Decision Making as Epiphany: A Search Based Model of Limited Rationality

Derek Laing
Syracuse University

September 22, 2009

(***Preliminary Version – please do not cite without permission***)

Abstract

In this paper we propose a dynamic model of boundedly rational agents in which, because of cognitive and/or and memory limitations, agents fail to comprehend all of the actions that are available to them at a given point in time. However, they can formulate new "ideas" that reveal alternative courses of action via a costly process of ratiocination – "thinking." Agents are aware of the distribution of potential rewards that accrue from each new idea : they are fully appraised of the economically essential aspects of their own ignorance. We establish that their behavior is characterized by a reservation utility, and we analyze the determinants of their optimal thinking "intensity." Our approach is notable in that it provides a useful, but albeit rudimentary, characterization of an agent's innate intelligence. Thus, one agent is more intelligent than another if, *ceteris paribus*, he thinks of ideas at a faster rate and/or thinks of better quality ideas. We apply the framework to a variety of economic settings.

Epiphany. *A sudden, intuitive perception of or insight into the reality or essential meaning of something, usually initiated by some simple, homely, or commonplace occurrence or experience*. [Random House, Webster's unabridged dictionary, 2002].

Keywords: Search; bounded Rationality; learning; strategy and tactics; authority.

JEL classification: D11; D81; D83; M37 ; M55

Acknowledgements: I appreciate helpful discussions with Dale Mortensen, and Ping Wang and comments from seminar participants at the S.E.D meetings in Florence, UIUC, Clemson, and Syracuse.

Correspondence address: Derek Laing, Department of Economics, 123 Eggers Hall, Syracuse University, Syracuse NY 13210 (E-mail: dlaing@maxwell.syr.edu.)

"If one wished to be unfair to economists in general, he might select ... a certain well-known though fictitious character whose idiosyncrasies furnish alternate joy and irritation... He is a somewhat inhuman individual who, inconsistently enough, carries the critical weighing of hedonistic values to the point of mania. So completely is he absorbed in his irrationally rational passion for impassionate calculation that he often remains a day laborer at pitifully low wages from sheer devotion to the fine art of making the most out of his scanty income ...His enemies consider him eccentric... They are right. He is eccentric. Indeed, he is not even a good hedonist... A good hedonist would stop calculating when it seemed likely to involve more trouble than it was worth...." [Clark (1918, p.26)]

1 Introduction

Over the two centuries or so that have elapsed since the publication of Adam Smith's *The Wealth of Nations*, much of the progress made by economic science has been made possible by a strict adherence to a particular conception of rational economic behavior. The crystallized embodiment of this viewpoint is located in the canonized being of rational economic man himself — *homo economus*.[1] He has consistently belonged to the vanguard of countless battalions of economists as they conquer the traditional territories of economics, and even ones that lie beyond them. There is little doubt that *homo economicus* will continue to remain an integral part of future spearheads, as subsequent legions of economists carve out and dispatch new and exciting areas of inquiry. Nevertheless, several areas of legitimate economic discourse have proven to be stubbornly recalcitrant, and, hence, unyielding to the knight-errant's prowess. The difficulties alluded to do not stem from some obvious defect that inheres in *homo economus* himself; on the contrary, they arise precisely because of his towering intellect.

Thus, the conception of rationality, as it is used by economists, hinders any satisfactory characterization of a player's innate *intelligence* – and, as an immediate corollary, his potential lack thereof. Rational men are identical in this regard, standing atop the highest pinnacles of intellectual acuity. Yet, differences in intelligence levels are arguably quintessential aspects of the world in which we live. How else can we explain, for example, why men and women of vision become leaders, while others become followers? Why some researchers are lackluster, while others excel for their efforts? More insidiously, lacking a suitable compass for intellectual acuity, it is difficult to analyze strategic situations in which, for example, the strong-minded seek to manipulate those who are intellectually weaker than themselves.

Moreover, *homo economus* has little use for the endless stream of (non-informative) advertising he sees on TV and in other media. Indeed, a man of such brilliance simply

[1] Of course, the decision maker is not necessarily a male. Instead, by an appeal to the synechdoche, we use "he" to mean "he or she."

selects the option that is best for him, given he understands both his own preferences and the constraints that limit his choices. Consequently, he is a bulwark against the apparently irrational behavior of firms, as they spend billions of dollars each year trying to influence him, by telling him what he already knows.[2] Finally, rational man has little need for the core institutions of modern economies, such as firms, unions, and even his family. These institutions call for him to cede at least part of his decision making authority to others, who then subsequently use their power to control, and hence limit his future actions. His omniscience precludes any apparent use for them, since he can, within the rule of law, use his intellect to judiciously craft contracts that replicate their behavior and, in the process, render them all equally otiose.

Thus, rational economic man's striking intellect evidently is an impediment to progress in some critical areas of economics. Some sixty years ago Simon (1955) advocated that economists should model the cognitive limitations of actual human beings and proposed a means whereby they could do so. This research program has garnered a several names since its inception — with limited or bounded rationality being the most common among them.[3] The field has recently witnessed renewed vigor after, perhaps, languishing in the recesses of the profession for some time.[4] An important theme in the new literature, is that much progress can be made by recognizing the costs of decision making. Conlisk (1996c) is a forceful advocate for this position:[5]

> *"If rationality is scarce, good decisions are costly. There is a trade-off between effort devoted to deliberation and effort devoted to other activities... A model of the trade-off requires some form of "deliberation technology" by which a decision maker turns scarce cognitive and other resources into better decisions*." [Conlisk (1996c, p.682).]

In this paper, we develop a dynamic model of boundedly rational agents that explicitly incorporates deliberation costs. We consider an environment in which, at any given instant in time, agents are unable to comprehend all of the actions, a, that are actually available to them: $\mathcal{A}$.[6] Instead, we assume that they only perceive those actions that belong to their

---

[2] We do not want to claim too much: There are several fine extant models of informative, and non-informative advertising (see Bagwell (2005)). Instead, our claim is that models populated by boundedly rational agents may play a useful complementary role in providing a more complete picture of the economics of advertising.

[3] Simon (1987, p.266) himself defines bounded rationality as follows:

> *"The term 'bounded rationality' is used to designate rational choice that takes into account the cognitive limitations of the decision maker—limitations of both knowledge and computational capacity."*

[4] Conlisk (1996c), Lipman (1995), and Rubinstein (1998) offer especially useful surveys.

[5] Despite its importance, he notes: *"[V]ery few explicit models of deliberation technology and deliberation cost have appeared."* [Conlisk (1996c, p.682)]. Conlsik has constructed several interesting models that involve a deliberation technology [see: Conlisk (1988), Conlisk (1996a), Conlisk (2001) and Conlisk (2003)]. None, however, adopt the search perspective proposed in this paper.

[6] In the terminology of Lipman (1999) the decision maker is not *logically* omniscient.

sparser *active* "awareness" sets $A(t) \subset \mathcal{A}$. Reasonably enough, a given agent can select an action $a \in \mathcal{A}$ – and in doing so derive utility $u(a)$ – if and only if he is aware of its existence (i.e., $a \in A(t)$).

At one extreme, this posited lack of "awareness," may simply reflect the difficulty of recalling known facts from memory. At the other, as emphasized by Aragones, Gilboa, Postlewaite, and Schmeidler (2005), from the practical difficulties associated with "fact-free learning ;" an endeavor that calls for agents to draw valid inferences from the information they already possess. This latter difficulty is no doubt as ubiquitous as it is important. As Aragones et al remark,

> *"[P]utting wheels at the bottom of a suitcase allows it to roll easily. This idea was quite original when it was first introduced. But, since it only selected and combined facts that everyone had already known, it appears obvious in hindsight."* [p.1355]

Military history is also replete with such apparent instances. Erwin Rommel's brilliant *tactical* use — on-the-fly —of the 88-FlaK *anti-aircraft* gun as an *anti-tank* weapon during the North-Africa campaign provides a notable case in point.[7] Thus,

> *"[D]uring a single action, the attack on Sidi Omar 22 November 1941, a British brigadier with 51 thick-skinned infantry tanks lost 47, most of them to 88-mm. antitank fire. By the end of the Winter Battle, out of 1,276 tanks sent to Libya, 674 were damaged and 274 were destroyed. Rommel's Afrika Korps had so crippled the armor that the British could not resume the offensive until May 1942."* [8]

This example is instructive on at least three counts. First, the use of the FlaK weapon in its *anti-tank* role provides a monstrous example of fact-free learning: its use in this capacity is obvious *with hindsight*. (The British were well aware that the Germans had batteries of *anti-aircraft* guns.) Had the British *recognized* this possibility, they (presumably) would have not pursued a tactic which led to the loss of upwards of 80% of their available armor. Second, not only does it provide an illustration of "fact free learning," it may also exemplify the difficulty of recalling already known facts from memory: the British had evidently forgotten that, during the earlier 1941 campaign against France, Rommel had already employed this tactic with deadly effect. Finally, in an observation that is particulary germane to this paper, it hints at the potential strategic considerations that arise in settings wherein a player faces an opponent who is capable of carrying out significant tactical innovations on-the-fly.

More subtly, the possibility of "fact-free" learning also has profound consequences for economists' usual Knightian distinction between situations of "risk and uncertainty." Thus, the outcome of a "fair" coin toss is usually deemed to be one of risk, while the outcome of a coin toss with "unknown probabilities" (or for that matter unknown outcomes) is one of

[7]Fliegerabweherkanone.

[8]Source: http://www.army.mil/.

uncertainty. Nevertheless, both of these concepts are predicated on the implicit assumption that the decision maker is capable of comprehending that a coin is being tossed in the first place. Thus, in the case of Rommel's use of FlaK guns as tank destroyers, presumably the problem was that the British command failed to *imagine* that the weapon could be used in this manner. *If* the idea had dawned on them, then they would have immediately recognized the *risks* it posed to their armor.

While the focus of Aragones et al's paper was characterizing the inherent computational complexity of drawing inferences from already known facts, the focus of this one is examining the behavior of a decision maker who is aware that he suffers from this malady. Indeed, our main point of departure from their work is that we assume that agents recognize that they can uncover possible actions via a process of ratiocination — "thinking." We model the thinking process using a search theoretic tools. Thus, at random points in time, and at rate governed by their search intensities, they have new "ideas" concerning feasible actions. As agents gradually uncover elements of $\mathcal{A}$ they augment their active awareness set $A(t)$. We assume that they can influence the rate, and the quality, of their ideas by exerting costly effort. This formulation allows us to model the speed with which an agent will arrive at a decision, and to relate it to both internal features of his environment (i.e., his innate intelligence level) and to external ones that arise because of, for example, the interposition of other players.

In our framework, players recognize that they can either select an action from $A(t)$, or else *wait* until they have a better idea. Although they do not know all of the elements of $\mathcal{A}$, they are assumed to understand the consequences of further deliberation. More specifically, we assume that they recognize the distribution of *utilities* associated with each new idea. In other words, each agent is assumed to be appraised of the economically essential aspects of his own ignorance. This feature imposes sufficient coherence on the decision maker's problem to render it highly tractable. In particular, in the environments we consider, an agent's behavior is characterized by a unique, and stationary, reservation utility s^*. The agent makes a decision if and only if he has an "epiphany;" an idea, $a \in \mathcal{A}$, that provides him with utility that exceeds his reservation value: $u(a) \geq s^*$.

In Sections 2-3 we develop the theoretical framework, while Sections 4–5 present some simple applications of the approach. More specifically, Section 4 considers the possibility of strategic manipulation, in which one player (the "speaker") seeks to influence the actions taken by another (the "listener") for her own benefit.[9] However, unlike *games of persuasion*, in which the speaker must select the *information* she intends to communicate to the listener, in this paper the speaker truthfully reveals everything she knows. Nevertheless, cognitive differences between the two parties can lead to circumstances under which the presence of the speaker actually makes the listener worse off; this despite his ability to veto any suggestion she makes. In Section **??** we examine the economics of *non-*

[9]See, for example, Glazer and Rubinstein (2006).

informative advertising, which arises when firms spend advertising dollars to apparently remind consumers of what they already know. Although there are no *market* frictions – we assume that consumers can instantly purchase the goods they desire – we show that trade is costly and time consuming. The reason is that cognitive limitations imply that it takes time for agents to formulate their consumption plans. Interestingly, we show that the Diamond paradox emerges only if the market is populated by a large number of firms; (paradoxically) it fails to emerge in the case of pure monopoly, in which there is a single firm.[10] Section 5 then applies that framework to study strategic situations of "conflict" between two parties. The point of departure from the literature is that we examine the consequences of allowing the parties to carry out *tactical* innovations "on the fly." Section 7 offers some concluding comments.

Related Literature

This paper is obviously also related to the voluminous search literature.[11] Indeed, a large body of work has constructed and analyzed *dynamic* models of boundedly rational decision making in general, and *search* models in particular.[12] However, "search" has been interpreted in several distinct ways in the literature. In broad outline, it is possible to delineate among four broad categories of model that make use of it. They are: (*i*) information acquisition, (*ii*) "rules of thumb," (*iii*) search *by* bounded rational agents and (*iv*) cognition *as* a process of search.

(i) The Acquisition of Information. In 1955 Simon proposed, and analyzed, the first sequential search model. He considered an environment in which an individual searches for a house, and must sequentially gather price information before he purchases one.[13] In Simon's model, and other search papers that analyze the acquisition of information, agents "understand" everything about their environments: there are no limitations placed on their abilities to process the information they possess. Instead, "frictions" inhere in the market, and hinder agents' abilities to acquire factual information which, in turn, implies that transactions are costly and time consuming.[14]

[10] The Diamond paradox arises when search frictions induce all firms to set the monopoly price.

[11] Rogerson, Shimer, and Wright (2005), provide an exceptional review of this literature.

[12] Some of them, include: Conlisk (2003), Conlisk (1996b), Cross (1973), Day (1967), Day and Tinney (1968), de Palma, Myers, and Papageorgiou (1994), Dow (1991), Lettau and Uhlig (1999), Mullainathan (2002), Osborne and Rubinstein (1998), Rosenthal (1993), Rubinstein (1993), Radner (1975b), Radner (1996), Radner (1975a), Radner and Rothschild (1975), Simon (1955), Winter (1971), and Wilson (2003).

[13] See, Radner (1975b), Radner (1996), Radner (1975a), and Radner and Rothschild (1975) for models of costly information acquisition as they pertain to boundedly rational agents. The field was more formally developed, and popularized, by Stigler (1961), in the context of the search for an automobile. Rogerson, Shimer, and Wright (2005) provide an excellent survey of the most recent literature. The earlier generations of search literature are surveyed by Lippman and McCall (1976), and Mortensen (1986).

[14] Burdett, Shi, and Wright (2001) consider an environment in which agents know prices and product loca-

In this paper, we seek to broaden the compass of the search approach by using it to study the implications of "fact free" learning. The distinction we make between the use of search-theoretic tools to model the acquisition of costly information and their use, proposed here, to model the costly process of ratiocination, is much more pronounced than the "suspenders versus braces" variety. Thus, it is conceivable that markets may be characterized by significant trade frictions, even if it is costless to gather additional information. For example, (at a rather mundane level) it make take an individual an inordinate amount of time to purchase a replacement box of his favorite breakfast cereal because he persistently forgets he wants to do so, rather than because he is unsure about where to buy one, and at what price. Moreover, during the lengthy process of deliberation, he may be susceptible to the "suggestions" — and hence influence — of others. For example, a judiciously placed cereal (or even milk) commercial may remind him of his desire to purchase a new box of cereal — even if the commercial reminds him of something he already knows, and thus fails to provide him with new information *per se*.

