a representative (yet, by no means comprehensive) review, see Benartzi and Thaler [4], Carbone and Hey [10], Charness and Rabin [11], Forsythe *et al.* [21], Henrich *et al.* [35], Hey [36], Lee and Hey [37], Kahneman and Tversky [43]; Loewenstein and Elster [49], and Thaler [78], [79].

limited foresight as a restriction on the forward depth of search in the decision-tree that depicts the problem at hand. For a given depth, the collection of all chains of information sets with length equal or smaller to the depth defines the horizon of search. For a given sequential decision problem of finite length, the decision-tree that depicts its extensive form defines a finite sequence of nested horizons of search. In my analysis, searching through a given decision-tree for the optimal alternative proceeds to encompass new information sets in the same way a filtration evolves to include new events.

To abstract from sources of departures from rational decision-making not related to costly search (such as errors in choice, misperceptions of payoffs, behavioral or psychological biases in preferences or information-processing, complete misunderstanding of the problem, etc.), the agent is taken to perceive correctly the payoffs at terminal contingencies that fall within her horizon of search and to be unrestricted in her ability to optimize across these payoffs. Beyond the horizon of search, however, the structure of the continuation problem or the possible terminal payoffs from different plans of action may not be known exactly. The agent is assumed to know only the worst possible terminal payoff, that may realize beyond the current horizon of search, from each plan of action available within the horizon. Her belief about the expected continuation payoff beyond the horizon, from a plan of action available within the horizon, is measured relative to the corresponding worst continuation payoff. The ratio of the expected to the worst continuation payoff is a function of the foresight of the horizon and the plan of action under consideration: the $\alpha_h(\cdot)$ function, defined formally in Section 2.2.

The suggested behavior under limited foresight amounts to a two-stage optimization. In the first stage, optimal plans of action are chosen for each value of foresight to which search can extend on the decision-tree. In the second stage, the agent chooses the optimal level of foresight. This process is, in principle, indistinguishable from rational decision-making. It can lead to quite different choices, however, if the cognitive, computational, or information-processing limitations that necessitate search also render it costly. Departures from rationality obtain when these costs constrain search to a horizon of inadequate foresight. By constructing the $\alpha_h(\cdot)$ function appropriately, any reasonable observed choice in a given decision problem corresponds to an optimal plan of action for the agent under some value for her foresight.[2] Observed choices that do not agree with the predictions of the rational paradigm will be optimal under a foresight which is not sufficient to account correctly for the entire set of future contingencies that may occur in the problem.

Under this framework, bounded rationality is modeled as optimization under informational constraints due to limited search. More importantly, observed boundedly-rational choices can be mapped to optimal choices under a particular foresight of search. The framework provides a platform, therefore, for assigning welfare values to observed choices, with respect to some underlying utility representation of preferences over terminal outcomes in the decision-tree under consideration.

[2]For what constitutes a "reasonable" choice in my analysis, see Definition 2 and the relevant discussion in Section 2.3.

By comparing these values to that corresponding to the choice prescribed by rational decision-making, we obtain the implied complexity costs of bounded rationality. These costs measure the welfare loss, relative to rational choices, incurred by the given decision maker.

In my analysis, the extent of bounded rationality is not assumed exogenously, but emerges endogenously: as problems become simpler or as the benefits of deeper search become larger relative to its costs, the choices more closely resemble those of a rational decision maker. In fact, for a given agent, if the problem is sufficiently simple or the benefits of search outweigh its costs, her choices become indistinguishable from those under rationality. For a fixed problem, the costs will vary across agents, explaining why the disparity between observed choices and those prescribed under rationality differs across decision makers.

Under certain assumptions, the implied search costs of bounded rationality can be calibrated from observed choices. Specifically, if the structure of the decision problem under study allows the foresight of search to be revealed from observed choices, the costs of search can be calibrated from observed responses when the problem is presented across a population of decision makers. More precisely, the marginal costs of searching deeper into the problem can be partially identified: the marginal cost of extending the current foresight of search by one step is bounded below by the marginal benefit that would be accrued from doing so, whereas the savings in search costs from reducing the current foresight by one step is bounded above by the marginal benefit that would be lost from doing so (see Proposition 2 and Corollary 1). By varying the marginal benefits of further search along the decision tree, these bounds can be sharpened until an estimated value for the corresponding marginal search costs is identified. Under a particular functional specification for the search costs, searching deeper into the problem terminates as soon as the marginal cost of further search begins to outweigh its marginal benefit (see Section 3.3).

Calibration of the implied costs of search suggests testable predictions and implications of the model. For a given decision problem, an estimate of the average marginal complexity costs across a population of decision makers allows for predictions. We can predict how changes in the terminal payoffs affect how far ahead into the problem an individual thinks and, consequently, what her responses are. More importantly perhaps, the model implies an increasing relation between the benefits of deeper search and the depth of search. For any given pair of a decision maker and a decision problem and no matter how large the implied costs of current search, the minimum payoff to deeper search required to induce searching further into the problem is finite. That is, there is a finite critical level for the benefits of further search which would necessarily induce searching further into the problem. This result offers a powerful test of my approach, rendering it falsifiable. Moreover, it allows us to transport the search costs calibrated in one decision-making situation to others of similar complexity and to make predictions regarding the depth of search of the decision maker under study.

Through three applications, I show that this model of endogenous bounded rationality is consistent with violations of timing independence in temporal framing problems, dynamic inconsistency

and diversification bias in sequential versus simultaneous choice problems, and with plausible but contrasting risk attitudes across small- and large-stakes gambles. Each case corresponds to a well-known paradox in decision-making. The three paradoxes are seemingly unrelated and obtain for obviously different reasons. Accordingly, a variety of models and arguments have been suggested in the literature to explain them. Yet, the complexity-costs approach provides a common framework for depicting the underlying limitations that force choices to depart from those prescribed by rationality. In each case, paradoxical behavior results when the costs of search are large relative to its benefits, preventing a boundedly-rational agent from considering the problem in its entirety.

Turning to a more specific description of the model, I consider sequential decision problems of finite length with uncertainty, in their extensive form. The payoffs to terminal nodes and the decision maker's preferences over these payoffs are given (thus, we know and agree upon what her choices in the given problem should be, if she were fully-rational). For simplicity, the agent is presumed to have expected utility preferences over the terminal payoffs.[3] To illustrate, consider the decision problem depicted in Figure 1. Decision and chance nodes are depicted by squares and circles, respectively. The agent's information set $i.3$ is a non-singleton (depicted by the shaded oval). The terminal nodes are indexed $\{\mathrm{A}, \ldots, \mathrm{K}\}$ while the values in brackets represent the corresponding terminal utility payoffs. The decision maker must choose a plan of action, at her initial node $i.1$, given a specification f of distributions at each chance node. For example, suppose f assigns probability $1/2$ to each branch at $n.3$ and $n.4$, and probability 1 to the branches leading to the terminal nodes B and K at $n.1$ and $n.2$, respectively. Under unbounded rationality and, thus, perfect foresight, the agent's plan of action would clearly choose the middle branch at $i.1$, the upper branch at the information set $i.3$, and the branches leading to the terminal nodes I and D at $i.2$ and $i.4$, respectively. Her expected utility payoff would be $45/4$.

In my approach, the search process truncates the decision-tree into horizons, with a horizon of foresight t consisting of all chains of information sets with length at most t. The decision-tree of Figure 2 defines a finite increasing sequence of values for the foresight of search, $t \in \{1, 2, 3, 4\}$, and a corresponding increasing sequence of nested horizons (depicted, in the figure, by combining adjacent hatched regions from left to right). For instance, the horizon of foresight 1 starting at $i.1$, denoted $i.1(1)$, is given by the collection $\{n.1, n.2, n.3\}$, while the horizon of foresight 2, denoted $i.1(2)$, by $i.1(1) \cup \{i.2, i.3\}$. In this example, having foresight $t = 4$ suffices for rational decision-making. For the horizon $i.1(1)$, the agent knows her utility payoffs at the terminal nodes A,B, and K, as they directly succeed nodes within the horizon. Thus, if Nature's play f is as above, the agent can correctly assess the payoffs to choosing the upper or the lower branch at $i.1$; she knows that they are 5 and 1, respectively. In contrast, she cannot necessarily assess correctly her payoff from choosing the middle branch at $i.1$ since she does not know with certainty the exact terminal consequences of her possible future actions. She assigns, therefore, an expected utility payoff according to some subjective probability distribution over the terminal nodes of the continuation problem (i.e. over

[3]Non-expected utility preferences can also easily be accommodated (see Section 3.2).

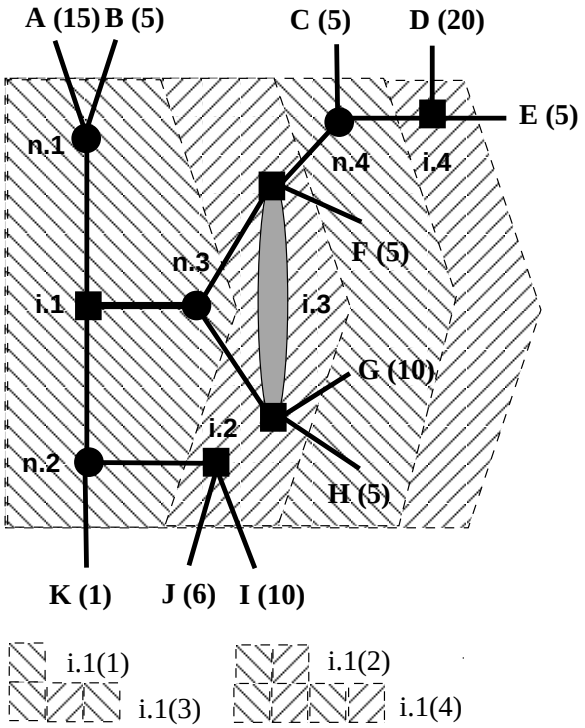


Figure 1: Foresight and Horizons of Search

the utility payoffs at C,D, . . . , H). A horizon of longer foresight, such as $i.1(2)$, gives the agent more precise information about some of the consequences of her choices (for example, under $i.1(2)$, she will know her expected payoff from choosing the lower branch at $i.3$), but could also be more costly to consider.

Suppose that an agent chooses the upper branch at $i.1$. Such a choice can be rationalized under bounded rationality and limited foresight. The terminal nodes A, B, and K fall within the limits of the horizon $i.1(1)$; even under the most myopic foresight, $t = 1$, the agent can correctly account for her expected payoff from choosing the upper and lower branches at $i.1$. With respect to the middle branch, however, she can only have some belief about her continuation payoff. Her observed choice at $i.1$ (i) reveals that she believes her expected continuation payoff to be less than 5 and (ii) implies that the increase in her complexity costs of extending her foresight from $t = 1$ to $t = 4$ is equal to $25/4$, the corresponding welfare loss from not doing so.

Consider now a decision maker who chooses the middle branch at $i.1$ but plans to take the action leading to the terminal node D at $i.4$ with probability less than 1. Her strategic behavior is closer to rational decision-making than that of the agent in the preceding paragraph. The increase in complexity costs, when extending her foresight from $t = 1$ to $t = 2$ at $i.1$, must be lower than that for the previous agent; it can even be zero now. Still, the new agent fails at $i.1$ to foresee that the terminal node D must be chosen with probability 1 at $i.4$. Thus, she must face non-zero

additional complexity costs in order to extend her foresight from $t = 2$ to $t = 4$. Yet, since she intends to choose the upper branch at $i.4$ with non-zero probability, she is planning to choose the upper branch at $i.3$. Consequently, her belief about her continuation payoff from the middle branch at $i.1$ must be at least $15/2$. Therefore, the additional search costs to extend her foresight from $t = 2$ to $t = 4$ can be at most $15/4$, the implied welfare loss from not doing so.

My approach in modeling bounded rationality fits into a large literature in economics. In most of this literature, however, aspects of the limitations leading to bounded rationality are taken to be exogenously-fixed. MacLeod [51] develops a search cost model, drawing on work in computer and cognitive sciences, in which both the payoffs and the costs of search are taken together to form an exogenously-given net-value function while search proceeds according to a heuristic algorithm (see also Eboli [19]). Models of limited foresight in repeated games (Rubinstein [63]; Jehiel [39], [39], [41],[42]) take the foresight to be exogenously given, implying that the complexity of the stage-game matters equally throughout the repeated game. When the set of contingencies that define the stage-game is not a singleton, however, the number of contingencies within a horizon increases in a combinatorial manner with the number of periods. In contrast, the notion of complexity introduced here incorporates both the number of periods and the number of contingencies within a given period. More importantly, the foresight of search is determined endogenously, which is crucial for explaining behavior in the applications I examine. In the work on repeated games played by finite automata (Neyman [52], [53]; Rubinstein [63], [63]; Abreu and Rubinstein [1], Piccione and Rubinstein [56], Kalai and Stanford [45], Sabourian [65]), the number of states of a machine is often used to measure complexity, leading to a decision-making framework that can be incorporated within my approach (see Section 3.5.3). Another strand of finite automata research focuses on the complexity of response rules (Eliaz [20], Chatterjee and Sabourian [12], Gale and Sabourian [26]).[4] While this notion of complexity is more comprehensive than the one presented here, it is unclear how it can be applied beyond repeated games, sequential matching and bargaining.[5]

The paper is organized as follows. Section 2 presents the general model and the main results. Section 3 develops some further results, including an endogenous stopping rule for search, and discusses the main implications of the analysis as well as some of the related literature in more detail. The applications are examined in Section 4 and Section 5 concludes. Finally, Appendix A offers an index for the notation introduced in Section 2 while Appendix B extends the analysis of Section 4.3 to a more general setting.

[4]Response complexity compares two response rules that are otherwise identical but for their prescriptions on some partial history of play. If, for this partial history, one assigns the same action irrespective of the current state of the machine while the other has the action dependent upon the state, the former is regarded as corresponding to lower complexity.

[5]With respect to machine complexity, response complexity comes closer to a complete ordering: two machines can have the same number of states but differ in response complexity.

2 Modeling Bounded Rationality

A decision problem can be viewed as a game between the agent, who is strategic, and Nature. As a game, it is defined by the set of observable events, the sets of possible plans of action for the agent and Nature, and the agent's utility payoff function over the terminal contingencies. It can be depicted by a directed graph $\Gamma = \langle \mathcal{D}, \mathcal{A} \rangle$ with $\mathcal{D}$ and $\mathcal{A}$ being the set of nodes and arcs, respectively. The elements of $\mathcal{A}$ correspond injectively to pairs of nodes (d, d') such that $d, d' \in \mathcal{D}$ with $d \neq d'$ and d' an immediate successor of d on the graph. Letting Ω denote the set of terminal nodes and $\mathcal{H}$ be the collection of the sets of nodes at which the same player chooses from the same set of actions, $\mathcal{H}$ forms a partition of $\mathcal{D} \setminus \Omega$ into information sets. H will denote the collection of the information sets where the agent acts. $\mathcal{H} \setminus H$ is the collection of Nature's chance nodes. Let also $D(h)$ denote the collection of nodes comprising the information set $h \in \mathcal{H}$. I will be referring to h as an *immediate successor of* some $d \in \mathcal{D}$, denoted by $d \rhd h$, if there exists $d' \in D(h)$ that immediately-succeeds d. Similarly, $h' \in \mathcal{H}$ with $h' \neq h$ is an *immediate successor* of h, denoted by $h \rhd h'$, if there exists $d \in D(h)$ and $d' \in D(h')$ with $d \rhd d'$.

2.1 Limited Foresight

When decision-making is constrained by limited foresight, the agent chooses optimally, from her set of strategic alternatives against Nature's play, taking into account, however, only a subset of the set of possible future contingencies. Specifically, she considers only some part of the decision-tree that lies ahead of her current decision node. Since a pair of plans of action for the agent and Nature defines a sequence of immediately-succeeding information sets that may occur with positive probability, I will restrict attention to parts of the decision-tree consisting of chains of arc-connected information sets. In what follows, I introduce the necessary concepts and notation to describe the precise structure of these parts and, thus, how search is directed through the graph.

For a positive integer t, a **path from the node d to the information set h' of length t** is a sequence of arc-connected information sets $\{h_\tau\}_{\tau=1}^{t} \subseteq \mathcal{H}$ such that $h_t = h'$, $d \rhd h_1$, and $h_k \rhd h_{k+1}$ for $k \in \{1, ..., t-1\}$. It will be denoted by $d \rhd \{h_\tau\}_{\tau=1}^{t}$. The union of all paths of length at most t that start from some $d \in D(h)$

$$h(t) = \bigcup_{d \in D(h)} \bigcup_{k \in \{1,...,t\}} \left\{ d \rhd \{h_\tau\}_{\tau=1}^{k} \right\}$$

defines the **horizon of foresight t starting at the information set** h. This is the collection of all chains of information sets starting at h whose length is at most t. The **continuation problem starting at** h, denoted by Γ_h, is the union of all horizons on the graph starting at h.[6] Observe

[6] This is the standard definition of a continuation problem which is to be distinguished from a subgame: a subgame is a continuation problem but not vice versa. A non-terminal node $d \in \mathcal{D} \setminus \Omega$ can be the root of a subgame Γ_d if and only if either $h \cap \Gamma_d = \varnothing$ or $h \in \Gamma_d$, $\forall h \in \mathcal{H}$. Any player who moves at d or at any information set of a subgame must know that d has occurred, which requires that only singleton information sets be starting points for subgames.

that the horizons can be ordered by proper set-inclusion and since

$$t' > t \quad \Longleftrightarrow \quad h(t) \subset h\left(t'\right) \qquad \forall t, t' \in \mathbb{N} \setminus \{0\}$$

this ordering can be equivalently represented by the standard ordering on the foresight length. Notice also that, in finite trees, $\Gamma_h = h\left(T_h\right)$ for some positive integer T_h.

Let Ω_h be the set of terminal nodes that may occur conditional on play having reached some $d \in D\left(h\right)$. If $h\left(t\right) \subset \Gamma_h$, then $h\left(t\right)$ defines a partition of Ω_h into the collection of terminal nodes that are immediate successors of information sets within the horizon and its complement. The collection is given by

$$\Omega_h\left(t\right) = \left\{y \in \Omega_h : \ \exists h' \in h\left(t\right) \wedge d' \in D\left(h'\right) \wedge d' \rhd y\right\}$$

A terminal node $y \in \Omega_h\left(t\right)$ will be called **observable within** $h\left(t\right)$.

The set of available actions at each of the nodes in $D\left(h\right)$ will be denoted by $A\left(h\right)$. A pure strategy is a mapping $s : H \to \times_{h \in H} A\left(h\right)$ assigning to each of the agent's information sets an available action. Let S be the set of such mappings in the game. A mixed strategy for the agent will be denoted by $\sigma \in \Sigma = \Delta\left(S\right)$. A deterministic play by Nature is a mapping $q : \mathcal{H} \setminus H \to \times_{d' \in \mathcal{H} \setminus H} A\left(d'\right)$ assigning to each chance node in $\mathcal{H} \setminus H$ a move by Nature. Q is the set of such mappings in the game and $f \in \Delta\left(Q\right)$ denotes a mixed strategy by Nature. Let also $p\left(d|\sigma, f\right)$ be the probability that the node d is reached under the scenario (σ, f). Given that play has reached d, the conditional probability that node d' will be reached under (σ, f) will be demoted by $p\left(d'|d, \sigma, f\right)$ with $p\left(d'|d, \sigma, f\right) = 0$ if d' cannot follow from d or $p\left(d|\sigma, f\right) = 0$,.

The set of terminal nodes that may occur under (σ, f), if play starts at a node of the information set h, is given by

$$\Omega_h\left(\sigma, f\right) = \{y \in \Omega_h : \ \exists d \in D\left(h\right) \wedge p\left(y|d, \sigma, f\right) > 0\}$$

Finally,

$$\Omega_h\left(t, \sigma, f\right) = \Omega_h\left(t\right) \cap \Omega_h\left(\sigma, f\right)$$

is the collection of terminal nodes which are observable within $h\left(t\right)$ and may occur under (σ, f) if play starts at a node of h.

2.1.1 An Example

Consider the decision-tree depicted in Figure 2 where decision and chance nodes are depicted by squares and circles, respectively, while the agent's non-singleton information set $i.3$ is depicted by the shaded oval. The terminal nodes are indexed $\{\mathrm{A}, \ldots, \mathrm{K}\}$ while the values in brackets represent the corresponding terminal utility payoffs. At the initial node $i.1$, the decision maker must choose a plan of action, given a specification f of distributions at each chance node. For

example, suppose that f assigns probability $1/2$ to each branch at $n.3$ and $n.4$, and probability 1 to the branches leading to the terminal nodes B and K at $n.1$ and $n.2$, respectively. Under unbounded rationality and, thus, perfect foresight, the agent can account for every possible contingency within the continuation problem (the entire tree here). Assuming for simplicity that she has expected utility preferences over final outcomes, her plan of action would choose the middle branch at $i.1$, the upper branch at the information set $i.3$, and the branches leading to the terminal nodes I and D at $i.2$ and $i.4$, respectively.[7]

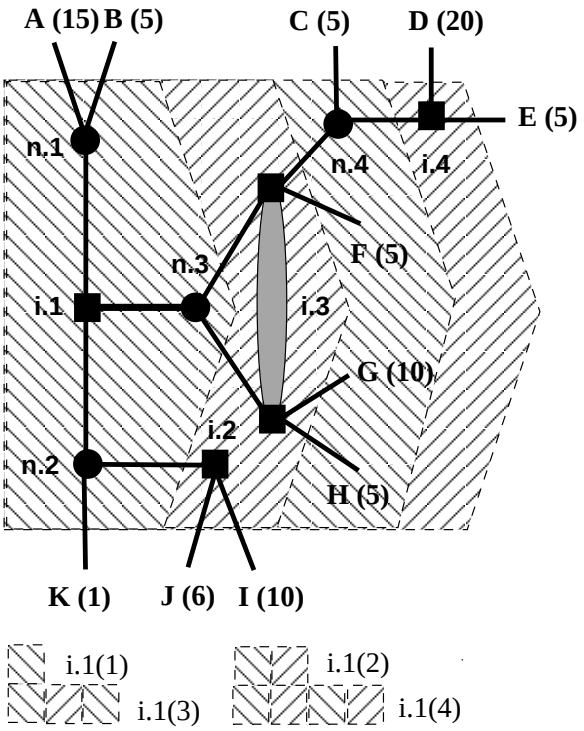


Figure 2: Foresight and Horizons of Search

Under bounded rationality, the search process is structured in my approach such that it truncates the decision-tree into horizons, with a horizon of foresight t consisting of all chains of information sets with length at most t. The decision-tree of Figure 2 defines a finite increasing sequence of values for the foresight of search, $t \in \{1, 2, 3, 4\}$, and a corresponding increasing sequence of nested horizons (depicted, in the figure, by combining adjacent hatched regions from left to right).[8] For $t = 1$, the search horizon is given by $i.1\,(1) = \{n.1, n.2, n.3\}$. For $t = 2$, there are two chains of information sets of that length, $i.1 \rhd \{n.2, i.2\}$ and $i.1 \rhd \{n.3, i.3\}$. The collection $\{n.1, n.2, n.3, i.2, i.3\}$ defines the horizon $i.1\,(2)$. Since only one chain of length $t = 3$ starts at $i.1$, $i.1 \rhd \{n.3, i.3, n.4\}$,

[7]The analysis can accommodate also non-expected utility preferences (see Section 3.2).

[8]In this framework, with respect to the length of its foresight, the search for the optimal alternative proceeds through the decision-tree so as to encompass new information sets in exactly the same way as a filtration evolves, with respect to time, to include new events.

the horizon $i.1\,(3)$ is given by $\{n.3, i.3, n.4\} \cup i.1\,(2)$. Finally, $\{n.3, i.3, n.4, i.4\} \cup i.1\,(3)$ defines $i.1\,(4)$ (giving the entire tree). The set of terminal nodes that can be reached from $i.1$ is, of course, $\Omega_{i.1} = \cup_{j=\mathrm{A}}^{\mathrm{K}} \{j\}$. The sets, though, that are observable within the horizons of search are given, respectively, by $\Omega_{i.1}\,(1) = \{\mathrm{A,B,K}\}$, $\Omega_{i.1}\,(2) = \Omega_{i.1}\,(1) \cup \{\mathrm{F}\ ,\mathrm{G,H,I,J}\}$, $\Omega_{i.1}\,(3) = \Omega_{i.1}\,(2) \cup \{C\}$, and $\Omega_{i.1}\,(4) = \Omega_{i.1}$.

Let σ prescribe positive probabilities for each of the available actions at $i.1$, $i.2$, and $i.4$, but probability 1 to Down at $i.3$. The sets of terminal nodes that may occur under the scenario (σ, f) and are observable within the horizons of search are given by $\Omega_{i.1}\,(1, \sigma, f) = \{\mathrm{B,K}\}$ and $\Omega_{i.1}\,(t, \sigma, f) = \{\mathrm{F,H,B,K}\} = \Omega_{i.1}\,(\sigma, f)$ for $t \geq 2$. If $s \in \mathrm{support}\{\sigma\}$ is the pure strategy that plays the middle branch at $i.1$, then $\Omega_{i.1}\,(1, s, f) = \varnothing$ while $\Omega_{i.1}\,(t, s, f) = \{\mathrm{F,H}\} = \Omega_{i.1}\,(s, f)$ for $t \geq 2$.

