A New Challenge for Economics: "The Frame Problem"

Xavier Gabaix
MIT

David Laibson
Harvard University and NBER*

May 3, 2001

Abstract

Most real-world decision problems are complex in the sense that the optimum is time-consuming to calculate. However, the optimal solution can often be approximated using a few sensible rules of thumb or good judgment. Sometimes these simplifying heuristics generate a poor decision, but this is the occassional price of approximation. The frame problem of cognitive science describes the challenge of designing intelligent systems that can effectively and quickly identify good responses to complex problems. This essay summarizes a modeling approach that addresses some of the issues raised by the frame problem. This model uses option value calculations to selectively allocate attention and cognition, thereby simplifying analysis of complex problems.

JEL classification: C70, C91, D80

*Gabaix: Department of Economics, MIT, Cambridge, MA 02142, xgabaix@mit.edu. Laibson: Department of Economics, Harvard University, and NBER, Cambridge MA 02138, dlaibson@harvard.edu. We are grateful for the suggestions of Isabelle Brocas and participants at the Economics and Psychology Conference at ECARE. Numerous research assistants contributed to this paper. We are particularly indebted to John Friedman, Guillermo Moloche, Rebecca Thornton, and Stephen Weinberg. Laibson acknowledges financial support from the National Institutes of Health (R01-AG-16605), the MacArthur Foundation, and the Olin Foundation. Gabaix and Laibson acknowledge support from the National Science Foundation.

Keywords: directed cognition, optimization, bounded rationality, heuristics, biases, experimental economics, simplification

1 Introduction: The Frame Problem

Consider the problem of a travelling salesman who must map out a route which includes visits to K cities. If an economist were asked to predict the path chosen by the traveller, the economist might begin by conceptualizing an optimization problem. 'Order the cities so that the resulting route minimizes total trip time.' However, the economist would soon realize that the number of possible paths grows so quickly in K that this optimization problem is effectively insolvable. For example, with $K = 30$, the problem yields $3{*}10^{32} = K!$ distinct paths, enough to overwhelm any supercomputer.[1]

Real decision makers confront versions of the travelling salesman's problem every day and most of the time the problem does not overwhelm *them.* They use numerous simplifications to "solve" this problem. For example, they may rule out certain classes of paths right from the start. They may use heuristics to identify promising candidate paths. They may stop their analysis when their current best path is unlikely to be substantially improved by further analysis.

Three features of the traveller's problem interest us. First, the problem is complex in the sense that the theoretical optimum is time-consuming to calculate. Second, the optimal solution can be approximated using a few sensible rules of thumb. Third, sometimes these simplifying heuristics actually generate a suboptimal choice. Such poor choices represent the price of simplification.

In the past twenty years, economists and psychologists have made substantial progress toward understanding the heuristics that humans use to analyze decision problems. The work of Daniel Kahneman and Amos Tversky stands out within this literature.[2] They showed that a few basic heuristics (e.g., representativeness, availability, and anchoring) lie behind many of our intuitive inferences. We view such heuristics as powerful simplifying tools that enable computationally limited organisms to make sensible judgments. However, such simplifications naturally distort our inferences, invari-

[1] The most powerful current supercomputers execute one trillion calculations per second (one teraflops). If each path is evaluated with a single calculation, then analysis of the $K = 30$ case would require 10^{13} years.

[2] For a few examples of other seminal contributions to our understanding of bounded rationality see Conlisk 1996, Payne, Bettman and Johnson 1993, Shugan 1980, Simon 1959, and Thaler 1991, 1994. All of these authors discuss both the costs and benefits of simplification.

ably causing deviations from perfect rationality and sometimes generating decisions that dramatically lower welfare relative to the rational benchmark.

Our current research studies the process by which humans simplify and approximately solve complex decision problems. We want to understand both the set of heuristics that decision makers use and the ways that those heuristics are applied to complex problems. For example, how does a decision-maker represent a complex problem? What information does the decision-maker use and what information is overlooked? How is the analyzed information cognitively manipulated? When does the decision-maker decide to stop working on a complex problem and act on her best guess? More generally how does the decision-maker decide how deeply to analyze a problem?

These questions relate closely to a set of questions in cognitive science known as the 'frame problem.' The frame problem arises in any model of cognition, but it is sometimes motivated as a challenge for the designers of a robot. How should the robot be programmed to decide which information to analyze and which information to overlook? Failing to analyze relevant information can lead to bad choices. Analyzing all of the potentially relevant information takes too long and effectively paralyzes the system. In a famous essay, Daniel Dennett motivates the frame problem with the following story:

> Once upon a time there was a robot, named R1 by its creators. Its only task was to fend for itself. One day its designers arranged for it to learn that its spare battery, its precious energy supply, was locked in a room with a time bomb set to go off soon. R1 located the room, and the key to the door, and formulated a plan to rescue its battery. There was a wagon in the room, and the battery was on the wagon, and R1 hypothesized that a certain action which it called PULLOUT (WAGON,ROOM) would result in the battery being removed from the room. Straightaway it acted, and did succeed in getting the battery out of the room before the bomb went off. Unfortunately, however, the bomb was also on the wagon. R1 knew that the bomb was on the wagon in the room, but didn't realize that pulling the wagon would bring the bomb out along with the battery. Poor R1 had missed that obvious implication of its planned act.
>
> Back to the drawing board. 'The solution is obvious,' said the designers. 'Our next robot must be made to recognize

> not just the intended implications of its acts, but also the implications about their side-effects, by deducing these implications from the descriptions it uses in formulating its plans.' They called their next model, the robot-deducer, R1D1. They placed R1D1 in much the same predicament that R1 had succumbed to, and as it too hit upon the idea of PULLOUT (WAGON,ROOM) it began, as designed, to consider the implications of such a course of action. It had just finished deducing that pulling the wagon out of the room would not change the color of the room's walls, and was embarking on a proof of the further implication that pulling the wagon out would cause its wheels to turn more revolutions than there were wheels on the wagon - when the bomb exploded. (Dennett, 1987, pp. 41-2)

Finding an effective middle ground between too little and too much cognitive analysis has proved to be a difficult challenge for cognitive scientists. How can artificially intelligent machines be designed to quickly and *effectively* identify the right depth of analysis? Humans do it all the time. What kind of model can describe or replicate this remarkable ability? The challenge of developing such a model has been described as "the most important problem in cognitive science (Dietrich and Fields, 1996, p. 13)."

To further illustrate the frame problem, consider an even more complex version of the travelling salesman problem. Real travellers consider many more attributes than the distance between two points. If a traveller had infinite cognitive capacity he would consider the financial costs of road tolls, the variability of travel times, the probability of an accident, the depreciation to his car caused by road disrepair, the presence of amenities during the trip (e.g., road-side diners or motels), the option value of alternative routes in case the chosen route turns out to be congested, the relationship between travel time and weather conditions (e.g., some roads are bad in snow), the proximity to repair services if the car breaks down, and so on.

Somehow, decision makers wrestle with such highly complex problems and come up with useful solutions. Moreover, people are not mechanistic in their analysis of such problems. They know when to work hard to come up with a sophisticated decision (e.g., when they are picking a route to drive from New York to San Francisco) and when to go with their first instinct (e.g., when travelling a few blocks in a familiar neighborhood).

Even when decision makers analyze a decision in depth, they still have to decide which information to consider and which information to ignore. For example, the Rocky Mountains receive 300 inches of snow per year. A winter coast-to-coast trip through the Rockies may generate snow delays. That's worth considering. By contrast, no normal traveller would incorporate the incremental effects of global warming in their estimate of the likelihood of snow delays. A solution to the frame problem must avoid analysis of every conceivable issue that might bear, however remotely, on our payoffs.