(ii) "Rules of Thumb." Many papers have argued that real world decision makers are incapable of solving the apparently complicated dynamic problems that economists pose for them. Simon (1955) suggested that in this sort of setting they will use simple "rules of thumb," in order to simplify their decision making: They will "satisfice" rather than "optimize." A focal point of this literature has been addressing the question: how well do rules of thumb perform relative to optimizing behavior?[15] The answer is often "remarkably well." In an interesting recent paper, Conlisk (2003) demonstrates that a simple adaptive procedure converges rapidly to the (unique) optimal stopping rule of sequential search theory.

Nevertheless, a basic problem with this approach is that restrictions on optimizing behavior are inherently *ad hoc*, and hence undesirable. Lipman's admonition concerning this point is quite apposite:

> *"It makes no sense to ask the question 'how can the agent carry out this complex calculation' when the 'complex task' is simply our representation of whatever it is the agent in fact manages to do"* [Lipman (1999, p.342-343)][16]

In view of this weakness, in this paper we are careful to describe the world *as perceived by the decision maker*. We then unabashedly use the tools of optimization to characterize the agent's behavior, implying that he behaves optimally in the environment as he comprehends it. Yet, despite being a "maximizer," the *constraints* the hinder the decision maker's

tions. However, frictions arise because they do not know which store other players in the game intend to visit (and goods can be in "short supply").

[15]Early work in this area includes: Baumol and Quandt (1964), Day (1967), Day and Tinney (1968), Rosenthal (1993), Simon (1955), Winter (1971). More recent contributions include papers by, Conlisk (2003), Conlisk (1996b), Cross (1973), de Palma, Myers, and Papageorgiou (1994), Lettau and Uhlig (1999).

[16]See also, Lipman (1991) and Lipman (1995) for related viewpoints.

ability to comprehend the actions that are available to him impose serious limitations on his *de facto* rationality.

(iii) Search by Bounded Rational Agents. Dow (1991) considers an environment in which a boundedly rational player sequentially searches for a good. (The player sequentially observes the prices of the same good offered in two different two stores).[17] Although the agent can revisit the first store, the snag is that he is capable of only remembering a range in which its price offer fell. Dow shows that the agent optimally "remembers," by partitioning of the range of prices into intervals. Nevertheless, although Dow uses search theory to model that sampling of prices it is not central to his approach. Put another way the agent's decision problem would remain unchanged if he lived in a world without search, but observed two prices at different points in time.[18]

(iv) Cognition as a Process of Search. From its very inception, search has been deemed to constitute an essential element of bounded rationality. Thus, according to Simon (1979):

> *"[T]he failures of omniscience are largely failures of knowing all the alternatives, uncertainty about relevant exogenous events, and inability to calculate consequences. ... If the alternatives for choice are not given initially to the decision maker, then he must search for them. Hence, a theory of bounded rationality must incorporate a theory of search"* [Simon (1979, p.502-503)]

4 Despite this auspicious beginning, few papers have actually used search theory to model the process of cognition itself. An important exception is a series of recent interesting papers by Macleod (MacLeod (1996), MacLeod (2002a), and MacLeod (2002b)), which are clearly germane to our work.

The hallmark of MacLeod's work is that he drops Savage's "small world assumption," which asserts that the decision maker understands all of the consequences of his actions, and can assign a probability to each state of nature. Instead, he constructs models of decision making in complex environments, drawing upon recent developments in cognitive psychology.[19] Like MacLeod, we also drop the small world assumption. However, MacLeod focuses on the problems that arise when players find it difficult to calculate the payoffs of *known* actions. In this paper we examine situations in which the decision maker fails to comprehend some actions at all.

[17]See also: Rubinstein (1993).

[18]Several authors have recently constructed interesting models that stress the importance of memory limitations See for example: Mullainathan (2002), and Wilson (2003). However, the structure of these models is quite different to the one proposed here.

[19]Specifically, in his (1996) paper he uses Hölmström and Milgrom's multi-task principal agent framework to show how remarkably complicated the environment becomes as the number of tasks increases. For example, suppose that there are k different tasks, and that each of them is associated with m possible realized benefits and n realized costs. If there are fifteen tasks and the benefits and costs can each take five distinct values, then MacLeod shows that, at one cent apiece, including every contingency in a contract would cost some ten million trillion dollars (MacLeod (1996, p.797)). MacLeod argues that this complexity leads endogenously to the emergence of incomplete contracts

2 The Model

Before presenting the details of the formal model, we first offer a simple illustration of the type of decision problem we have in mind.

Composition

Consider a budding author – Dougal – who has reached an impassé as he seeks to complete the sentence:

> *"A wise and shrewd politician, Smith recognized the enormous political capital that he would amass if he **** (in) implementing the sweeping set of reforms he had promised during the elections."*

Dougal seeks a word to insert in place of **** that captures "intentionally delays." Clearly, the "value" of the sentence depends upon inserting just the right word. For the sake of the current argument, assume that the (universal) set of possibilities is:[20]

$$\mathbb{A} = \{\textit{late, delay, procrastinate, tarry, cunctatious, gele}\}$$

He might behave in the following manner,

- He immediately (i.e., at $t = 0$) thinks of "late" and "delay."
- Unhappy with these choices, he thinks of "procrastinate" at time $t_1 > 0$.
- Not satisfied with {*delay, late, procrastinate*}, he then looks for his dictionary. At time t_2 he finds it. He is *reminded* of the word "*tarry*" (at t_3), and *discovers* the new word {*cunctatious*} at t_4.
- He ends his decision problem at time $t = t_5$, by selecting *cunctatious*, which is the most apposite word available in the set: {*delay, late, procrastinate, tarry, cunctatious*}.

Although highly stylized, the example is a useful aid for helping to clarifying several different categories of knowledge that are important in our paper.

The first thing to note is that Dougal's perception of the world is limited (i.e., it is not all of $\mathbb{A}$). At date $t = 0$ his *latent or percipient* knowledge is:

$$\mathcal{A} = \{\textit{late, delay, procrastinate, tarry}\}$$

[20]According to the *Oxford English Dictionary* (2nd Ed.), the last recorded usage of *gele* was in 971 A.D.

Percipient knowledge describes what he could conceivably know at date $t = 0$, given his experiences up to that point. Thus, even though it did not initially occur to him at $t = 0$, he is quite familiar with the meaning of "procrastinate." In contrast, it is inconceivable for him to even "think of" *cunctatious* given the state of his initial knowledge.[21]

In what follows, we denote the agent's *active* knowledge set by $A(t)$. This set includes the actions the agent recognizes are feasible at date t. In this setting, "recognize" simply means that the decision maker *understands* that he can: (i) select one of the actions $a \in A(t)$, and (ii) by doing so, he can end his decision problem, and accrue an (expected) payoff of $u(a)$. Although his latent knowledge $\mathcal{A}$ is quite extensive, his active knowledge (A_0) is initially quite meagre; it consists only of $A_0 = \{$*delay, late*$\}$. However, by the time he makes his decision at $t = t_5$, his active knowledge is $A(t_5) = \{$*delay, late, procrastinate, tarry, cunctatious* $\}$.

In the example, Dougal augments $A(t)$ via two sources. First, he can "think" of alternative words through a (costly) process of "fact free learning." Alternatively, he can employ a technology – a dictionary – that can remind him of words he already knows (*tarry*) and *discover* new words altogether (*cunctatious*).[22]

Dougal could end his decision problem at its inception by picking either *late* or *delay.* However, he chooses not to do so, and, instead, searches for alternative words by "thinking" and by acquiring more information from his dictionary. Although both of these activities are costly, the reason he undertakes them is because he understands there are obvious benefits to be gleaned from doing so. Hence, although (for reasons that are all too obvious) he cannot know the words (i.e., actions) he will perceive prior to thinking of them — and which one he will ultimately choose — he acts as if he knows the *payoffs* he expects will accrue from search. In other words Dougal is appraised of economically relevant aspects of his own ignorance. In the example, Dougal makes a decision, at date: t_5 and ends his decision problem (he picks *cunctatious* $\in A(t_5)$). At that point, he deems the (expected discounted) costs of further search to exceed the expected benefits

[21] Of course, like the monkey at the proverbial typewriter, he might stumble across the string of symbols "*cunctatious*" quite by accident. In no way can it be said, however, that he has latent knowledge of its meaning. Thus, for Dougal, "cunctatious" is as likely to be synonymous with "procrastinate" as it is with any other (unrecognized) word in the English language.

[22] We use "discover" in its modern sense:

> *"To obtain sight or knowledge of (something previously unknown) for the first time; to come to the knowledge of; to find out.""* [entry 8, The Oxford English Dictionary, 2nd ed., 2002.]

2.1 The Formal Model

We analyze a simple decision problem in which an agent accrues a payoff – and his decision problem ends – after he selects a single action.[23] The main assumptions are presented below.

The Environment.

Time is continuous, and is indexed by $t \geq 0$. We consider an infinitely-lived decision maker who discounts the future at the rate ρ. We interpret discounting to mean that the agent recognizes that he is indifferent between two options that give him utility: $u = u(t)$ at date $t \geq t'$ and utility: $\exp[-\rho(t - t')]u$ at date t'. The basic obstacle that confronts the agent is that he is not omniscient. This means that, unlike the fabulous *homo economus*, he is incapable of grasping all of the logical implications of the knowledge he already possesses.

Let $\mathbb{A}$ represent all of the feasible actions available to the individual.[24] The agent *makes a decision* if he picks one of the actions belonging to $\mathbb{A}$. After making a decision, his payoff is realized and his decision problem is assumed to end. We assume that the individual has a well-defined utility function: $u\colon \mathbb{A} \to \mathbb{R}$, and denote the utility derived by selecting action $a \in \mathbb{A}$ by $u(a)$.

The actions we have in mind encompass a wide array of economic possibilities. They might represent a feasible consumption bundle, a strategy, or even the maximal utility derived by consuming a bundle of goods from some feasible set. Alternatively, the action might represent the choice of a lottery, in which case $u(\cdot)$ is the individual's expected utility. Specific interpretations of potential actions a, and the action space $\mathbb{A}$ are given in subsequent applications of the model.

Our goal is to study environments in which the decision maker cannot comprehend all the feasible actions that are available to him. In other words, he does not "know" every element that belongs to the set $\mathbb{A}$. To this end, let $\mathcal{A} \subset \mathbb{A}$ denote the agent's latent knowledge. This set includes actions that the individual could *possibly* think of given the information he possesses and his experiences at date $t = 0$. At a given point in time, the agent may understand only a small part of $\mathcal{A}$. Accordingly, let $A(t) \subset \mathcal{A}$ denote the actions that the individual *actually* does perceive at date t, and $A_0 \equiv A(0)$ the initial perceptions of his environment. The decision maker *perceives* an action a if: (i) he understands that he can select a, and (ii) he (expects) a payoff $u(a)$ from doing so.

[23] It is simple enough to extend our framework to examine the case in which the agent makes a sequence of decisions.

[24] In practice, $\mathbb{A}$ may depend upon the state of nature. Indeed, many papers have examined the difficulties of writing complete contracts in situations in which there are of unforseen states of the world (see for example, Hart and Moore (1999)). In the environment considered here, we assume (for simplicity) that there is a single state of the world. Nevertheless, each agent faces a formidable decision problem since he is unaware of his complete action space.

In order to focus on "fact-free" — as opposed to "factual" — learning, in this paper we impose:

Assumption 1 *The individual's latent knowledge, $\mathcal{A}$ equals his action space: $\mathcal{A} = \mathbb{A}$.*

The force of Assumption 1 is that the agent's decision making process is not hindered by a lack of information *per se*; in principle he knows everything he "needs" to know to determine the action that maximizes his utility. The trouble is that it is difficult for him to process this information, and to comprehend the choices that are actually available to him. Under Assumption 1, rational man would instantly recognize all of the actions in $\mathbb{A}$ and immediately select the best one.

In general, what an agent perceives, and what he could feasibly perceive depend in a complex manner upon his lifetime experiences up to that point. In this paper, we eschew modelling this process directly. Instead, we draw a Rubicon. On one side of the line, we content ourselves with an *exogenous* description of what the individual: *does perceive* (A_0), and is *capable* of perceiving ($\mathcal{A}$). On the other side if it, we posit an *endogenous* process whereby the individual gradually updates the actions that he comprehends. Thus, the agent's *initial* perceptions of possible courses of action are important *exogenous* characterizations of both him and his environment. Nevertheless, the agent is confronted with an important, and, indeed, non-trivial, decision problem regarding how he will update his perceptions of the world in which he lives. Critical to our approach is the viewpoint that the "thinking process," whereby the decision maker recognizes possible actions is costly, time consuming, and desultory. This description is, of course, in accordance with common experience and observation. Everyone finds the task of deducing valid conclusions from "known" premises to be a time consuming endeavour; some even regard it to be a painful one. Even the apparently rudimentary task of recollecting past events often entails running a similar gauntlet.

We assume that at each instant, t, the individual always understands, that the actions: $a(t)_0 = 1$ ("making a decision") and $a(t)_0 = 0$ ("deferment") are available to him. We assume that the specific choice of activity $a(t)_0 \in \{0,1\}$ is without direct bearing on his utility.[25] By setting: $a(t)_0 = 1$, the agent picks a feasible action: $a \in A(t)$. After doing so, his decision problem ends, and he accrues a utility $u(a)$. In contrast, if he sets: $a(t)_0 = 0$, he elects to *not* select an action from $A(t)$ and hence, defers making a decision. Obviously, he *must* set $a(t)_0 = 0$ if $A(t) = \varnothing$., indicating that he perceives there are currently no other options available to him. However, he also *might optimally* set $a(t)_0 = 0$. Indeed, the central theme explored in this paper is that by setting $a(t)_0 = 0$ — and hence postponing making an immediate decision — the agent can uncover new ideas, which serve to enhance the set of actions that he perceives are feasible: $A(t)$.

[25]Setting the agent's instantaneous utility to zero from picking the action: $a(t)_0 = 0$, is a normalization. Moreover, any cost, or benefit, of setting: $a(t)_0 = 1$ can be rolled into the payoff $u(a)$.

In order to avoid tedious issues that will subsequently arise (essentially) because of the difference between statistical sampling "with" and "without" replacement, consider:

Assumption 2 *(The sets $\mathcal{A}$ and A_0.) The agent's latent knowledge $\mathcal{A}$ contains an uncountable number of distinct elements. The set A_0 includes a finite number of distinct elements. Furthermore, provided $A \neq \emptyset$, then $a(0)^* \equiv \arg\max_{a \in A_0}$ is unique.*

The individual initially "perceives" that only those actions that belong to A_0 are feasible. The extent of the decision maker's initial *ignorance* is captured by: $\mathcal{A} \backslash A_0$, which is the difference between the "reality" of the situation and the world as he perceives it. In view of Assumption 2 the agent initially comprehends only a tiny (negligible) part of the reality that confronts him. Provided that the initial perceived set of actions, A_0, is not empty, then $u^*(0) \equiv \max_{a \in A_0} u(a)$ exists. The assumption that there is a unique maximizing choice, $a(0)^*$, is inconsequential. In addition to Assumption 2, consider:

Assumption 3 *(The Utility Possibility Set U) The set of feasible utilities $U = \{u(a) : a \in \mathcal{A}\}$ equals the compact interval: $U = [0, \bar{u}]$, where: $\bar{u} \equiv \max_{a \in \mathcal{A}} u(a) > 0$.*

Notice that since U is a compact interval of real numbers, defining a probability measure on it (as we do later) is routine.