Similarly, for the decision to be made at the information set $i.3$, we have $\Omega_{i.3}\,(1) = \{\mathrm{C,F,G,H}\}$, $\Omega_{i.3} = \{\mathrm{C,D,E,F,G,H}\} = \Omega_{i.3}\,(2)$, and $\Omega_{i.3}\,(1, \sigma, f) = \{\mathrm{F,H}\} = \Omega_{i.3}\,(\sigma, f)$.

2.2 Behavior under Limited Foresight

Starting at $h \in H$, the agent may consider a horizon $h\,(t) \subseteq \Gamma_h$ instead of the entire continuation problem. Given its limited foresight, her decision-making is sophisticated: with respect to the terminal nodes that are observable under her current foresight, she can correctly optimize according to her underlying preferences over terminal consequences. Regarding outcomes beyond her horizon, she is not able to do so, even though she is aware that there may exist contingencies which could occur, under Nature's $f \in \Delta\,(Q)$ and her choice $\sigma \in \Sigma$, but are currently unforeseen within $h\,(t)$. I will assume that the agent evaluates the set of possible but currently unforeseen terminal contingencies according to the expectation operator $\mathbb{E}_\mu$, for some probability distribution μ on this set.

Formally, let $u : \Omega \to \mathbb{R}_{++}$ represent the agent's preferences over terminal outcomes.[9] Abusing notation slightly, for any set Y of terminal nodes, let

$$u\,(Y) = \sum_{y \in Y} u\,(y) \qquad \text{and} \qquad u\,(\varnothing) = 0$$

Given a pure strategy profile $(s, q) \in S \times Q$, $u\,(\Omega_h\,(s, q))$ is the agent's total utility payoff from the continuation problem Γ_h. Preferences over lotteries on terminal consequences admit an expected utility representation. The agent's expected utility payoff under $(\sigma, f) \in \Sigma \times \Delta\,(Q)$, given that play reached h, is given, therefore, by

$$u_h\,(\sigma, f) = \sum_{s \in S} \sum_{q \in Q} \sigma\,(s)\, f\,(q)\, u\,(\Omega_h\,(s, q))$$

[9] Although at no loss of generality, the non-negativity of the terminal utility payoffs is essential for determining precisely what unbounded rationality corresponds to in my approach (Lemma 1). I make the stronger assumption of strict positivity so as to represent some of the defining features of my model in terms of percentage changes in the expected continuation utility payoffs across immediately-succeeding horizons (Proposition 2).

while

$$R_h(f) = \arg\max_{\sigma \in \Sigma} u_h(\sigma, f)$$

is the set of best-responses under rational decision-making (unlimited foresight).

Under limited foresight, the notation must be enriched to distinguish the payoffs associated with terminal nodes that are observable within the horizon from those corresponding to terminal nodes beyond. Since the agent perceives correctly the former collection, the payoff from (s,q) at terminal nodes that are observable under foresight t is given by $u(\Omega_h(t,s,q))$. Within the unobservable set $\Omega_h \setminus \Omega_h(t)$, however, the terminal contingencies can be reached via any strategy profile $(\sigma', f') \in \Sigma \times \Delta(Q)$ that results in the same play as (s,q) within $h(t)$. That is, (σ', f') is another possible continuation scenario from the perspective of the horizon $h(t)$. By Kuhn's Theorem, we can define the following concept:

Definition 1 *Let $h \in H$ and $h(t) \subseteq \Gamma_h$. Two strategy profiles $(\sigma, f), (\sigma', f') \in \Sigma \times \Delta(Q)$ are $h(t)$-**equivalent**, denoted by $(\sigma, f) \underset{h(t)}{\simeq} (\sigma', f')$, if both prescribe the same probability distribution at each node in $D\left(\widetilde{h}\right)$, for any information set $\widetilde{h} \in h(t)$. That is,*

$$(\sigma, f) \underset{h(t)}{\simeq} \left(\sigma', f'\right) \quad \textit{iff} \quad (\sigma, f)|_{h(t)} \equiv \left(\sigma', f'\right)|_{h(t)}$$

Here, $\left(\widetilde{\sigma}, \widetilde{f}\right)|_{h(t)}$ denotes the restriction to $h(t)$ of the unique behavior-strategy profile that corresponds to $\left(\widetilde{\sigma}, \widetilde{f}\right)$.

Regarding outcomes beyond the truncated tree, the set of terminal nodes that may follow from any $d \in D(h)$ when (s,f) is played within $h(t)$ is given by[10]

$$\overline{\Omega}_h(t,s,f) = \bigcup_{(s',f') \in S \times \Delta(Q) : (s',f') \underset{h(t)}{\simeq} (s,f)} \Omega_h\left(s',f'\right) \setminus \Omega_h(t,s,f)$$

The set of the corresponding terminal utility payoffs will be denoted by

$$\mathcal{U}_h(t,s,f) = \left\{u(y) : y \in \overline{\Omega}_h(t,s,f)\right\} \subset \mathbb{R}_{++}$$

The agent evaluates the set $\overline{\Omega}_h(t|s,f)$ of possible but currently unforeseen terminal contingencies according to the expectation operator $\mathbb{E}_\mu[\mathcal{U}_h(t,s,f)]$ for some probability measure μ on $\overline{\Omega}_h(t|s,f)$. Observe that, letting $\inf \mathcal{U}_h(t,s,f)$ and $\sup \mathcal{U}_h(t,s,f)$ denote the minimum and maximum of this set, respectively, this criterion can be represented as

$$\mathbb{E}_\mu[\mathcal{U}_h(t,s,f)] = \alpha_h(t,s,f) \inf \mathcal{U}_h(t,s,f) \tag{1}$$

for some $\alpha_h(t,s,f) \in \left[1, \frac{\sup \mathcal{U}_h(t,s,f)}{\inf \mathcal{U}_h(t,s,f)}\right]$. In Section 3, I discuss the robustness of my results with respect to using the this particular representation. The total expected utility payoff under horizon

[10] Observe that $\Omega_h(t,s',f') \equiv \Omega_h(t,s,f)$, for any $(s',f') \in S \times \Delta(Q) : (s',f') \underset{h(t)}{\simeq} (s,f)$.

$h(t)$, when (σ, f) is played within the horizon, is given by the mapping $U_h(\cdot) : \{1, ..., T_h\} \times \Sigma \times \Delta(Q) \to \mathbb{R}_{++}$ with

$$\begin{aligned} U_h(t, \sigma, f) &= u(\Omega_h(t, \sigma, f)) + \sum_{s \in S} \sigma(s) \mathbb{E}_\mu [\mathcal{U}_h(t, s, f)] \\ &= u(\Omega_h(t, \sigma, f)) + \sum_{s \in S} \sigma(s) \alpha_h(t, s, f) \inf \mathcal{U}_h(t, s, f) \end{aligned} \tag{2}$$

The second term above corresponds to the agent's expected utility payoff from the possible but currently unforeseen future terminal contingencies when (σ, f) is played within $h(t)$. It will be called, henceforth, the **continuation value** of the horizon $h(t)$ under (σ, f).

A binary relation can now be defined on $\{1, ..., T_h\}$ by

$$t \succsim_h t' \quad \text{iff} \quad \max_{\sigma \in \Sigma} U_h(t, \sigma, f) \geq \max_{\sigma \in \Sigma} U_h(t', \sigma, f)$$

This is the agent's preference relation over the horizons in Γ_h when Nature's strategy is f. By construction, it admits the utility representation

$$V_h(t, f) = \max_{\sigma \in \Sigma} U_h(t, \sigma, f)$$

while the correspondence

$$R_h(t, f) = \arg\max_{\sigma \in \Sigma} U_h(t, \sigma, f)$$

gives the optimal choices from Σ against f under the horizon $h(t)$. The optimization problem $\max_{t \in \{1, ..., T_h\}} V_h(t|f)$ determines optimal horizon - strategy set pairs $\{h(t^*), R_h(t^*, f)\} \in \Gamma_h \times \Sigma$.

2.3 Bounded Rationality

The behavior under limited foresight described above can be used to compare boundedly-rational with rational decision-making once some minimal structure is imposed on the underlying choice rules.

Definition 2 *A choice rule $\mathcal{C}_h(S, f) \subseteq \Sigma$, against $f \in \Delta(Q)$ at h, is **admissible** if there exists $t \in \{1, ..., T_h\}$ such that $\mathcal{C}_h(S, f) \equiv R_h(t, f)$.*

The criterion in Definition 2 admits really any reasonable choice. It refuses only choice rules which ignore strict dominance to an extent that it is impossible to draw a line between bounded rationality and irrationality. To see this, let $s^\star \in \mathcal{C}_h(S, f)$ for some inadmissible $\mathcal{C}_h(S, f)$. By definition, we cannot find any $t \in \{1, ..., T_h\}$ and an corresponding array $[\alpha_h(t, s, f)]_{s \in S}$ such that $U_h(t, s, f) \leq U_h(t, s^*, f)$ for all $s \in S$. Fixing $t \in \{1, ..., T_h\}$, this implies the existence of some $\widetilde{s} \in S$ such that $U_h(t, \widetilde{s}, f) > U_h(t, s^*, f)$ for any $\alpha_h(t, \widetilde{s}, f) \in \left[1, \frac{\sup \mathcal{U}_h(t, \widetilde{s}, f)}{\inf \mathcal{U}_h(t, \widetilde{s}, f)}\right]$ and any

$\alpha_h(t,s^*,f) \in \left[1, \frac{\sup \mathcal{U}_h(t,s^*,f)}{\inf \mathcal{U}_h(t,s^*,f)}\right]$. Taking $\alpha_h(t,\widetilde{s},f) = 1$ and $\alpha_h(t,s^*,f) = \frac{\sup \mathcal{U}_h(t,s^*,f)}{\inf \mathcal{U}_h(t,s^*,f)}$, by (1) and (2), we get

$$\begin{aligned}
& \min_{(s',f') \in S \times \Delta(Q):\ (s',f') \underset{h(t)}{\simeq} (\widetilde{s},f)} u_h\left(s',f'\right) \\
= \ & u\left(\Omega_h(t,\widetilde{s},f)\right) + \min_{(s',f') \in S \times \Delta(Q):\ (s',f') \underset{h(t)}{\simeq} (\widetilde{s},f)} \left\{u_h\left(s',f'\right) - u\left(\Omega_h(t,\widetilde{s},f)\right)\right\} \\
> \ & u\left(\Omega_h(t,s^*,f)\right) + \max_{(s',f') \in S \times \Delta(Q):\ (s',f') \underset{h(t)}{\simeq} (s^*,f)} \left\{u_h\left(s',f'\right) - u\left(\Omega_h(t,s^*,f)\right)\right\} \\
= \ & \max_{(s',f') \in S \times \Delta(Q):\ (s',f') \underset{h(t)}{\simeq} (s^*,f)} u_h\left(s',f'\right)
\end{aligned}$$

That is, an inadmissible choice rule opts for the plan of action s^* against f even though there exists a plan $\widetilde{s}$ under which the worst possible outcome, from the perspective of the horizon $h(t)$, is strictly better than the best possible outcome under s^*. Moreover, since the foresight t was taken arbitrarily, this applies even for every horizon in Γ_h.

The notion of admissibility of Definition 2 imposes some very weak restrictions on the extent to which the choice behavior under study departs from rationality. I will now demonstrate that, within the realm of admissibility, the behavior under limited foresight in 2 can be used as a platform onto which both rational as well as boundedly-rational decision-making can be represented. For this, I must make the following behavioral assumption.

A.1 $\forall h \in H$ *and* $\forall (s,f) \in S \times \Delta(Q)$ $\quad \alpha_h(t,s,f)$ *is non-decreasing on* $\{1,...,T_h\}$

Recall that the mapping $\alpha_h(t,s,f)$ depicts the agent's belief about the continuation value of the horizon $h(t)$ under the profile (s,f) as a combination between the best and worst utility payoff out of the set of currently unforeseen but possible future terminal contingencies when (s,f) is played within $h(t)$. The assumption amounts simply to requiring that the agent does not become more pessimistic about the continuation value of her horizon of search, with respect to a given profile (s,f), as her search deepens. This suffices for the agent to be exhibiting a weak preference for deepening her horizons of search.

Proposition 1 *Under Assumption A.1, the relation* $\succsim_h$ *exhibits a weak preference for larger horizons:*

$$t > t' \implies t \succsim_h t'$$

Notice that a weak preference for deeper search can be equivalently expressed as $\Delta_t V_h(t,f) \geq 0$ $\forall t \in \{1,...,T_h\}$. Therefore,

$$\max_{t \in \mathbf{T}_h} V_h(t,f) = V_h(T_h,f) = \max_{\sigma \in \Sigma} u_h(\sigma,f)$$

At any $h \in H$, an agent with preferences $\succsim_h$ on $\{1,...,T_h\}$ will choose the maximal horizon and, thus, the same set of best responses from Σ against $f \in \Delta(Q)$ as a rational decision maker.

Under assumption A.1, the suggested behavior under limited foresight is indistinguishable from rationality. It can be taken, therefore, as a benchmark: departures from rational decision-making can be equivalently depicted as departures from this behavior under limited foresight. Within the context of admissible choice rules, I will restrict attention to departures that result in a welfare loss relative to rational decision-making.

Definition 3 *An admissible choice rule $\mathcal{C}_h(S,f)$ is* ***boundedly-rational*** *if $\mathcal{C}_h(S,f) \equiv R_h(t^*,f)$ and $V_h(t^*,f) < V_h(T_h,f)$, for some $t^* \in \{1,...,T_h-1\}$.*

Since it is admissible, a boundedly-rational choice rule should give a best-response set that is optimal for the optimization over $\mathbf{T}_h$ of some appropriate transformation of the value function $V_h(t,f)$. This transformation could represent a variety of psychological, behavioral, cognitive or ability factors that force departures from rationality. However, boundedly-rational choices ought to be admissible with respect to every foresight $t \in \{1,...,T_h\}$. Moreover, preferences over lotteries on the set of terminal consequences should admit an expected utility representation. Clearly, we must restrict attention to strictly-increasing, affine transformations of $U_h(t,\cdot,f)$. Let, therefore, $\widetilde{U}_h : \{1,...,T_h\} \times \Sigma \times \Delta(Q) \to \mathbb{R}_{++}$ be defined by

$$\widetilde{U}_h(t,\sigma,f) = \gamma_h(t,f)\, U_h(t,\sigma,f) + \gamma_h^0(t,f)$$

for some functions $\gamma_h : \mathbf{T}_h \times \Delta(Q) \to \mathbb{R}_{++}$ and $\gamma_{0h} : \mathbf{T}_h \times \Delta(Q) \to \mathbb{R}$. The optimization problem now becomes

$$\max_{t \in \{1,...,T_h\}} \gamma_h(t,f)\, V_h(t,f) + \gamma_h^0(t,f) \tag{3}$$

Let $t_h(f)$ denote the smallest optimal foresight in (3). The corresponding set of maximizers from Σ will be called the boundedly-rational best-response set, denoted by $BR_h(f)$.

The formulation in (3) embeds rational decision-making as a special case. More importantly, it does so in a way that establishes the functions γ_h and γ_h^0 as *indicators* of bounded rationality. For the agent to be boundedly-rational against f at h, at least one of two functions must not be constant on the set $\{t_h(f),...,T_h\}$. Put differently, rational decision-making against *any* $f \in \Delta(Q)$ at h is consistent only with both functions being constant on $\left\{\min_{f\in\Delta(Q)} t_h(f),...,T_h\right\} \times \Delta(Q)$.

Proposition 2 *Let $f \in \Delta(Q)$. The following are equivalent:*

(i) The optimization in (3) represents $\succsim_h$ on the set $\{t(f),...,T_h\}$

(ii) The functions γ_h and γ_h^0 are constant on the set $\{t(f),...,T_h\}$.

For the remaining of the paper, I will normalize γ_h^0 to be the zero function on $\{1,...,T_h\} \times \Delta(Q)$. This allows for an empirically testable and robust relation between the terminal utility payoffs of the decision problem at hand and the agent's optimal horizon. To see this, observe first that

$$\begin{aligned} \Delta_t\left[\gamma_h(t,f)\, V_h(t,f)\right] &= \gamma_h(t+1,f)\, V_h(t+1,f) - \gamma_h(t,f)\, V_h(t,f) \\ &= \gamma_h(t+1,f)\, \Delta_t V_h(t,f) + V_h(t,f)\, \Delta_t \gamma_h(t,f) \end{aligned} \tag{4}$$

while, at the optimal foresight, we must have $\Delta_t \left[\gamma_h \left(t\left(f\right), f\right) V_h \left(t\left(f\right), f\right)\right] \leq 0$. Since $\Delta_t V_h \left(t, f\right) \geq 0$ on $\{1, ..., T_h\}$ (recall Proposition 1),

Remark 1 *For the representation in (3), let γ_h^0 be the zero function on $\{1, ..., T_h\} \times \Delta\left(Q\right)$. Then, $\Delta_t \gamma_h \left(t\left(f\right), f\right) \leq 0$.*

Moreover, all utility payoffs of terminal outcomes are strictly positive (i.e. $V_h \left(t, f\right) > 0$ for all $\left(t, f\right) \in \mathbf{T}_h \times \Delta\left(Q\right), h \in H$) and so is the function γ_h. Clearly, we must have

$$\frac{\Delta_t V_h \left(t\left(f\right), f\right)}{V_h \left(t\left(f\right), f\right)} \leq -\frac{\Delta_t \gamma_h \left(t\left(f\right), f\right)}{\gamma_h \left(t\left(f\right) + 1, f\right)}$$

Since the transformation function γ_h does not vary with the terminal utility payoffs, the validity of the above inequality is completely determined by the left-hand side quantity. If this increases sufficiently, other things remaining unchanged, the inequality will no longer be valid.[11] Similarly, by the optimality of (f), we have $\Delta_t \left[\gamma_h \left(\left(f\right) - 1, f\right) V_h \left(\left(f\right) - 1, f\right)\right] \geq 0$ and, thus,

$$\frac{\Delta_t V_h \left(\left(f\right) - 1, f\right)}{V_h \left(\left(f\right) - 1, f\right)} \geq -\frac{\Delta_t \gamma_h \left(\left(f\right) - 1, f\right)}{\gamma_h \left(\left(f\right), f\right)}$$

If the quantity $\frac{\Delta_t V_h \left(\left(f\right) - 1, f\right)}{V_h \left(\left(f\right) - 1, f\right)}$ decreases sufficiently, ceteris paribus, the inequality will cease to hold.[12]

Remark 2 *In the representation in (3), let γ_h^0 be the zero function on $\{1, ..., T_h\} \times \Delta\left(Q\right)$. Suppose that the specification of the terminal utility payoffs in $\cup_{s \in S} \Omega_h \left(s, f\right)$ changes and denote by $t^{old}\left(f\right)$ and $t^{new}\left(f\right)$ the optimal horizons under the original and new specifications, respectively. Let also $V_h^{new}\left(t, f\right)$ denote the value function in (3) under the new specification. Then,*

1. *$t^{new}\left(f\right) = t^{old}\left(f\right)$ if the changes in the terminal utility payoffs are such that*

$$\frac{\Delta_t V_h^{new}\left(t^{old}\left(f\right), f\right)}{V_h^{new}\left(t^{old}\left(f\right), f\right)} \in \left[0, -\frac{\Delta_t \gamma_h \left(t^{old}\left(f\right), f\right)}{\gamma_h \left(t^{old}\left(f\right), f\right)}\right] \quad \textit{and}$$
$$\frac{\Delta_t V_h^{new}\left(t^{old}\left(f\right) - 1, f\right)}{V_h^{new}\left(t^{old}\left(f\right) - 1, f\right)} \in \left[-\frac{\Delta_t \gamma_h \left(t^{old}\left(f\right), f\right)}{\gamma_h \left(t^{old}\left(f\right), f\right)}, +\infty\right)$$

(2.i) $t^{new}\left(f\right) > t^{old}\left(f\right)$ if the changes in the terminal utility payoffs are such that

$$\frac{\Delta_t V_h^{new}\left(t^{old}\left(f\right), f\right)}{V_h^{new}\left(t^{old}\left(f\right), f\right)} \notin \left[0, -\frac{\Delta_t \gamma_h \left(t^{old}\left(f\right), f\right)}{\gamma_h \left(t^{old}\left(f\right), f\right)}\right] \quad \textit{and}$$
$$\frac{\Delta_t V_h^{new}\left(t^{old}\left(f\right) - 1, f\right)}{V_h^{new}\left(t^{old}\left(f\right) - 1, f\right)} \in \left[-\frac{\Delta_t \gamma_h \left(t^{old}\left(f\right), f\right)}{\gamma_h \left(t^{old}\left(f\right), f\right)}, +\infty\right)$$

[11] *Other things remaining unchanged* refers here to the quantities $\frac{\Delta_t V_h(t,f)}{V_h(t,f)}$, for $t \in \{1, ..., T_h\} \setminus \{t\left(f\right)\}$.

[12] *Ceteris paribus* refers to the quantities $\frac{\Delta_t V_h(t,f)}{V_h(t,f)}$, for $t \in \{1, ..., T_h\} \setminus \{t\left(f\right) - 1\}$.

(2.ii) $t^{new}(f) < t^{old}(f)$ *if the changes in the terminal utility payoffs are such that*

$$\frac{\Delta_t V_h^{new}\left(t^{old}(f), f\right)}{V_h^{new}\left(t^{old}(f), f\right)} \in \left[0, -\frac{\Delta_t \gamma_h\left(t^{old}(f), f\right)}{\gamma_h\left(t^{old}(f), f\right)}\right] \quad \textit{and}$$
$$\frac{\Delta_t V_h^{new}\left(t^{old}(f)-1, f\right)}{V_h^{new}\left(t^{old}(f)-1, f\right)} \notin \left[-\frac{\Delta_t \gamma_h\left(t^{old}(f), f\right)}{\gamma_h\left(t^{old}(f), f\right)}, +\infty\right)$$

The suggested normalization guarantees also that the optimal horizon against $f \in \Delta(Q)$ at $h \in H$ in (3) is not affected by affine transformations in the utility representation of the agent's preferences over terminal consequences.

Claim 1 *For the representation in (3), let* γ_h^0 *be the zero function on* $\{1, ..., T_h\} \times \Delta(Q)$ *and consider an affine transformation,* $\lambda u(\cdot)$ *with* $\lambda > 0$*, of the utility representation of preferences over the terminal consequences. Then,* $t^{old}(f) = t^{new}(f)$*.*

2.4 Complexity and its Costs

In general, different search horizons $h(t) \in \Gamma_h$ impose different degrees of informational limitations on the agent's decision-making. The particular search horizon the agent considers should, therefore, matters for her response against any $f \in \Delta(Q)$. Yet, the proof of Proposition 1 shows that this is not always the case since the preference for larger horizons is not necessarily strict. Depending on the structure of the problem and Nature's play, the marginal benefit from deeper search, $\Delta_t V_h(t, f)$, may be zero. In such a case, extending the search horizon from $h(t)$ to $h(t+1)$ does not affect the set of optimal responses against f. This corresponds to a non-complex decision-making situation.

Definition 4 *For any* $h \in H$ *and* $t \in \{1, ..., T_h - 1\}$*, we say that* $f \in \Delta(Q)$ *is* $h(t)$-***complex*** *if* $\Delta_t V_h(t, f) > 0$*.* f *is* h-***complex*** *if it is* $h()$*-complex for some* $t \in \{1, ..., T_h\}$*. The decision problem is* ***complex*** *at* h*, if there exists* $f \in \Delta(Q)$ *that is* h*-complex.*

Example 2.4: For the decision problem of Figure 2, suppose that Nature may move in either direction at $n.3$ and $n.4$ with equal probability, left at $n.1$ with any probability p, and deterministically down at $n.2$. At $i.1$, if the agent chooses the upper branch, she expects a utility payoff of $5 + 10p$. If she chooses the lower branch, her payoff will be 1. The middle branch corresponds to a continuation value of $5\alpha_{i.1}(1, m, f)$ for $t = 1$ (the worst case scenario is her moving Down at $i.3$), $\frac{10}{2} + \frac{5}{2}\alpha_{i.1}(2, m, f)$ for $t = 2$, $\frac{10}{2} + \frac{5}{4} + \frac{5}{4}\alpha_{i.1}(3, m, f)$ for $t = 3$ (for $t \in \{2, 3\}$, the worst case scenario entails moving Up at $i.3$ but Down at $i.4$), and $\frac{45}{4}$ for $t = 4$. Notice that $\alpha_{i.1}(1, m, f) \in \left[1, \frac{\frac{10}{2} + \frac{20}{2}}{5}\right]$ whereas $\alpha_{i.1}(2, m, f), \alpha_{i.1}(3, m, f) \in \left[1, \frac{20}{5}\right]$.