A solution to the frame problem must also avoid the paralysis that would arise if we had to separately identify and decide to ignore all of the inconsequential issues that do *not* bear on our decision problems. For every relevant fact, like snowfall in Colorado, there are untold irrelevant facts, like Colorado's date of statehood. Identifying everything that we know about Colorado and determining the relevance of all of those facts is a cognitive waste.

For the purposes of this essay, we conceptualize the frame problem as the simultaneous challenge of thinking effectively *and* quickly about a complex problem. This definition matches the approach that Dennett has taken in a series of influential essays over the past two decades: "A creature that can solve any problem given enough time — say a million years — is not in fact intelligent at all. We live in a time-pressured world and must be able to think quickly before we leap (1987, p. 49)." "I see the frame problem as arising most naturally and inevitably as a problem of finding a *useful*, compact representation of the world — providing actual *anticipations in real time* for purposes of planning and control (1996, p. 1, original emphasis)."[3]

Five motivations heighten our interest in the frame problem. First, the frame problem represents a basic economic question. In essence, the frame problem describes the way that individuals allocate particularly important scarce resources: attention and cognition.

Second, without a solution to the frame problem, economic models are

[3] Other authors use the label 'frame problem,' to describe a narrower set of issues, particularly representations of information in dynamic environments in which an agent changes the environment through his actions. See Pylyshyn (1987), and Ford and Pylyshyn (1996) for discussions of various conceptualizations of the frame problem. For example, Glymour (1996), writes, "The central questions about planning and the frame problem are: (a) How can causal structure be learned reliably, efficiently and feasibly? and (b) How can complete or partial knowledge of causal structure be used to reliably and feasibly predict the effects of interventions in the causal system (p. 32)." See McCarthy and Hayes (1969) for the first description of the frame problem. See also Shanahan (1997) for a thorough treatment in a particular context.

necessarily *incomplete*. For example, the standard economic maximization model makes no practical quantitative prediction in the traveller's problem, since the maximization problem can not be solved on any modern computer.

Third, a century of cognitive psychology research suggests that decision-makers use short-cuts to solve complex problems. We want to understand that simplification process.

Fourth, solving the frame problem is an important input to the study and development of artificial intelligence. To build intelligent machines we need to understand how to efficiently simplify complex problems.

Fifth, and perhaps most importantly, the frame problem provides a microfoundation for many of the systematic errors that humans make when solving hard problems. Sensible cognitive shortcuts necessarily generate imperfect judgments. Ultimately, we believe that understanding endogenous mental shortcuts will lead to a general framework for understanding many of the ways that humans deviate from the rational benchmark.

The remainder of this essay summarizes a modeling approach that addresses some of the issues raised by the frame problem. We discuss an economic model of attention allocation which is based on standard *option value* calculations (Gabaix and Laibson 2000b).[4] The option value approach quantifies the economic value of the option to collect additional information.[5] In our model, actors think more deeply about problems when additional analysis is likely to reveal a large amount of new information or when no obvious winner has emerged from a class of competing alternatives. We describe the intuition behind this model, sketch the formal structure of the model, and review an experiment which we have conducted to test the model.

2 Directed cognition

Our 'directed cognition' model applies two simple economic ideas that derive from the option value literature. These ideas provide a framework for deciding how deeply to analyze a complex decision problem.

First, when cognitive analysis yields little new insight, the option value of continued analysis declines. Second, when many different choices are being

[4] An earlier paper (Gabaix and Laibson 2000a) analyzes the same phenomena, but applies a mechanistic model, which does not provide a theory of endogenous simplification.

[5] See Dixit and Pindyck (1994) for an introduction to the option value approach to investment under uncertainty.

compared and a particular choice gains a large edge over the available alternatives, the option value of continued analysis also declines. Our directed cognition model quantifies these two effects and integrates them in a formal framework that makes sharp quantitative predictions about boundedly rational choices. We have successfully tested these predictions in an experiment conducted on Harvard undergraduates.

To illustrate the two basic ideas in our model, consider the complex decision tree in Figure 1.[6] This is one of twelve randomly generated trees that we asked our subjects to analyze. Each starting box in the left-hand column leads probabilistically to boxes in the second column. Branches connect boxes and each branch is associated with a given probability. For example, the first box in the last row contains a 5. From this box there are two branches with respective probabilities .65 and .35. The numbers inside the boxes represent flow payoffs. Starting from the last row, there exist 7,776 outcome paths. For example, the outcome path that follows the highest probability branch at each node is $\langle 5, 4, 4, 3, 1, -2, 5, 5, -3, 4 \rangle$. Integrating with appropriate probability weights over all 7,776 paths, the expected payoff of starting from row five is 4.12 (where the subject receives the payoff in every box through which he travels).

We asked undergraduate subjects to choose one of the boxes in the first column of Figure 1. We told the subjects that they would be paid the expected value associated with whatever starting row they chose.

The directed cognition model can be used to analyze trees like those in Figure 1. We informally describe this application of the model in this essay. We refer interested readers to Gabaix and Laibson (2000b) for a detailed description of the model.

Our application is built on a basic cognitive operation: extending a partially examined path one column deeper into the tree. Such extensions enable the decision-maker to improve her forecast of the expected value of any given starting row.[7] For example, consider again the last row of Figure 1. Imagine that a decision-maker is trying to approximate the expected value of that starting row, and that he has calculated only the expected value

[6] Decision trees are difficult to solve, but they are not in the same class of problems (NP-problems) as the travelling salesman problem.

[7] Such forward induction is motivated by experimental work by Colin Camerer et al (1994) and theoretical work by Philippe Jéhiel (1995), which argue that decision-makers may solve problems by looking forward. Our framework can generalize to also include cognitive operations based on backwards induction.

of the two truncated paths leading from column 1 to column 2. Hence, the estimated value would be: $a = 5 + (.65)(4) + (.35)(-5)$. Following the highest probability path (i.e., the upper path with probability .65), the decision maker could look ahead one additional column and come up with a revised estimated value:

$$a' = a + .65[.15(-3) + .5(4) + .2(-5) + .15(-4)].$$

Each path extension refines the decision-maker's expectations about the value of a starting row. We assume that each path extension requires some time and generates cost q.

We will also allow concatenated path extensions. Specifically, if the decision-maker executes a two-step path extension from a particular node, he twice follows the branch with the highest probability. Hence, starting from the last row of Figure 1, his updated estimate would be,

$$a'' = a' + (.65)(.5)[.3(0) + .05(3) + .65(3)].$$

Such concatenated path extensions generalize naturally. A τ-step path extension looks τ columns more deeply into the tree. At each intermediate node the extension follows the branch with highest probability.

A given path from a particular starting node can be extended in many different ways, since a path will generally have many subpaths. Let f represent one such path extension; f embeds information about the starting row, the subpath which is to be extended, and the number of steps in that extension. Let σ represent the standard deviation of the updated estimate resulting from application of f:

$$\sigma^2 = E(a' - a)^2,$$

where a' represents the updated value of a after application of path extension f. We assume that the decision maker thinks about the update as if

$$a' = a + x$$

where x has a normal distribution with mean zero and standard deviation σ.

We now derive the ex-ante expected value of path extension f. Consider the simple case in which the decision-maker faces a choice between two rows: A and B. Assume that the agent knows that choosing row B will generate a certain payoff of b.

The payoff from choosing row A is uncertain. Call this expected payoff a. The agent can learn more about this expected payoff if she executes a path extension f. Specifically, executing the path extension will enable her to update the expected payoff of row A from a to $a' = a + x$.