The Thinking Process

By exerting effort, the agent can think of new ideas that point to alternative actions that augment his active action set $A(t)$. The snag is that thinking is a costly, time consuming, and inherently desultory process: each new idea arrives at random points in time, and utilities can differ markedly from one idea to the next. Assumption 4 describes the agent's thinking process,

Assumption 4 *(Thinking).*
(i) The individual can choose the intensity ("effort"), $e(t) \geq 0$ which allows him to think of ideas belonging to $\mathcal{A}$.
(ii) By devoting effort e to thinking, the individual realizes new ideas as independent draws from a Poisson distribution with parameter $\lambda(e)$. We assume that $\lambda(e)$ is non-decreasing, and concave in e. Furthermore λ satisfies the boundary conditions: $\lambda(0) = \lambda_0 \geq 0$ and $\lim_{e \to 0} \lambda'(e) = \infty$.
(iii) The flow cost of formulating new ideas is: $c = c(e)$, which is an increasing and strictly convex function of e, and which satisfies the boundary condition: $c(0) = 0$.

According to parts (i) and (ii), by exerting greater effort $e(t)$, (i.e., by "concentrating"), the decision maker can think more rapidly of alternative actions. Nevertheless, according to part (iii) cognition is a costly process.[26] The restrictions on $\lambda(\cdot)$, and $c(\cdot)$ simply ensure a unique interior solution for the optimal choice of effort e.

Given a constant value of effort: $e \equiv e(\tau)$, the time between successive ideas is exponentially distributed with a constant mean equal to $1/\lambda(e)$. Each new idea provides the agent with a prospective payoff $u(a)$. The Poisson distribution succinctly captures the agent's limited rationality. The agent is capable of having only one idea at a time, and, generically, after having one idea he must wait a positive time interval before having the next.

Given his effort intensity level: $e(\tau)$, the decision maker can think of at most a countable number of new actions over a given finite interval of time $\tau \in [0,t]$. It follows from this, and from Assumption 2, that $A(t)$ can have at most a countable number of elements. Inconsequently, we assume that: $A(t) \subseteq A(t')$, whenever $t' \geq t$, implying that delay does not restrict the set of actions available to the individual.[27]

However, although the individual's rationality is seriously compromised by his failure to comprehend $\mathcal{A}$ we assume that he understands the (expected) value of formulating new ideas. Indeed, we assume that he correctly understands all of the relevant aspects of the *distribution of utilities* he can expect from uncovering alternative actions. In other words, the agent acts *as if* he is appraised of the economically relevant attributes of his own ignorance. To operationalize this notion, we assume that as the agent "thinks," he believes he is sampling *utilities* from a known distribution that is defined on the support U. Thus,

Assumption 5 *(The rewards from thinking).*
(i) The decision maker perceives that the utility, u, he generates by thinking of each new idea is an i.i.d. draw from the given distribution function: $u \sim F(u)$, which is defined on the support: $u \in U$. The distribution function, $F(u)$ satisfies: $F(0) = 0$ and $F(\bar{u}) = 1$. The function $F(u)$ is differentiable with respect to u and satisfies: $f(u) \equiv F'(u) > 0$, everywhere on the interior of the interval U.
Remark. *For the moment, we assume that the distribution $F(\cdot)$ does not depend on the agent's effort level e. However, it is plausible that by thinking more intensively the agent not only generates ideas at a faster rate, but also ones that are better "on average." Section 3.1 considers this possibility.*

Throughout, we assume that the agent acts as if he "knows" the true distribution $F(u)$. It is to this extent that the individual acts as if he is fully aware of all of the economically

[26]As Clark (1918, p.23) put it, almost ninety years ago,

> *"Decision involves effort of attention, and this effort cannot be sustained beyond a few seconds at a time, nor repeated without limit — a fact which suggests the using up by fatigue of a limited capacity for this kind of mental act. [T]he exhausting character of choice, so exhausting that it becomes the part of wisdom to choose to yield up our prerogative of choice, save in the things we hold most important."*

[27]This would not be the case if the agent either possesses an imperfect memory, or if his action set is explicitly time dependent.

relevant aspects of *what he does not know*. The technical assumptions on the distributions are designed to simplify the subsequent arithmetic.

Together, the function: $\lambda(\cdot)$ and the distribution $F(\cdot)$ provide a useful and natural characterization of an agent's intelligence and his creativity. Consider two agents, indexed $j = 1,2$. We define agent 1 to be "more intelligent" than agent 2 if:

- For each $u \in U$, $F(u)_1 = F(u)_2$, and for each $e > 0, \lambda_1(e) > \lambda_2(e)$,

and to be *more creative* if:

- for each $e \geq 0$, $\lambda_1(e) = \lambda_2(e)$, and $F_1(u) \leq F_2(u)$ for all $u \in U \equiv U_1 \cup U_2$.

The first condition says that, "on average," both agents think of the same quality ideas, but player 1 thinks of them faster. The second, says that although both of them think at the same speed, player $1's$ ideas are "better" (in the sense of first-order stochastic dominance).

The decision maker described here, should be compared with traditional "rational man." In our approach, the agent perceives only a small part of what he can potentially "understand" given the information available to him. However, he recognizes that by "thinking" that he can slowly enlarge the set of choices he perceives are available to him, and he understands the rewards (and costs) of doing so. In this sort of environment, it is conceivable that agents may differ markedly in their ability to generate new and valuable ideas under ostensibly identical economic circumstances. Moreover, (as we shall see when we discuss advertising) during this slow process of deliberation *other agents* can "suggest" profitable courses of action that the individual could have - but did not - think of. In other words, other players can strategically "influence" people. In contrast, *homo economus* perceives that: $A_0 = \mathcal{A}$ right off the bat, and immediately picks the "best" option available to him. In this setting, there is no obvious criterion whereby one *homo economus* agent can be ascribed to be more intelligent than another. Likewise, there is little scope, for players to "influence" each others' behavior without actually proffering *useful* (i.e., new) information.

Despite the decision maker's limited comprehension of the world in which he lives, as formulated his decision problem is actually quite simple. Formally it is a sequential "search problem," with an optimal stopping rule. Thus, at each instant he must decide whether he will either: (a) make a decision, by selecting an option from $A(t)$ or (b) continue thinking of new ideas that belong to: $\mathcal{A}\backslash A(t)$. Evidently, as discussed in the next section, the solution to this problem depends upon a variety of factors, such as his impatience ρ, the value of the best action available to him in $A(t)$, the rate at which he can think of/discover ideas, and the (distribution of) values associated with any such action he takes.

3 Optimally " Thinking."

In this Section, we analyze the agent's optimal "thinking process," whereby he generates new ideas concerning the actions available to him. At date $t = 0$, when his decision problem begins, the agent's knowledge is characterized by the pair ($\mathcal{A}$,A_0), where $\mathcal{A}$ is the set of actions that the agent can possibly conceive of in the light of his experiences up to that point, and A_0 is the set of actions he actually does perceive to be possible. Provided that $A(t) \neq \varnothing$, the agent recognizes that, at each point in time $t \geq 0$, he can set: $a(t)_0 = 1$, and terminate his decision problem, by picking some action, $a \in \mathcal{A}$ In what follows, let:

$$u^*(t) = \max_{a \in A(t)} u(a), \qquad \textit{and,} \qquad a^*(t) \equiv \arg\max_{a \in A(t)} u(a)$$

which correspond to the maximum utility and optimal action in the set $A(t)$.[28]

Alternatively, by setting $a(t)_0 = 0$ he can defer making an immediate decision, and "think" about alternative courses of action that he does not currently see: $\mathcal{A} \backslash A(t)$. By exerting effort, e, (at an instantaneous cost: $c(e)$) he is cognizant that he will uncover new ideas at the rate $\lambda(e)$, and that the realized value of each new idea is an *i.i.d.* draw from a distribution, $F(u)$ defined on the support U. Thus, at each point in time the decision maker has two basic choices. On the one hand, he can "stop" thinking and make a decision, by selecting $a(t)_0 = 1$ (and choose some $a \in A(t)$). On the other hand, he can "continue" thinking, by choosing $a(t)_0 = 0$, by selecting an effort level $e(t) \geq 0$.

Since, under Assumption 2, the perceived action set: $A(t)$, contains at most a finite number of elements, and since $\mathcal{A}$ contains an *uncountable* number, the individual under consideration is, and remains, in a strict sense "quite ignorant" about the true nature of reality. Yet, despite this, at each point in time he is fully aware of all of the economically relevant aspects of his environment.[29] Consequently, following Lipman (1991), we take the view that the agent maximizes his expected discounted utility. In other words, he behaves optimally in the environment he perceives.

In the following Lemma, we use $a^*(t)$ and $u^*(t)$ to characterize the individual's behavior.

Lemma 1 *(Optimal Decision Making). Under Assumptions 1–5, the decision maker's behavior is completely characterized by a unique and constant reservation utility level: $s^* \in (0, \bar{u})$, and a unique and constant level of effort, $e^* \geq 0$, in which:*
(i) if $u^*(0) \geq s^*$ he chooses $a^*(0)_0 = 1$ and the action $a^* \in A(0)$.
(ii) if $u^*(0) < s^*$ he chooses $a^*(t)_0 = 0$ and $e(t) = e^* > 0$ for all $t \in [0, \tau)$, where τ is the first time that he has an idea that generates utility: $u^*(\tau) \geq s^*$. At time τ he sets: $a^*(\tau)_0 = 1$ and picks the (unique) optimal action $a^* \in A(\tau)$ defined by: $u(a^*) = u^*(\tau)$.

Proof. All proofs are presented in the Appendix. ∎

The choice of s^* resembles the familiar "reservation wage property" encountered in sequen-

[28]The maximum is always well defined, since (i) $A(t)$ contains at most a finite number of elements, and (ii) if $A(t) = \varnothing$, we have normalized: $u(a) = 0$. Moreover, under Assumption 2 the choice of $a^*(0)$ is unique.

[29]In particular, he knows the maximum payoff $u^*(t) \equiv \arg\max_{a \in A(t)} u(t)$ available to him from choosing $a(t)_0 = 1$, as well as the (expected) discounted costs and benefits of picking $a(t)_0 = 0$ and selecting any $e(t) \geq 0$.

tial (job) search theory.[30] In this framework, the individual sequentially gathers *information* about wage offers until he encounters the first "acceptable" offer, which is one that exceeds his reservation value w^*. In this paper, the individual does not acquire additional information *per se*. Instead, he engages in a process of "fact free" learning, and terminates his search efforts only after formulating an idea that provides him with a level of utility that exceeds his reservation value.

The individual may immediately reach a decision without attempting to think of new ideas. Intuitively, this will the case if the initial set of possible actions, A_0 includes an action that provides a "sufficiently high" (i.e., exceeding s^*) level of utility. In contrast, if the prospects available in A_0 are suitably "bleak," the agent actively seeks to think of more lucrative possibilities. He does this by "thinking" at the constant rate $\lambda(e^*)$ and terminates his thought process, by making a decision, once he has stumbled upon a sufficiently good idea (i.e., one that provides him with a payoff in excess of the reservation value s^*). Notice that the agent is an "optimizer," even though to the casual observer it might look at first glance that, in Simon's terminology, he is a "satisficer." I.e., he accepts the first action that gives him a utility of at least s^* - which is his *aspiration level*. Critically, in our approach the agent constantly weighs the costs of postponing his decision against the benefits he can expect from doing so. What is more, the aspiration level, s^*, is endogenous to the model.

During the time he is thinking of new ideas (i.e., over the interval $[0,\tau)$) his choice of effort is constant: $e(t) = e^*$. This is despite the fact that, as formulated, his decision problem is *non-stationary.* (More specifically, $u^*(t) \equiv \max_{a \in A(t)} u(a)$ weakly increases with t as new elements are added to $A(\cdot)$ and none are removed). The reason, of course, is that if: $u^*(t) < s^*$, then it is 'as if' the agent seeks to maximize his expected discounted utility by choosing $e(t) \geq 0$ subject to the constraint that $A(t) = \varnothing$. From the properties of the distribution function, $F(u)$, this problem *is* stationary. Intuitively, even though the options in $A(t)$ (weakly) improve through time, if none of them is currently "good enough," then they are all equally irrelevant.

It is possible, and desirable, to formulate the individual's decision problem recursively as this provides greater insight into its structure and major properties. To this end, let $e > 0$ and $s \in (0,\bar{u})$ be given, where e is effort and s a stopping value. The individual arrives at a decision (i.e., terminates his thought process) if and only if he thinks of an idea that gives him a utility of at least s. For convenience let:

$$E(s) \equiv \frac{\int_s^{\bar{u}} udF(u)}{1 - F(s)},$$

which is the mean of u conditional upon the event that $u \geq s$. Let $v(e,s)$ denoted the agent's expected discounted utility given the choices of e and s. The value, $v(e,s)$ satisfies

[30] Rogerson, Shimer, and Wright (2005) provide an excellent survey of the most recent literature. The earlier generations of search literature are surveyed by Lippman and McCall (1976), and Mortensen (1986).

the Bellman equation:

$$\rho v(e,s) = -c(e) + \lambda(e)(1 - F(s))\{E(s) - v(e,s)\} \tag{1}$$

The discount rate, ρ, times the value $v(\cdot)$ is the flow utility the agent derives by continuing to think of new actions. It equals the flow disutility from effort, $-c(e)$, plus the the likelihood that he thinks of an *acceptable* idea $\lambda(e)(1 - F(s))$, times the expected capital gain in this event $(\{E(s) - v(e,s)\})$.[31] Rearranging equation (1) and solving for $v(e,s)$ gives:

$$v(e,s) = \frac{-c(e) + \lambda(e)(1 - F(s))E(s)}{(\rho + \lambda(e)(1 - F(s)))} \tag{2}$$

Here, $1/(\rho + \lambda(1 - F))$ corresponds to an "effective" discount rate. Intuitively, the individual places less weight on the future the: (i) higher is his rate of time preference, ρ, (ii) the *faster* he expects to reach a decision (which increases with his thinking speed λ and as he becomes less "picky" - i.e., as F falls).

Using equation (2), the optimal choice of effort, e^*, is governed by the condition:

$$c'(e^*) = \lambda'(e^*)(1 - F(s^*))\{E(s^*) - s^*\} \tag{3}$$

The left-hand side of equation (3) is the marginal cost (disutility) of effort; the right-hand side is its marginal benefit. It equals the *incremental* flow probability that additional effort will lead to an acceptable idea, $\lambda'(e^*)(1 - F(s^*))$, times the expected capital gain in this event $(E(s^*) - s^*)$. In other words, the agent equates the costs and benefits of effort at the margin.

Let s^* and e^* denote the agent's optimal choices of s and e. Define the corresponding maximized expected discounted value to the agent of continuing to think of new and better ideas by: $V = v(e^*,s^*)$. Since V is the value of continuing his thought process, the individual arrives at a decision, and takes an action, if and only if it provides him with utility: $u' \geq V$. Yet, s^* is the optimal stopping value. It too says that the individual will arrive at a decision iff it provides utility $u' \geq s^*$. It follows that s^* is implicitly defined by the fixed-point: $s^* \equiv v(e^*,s^*) = V$. Hence, at a maximum:

$$s^* \equiv v(e^*,s^*) = \frac{-c(e^*) + \lambda(e^*)(1 - F(s^*))E(s^*)}{(\rho + \lambda(e^*)(1 - F(s^*)))} \in (0,\bar{u}) \tag{4}$$

Notice that the decision maker's inability to immediately perceive all of the opportunities available to him, $\mathcal{A}$, is costly in three separate respects. First, he must exert costly effort to think of possible actions.[32] Second, even in the absence of this direct effort cost, the

[31] The capital gain equals the expected value of the idea, $E(\cdot)$, *minus* the option value of making a decision and, hence, *discontinuing* the search for better ideas.