If $p \geq \frac{5}{8}$, the expected value of every horizon is $5 + 10p$; Nature's play is not complex at $i.1$. For $p < \frac{5}{8}$, the expected values are $5 \max\{1 + 2p, \alpha_{i.1}(1, m, f)\}$ for $t = 1$, $5 + 5 \max\left\{2p, \frac{1}{2}\alpha_{i.1}(2, m, f)\right\}$

for $t = 2$, $5 + 5\max\left\{2p, \frac{1}{4} + \frac{1}{4}\alpha_{i.1}(3, m, f)\right\}$ for $t = 3$, and $\frac{45}{4}$ for $t = 4$. Nature's play can now be complex at $i.1$ (for instance, let $p < \frac{1}{4}$). $\square$

Consider now the function $C_h : \{1, ..., T_h\} \to \mathbb{R}$ defined by

$$C_h(t, f) = [1 - \gamma_h(t, f)] V_h(t, f)$$

This function acquires some interesting properties when the notion of complexity defined above is linked with the representation in (3). Since $t_h(f)$ denotes an optimal horizon in (3), we have

$$\begin{aligned} V_h(t_h(f), f) - C_h(t_h(f), f) &\geq V_h(t, f) - C_h(t, f) \Longrightarrow \\ C_h(t, f) - C_h(t_h(f), f) &\geq V_h(t, f) - V_h(t_h(f), f) \\ &\geq 0 \quad \text{(by Proposition 1)} \quad \forall t \in \{t(f), ..., T_h\} \end{aligned}$$

If we normalize $\gamma_h(t_h(f), f) = 1$, then $C_h(t_h(f), f) = 0$. Since now $C_h(t, f) \geq 0$, for any $t \in \{t(t) T_h\}$ and $f \in \Delta(Q)$, the function C_h becomes a cost function on $\{t(f), ..., T_h\}$. If, moreover, we restrict g_h to take values in the set $(0, 1]$, C_h becomes a cost function on the entire set $\{1, ..., T_h\}$. In either case, it will be called, henceforth, the **complexity cost** function.

By the very definition of the complexity-cost function, when γ_h^0 is the zero function on $\{1, ..., T_h\}$, the decision-making under bounded rationality given in (3) admits equivalently the following representation

$$\max_{t \in \{1, ..., T_h\}} V_h(t, f) - C_h(t, f) \tag{5}$$

This representation illustrates the economic meaning of the complexity costs. Notice that, if $\Delta_t C_h(t, f) = 0$ for all $t \in \{1, ..., T_h\}$, the representation is equivalent to rational decision-making. The function $C_h(\cdot, f)$ depicts, therefore, the discrepancy between the payoff, $\gamma_h(t, f) V_h(t, f)$, that the boundedly-rational agent assigns to her best response, $R_h(t, f)$, against f when her horizon is limited to $h(t)$ and the actual payoff, $V_h(t, f)$, from $R_h(t, f)$ against f when the foresight is restricted to t. It is trivial to verify the following relations between the marginal benefits of extending search around the optimal foresight against f and the corresponding marginal complexity costs. In Section 3.4, I show that they lead to important testable implications of my model.

Remark 3 *For the representation in (3), let g_h^0 be the zero function on $\{1, ..., T_h\}$. For any $f \in \Delta(Q)$ and $h \in H$, we have*

$$\Delta_t C_h(t(f), f) \geq \Delta_t V_h(t(f), f) \quad \textit{and} \quad \Delta_t C_h(t(f) - 1, f) \leq \Delta_t V_h(t(f) - 1, f)$$

For the representation in (6), discrepancies between boundedly-rational and rational decision-making can exist only in complex situations. That is, the complexity costs do not affect the behavior under limited foresight in non-complex decision-making situations.

Claim 2 *For the representation in (3), let γ_h^0 be the zero function on $\{1, ..., T_h\}$. If $f \in \Delta(Q)$ is not $h(t)$-complex, then it cannot entail any marginal complexity costs on the horizon $h(t)$. That is,*

$$\Delta_t V_h(t, f) = 0 \Longrightarrow \Delta_t C_h(t, f) = 0$$

The complexity-cost function affects choices only in complex problems and only where the agent would strictly prefer to expand her horizon if she were not constrained by her bounded rationality. In this sense, the function depicts the presence of (and only of) bounded rationality. To see this, consider again Nature's mixed strategy in Example 2.4 with $p \geq \frac{5}{8}$. An agent who chooses the upper branch at $i.1$ may or may not be boundedly-rational. Such a stochastic play by Nature is not complex enough to reveal bounded rationality from observed choices in this problem. Accordingly, expanding the agent's foresight does not entail any marginal complexity costs.

Complexity has been defined here with respect to the structure of the decision problem. It depends, therefore, upon the particular conjecture f about Nature's play. This dependence is reflected in the function C_h allowing the model to exhibit an important endogenous form of bounded rationality. When playing a complex game against a strategic opponent, the perceived ability of the opponent seems to obviously matter when the agent decides how deeply into the game-tree to plan for. Even though Nature is not a strategic player, there is an analogue of this intuition in a decision-theoretic setting. Different conjectures about Nature's play correspond to different sets of unforeseen contingencies that may occur beyond the current horizon. Hence, the difficulty of a given decision problem varies with the set of Nature's stochastic plays that the agent considers. In this model, the conjectures about Nature's play matter not only through the expected terminal payoffs but also directly through the complexity-cost function. This may lead a boundedly-rational agent to choose different horizons in response to different conjectures.

Allowing the complexity costs function to vary on $\Delta(Q)$ offers also a built-in flexibility to differentiate between the length and breadth of a horizon, in terms of complexity. The cost of considering a horizon $h(t)$ against a strictly-mixed f may well be strictly higher than the corresponding cost against any pure strategy in the support of f. This could reflect the larger number of contingencies which may happen within the horizon against the mixture. While the foresight is t in either case, responding against f requires accounting for more contingencies. Finally, the model allows the function C_h to vary across the information sets $h \in H$ at which decisions are to be made against f. It can accommodate, therefore, decision-making situations where the complexity of the continuation problem varies at different stages of the problem.

3 Discussion and Related Literature

The preceding analysis views bounded rationality as resulting from costs implied by the complexity of the decision problem. The costs obtain in such a way that decision-making under bounded

rationality admits the following representation

$$\max_{t\in\mathbf{T}_h} V_h(t,f) - C_h(t,f) \tag{6}$$

In this section, I discuss some of the strengths and limitations of this representation and how it relates to some of the literature.

3.1 Strictly Dominated Strategies Matter

Decision-making under limited foresight is based upon comparing plans of action that lead to terminal outcomes occurring within the agent's current foresight against her beliefs about the continuation value of plans that do not. Such decision-making is affected by strictly dominated strategies, if the agent's foresight is not long enough to reveal the dominance.

In finite, non-repeated problems, any $(s,q) \in S \times Q$ will eventually reach a unique terminal node. That is, one of the two terms of the partition $\Omega_h(t,s,q) \cup \overline{\Omega_h}(t,s,q)$ is the empty set, for any $t \in \mathbf{T}_h, h \in H$. There exists, therefore, $t_h(s,q) \in \mathbf{T}_h$ such that $\widetilde{U}_h(t_h(s,q),s,q) = u(h,s,q)$. Suppose now that $\widetilde{s} \in S$ does strictly worse than s against q conditional on play having reached h: $u(h,s,q) > u(h,\widetilde{s},q)$. The presence of $\widetilde{s}$ does not matter for the objective in (3) as long as $t \geq t_h(s,q)$. For $t \geq \max_{q\in Q} t_h(s,q)$, this is true for any strictly dominated strategy $\widetilde{s}$. For finitely-repeated problems, this argument suggests that the presence of $\widetilde{s}$ does not matter for the optimization in (3), if there exists $s \in S$ that does strictly better than $\widetilde{s}$ against q in the stage game and the agent's foresight suffices to cover the stage game.

To illustrate, consider removing the option of moving Down at $i.3$ in Figure 2. Any strategy that uses this action with positive probability is strictly dominated by any other that assigns probability 1 to Up at $i.3$. The deletion will not affect the valuation $V_{i.1}(t,f)$, for $t \geq 2$, because the dominance is revealed within the horizon. This is not true, however, for $t = 1$. Consider the play f of Example 2.4. The continuation value of the middle branch at $i.1$ is now the one corresponding to the horizon $i.1\,(2)$ in that example, $\frac{10}{2} + \frac{5}{2}\alpha_{i.1}(2,m,f)$. For $\alpha_{i.1}(2,m,f) > 4p$, an agent with foresight $t = 1$ will choose Middle at $i.1$. This inequality can be satisfied even for $p \geq \frac{5}{8}$; that is, in cases where the agent would choose Up in Example 2.4. The removal of a strictly-dominated option at $i.3$ makes her here ex-ante, strictly better-off at $i.1$.

It is equally straightforward to construct a strictly dominated strategy whose introduction makes a boundedly-rational decision maker better off. Following Nature moving left at $n.1$ in Figure 2, suppose that the agent can now choose between two moves, say Left and Right, for a utility payoff of 15 and 1, respectively. The new problem is equivalent to the old for a rational agent. Yet, for Nature's f as above, the continuation payoff of choosing Up at $i.1$, for $t = 1$, becomes $p\alpha_{i.1}(1,u,f) + 5(1-p)$, where $\alpha_{i.1}(1,u,f) \in [1,15]$. For $\alpha_{i.1}(1,u,f) < 5$, an agent with foresight $t = 1$ will choose the middle branch at $i.1$ for any $p \in [0,1]$. Introducing a strictly dominated option at $n.1$ makes her ex-ante strictly better-off at $i.1$.

Notice that the dependence on strictly dominated strategies obtains only as long as the continuation value of the horizon depends on the length of foresight. The continuation value of any $(s,f) \in S \times \Delta(Q)$ does not vary with t if and only if $\alpha_h(t,s,f) = \frac{\sup \mathcal{U}_h(t,s,f)}{\inf \mathcal{U}_h(t,s,f)}$, for all $t \in \mathbf{T}_h$. If the agent correctly identifies the best possible future outcome of (s,f) under any horizon $h(t)$, her limited foresight poses no informational limitation for decision-making. She will always choose the same best-response set as if her foresight were unlimited. Planning for future contingencies that lie beyond even the shortest horizon, $h(1)$, is not necessary.

3.2 Robustness

The complexity cost approach is based on measuring the extent of bounded rationality via the implied welfare loss when observed choices differ from those prescribed under rational decision-making. This framework for modeling endogenous bounded rationality is robust to many different ways of evaluating a given (s,f) with respect to the horizon $h(t) : t \in \mathbf{T}_h$. Specifically, it can accommodate any functional $U_h(t,s,f)$. A different valuation criterion would define a new value function $V_h(t,f)$ which, in turn, require a new transformation $\widetilde{V}_h(t,f)$ to represent the discrepancies between $V_h(t,f) : t < T_h$ and the criterion of rational decision-making, $V(T_h,f)$. Nevertheless, this representation of endogenous bounded rationality can still be given as the trade-off between $V_h(t,f)$ and the costs of complexity, $C_h(t,f) = V_h(t,f) - \widetilde{V}_h(t,f)$.

The functional defined in (2) is based upon an intuitive interpretation of limited foresight. Namely, that it imposes informational and/or other decision-making constraints only with respect to contingencies that may occur beyond the reach of the current horizon. An obvious way to evaluate the set of such contingencies is via the expectation operator on the set of the corresponding terminal utility payoffs according to some subjective probability distribution the decision maker places on this set. This valuation criterion produces a ranking of the available plans of action that is independent of affine transformations of the utility representation of the preferences over terminal consequences. If we restrict attention, moreover, to strictly positive utility payoffs, it can be equivalently given by the second term on the right-hand side of (2).

Assumption A.1 ensures that the value function $V_h(t,f)$ is non-decreasing in t, for any $f \in \Delta(Q)$. This monotonicity is crucial for an important feature of my approach to modeling bounded rationality. By specifying precisely how rational decision-making corresponds to a limiting case of the representation, the property provides the necessary structure for the optimal horizon to be obtained endogenously and for the implied complexity costs to be inferred from observed choices. The assumption is restrictive in the following sense. As the foresight t increases, the agent becomes better-informed about the possible paths of play under (s,f) and drops plans $f' \neq f$ such that (s,f') is $h(\tau)$-equivalent to (s,f) for some $\tau < t$ but not for $\tau = t$. Assumption A.1 requires that $\alpha_h(\cdot,s,f)$ is non-decreasing even when $u(h,s,f') > u(h,s,f)$.

This requirement on the agent's beliefs becomes quite weak in the presence of another require-

ment, on the consistency of her forecasts about Nature's play. Specifically, suppose that the agent's forecast is always correct; that is, it is correct from the perspective of every horizon in Γ_h. The set $\mathcal{U}_h(t,s,f)$ of terminal utility payoffs that can realize beyond the horizon $h(t)$ under (s,f) is now given by

$$\mathcal{U}_h(t,s,f) \equiv \left\{ x = u\left(h,s',f\right) - u\left(\Omega_h\left(t,s',f\right)\right) : \; s' \in S, \left(s',f\right) \underset{h(t)}{\simeq} (s,f) \right\} \subset \mathbb{R}_{++}$$

As the foresight $t \in \mathbf{T}_h$ increases, $\{\mathcal{U}_h(t,s,f)\}_{t\in\mathbf{T}_h}$ forms a decreasing nested sequence. Since the agent optimizes over S, Assumption A.1 requires now merely that the agent does not become more pessimistic about her abilities to respond optimally against f as her foresight increases.

In general, $\min_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s,f)} u(h,s,f)$ is non-decreasing it the foresight t (recall relation (i) in the proof of Lemma 1). Under the consistency requirement of the preceding paragraph, $\max_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s,f)} u(h,s,f)$ is non-increasing in t. This enables my approach to model also decision makers whose limited foresight changes along the decision-tree in a way that may not agree with the sequence $\{h(t)\}_{t\in\mathbf{T}_h}$. Optimizing under horizon $h(t)$ requires that the agent considers all contingencies in the t-length chains of information sets for all profiles $(s,f) : s \in S$. Consider, however, an agent who examines the contingencies, that may occur against f, when she follows the plans s and s^* along chains of information sets of length t and $t^* > t$, respectively. If she prefers s to s^*, we can model her choice by choosing appropriate values for $\alpha_h(\cdot,f)$ such that $U_h(t,s,f) \geq U_h(t,s^*,f)$. If she chooses s^* over s, we can construct the valuation criterion such that $U_h(t^*,s,f) \leq U_h(t^*,s^*,f)$. It is easy to check that, for none of the two inequalities above to be valid, the following two conditions must hold simultaneously:

(i) $\max_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s,f)} u(h,s,f) < \min_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s^*,f)} u(h,s^*,f)$

(ii) $\min_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s,f)} u(h,s,f) > \max_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s^*,f)} u(h,s^*,f)$

But this cannot be since (i) leads to a contradiction of (ii):

$$\begin{aligned}
& \min_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s,f)} u(h,s,f) \\
\leq \; & \max_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s,f)} u(h,s,f) \\
\leq \; & \max_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s,f)} u(h,s,f) \overset{\text{(i)}}{<} \min_{s'\in S:(s',f)\underset{h(t)}{\simeq}(s^*,f)} u(h,s^*,f) \leq \min_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s^*,f)} u(h,s^*,f) \\
\leq \; & \max_{s'\in S:(s',f)\underset{h(t^*)}{\simeq}(s^*,f)} u(h,s^*,f)
\end{aligned}$$

By choosing the appropriate foresight and values for $\alpha_h(\cdot,f)$, the agent's actual ranking of s and s^* can be depicted by the model under one of the horizons $h(t)$, $h(t^*)$. As long as the agent's choices come from an admissible rule, they can be represented using the interpretation of limited foresight

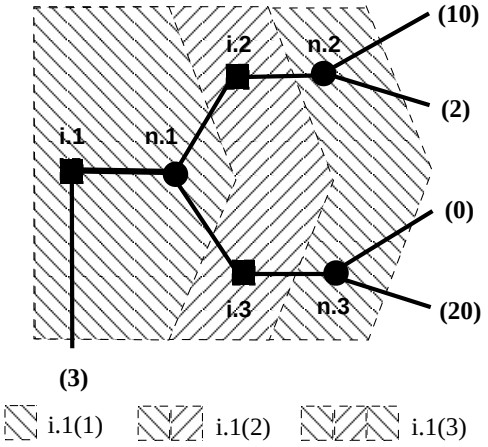


Figure 3: A Non-Complex Problem

presented here even though the true underlying decision-making criterion might not examine the same length of chains of information sets across all available plans of action.

Consistent forecasts provide my approach with another desirable feature. Lemma 2 dictates now that any contingency which lies beyond the current horizon and does not require strategic decisions by the agent cannot add to the current complexity costs. The following example illustrates two immediate implications of this result. First, the benefits and costs of further search are zero along dimensions where further search is futile. Second, calculating expectations with respect to Nature's probability distributions is not costly; it cannot account for departures from rationality. Discrepancies from rational decision-making can result only from the the agent's limited ability to handle the complexity of making optimal decisions.

EXAMPLE 3.2: For the problem of Figure 3, consider the stochastic play f by Nature that places equal probability at the branches of every chance node. There are infinitely many plays, $f' \neq f$, such that (s, f) and (s, f') are $i.1\,(2)$-equivalent to f, for any $s \in S$. If the set $\mathcal{U}_h\,(t, s, f)$ of terminal utility payoffs beyond the horizon is defined as in Section 2.2, the agent's continuation value, for $t \leq 2$, depends on her conjecture about the likelihood of the plays f'. As long as the agent believes that such an f' is played with non-zero probability, $\Delta_t V_{i.1}\,(t, f) \neq 0$ for $t \leq 2$.

Let now $\mathcal{U}_h\,(t, s, f)$ be defined as above. Since the agent has only one available move at $i.2$ and $i.3$, the problem is not complex at any of her decision nodes. The value of $i.1\,(1)$ is 8 (playing down at $i.1$ terminates the game within the reach of the horizon for a utility payoff of 3; playing right corresponds to an expected continuation payoff of 8). The best response and the value of the horizon $i.1\,(t)$ remain unchanged for $t \geq 2$: $\Delta_t V_{i.1}\,(t, f) = 0$ and, thus, $\Delta_t C_{i.1}\,(t, f) = 0$. Notice that the decision-tree here can be reduced so that it consists of a single decision node, at $i.1$, without

changing the complexity of the problem at hand.

3.3 Uniqueness of Optimal Foresight

Recall that $t_h(f)$ denotes the smallest optimal foresight against f at h in (3). This section examines specifications for the transformation function γ_h under which the optimal foresight is unique in (3) - with γ_h^0 being the zero function on $\{1,...,T_h\}$.

Consider the case of a strategy $s^* \in S$, that an agent chooses against f at h, being optimal for more than one lengths of foresight with respect to the criterion V_h. For example, let $U_h(\tau, s^*, f) = U_h(\tau', s^*, f)$, for $\tau, \tau' \in \mathbf{T}_h : \tau < \tau'$. Taking $\gamma_h(\tau, f) > \gamma_h(\tau', f)$ gives $\widetilde{U}_h(\tau, s^*, f) > \widetilde{U}_h(\tau', s^*, f)$, leaving s^* optimal with respect to only the smaller horizon in (3). It is trivial to provide sufficient conditions for a stronger result, to guarantee that the representation in (3) is single-peaked on $\mathbf{T}_h$.

Remark 4 *For the representation in (3), let γ_h^0 be the zero function on $\{1,...,T_h\}$. For $t^* \in \{1,...,T_h - 1\}$, suppose that the following conditions hold:*

(i) $\Delta_t \gamma_h(t, f) > 0, \ \forall t \in \{1,...,t^* - 1\}$

(ii) $\gamma_h(t+1, f)\,\Delta_t V_h(t, f) + \Delta_t \gamma_h(t, f)\, V_h(t, f) < 0, \ \forall t \in \{t^*,...,T_h - 1\}$

Then, $t^* = t(f)$.

Under conditions (i)-(ii), my approach gives an endogenous stopping rule regarding the search for optimal alternatives against f. Specifically, the optimal foresight $t_h(f)$ is obtained by considering the nested sequence of horizons $\{h(t)\}_{t \in \mathbf{T}_h}$ until the marginal benefit from enlarging the search horizon falls short of the marginal complexity cost of doing so:

$$t_h(f) = \min\{t \in \{1,...,T_h\} : \Delta_t V_h(t, f) < \Delta_t C_h(t, f)\}$$

This stopping-rule for search can be interpreted as a *satisficing* condition, to use the terminology of Simon [73]. In contrast to popular interpretations, satisficing here is not the phenomenon of giving up when the decision maker has achieved a predetermined level of utility. As in MacLeod [51], satisficing instead refers to stopping search when it is believed that further search is unlikely to yield a higher net payoff.[13] As an example, let $\gamma_h : \mathbf{T}_h \times \Delta(Q) \to \mathbb{R}_{++}$ be defined by

$$\gamma_h(t, f) = e^{-k_h(f)[t - t^*]^2}, \qquad k_h(f) \in \mathbb{R}_{++} \tag{7}$$

Observe that $\gamma_h(t, f) \leq 1$ on $\{1,...,T_h\}$ and

$$\Delta_t \gamma_h(t, f) = \gamma_h(t+1, f) - \gamma_h(t, f) = \gamma_h(t, f)\left(e^{-k_h(f)[1 + 2(t - t^*)]} - 1\right)$$

[13] This agrees closely with the intuition of Stigler [75] who described the notion of an optimal stopping rule for search of satisfactory alternatives via the example of a person who wants to buy a used car and stops searching when the cost of further search would exceed the benefit from it.

Clearly, $\Delta_t \gamma_h(t,f) > 0$ if $t < t^*$. For $t \geq t^*$, we have

$$\begin{aligned}\Delta_t \left[\gamma_h(t,f) V_h(t,f)\right] &= \gamma_h(t,f)\left[e^{-k_h(f)[1+2(t-t^*)]} V_h(t+1,f) - V_h(t,f)\right] \\ &< \gamma_h(t,f)\left[e^{-k_h(f)} V_h(t+1,f) - V_h(t,f)\right]\end{aligned}$$

Choosing $k_h(f) \geq \max_{t\in\{t^*,...,T_h-1\}} \ln\left(\frac{V_h(t+1,f)}{V_h(t,f)}\right)$, this specification meets the requirements of the corollary. Since all terminal payoffs are finite, there clearly exists $k_h(f)$ sufficiently large so that the corollary applies with $t_h(f) = t^*$.

The conditions (i)-(ii)allow also the complexity cost function to depict the extent of bounded rationality as this varies across agents for a given decision problem.

Definition 5 *Let $\mathcal{C}_h(S,f) = BR_h(t,f)$ and $\mathcal{C}'_h(S,f) = BR_h(t',f)$ be two boundedly-rational choice rules against $f \in \Delta(Q)$ at h. $\mathcal{C}_h(S,f)$ is **more boundedly-rational** than $\mathcal{C}'_h(S,f)$ if $t < t'$ and $BR_h(t,f) \neq BR_h(t',f)$.*

By Proposition 1, the optimal responses against f under foresight t correspond to strictly lower expected payoff than those under foresight t': $V_h(t,f) < V_h(t',f)$. This definition is based, therefore, upon the premise that the more boundedly-rational agent will do worse against f as a consequence of her more myopic foresight. If the optimal foresight is unique, then the behavior of the more myopic decision maker corresponds to higher marginal complexity costs at values of foresight that lie between the two respective optima. More precisely, if $C_h(\cdot,f)$ and $C'_h(\cdot,f)$ are complexity costs functions for the choice rules $\mathcal{C}_h(S,f)$ and $\mathcal{C}'_h(S,f)$, respectively, then

$$\Delta_t C'_h(\tau,f) < \Delta_t V_h(\tau,f) < \Delta_t C_h(\tau,f) \quad \forall \tau \in \{t,...,t'-1\}$$

Remark 5 *Suppose that $\mathcal{C}_h(S,f) = BR_h(t,f)$ is a more boundedly-rational choice rule than $\mathcal{C}'_h(S,f) = BR_h(t',f)$ against f at h. Suppose also that $C_h(\cdot,f)$ and $C'_h(\cdot,f)$ are the corresponding complexity cost functions. Then*

$$\Delta_t C_h(\tau,f) > \Delta_t C'_h(\tau,f) \quad \forall \tau \in \{t,...,t'-1\}$$

3.4 Empirical Testing

Proposition 2 and Corollary 1 suggest a procedure for empirical inference in decision problems where observed choices reveal the limited foresight of the decision maker. Suppose that, against f at h, the agent's optimal foresight is t^*: $\frac{V_h(t^*+1,f)}{V_h(t^*,f)} \leq \frac{\gamma_h(t^*,f)}{\gamma_h(t^*+1,f)}$. Let, moreover, the terminal utility payoffs change so that $V_h(t^*+1,f)$ increases, say to $V_h(t^*+1,f)+v$, other things being equal. For a sufficient increase, the agent will find it optimal to extend her horizon beyond $h(t^*)$. The smallest required such increase is $v_0 = \frac{\gamma_h(t^*,f)}{\gamma_h(t^*+1,f)} V_h(t^*,f) - V_h(t^*+1,f)$. In decision problems where different strategic choices by a boundedly-rational agent correspond to different levels of optimal foresight, v_0 can be used to identify the quantity $\frac{\gamma_h(t^*,f)}{\gamma_h(t^*+1,f)}$ and, thus, the objective in (3), at various levels of t, when $g_{0h}(\cdot,f)$ is the zero function on $\mathbf{T}_h$.