If the agent doesn't execute the path extension, she'll need to pick between A and B without observing x. So her expected payoff will be

$$\max(E[a'], b) = \max(a, b).$$

If the agent executes the path extension, her expected payoff will be

$$E[\max(a', b)] - q_f,$$

where q_f is the cost of executing the path extension. The value of executing the path extension is the difference between the previous two expressions:

$$E[\max(a', b)] - \max(a, b) - q_f.$$

This can be rewritten as,

$$\int_{-\infty}^{-|a-b|} \frac{|x|}{\sigma} \phi\left(\frac{x}{\sigma}\right) dx - q_f. \tag{1}$$

Here $\phi(\cdot)$ is the standard normal density function.[8]

To gain intuition for equation 1, begin by assuming (without loss of generality) that $a \geq b$. In this case, the decision-maker only learns something useful when the new analysis leads to a revised value a' that lies below the next best alternative b. Since x represents the new information (i.e., $a' = a + x$), equation 1 integrates over x values that are less than $-|a-b|$. Finally, $\frac{1}{\sigma}\phi\left(\frac{x}{\sigma}\right)$ represents the density of x, since $x \sim N(0, \sigma)$.

For additional intuition consider Figure 2, which represents $\int_{-\infty}^{-|a-b|} \frac{|x|}{\sigma}\phi(\frac{x}{\sigma})dx$ graphically. The shaded area under the density in Figure 2 represents the probability densities of states of the world in which row a has been revealed to be inferior to row b: i.e., $a' < b$. Naturally these are the states of the world

[8] Integrating by parts yields the alternative expression,

$$-|a-b|\Phi\left(-\frac{|a-b|}{\sigma}\right) + \sigma\phi\left(\frac{a-b}{\sigma}\right) - q_f,$$

where $\Phi(\cdot)$ is the cumulative of the standard normal distribution.

in which it is ex-post useful to have analyzed row a. In these states the decision-maker learns that row b is a better choice than row a.

Three fundamental comparative statics are captured in this option value framework. First, the value of a path exploration declines as the cost q_f of the exploration rises.

Second, the value of a path exploration falls with the variability of the information that will be obtained: σ. In other words, the less information that is likely to be revealed by a path exploration, the less valuable such a path exploration becomes. This can be seen graphically by comparing Figures 2 and 3. Figure 3 represents $\int_{-\infty}^{-|a-b|} \frac{|x|}{\sigma^-}\phi(\frac{x}{\sigma^-})dx$, for $\sigma^- < \sigma$. As σ falls, so does the probability of reaching states in which $x < -|a-b|$.

Third, the value of a path exploration falls the larger the gap between a and b. As the gap rises, the interval of states for which $x < -|a-b|$ grows smaller. This can be seen graphically by comparing Figures 2 and 4. Figure 4 represents $\int_{-\infty}^{-|a-b^-|} \frac{|x|}{\sigma}\phi(\frac{x}{\sigma})dx$, for $b^- < b$. As b falls, so does the measure of states in which $x < -|a-b|$.

In Gabaix and Laibson (2000b) we apply these option theoretic considerations to twelve tree games including the game reproduced in Figure 1. We use the option value calculation in Equation 1 to measure the ex-ante expected value of a wide range of path extensions. The directed cognition algorithm then executes the path extension with the highest expected value net of cognition cost, q. The algorithm continues in this way — evaluating the expected value of path extensions and executing the most promising path extension — until no remaining path extension has a positive expected value net of cognition cost.

This procedure radically simplifies analysis of our decision trees. Each unsimplified tree has approximately 100,000 paths leading from the first column to the last column. Our directed cognition algorithm, restricts attention to only a handful of these paths. For example, Figure 5 plots the path extensions generated by the directed cognition model for the game in Figure 1. Using ex-ante option-theoretic considerations, the algorithm restricts analysis to $\frac{7}{100{,}000}$ of the possible paths in this tree. After these seven path extensions are executed, the algorithm stops. In this sense, the directed cognition provides a partial solution to the frame problem. Option-theoretic intuitions enable the decision-maker to simplify his problem by restricting attention a tiny fraction of the available information.

We have tested the directed cognition model by comparing the predictions of the model to the actual experimental choices made by 252 Harvard undergraduates. We measured the Euclidean distance between the model predictions and the empirical choices of our subjects. We then measured the Euclidean distance between the predictions of the rational choice model (with zero cognition cost) and the empirical subject choices. Using this metric, the directed cognition model outperformed the rational choice model with an associated t-statistic of 5.2.

We also compared the directed cognition model to three variants of the rational model that prune the decision trees by discarding information in the right-most columns of the trees: the *column cutoff model* the *column discounting model*, and the *follow the leaders model*.

The column cutoff model assumes that decision-makers calculate perfectly rationally but pay attention to only the first Q columns of the C-column tree, completely ignoring the remaining $C - Q$ columns. The parameter Q is estimated to maximize the fit of the model. The column discounting model assumes that decision-makers follow all paths, but exponentially discount payoffs according to the column in which those payoffs arise. The discount factor is estimated to maximize the fit of the model. Finally, the follow the leaders model only follows branches that have a marginal probability greater than or equal to some threshold probability. The threshold probability is estimated to maximize the fit of the model.[9]

These three algorithms are designed to capture the idea that decision-makers ignore information that is relatively less likely to be useful. For example, the discounting model can be interpreted as a model in which the decision-maker has a reduced probability of seeing payoffs in "later" columns. The directed cognition model significantly outperforms the column cutoff model and the column discounting model. The directed cognition model statistically ties the follow the leaders model.

The directed cognition model simplifies analysis of complex problems and partially replicates the decisions that subjects make. The directed cognition model exploits three basic principles. First, cognition is costly. Second, cognition is more useful when it reveals a lot of new information. Third, cognition is more useful when it reveals information about choices that are

[9] From any starting box b follow all branches that have a probability greater than or equal to p. Continue in this way, moving from left to right across the columns in the decision tree. If a branch has a probability less than p, consider the box to which the branch leads but do not advance beyond that box. Weight all boxes that you consider with their associated cumulative probabilities, and calculate the weighted sum of the boxes.

relatively close competitors. Using these basic ideas we were able to dramatically simplify the analysis of complex decision trees, generating behavioral predictions that match subject choices.

In addition, the directed cognition model can be easily generalized. In the version of the model described above, the cognitive operation in which the decision maker engages is to look forward into the tree. To generalize the model, other cognitive operations could also be considered. For example, if the middle of the decision tree contained occasional large outlier payoffs, it might make sense for the decision maker to start at those large payoff boxes and then backwards induct back toward the first column. Such backward induction is a cognitive operation that can be evaluated using the same option value calculations described above.

Although the directed cognition model can be applied to a wide range of choice problems (e.g., all decision trees and sequential extensive form games), the model has several critical limitations, which jointly imply that the model is *not* a general solution to the frame problem. Most importantly, the directed cognition model relies on option values that may take substantial time to calculate, and which can only be calculated with sophisticated knowledge about the decision problem. For example, the model assumes that the decision maker has rational expectations about σ, the standard deviation of signals arising from additional analysis of the problem. Real actors may not have such prior knowledge, though we believe that they intuitively crudely infer such standard deviations from their previous experience with analogous problems. We imagine that our model's formal option value calculations represent proxies for the informal option value intuitions that guide real decision makers. These informal intuitions are undoubtedly more experienced-based and backward looking than the rationally forward-looking calculations implied by the model.