[32] The expected discounted sum of these costs is: $c/(\rho + \lambda(1 - F))$.

slow process of deliberation entails costly delay since the agent is impatient.[33] Finally, when combined, the first two costs lead him to settle for, and hence expect, a payoff that is strictly less than the maximum feasible payoff $\bar{u}$.[34]

Summarizing, the individual's decision problem is characterized by the pair (s^*,e^*) which uniquely solve equations (4) and (3). The individual immediately makes a decision iff $u^*(0) \geq s^*$. Otherwise, he thinks of alternative actions at the rate $\lambda(e^*)$, and continuous to do so until he finds one to his satisfaction.

Properties of the solution.

Suppose that: $\lambda(e) = \lambda_0\hat{\lambda}(e)$, where $\lambda_0 > 0$. The solution to the individuals problem is characterized by a reservation utility level, $s^* = s(\rho,\lambda_0)$ and an optimal level of effort, $e^* = e(\rho,\lambda_0)$. From previous remarks, the parameter λ_0 provides a measure of the innate agent's intelligence. *Ceteris paribus* the greater the value of λ_0 the greater is the rate at which the agent can think of new ideas. The parameter ρ measures the extent of the agent's impatience. The agent's expected decision making (i.e., "computational") speed is:

$$\sigma^* = \sigma(\rho,\lambda_0) = \lambda(e^*)(1 - F(s^*))$$

where: $1/\sigma^*$ is the expected time it takes for him to make a decision. Intuitively, the speed with which an individual arrives at a decision depends positively upon how quickly he can think of new ideas, $\lambda(.)$, and negatively upon how "picky" he is, once he has had one, $1 - F^*$. The (average) value of the action the agent ultimately selects is: $E(s^*) = E\{u|u \geq s^*\}$. The effects of differences in temperament ("patience"), ρ, and innate intelligence, λ_0, are easy enough to work out. Using subscripts to denote partial derivatives, consider:

Proposition 1 *(Impatience, ρ, and intelligence, λ_0). The individual's decision problem satisfies:*
(i) Impatience: $s^*_\rho < 0, e^*_\rho > 0$, $\sigma^*_\rho > 0$, $E(s^*)_\rho < 0$.
(ii) Intelligence: $s^*_{\lambda_0} > 0$, $e^*_{\lambda_0} > 0$, *and, if the density $f(u)$ is log-concave, then:,* $\sigma^*_{\lambda_0} > 0$.

According to part (i), more *impatient* (higher ρ) players are *less* picky (their s^* is lower) and think harder (their e^* is higher) than less impatient ones. This is because their impatience drives them to seek a quick resolution to their problem (recall the average decision speed, σ^*, increases in e^* and decreases in s^*). However, the effect of the rush to make a decision is that the (average) *ex post* payoff of impatient players, $E(s^*)$, is lower than that of more patient ones.

We now turn to the effects of differences in the player's innate "intelligence" level, which is parameterized by λ_0. An increase in λ_0 increases the the player's "pickiness," by raising

[33]This is captured by: $\delta \equiv \lambda(1-F)/\{\rho + \lambda(1-F)\} < 1$.
[34]Hence, $E(s) < \bar{u}$.

his reservation value s^*. (This implies their expected *ex post* payoffs $E(s^*)$ are also higher). More intelligent players also devote more effort toward thinking of new ideas (e^* increases with λ_0).

The agent's expected decision making speed is: $\sigma^* \equiv (1-F^*)\lambda_0\hat{\lambda}(e^*)$. Consequently, in view of their "pickiness" there is a tendency for more intelligent players to take longer to arrive their decisions (i.e., $F^* \equiv F(s^*)$ *increases* with s^*). However, more intelligent players also devote more effort towards reaching a decision, which speeds up their decision making process (i.e., $\lambda^* \equiv \lambda_0 \cdot \lambda(e^*)$ rises). Provided that their payoff density $f(\cdot)$ is log-concave, then according to the Proposition, this latter effect dominates so that the more intelligent the agent the faster he arrives at a decision.[35]

In contrast, if $f(.)$ is *not* log-concave, it is conceivable that: $\sigma^*_{\lambda_0} < 0$, indicating that highly intelligent agents might take the longest time to think, before they arrive at their decisions. The explanation is as follows. For simplicity, suppose that that $\lambda_0\lambda(e) \equiv \lambda_0 > 0$, which is independent of e. Suppose further that $f(.)$ is log convex and that: $E_{s^*} \equiv \partial E(u|u \geq s^*)/\partial s^* > 1$, indicating that a one unit increase in the reservation value s^* increases the expected value of reaching a decision by *more* than a unit. The value of continuing to think of new ideas is:

$$V = \delta E(s^*)$$

where,

$$\delta \equiv \frac{\lambda_0(1-F^*)}{(\rho + \lambda_0(1-F^*))}$$

Here, δ is akin to an effective "discount factor." Notice that, *given* s^*, the term δ increases with λ_0. Hence — given s^* — an increase in λ_0 raises the expected present value of the reward $E(s^*)$. However, if $E_{s^*} > 1$, the agent could conceivably respond to an increase in λ_0, by raising s^* so much — to take advantage of the fact that $E_{s^*} > 1$) — that his average decision speed: $\sigma^* \equiv (1-F^*)\lambda_0$ declines.

Under the conditions of the Proposition, more intelligent players reach better decisions (their $E(s^*)$ is higher) and do so more rapidly (their σ^* is higher) than do less intelligent ones. This raises interesting questions concerning the limiting properties of innate intelligence levels on the agent's optimal behavior. More specifically, we know that the reservation value s^* increases with λ_0. We also know that it is bounded above by $\bar{u}$, implying that it converges to some limit: $\lim_{\lambda_0 \to \infty} s^*(\lambda_0) = \bar{s}^*$. Heuristically, in the limit, as $\lambda_0 \to \infty$, the decision maker can think "infinitely quickly" of alternatives actions. In this case, does his expected payoff, $E(s^*)$, converge to $\bar{u}$ and does he make a decision instantly?

[35] In the context of a job search model (with an endogenous reservation wage), Burdett and Ondrich (1985) show that log-concavity of the wage offer distribution is sufficient for an increase in the job contact rate to induce workers to *accept* jobs at a faster rate. However, they do not consider the effects of changes in workers' search intensities.

Consider,

Proposition 2 *(The limiting properties of $s^*(\lambda_0)$ and $\sigma^*(\lambda_0)$. If $\lambda(e) = \lambda_0\hat{\lambda}(e)$ and if $f(.)$ is log-concave then:* $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{s}^* = \bar{u}$ *and* $\lim_{\lambda_0\to\infty} \sigma^*(\lambda_0) = \infty$.

From this vantage point, the limit $\lambda_0 \to \infty$ is one in which the agent can think of possibilities so quickly that his behavior (and utility) converges (in expectation) to that of "rational economic man."

3.1 The Quality of Ideas

It is conceivable that by exerting more effort, the decision maker may not only think of ideas faster, he might also think of ones that are better on average. In order to more clearly differentiate between the two effects of "speed" and "quality," assume that: $\lambda(e) = \lambda_0$ (independent of e). Suppose now that the distribution function (defined on the support U is $F(u\,|\,e)$, which depends upon the agent's effort, $e \geq 0$. In addition to the properties already laid out in Assumption 5 consider:

Assumption 6 *(The properties of $F(u\,|\,e)$).*
(i) (First-order-stochastic dominance (FOSD), and convexity of the distribution function (CXDF)). Let $F_e \leq 0$ and $F_{ee} \geq 0$, where: $F_e \equiv \partial F(u\,|\,e)/\partial e$ and $F_{ee} \equiv \partial^2 F(u\,|\,e)/\partial e^2$, for all $u \in U$. Assume that for any $s < \bar{u}$: $\int_s^{\bar{u}} F_e du < 0$ and $\int_s^{\bar{u}} F_{ee} du > 0$.
(ii) (Log concavity) The density function, $f(u\,|\,e)$ is log-concave (LC).
(iii) The distribution function, $F(.)$, satisfies the monotone likelihood ratio property (MLRP): $f_e/(1-F)$ is decreasing in u.

The individual's decision problem is once again characterized by a reservation utility, s^*, and a level of effort e^*. This time, these values satisfy the asset-value equation:

$$\rho s^* \equiv -c(e^*) + \lambda_0(1 - F(s^*|e^*))(E(s^*,e^*) - s^*) \tag{5}$$

and the first-order condition:

$$0 \equiv -c'(e^*) + \lambda_0 E(s^*,e^*)_e \tag{6}$$

where the short-hand:

$$E(s^*,e^*) \equiv E\{u|u \geq s^*,e^*\} \equiv \frac{\int_{s^*}^{\bar{u}} udF(u|e^*)}{\int_{s^*}^{\bar{u}} dF(u|e^*)}$$

and $E(s^*,e^*)_e \equiv \partial E(s^*,e^*)/\partial e$ is used. Part (i) of Assumption 6 is sufficient to ensure that equations (5) and (6) uniquely characterize s^* and e^*.

According to equation (5), the right hand side is the net flow value of continuing to thinking of new ideas. It equals minus the flow cost of search, $-c(s^*)$, plus the flow probability that the agent has an acceptable idea: $\lambda_0(1-F(s^*|e^*))$ times the capital gain in this event: $E(s^*,e^*)-s^*$. Equation (6) indicates that the agent equates the marginal costs and benefits of additional effort at the maximum.

The following Proposition characterizes the main features of the solution to the decision maker's problem.[36]

Proposition 3 *(The solution $s^*(\lambda_0), e^*(\lambda_0)$, $\sigma^*(\lambda_0)$).*
(i) If Part (i) of Assumption 6 holds, then: $s^*_{\lambda_0} > 0$, and $e^*_{\lambda_0} > 0$.
(ii) If Parts (i)-(ii) of Assumption 6 hold, then: $\sigma^*_{\lambda_0} > 0$, indicating that the agent's decision speed: $\sigma^* \equiv \lambda_0(1-F(s^*|e^*))$ increases with λ_0.
(iii) If Parts (i)-(iii) of Assumption 6 hold, then: $\partial E(s^*,e^*)/\partial\lambda_0 > 0$.

An increase in the rate at which the individual can think of alternative actions, λ_0, makes him more selective in the ones he is willing to act upon. This explains why: $s^*_{\lambda_0} > 0$. The increase in λ_0 also raises the effort, e^*, he devotes to thinking for two simple reasons. First, there is a direct benefit of the increase in λ_0, as the agent anticipates arriving at a decision more rapidly (and he is impatient). Second, the increase in his reservation value s^* implies that his expected utility rises, conditional upon making a decision. Both of these effects encourage additional effort at the margin.

The agent's *optimal* decision making speed, σ^*, depends upon his innate intelligence λ_0, and the choices he makes regarding both his effort, e^*, and the cutoff value, s^*. Ceteris paribus, an increase in λ_0 raises σ^* directly. However, the increase in λ_0 raises both e^* and s^*. The former of these effects tends to increase σ^*, the latter to reduce it (the agent becomes more "picky"). If the density $f(.)$ is log-concave, then the direct effect dominates; the agent's decision making speed, σ^*, increases with his intelligence λ_0. Finally, if in addition to FOSD,CXDF, and LC the distribution function satisfies the monotone likelihood ratio property (MLRP), then the utility the agent can expect — *conditional on making a decision*, $E(s^*,e^*)$, — also increases with his innate intelligence λ_0. From part (*ii*) of the Proposition, both s^* and e^* increase with λ_0. *Ceteris paribus*, the former effect unambiguously tends to raise $E(s^*,e^*)$, which is, after all, the conditional expectation that $u \geq s^*$. Under the MLRP $\partial E(.)/\partial e > 0$, implying that the latter effect - operating through e^* - increases it as well.

The reason that it is necessary to wheel in the technical "heavy artillery" — i.e., the MLRP — to prove this result is that it is conceivable that the expected quality $E(.)$ *decreases* with e under FOSD. The agent chooses his effort level, e^*, with one eye on the "benefit:" $b(e) \equiv (1-F(s^*|e))(E(s^*,e)-s^*)$ and the other on the cost $c(e)$. FOSD simply ensures that

[36]The effects of changes in ρ are not reported. They parallel the findings already reported in Proposition 1, where $\lambda = \lambda(e)$.

$b'(e) > 0$, which, given $F_e < 0$, could be true even if $E(\cdot)$ *decreases* with e. Together, MLRP and FOSD imply that $E(\cdot)$ *increases* with e, which is then sufficient to prove the result.

As for the "limiting" properties with respect to λ_0, consider:

Proposition 4 *(The limiting properties of $s^*(\lambda_0)$, $e^*(\lambda_0)$, and $\sigma^*(\lambda_0)$). In the limit, $\lambda_0 \to \infty$, the individual's behavior is characterized by:* $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{u}$; $\lim_{\lambda_0\to\infty} e^*(\lambda_0) = 0$*; and* $\lim_{\lambda_0\to\infty} \sigma^*(\lambda_0) = \infty$.

Thus, (in expectation), he instantly selects an action that gives him $\bar{u}$. Consequently, as the agent becomes "infinitely smart," his behavior does indeed converge in expectation to that of "rational economic man."

4 Application 1: Strategic Manipulation

"Men are so simple and yield so readily to the desires of the moment that he who will trick will always find another who will suffer to be tricked." [Nicolo Machiavelli]

In this Section we use a simplified version of our framework to understand aspects of strategic environments in which one player (a "speaker") seeks to manipulate another player (a "listener") and bend him to her will, by simply proffering *suggestions*.[37] Significantly, unlike persuasion games, we assume that the suggestions tendered by the speaker are both truthful *and* complete.

We assume that, with a flow cost $c(e)$, the listener thinks of actions at the rate $\lambda_0 \cdot \lambda(e)$. The function $c(\cdot)$ is increasing and strictly convex in e, while $\lambda(\cdot)$ is strictly increasing and concave. Both functions satisfy the boundary conditions, $c(0) = \lambda(0) = 0$. As for payoffs, we assume that conditional upon formulating an idea, with probability $x \in [0,1]$, the listener thinks of one that provides him with a utility Z. With complementary probability $(1-x)$ he thinks of one that provides him with a "low" level of utility, which we normalize to zero. As for the speaker, with a flow cost $C(E)$ she thinks of ideas at the rate $\mu_0 \cdot \mu(E)$. The functions $C(E)$ and $\mu(E)$ satisfy the same conditions as those of the listener.

Upon having an idea, the speaker can costlessly and credibly communicate it to the listener. The listener can always *reject* the speaker's advice. However, once the "cat's out the bag" he cannot *ignore* it. In order to capture the essential aspects of manipulation, we assume that the speaker accrues a payoff only if the listener accepts her idea. In this event, the *listener's* payoff is z and the speaker's $Z - z$.[38] The details concerning the proposal, z, are presented below.

[37] We follow the terminology of Glazer and Rubinstein (2006).

[38] This formulation implies that the speaker always thinks of "valuable" ideas. Nevertheless, it is a harmless restriction. To see this, suppose that conditional upon having an idea, the probability that the speaker thinks of one that gives a payoff of Z is $y \in [0,1]$. In this case, the flow probability of having an *acceptable* idea is $\mu_0 \cdot \mu(E) \cdot y$. However, by simply defining, $\mu_0 \equiv \hat{\mu}_0/y$, we could carry out the analysis with $\hat{\mu}$ in place of μ, and the analysis would be identical to that presented below.