To illustrate, consider the decision-making situation of Example 2.4 with $p < \frac{5}{8}$. If an agent plays Up at $i.1$, $\frac{g_{i.1}(1,f)}{g_{i.1}(2,f)} \geq 1 + \frac{\max\left\{0, \frac{1}{2}\alpha_{i.1}(2,m,f) - 2p\right\}}{1+2p}$. If she chooses the middle branch with an intention to play Right (Up) at $i.4$, $\frac{g_{i.1}(1,f)}{g_{i.1}(2,f)} \leq 1 + \frac{\max\left\{0, \frac{1}{2}\alpha_{i.1}(2,m,f) - 2p\right\}}{1+2p}$ and $\frac{g_{i.1}(3,f)}{g_{i.1}(4,f)} > (\leq)\, 1 + \frac{1-\frac{1}{4}\alpha_{i.1}(3,m,f)}{\frac{5}{4}+\frac{1}{4}\alpha_{i.1}(3,m,f)}$. We can proceed now in two ways. We ask an agent who chose Up at $i.1$ to consider the problem again and again, with ever increasing payoffs at $n.4$ and $i.4$, until the first instant where she chooses the middle option. This exercise would give an upper bound for $\frac{g_{i.1}(1,f)}{g_{i.1}(2,f)}$ and an increasing array of lower bounds. If the agent plans at $i.1$ to play Right at $i.4$, we ask her to consider the problem more times while increasing the payoffs at $i.4$ until the first instant where she announces the intention to move Up at $i.4$. This would give a lower bound for $\frac{g_{i.1}(3,f)}{g_{i.1}(4,f)}$ and a decreasing array of upper bounds. By refining our process of payoff-increases, we can reduce the size of the resulting intervals for the quantities $\frac{g_{i.1}(1,f)}{g_{i.1}(2,f)}$ and $\frac{g_{i.1}(3,f)}{g_{i.1}(4,f)}$ to acceptable estimates.

Alternatively, we can present the problem, with a randomly-drawn payoff structure, to a population of agents and estimate the infimum for (i) $\frac{V_{i.1}(2,f)}{V_{i.1}(1,f)}$ such that the middle option is chosen at $i.1$ and (ii) $\frac{V_{i.1}(4,f)}{V_{i.1}(3,f)}$ such that the intention to move Up at $i.4$ is announced at $i.1$, given that the middle branch is chosen at $i.1$. Observe that, under either method, the model is falsified if agents are observed to switch from the middle to the upper branch at $i.1$ or from Up to Right at $i.4$ while presented with increasing $\frac{V_{i.1}(2,f)}{V_{i.1}(1,f)}$ or $\frac{V_{i.1}(4,f)}{V_{i.1}(3,f)}$, respectively.

In terms of the representation in (6), Proposition 2 implies that the benefits and the implied complexity costs of search are related in a particular way: when the benefits of further search are increased, the costs of search also increase but not as quickly. Empirical inference of the complexity cost function can be based on the fact that higher benefits to search offer an incentive for the decision-maker to extend her search horizon. Notice that these are costs of bounded rationality *implied* by the model. They determine the cost-benefit trade-offs required for it to replicate the observed choices. Thus, my analysis abstracts from some well-known issues associated with modeling limited search as optimization with decision costs. First, reliable estimates of the true underlying benefits and costs (such as opportunity or computational costs) can demand large degrees of knowledge. Second, the knowledge and the computations involved can be so massive that one is forced to assume that ordinary people have the computational abilities and statistical software of econometricians (Conlisk [14], Sargent [67]). And of course, there is the issue of infinite regress: the cost-benefit computations themselves are costly and demand some meta-level cost-benefit computations and so on (Conlisk [14]).

In terms of cost-constrained decision-making, there must indeed exist some regression problem but it need not be infinite. If decision makers do follow some sequence of meta-level cost-benefit computations, the very fact that they do deliver strategic responses means that somehow this sequence has converged to a stopping point after a finite number of iterations (Bewley [6] presents a setting where this is actually optimal). My *as if* approach to modeling departures from rationality is based on a simple observation. Whatever goes on inside agents' minds in decision-making situations

(what many authors have referred to as the "black box" of bounded rationality), it does result in their choices. These choices reveal a particular cost-benefit relation when they are compared to the behavior under limited foresight presented here. And this is sufficient for a representation.

3.5 Related Literature

The analysis presented here can be viewed as a generalization of an interesting class of limited foresight models (Jehiel [39], [40]). In these, the agent's foresight is exogenously-given, at some $t_h \in \mathbf{T}_h \setminus \{T_h\}$ for all $f \in \Delta(Q)$, and attention is restricted to situations where the complexity of the problem increases along a single dimension: the length of foresight determines the unique chain of singleton information sets the agent considers. This can be presented as a special case of the specification in (7) with $t^* = t_h$ and $k_h(\cdot) = k \geq \max_{(t,f)\in\{t_h,...,T_h-1\}\times\Delta(Q)} \ln\left(\frac{V_h(t+1,f)}{V_h(t,f)}\right)$. When the foresight is fixed, the representation in (3) reduces to $V_h(t_h, f)$, for any $f \in \Delta(Q)$, which is more tractable but also harder to falsify. In the face of incorrect predictions, one cannot easily distinguish between the limited foresight approach being inadequate or having assumed the wrong length of foresight. This identification issue becomes intractable in general problems, like the one depicted in Figure 1, where the complexity changes also along the depth dimension of the tree and the change in the cardinality, $\Delta_t|h(t)|$, of the horizon is not necessarily constant on $\mathbf{T}_h$.

A sizeable literature on information theory and computational complexity has avoided the assumption of fixed foresight. Instead, boundedly-rational decision-making in complex problems is modeled via algorithms for searching through the tree. The central issue, of which direction and how far into the tree one should look ahead, is usually addressed by various interpretations of the following principle. A search path is evaluated by assigning (i) the utility payoff given by the underlying preferences over terminal consequences, for those terminal nodes that can be reached within the path, (ii) some continuation value for the remaining problem beyond the path, and (iii) some computational costs for considering the contingencies encountered by the path. McLeod [51] prefers to combine (ii) and (iii) into one heuristic function since their decomposition often leads to underestimating the difficulty of the problem: it assumes implicitly that one can costlessly create probability distributions over uncertain events. He argues that the complexity of a problem is very sensitive to not only the size of the search graph but also the structure of the heuristic function.

The similarities between this literature and my approach are obvious. With respect McLeod's criticism, the implied complexity costs of search are derived here via a framework that does assume costless assignment of probability distributions over uncertain events. Its advantage, however, is that it allows for the discrepancies, between the criterion defined by (i)-(ii) and rational decision-making, to be measured by these costs alone and, therefore, for the costs to be inferred from observed choices. The heuristic function, as well as exogenously-given decompositions of it in (ii)-(iii), impose exogenously the interdependency between two sources of bounded rationality: limited foresight and the perception of the net benefits of further search. As in the models of fixed limited

foresight, this hinders inference of the costs of search from observed behavior.

In the spirit of the search-cost algorithms literature, Gabaix and Laibson [25], propose a representation that corresponds to (3) with $\gamma_h(t,f)=1$ and $\gamma_h^0(t,f)=kt$ on $\{1,...,T_h\}\times\Delta(Q)$, for some positive constant k. This is a parsimonious and tractable model that lends itself easily to the authors' experimental analysis. However, Remark 1 and Claim 2 highlight two important issues. Normalizing $\gamma_h(\cdot)$ makes the predictions about the optimal foresight of search and the corresponding optimal choices highly sensitive to the units used for the utility representation of preferences over terminal consequences. Moreover, the empirical validity of any form of the representation in (3) cannot be established just by showing that it explains observed choices better than rational decision-making. Departures in choice behavior from some rational paradigm can be modeled in many ways, including ones that do not view search-costs as the source of bounded rationality. For example, there is a vast literature which prefers to parameterize bounded rationality directly in the preferences over terminal consequences or in the formation of beliefs. One of the main results of this paper implies that the fundamental test of the validity of the search-costs approach should be whether or not the experimental data supports a systematic relation between the extent of search and the expected benefits of further search.

A considerable body of work on bounded rationality has modeled complexity in repeated games via the use of finite automata. These are machines with a finite number of states whose operating instructions are given by (i) an output function that assigns an action at each state and (ii) a transition function that assigns a succeeding state to every pair of a current state and an action by the opponent (Neyman [**?**], Rubinstein [63], Abreu and Rubinstein [1], Piccione and Rubinstein [56], Kalai and Stanford [45]). An automaton is meant to be an abstraction of the process by which a player implements a rule of behavior. The complexity of a machine M is given by the number of its states, $,M,$. A player's preferences over machines are lexicographic along two dimensions: they are increasing in her payoff in the repeated game and decreasing in complexity, the former being the dominant relation.

A machine M defines a plan of action for a finite number of contingencies and a strategy $q(M)\in Q$ which repeats this plan cyclically. Each cycle corresponds to a horizon of foresight $t(M)=r,M,$ (r being the number of actions available to the agent in the stage-game) while, by being a cycle, play returns to the initial state at the end of it. Diasakos [18] shows that, when facing a machine M, the decision-making rule of the implementation complexity approach is equivalent to $V_h(t(M),q(M))$ with $\alpha_h(\cdot)=1$ on $\mathbf{T}_h\times\Delta(Q)$. Since play returns to the initial state at the end of the horizon, the worst continuation payoff beyond the horizon, $\min_{(s',q')\underset{h(t(M))}{\simeq}(s,q(M))}u(s',q')$, is independent of the $s\in S$ played within the horizon. Hence, $V_h(t(M),q(M))$ chooses $s(M)\in S$ that responds optimally against $q(M)$ within $h(t(M))$ and is $h(t(M))$-equivalent to the optimal machine against M.

The complexity cost approach agrees also with most of the "general principles" of decision-making in complex problems given by Geanakoplos and Gray [27]. The nested sequence of horizons

$\{h(t)\}_{t=1}^{T_h}$ guarantees that the *proximity* principle (thinking about the immediate consequences of your actions first) is met, but leads generally to at least as good choices as the *relevance* principle (examining the most promising alternatives first).[14] Assigning continuation values beyond the horizon leads, in general, to a set of options that are considered good rather than singling out one possibility as excellent (the *family* principle). By appropriate choice of the $\alpha_h(\cdot)$ function, this can be consistent with the *clear sight* (if possible, one should base decisions on clear-cut results: observations with the lowest error-probability should typically receive the most weight) and *stability* (consistency increases confidence so that, if our guesses are accurate, we should guess values along the intended path that change very little) principles.

4 Applications

This section illustrates several important aspects of complexity costs and demonstrates that the model is consistent with three well-known paradoxes in decision-making. The paradoxes obtain, for obviously different reasons, in situations that are seemingly simple and unrelated. Accordingly, a variety of models and arguments have been suggested in the literature to explain them. Yet, the complexity-costs approach provides a common framework for depicting the underlying limitations that might force choices to depart from those prescribed by rationality. For the first application, in particular, the model reconciles the findings of two experimental studies that seem to be at odds under other theoretical explanations.

In each application, paradoxical behavior results when the complexity costs are large relative to the benefits of search, preventing a boundedly-rational agent from considering the problem in its entirety. For simplicity, the criterion $U_h(\cdot)$ in (2) is taken with $\alpha_h(\cdot) = 1$ on $\mathbf{T}_h \times S \times \Delta(Q)$. Even though the agent is assumed to use only the worst-case scenario as her continuation payoff beyond her current horizon, all three cases correspond to complex decision-making situations. Moreover, in each setting, the best response set under rationality is a singleton and we can construct $g_h(\cdot)$ on $\mathbf{T}_h \times \Delta(Q)$ such that any other choice set can be optimal in (3), for some $t^* < T_h$. Notice also that, in all cases, Nature's moves are dictated by distributions that are known when the agent makes her choice. Clearly, the consistency of forecasts requirement of Section 3.2 is met.

4.1 Temporal Framing

The two problems of Figure 4 differ only in their temporal framing. In Γ_{np}, a random draw takes place first. With probability $1-p$, the game ends immediately and the decision-maker gets $\$x_3$. Otherwise, the agent is asked to choose between ending the game and receiving $\$x_2$ ($x_2 > x_3$) or

[14] The relevance principle dictates that the path defined by the most promising alternative is followed deeper into the tree than the remaining alternatives. In contrast, the horizon of search applies here the same depth to all alternatives and, thus, leads to at least as good choices. However, the discussion of Section 3.2 indicates that choice-rules obeying the relevance principle can be represented in my framework under the consistency requirement.

allowing for another random draw in which she can get $\$x_1$ ($x_1 > x_2$) with probability q and $\$x_3$, otherwise. The setup of Γ_p is the same but for pre-commitment: the agent must decide, before any resolution of uncertainty, whether or not she will terminate the game if it does not end at the first draw.

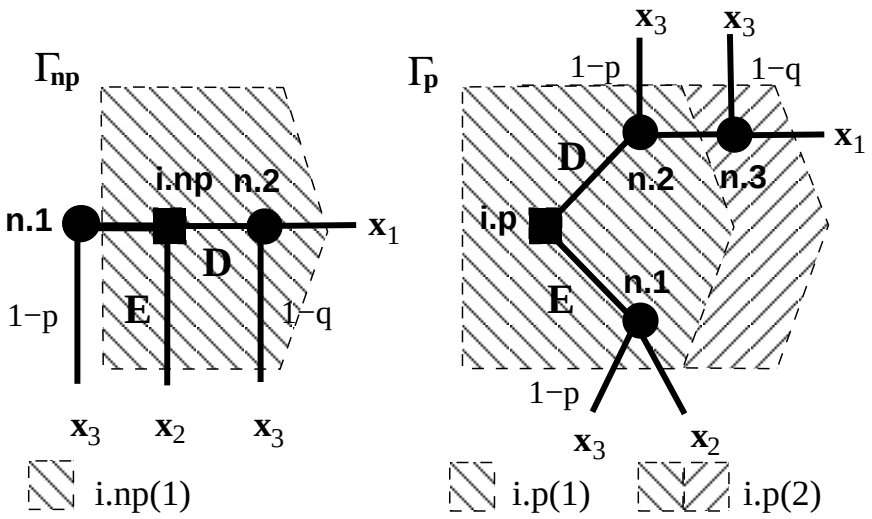


Figure 4: Temporal Framing (Extensive Forms)

Standard economic theory offers two methods for solving such decision problems: the strategy method and backward induction. Under the strategy method, the decision-maker looks at the normal form representation of the game, calculates the implications of each available strategy for the probability distribution over final outcomes, and chooses the optimal strategy with respect to her underlying preferences over final outcomes. For this method, the two problems are identical as they have the same normal form representation. In both, the agent chooses only once from two available actions: "Take the Draw" (D) and "End the Game" (E). In both, choosing E versus D corresponds to a difference of $\delta = p\left[u(x_2) - qu(x_1) - (1-q)\,u(x_3)\right]$ in unconditional ex-ante expected utility. Let $\delta > 0$, so that E is optimal, and denote by f Nature's play.[15]

[15]The normal form representation of either problem is depicted in Table 1. In terms of the notation introduced in Section 2, let X and Y denote the outcomes "Nature's last draw ended the game paying x_3" and "not X", respectively. Then $Q = \{X, Y\} \times \{X, Y\}$ where $(x, y) \in Q$ is ordered according to the outcome of the first draw and second draw, respectively. The two independent draws in Γ_{np} and Γ_p are given by the probability distribution

Table 1: Temporal Framing (Normal Form)

	UU	UD	DU	DD
D	x_1	x_3	x_3	x_3
E	x_2	x_2	x_3	x_3

In contrast, the complexity-costs approach can allow for different choices across the two problems. In Γ_{np}, the agent does not have to make any decisions prior to the first draw. Starting at $i.np$, the horizon $i.np\,(1)$ covers the continuation problem. Complexity costs $C_{i.np}\,(1,f)$ suffice, therefore, for her to figure out that choosing between E and D corresponds to an expected utility difference of $\frac{\delta}{p}$. The net expected utility difference is $\frac{\delta}{p} - C_{i.np}\,(1,f)$. In problem Γ_p, she has to commit in advance to an action and it is the horizon $i.p\,(2)$ that covers the continuation problem (once again, an increasing sequence of nested horizons is depicted, in the figure, by combining adjacent hatched regions from left to right). However, Nature's plan f is known here in advance and, thus, the agent's forecast has to be consistent. By Lemma 2, therefore, enlarging the horizon is not costly if it introduces only information sets where no decisions are made: $C_{i.p}\,(2,f) = C_{i.p}\,(1,f)$. The net expected utility difference from considering this problem is $\delta - C_{i.p}\,(1,f)$.

If

$$\delta - C_{i.p}\,(1,f) < 0 \leq \frac{\delta}{p} - C_{i.np}\,(1,f) \tag{8}$$

the two problems are not equivalent. The pre-commitment problem entails complexity that the agent finds more costly than the expected benefit from fully considering her two choices. Since it is not optimal to incur the costs of complexity, she will be indifferent, at $i.p$, between continuing or terminating the game after the first draw. In problem Γ_{np}, on the other hand, incurring the complexity costs is optimal and she will choose to terminate the game at $i.np$.

Cubitt *et al.* [17] investigated such choices experimentally. They presented their subjects with (non-graphical) descriptions of Γ_{np}, Γ_p, and a third problem Γ_s. In Γ_s, the agent must choose between the simple lottery that follows E and the compound one that follows D, respectively, in the tree depicting Γ_p in Figure 4. Their study found that play in the first two problems differs in a statistically significant way: up to 70% of their subjects chose E in Γ_{np} and 57% chose D in Γ_p. They did not observe any significant difference in play across the last two problems.[16]

Since Nature is a non-strategic player, Γ_s is equivalent in extensive form to Γ_p and, consequently, in normal form to both Γ_{np} and Γ_p. This means that the strategy method cannot account for the findings of the study. Under backward induction, however, the agent looks at the extensive

$f = \{pq, p(1-q), q(1-p), (1-p)(1-q)\}$ on $\{(Y,Y),(Y,X),(X,Y),(X,X)\}$.

[16] Problems 2, 3 and 4 in Cubitt *et al.* [17] correspond to Γ_{np}, Γ_p, and Γ_s, respectively. Problems 1 and 5 are, respectively, scaled versions of 2 and 4. The authors run their experiments with $x_1 = 16$, $x_2 = 10$, $x_3 = 0$, $p = 0.25$, and $q = 0.8$.

form representation and may not be indifferent across the three problems depending on how the pre-commitment problem Γ_p is evaluated. There are two possibilities. The agent may consider her decision as being made at the beginning of the game, in which case Γ_p is equivalent to Γ_s. Alternatively, she may regard the decision as to be made when she actually gets to play. Now, Γ_p is equivalent to Γ_{np} (regardless of whether backward induction is combined with reduction of compound lotteries or substitution with certainty equivalents).

In either case, under the standard setting of expected utility optimization, E ought to be chosen. The authors turn, therefore, to non-expected utility theories for an explanation of their experimental finding. They consider those of Machina [50], Segal [69], Kahneman and Tversky [44], and Karni and Safra [46]. Except for that of Karni and Safra, all theories are rejected as they require timing independence in decision-making which forces the same predictions for Γ_{np} and Γ_p. In contrast, the behavioral consistency theory of Karni and Safra does not require timing independence and views Γ_p as equivalent to Γ_s, but not to Γ_{np}. Because it is based on generalized expected utility, the theory permits the common ratio effect to obtain, accommodating different responses across Γ_{np} and Γ_p.[17]

An alternative explanation can be supported by the model of Koszegi and Rabin [47]. An agent can opt for different choices across the two problems if her risk attitudes depend upon her reference point. Under pre-commitment, her reference point can be her current endowment, relative to which Γ_p involves no potential losses: E corresponds to gaining x_2 with probability p or remaining at the reference point otherwise, D results in a gain of x_1 with probability pq or remaining at the reference point otherwise. Problem Γ_{np}, on the other hand, does entail the possibility of losses. Choosing E at *i.np* guarantees x_2 which can be incorporated into the agent's reference point. D corresponds now to a gamble between a gain of $x_1 - x_2$ with probability pq and a loss of $x_2 - x_3$ otherwise.

Explaining the findings of Cubitt *et al.* [17] based on either the common ratio effect or reference-dependence loss aversion seems to be at odds, however, with the findings of another study. Instead of choices between continuing or terminating the game after the first draw, Hey and Paradiso [38] elicit subjects' preferences for the three problems via their willingness to pay for playing them.[18] Since all valuations are made ex-ante, the reference point in Γ_p and Γ_{np} is the same and the problems are equivalent under reference-dependence loss aversion. Yet, only 56% of the subjects revealed that they were indifferent between the two. Under the common ratio effect, Γ_p should be preferred to Γ_{np}: the ex-ante expected value of choosing D in Γ_p is $pqu(x_1) + (1 - pq)u(x_3)$ while that of E in Γ_{np} is $pu(x_2) + (1 - p)u(x_3)$. Nevertheless, only a third of the subjects were willing to pay a significantly higher amount for playing Γ_p.

[17] Consider lotteries defined on the set $\{x_1, x_2, x_3\}$ of monetary consequences. A simple prospect can be denoted by the vector $(\pi_1, \pi_2, 1 - \pi_1 - \pi_2)$ of probabilities assigned, respectively, to x_1, x_2, and x_3. The *common ratio effect* occurs when the option $(0, 1, 0)$ is chosen over $(q, 0, 1 - q)$ while $(pq, 0, 1 - pq)$ is chosen over $(0, p, 1 - p)$, for some $q, p \in (0, 1)$. Notice that the former pair of lotteries corresponds to the decision the agent faces at *i.np* while the latter pair to the one at *i.p*.

[18] Hey and Paradiso use $x_1 = 50$, $x_2 = 30$, and $x_3 = 0$ while p and q are the same as in Cubitt *et al* [17].

The complexity-costs approach can accommodate the findings of Paradiso and Hey. For the majority of their subjects, the complexity costs seem to be trivial and the three problems are equivalent; each one has an ex-ante expected value of $pu(x_2) + (1-p)u(x_3)$. For a third of the subjects, though, the costs are non-trivial but incurring them is optimal. The expected utility payoff from any problem is $pu(x_2) + (1-p)u(x_3)$. If $C_{i.p}(1,f) < C_{i.np}(1,f)$, however, the net expected value of the pre-commitment problem Γ_p is higher than that of Γ_{np}. This relation for the complexity costs across the two problems can indeed obtain here since the value of the horizon is always an upper bound for its complexity costs.[19] These upper bounds for $C_{i.np}(1,f)$ and $C_{i.p}(1,f)$ are $u(x_2)$ and $pu(x_2) + (1-p)u(x_3)$, respectively.

The complexity costs depict the expected utility difference between fully- and boundedly-rational choices. Obviously, the higher the expected benefit from thinking further into a decision problem, the higher ought to be the penalty for not doing so. In this example, the boundedly-rational choice is always to not terminate the game after the first draw. In problem Γ_p, failing at $i.p$ to choose E corresponds to an expected loss of $-p\mathbb{E}[\Delta u(x)]$. In Γ_{np}, on the other hand, failing to choose E at $i.np$ carries an expected loss of $-\mathbb{E}[\Delta u(x)]$. Ex-ante, the expected loss from making the wrong choice is the same in both problems. At the time of making the decision, however, it is higher at $i.np$.

When the entire decision-tree is costlessly considered, the ex-ante value of a problem does not depend upon the timing of decisions. Under limited foresight, though, the expected benefit of further search does. Since the complexity costs measure the discrepancies between the two, the timing of decisions matters for the costs of complexity. Moreover, it matters even when the costs do not prevent the decision-maker from considering the entire problem and choosing the plans of action that are optimal under rational decision-making. The complexity-costs approach allows for violations of both timing independence and dynamic consistency in decision-making.[20] The timing of the resolution of uncertainty is important, therefore, beyond the way in which it is when dynamic consistency is required. Under dynamic consistency, the timing of the resolution of uncertainty matters for the ex-ante value of a decision problem only if it affects the choice of optimal plans. Under complexity costs, it matters even when the optimal choice remains the same allowing the approach to reconcile the findings of the two experimental studies.

The pre-commitment problem requires essentially thinking at $i.p$ about the strategic position at $i.np$, after some partial resolution of uncertainty. Relative to its ex-ante expected utility payoff, this is a more complex undertaking than just waiting to see in which strategic position one ends up and deciding then, as in problem Γ_{np}. For an agent for whom such complexity is too costly relative to the benefit of thinking about pre-commitment, her choices may differ in the two problems. For

[19] Recall that $C_h(t,f) = [1 - g_h(t,f)]\, V_h(t,f)$, with $g_h(t,f) \in (0,1]$ for all $t \in \mathbf{T}_h$.