Finally, we believe that decision makers use a wide range of complementary tools to limit the number of option-value intuitions/calculations that they execute. For example, agents may only execute option value analysis on concepts called to mind by environmental cues. We don't think about global warming (while planning a cross-country trip), unless an explicit cue calls this issue to our attention. Without such attention filters our agent might be paralyzed by myriad option value calculations, which would have the unintended consequence of slowing down his analysis instead of speeding it up.

A true resolution of the frame problem would need to explicitly model such attention filters. We didn't need attention filters in our decision tree analysis, since the number of option value calculations were limited by the forward induction framework that we adopted. Branching paths leaving a particular starting row in Figure 1 were endogenously pruned by the model before they had branched too finely. In more general problems, the analysis isn't so well-ordered and the potential option value calculations aren't so tightly bounded. For these general problems, (e.g., walking across Central Park in Manhattan), attention filters are probably necessary equipment.

Zenon Pylyshyn, one leading contributor to the frame problem literature, recounts a shaggy dog story in which an engineer boasts that he has invented a perpetual motion machine. The complicated machine is revealed with great fanfare but no parts are moving and one onlooker asks why. The inventor responds, "Well the machine is all finished except for one minor little piece that is still on order and is expected any day now. It's just a small ratchet that fits in here and goes back and forth forever (1996, p. xii)." Pylyshyn suggests that solutions to the frame problem suffer from the same problem. "It usually ends up with the critical element being on order." Attention filters are one of the missing ratchets in our machine.

We hope that our model's shortcomings will serve primarily to entice other economic researchers to enter the fray. The frame problem represents a fundamental challenge to researchers who build models of intelligent agents: How do agents allocate scarce resources like cognition and attention? We imagine that economists may have something useful to say about this interesting and important allocation problem.

3 Conclusion

The frame problem is the challenge of developing intelligent machines that can think quickly *and* effectively about complex problems. A solution to the frame problem would provide a theory of endogenous simplification of complex problems.

The option value approach provides a small step in this direction. Our directed cognition model provides a framework for deciding how deeply to analyze decision problems. When cognitive analysis is likely to yield little new insight, the option value of continued analysis declines. This option value also declines when a particular choice gains a large edge over the available alternatives. Our directed cognition model quantifies these two effects and integrates them in a formal framework that makes sharp quantitative predictions about behavior. We have successfully tested these predictions in an experiment conducted on Harvard undergraduates.

Future work will sharpen this analysis by actually measuring the allocation of *attention* within decision problems. Computer based experiments enable researchers to determine exactly what information subjects are analyzing moment by moment.[10] With these kind of real time observations, we will be able to directly test the attention allocation predictions of the option value approach.

Ultimately, we expect that theories of endogenous simplification will explain many deviations from pure rationality. In Gabaix and Laibson (2000) we use the directed cognition model to explain psychological phenomena like myopia, salience, and anchoring. Understanding how decision-makers simplify complex problems will reveal why they occassionally make mistakes, providing a possible microfoundation for anomalous decision-making.

4 References

Camerer, Colin, E. J. Johnson, T. Rymon, and S. Sen, "Cognition and Framing in Sequential Bargaining for Gains and Losses," in Ken Binmore, Alan Kirman, and P. Tani, eds., Frontiers of game theory. Cambridge, MA: MIT Press, 1994, pp. 27-47.

Conlisk, John. "Why bounded rationality?" *J. Econ. Literature*, June 1996, *34* (2), pp.669-700.

Dennett, Daniel, "Cognitive Wheels: The Frame Problem of AI," in Z. W. Pylyshyn (Ed.) *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Ablex Publishing Corporation: Norwood N.J., 1987, pp. 41-64.

Dennett, Daniel, "Producing Future by Telling Stories," in K. M. Ford Z. W. Pylyshyn (Eds.) *The Robot's Dilemma Revisited: The Frame Problem*

[10] We are currently running experiments that ask subjects to analyze problems on a computer screen. By using Mouselab (Payne et al 1993), a programming language that enables experimenters to record the exact location of the mouse cursor, and by only revealing information in a neighborhood of the cursor, we can determine what information subjects are using from moment to moment.

in Artificial Intelligence, Ablex Publishing Corporation: Norwood N.J., 1996.

Dietrich, Eric and Chris Fields, "The Role of the Frame Problem in Fodor's Modularity Thesis: A Case Study of Rationalist Cognitive Science," in K. M. Ford Z. W. Pylyshyn (Eds.) *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*, Ablex Publishing Corporation: Norwood N.J., 1996.

Dixit, Avinash K. and Robert S. Pindyck, *Investment under uncertainty*, Princeton Unveristy Press: Princeton N.J., 1994.

Ford, Kenneth M. and Zenon W. Pylyshyn (Eds.) *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*, Ablex Publishing Corporation: Norwood N.J., 1996.

Gabaix, Xavier and David Laibson. "A boundedly rational decision algorithm" *AER Papers and Proceedings*, May 2000a.

Gabaix, Xavier and David Laibson, "Bounded rationality and directed cognition," MIT and Harvard mimeo, 2000b.

Glymour, Clark, "The Adventures Among the Asteroids of Angela Android, Series 8400XF with an Afterword on Planning, Prediction, Learning the Frame Problem, and a Few Other Subjects," in K. M. Ford and Z. W. Pylyshyn (Eds.) *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*, Ablex Publishing Corporation: Norwood N.J., 1996, pp. 25-34.

Jéhiel, Philippe. "Limited horizon forecast in repeated alternate games," *Journal of Economic Theory*, 1995, 67: 497-519.

Kahneman, Daniel and Tversky, Amos. "Judgment Under Uncertainty: Heuristics and Biases". *Science*, 1974.

McCarthy, J. and P. J. Hayes. "Some philosophical problems from the standpoint of artificial intelligence," In B. Meltzer and D. Michie (eds.) *Machine Intelligence,* Vol 4, pp. 463-502 Edinburgh: University of Edinburgh Press, 1969.

Payne, John W., James R. Bettman, and Eric J. Johnson, *The Adaptive Decision Maker*, 1993, Cambridge: Cambridge University Press.

Pylyshyn, Zenon W. (Ed.) *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Ablex Publishing Corporation: Norwood N.J., 1987.

Shanahan, Murray, *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*, Cambridge: MIT Press, 1997.

Shugan, Steven M. "The Cost of Thinking," *Journal of Consumer Research,* 7(2), 1980, pp. 99-111.

Simon, Herbert. "A behavioral model of rational choice," *Quart. J. Econ.*, Feb 1955, 69 (1).

Thaler, Richard. *Quasi Rational Economics*, New York: Russell Sage, 1991.

Thaler, Richard. *The winner's curse.* New York: Russell Sage, 1994.