Analysis

The analysis of the model's properties is relatively straightforward. Temporarily, assume that the speaker's offer, z, and her effort, E, are given. For convenience, define the rate at which the speaker makes suggestions by $m \equiv \mu_0 \cdot \mu(E)$. Let v denote the listener's value. The relevant Bellman equation is:

$$\rho v = -c(e) + x \cdot \lambda_0 \lambda(e) \cdot [Z - v] + m \cdot I \cdot [z - v] \tag{7}$$

Here, $\rho{\cdot}v$ is the listener's flow value of contemplating possible actions. The first term on the right-hand side is the flow cost of the agent's effort. The second is the flow probability the agent formulates his own *acceptable idea* $(x \cdot \lambda_0 \cdot \lambda(e))$, multiplied by the capital gain in this event $(Z - v)$. The final term captures the influence of the speaker. With flow probability m the speaker makes a suggestion that *could* permit the listener to enjoy the capital gain $[z - v]$. Here, $I \in \{0, 1\}$ reflects whether or not the listener accepts $(I = 1)$ or rejects $(I = 0)$ the speaker's suggestion. The listener's problem is to select both I and e to maximize his welfare. As for the speaker, he understands the listener's behavior. His objective is to select E and z to maximize his own welfare.

The fundamental problem we must broach is, will the listener ever set $I = 0$, and reject the speaker's proposal? In order to answer this question, let $m \geq 0$ be given and *assume* that $I = 1$, implying that the listener *always accepts* the speaker's proposal. The listener's value (which depends upon m and z) is:

$$v(e;\ m,\ z) = \frac{-c(e) + \lambda_0 \cdot \lambda(e) Z + m \cdot z}{r + \lambda_0 \cdot \lambda(e) + m} \tag{8}$$

The listener picks a level of effort to maximize his value $v(e;\ m,\ z)$. Let $e^*(m,z)$ denote his optimal effort level and define his optimal value by $V(m,z) \equiv \max_e v(e;\ m,\ z)$.

Now suppose that the listener always *rejects* the speaker's proposal, by setting $I = 0$. Notice from equation (7) that this is identical to the situation in which he would always *accept* her proposals (by setting $I = 1$), but never actually receives any from her — $m = 0$. It follows that equation (8) nests the listener's optimal valuation in the case he always rejects her advice, under the parameter restriction that $m = 0$. Accordingly denote his value of *rejecting* the speaker's proposal by $V(0,z) \equiv V_0$.

Figure, 1 depicts the value of V_0 against the speaker's offer z. It is (obviously) a horizontal line, since (by assumption) the listener is assumed to reject the offer. The 45^o line depicts the possible value of offers, z, that are made by the speaker.

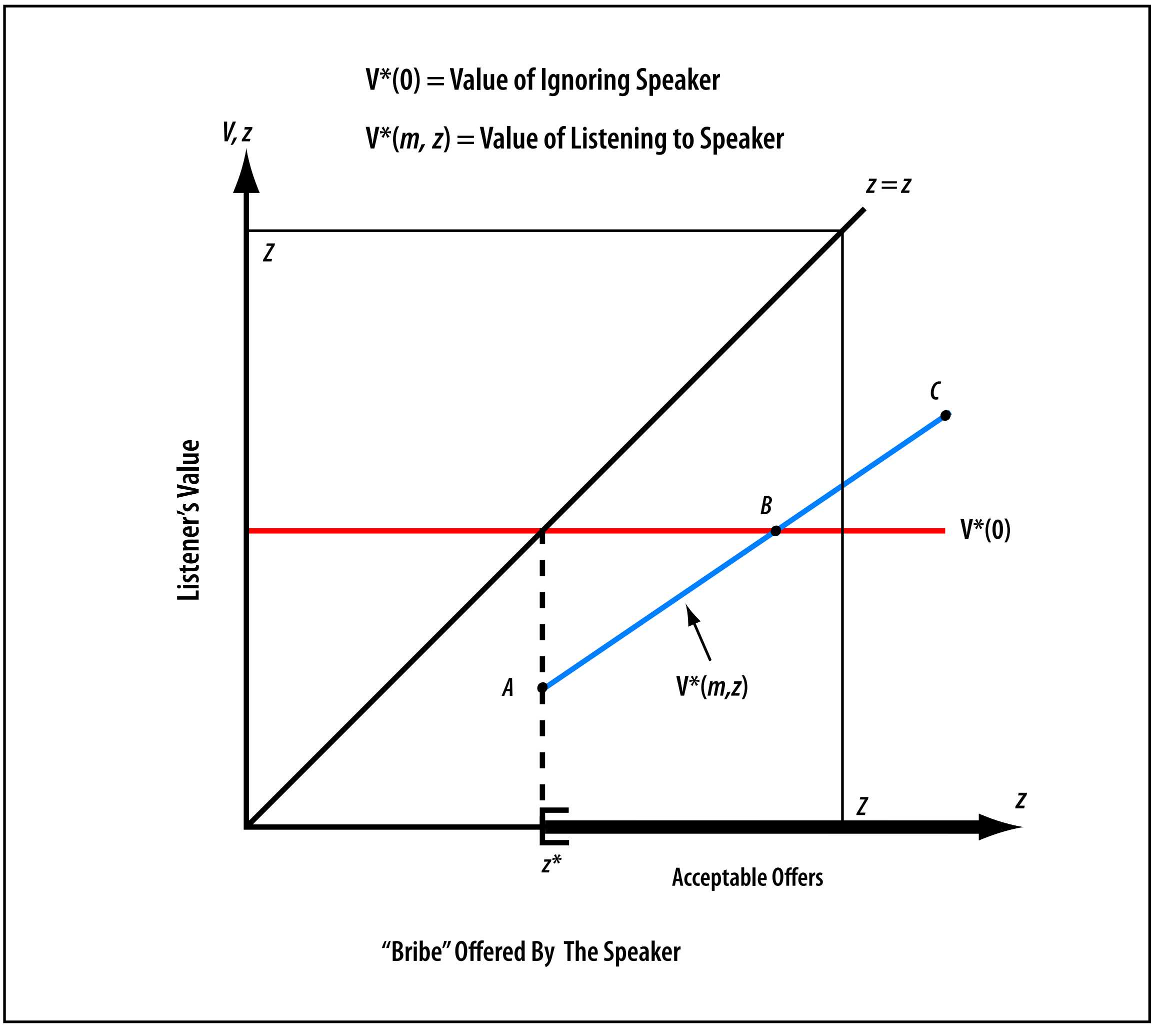

Figure 1: Strategic Manipulation

Notice that if the speaker's offer is *less* than z^* it is optimal for the listener to ignore her suggestions. The reason is that, by doing so, his expected payoff is V_0, which exceeds the value of her offer $z < z^*$. However, the snag is that it is *incredible* for him to *reject* her proposal if it lies in the range $z \geq z^*$. The reason is that having such an offer "in hand" exceeds the (expected discounted) value of Z "in the bush," so to speak.

Yet, there is more. The listener recognizes in any subgame in which the speaker offers $z \geq z^*$ he will "cave in" and accept the offer. He accordingly adjusts his own optimal thinking intensity $e^*(m,z)$ to reflect the reality of his situation. Conditional upon his optimal *credible* behavior, his valuation along the equilibrium path, $V(m,z)$, is given by the maximized value of equation 8. Consider,

Lemma 2 *(The Listener's Optimal Behavior.)*
(i) The listener's value is $V_0 \in (0,Z)$ *if he rejects the speaker's proposals.*
(ii) The listener credibly rejects any proposal made by the speaker that lies in the range $z < z^* \equiv V_0$.
(iii) It is incredible for the listener to reject proposals that lie in the range, $z \geq z^*$. *In this case his valuation is* $V^*(m,z)$.
(iv) If $m > 0$*, then: (a)* $V^*(m,z^*) < V_0$ *and (b)* $\partial V^*(m,z)/\partial z > 0$ *for* $z > z^*$

Figure 1 depicts the listener's valuation $V^*(m,z)$. It increases with z, and initially lies *below* $V^*(0)$. Crucially, it is predicated on the listener's credible beliefs concerning his own ability to reject the speaker's suggestions.

We now analyze the speaker's behavior. For simplicity, and in order to place the argument in the starkest relief, suppose that she makes a take-it-or-leave it offer to the listener concerning the value of z.[39] Let $w(E,z)$ denote her value. It satisfies the Bellman equation,

$$\rho \cdot w(E,z) = -c(E) + \mu_0 \cdot \mu(E) \cdot I \cdot [Z - z - w(E,z)] + \lambda_0 \lambda(e^*(\cdot))[-w(E,z)] \tag{9}$$

The first term on the right-hand side of the expression is her flow effort cost. The second represents the flow probability that she makes an *acceptable* proposal ($\mu_o \cdot \mu(E)$) times her capital gain in this event (($Z - z - w(E,z)$)). The third, and final, term captures the capital loss ($-w(E,z)$) she suffers if the listener beats her to the post, by independently formulating his own plan.

The listener's problem is to pick E and a value of z that is *acceptable* to the listener to maximize her own value, $w(E,z)$. Consider,

Proposition 5 *(Strategic Manipulation)*
(i) The speaker exerts positive effort $E^* > 0$. *Define, her optimal thinking rate by:* $m^* \equiv \mu_0 \cdot \mu(E^*)$.
(ii) Upon having an idea, she offers the listener a proposal that is worth $z^* \equiv V_0$ *to him.*
(iii) The listener's equilibrium valuation is $V(z^*,m^*) < V_0$.

Quite remarkably, according to part (*iii*) of the Proposition, the presence of the listener actually reduces the speaker's equilibrium payoff ($V^*(z^*,m^*)$) below that which he would attain in her absence ($V^*(0,m^*)$). This despite the assumption that he *can* reject any and all of her suggestions. Essentially, the speaker manipulates the listener by giving him suggestions that she knows he cannot credibly ignore. He too knows this, and responds by adjusting (lowering) his optimal thinking speed accordingly. The result of this distortion is that his welfare is lower than it would be in the speaker's absence.

[39]Indeed, since this is not a problem of imperfect information (i.e., the speaker honestly reveals all that she knows) imagine that she proposes a binding contract that promises to pay the speaker z.

The distortion itself arises directly from the listener's bounded cognitive abilities. To see this, consider the limit $\lim_{\lambda_0 \to \infty}$ in which the listener thinks of ideas "infinitely quickly,"

Proposition 6 *(Limiting Behavior)*
In the limit, $\lim_{\lambda_0 \to \infty}$, *we have:* $\lim_{\lambda_0 \to \infty} z^*(\cdot) = Z$; $\lim_{\lambda_0 \to \infty} E^*(\cdot) = 0$; *and* $\lim_{\lambda_0 \to \infty} V^*(m^*, z^*) = Z$.

Thus, in this setting intelligence – as measured by λ_0 – provides the listener with a bulwark against manipulation by the speaker. Indeed, as the listener becomes an infinitely fast thinker, the speaker's value converges to zero, and the listener himself captures all of the economic surplus (Z).

5 Application 2: Strategy Versus Tactics (Sketch)

- Military theory grounded upon the following triumvirate: (i) Strategy, (ii) Tactics, and (iii) Logistics
- Modern game theory provides an excellent forum for thinking about *strategic* aspects of conflict. The tactical (logistical) aspects are usually relatively trivial.
- The reason: agents *know* their action spaces.
- However, tactical concerns can become pertinent among boundedly rational agents who do not know their action sets.
- In this case, a lá Rommel, an agent has the opportunity of inflicting a "tactical surprise."

Consider the following framework.

- Time is continuous, beginning at date $t = 0$. There are two players: $i \in \{1,2\}$ who discount the future at the rate ρ.
- The players recognize that at future date T they will play a *game* (to be described).
- At the beginning ($t = 0$) of the game, each player has a privately known plan of action z_{10} and z_{20}.
- By exerting effort $e_i \in \{0,1\}$ each player can think of alternative actions at the rate: $\lambda_i \cdot e_i$.
- The cost of effort is $c \cdot e_i$
- Each idea has a "quality" $z \in [0,\infty)$
- Assume that $z_i \sim G(z|m_i)$, where m_i is agent i's *creativity*. For concreteness, assume that, $G(z|m_i) = 1\text{-} \exp[\text{-}z/m_i]$

Analysis

- Suppose that agent i has iN new ideas by date T.
- Define the best idea: $Q_i = \max\{z_{i1}, z_{i2}, ..., z_{iN}\}$.
- At the end of period T agents play a game, using their (privately known) best ideas $Q_i^* \equiv \max\{z_{i0}, Q_i\}$.

5.1 A Simple Game of Military Conflict

- At date $t = 0$ player 2 — the belligerent — can either initiate hostilities (**H**) against 1 or not (**W**) (withdraw).
- In the latter case, the players accrue payoffs v_i
- In the former, the players have a time period T in which they can do one of two things: stick with their initial plan (**N**) or formulate new battle one (**P**), by exerting costly effort $e_i = 1$

The Battle is joined

- The (discounted) payoff from victory is $100. (constant sum).
- The probability of victory by player 1 is: $I(Q_1^*, Q_2^*) \in \{0,1\}$, where: $I = 1$ if $Q_1^* \geq Q_2^*$ and zero otherwise.
- Let π_{jk}, where $j,k \in N,P$ denote the *probability of victory by player 2* given he picks j and player 1 k.
- Player 2's expected (gross) payoff is $100 \cdot \pi_{j,k}$.
- The *net* payoffs are calculated by subtracting the computational costs $c \geq 0$.
- The first goal is to determine the distribution of Q given e for each of the players: $F(Q|m,e)$
- It can be shown that:

$$F(Q|m,e) = 1\text{-}F(\text{-}\lambda_0\lambda(e)T \exp[\text{-}Q/m])$$

- which corresponds to a Gumbel (extreme value) distribution.
- For "large" T, the distribution converges to a limiting distribution that is "independent" of $G(Q|m)$ – the quality of idea distribution.

- The model permits a potentially rich analysis of helping is understand strategic issues that arise in settings in which one's opponent can carry out tactical innovations on the fly.
- Assume the players beliefs are congruent in the case that they select N,N and π_{NN} is the (common) prior belief of 2's victory if they adhere to their initial plans.
- Now suppose that player 1 sticks to her initial plan, and that 2 adapts his. What then?
- With probability $F(z_{20}|e_2,m_2)$ player 2 fails to think of an idea that is better than her initial plan z_{20}.
- With complementary probability 1-$F(z_{20}|e_2,m_2)$ she does think of a better idea.
- The conditional distribution is: $\hat{f}(Q\,|e_2,m_2)$
- Player 2 achieves victory if $Q \geq z_{10}$.
- Let $F_{0,1}(z_{10})$ denote 2's prior beliefs about z_{10}
- Given Q the *probability* of victory by 2 is (1-$F_{0,1}(z_{0,1})$).
- Thus, the probability of victory is: $\pi_{PN}(\cdot) \equiv$

$$F(z_{20}|\cdot)\cdot\pi_{NN} + (1 - F(z_{20}|\cdot))\cdot\int_{z_{20}}^{\infty}\hat{f}(Q\,|e_2,m_2)\cdot(1\text{-}F_{0,1}(Q))dQ$$

- Player 1 chooses e_2 to maximize his *ex ante* payoff: $(100)\pi_{PN}(\cdot) - c(e_2)$
- Denote the solution by: π^*_{PN}.
- Similar remarks apply to the other cases.