[20] It is easy to see how dynamic inconsistencies can obtain here. Let Γ_p be a problem about advanced planning instead of pre-commitment. If (8) holds, the agent cannot decide what her best move should be at $i.p$; she is indifferent between her two actions. When she gets to $i.np$, however, she will prefer to not take the second draw.

an agent for whom the complexity costs are not too high yet non-trivial, the choice will be the same but the ex-ante expected net value of Γ_p can exceed that of Γ_{np}. The latter agent can be viewed as sophisticated about her bounded rationality. She realizes that she may err in her choices and this makes the amalgamation of risks offered by the pre-commitment problem Γ_{np} valuable as it mitigates the expected cost of her errors. She prefers to pre-commit as she knows that erring at $i.np$ will be more costly than it appears to be at $i.p$.

4.2 Diversification Bias

Consumers have been repeatedly found to be more variety-seeking when they plan for future consumption than when they make separate sequential choices preceding each consumption period. This effect was first demonstrated by Simonson [74], who gave students the opportunity to select among six snacks in one of two conditions: (i) they picked one snack at each of three weekly class meetings (sequential choice), (ii) on the first meeting, they selected three snacks to be consumed one per week over the meetings (simultaneous choice). The subjects displayed significantly more variety-seeking under simultaneous choice where 64% chose three different snacks. Under sequential choice, only 9% did so. Simonson suggested that this behavior might be explained by variety-seeking serving as a choice heuristic (when asked to make several choices at once, people tend to diversify). Thaler [77], on the other hand, regards this behavior as resulting from a failure of predicted utility to accurately forecast subsequent experienced utility.

Read and Loewenstein [60] replicated Simonson's diversification effect (they refer to it as the "diversification bias") for snacks in several experiments. In one study, conducted on Halloween night, their subjects were young trick-or-treaters who approached two adjacent houses. In one setting, the children were offered a choice between two candies at each house. In another, they were told at the first house they reached to "choose whichever two candy bars you like" (with large piles of both candies being displayed so that the children would not think it rude to take two of the same). The results showed again a strong bias for diversification under simultaneous choice: every child selected one of each candy. In contrast, only 48% of those given sequential choice picked different candies. Evidence of diversification bias has been found also in studies of defined contribution savings plans, supermarket purchases, choices of lottery tickets, and audio tracks (Read *et al.* [61]).

In what follows, I present an example of the complexity-costs model which is consistent with the findings of the experiment by Read and Loewenstein. Although the set-up does not explain why people err in the direction of diversification, it accounts for why they do not diversify in the sequential setting.[21] Consider a decision-maker who must choose one of two snacks, a and b, on

[21]In the Read and Loewenstein experiment, the decision problem is presented to subjects in a way that allows for erring in only one direction, choosing a diversified plan. Bias towards diversification has been demonstrated also in other experiments that allow for making errors along more than one dimensions (Read *et al.* [61]). The one-dimensional character of Read and Loewenstein, however, allows for a description of the problem by a simple model

each of two consecutive periods, 1 and 2. The objects of choice, therefore, are two-period plans in $X = \{a, b\} \times \{a, b\}$. Let $(x, y) \in X$ denote the plan in which x is chosen in the first period and y in the second. Experiments of diversification bias use objects that are familiar to the subjects so that consumption per se provides no new information to the decision-maker about the product itself. Nevertheless, even without any objective uncertainty about static preferences across the different items, the agent has subjective uncertainty about which plan she would prefer over the dynamic setting she must consider. Let this kind of uncertainty be completely described by a finite state space Ω and a probability distribution μ on Ω. For $((x, y), \omega) \in X \times \Omega$, the state-dependent utility is given by some function $u : X \times \Omega \to \mathbb{R}_{++}$.

Since the aim of this example is to illustrate how the temporal nature of a decision problem matters for observed choices, the analysis can be simplified by abstracting from the static preferences across snacks. I will assume that the choice of a particular snack matters only with respect to whether or not it adds variety in consumption across periods:

Assumption 4.2.A

$$u(a, b, \omega) = u(b, a, \omega) \quad \text{and} \quad u(a, a, \omega) = u(b, b, \omega) \qquad \forall \omega \in \Omega$$

Notice also that, temporal issues aside, diversification should *not* be ex-ante preferable; otherwise, the bias is not really paradoxical for we could have the following scenario. In the simultaneous-choice setting, the diversified plan is chosen since it is optimal according to the ex-ante expected-utility ranking. Under sequential-choice, choosing $x \in \{a, b\}$ in period 1 defines the period-2 choice set to be $\{(x, a), (x, b)\}$. The choice from this set will be according to the state-dependent utility ranking for the state that obtains in period 2. Diversification would be optimal now at any of the states in $\Omega_0 = \{\widetilde{\omega} \in \Omega : u(x, y, \widetilde{\omega}) > u(x, x, \widetilde{\omega})\}$, for $y \neq x$. Since Ω_0 occurs with probability $\mu(\Omega_0)$, an outside observer would record a diversified plan as chosen always in the simultaneous setting but only $100 \times \mu(\Omega_0)\%$ of the time in the sequential one.

Assumption 4.2.B

$$\mathbb{E}_\Omega [u(x, y, \omega)] < \mathbb{E}_\Omega [u(x, x, \omega)] \qquad \forall x, y \in \{a, b\} : x \neq y$$

Under these two assumptions, rational decision-making under expected utility is not consistent with the diversification bias: an *undiversified* plan is chosen always in the simultaneous setting and $100 \times (1 - \mu(\Omega_0))\%$ of the time in the sequential one. However, this is not necessarily the case in the complexity-costs model. The temporal framing of the decision problem matters for choices, by affecting the choice set of the agent and, consequently, whether or not dynamic inconsistency matters for final outcomes.

which explains, without loss of generality, why choices differ across the two settings.

Under sequential choice, opting for x in period 1 leaves $\{(x,a),(x,b)\}$ as the feasible set for period 2. This allows for any dynamically inconsistent choice of plan, made in period 1 (due to complexity being too costly), to be reversed in period 2. Under simultaneous choice, opting for (x,y) in period 1 makes it the singleton feasible set for period 2. The diversification bias can obtain because the choice of plan made in period 1 is now irreversible.

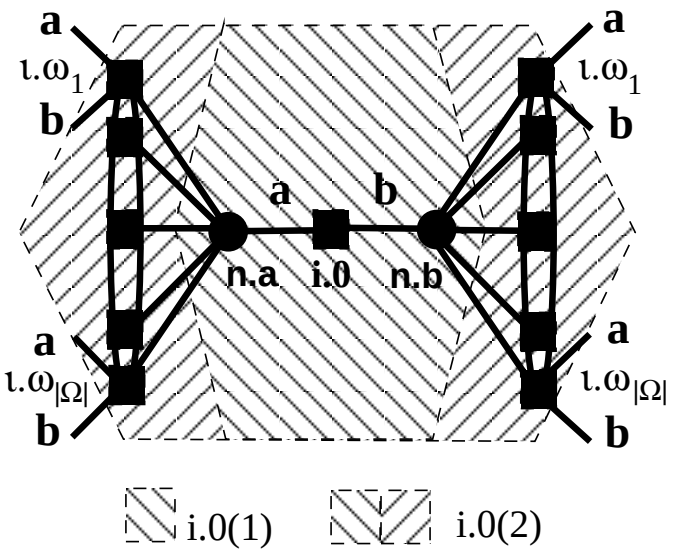


Figure 5: Simultaneous Choice

Consider first the simultaneous choice problem depicted by Figure 5. The period-2 realizations of the states in Ω are represented by the terminal nodes $i.\omega$. As in the preceding figures, a sequence of nested horizons is depicted by combining adjacent hatched regions. In period 1, the agent has to choose and commit to one of the four plans in X. At $i.0$, when thinking about the possible terminal contingencies, she faces two paths. Each has length $t = 2$ and is associated with the period-2 realization of one of the states in Ω, following a period-1 choice of snack. Accounting for all terminal contingencies requires, therefore, complexity costs $C_{i.0}(2,\mu)$. Since there is a unique prior about the terminal nodes, the decision-making rule reduces to that of the standard expected utility ranking between the four plans across the states in Ω. The value of the horizon $i.0(2)$ is

given by

$$V_{i.0}(2,\mu) = \max_{(x,y)\in X} \sum_{\omega\in\Omega} \mu(\omega)\, u(x,y,\omega) = \max_{(x,y)\in X} \mathbb{E}_\Omega [u(x,y,\omega)]$$

The decision-maker will choose either of the two *undiversified* plans.

Suppose, however, that the agent limits her foresight to $t = 1$. The corresponding complexity cost is $C_{i.0}(1,\mu)$. Since the horizon does not reach the terminal nodes, the decision problem becomes

$$V_{i.0}(1,\mu) = \max_{x\in\{a,b\}} \min_{y\in\{a,b\}} \sum_{\omega\in\Omega} \mu(\omega)\, u(x,y,\omega) = \min_{(x,y)\in X} \mathbb{E}_\Omega [u(x,y,\omega)]$$

She will choose either of the two *diversified* plans. In other words, the following condition

$$\Delta_t C_{i.0}(1,\mu) > \mathbb{E}_\Omega [u(x,x,\omega)] - \mathbb{E}_\Omega [u(x,y,\omega)] \qquad \{x,y\} \in X : x \neq y \tag{9}$$

suffices for the agent to choose a diversified plan when she has to commit to a decision about a two-period plan at the beginning of period 1.

Under sequential choice, if $x \in \{a,b\}$ was picked in period 1, the period-2 choice will be from $\{(x,b),(x,a)\}$. The decision-tree is the same as that of Figure 5, albeit for each state $i.\omega$ that follows from $n.a$ or $n.b$ defining a singleton information set. At each $i.\omega$, the decision problem is trivial in terms of complexity costs: $V_{i.\omega}(\mu) = \max_{y\in\{a,b\}} u(x,y,\omega)$. At $i.0$, though, the problem is similar to that of Figure 5 in terms of complexity. There are two horizons the agent might consider: $i.0\,(1)$, and $i.0\,(2)$. The corresponding complexity costs $C_{i.0}(1,\mu)$, $C_{i.0}(2,\mu)$ and the values $V_{i.0}(1,\mu)$, $V_{i.0}(2,\mu)$ are the same as before.

If (9) holds, the agent will again find it not worthwhile, at $i.0$, to think about her period-2 choice once she has chosen something in period 1. She will choose either of the two diversified plans. If (9) does not hold, on the other hand, she will choose either of the two undiversified ones. Nevertheless, a choice of a plan $(x,y) \in X$ at $i.0$ is now merely an *intention* to pick y in period 2, not a commitment. Having chosen x in period 1, when she chooses again at $i.\omega$, it will be out of the set $\{(x,a),(x,b)\}$. Picking $(x,y) : y \neq x$ will then be optimal at any state in Ω_0. Since this occurs with probability $\mu(\Omega_0)$, an outside observer would record a diversified plan only $100\mu(\Omega_0)\%$ of the time.

When (9) holds, the predictions of the model are in agreement with the observed bias towards diversification. Under simultaneous choice, at the beginning of period 1 the agent must commit to a plan which will irrevocably determine her utility for both subsequent periods. Clearly, there is an advantage in considering the entire time-horizon of the problem. If the corresponding complexity is too costly, however, diversification results.

Under sequential choice, at the time of making the period-2 decision, the uncertainty about preferences will have been resolved. Thinking about this resolution of uncertainty in period 1 is still a complex undertaking but the agent is not penalized for failing to do so. If the complexity costs are too high so that (9) applies, she will have the wrong intentions in period 1 about period-2

play. Yet, this can be corrected by period-2 actual decisions. Diversification as a final outcome obtains only in the states of Ω_0, when it is actually called for.

4.3 Risk Aversion in the Small and Large

This application presents an intertemporal model of consumption and savings in which larger amounts of wealth to be allocated across time induce longer planning horizons. Such a relation between the wealth and the length of the optimal horizon has significant implications for risk attitudes in the small and large. Here, I show that the model can accommodate plausible but contrasting preferences between small- and large-stakes gambles that cannot be explained under standard settings.

Consider a decision-maker who has an endowment w_0 and a choice between accepting or not a lottery $\widetilde{z}$ which would result in her wealth being $\widetilde{w}_0 = w_0 + \widetilde{z}$. Rabin [58] and Safra and Segal [66] show that, under expected as well as non-expected utility maximization, rejecting certain small gambles implies that ridiculously favorable gambles should also be turned down. More precisely, let $V(\widetilde{w})$ denote the value functional for wealth prospects $\widetilde{w}$. If $V(\widetilde{w}_0') \leq V(w_0)$ for a small-stakes gamble $\widetilde{w}_0'$, then there exist very favorable gambles $\widetilde{w_0}''$ for which $V(\widetilde{w}_0'') < V(w_0)$. In fact, this is the case for gambles $\widetilde{w_0}''$ involving small losses, very large winnings, and winning probabilities such that very few people, if any, would reject them. Such predictions contradict plausible patterns of behavior since many people would turn down the small gambles.

Suppose, however, that the agent also faces the problem of allocating her wealth for consumption and savings across a lifetime of T periods. Under the complexity-costs approach, her optimal horizon, $t(\widetilde{w}) \in \{1, ..., T\}$, for the temporal allocation of any wealth prospect $\widetilde{w}$ enters into the value function $V(t(\widetilde{w}), \widetilde{w})$. In what follows, I construct an example in which the optimal $t(\cdot)$ is consistent with rejecting small gambles while accepting sufficiently favorable large ones.

Rabin [58] offers a strong argument for why expected utility theory is unable to provide a plausible account of risk aversion over modest stakes. Turning down a modest-stakes gamble means that the marginal utility of wealth diminishes quickly even for small changes in wealth. His examples indicate that seemingly innocuous risk aversion with respect to small gambles implies an absurd rate of decrease in the marginal utility of wealth. Safra and Segal [66] extend this criticism to several non-expected utility theories.

With complexity costs, two competing effects govern choices: the decreasing marginal utility of wealth, and the fact that a larger wealth level leads to a jump in the value of wealth when it induces a longer optimal horizon. The value of any amount of wealth is higher when this amount gets allocated over a longer horizon, because it is better allocated. The issue that Rabin, Safra and Segal point out applies, other things being equal, within a fixed horizon. The following example suggests that (i) it can be outweighed when the value of money increases sufficiently with the length of the optimal horizon and (ii) the increase in the value of money from extending the horizon increases

with the amount of wealth to be allocated.

Apart from being consistent with contrasting risk attitudes in the small and large, an increasing relation between wealth and the length of the optimal horizon also explains why risk aversion would be, in particular, large in the small but small in the large. The possible realizations of a lottery define different contingencies with respect to the temporal wealth-allocation problem. Gambles with small stakes correspond to contingencies in which the resulting wealth levels are similar enough to the endowment to be allocated under the same horizon as the endowment.

When the stakes are large, however, the contingency of winning big motivates the agent to think further into the future. A longer planning horizon means that a larger fraction of the wealth endowment is allocated optimally and, thus, used to evaluate the prospect. With non-increasing relative risk aversion utility for wealth, this results in risk aversion being smaller in the large than in the small. Moreover, it can produce realistically high degrees of risk aversion in the small even when the endowment is too large for this to obtain under standard settings. This is the case when the complexity costs of wealth allocation are high enough for the agent to be sufficiently myopic in her planning relative to her lifetime. Then, the corresponding part of her lifetime endowment that is used for optimal allocation is small, and her risk aversion large.

In the complexity-costs model, differences in the corresponding optimal horizons can produce sharp contrasts in the way agents evaluate risky prospects in the small and large. This can also account for acceptance of gambles with very large winnings and very small losses, even when the probability of winning is extremely low. That is, it can explain why people purchase lottery tickets. In the standard settings, accepting bets with extremely low expected gains requires convex utility functions for wealth. In the presence of complexity costs, this can result from players' attitudes towards lotto-type bets being more sensitive to the large jackpot than to the minuscule probability of winning. The case of lotto-type bets is examined in Appendix A.2. In what follows, I present a simplified version that focuses on risk aversion in the small and large.

Formally, consider an agent who will lives for 3 periods, denoted by $\tau \in \{1, 2, 3\}$ (the general, T-period version is analyzed in Appendix A.1). She has a wealth endowment w_0 and must allocate her wealth for consumption across her lifetime. Her preferences are represented by a time-separable function, with utility for consumption in period τ given by $u_\tau(c_\tau) = \log c_\tau$. I will assume also that she has a belief regarding the worst possible level of consumption in period τ, given by $\underline{c}_\tau > 0$.[22] To complete the description, any wealth not consumed in period $\tau \in \{1, 2\}$ can be invested in a safe asset that generates a period-τ rate of return, r_τ, while the discount factor between two consecutive periods is $\delta \in (0, 1)$.

Suppose the agent is considering a horizon of length $t \in \{1, 2, 3\}$. Her decision problem for

[22] An obvious interpretation of such a belief is some subsistence level of consumption.

allocating any amount of wealth w is given by

$$V(t,w) = \max_{\mathbf{c}\in\mathbf{C}} \left\{ \sum_{\tau=1}^{t} \delta^{\tau-1} \log c_\tau + \min_{\mathbf{c}'\in\mathbf{C}:\ \mathbf{c}'\underset{t}{\simeq}\mathbf{c}} \sum_{\tau=t+1}^{3} \log c'_\tau \right\}$$

where $\mathbf{C}$ denotes the set of feasible lifetime consumption plans. These are plans $\mathbf{c} = \{c_\tau\}_{\tau=1}^{3}$ satisfying the following system of constraints:

$$s_0 = w, \quad s_\tau = s_{\tau-1}(1 + r_{\tau-1}) - c_\tau, \quad s_\tau \geq 0, \quad c_\tau \geq \underline{c}_\tau \quad \tau \in \{1,2,3\} \tag{10}$$

where s_τ denotes the wealth available at the beginning of period τ while consumption takes place at the beginning of each period. Of course, the initial wealth must suffice for the plan $\{\underline{c}_\tau\}_{\tau=1}^{3}$ of least consumption:[23]

$$w \geq \sum_{\tau=1}^{3} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1 + r_i)^{-1} \tag{11}$$

With respect to foresight t, two lifetime consumption plans $\mathbf{c}, \mathbf{c}'$ are equivalent ($\mathbf{c} \underset{t}{\simeq} \mathbf{c}'$) iff $c_\tau = c'_\tau$ for all $\tau \leq t$. The worst outcome for the remainder of her life beyond the horizon would obtain if she were to consume at $\{\underline{c}_\tau\}_{\tau=t+1}^{3}$. For a horizon of length t, the decision problem can be re-written as follows

$$V(t,w) = \max_{\mathbf{c}\in\mathbf{C}} \left\{ \sum_{\tau=1}^{t} \delta^{\tau-1} \log c_\tau + \sum_{\tau=t+1}^{3} \delta^{\tau} \log \underline{c}_\tau \right\}$$

For $\mathbf{c} \in \mathbf{C}$, let $\mathbf{c}^t = \{c_\tau\}_{\tau=1}^{t}$ denote the restriction of $\mathbf{c}$ to the horizon of length t. If the agent chooses the consumption plan $\mathbf{c}^t$ within the horizon, her available wealth at the end of period t will be

$$s_t(\mathbf{c}^t) = w \prod_{i=1}^{t} (1 + r_i) - \sum_{\tau=1}^{t} c_\tau \prod_{i=\tau}^{t} (1 + r_i) \qquad t \in \{1,2\}$$

Along with the system of constraints in (10), this implies that

$$w = \sum_{\tau=1}^{3} c_\tau \prod_{i=1}^{\tau-1} (1 + r_i)^{-1} \tag{12}$$

is met by any solution $\mathbf{c} \in \mathbf{C}$. Notice also that, since $V(2,w) = V(3,w)$, rational decision-making requires $t \geq 2$.[24]

[23] Abusing notation slightly, in what follows, the term $\prod_{i=1}^{0} x_i$ will denote the number 1, for any $x_i > 0$.

[24] For $t = 2$, a given plan $\{c_1, c_2\}$ chosen for consumption within the horizon corresponds to a unique level of consumption in the final period as determined by (12). Planning for the first two periods implicitly accounts also for consumption in the last period.

Let the agent consider foresight $t = 1$. Given that she anticipates her consumption beyond the horizon to be at the worst level, the choice of period-1 consumption is given trivially by the budget constraint (12):

$$c_1 = w - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)$$

The value of the horizon is

$$V(1,w) = \log\left[w - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)\right] + \delta\log\underline{c}_2 + \delta^2\log\underline{c}_3$$

For $t \geq 2$, the problem becomes a standard one. The optimal consumption vector is $c_1^* = \frac{w}{1+\delta+\delta^2}$, $c_2^* = \delta c_1^*(1+r_1)$, and $c_3^* = \delta^2 c_1^*(1+r_1)(1+r_2)$. The value of the problem is

$$\begin{aligned} V(2,w) &= \left(1+\delta+\delta^2\right)\log w + \delta\log\left[\delta(1+r_1)\right] \\ &\quad - \left(1+\delta+\delta^2\right)\log\left[1+\delta+\delta^2\right] + \delta^2\log\left[\delta^2(1+r_1)(1+r_2)\right] \end{aligned}$$

Consider now the agent facing a lottery $\widetilde{z}$, which would result in her wealth being $\widetilde{w}_0 = w_0 + \widetilde{z}$. For simplicity, I will restrict attention to 50-50 lose $\$l$/gain $\$g$ lotteries ($g > l$), denoted by $\left(g, \frac{1}{2}; -l, \frac{1}{2}\right)$. Under such gambles, the expected values of the horizons are given by

$$\begin{aligned} &\mathbb{E}_{\widetilde{w}_0}\left[V(1,\widetilde{w}_0)\right] \\ = \;& \frac{1}{2}\left(\log\left[w_0 + g - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)\right] + \log\left[w_0 - l - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)\right]\right) \\ &+\delta\left(\log\underline{c}_2 + \delta\log\underline{c}_3\right) \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}_{\widetilde{w}_0}\left[V(2,\widetilde{w}_0)\right] &= \frac{\left(1+\delta+\delta^2\right)}{2}\log\left[(w_0+g)(w_0-l)\right] + \delta\log\left[\delta(1+r_1)\right] \\ &\quad - \left(1+\delta+\delta^2\right)\log\left[1+\delta+\delta^2\right] + \delta^2\log\left[\delta^2(1+r_1)(1+r_2)\right] \end{aligned}$$

Let $t(\widetilde{w}_0)$ and $t(w_0)$ denote the corresponding optimal horizons for the wealth allocation problem under the lottery $\widetilde{w}_0$ and the endowment, respectively. The agent will accept the lottery if and only if $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t(\widetilde{w}_0), \widetilde{w}_0\right)\right] > V\left(t(w_0), w_0\right)$. That is, the relative lengths $t(\widetilde{w}_0)$ and $t(w_0)$ of the optimal horizons matter for the decision of whether or not to accept the lottery. The following two claims establish precisely how they matter.

Claim 3 *For a gamble $\left(g, \frac{1}{2}; -l, \frac{1}{2}\right)$, suppose that the induced optimal horizon is not longer than that induced under the certain prospect, $t(\widetilde{w}_0) \leq t(w_0)$. If*

$$w_0 - \mathbf{1}_{t(\widetilde{w}_0)=1}\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right) \leq \frac{gl}{g-l} \tag{13}$$

then the gamble will be rejected. For $t(\widetilde{w}_0) = t(w_0)$, $\left(g, \frac{1}{2}; -l, \frac{1}{2}\right)$ is rejected only if (13) holds.

Proof. Observe first that $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t,\widetilde{w}_0\right)\right]$ and $V\left(t,w_0\right)$ have all but their respective first terms equal. If both problems induce fully-rational planning, $t\left(\widetilde{w}_0\right)>1$, the claim is immediate by comparing $\mathbb{E}_{\widetilde{w}_0}\left[V\left(2,\widetilde{w}_0\right)\right]$ and $V\left(2,w_0\right)$. Otherwise, we have

$$\begin{aligned}&\left(w_0+g-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)\left(w_0-l-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)\\ =\;&\left(w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)^2+(g-l)\left(w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)-gl\end{aligned}$$

By (13), this implies $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t\left(\widetilde{w}_0\right),\widetilde{w}_0\right)\right]\leq V\left(t\left(\widetilde{w}_0\right),w_0\right)$. Since $t\left(\widetilde{w}_0\right)\leq t\left(w_0\right)$, Lemma 2 implies $V\left(t\left(\widetilde{w}_0\right),w_0\right)\leq V\left(t\left(w_0\right),w_0\right)$. ■

Claim 4 *For a gamble* $\left(g,\frac{1}{2};-l,\frac{1}{2}\right)$*, suppose that it induces a longer optimal horizon than the certain prospect,* $t\left(\widetilde{w}\right)>t\left(w_0\right)$*. If*

$$\left(w_0+g\right)\left(w_0-l\right)>Z\left(\delta\right)^2\left(w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)^2 \tag{14}$$

where

$$Z\left(\delta\right)=\left(1+\delta+\frac{1}{\delta}\right)^{\delta+\delta^2}\left(\frac{1+\delta+\delta^2}{\delta^{\delta^2}}\right)$$

then the gamble will be accepted.