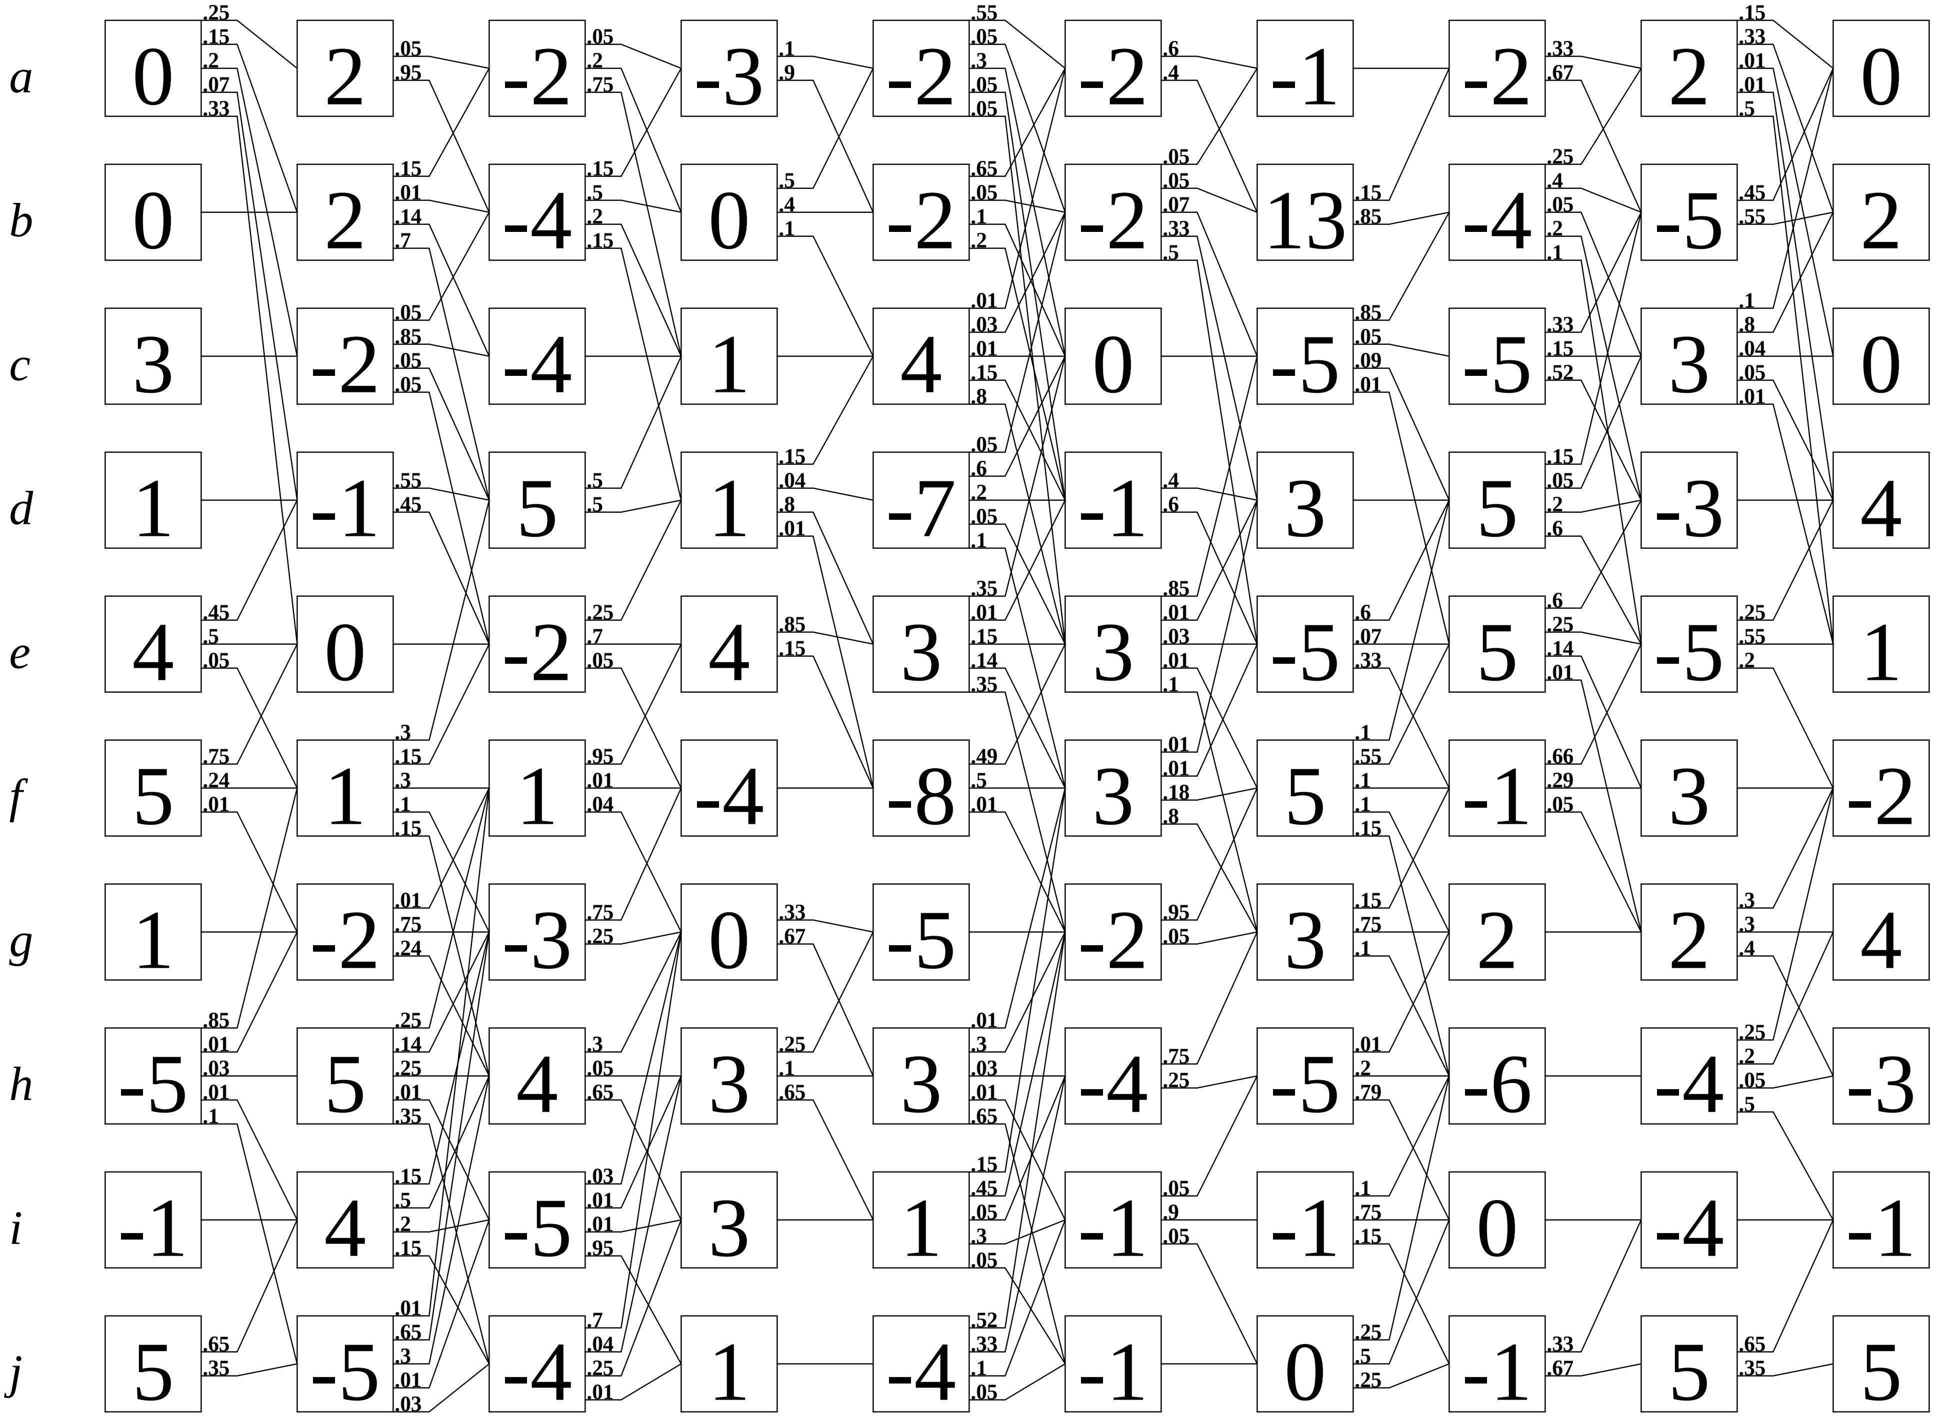

a 0 2 -2 -3 -2 -2 -1 -2 2 0
b 0 2 -4 0 -2 -2 13 -4 -5 2
c 3 -2 -4 1 4 0 -5 -5 3 0
d 1 -1 5 1 -7 -1 3 5 -3 4
e 4 0 -2 4 3 3 -5 5 -5 1
f 5 1 1 -4 -8 3 5 -1 3 -2
g 1 -2 -3 0 -5 -2 3 2 2 4
h -5 5 4 3 3 -4 -5 -6 -4 -3
i -1 4 -5 3 1 -1 -1 0 -4 -1
j 5 -5 -4 1 -4 -1 0 -1 5 5