Payoff structure

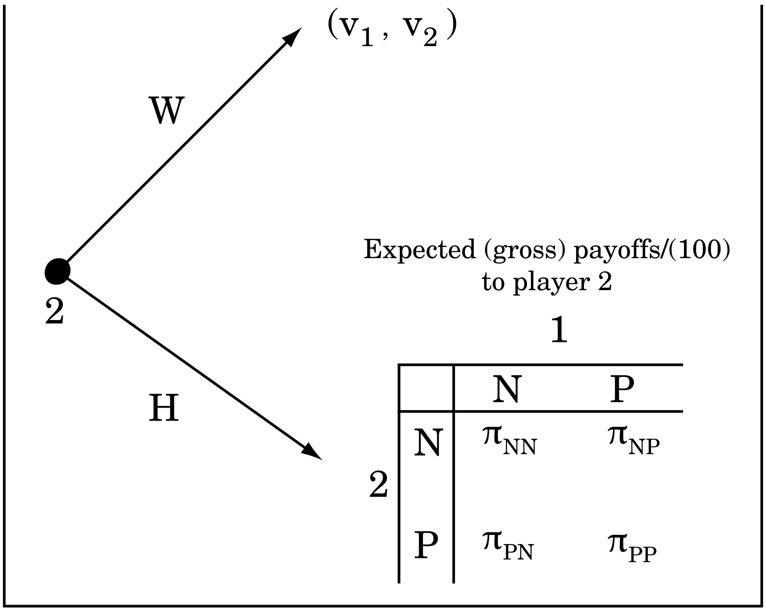

Figure 2: Payoff Structure

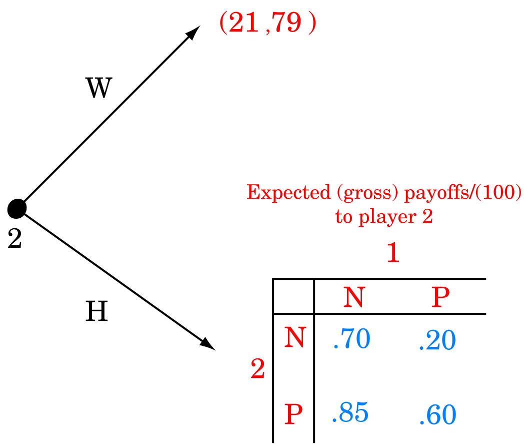

Figure 3: Ex Post Payoffs

COSTS

2 \ 1	N	P
N	(0,0)	(0,41)
P	(41,0)	(41,41)

Expected Payoffs

2 \ 1	N	P
N	(70,30)	(20,39)*
P	(44,15)	(19,-1)

Figure 4: Ex Ante Payoffs: Nash Equilibrium

- Despite the symmetric costs *and* the propitious odds for player 2, in the subgame perfect equilibrium she performs poorly.
- Anticipating this she opts *not* to engage in hostilities (H) and accrues 21 – which is not much better.

6 Application 3: Authority

In this Section we build upon Aghion and Tirole (1997) by examining how the allocation of authority in organizations depends upon innate intelligence and creativity.

- Consider a principal (P) who possesses a technology that calls for the employment of two agents (A1,A2).
- The agents are to formulate and implement a "plan of attack." for some project.
- If the value of a project is z, the principal accrues a payoff $(1-\alpha)\cdot z$. The remaining α accrues to the agents (in a manner described below).
- The principal is unable to distinguish among the recommendations of the agents, and hence rubber-stamps any proposal that comes before her.
- Assume that the two agents differ in their *innate* intelligence levels and in their *innate* their creativities.
- The basic economic problem is to determine how to allocate the decision making authority between the two agents.
- At the beginning of the game, the players pick the authority relationship that maximizes *ex ante* welfare.

There are three cases of interest:

1. **A1 authority** Only A1 has the right to send proposals to P. In particular, while he can entertain suggestions from A2 he has the right to veto them.
2. **A2 authority** The converse of the above.
3. **Balanced authority (B)** Either player can recommend a proposal to the principal.

The Analysis

- Time is continuous, and all players discount the future at the rate ρ Agent i thinks of projects at the rate $\lambda_i \cdot \lambda(e)$ and at a cost $c(e_i)$.
- The quality of each of A1's projects is Z. If the project is implemented, A1's payoff is $\alpha \cdot a_1 \cdot Z$. Player 2's (expected) payoff is $\alpha \cdot (1\text{-}a\,1) \cdot Z$
- Here, the parameter $a_1 \geq 0$ captures the *affinity* between A1 and A2. The quality of A2's proposal is z. The corresponding affinity parameter is a_2.
- If $a_i > 1$ there is a strict conflict of interest between the agents.

Information

- Hard: upon thinking of plan, the agent can credibly communicate it to the other agent.
- Soft: upon thinking of a plan, the agent can make an informal "suggestion" to the other player.
- For the purpose of this discussion, assume that information is *soft*, and that $a_1 = a_2 = 1$.

A-formal Authority

- Given that a2 = 1, A1 will veto any proposal advanced by 2. Lacking any benefit from effort, he will set $e_2 = 0$.
- As for player A1, his value, V_1 is:

$$rV_{A1} = -c(e_1) + \lambda_1 \lambda(e_1) \cdot [\alpha \cdot Z - V_{A1}]$$

- The optimal choice of effort is governed by the condition:

$$\frac{\alpha \cdot Z \lambda_1 \cdot \lambda'(\cdot) \cdot (r + c(\cdot))}{\lambda_1 \cdot \lambda + r} - c'$$

Balanced Authority

- In this case, player A1 worries about finishing last. Conditional upon e_2 his payoff V_{1B} is:

$$rV_{1B} = -c(e_1) + \lambda_1\lambda(e_1) \cdot [\alpha \cdot Z - V_{1B}] + \lambda_2\lambda(e_2) \cdot [0 - V_{1B}]$$

- The solution to this program is a best response function, $e^*_{1B}(\cdot, e_{2B})$.
- Similar but opposite remarks for A2.
- A Nash equilibrium is a pair: e^*_{1B} e^*_{2B} that satisfy these best-response functions.

Analysis

- Let V_P denote the principal's payoff.
- Ex ante welfare is: $W = V_P + V_1 + V_2$.
- As shown in Figure 5, depending upon the values of z, Z, λ_i either of the three authority schemes may be optimal.

A Simulation

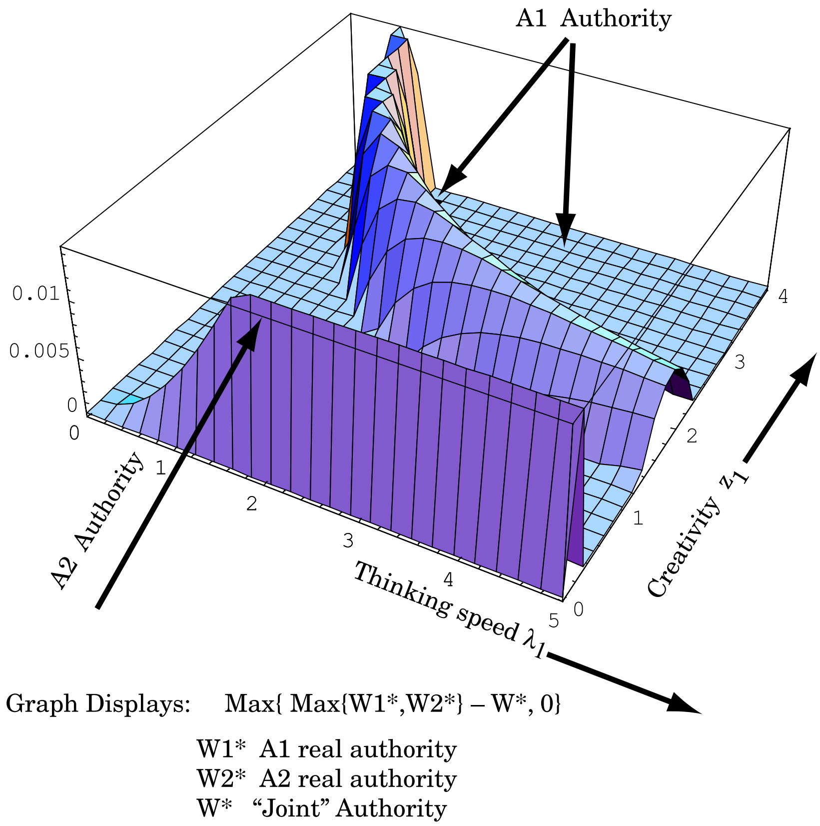

Figure 5: The Benefits and Costs of Delegation

7 Concluding Comments

We have introduced a simple model of boundedly rational agents. The heart of our framework is considering environments in which agents do not know their action spaces. However, via a process of "thinking" they can uncover feasible actions. The approach potentially has many applications. Three have been examined. A Third is to non-informative advertising.

A fourth would be to examine the micro-foundations of human capital accumulation, in which agents can make investments that help them think "faster," or deeper."

8 Appendix

In this Appendix we provide proofs of the lemmas and propositions presented in the paper.

Define the value: $v(e,s)$ by:

$$v(e,s) = \frac{-c(e) + \lambda(e)I(s)}{(\rho + \lambda(e)(1 - F(s)))}$$

where: $I(s) \equiv \int_s^{\bar{u}} udF(u)$. Here $v(e,s)$ is the value that accrues to the agent if he chooses effort e and *acts* according to the cutoff rule s. (Of course, $v(\cdot)$ is time consistent iff $v(e,s) = s$). Despite this, it is useful to work directly with $v(e,s)$, deferring for the moment issues of sequential rationality.

We first characterize the unique optimal value of e *given* s. To this end, pick any $s = s_0$ for which $I(s_0) > 0$. Notice that if $e = 0$, then: $v(0,s_0) = \lambda_0 I(s_0)/\rho + \lambda_0) \geq 0$, where: $\lambda_0 \equiv \lambda(0) \geq 0$. Given $I(s_0)$, define $\bar{e}(s_0) > 0$, by: $v(\bar{e}(s_0),s_0) \equiv 0$. Assumption 4 implies such a value always exists. Moreover, $v(e,s_0) < 0$, whenever $e > \bar{e}(s_0)$. Finally, note that for any $e \geq 0$:

$$v(e,s_0) = \frac{-c(e) + \lambda(e)(1 - F(s_0))E[u|u \geq_0]}{(\rho + \lambda(e)(1 - F(s_0)))} < E[u|u \geq s_0]$$

Differentiating $v(e,s_0)$ with respect to e yields:

$$v_e(e,s_0) = \frac{-c'(e) + \lambda'(e)(1 - F(s_0))}{E[u|u \geq s_0] - v(e,s_0)\}}\Delta_1$$

where: $\Delta_1 = \rho + \lambda(e)(1 - F(s_0)) > 0$. From the properties of c' and λ', we have: $\lim_{e\to 0} v(e,s_0) = \infty$. Hence, $v(e,s_0)$ possesses an interior global maximum on $(0,\bar{e}(s_0))$. Let $e(s_0) \in (0,\bar{e}(s_0))$ denote *any* stationary point of $v(e,s_0)$. By definition it is characterized by: $v_e(e(s_0),s_0) \equiv 0$. Differentiating $v_e(e,s_0)$ with respect to e around this point gives:

$$v_{ee}(e(s_0),s_0) = (-c'' + \lambda''(1 - F(s_0))\{E[u|u \geq s_0] - v(e,s_0)\})/\Delta_1 < 0.$$

Hence, any such stationary point is a maximum. It follows that $e(s_0)$ is *the* maximizing choice and that it is uniquely characterized by the condition: $v_e(e(s_0),s_0) \equiv 0$. Turning now to the properties of $v(e(s_0),s_0)$. If $s_0 = 0$, then $e(s_0) > 0$ and $v(e(0),0) > 0$. In contrast, if $s_0 = \bar{u}$, then $e(\bar{u}) = 0$ and $v(0,\bar{u}) = 0$. Since $v(e(s_0),s_0)$ is continuously differentiable in s_0, then, by the intermediate value theorem, there exists at least one fixed point: $s^* = v(e(s^*),s^*) \in (0,\bar{u})$. We now show that this fixed point is unique. With two or more fixed points, at least one of them must satisfy: $dv(e(s),s)/ds \neq 0$ (evaluated around the fixed point). Contrary to claim, let s^* represent *any* of two or more fixed points. Differentiate $v(e(s),s)$ around this point $s = s^*$ yields,

$$\lambda(e(s^*))f(s^*)\frac{(v(e(s^*),s^*) - s^*)}{\Delta_2} = 0$$

where: $\Delta_2 = (\rho + \lambda(e(s^*))(1 - F(s^*))) > 0$. The equality follows as $\lambda f(s^*) > 0$, and the definition of a fixed point $v(e(s^*),s^*) = s^*$.

8.0.1 Comparative Static Properties:

- Assume that: $\lambda(e) = \lambda_0 + \hat{\lambda}(e)$, where: $\hat{\lambda}(0) = 0$.

The individual's problem is characterized by the pair $s^* \in (0, \bar{u})$, $e^* > 0$, which uniquely solve:

$$\begin{aligned} s^* - \frac{-c^* + \lambda^* \int_{s^*}^{\bar{u}} udF(u)}{(\rho + \lambda^*(1-F^*))} &\equiv 0 \\ c'(e^*) - \lambda'(e^*)(1-F^*)\{E(u|u &\geq s^*) - s^*\} \equiv 0 \end{aligned}$$

Totally differentiating this system with respect to (s,e,λ_0) gives:

$$\begin{aligned} \{1\}ds + \{0\}de - \{(1-F^*)(E(u|u &\geq s^*))/\Delta_2\}d\lambda_0 = 0 \\ -\{v_{es}\}ds - \{v_{ee}\}de + \{0\}d\lambda_0 &= 0 \end{aligned}$$

$v_{es} = -\lambda'(1-F^*) < 0$ and $v_{ee} = -\{c'' - \lambda''(1-F^*)(E(u|u \geq s^*) - s^*)\} < 0$. Or,

$$\begin{bmatrix} 1 & 0 \\ -v_{es} & -v_{ee} \end{bmatrix} \begin{pmatrix} ds^*/d\lambda_0 \\ de^*/d\lambda_0 \end{pmatrix} = \begin{pmatrix} \{(1-F^*)(-s^* + E(u|u \geq s^*))/\Delta_2\} \\ 0 \end{pmatrix}$$

Or $Ax^* = d$. Using Cramer's rule:

$$\begin{aligned} ds^*/d\lambda_0 &= \frac{Det(A_s)}{Det(A)} > 0 \\ de^*/d\lambda_0 &= \frac{Det(A_e)}{Det(A)} < 0 \end{aligned}$$

where:

- $Det(A) = -v_{ee} > 0$,
- $Det(A_s) \equiv \begin{vmatrix} \{(1-F^*)(-s^* + E(u|u \geq s^*))/\Delta_2\} & 0 \\ 0 & -v_{ee} \end{vmatrix} = -v_{ee}\{(1-F^*)(-s^* + E(u|u \geq s^*))/\Delta_2\} > 0$
- $Det(A_e) = \begin{vmatrix} 1 & \{(1-F^*)(E(u|u \geq s^*))\} \\ -v_{es} & 0 \end{vmatrix} = v_{es}\{(1-F^*)(-s^* + E(u|u \geq s^*))/\Delta_2\} < 0.$

Consequently, fast thinkers are more picky (their s^* is higher), but lazier (their e^* is smaller). Define the individual's computational *speed* by:

$$\sigma(\lambda_0) \equiv \lambda(e^*(\lambda_0))1 - F(s^*(\lambda_0))$$

which is the flow rate of contemplating new ideas, λ, times the probability that one is acceptable: $1 - F^*$. Differentiating with respect to λ_0 yields,

$$1 > \sigma'(\lambda_0) \equiv (1-F^*) + \{\lambda' e^*_{\lambda_0}(1-F^*) - \lambda_0 f^* s^*_{\lambda_0}\}$$

Conditional upon the cutoff value, s^*, the effect of a unit increase in the flow rate of ideas is the raise the speed with which the agent arrives at a decision. However, the increase in λ_0 is associated with two indirect effects that retard the speed with which he arrives at a decision. First, the individual reduces his thinking effort, which lowers σ. Secondly, the individual becomes more choosy in his ratiocination process, raising the cutoff value s^*.