Proof. For the gamble to be accepted, it must be $\mathbb{E}_{\widetilde{w}_0}\left[V\left(2,\widetilde{w}_0\right)\right]>V\left(1,w_0\right)$. By (12), $\widetilde{w}_0>\frac{\underline{c}_2}{1+r_1},\frac{\underline{c}_3}{(1+r_1)(1+r_2)}$, for both realizations of $\widetilde{w}_0$. That is,

$$\begin{aligned}\left(\frac{\delta+\delta^2}{2}\right)\log\left[\left(\widetilde{w}_0+g\right)\left(\widetilde{w}_0-l\right)\right]\;>\;&\delta\log\underline{c}_2+\delta^2\log\underline{c}_3\\ &-\delta\log\left[1+r_1\right]-\delta^2\log\left[\left(1+r_1\right)\left(1+r_2\right)\right]\end{aligned}$$

It suffices, therefore, to have

$$\begin{aligned}&\log\left[\left(w_0+g\right)\left(w_0-l\right)\right]+2\left[\delta\log\delta-\left(1+\delta+\delta^2\right)\log\left[1+\delta+\delta^2\right]+\delta^2\log\delta^2\right]\\ >\;&\log\left[w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]^2\end{aligned}$$

which is equivalent to (14). ■

For complexity costs to account for different risk attitudes in the small and large, it must be that small gambles correspond to shorter optimal horizons than large ones. In what follows, I present sufficient conditions for this to be the case. Specifically, these conditions guarantee that the optimal horizon induced by the certain prospect, w_0, is at least as long as that for a small gamble

but shorter than that for a large one. To facilitate comparisons of optimal foresight, observe that

$$\begin{aligned}
& 2\left(\mathbb{E}_{\widetilde{w}_0}\left[\Delta_t V\left(1,\widetilde{w}_0\right)\right]-\Delta_t V\left(1,w_0\right)\right) \\
= \;& 2\mathbb{E}_{\widetilde{w}_0}\left[V\left(2,\widetilde{w}_0\right)-V\left(1,\widetilde{w}_0\right)\right]-2\left[V\left(2,w_0\right)-V\left(1,w_0\right)\right] \\
= \;& \left(1+\delta+\delta^2\right)\log\left[\frac{\left(w_0+g\right)\left(w_0-l\right)}{w_0^2}\right] \\
& -\log\left[\frac{\left(w_0+g-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)\left(w_0-l-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)}{\left(w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right)^2}\right]
\end{aligned} \tag{15}$$

Before proceeding, I must specify what are the costs of complexity in the context of this application as well as what gambles should be regarded as small in the presence of complexity costs. The complexity is that of the decision problem of allocating any given amount of wealth w for consumption and savings across the agent's lifetime. For a horizon $t \in \{1,2,3\}$ in the allocation problem, the complexity costs will be given by some function $C\left(t,w\right)$.

When considering a gamble, the agent must address the allocation problem for each amount of wealth resulting from the possible realizations of the gamble. In this context, it is natural to regard as small any gamble such that all of its possible realizations are small enough for the problem of allocating the corresponding wealth to not differ, in terms of complexity, from that of allocating the endowment.[25]

Definition 6 *A possible realization z of a lottery $\widetilde{z}$ is a **small stake** if*

$$C\left(t,w_0+z\right)=C\left(t,w_0\right) \quad t\in\{1,2\}$$

*A lottery $\widetilde{z}$ is **small** if every z in the support of $\widetilde{z}$ is a small stake.*

This definition requires that the complexity-cost function is constant with respect to wealth, for a given foresight t, over some range $(w_0-M,w_0+M):M>0$ around the endowment level $(,z,\leq M$ for a small stake). Equivalently,

$$\frac{g\left(t,w_0+\widetilde{z}\right)}{g\left(t,w_0\right)}=\frac{V\left(t,w_0\right)}{\mathbb{E}_{w_0+\widetilde{z}}\left[V\left(t,w_0+\widetilde{z}\right)\right]} \quad t\in\{1,2\} \quad \text{for small } \widetilde{z} \tag{16}$$

By (15), the following condition

$$\begin{aligned}
& \left[\frac{w_0^2}{\left(w_0+g'\right)\left(w_0-l'\right)}\right]^{1+\delta+\delta^2} \geq \\
& \frac{\left[w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]^2}{\left[w_0+g'-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]\left[w_0-l'-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]}
\end{aligned} \tag{17}$$

[25]The definition does not specify a relation for $t=3$ as this would be redundant. Since $V\left(2,w\right)=V\left(3,w\right)$, by Lemma 2, $C\left(2,w\right)=C\left(3,w\right)$.

is equivalent to $\mathbb{E}_{\widetilde{w}_0'}\left[\Delta_t V\left(1,\widetilde{w}_0'\right)\right] \leq \Delta_t V\left(1,w_0\right)$ and, thus, implies

$$\mathbb{E}_{\widetilde{w}_0'}\left[\Delta_t V\left(1,\widetilde{w}_0'\right)\right] - \Delta_t C\left(1,\widetilde{w}_0'\right) \leq \Delta_t V\left(1,w_0\right) - \Delta_t C\left(1,w_0\right)$$

for a small gamble.[26] Under a specification for the function $g\left(t,\cdot\right)$ that meets the requirements of Corollary 2, such as the one given by (7), (17) suffices for a small gamble $\left(g',\frac{1}{2};-l',\frac{1}{2}\right)$ to give $t\left(\widetilde{w}_0'\right) \leq t\left(w_0\right)$ and, thus, for Claim 4 to establish then a rejection condition.

Let us now turn to large gambles. Since the focus of this application is on explaining why very favorable gambles could be accepted even though small gambles are rejected, I restrict attention to gambles offering large gains but small losses: $\left(g'',\frac{1}{2};-l'',\frac{1}{2}\right) : g'' > l''$ where l'' is a small stake but g'' is not. The complexity costs argument requires that such gambles induce longer optimal horizons than small ones. If (17) holds for a small gamble, this is the case when the optimal horizon under the large gamble corresponds to rational planning, $t\left(\widetilde{w}_0''\right) = 2$, while that induced under the endowment does not, $t\left(w_0\right) = 1$. Claim 4 gives then a sufficient condition for the large gamble to be accepted. Hence, we require that the expected marginal benefit of extending the horizon from $t = 1$ to $t = 2$, under the large gamble, exceeds the corresponding marginal benefit under the certain prospect: $\mathbb{E}_{\widetilde{w}_0''}\left[\Delta_t V\left(1,\widetilde{w}_0''\right)\right] > \Delta_t V\left(1,w_0\right)$.[27] For this, it suffices that

$$\left[\frac{w_0^{1+\delta+\delta^2}}{w_0 - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)}\right]^2 < \frac{\left[\left(w_0+g''\right)\left(w_0-l''\right)\right]^{1+\delta+\delta^2}}{\left[w_0+g''-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]\left[w_0-l''-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)\right]} \tag{18}$$

In this inequality, only the right-hand side depends upon the realizations of the gamble. Since $\frac{\left(w_0-l''\right)^{1+\delta+\delta^2}}{w_0-l''-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)} > \left(w_0-l''\right)^{\delta+\delta^2}$ and the gamble involves small losses, $l'' \leq M$, the right-hand side is bounded below by $\underline{G} = \frac{\left(w_0+g''\right)\left[\left(w_0+g''\right)\left(w_0-M\right)\right]^{\delta+\delta^2}}{w_0+g''-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)}$. Yet, the gains are not small and $\lim_{g''\to+\infty}\underline{G} = +\infty$. Clearly, (18) holds for gambles with small losses but sufficiently large gains.[28]

[26] Let $x = \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)$ be the present value of the worst possible future consumption stream. As $x \to w_0$, the right-hand side of (i) tends to 0^+ while the strictly positive left-hand side remains unchanged. That is, (i) holds, for any gamble $(g',\frac{1}{2};-l',\frac{1}{2})$, as long as x is sufficiently close to the initial endowment w_0.

[27] Let the myopic foresight be optimal under the certain prospect $\left(t\left(w_0\right)=1\right)$ and $\mathbb{E}_{\widetilde{w}_0''}\left[V\left(1,\widetilde{w}_0''\right)\right] < V\left(1,w_0\right)$ (the large gamble is rejected under myopic planning). Set $\frac{g\left(2,\widetilde{w}_0''\right)}{g\left(1,\widetilde{w}_0''\right)} = \frac{g\left(2,w_0\right)}{g\left(1,w_0\right)} = \frac{V\left(2,w_0\right)}{V\left(1,w_0\right)}$. Then, $\mathbb{E}_{\widetilde{w}_0''}\left[\Delta_t V\left(1;\widetilde{w}_0''\right)\right] > \Delta_t V\left(1,w_0\right)$ implies $\frac{\mathbb{E}_{\widetilde{w}_0''}\left[V\left(2,\widetilde{w}_0''\right)\right]}{\mathbb{E}_{\widetilde{w}_0''}\left[V\left(1,\widetilde{w}_0''\right)\right]} > \frac{V\left(2,w_0\right)}{V\left(1,w_0\right)} = \frac{g\left(2,\widetilde{w}_0''\right)}{g\left(1,\widetilde{w}_0''\right)}$. By Proposition 2(i), the large gamble induces rational planning: $t\left(\widetilde{w}_0''\right) > 1$.

[28] My analysis here is about the expected marginal benefit of extending the planning horizon in absolute terms, not relative to wealth. It is important to emphasize this as one could argue that the smaller the agent's wealth, the more important it is that she allocates it optimally; hence, the more important that her horizon is long. This is a statement as to why the expected marginal benefit of further planning *relative* to wealth should be falling with wealth.

Conditions (17) and (18) will hold under plausible assumptions; namely, if the large gamble offers sufficiently large winning stakes and the present value of the worst possible future consumption stream is sufficiently close to the agent's wealth endowment. Claims 2 and 3 allow then for a large favorable gamble to be accepted when a small one is turned down. To illustrate, consider a function $g(t,\cdot)$ such that the problem of allocating an initial endowment w_0 for consumption and savings across one's lifetime is sufficiently complex to induce myopic planning: $t(w_0)=1$. Let $\left(g',\frac{1}{2};-l',\frac{1}{2}\right)$ be a small gamble that satisfies (i), inducing, therefore, the same myopic horizon as the certain prospect, $t(\widetilde{w}_0')=t(w_0)=1$. If (13) holds, this gamble ought to be rejected. Let also $\left(g'',\frac{1}{2};-l'',\frac{1}{2}\right)$ be a gamble with small losses but sufficiently large gains for (ii) to hold, inducing, thus, rational planning, $t(\widetilde{w}_0'')=2$. In this case, (14) suffices for the large gamble to be accepted. Observe also that, since (13) holds, an even stronger acceptance condition is given by

$$\left(w_0+g''\right)\left(w_0-l''\right)>Z(\delta)^2\left(\frac{g'l'}{g'-l'}\right)^2 \tag{19}$$

Clearly, $G=3^6\left(\frac{g'l'}{g'-l'}\right)^2$ is an upper bound for the quantity on the right-hand side of (19).[29]

Rabin [58] presents values for g'', l'', g', l', and w_0 such that many people would turn down the small bet but few - if any - would reject the large one (Tables I and II; the discussion in pp. 1282-1285). For example, let $g''=635,670$, $l''=4,000$, $g'=105$, $l'=100$, and $w_0=340,000$ or $g''=36\times10^9$, $l''=600$, $g'=125$, $l'=100$, and $w_0=290,000$. Observe that

$$\frac{g''l''}{g''-l''}>\frac{g'l'}{g'-l'} \tag{20}$$

holds in both cases. If (13) obtains for the small gamble, it will be rejected under the horizon of length $t=1$ and so will be the large gamble. However, (19) is also satisfied in both cases. The corresponding upper bounds for its right-hand side are $G=32.15\times10^8$ and $G=1.8\times10^8$ whereas the values for its left-hand side are given by 32.78×10^{10} and 1.04×10^{16}, respectively. Hence, the large gambles will be accepted under the horizon of length $t=2$. Similar results can be obtained for all large bets in Rabin's Table II.[30]

For a small gamble to be turned down, condition (13) requires that the part of the endowment w_0 which is allocated optimally under the myopic horizon, $w_0-\left(\frac{1}{1+r_1}\right)\left(\underline{c}_2+\frac{\underline{c}_3}{1+r_2}\right)$, is small. This

[29]The quantity $Z(\delta)$ was defined in the proof of Claim 2. We have

$$\begin{aligned}\log Z(\delta) &= \left(1+\delta+\delta^2\right)\log\left[1+\delta+\delta^2\right]-\left(\delta+2\delta^2\right)\log\delta\\ &< \left(1+\delta+\delta^2\right)\left(\log\left[1+\delta+\delta^2\right]-\log\delta\right)=\left(1+\delta+\delta^2\right)\log\left[1+\delta+\delta^{-1}\right]\end{aligned}$$

It is trivial to check that the last quantity above is increasing in δ; $3\log3$ is an upper bound for $\log Z(\delta)$.

[30]All bets in that table qualify for the sufficient condition. Yet, only the ones with gains large enough to induce long optimal horizons should be taken in consideration. Notice also that, with respect to Table I, my predictions hold for large gambles and for sufficiently large levels of initial wealth. For example, with $w_0<20,000$, large bets will be turned down even though they may induce longer horizons than small bets.

is in contrast to the standard expected log-utility setting, which requires the entire endowment to be small.[31] For the two small gambles above to be rejected, it must be that not more than \$2100 and \$500, respectively, are allocated optimally. Of course, these amounts may be unrealistically small but the two sides of the acceptance condition (19) differ by several orders of magnitude. Hence, the examples do make the point that the model offers considerable room for maneuver in the selection of gambles that produce the desired contrasting risk attitudes in the small and large.

Another possible response to Rabin's argument is that decision-makers do not reject small lotteries even when their expected payoff is almost zero. My analysis abstracts from this issue by assuming sufficient conditions so that small gambles are rejected. This allows the exposition to focus on the essence of Rabin's criticism of expected utility: the rapid deterioration in the value of wealth that rejecting small gambles reveals. For the log-utility specification, this is highlighted by condition (20). For any w_0, once $\left(g', \frac{1}{2}; -l', \frac{1}{2}\right)$ is rejected over a horizon of length $t \in \{1, 2\}$, any gamble $\left(g'', \frac{1}{2}; -l'', \frac{1}{2}\right)$ satisfying (20) will also be turned down over the same horizon, no matter how favorable. In the previous numerical example, (20) holds and both gambles are rejected under myopic planning, $t = 1$, but the large gamble is ultimately accepted if it induces a longer optimal horizon. The example used the acceptance condition (19) which is not compatible with the small gamble being rejected under the standard setting when (20) holds. (14), however, is weaker than (19) allowing for $w_0^2 > (w_0 + g'')(w_0 - l'') > Z(\delta)^2 \left(w_0 - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)\right)^2$. Thus, the model can also accommodate pairs of gambles such that both are turned down under the standard setting. In this case, both gambles will be rejected under myopic planning but the large one will be accepted if it induces an optimal horizon of length $t = 2$.

In the complexity-costs model, under the myopic horizon $t = 1$, the amount of wealth the agent is using to evaluate risky prospects consists of the portion of her endowment that she intends to spend within her horizon, namely $w_0 - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)$. The analysis indicates that the smaller this part, other things remaining unchanged, the more likely that small gambles will be turned down and large gambles will be accepted - recall conditions (13) and (14). This is a qualitative prediction regarding the heterogeneity of risk attitudes in the small and large across agents: controlling for initial endowments, *the more optimistic is an agent about her worst possible future consumption stream, the more likely that her risk attitudes will differ in the small and large.* The statement generalizes to the T-period version of the model where, for a given initial endowment w_0, the more myopic the agent's planning horizon relative to her lifetime, the more likely that conditions (25), (26) and (28) of Appendix A.1 will hold.

Appendix A.2 extends the model to lotto-type large gambles, offering very large winnings, very small losses and minuscule winning probabilities. I show that, for sufficiently large amounts of wealth, the resulting jump in welfare from extending the allocation horizon outweighs the deterioration in the marginal value of wealth even under extremely small probabilities of winning. As

[31] A gamble $\left(g, \frac{1}{2}; -l, \frac{1}{2}\right)$ is rejected, under the standard setting, if $w_0 \leq \frac{gl}{g-l}$.

in this example, my argument depends on the assumption that considering how to allocate one's wealth for consumption in the winning contingency corresponds to a marginal benefit from extending the allocation horizon that outweighs the marginal complexity costs, inducing a longer optimal horizon than that under the endowment. My intuition, that accepting lotto-type gambles is fundamentally based upon inducing one to think how to allocate the winnings, seems to be in agreement with common marketing practises for such products. Large-prize lotteries are almost always marketed by pointing out various ways the prize could be spent. Advertisements usually describe luxury homes, boats, vacations and retirement locations that could be afforded with the winnings before inviting one to play. Actually, in exactly this spirit of getting one thinking about the contingency of winning, the National Lottery Organization of Greece ran a very successful advertising campaign in 2004-05 with the slogan "But if you win?".

The prediction given above remains valid for the analysis of Appendix A.2: the complexity-costs model asserts that, other things being equal, more myopic agents are more likely to accept lotto-type bets. Moreover, *poorer agents (with smaller initial endowments) are more likely to accept lotto-type bets*, other things remaining unchanged - see conditions (30), (A.2.ii) and the subsequent discussion. Observe also that (14) becomes easier to satisfy as the gains g'' increase while losses remain small. This result, that higher potential winnings make it more likely that the bet will be taken other things remaining unchanged, holds also for the lotto-type gambles. Cook and Clotfelter [15] show that, across states in the US, lottery-ticket sales are strongly correlated with the size of the state's population, which is in turn positively correlated with the size of the jackpot. Within a state, ticket sales each week are strongly positively correlated with the size of the rollover.

A natural conclusion from Rabin's criticism is that expected utility should be replaced with more general theories, especially those that exhibit first order risk aversion, those based on uncertainty aversion or those derived from preferences over gains and losses rather than final wealth levels. As observed first by Samuelson [68], expected-utility theory makes also the powerful prediction that agents don't see an amalgamation of independent gambles as significant insurance against the risk of those gambles. He showed that, under expected utility, if (for some sufficiently wide range of initial wealth levels) someone turns down a particular gamble, she should also turn down an offer to play $n > 1$ independent plays of this gamble. In his example, if Samuelson's colleague is unwilling to accept a $\left(100, \frac{1}{2}; -200, \frac{1}{2}\right)$ bet, he should be unwilling to accept 100 of those taken together. Samuelson's paradox is weaker than Rabin's but it makes the point that adding together many independent risks should not alter attitudes towards those risks in any meaningful way for an expected-utility maximizer. Although rather cumbersome mathematically, it is easy to replace the large 50-50 gamble here by n small ones taken together. For large enough n, this is again a large stakes bet and it will not be turned down if it induces a sufficiently long planning horizon.

Rabin [58] and Rabin and Thaler [59] argue that the way to resolve these paradoxes is loss aversion. The calibration results of Safra and Segal [66], however, apply to a large collection of non-expected utility theories while both the Rabin and the Samuelson paradoxes can be reframed

to apply to pure gains.[32] In other words, there seems to be a fundamental discrepancy in the way agents evaluate risky prospects in the small and large, irrespective of the underlying preferences over final outcomes. Indeed, the problems with assuming that risk attitudes over modest and large stakes derive from the same utility function relate to a long-standing debate in economics. Here, the discrepancy arises using the same, simplistic underlying preferences in the small and large. Instead, my argument for why an agent considers a different objective function to evaluate small and large risky prospects derives from considering different horizons for the allocation problem. This results purely from the interaction between the costs and benefits of further search. As long as complexity costs are present, it would apply whatever the valuation functional $V(t, w)$ for the underlying preferences over final wealth outcomes.

Another example of treating differently small and large bets in an intertemporal consumption-savings problem is suggested by Fudenberg and Levine [22]. They model a "dual-self" agent, consisting of a patient long-run self with infinitely-long time-horizon, and a sequence of myopic short-run selves who live for only one period. The long-run self wants to restrain the myopic self from overspending and a commitment device is available in the form of a cash-in-advance constraint. That is, the long-run self decides on how much "pocket cash" to allow the short-run self to have for allocation between consumption and savings during the period. Differences in risk attitudes for small and large bets arise due to a wedge between the propensity to consume out of pocket cash and out of wealth. Winnings from sufficiently small gambles are spent in their entirety, and so are evaluated according to the short-run self's preferences. But when the stakes are large, self-restraint kicks in because part of the winnings will be saved and spread over the lifetime. Since small and large gambles are evaluated relative to pocket money and wealth, respectively, under constant relative risk aversion, the "dual-self" decision-maker appears to be less risk averse when facing large gambles than when facing small ones.

This intuition is close to the one arising from the complexity costs. Both approaches attribute differences in risk attitudes to differences in the amount of wealth used as reference to evaluate a given gamble. For small gambles, this is restricted to be pocket-cash and $w_0 - \left(\frac{1}{1+r_1}\right)\left(\underline{c}_2 + \frac{\underline{c}_3}{1+r_2}\right)$, respectively. As with the pocket-cash interpretation, the smaller is the fraction of the endowment allocated optimally under the certain prospect, the more likely that large bets are accepted even though small ones are not. If the current endowment w_0 is interpreted as the agent's lifetime income, the fraction in question would be the part that is currently available. That is, given one's lifetime income, the more cash constrained the agent is the more pronounced the divergence in her risk attitudes in the small and large. In the general T-period version, this is the case the more

[32] For the Samuelson paradox, consider an agent who (over a sufficiently wide range of initial wealth levels) turns down a 50-50 win \$300 or nothing bet for the certain prospect of \$100. Then she should also prefer \$100n to n plays of the gamble. But as n increases, the expected payoff from the risky prospect becomes infinitely larger than that of the certain one while the probability of winning more by choosing it approaches 1. For the Rabin paradox, if an agent prefers the certain prospect of \$$s$ to a 50-50 bet of winning \$$g$ or nothing, she should also choose \$$\frac{s^2}{g-2s}$ over a 50-50 bet of winning infinity or nothing.

myopic is the agent's planning in allocating her endowment.

Chetty and Szeidl [13] offer another interesting analysis of risk preferences. Their model has two consumption goods, one of which involves a commitment in the sense that an adjustment cost must be paid whenever the per-period quantity consumed is changed. Consumption commitments affect risk preferences by amplifying risk aversion against moderate-scale shocks and by creating a certain gambling motive. When the agent cannot adjust commitment consumption easily, small shocks to wealth must be handled by changing adjustable consumption; the marginal utility of wealth increases more rapidly when the agent has commitments. The gambling motive obtains when the expected decrease in utility from the adjustable good in the case of a loss is offset by the large expected increase in utility from changing commitment consumption under a gain.

Their analysis offers a convincing rationalization of why people buy insurance against moderate shocks, change their consumption patterns significantly in the face of unemployment shocks when these lead to large income losses but not when they lead to small ones, or why wealthier individuals are less likely to play the lottery. An important limitation, however, is approximate risk neutrality for small stakes (see their Table II). High risk aversion over moderate stakes offers an explanation for why moderate-stakes gambles would be rejected. It does not explain why risk attitudes should differ in the small and large for a given wealth endowment. Moreover, the gambling motive obtains only for small gambles and for initial wealth levels around the points where the agent is indifferent between adjusting or not commitment consumption.[33] Rather than explaining why large-stakes gambles would be accepted, this suggests that small gambles shouldn't be rejected.

5 Concluding Remarks

This analysis developed a model of bounded rationality due to search costs arising from the complexity of the decision problem. By inferring these costs from revealed preferences through choices, bounded rationality and its extent emerge endogenously. Revealed bounded rationality encompasses variation across decision-makers as well as decision problems. Under additional assumptions, calibration of search costs suggests predictions and testable implications of the model. Applications to seemingly disparate problems illustrate the flexibility of this approach. Endogenous complexity costs can be consistent with violations of timing independence in temporal framing problems, dynamic inconsistency and diversification bias in sequential versus simultaneous choice problems, and with plausible but contrasting risk attitudes in small- and large-stakes gambles.

My approach is based fundamentally on using the behavior under limited foresight, depicted by the representation in (2), as a common framework for comparing the revealed preferences from observed boundedly-rational choices to those prescribed by the rational paradigm. For this to be a complete model of bounded rationality, however, it requires a theoretical foundation, an

[33] See their Proposition 1(iii). Since the value function is locally convex only at these points, the wealth outcomes need to remain within the corresponding neighborhoods.

axiomatization of the actual preferences that lie behind the representations in (2) and (3). Current work in progress is attempting to exploit the insights of Ghirardato [28] and Ahn and Ergin [2] to this end. The former paper seems to provide promising avenue for further research in constructing preferences over unforeseen contingencies that could be represented by the utility function in (2). The latter develops a theoretical framework for studying how preferences might vary as the way or extent to which the future might be unforeseen varies.