Figure 2: Calculation of option value

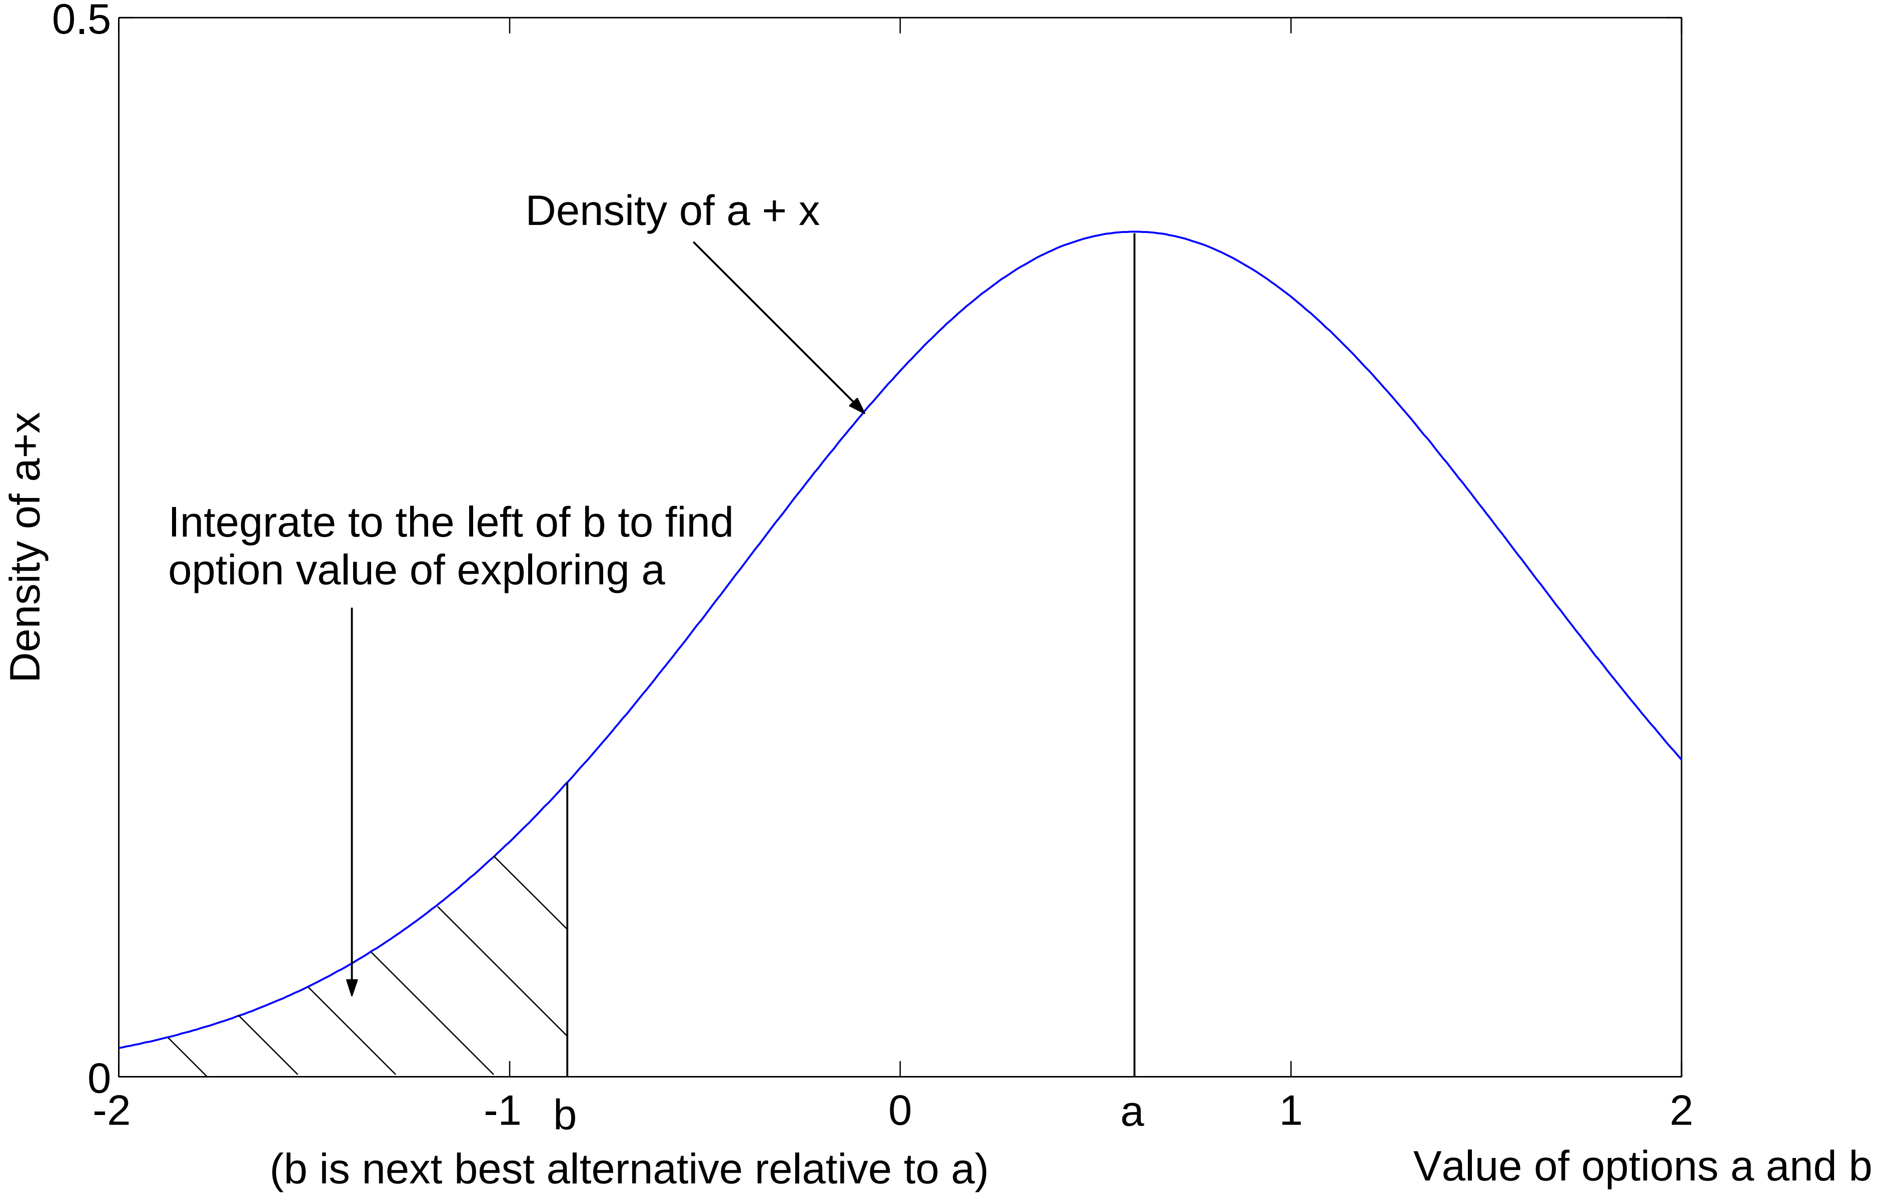

Figure 3: Calculation of option value (low variance case)