A special case: f is log-concave and $\lambda=\lambda_0$. More definite results can be obtained in the case in which $F(u)$ is log-concave and in which $\lambda(e)=\lambda_0$ (which is independent of e. In this case, $e^*=0$ and $c(e^*)=0$. Consider:

Proposition 7 *(Properties of $\sigma(\lambda)$). Assume that $F(u)$ is log-concave and that $\lambda(e)=\lambda_0$. In this case,*

$$\sigma'(\lambda_0)=(1-F(s^*(\lambda_0))-\lambda f s^*_{\lambda_0}$$

Proof. See Burdett and Ondrich (1985) and Flinn and Heckman (1983, p.39-41.). The essential argument is as follows. First the reservation utility is:

$$\rho u^*=\lambda_0(1-F(s^*))\{E[u{:}u\geq s^*]-s^*\}$$

This can be written:

$$\rho s^*=\lambda_0\int_{s^*}^{\bar{u}}(1-F(u))du$$

This follows as:

$$\begin{aligned}(1-F(s^*))\{E[u \quad : \quad u\geq s^*]-s^*\} &= \int_{s^*}^{\bar{u}}\{u-s^*\}dF(u)= \\ \bar{u}-s^*-\int_{s^*}^{\bar{u}}F(u)du \quad &= \quad \int_{s^*}^{\bar{u}}(1-F(u))du\end{aligned}$$

Differentiating s^* wrs ρ yields, s^*_λ, defined by:

$$\rho s^*_\lambda=\ -\lambda_0(1-F^*)s^*_\lambda+\int_{s^*}^{\bar{u}}(1-F(u))du$$

which upon re-arrangement yields:

$$s^*_\lambda=\frac{\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)}$$

As for the change in the speed of thinking, this is:

$$\sigma'=(1-F^*)-\lambda_0 f^* s^*_\lambda$$

Using the solution for s^*_λ gives:

$$\sigma'=(1-F^*)-\lambda_0 f^*\frac{\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)}=(1-F^*)-\frac{f^*}{1-F^*}\frac{\lambda_0(1-F^*)\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)}$$

Write this as,

$$\begin{aligned}
&(1-F^*) - \frac{f^*}{1-F^*}\frac{\lambda_0(1-F^*)\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)} = \\
&(1-F^*) - \frac{f^*}{1-F^*}\frac{\lambda_0(1-F^*)\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)} \\
&- \frac{f^*}{1-F^*}\frac{\rho\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)} + \frac{f^*}{1-F^*}\frac{\rho\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)} \\
= \quad &\left\{(1-F^*) - f^*\left(\frac{\int_{s^*}^{\bar{u}}(1-F(u))du}{1-F^*}\right)\right\} + \frac{f^*}{1-F^*}\frac{\rho\int_{s^*}^{\bar{u}}(1-F(u))du}{\rho+\lambda_0(1-F^*)}
\end{aligned}$$

In order that $\sigma' > 0$, it is sufficient that the term in braces is non-negative. If $f(.)$ is log-concave, this is indeed the case.

Proposition 8 *(An alternative proof that if f is log-concave that $\sigma' > 0$).*

Proof. The proof is based on the fact the if f is LC, then: $dE(u|u \geq s^*)/ds^* \equiv E_s < 1$. To show this, note that:

$$E = \frac{\int_{s^*}^{\bar{u}} uf(u)du}{\int_{s^*}^{\bar{u}} f(u)du}$$

Differentiate with respect to s^*,

$$E_s = \frac{f^*}{1-F^*}[E - s^*]$$

We must show that with f LC, that: $E_s < 1$. Suppose not, then: $E_s \geq 1$. In which case:

$$\frac{f^*}{1-F^*}[E - s^*] \geq 1$$

and,

$$\frac{[E - s^*](1-F^*)}{1-F^*} \geq \frac{(1-F^*)}{f^*}$$

Whence, using the definition of E and integrating by parts,

$$\frac{\int_{s^*}^{\bar{u}}(1-F(u))du}{1-F^*} \geq \frac{(1-F^*)}{f^*}$$

Define: $H(s^*) \equiv \int_{s^*}^{\bar{u}}(1-F(u))du$. Since, f is log-concave, then $H(.)$ is log-concave. Thus, $\ln[H^*]$ is concave. Consequently, $d[H^{*\prime}/H^*]/ds^* = H^{*\prime\prime}H^* - (H^{*\prime})^2 < 0$. But, $H^{*\prime\prime} = f(s^*)$ and $H^{*\prime} = -(1-F^*)$. Hence,

$$f^*H^* - (1-F^*)^2 < 0.$$

Yet,

$$\frac{H^*}{1-F^*} \geq \frac{(1-F^*)}{f^*}$$

which is a contradiction. Hence, f LC$\Longrightarrow E_s < 1$. Now let us show that: $\sigma' = d\sigma/d\lambda_0 > 0$. From the definition of s^* we have:

$$\rho s^* = \lambda_0(1-F^*)(E-s^*)$$

Differentiating with respect to λ:

$$\rho s^*_\lambda = -(\lambda_0 f^*(E-s^*) + (1-E_s)(1-F^*))s^*_\lambda + (1-F^*)(E-s^*)$$

Re-arranging, and solving for s^*_λ:

$$s^*_\lambda = \frac{(1-F^*)(E-s^*)}{\rho + (\lambda_0 f^*(E-s^*) + (1-E_s)(1-F^*))} > 0$$

which follows as: $E_s < 1$ and $E > s^*$. The rate at which the agent thinks of ideas *and* reaches a decision is:

$$\sigma = \lambda_0(1-F(s^*))$$

Differentiating this with respect to λ gives:

$$\sigma' = (1-F^*) - \lambda_0 f^* s^*_\lambda$$

Substituting for s^*_λ yields:

$$\begin{aligned} \sigma' &= (1-F^*) - \frac{\lambda_0 f^*(1-F^*)(E-s^*)}{\rho + (\lambda_0 f^*(E-s^*) + (1-E_s)(1-F^*))} = \\ &\quad (1-F^*)\left\{1 - \frac{\lambda f^*(E-s^*)}{\rho + (\lambda_0 f^*(E-s^*) + (1-E_s)(1-F^*))}\right\} \end{aligned}$$

From which:

$$\sigma' = (1-F^*)\left\{\frac{\rho + (1-E_s)(1-F^*))}{\rho + (\lambda_0 f^*(E-s^*) + (1-E_s)(1-F^*))}\right\} > 0$$

as $E_s < 1$.|| ■

Limiting Properties of λ_0. The cutoff level s^* increases with λ_0. It is also bounded above by $\bar{u}$. It therefore converges to a limit: $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{s}^*$. Heuristically, in this limit the decision maker can think "infinitely quickly" of alternatives. In this case, does his expected payoff converge to $\bar{u}$? Consider,

Proposition 9 *(The limiting properties of $s^*(\lambda_0)$). Given that: $\lambda(e) = \lambda_0 + \hat{\lambda}(e)$, then:* $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{s}^* = \bar{u}$.

The proof is straightforward. Suppose that contrary to claim, $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{s}^* < \bar{u}$. The value of effort e^* is decreasing and bounded below by zero. It converges to some limit: $\lim_{\lambda_0\to\infty} e^*(\lambda_0) = e^*_L \geq 0$. For every $\lambda_0 \geq 0$,

$$s^* - \frac{-c^* + \lambda^* \int_{s^*}^{\bar{u}} u dF(u)}{(\rho + \lambda^*(1-F^*))} \equiv 0$$

where: $s^* = s^*(\lambda_0)$ and $e^* = e^*(\lambda_0)$. In preparation for using l'hôpital's rule, write:

$$\bar{s}^* - \lim_{\lambda_0\to\infty} \frac{c^*(e_L^*)}{(\rho + (\lambda_0 + \lambda(e_L^*))(1 - F(\bar{s}^*))} - \lim_{\lambda_0\to\infty} \frac{(\lambda_0 + \lambda(e_L^*)) \int_{\bar{s}^*}^{\bar{u}} udF(u)}{(\rho + (\lambda_0 + \lambda(e_L^*))(1 - F(\bar{s}^*))} = 0$$

which yields,

$$\bar{s}^* - \frac{\int_{\bar{s}^*}^{\bar{u}} udF(u)}{(1 - F(\bar{s}^*)))} < 0$$

which is a contradiction. Hence, $\lim_{\lambda_0\to\infty} s^*(\lambda_0) = \bar{s}^* = \bar{u}$.||

8.0.2 *Special case II.* $\lambda(e) = \lambda_0\hat{\lambda}(e)$

Here, all of the results are determinate if f is log concave.

Proposition 10 *(Special case II): if $\lambda(e) = \lambda_0\hat{\lambda}(e)$ and if $f(\cdot)$ is log-concave, then: $\partial s^*/\partial\lambda_0 > 0$, $\partial e^*/\partial\lambda_0 > 0$, and $\partial\sigma^*/\partial\lambda_0 > 0$.*

Proof. Begin with the condition:

$$\rho s^* = -c^* + \lambda^*(1 - F^*)(E - s^*)$$

where it is understood that $E = E[u{:}u \geq s^*]$. Under LC, we have: $E_{s^*} < 1$. Totally differentiating the above yields:

$$\rho s_\lambda^* = -\lambda^* f^* \Delta E s_\lambda^* - \lambda^*(1 - F^*)(1 - E_s)s_\lambda^* \hat{\lambda}^*(1 - F^*)\Delta E$$

where: $\Delta E \equiv E - s^* > 0$. (Terms in e^* an be ignored by the envelope theorem). Re-arranging gives:

$$s_\lambda^* = \frac{\hat{\lambda}^*(1 - F^*)DeltaE}{\rho + \lambda^* f^* \Delta E + \lambda^* 1 - F^*)(1 - E_s)} > 0$$

which follows as: $1 - E_s > 0$. The corresponding FOC for e^* is:

$$-c^{*\prime} + \lambda_0\hat{\lambda}^{*\prime}(1 - F^*)\Delta \equiv 0$$

Totally differentiating this expression yields:

$$-\Delta_3 e_\lambda^* + \hat{\lambda}^{*\prime}(1 - F^*)\Delta E - \{\lambda_0\hat{\lambda}^{*\prime} f^* \Delta E + \lambda_0\hat{\lambda}^{*\prime}(1 - F^*)(1 - E_s)\}s_\lambda^* = 0$$

where: $\Delta_3 = c^{*\prime\prime} - \lambda^{*\prime\prime}(1 - F^*)\Delta E > 0$. Re-arranging in terms of e_λ^* gives:

$$\Delta_3 e_\lambda^* = \hat{\lambda}^{*\prime}\{(1 - F^*)\Delta E - \{\lambda_0 f^* \Delta E + \lambda_0(1 - F^*)(1 - E_s)\}s_\lambda^*\}$$

Substituting for s_λ^*:

$$\begin{aligned} \Delta_3 e_\lambda^* &= \hat{\lambda}^{*\prime}\{ \\ &\quad (1 - F^*)\Delta E - \{\lambda_0 f^* \Delta E + \lambda_0(1 - F^*)(1 - E_s)\}\frac{\hat{\lambda}^*(1 - F^*)\Delta E}{\rho + \lambda^* f^* \Delta E + \lambda^*(1 - F^*)(1 - E_s)}\Big\} \end{aligned}$$

Whence:

$$\Delta_3 e_\lambda^* = \frac{\hat{\lambda}^{*\prime}}{\Delta_4}\{(1-F^*)\Delta E\{\rho+\lambda^* f^*\Delta E + \lambda^*(1-F^*)(1-E_s)\} - ((1-F^*)\Delta E)\{\lambda^* f^*\Delta E + \lambda^*(1-F^*)(1-E_s)\}$$

where: $\Delta_4 = \rho + \lambda^* f^* \Delta E + \lambda^*(1-F^*)(1-E_s) > 0$. Thus,

$$\Delta_3 e_\lambda^* = \rho\frac{\hat{\lambda}^{*\prime}}{\Delta_4}(1-F^*)\Delta E > 0$$

The speed of thinking is given by:

$$\sigma \equiv \lambda^*(1-F^*)$$

It follows that:

$$\sigma' = -\lambda^* f^* s_\lambda^* + \frac{\lambda^*}{\lambda_0}(1-F^*) + (1-F^*)\lambda^{*\prime} e_{\lambda_0}^* > 0$$

The second and last term are positive. It is therefore sufficient to show that, taken together, the first two are *non-negative*. Consider,

$$-\lambda^* f^* s_\lambda^* + \frac{\lambda^*}{\lambda_0}(1-F^*) = \frac{\lambda^*}{\lambda_0}[(1-F)^* - \lambda_0 f^* s_\lambda^*]$$

Using the expression for s_λ^* gives:

$$\begin{aligned}
& \frac{\lambda^*}{\lambda_0}\left[1-f^*\frac{\lambda^*\Delta E}{\rho+\lambda^* f^*\Delta E + \lambda^*(1-F^*)(1-E_s)}\right] \\
= \; & \frac{\lambda^*}{\Delta_4\lambda_0}(1-F)^*[\rho+\lambda^* f^*\Delta E + \lambda^*(1-F^*)(1-E_s) - f^*\lambda^*\Delta E] \\
= \; & \frac{\lambda^*}{\Delta_4\lambda_0}(1-F)^*[\rho+\lambda^*(1-F^*)(1-E_s)] > 0
\end{aligned}$$

which follows as, $E_s < 1$. ■

Under the conditions of the proposition, an increase in "intelligence" — parameterized by λ_0 — implies that the agent exerts more effort at thinking and that he is more judicious in his acceptable actions. Even though they are more "picky," intelligent individuals also reach decisions more rapidly than less intelligent ones. Notice the importance of log-concavity. Suppose that: $\varepsilon_0 \geq \rho > 0$, and (for simplicity) $\hat{\lambda}' = 0$, then it is easy to see that there exists and $\varepsilon_0 > 0$, for which: $\lim_{\rho\to 0}\sigma_{\lambda_0} < 0$. Intuitively, if $E_s > 1$, then the returns from increasing s^* are quite high. For a sufficiently low discount rate, the decision maker takes advantage of this by thinking longer as λ_0 increases.

8.0.3 *The Effects of Patience* ($\lambda = \lambda(e)$).

- Infinitely (im)patient players.