Another obvious next for further research would be to extend this approach to strategic games. That is, to attempt study extend this approach to strategic games. Current work in progress establishes existence of equilibrium (in the standard sense of no player having an incentive to deviate when her conjectures happen to be correct) for extensive-form games when players are boundedly-rational. Since, however, the standard notion of equilibrium is problematic with currently unforeseen contingencies, I examine which conjectures can be justifiable under limited foresight. The fact that the degree of limited foresight obtains endogenously implies a relation between choice of foresight and justifiable conjectures that can be viewed as the bounded rationality analogue to rationalizability. The model can explain some game-theoretical paradoxes in which observed or even intuitively correct responses are at odds with the predictions of backward induction. I show that the set of "rationalizable"-under limited-foresight strategies includes ones consistent with observed play.

As a tribute to those most commonly observed humans who regularly demonstrate impressive ingenuity yet even more often prove themselves capable of rather foolish acts, this research attempted to explain why, in the latter situations, a frustrated observer would exclaim "you were *not* thinking". The sequel papers are looking for an answer to the observer's question "*what* were you thinking?".

References

[1] Abreu D. and A. Rubinstein (1988): "The Structure of Nash Equilibria in Repeated Games with Finite Automata", *Econometrica*, 56(6):1259-82.

[2] Ahn D.S. and H. Ergin: "Framing Contingencies", Preprint(2007)

[3] Anderlini L. and L. Felli (2004): "Bounded Rationality and Incomplete Contracts", *Research in Economics*, 58(1):3-30.

[4] Benartzi S. and R.H. Thaler (2001): "Naive Diversificaton Strategies in Defined Contribution Savings Plans", *American Economic Review*, 91(1):79-98.

[5] Bewley T.F. (2002): "Knightian Uncertainty Part I" *Decisions in Economics and Finance*, 25(2):79-110.

[6] Bewley T.F. (1987): "Knightian Uncertainty Part II: Intertemporal Problems", *Cowles Foundation Discussion Paper* No.835, Yale University.

[7] Binmore K., J. McCarthy, G. Ponti, L. Samuelson, and A. Shaked (20004): "A Backward Induction Experiment", *Journal of Economic Theory*, 104(1):48-88.

[8] Bone J.D., J.D. Hey, and J.R. Suckling (2003): "Do People Plan Ahead?", *Applied Economic Letters*, 10(5):277-80.

[9] Camerer C., T.H. Ho and J.K. Chong (2004): "A Cognitive Hierarchy Model of Games", *Quarterly Journal of Economics*, 119(3):861-98.

[10] Carbone E., and Hey J.A. (2001): "A Test of the Principle of Optimality", *Theory and Decision*, 50(3):263-81.

[11] Charness G. and Rabin M. (2002): "Understanding Social Preferences with Simple Tests", *The Quarterly Journal of Economics*, 117(3):817-69.

[12] Chatterjee K. and H. Sabourian (2000): "Multiperson Bargaining and Strategic Complexity", *Econometrica*, 68(6):1491-1509.

[13] Chetty R. and A. Szeidl (2006): "Consumption Committments and Risk Preferences", *Quarterly Journal of Economics*, forthcoming(2006)

[14] Conlisk J. (1996): "Why Bounded Rationality?", *Journal of Economic Literature*, 34(2):669-700.

[15] Cook P.J. and C.T. Clotfelter (1993): "The Peculiar Scale Economies of Lotto", *American Economic Review*, 83(3):634-43.

[16] Crawford V.P. and N. Iriberri: "Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?", *Econometrica*, (forthcoming).

[17] Cubitt R.P., C. Starmer, and R. Sugden (1998): "Dynamic Choice and the Common Ratio: an Experimental Investigation", *The Economic Journal*, 108:1362-80.

[18] Diasakos T.: "An Embeding of Several Well-known Models of Bounded Rationality", Preprint (2008).

[19] Eboli M. (2003): "Two Models of Information Costs Based on Computational Complexity", *Computational Economics*, 21:87-105.

[20] Eliaz K. (2003): "Nash Equilibrium when Players account for the Complexity of their Forecasts", *Games and Economic Behavior*, 44(2):286-310.

[21] Forsythe R., Horowitz J., Savin N.E., and M. Sefron (1994): "Fairness in Simple Bargaining Experiments", *Games and Economic Behavior*, 6(3):347-69

[22] Fudenberg D. and D.K. Levine (2006): "A Dual Self Model of Impulse Control", *American Economic Review*, 96(5):1449-76.

[23] Fudenberg D. and J. Tirole: *Game Theory*, MIT Press (1993)

[24] Gabaix X., D. Laibson, G. Moloche, and S. Weinberg (2006): "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model", *American Economic Review*, 96(4):1043-68.

[25] Gabaix X. and D. Laibson: "Bounded Rationality and Directed Cognition", Preprint(2005)

[26] Gale D. and H. Sabourian (2005): "Complexity and Competition", *Econometrica*, 73(3):739-69.

[27] Geanakoplos J. and L. Gray (1991): "When Seeing Further is Not Seeing Better", *Bulletin of the Santa Fe Institute*, 6(2).

[28] Ghirardato P. (2001): "Coping with ignorance: unforeseen contingencies and non-additive uncertainty", *Economic Theory*, 17:247-76.

[29] Gigerenzer G. (2001): "Rethinking Rationality", in Gigerenzer G. and R. Selten (eds): *Bounded Rationality: The Adaptive Toolbox*, MIT Press.

[30] Gilboa I. and D. Schmeidler (1989): "Maxmin Expected Utility with a Non-unique Prior", *Journal of Mathematical Economics*, 18:141-53.

[31] Goldman S.M. (1980): "Consistent Plans", *Review of Economic Studies*, 47:533-7.

[32] ————- (1979): "Intertemporally Inconsistent Preferences and the Rate of Consumption", *Econometrica*, 47:621-6.

[33] ————- (1969): "Sequential Planning and Continual Planning Revision", *Journal of Political Economy*, 77(4):653-64.

[34] ————- (1968): "Optimal Growth and Continual Planning Revision", *Review of Economic Studies*, 35(2):145-54.

[35] Henrich J., Boyd R., Bowles S., Camerer C., Gintis H., McElreath R., and E. Fehr (2001): "In Search of Homo Economicus: Experiments in 15 Small-scale Societies" *American Economic Review*, 91(2):73-9.

[36] Hey J.D. (2005): "Do People (Want to) Plan?", *Scottish Journal of Political Economy*, 52(1):122-38.

[37] Hey J.D. and J. Lee (2005): "Do Subjects Separate (or Are they Sophisticated)?", *Experimental Economics*, 8(3):233-65.

[38] Hey J.D. and M. Paradiso (2006): "Preferences over Temporal Frames in Dynamic Decision Problems: An Experimental Investigation", *The Manchester School*, 74(2):123-37.

[39] Jehiel P. (2001): "Limited Foresight May Force Cooperation", *Review of Economic Studies*, 68(2):369-91.

[40] ——— (1998): "Repeated Games and Limited Forecasting", *European Economic Review*, 42(3-5):543-51.

[41] ——— (1998): "Learning to Play Limited Forecast Equilibria", *Games and Economic Behavior*, 22(2):274-98.

[42] ——— (1995): "Limited Horizon Forecast in Repeated Alternate Games", *Journal of Economic Theory*, 67(2):497-519.

[43] Kahneman D. and A. Tversky (2000): *Choices, Values, and Frames*, Cambridge University Press.

[44] Kahneman D. and A. Tversky (1979): "Prospect Theory: an Analysis of Decision under Risk", *Econometrica*, 47(2):263-91.

[45] Kalai E. and W. Stanford (1988): "Finite Rationality and Interpersonal Complexity in Repeated Games" *Econometrica*, 56(2):397-410.

[46] Karni E. and Z. Safra (1990): "Behaviorally Consistent Optimal Stopping Rules", Journal of Economic Theory, 51(2):391-402.

[47] Koszegi B. and M. Rabin (2006): "A Model of Reference-Dependent Preferences", *Quarterly Journal of Economics*, 121(4):1133-65.

[48] Lee J. and H. Sabourian (2006): "Coase Theorem, Complexity and Transaction Costs", *Journal of Economic Theory*, 135(1):214-35.

[49] Loewenstein G. and J. Elster (ed.): *Choice over Time*, Russell Sage, New York (1992)

[50] Machina M.J. (1989): "Dynamic Consistency and Non-expected Utility Models of Choice under Uncertainty", *Journal of Economic Literature*, 27:1622-28.

[51] MacLeod W.B. (2002): "Complexity, Bounded Rationality and Heuristic Search", *Contributions to Economic Analysis and Policy* 1(1)#8, B.E. Journals in Economic Analysis & Policy, Berkeley Electronic Press.

[52] Neyman A. (1998): "Finitely Repeated Games with Finite Automata", *Mathematics of Operations Research*, 23(3):513-52.

[53] Neyman A. (1985): "Bounded Complexity Justifies Cooperation in Finitely Repeated Prisoners' Dilemma", *Economic Letters*, 19:227-29.

[54] Osborne M.J. and A. Rubinstein (1998): "Games with Procedurally Rational Players", *American Economic Review*, 88(4):834-47.

[55] Peleg B. and M.E. Yaari (1973): "On the Existence of a Consistent Course of Action when Tastes are Changing", *Review of Economic Studies*, 40:391-401.

[56] Piccione M. and A. Rubinstein (1993): "Finite Automata Play a Repeated Extensive-Form Game", *Journal of Economic Theory*, 61(1):160-68.

[57] Pollak R.A. (1968): "Consistent Planning", *Review of Economic Studies*, 35:201-8.

[58] Rabin M. (2000): "Risk Aversion and Expected Utility Theory: A Calibration Theorem", *Econometrica*, 68(5):1281-92.

[59] Rabin M. and R. Thaler (2001): "Anomalies: Risk Aversion", *Journal of Economic Perspectives*, 15(1):219-32.

[60] Read D. and G. Loewenstein (1995): "Diversification Bias: Explaining the Discrepancy in Variety-seeking between Combined and Separate Choices", *Journal of Experimental Psychology: Applied*, 1:34-49.

[61] Read D., Loewenstein G., and M. Rabin (1999): "Choice Bracketing", *Journal of Risk and Uncertainty*, 19: 171-97.

[62] Rosenthal R. (1989): "A Bounded-Rationality Approach to the Study of Noncooperative Games", *Int. Journal of Game Theory*, 18(3):273-92

[63] Rubinstein A.: *Modelling Bounded Rationality*, MIT Press (1998)

[64] ———— (1986) "Finite Automata Play the Repeated Prisoners' Dilemma", *Journal of Economic Theory*, 39(1):83-96

[65] Sabourian H. (2004): "Bargaining and Markets: Complexity and the Competitive Outcome", *Journal of Economic Theory*, 116(2):189-228

[66] Safra Z. and U. Segal: "Calibration Results for Non-Expected Utility Theories", Preprint(2006).

[67] Sargent T.J.: *Bounded Rationality in Macroeconomics*, Oxford University Press (1993).

[68] Samuelson P. (1963): "Risk and Uncertainty: A Fallacy of Large Numbers", *Scientia*, 98:108-13.

[69] Segal U. (1987): "The Ellsberg Paradox and Risk Aversion: an Unanticipated Utility Approach", *International Economic Review*, 28:175-202.

[70] Selten R. (2001): "What is Bounded Rationality?", in Gigerenzer G. and R. Selten (eds): *Bounded Rationality: The Adaptive Toolbox*, MIT Press.

[71] Simon H.A. (1957): *Models of Man* Wiley, New York.

[72] ———- (1956): "Rational Choice and the Structure of the Environment", *Psychological Review*, 63:129-38.

[73] ———-(1955): A Behavioral Model of Rational Choice, *Quarterly Journal of Economics*, 69(1):99-118.

[74] Simonson I. (1990): "The Effect of Purchase Quantity and Timing on Variety-Seeking Behavior", *Journal of Marketing Research*, (27):150-62.

[75] Stigler G.J. (1961): "The Economics of Information", *Journal of Political Economy*, 69(3):213-25.

[76] Strotz R.H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies*, 23(3):165-80.

[77] Thaler R.H. (1999): "Mental Accounting Matters", *J. of Behavioral Decision Making*, 12:183-206.

[78] Thaler R. (ed.): *The Winner's Curse: Paradoxes and Anomalies of Economic Life*, Free Press, New York, (1992)

[79] Thaler R. (ed.): *Quasi-Rational Economics*, Russell Sage, New York, (1991)

A Notation

- $\mathbb{N} \setminus \{0\}$: the set of positive integers.
- $\Gamma = \langle \mathcal{D}, \mathcal{A} \rangle$: a directed graph.

- $\mathcal{D}$: the set of nodes in the graph
- $\mathcal{A}$: the set of arcs in the graph. The elements of $\mathcal{A}$ correspond injectively to pairs of nodes $(d, d') : d \neq d'$, for $d, d' \in \mathcal{D}$.
- $(d, d') \in \mathcal{A}$: d' is an immediate successor of d.
- Ω: the set of terminal nodes
- $\mathcal{H}$: the collection of the sets of nodes at which the same player chooses from the same set of actions. It forms a partition of $\mathcal{D} \setminus \Omega$ into information sets.
- $H \subset \mathcal{H}$: the collection of information sets where the agent acts.
- $\mathcal{H} \setminus H$: the collection of Nature's chance nodes.
- $D(h)$: the collection of nodes comprising the information set $h \in \mathcal{H}$.
- $d \rhd h$: the information set h is an immediate successor of the node d.
- $h \rhd h'$: h' is an immediate successor of h.
- $h(t) = \bigcup_{d \in D(h)} \bigcup_{k \in \{1,\ldots,t\}} \left\{ d \rhd \{h_\tau\}_{\tau=1}^{k} \right\}$: the horizon of foresight t starting at the information set h. It is the collection of all chains of information sets that start from h and are of length at most t.
- $\Omega_h(t) = \{y \in \Omega_h : \exists h' \in h(t) \wedge d' \in D(h') \wedge d' \rhd y\}$: the collection of terminal nodes that succeed immediately from information sets within the horizon $h(t)$.
- $A(h)$: the set of available actions at each of the nodes in h.
- $p(d|\sigma, f)$: the probability that the node d is reached under the strategy-profile $(\sigma, f) \in \Sigma \times \Delta(Q)$.
- $p(d, d'|\sigma, f)$: the conditional probability that d' is reached under (σ, f), if play started at d.
- $\Omega_h(\sigma, f) = \{y \in \Omega_h : \exists d \in D(h) \wedge p(d, y|\sigma, f) > 0\}$: the set of terminal nodes that may occur under (σ, f), if play starts at somewhere in the information set h
- $u(h|\sigma, f) = \sum_{s \in S} \sum_{q \in Q} \sigma(s) f(q) u(\Omega_h(s, q))$: the expected utility payoff under (σ, f), given that play reached h and the agent is fully-rational.
- $R_h(f) = \arg\max_{\sigma \in \Sigma} u(h|\sigma, f)$: the best-response set under rational decision-making (unlimited foresight).
- $(\sigma, f) \underset{h(t)}{\simeq} (\sigma', f')$: two strategy-profiles that are equivalent with respect to the horizon $h(t)$. They agree exactly at every information set within the horizon $h(t)$.

- $\overline{\Omega_h}(t|s,f) = \bigcup_{(s',f')\in S\times\Delta(Q):(s',f')\underset{h(t)}{\simeq}(s,f)} \Omega_h(s',f') \setminus \Omega_h(t|s,f)$: the set of terminal nodes that may follow from any node of the information set h, when (s,f) is played within the horizon $h(t)$, but lie beyond the horizon.

- $\mathcal{U}_h(t|s,f) \equiv \left\{x = u(h|s',f') - u(\Omega_h(t|s,f)) : (s',f') \in S\times\Delta(Q), (s',f') \underset{h(t)}{\simeq} (s,f)\right\}$: the set of terminal utility payoffs that may realize when (s,f) is played within the horizon $h(t)$, but are currently unforeseen.

- $U_h(t|\sigma,f) = u(\Omega_h(t|\sigma,f)) + \sum_{s\in S}\sigma(s)\alpha_h(t|s,f)\inf\mathcal{U}_h(t|s,f)$: the total expected utility payoff to an agent with current horizon $h(t)$, when (σ,f) is played within the horizon.

- $\sum_{s\in S}\sigma(s)\alpha_h(t|s,f)\inf\mathcal{U}_h(t|s,f)$: the continuation value of the horizon $h(t)$, under (σ,f). It is the agent's expected outcome from possible but currently unforeseen future terminal contingencies given that (σ,f) is played within $h(t)$.

- $V_h(t|f) = \max_{\sigma\in\Sigma} U_h(t|\sigma,f)$

- $R_h(t|f) = \arg\max_{\sigma\in\Sigma} U_h(t|\sigma,f)$: the optimal choices from Σ against f under the horizon $h(t)$.

B Risk Aversion in the Small and Large

Here, the analysis of section 4.3 is extended to the general T-period case. For a horizon of length $t\in\{1,...,T\}$, the problem of allocating an amount of wealth w becomes

$$V(t|w) = \max_{\mathbf{c}\in\mathbf{C}}\left\{\sum_{\tau=1}^{t}\delta^{\tau-1}\log c_\tau + \min_{\mathbf{c}'\in\mathbf{C}:\ \mathbf{c}'\underset{t}{\sim}\mathbf{c}}\sum_{\tau=t+1}^{T}\delta^{\tau-1}\log c'_\tau\right\}$$

where $\mathbf{C}$ denotes the set of lifetime consumption plans $\mathbf{c} = \{c_\tau\}_{\tau=1}^{T}$ satisfying

$$s_0 = w \quad s_\tau = s_{\tau-1}(1+r_{\tau-1}) - c_\tau, \quad s_\tau \geq 0,\ c_\tau \geq \underline{c}_\tau \quad \tau\in\{1,...,T\} \tag{21}$$

where s_τ denotes the wealth available at the beginning of period τ while consumption takes place at the beginning of each period. Of course, the initial wealth must suffice for the plan $\{\underline{c}_\tau\}_{\tau=1}^{T}$ of least consumption:

$$w \geq \sum_{\tau=1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}(1+r_i)^{-1} \tag{22}$$

With respect to foresight t, the worst case scenario beyond the horizon would obtain if the agent were to consume $\{\underline{c}_\tau\}_{\tau=t+1}^{T}$. The decision problem can be re-written, therefore, as follows

$$V(t|w) = \max_{\mathbf{c}\in\mathbf{C}}\left\{\sum_{\tau=1}^{t}\delta^{\tau-1}\log c_\tau + \sum_{\tau=t+1}^{T}\delta^{\tau}\log\underline{c}_\tau\right\}$$

For $\mathbf{c} \in \mathbf{C}$, recall that $\mathbf{c}^t = \{c_\tau\}_{\tau=1}^t$ denotes the restriction of $\mathbf{c}$ to the horizon of length t. If the agent chooses the consumption plan $\mathbf{c}^t$ within the horizon, her available wealth at the beginning of period $t+1$ will be

$$s_{t+1}\left(\mathbf{c}^t\right) = w\prod_{i=1}^{t}(1+r_i) - \sum_{\tau=1}^{t} c_\tau \prod_{i=\tau}^{t}(1+r_i) \qquad t \in \{1,...,T-1\}$$

Along with the system of constraints in (21), this implies

$$w = \sum_{\tau=1}^{T} c_\tau \prod_{i=1}^{\tau-1}(1+r_i)^{-1} \tag{23}$$

for any solution $\mathbf{c} \in \mathbf{C}$. Clearly, rational planning requires foresight $t \geq T-1$.

Let the agent consider a horizon of length t. Given that she anticipates her consumption beyond the horizon to be at the worst level, $\{\underline{c}_\tau\}_{\tau=t+1}^T$, the planning problem can be re-written as follows:

$$V(t|w) = \max_{\mathbf{c}^t \in \mathbf{C}} \left\{ \sum_{\tau=1}^{t} \delta^{\tau-1} \log c_\tau + \sum_{\tau=t+1}^{T} \delta^{\tau-1} \log \underline{c}_\tau \right\}$$

The first-order conditions are given by

$$\frac{c_\tau}{c_1} = \delta^{\tau-1} \prod_{i=1}^{\tau-1}(1+r_i) \quad 1 < \tau \leq t \tag{24}$$

The optimal vector $\mathbf{c}^*(w) = \left\{ \{c_\tau^*(w)\}_{\tau=1}^t, \{\underline{c}_\tau\}_{\tau=t+1}^T \right\}$ solves the system of t equations defined by (21) and (24) giving

$$c_1^*(w) = \left(\sum_{\tau=1}^{t} \delta^{\tau-1}\right)^{-1} \left(w - \sum_{\tau=t+1}^{T} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}(1+r_i)} \right)$$

The value of the horizon of length t is given by[34]

$$\begin{aligned} V(t|w) \;=\; & \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(w - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1}(1+r_i)^{-1} \right) + \sum_{\tau=t+1}^{T} \delta^{\tau-1} \log \underline{c}_\tau \\ & - \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(\sum_{\tau=1}^{t} \delta^{\tau-1} \right) + \sum_{\tau=2}^{t} \delta^{\tau-1} \log\left(\delta^{\tau-1} \prod_{i=1}^{\tau-1}(1+r_i) \right) \end{aligned}$$

[34] Abusing notation slightly, the term $\sum_{\tau=2}^{1} x_i$ is to be taken as zero, for any $x_i > 0$.

For a 50-50 lose \$$l$/gain \$$g$ gamble ($g > l$), we get

$$\begin{aligned}
&\mathbb{E}_{\widetilde{w}_0}\left[V\left(t|\widetilde{w}_0\right)\right] \\
= &\sum_{\tau=1}^{t}\frac{\delta^{\tau-1}}{2}\log\left[\left(w+g-\sum_{\tau=t+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)\left(w-l-\sum_{\tau=t+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)\right] \\
&+\zeta\left(\delta,\left\{r_i\right\}_{i=1}^{t-1},\left\{\underline{c}_i\right\}_{i=t+1}^{T}\right)
\end{aligned}$$

where

$$\begin{aligned}
&\zeta\left(\delta,\left\{r_i\right\}_{i=1}^{t-1},\left\{\underline{c}_i\right\}_{i=t+1}^{T}\right) \\
= &\sum_{\tau=t+1}^{T}\delta^{\tau-1}\log\underline{c}_\tau-\sum_{\tau=1}^{t}\delta^{\tau-1}\log\left(\sum_{\tau=1}^{t}\delta^{\tau-1}\right)+\sum_{\tau=2}^{t}\delta^{\tau-1}\log\delta^{\tau-1}+\sum_{\tau=2}^{t}\delta^{\tau-1}\log\left(\prod_{i=1}^{\tau-1}\left(1+r_i\right)\right)
\end{aligned}$$

Recall that $t\left(\widetilde{w}_0\right)$ and $t\left(w_0\right)$ denote the corresponding optimal horizons for the wealth allocation problem under the lottery $\widetilde{w}_0$ and the endowment, respectively. The agent will accept the lottery if and only if $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t\left(\widetilde{w}_0\right)|\widetilde{w}_0\right)\right] > V\left(t\left(w_0\right)|w_0\right)$. That is, the relative lengths $t\left(\widetilde{w}_0\right)$ and $t\left(w_0\right)$ of the optimal horizons matter for the decision of whether or not to accept the lottery. The following two claims establish precisely how they matter.

Claim 5 *For a gamble $\left(g,\frac{1}{2};-l,\frac{1}{2}\right)$, suppose that the induced optimal horizon does not exceed that induced under the certain prospect, $t\left(\widetilde{w}_0\right)\leq t\left(w_0\right)$. If*

$$w_0-\mathbf{1}_{t(\widetilde{w}_0)<T-1}\sum_{\tau=t(w_0)+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\leq\frac{gl}{g-l} \tag{25}$$

then the gamble is rejected. For $t\left(\widetilde{w}_0\right)=t\left(w_0\right)$, the gamble is rejected only if (25) holds.

Proof. $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t|\widetilde{w}_0\right)\right]$ and $V\left(t|w_0\right)$ have all but the respective first summations equal. If both problems induce rational planning, $t\left(\widetilde{w}_0\right)\geq T-1$, the claim is immediate by comparing $\mathbb{E}_{\widetilde{w}_0}\left[V\left(T-1|\widetilde{w}_0\right)\right]$ and $V\left(T-1|\widetilde{w}_0\right)$. Otherwise, notice that

$$\begin{aligned}
&\left(w_0+g-\sum_{\tau=t+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}}\right)\left(w_0-l-\sum_{\tau=t+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}}\right) \\
= &\left(w_0-\sum_{\tau=t+1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)^2+(g-l)\left(w_0-\sum_{\tau=t+1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)-gl
\end{aligned}$$

By (25), this implies $V\left(t\left(\widetilde{w}_0\right)|w_0\right)\geq\mathbb{E}_{\widetilde{w}_0}\left[V\left(t\left(\widetilde{w}_0\right)|\widetilde{w}_0\right)\right]$. Since $t\left(\widetilde{w}_0\right)\leq t\left(w_0\right)$, the monotonicity property gives $V\left(t\left(w_0\right)|w_0\right)\geq V\left(t\left(\widetilde{w}_0\right)|w_0\right)$. ■

Claim 6 *For a gamble* $\left(g, \frac{1}{2}; -l, \frac{1}{2}\right)$*, suppose that it induces a longer optimal horizon than the certain prospect,* $t\left(\widetilde{w}_0\right) = t\left(w_0\right) + k:\ 1 \leq k \leq T - t\left(w_0\right)$*. If*

$$\begin{aligned} & \exp\left(-\frac{2\zeta_k\left(\delta, t\left(w_0\right)\right)}{\sum_{\tau=1}^{t(w_0)} \delta^{\tau-1}}\right)\left(w_0 - \sum_{\tau=t(w_0)+1}^{T} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)^2 \\ < & \left(w_0 + g - \sum_{\tau=t(w_0)+k+1}^{T} \frac{\underline{c}_t}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)\left(w_0 - l - \sum_{\tau=t(w_0)+k+1}^{T} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right) \end{aligned} \tag{26}$$

where

$$\zeta_k\left(\delta, t\right) = \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(\frac{\sum_{\tau=1}^{t} \delta^{\tau-1}}{\sum_{\tau=1}^{t+k} \delta^{\tau-1}}\right) + \sum_{\tau=t+1}^{t+k} \delta^{\tau-1} \log\left(\frac{\delta^{\tau-1}}{\sum_{\tau=1}^{t+k} \delta^{\tau-1}}\right)$$

then the gamble will be accepted.