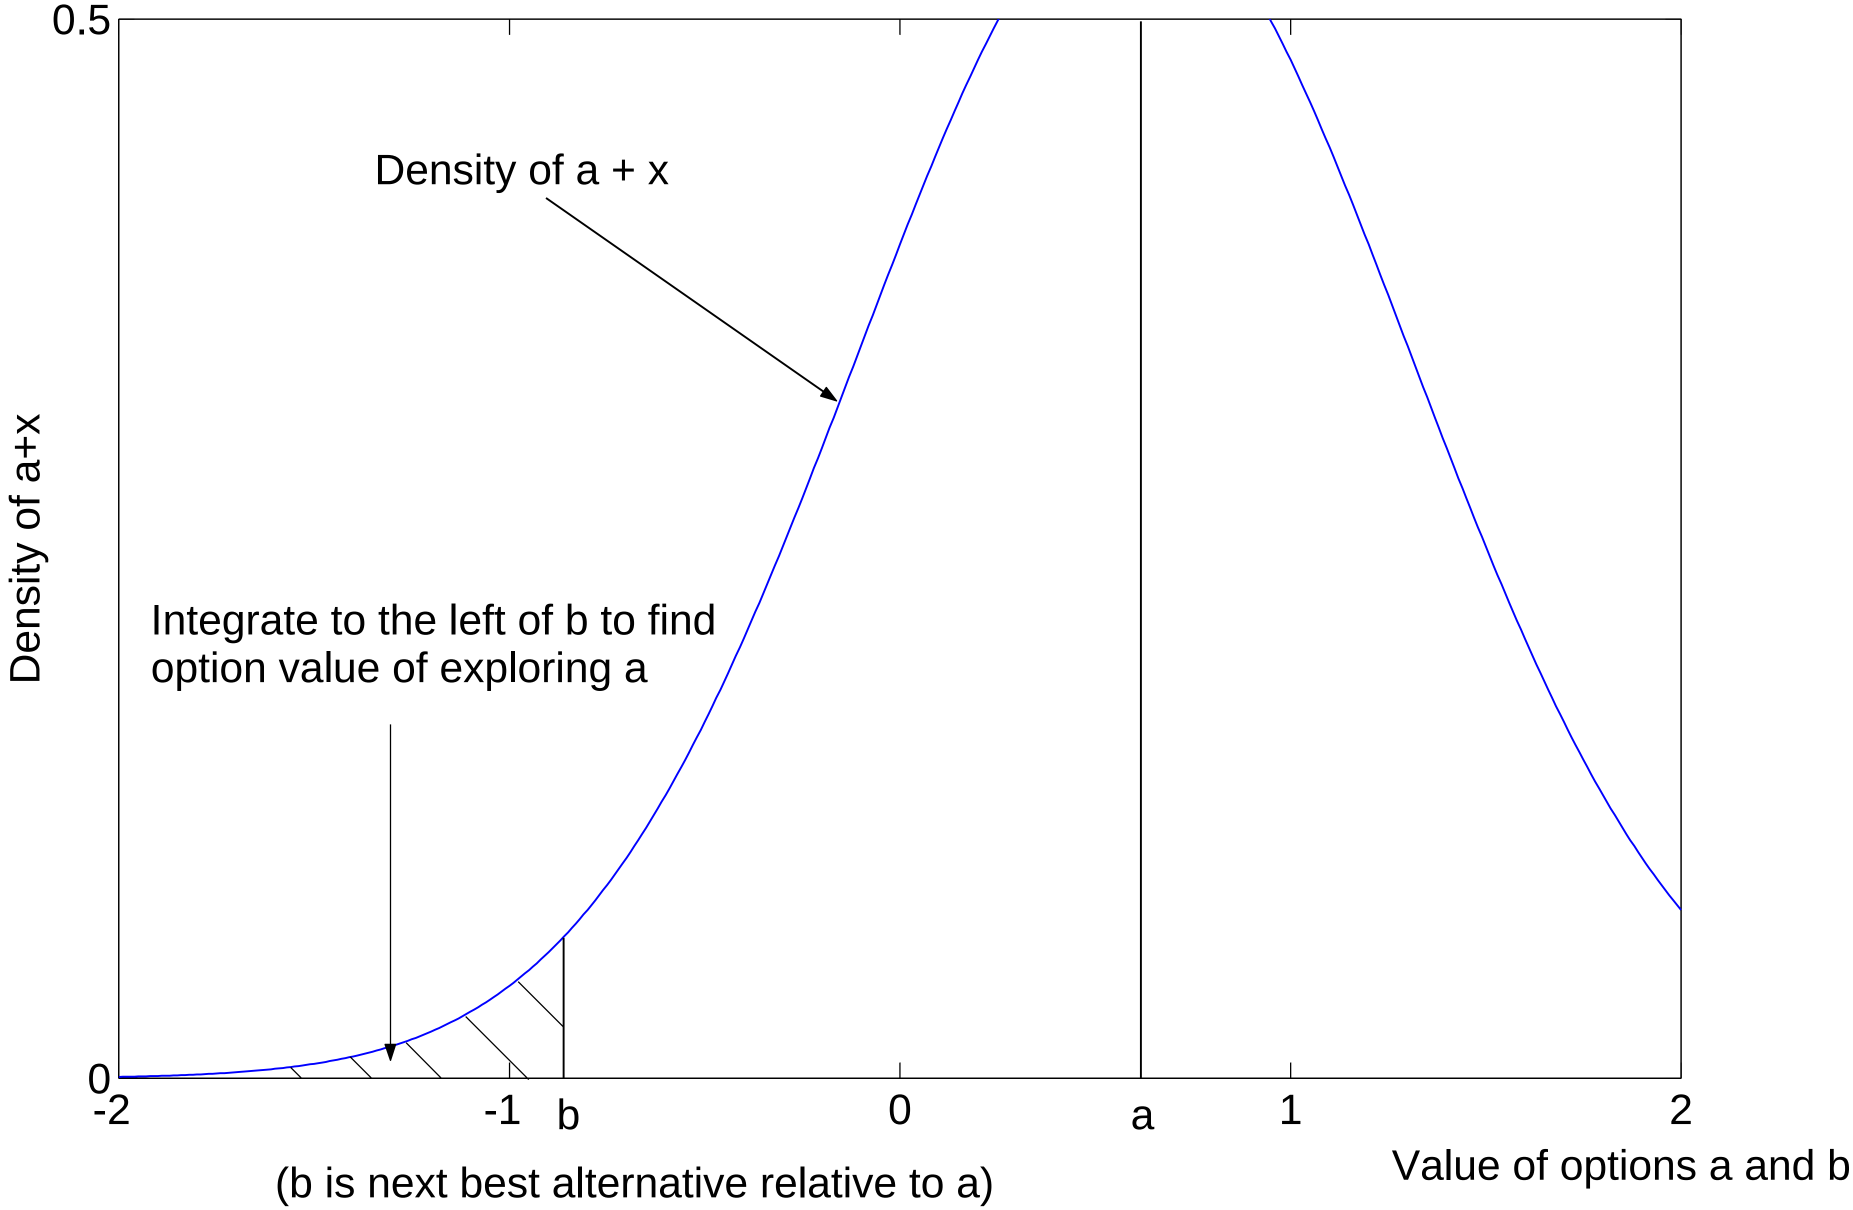

Figure 4: Calculation of option value (low b case)