Let us write $s^*(\rho)$ and $e^*(\rho)$. We are interested in the limiting properties: $\rho \to 0$ and $\rho \to \infty$. They uniquely solve:

$$\begin{aligned} s^* - \frac{-c^* + \lambda^* \int_{s^*}^{\bar{u}} u dF(u)}{(\rho + \lambda^*(1 - F^*))} &\equiv 0 \\ c'(e^*) - \lambda'(e^*)(1 - F^*)\{E(u|u &\geq s^*) - s^*\} \equiv 0 \end{aligned}$$

Totally differentiating this system with respect to (s,e,ρ) gives:

$$\begin{aligned} \{1\}ds + \{0\}de + \{s^*/\Delta_2\}d\lambda_0 &= 0 \\ -\{v_{es}\}ds - \{v_{ee}\}de + \{0\}d\lambda_0 &= 0 \end{aligned}$$

$v_{es} = -\lambda'(1 - F^*) < 0$ and $v_{ee} = -\{c'' - \lambda''(1 - F^*)(E(u|u \geq s^*) - s^*)\} < 0$. Or,

$$\begin{bmatrix} 1 & 0 \\ -v_{es} & -v_{ee} \end{bmatrix} \begin{pmatrix} ds^*/d\rho \\ de^*/d\rho \end{pmatrix} = \begin{pmatrix} -s^*/\Delta_2 \\ 0 \end{pmatrix}$$

Or $Ax^* = d$. Using Cramer's rule:

$$\begin{aligned} ds^*/d\rho &= \frac{Det(A_{s\rho})}{Det(A)} < 0 \\ de^*/d\rho &= \frac{Det(A_{e\rho})}{Det(A)} < 0 \end{aligned}$$

where:

- $Det(A) = -v_{ee} > 0$
- $Det(A_{s\rho}) = \begin{vmatrix} -s^*/\Delta_2 & 0 \\ 0 & -v_{ee} \end{vmatrix} = v_{ee}s^*/\Delta_2 < 0$
- $Det(A_{e\rho}) = \begin{vmatrix} 1 & -s^*/\Delta_2 \\ -v_{es} & 0 \end{vmatrix} = -v_{es}s^*/\Delta_2 > 0.$

Thus, more *impatient* players are *less* picky. However, they work harder. Both of these are to ensure a more rapid response of their decision problem. As for computational speed,

$$\sigma'(\rho) = \lambda' e_\rho^* - \lambda f s_\rho^* > 0$$

Proof of Proposition 3

It is conceivable that by exerting effort, the decision maker may think about better quality ideas, rather than just thinking faster. Assume that: $\lambda(e) = \lambda_0$ (independent of e) and that $c(e) = 0$. This enables us to more clearly differentiate between the effects of thinking faster, versus the effects of thinking "better."

Conditional upon s^* the value of e is:

$$\rho v(e,s^*) = -c(e) + \lambda_0(1 - F^*)(E - s^*) = -c(e) + \lambda_0 \int_{s^*}^{\bar{u}} (1 - F)du$$

The Properties of $F(u|e)$. From an economic standpoint, the first basic condition we require is that: $b(e) \equiv (1-F^*)(E-s^*)$ is increasing and concave in e. Differentiating with respect to e gives:

$$\begin{aligned} b_e &= -F_e^*\Delta E + (1-F^*)E_e > 0 \\ b_{ee} &= -F_{ee}^*\Delta E - 2F_e^* E_e + (1-F^*)E_e < 0 \end{aligned}$$

where: $\Delta E \equiv E\{u|u \geq s^*, e\}$. If $F_e \leq 0$ (strict for some u) then: $-F_e^*\Delta E > 0$ is the benefit of the *increased* likelihood that a given idea is acceptable. A sufficient condition for the inequality is that $E_e > 0$, which says that increased effort raises the average quality of acceptable ideas. Using the fact that

$$b(e) \equiv (1-F^*)(E-s^*) = \int_{s^*}^{\bar{u}} (1-F)du$$

implies that:

$$b_e = -\int_{s^*}^{\bar{u}} F_e du$$

Consider:

Lemma 3 *(The benefit function $b(e)$). If $F_e \leq 0$ for all $u \in [0,\bar{u}]$ and $F_e < 0$ for some u then: $b_e > 0$. (This is First-order stochastic dominance). If $F(\cdot)$ is convex in e then: $b_{ee} < 0$. If in addition: $LR \equiv f/(1-F)$ is decreasing in e then $E_e > 0$.*

From an economic standpoint, the first basic condition we require is that: $b(e) \equiv (1-F^*)(E-s^*)$ is increasing and concave in e. Differentiating with respect to e gives:

$$\begin{aligned} b_e &= -F_e^*\Delta E + (1-F^*)E_e > 0 \\ b_{ee} &= -F_{ee}^*\Delta E - 2F_e^* E_e + (1-F^*)E_e < 0 \end{aligned}$$

where: $\Delta E \equiv E\{u|u \geq s^*, e\}$. If $F_e \leq 0$ (strict for some u) then: $-F_e^*\Delta E > 0$ is the benefit of the *increased* likelihood that a given idea is acceptable. A sufficient condition for the inequality is that $E_e > 0$, which says that increased effort raises the average quality of acceptable ideas. Using the fact that

$$b(e) \equiv (1-F^*)(E-s^*) = \int_{s^*}^{\bar{u}} (1-F)du$$

implies that:

$$b_e = -\int_{s^*}^{\bar{u}} F_e du$$

Consider:

Lemma 4 *(The benefit function $b(e)$). If $F_e \leq 0$ for all $u \in [0,\bar{u}]$ and $F_e < 0$ for some u then: $b_e > 0$. (This is First-order stochastic dominance). If $F(\cdot)$ is convex in e then: $b_{ee} < 0$. If in addition: $LR \equiv f/(1-F)$ is decreasing in e then $E_e > 0$.*

The further properties of $F(.)$. Suppose that $F_e < 0$ as required. What are the implied properties of E (if any)? What further restrictions when placed on F lead to "intuitive" comparative static results?

To answer this question, note that:

$$b_e = -F_e^* \Delta E + (1-F^*)E_e = -\int_{s^*}^{\bar{u}} F_e du > 0$$

Alternatively,

$$(1-F^*)E_e = \frac{F_e^*}{1-F^*}\int_{s^*}^{\bar{u}}(1-F)du - \int_{s^*}^{\bar{u}} F_e du$$

Whence,

$$(1-F^*)E_e = \int_{s^*}^{\bar{u}}\{\frac{F_e^*}{1-F^*}(1-F) - F_e\}du$$

And,

$$(1-F^*)E_e = \int_{s^*}^{\bar{u}}(1-F)\{\frac{F_e^*}{1-F^*} - \frac{F_e}{(1-F)}\}du$$

Define:

$$\phi(u) \equiv \frac{F_e^*}{1-F^*} - \frac{F_e}{(1-F)}$$

Then, $\phi(u) = 0$. If $\phi'(u) > 0$, then:

$$(1-F^*)E_e = \int_{s^*}^{\bar{u}}(1-F)\{\phi(u)\}du > 0$$

Differentiating $\phi(u)$ wrs to u yields,

$$\phi'(u) = -\left(\frac{F_{eu}(1-F) + F_e F_u}{(1-F)^2}\right) = -\left(\frac{f_e(1-F) + F_e f}{(1-F)^2}\right) > 0$$

Given that $F_e < 0$, a sufficient condition is that $f_e < 0$. A weaker condition is MLRP. Thus, differentiate $LR \equiv f/(1-F)$ with respect to e to yield:

$$LR_e(1-F)^2 = f_e(1-F) + F_e f$$

The MLRP says this is negative. In which case, $\phi'(u) > 0$ as required. Consequently, if F satisfies: $F_e < 0$, then: $b_e > 0$. If, in addition, F satisfies MLRP, then: $E_e > 0$.

8.0.4 Characterization of behavior.

The reservation value is:

$$\rho s^* = -c(e^*) + \lambda_0 \int_{s^*}^{\bar{u}}(1 - F(u|e^*))du$$

and the first-order condition for e^* is:

$$0 = -c'(e^*) - \lambda_0 \int_{s^*}^{\bar{u}} F(u|e^*)_e du$$

Totally differentiating with respect to λ_0:

$$\rho s^*_{\lambda_0} = -\lambda_0 \,(1 - F(s^*|e^*))s^*_\lambda + \int_{s^*}^{\bar{u}} (1 - F(u|e^*))du = -\lambda_0 \,(1 - F(s^*|e^*))s^*_\lambda + \frac{\rho s^* + c^*}{\lambda_0}$$

This yields:

$$s^*_{\lambda_0} = \frac{\int_{s^*}^{\bar{u}} (1 - F(u|e^*))du}{\rho + \lambda_0(1 - F^*)} = \frac{(\rho s^* + c^*)/\lambda_0}{\rho + \lambda_0(1 - F^*)} > 0$$

As for the choice of e^* this (at the interior max.) is:

$$\{c'' + \lambda_0 \int_{s^*}^{\bar{u}} F^*_{ee} du\}e^*_\lambda = -\lambda_0 F^*_e s^*_\lambda + \int_{s^*}^{\bar{u}} F(u|e^*)du > 0$$

The agent becomes more picky, encouraging effort by the generation of better ideas. In addition, he thinks faster, encouraging effort at the margin. As for the speed of his thinking, this is:

$$\sigma(\lambda_0) \equiv \lambda_0(1 - F(s^*|e^*))$$

Differentiating wrs to λ_0:

$$1 - F^* + \lambda_0\{-f^* s^*_\lambda - F^*_e e^*_\lambda\}$$

with $F^*_e < 0$ and $e^*_\lambda > 0$, then the agent unambiguously reaches a faster decision if f is log-concave (using previous arguments). Furthermore, if $F_{eu} < 0$, then the average quality of the idea increases as well (i.e., $E(u|u \geq s^*, e^*)_{\lambda_0} > 0$).

References

AGHION, P., AND J. TIROLE (1997): "Formal and Real Authority in Organizations," *Journal of Political Economy*, 105(1), 1–29.

ARAGONES, E., I. GILBOA, A. POSTLEWAITE, AND D. SCHMEIDLER (2005): "Fact Free Learning," *American Economic Review*, 95(6), 1355–1368.

BAGWELL, K. (2005): "The Economic Analysis of Advertising," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3. Elsevier, North-Holland, Amsterdam.

BAUMOL, W. J., AND R. E. QUANDT (1964): "Rules of Thumb and Optimally Imperfect Decisions," *American Economic Review*, 54(1), 23–46.

BURDETT, K., AND J. ONDRICH (1985): "How Changes in Labor Demand Affect Unemployed Workers," *Journal of Labor Economics*, 3(1), 1–10.

BURDETT, K., S. SHI, AND R. WRIGHT (2001): "Pricing and Matching with Frictions," *Journal of Political Economy*.

CLARK, JOHN, M. (1918): "Economics and Modern Psychology," *Journal of Political Economy*, 26(1), 1–30.

CONLISK, J. (1988): "Optimization Cost," *Journal of Economic Behavior and Organization*, 9(3), 213–28.

——— (1996a): "Bounded Rationality and Market Fluctuations," *Journal of Economic Behavior and Organization*, 29(2), 233–50.

——— (1996b): "Bounded Rationality and Market Fluctuations," *Journal of Economic Behavior and Organization*, 29(2), 233–50.

——— (1996c): "Why Bounded Rationality?," *Journal of Economic Literature*, 34(2), 669–700.

——— (2001): "Costly Predation and the Distribution of Competence," *American Economic Review*, 91(3), 475–84.

——— (2003): "A Note on Convergence of Adaptive Satisficing to Optimal Stopping," *Journal of Political Economy*, 111(6), 1353–60, 0022-3808.

CROSS, JOHN, G. (1973): "A Stochastic Learning Model of Economic Behavior," *Quarterly Journal of Economics*, 87(2), 239–266.

DAY, R. H. (1967): "Profits, Learning and the Convergence of Satisficing to Marginalism," *Quarterly Journal of Economics*, 81(2), 302–311.

DAY, R. H., AND E. H. TINNEY (1968): "How to Co-operate in Business without Really Trying: A Learning Model of Decentralized Decision Making," *Journal of Political Economy*, 76(4), 583–600.

DE PALMA, A., G. M. MYERS, AND Y. Y. PAPAGEORGIOU (1994): "Rational Choice Under an Imperfect Ability To Choose," *American Economic Review*, 84(3), 419–440.

DOW, J. (1991): "Search Decisions with Limited Memory," *Review of Economic Studies*, 58(1), 1–14.

FLINN, C. J., AND J. J. HECKMAN (1983): “Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?,” *Journal of Labor Economics*, 1(1), 28–42.

GLAZER, J., AND A. RUBINSTEIN (2006): “A study in the pragmatics of persuasion: a game theoretical approach,” *Theoretical Economics*, 1, 395Ű410.

HART, O., AND J. MOORE (1999): “Foundations of Incomplete Contracts,” *Review of Economic Studies*, 66(1), 115–38.

HOLMSTROM, B., AND P. MILGROM (1991): “Mult-task Principal-agent Analysis: Incentive Contracts, Asset Ownership, and Job Design.”,” *Journal of Law, Economics, and Organization*, 7(1), 24–52.

LETTAU, M., AND H. UHLIG (1999): “Rules of Thumb versus Dynamic Programming,” *American Economic Review*, 89(1), 148–174.

LIPMAN, B. L. (1991): “How to Decide How to Decide How to. . . : Modeling Limited Rationality,” *Econometrica*, 59(4), 1105–25.

——— (1995): “Information Processing and Bounded Rationality: A Survey,” *Canadian Journal of Economics*, 28(1), 42–67.

——— (1999): “Decision Theory without Logical Omniscience: Toward an Axiomatic Framework for Bounded Rationality,” *Review of Economic Studies*, 66(2), 339–61.

LIPPMAN, S. A., AND J. J. MCCALL (1976): “The Economics of Job Search: a Survey,” *Economic Inquiry*, 14(4), 347–68.

MACLEOD, W. B. (1996): “Decision, Contract, and Emotion: Some Economics for a Complex and Confusing World,” *Canadian Journal of Economics*, 29(4), 788–810.

——— (2002a): “Complexity and Contract,” in *The Economics of Contract in Retrospect and Prospect*, ed. by E. Brosseau, and J.-M. Glachant. C.U.P., Cambridge, UK.

——— (2002b): “Complexity, Bounded Rationality and Heuristic Search,” *Contributions to Economic Analysis and Policy*, 1(1), na.

MORTENSEN, DALE, T. (1986): “Job Search and Labor Market Analysis,” in *Handbook in Labor Economics*, ed. by O. Ashenfelter, and R. Layard. Elsevier, Amsterdam, North Holland.

MULLAINATHAN, S. (2002): “A Memory-Based Model of Bounded Rationality,” *Quarterly Journal of Economics*, 117(3), 735–74.

OSBORNE, M. J., AND A. RUBINSTEIN (1998): “Games with Procedurally Rational Players,” *American Economic Review*, 88(4), 834–47.

RADNER, R. (1975a): “A Behavioral Model of Cost Reduction,” *Bell Journal of Economics*, 6(1), 196–215.

——— (1975b): “Satisficing,” *Journal of Mathematical Economics*, 2(2), 253–262.

——— (1996): “Bounded Rationality, Indeterminacy, and the Theory of the Firm,” *Economic Journal*, 106(438), 1360–1373.

RADNER, R., AND M. ROTHSCHILD (1975): “On the Allocation of Effort,” *Journal of Economic Theory*, 10(3), 358–76.

ROGERSON, R., R. SHIMER, AND R. WRIGHT (2005): "Search-Theoretic Models of the Labor Market: A Survey," *Journal of Economic Perspectives*, forthcoming.

ROSENTHAL, R. W. (1993): "Rules of Thumb in Games," *Journal of Economic Behavior and Organization*, 22(1), 1–13.

RUBINSTEIN, A. (1993): "On Price Recognition and Computational Complexity in a Monopolistic Model," *Journal of Political Economy*, 101(3), 473–84.

——— (1998): *Modeling bounded rationality*. MIT Press, Zeuthen Lecture Book Series. Cambridge and London.

SIMON, HERBERT, A. (1955): "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, 69(1), 99–118.

——— (1979): "Rational Decision Making in Business Organizations," *American Economic Review*, 69(2), 493–513.

——— (1987): "Bounded Rationality," in *The New Palgrave*, ed. by J. Eatwell, M. Milgate, and P. Newman. W.W. Norton, New York, NY.

STIGLER, G. (1961): "The Economics of Information," *Journal of Political Economy*, 69(3), 213–225.

WILSON, A. (2003): "Bounded Memory and Biases in Information Processing," .

WINTER, SIDBEY, G. (1971): "Satisficing, Selection, and The Innovating Remnant," *Quarterly Journal of Economics*, 85(2), 237–261.