Proof. We want $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t\left(w_0\right)+k|\widetilde{w}_0\right)\right] > V\left(t\left(w_0\right)|w_0\right)$. By (22), we have

$$\begin{aligned} & \left(w_0 + g - \sum_{\tau=t(w_0)+k+1}^{T} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)\left(w_0 - l - \sum_{\tau=t(w_0)+k+1}^{T} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right) \\ \geq & \left(\sum_{\tau=t+1}^{t(w_0)+k} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}\left(1+r_i\right)}\right)^2 \\ = & \ \zeta\left(\delta, \{r_i\}_{i=1}^{t(w_0)+k-1}, \{\underline{c}_i\}_{i=t(w_0)+k+1}^{T}\right) - \zeta\left(\delta, \{r_i\}_{i=1}^{t(w_0)-1}, \{\underline{c}_i\}_{i=t(w_0)+1}^{T}\right) + \zeta_k\left(\delta, t\left(w_0\right)\right) \end{aligned}$$

The result follows immediately. ■

For complexity costs to account for different risk attitudes in the small and large, it must be that small gambles correspond to shorter optimal horizons than large ones. In what follows, I present sufficient conditions for this to be the case. Specifically, I establish conditions that suffice for the optimal horizon induced by the certain prospect w_0 to be at least as long as the one for a small gamble but shorter than that for a large one. To facilitate comparisons of optimal foresight,

consider

$$\begin{aligned}
&\mathbb{E}_{\widetilde{w}_0}\Delta_t V\left(t|\widetilde{w}_0\right)-\Delta_t V\left(t|w_0\right)\\
=&\sum_{\tau=1}^{t}\frac{\delta^{\tau-1}}{2}\Delta_t\left(\begin{array}{c}\log\left(w_0+g-\sum_{\tau=t+1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)\\+\log\left(w_0-l-\sum_{\tau=t+1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)\\-\log\left(w_0-\sum_{\tau=t+1}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)^2\end{array}\right)\\
+&\frac{\delta^t}{2}\left(\begin{array}{c}\log\left(w_0+g-\sum_{\tau=t+2}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)\\+\log\left(w_0-l-\sum_{\tau=t+2}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)\\-\log\left(w_0-\sum_{\tau=t+2}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}\right)^2\end{array}\right)
\end{aligned}\tag{27}$$

As in Section 4.3, small gambles have gains and losses small enough so that the problem of allocating the corresponding wealth to not differ, in terms of complexity, from that of allocating the endowment.

Definition 7 *A possible realization z of a lottery $\widetilde{z}$ is a **small stake** if*

$$C\left(t|w_0+z\right)=C\left(t|w_0\right)\quad t\in\{1,...,T-1\}$$

*A lottery $\widetilde{z}$ is **small** if every z in the support of $\widetilde{z}$ is a small stake.*

For small stakes z, $C\left(t|\cdot\right)$ is taken to be constant for some range, $(w_0-M,w_0+M):M>0$, around the endowment level: $|z|\leq M$.[35] Equivalently, (16) is valid for $t\in\{1,...,T-1\}$ and for a small gamble $\widetilde{z}$.

Let the present value of the least-consumption stream, $\{\underline{c}_\tau\}_{\tau=t+k}^{T}$, be denoted by $x_{t+k}=\sum_{\tau=t+k}^{T}\underline{c}_\tau\prod_{i=1}^{\tau-1}\left(1+r_i\right)^{-1}$. Define also $F:\mathbb{R}_+^2\times\mathbb{R}_{++}\to\mathbb{R}_{++}$ by $F\left(g,l,x\right)=\frac{\left(w_0-x\right)^2}{\left(w_0+g-x\right)\left(w_0-l-x\right)}$. From (27),

$$\text{(B.i)}\qquad\delta^t\log F\left(g,l,x_{t+2}\right)\geq\sum_{\tau=1}^{t}\delta^{\tau-1}\log\left(\frac{F\left(g,l,x_{t+1}\right)}{F\left(g,l,x_{t+2}\right)}\right)$$

is equivalent to $\mathbb{E}_{\widetilde{w}_0}\Delta_t V\left(t|\widetilde{w}_0\right)\leq\Delta_t V\left(t|w_0\right)$ and implies $\mathbb{E}_{\widetilde{w}_0'}\left[\Delta_t V\left(t|\widetilde{w}_0'\right)\right]-\Delta_t C\left(t|\widetilde{w}_0'\right)\leq\Delta_t V\left(t|w_0\right)-\Delta_t C\left(t|w_0\right)$ for a small gamble.[36] Under a specification for the function $g\left(t|\cdot\right)$ that meets the require-

[35] The definition does not specify a relation for $t=T$ as this would be redundant. Recall that $t=T-1$ is sufficient for rational planning: $V\left(T-1|w\right)=V\left(T|w\right)$, for any wealth realization w. By Lemma 2, $C\left(T-1|w\right)=C\left(T|w\right)$.

[36] $\frac{\partial}{\partial x}F\left(g'',l'',x\right)>0$ is equivalent to $\left(w_0-x\right)^2\left[2\left(w_0-x\right)+g''-l''\right]>2\left(w_0-x\right)\left(w_0+g''-x\right)\left(w_0-l''-x\right)$ or $\left(w_0-x\right)\left(g''-l''\right)<2g''l''$. The rejection condition under the standard setting, $w_0^2>\left(w_0+g\right)\left(w_0-l\right)$, suffices for the last inequality to hold. Since $x_{t+2}>0$, if the gamble is rejected under the standard setting, then $F\left(g,l,x_{t+2}\right)>F\left(g,l,0\right)>0$. Observe now that $x_{t+1}>x_{t+2}$. Let $x_{t+1}\to w_0$ and suppose that $x_{t+1}-x_{t+2}$ remains bounded below by some $K>0$. The right-hand side of (A.1.i) tends to $-\infty$ whereas the left-hand side remains strictly positive.

ments of Corollary 2, such as the one given by (7), (B.i) suffices for a small gamble $\left(g', \frac{1}{2}; -l', \frac{1}{2}\right)$ to give $t\left(\widetilde{w}_0'\right) \leq t\left(w_0\right)$ and, thus, for Claim 5 to establish then a rejection condition.

Let us now turn to large gambles: $\left(g'', \frac{1}{2}; -l'', \frac{1}{2}\right) : g'' > l''$, where l'' is a small stake but g'' is not. The complexity costs argument requires that such gambles induce longer optimal horizons than small ones. If (B.i) holds for a small gamble, it suffices that the optimal horizon under the large gamble is longer than the one induced under the endowment. Thus, we require that the expected marginal benefit of extending the horizon, under the large gamble, exceeds the corresponding marginal benefit under the certain prospect.[37] That is, $\mathbb{E}_{\widetilde{w}_0''}\left[\Delta_t V\left(t|\widetilde{w}_0''\right)\right] > \Delta_t V\left(t|w_0\right)$ which is equivalent to

$$\text{(B.ii)} \qquad \delta^t \log F\left(g'', l'', x_{t+2}\right) < \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(\frac{F\left(g'', l'', x_{t+1}\right)}{F\left(g'', l'', x_{t+2}\right)}\right)$$

Recall that the losses l'' are bounded above by M but the gains are not; the left-hand side of (B.ii) tends to $-\infty$ as $g'' \to +\infty$. In footnote 3 of this section, I show that $F\left(g'', l'', \cdot\right)$ is increasing if the gamble is rejected under the standard setting: $\frac{g''l''}{g''-l''} > w_0$. Since $x_{t+1} > x_{t+2} > 0$, the right-hand side of (B.ii) remains strictly positive and the condition is valid for gambles with sufficiently large gains.

Conditions (B.i) and (B.ii) will hold under plausible assumptions; namely, if both gambles are rejected under the standard setting, the large gamble offers sufficiently large winning stakes and the present value of the worst possible future consumption stream beyond the horizon of the certain prospect, $t\left(w_0\right)$, is sufficiently close to the agent's wealth endowment. Claims 5 and 6 allow then for a large favorable gamble to be accepted when a small one is turned down.

For simplicity of exposition, suppose that the large gamble induces rational planning. For $k = T - t\left(w_0\right) - 1$, (26) reads

$$g'' - l'' \geq \frac{g''l'' - Z\left(\delta, t\left(w_0\right), w_0\right)}{w_0} \tag{28}$$

where

$$Z\left(\delta, t, w\right) = w^2 - \exp\left(-\frac{2\zeta_{T-t-1}\left(\delta, t\right)}{\sum_{\tau=1}^{t} \delta^{\tau-1}}\right)\left(w - x_{t+1}\right)^2$$

Observe that $Z\left(\delta, t\left(w_0\right), w_0\right) > 0$ iff $x_{t(w_0+1)} > \left[1 - \exp\left(\frac{\zeta_{T-t(w_0)-1}(\delta, t(w_0))}{\sum_{t=1}^{t(w_0)} \delta^{\tau-1}}\right)\right] w_0$.[38] For suffi-

[37] Let τ be the optimal foresight under the certain prospect, $t\left(w_0\right) = \tau$, and $\mathbb{E}_{\widetilde{w}_0''}\left[V\left(\tau|\widetilde{w}_0''\right)\right] < V\left(\tau|w_0\right)$ (the large gamble is rejected under foresight τ). Set $\frac{g\left(1+\tau|\widetilde{w}_0''\right)}{g\left(\tau|\widetilde{w}_0''\right)} = \frac{g(1+\tau|w_0)}{g(\tau|w_0)} = \frac{V(1+\tau|w_0)}{V(\tau|w_0)}$. Then, $\mathbb{E}_{\widetilde{w}_0''}\left[\Delta_t V\left(\tau; \widetilde{w}_0''\right)\right] > \Delta_t V\left(\tau|w_0\right)$ implies $\frac{\mathbb{E}_{\widetilde{w}_0''}\left[V\left(1+\tau|\widetilde{w}_0''\right)\right]}{\mathbb{E}_{\widetilde{w}_0''}\left[V\left(\tau|\widetilde{w}_0''\right)\right]} > \frac{V(1+\tau|w_0)}{V(\tau|w_0)} = \frac{g\left(1+\tau|\widetilde{w}_0''\right)}{g\left(\tau|\widetilde{w}_0''\right)}$. By Proposition 2(i), the large gamble induces a longer horizon than the certain prospect: $t\left(\widetilde{w}_0''\right) > \tau$.

[38] Consider $Z\left(\delta, t, w_0\right)$ as a quadratic in x_{t+1}. Then, $Z\left(\delta, t, w_0\right) > 0$ iff $\left[1 - \exp\left(\frac{\zeta_{T-t-1}(\delta,t)}{\sum_{\tau=1}^{t} \delta^{\tau-1}}\right)\right] w_0 < x_{t+1} < \left[1 + \exp\left(\frac{\zeta_{T-t-1}(\delta,t)}{\sum_{\tau=1}^{t} \delta^{\tau-1}}\right)\right] w_0$. By (23), the second inequality holds trivially.

ciently small δ, this holds for any values of $\{r_i\}_{i=1}^{T-1}, w_0$ and $\{\underline{c}_\tau\}_{\tau=t(w_0)}^{T}$.[39] Moreover, if $t(w_0)+1 > \frac{\sum_{\tau=1}^{t(w_0)}\tau\delta^\tau}{\sum_{\tau=1}^{t(w_0)}\delta^\tau}$, a sufficient condition for it to hold for all values of δ is $\sum_{\tau=t(w_0)+1}^{T}\frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}(1+r_i)} > \left(T^{\frac{T}{t(w_0)+1}} - t(w_0) - 1\right) T^{-\frac{T}{t(w_0)+1}} w_0$.[40]

Notice also that $Z(\delta, w_0, t)$ increases with x_{t+1}.[41] Ceteris paribus, the lower bound for $g''-l''$ in (28) falls with the present value of the amount of initial wealth allocated for consumption optimally under the certain prospect. Moreover, if

$$\frac{w_0 - x_{t(w_0)+2}}{w_0 - x_{t(w_0)+1}} > \exp\left(\frac{2\zeta(\delta, t(w_0)+1)}{\sum_{\tau=1}^{t(w_0)+1}\delta^{\tau-1}} - \frac{2\zeta(\delta, t(w_0))}{\sum_{\tau=1}^{t(w_0)}\delta^{\tau-1}}\right) \tag{29}$$

the lower bound decreases the shorter is the optimal horizon under the certain prospect, $t(w_0)$.[42]

---

[39] Since $\lim_{\delta\to 0^+}\log\left(\frac{\sum_{\tau=1}^{t}\delta^{\tau-1}}{\sum_{\tau=1}^{t+k}\delta^{\tau-1}}\right) = 0 = -\lim_{\delta\to 0^+}\left(\frac{\sum_{t=t+1}^{t+k}\delta^{\tau-1}}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\right)\log\left(\sum_{\tau=1}^{t+k}\delta^{\tau-1}\right)$ and $\lim_{\delta\to 0^+}\frac{\sum_{\tau=t+1}^{t+k}\delta^{\tau-1}\log\delta^{\tau-1}}{\sum_{\tau=1}^{t}\delta^{\tau-1}} = 0^-$, we have $\lim_{\delta\to 0^+}\frac{\zeta_k(\delta,t)}{\sum_{\tau=1}^{t}\delta^{\tau-1}} = 0^-$. Let $k = T - t(w_0) - 1$ for the result in the text.

[40] We have

$$\begin{aligned}
\frac{\partial}{\partial\delta}\left(\frac{\zeta_{T-t-1}(\delta,t)}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\right) &= \frac{\sum_{\tau=2}^{t}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{t}\delta^{\tau-1}} - \frac{\sum_{\tau=2}^{T}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{T}\delta^{\tau-1}} \\
&\quad + \frac{\sum_{\tau=t+1}^{T-1}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{t}\delta^{\tau-1}} - \frac{\sum_{\tau=t+1}^{T-1}\delta^{\tau-1}}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\frac{\sum_{\tau=2}^{T}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{T}\delta^{\tau-1}} \\
&\quad + \left(\sum_{\tau=1}^{t}\delta^{\tau-1}\right)^{-1}\sum_{\tau=t+1}^{T-1}(\tau-1)\delta^{\tau-2}\log\left(\frac{\delta^{\tau-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right) \\
&\quad - \left(\sum_{t=1}^{t}\delta^{\tau-1}\right)^{-2}\sum_{\tau=2}^{t}(\tau-1)\delta^{\tau-2}\sum_{\tau=t+1}^{T-1}\delta^{\tau-1}\log\left(\frac{\delta^{t-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right) \\
&= \frac{\sum_{\tau=2}^{T}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{t}\delta^{\tau-1}} - \frac{\sum_{\tau=2}^{T}(\tau-1)\delta^{\tau-2}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\left(1 + \frac{\sum_{\tau=t+1}^{T}\delta^{\tau-1}}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\right) \\
&\quad + \left(\sum_{\tau=1}^{t}\delta^{\tau-1}\right)^{-2}\left(\begin{array}{c}\sum_{\tau=1}^{t}\delta^{\tau-1}\sum_{\tau=t+1}^{T}(\tau-1)\delta^{\tau-2}\log\left(\frac{\delta^{\tau-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right) \\ -\sum_{\tau=2}^{t}(\tau-1)\delta^{\tau-2}\sum_{\tau=t+1}^{T}\delta^{\tau-1}\log\left(\frac{\delta^{\tau-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right)\end{array}\right) \\
&= \left(\sum_{\tau=1}^{t}\delta^{\tau-1}\right)^{-2}\left(\begin{array}{c}\sum_{\tau=t+1}^{T}(\tau-1)\delta^{\tau-2}\log\left(\frac{\delta^{\tau-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right) \\ +\sum_{\tau=2}^{t}\sum_{\tau=t+1}^{T}(t-\tau)\delta^{t+\tau-3}\log\left(\frac{\delta^{t-1}}{\sum_{\tau=1}^{T}\delta^{\tau-1}}\right)\end{array}\right)
\end{aligned}$$

This expression is not of the same sign across all $T \in \mathbb{N}\setminus\{0\}, t \in \{1,...,T-1\}$. A sufficient condition for it to be negative is $t+1 > \frac{\sum_{\tau=1}^{t}\tau\delta^\tau}{\sum_{\tau=1}^{t}\delta^\tau}$. In this case, $\inf_{\delta\in(0,1)}\exp\left(\frac{\zeta_{T-t-1}(\delta,t)}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\right) = \exp\left(\frac{\zeta_{T-t-1}(1,t)}{t+1}\right)$ where $\zeta_{T-t-1}(1,t) = (t+1)\log(t+1) - T\log T$.

[41] Given $\lambda_t > 0$ and $w_0 > x_{t+1}$, $-\lambda_t(w_0 - x_{t+1})^2$ is strictly increasing in x_{t+1}. In $Z(\delta, w_0, t)$, set $\lambda_t = \exp\left(-\frac{2\zeta_{T-t-1}(\delta,t)}{\sum_{\tau=1}^{t}\delta^{\tau-1}}\right)$. We have $w_0 > x_{t+1}$ by (23).

[42] $\Delta_t Z(\delta, w_0, t) = -\lambda_{t+1}(w_0 - x_{t+2})^2 + \lambda_t(w_0 - x_{t+1})^2 < -(w_0 - x_{t+1})[\lambda_{t+1}(w_0 - x_{t+2}) - \lambda_t(w_0 - x_{t+1})]$ where λ_t is defined in the preceding footnote. The inequality follows from x_{t+k} being decreasing in k.

B.1 Extension to Lotto-type Bets

Let us now consider gambles $(g^*, p; -l^*, 1-p) : p \to 0$ with again only l^* being a small stake. A minuscule probability of winning simplifies significantly the mathematical exposition of the argument and is also appropriate for lotto-type bets (where p is usually of the order 5×10^{-9} when one buys one lotto entry for $l^* = \$1$). We have

$$\begin{aligned}
&\mathbb{E}_{\widetilde{w}_0^*}\left[V\left(t|\widetilde{w}_0^*\right)\right] \\
= &\sum_{\tau=1}^{t} \delta^{\tau-1} \left(\begin{array}{c} \log\left(w_0 + g^* - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{p} \\ +\log\left(w_0 - l^* - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{1-p} \end{array} \right) \\
&+ \sum_{\tau=t+1}^{T} \delta^{\tau-1} \log \underline{c}_\tau - \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(\sum_{\tau=1}^{t} \delta^{\tau-1}\right) + \sum_{\tau=2}^{t} \delta^{\tau-1} \log\left(\delta^{\tau-1} \prod_{i=1}^{\tau-1} (1+r_i)\right)
\end{aligned}$$

and

$$\begin{aligned}
&\mathbb{E}_{\widetilde{w}_0^*}\left[\Delta_t V\left(t|\widetilde{w}_0^*\right)\right] - \Delta_t V\left(t|w_0\right) \\
= &\sum_{\tau=1}^{t} \delta^{\tau-1} \Delta_t \left(\begin{array}{c} \log\left(w_0 + g^* - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{p} \\ +\log\left(w_0 - l^* - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{1-p} \\ -\log\left(w_0 - \sum_{\tau=t+1}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right) \end{array} \right) \\
&+\delta^t \log \left[\frac{\left(w_0 + g^* - \sum_{\tau=t+2}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{p} \left(w_0 - l^* - \sum_{\tau=t+2}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}\right)^{1-p}}{w_0 - \sum_{\tau=t+2}^{T} \underline{c}_\tau \prod_{i=1}^{\tau-1} (1+r_i)^{-1}} \right]
\end{aligned}$$

Therefore, $\mathbb{E}_{\widetilde{w}_0^*}\left[\Delta_t V\left(t|\widetilde{w}_0^*\right)\right] > \Delta_t V\left(t|w_0\right)$ is equivalent to

$$\text{(B.1.i)} \qquad \delta^t \log F\left(g^*, l^*, x_{t+2}\right) < \sum_{\tau=1}^{t} \delta^{\tau-1} \log\left(\frac{F\left(g^*, l^*, x_{t+1}\right)}{F\left(g^*, l^*, x_{t+2}\right)}\right)$$

where $F : \mathbb{R}_+^2 \times \mathbb{R}_{++} \to \mathbb{R}_{++}$ is defined by $F\left(g, l, x\right) = \frac{(w_0 - x)^2}{(w_0 + g - x)^p (w_0 - l - x)^{1-p}}$. For given $p > 0$, as $g \to +\infty$, the left-hand side of (B.1.i) tends to $-\infty$ whereas the right-hand side remains bounded below by a finite negative number.[43] That is, for every p, there exists some large enough gain g^* so that (B.1.i) holds.

Therefore, the lotto-type bet can induce a longer optimal horizon than that of the certain prospect, $t\left(\widetilde{w}_0^*\right) = t\left(w_0\right) + k$ for some $k \geq 1$. In this case, the following claim establishes a sufficient condition for the gamble to be accepted.

[43] Since $x_{t+1} > x_{t+2}$, $\frac{F\left(g,l,x_{t+1}\right)}{F\left(g,l,x_{t+2}\right)} = \left(\frac{w_0 - x_{t+1}}{w_0 - x_{t+2}}\right)^2 \left(\frac{w_0 + g - x_{t+2}}{w_0 + g - x_{t+1}}\right)^p \left(\frac{w_0 - l - x_{t+2}}{w_0 - l - x_{t+1}}\right)^{1-p} > \left(\frac{w_0 - x_{t+1}}{w_0 - x_{t+2}}\right)^2 > 1.$

Claim 7 *For a gamble $(g^*, p; -l^*, 1-p)$, suppose that it induces a longer optimal horizon than the certain prospect, $t(\widetilde{w}_0^*) = t(w_0) + k : 1 \leq k \leq T - t(w_0)$. If*

$$e^{-\frac{2\zeta_k(\delta, t(w_0))}{\sum_{\tau=1}^{t(w_0)} \delta^{\tau-1}}} \left(w_0 - x_{t(w_0)+1}\right) \leq \left(w_0 + g^* - x_{t(w_0)+k+1}\right)^p \left(w_0 - l^* - x_{t(w_0)+k+1}\right)^{1-p} \tag{30}$$

then the gamble will be accepted.

Proof. We want $\mathbb{E}_{\widetilde{w}_0}\left[V\left(t\left(w_0\right)+k|\widetilde{w}_0^*\right)\right] > V\left(t\left(w_0\right)|w_0\right)$ for $1 \leq k \leq T - t(w_0)$. The argument is identical to that in the proof of Claim 4, since by (23) we get

$$\left(w_0 + g^* - x_{t(w_0)+k+1}\right)^p \left(w_0 - l^* - x_{t(w_0)+k+1}\right)^{1-p} > \sum_{\tau=t+1}^{t+k} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}(1+r_i)}$$

for any $1 \leq t \leq T-k$. ■

Given g^*, let $p \to 0$. (30) will hold if

$$\text{(A.2.ii)} \quad \left(e^{-\frac{2\zeta_k(\delta, t(w_0))}{\sum_{\tau=1}^{t(w_0)} \delta^{\tau-1}}} - 1\right) \left(w_0 - x_{t(w_0)+k+1}\right) \leq e^{-\frac{2\zeta_k(\delta, t(w_0))}{\sum_{\tau=1}^{t(w_0)} \delta^{\tau-1}}} \sum_{\tau=t(w_0)+1}^{t(w_0)+k} \frac{\underline{c}_\tau}{\prod_{i=1}^{\tau-1}(1+r_i)} - l^*$$

Since $\lim_{\delta \to 0^+} \frac{\zeta_k(\delta,t)}{\sum_{\tau=1}^{t} \delta^{\tau-1}} = 0^-$ (recall footnote 6 of the previous section of Appendix B), this inequality is valid for sufficiently low values of δ. For every g^*, there exist pairs (δ, p), with the discount factor and winning probability both being small enough, so that (30) holds.[44] Notice that (B.1.i) becomes easier to satisfy the lower is the initial endowment w_0, other things remaining unchanged.

[44] Since a larger g^* makes (30) more likely to hold, this argument is valid in conjunction with the one given for why (B.1.i) should hold. That is, for sufficiently small δ, there do exist pairs (p, g^*) with p very small but g^* sufficiently large such that both (30) and (B.1.i) are satisfied.