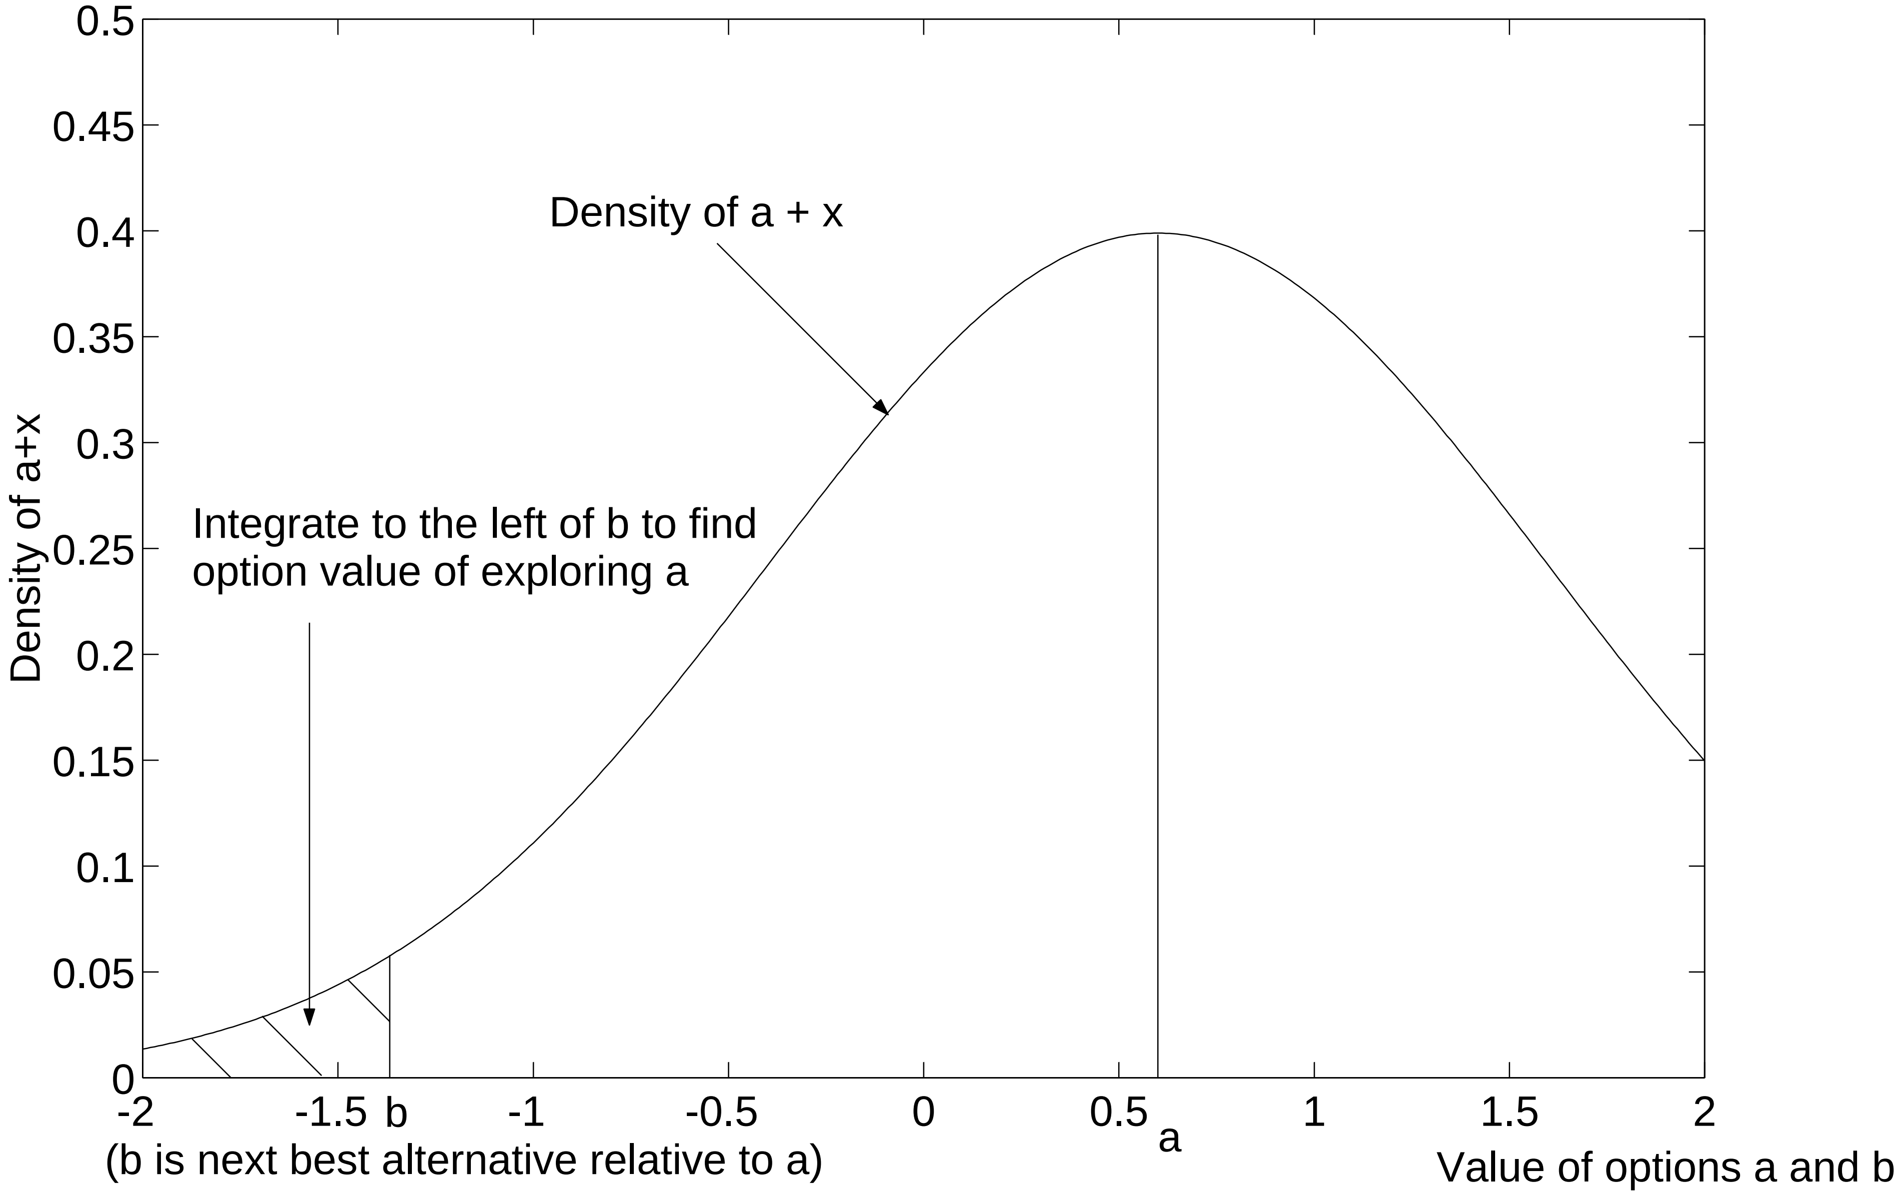

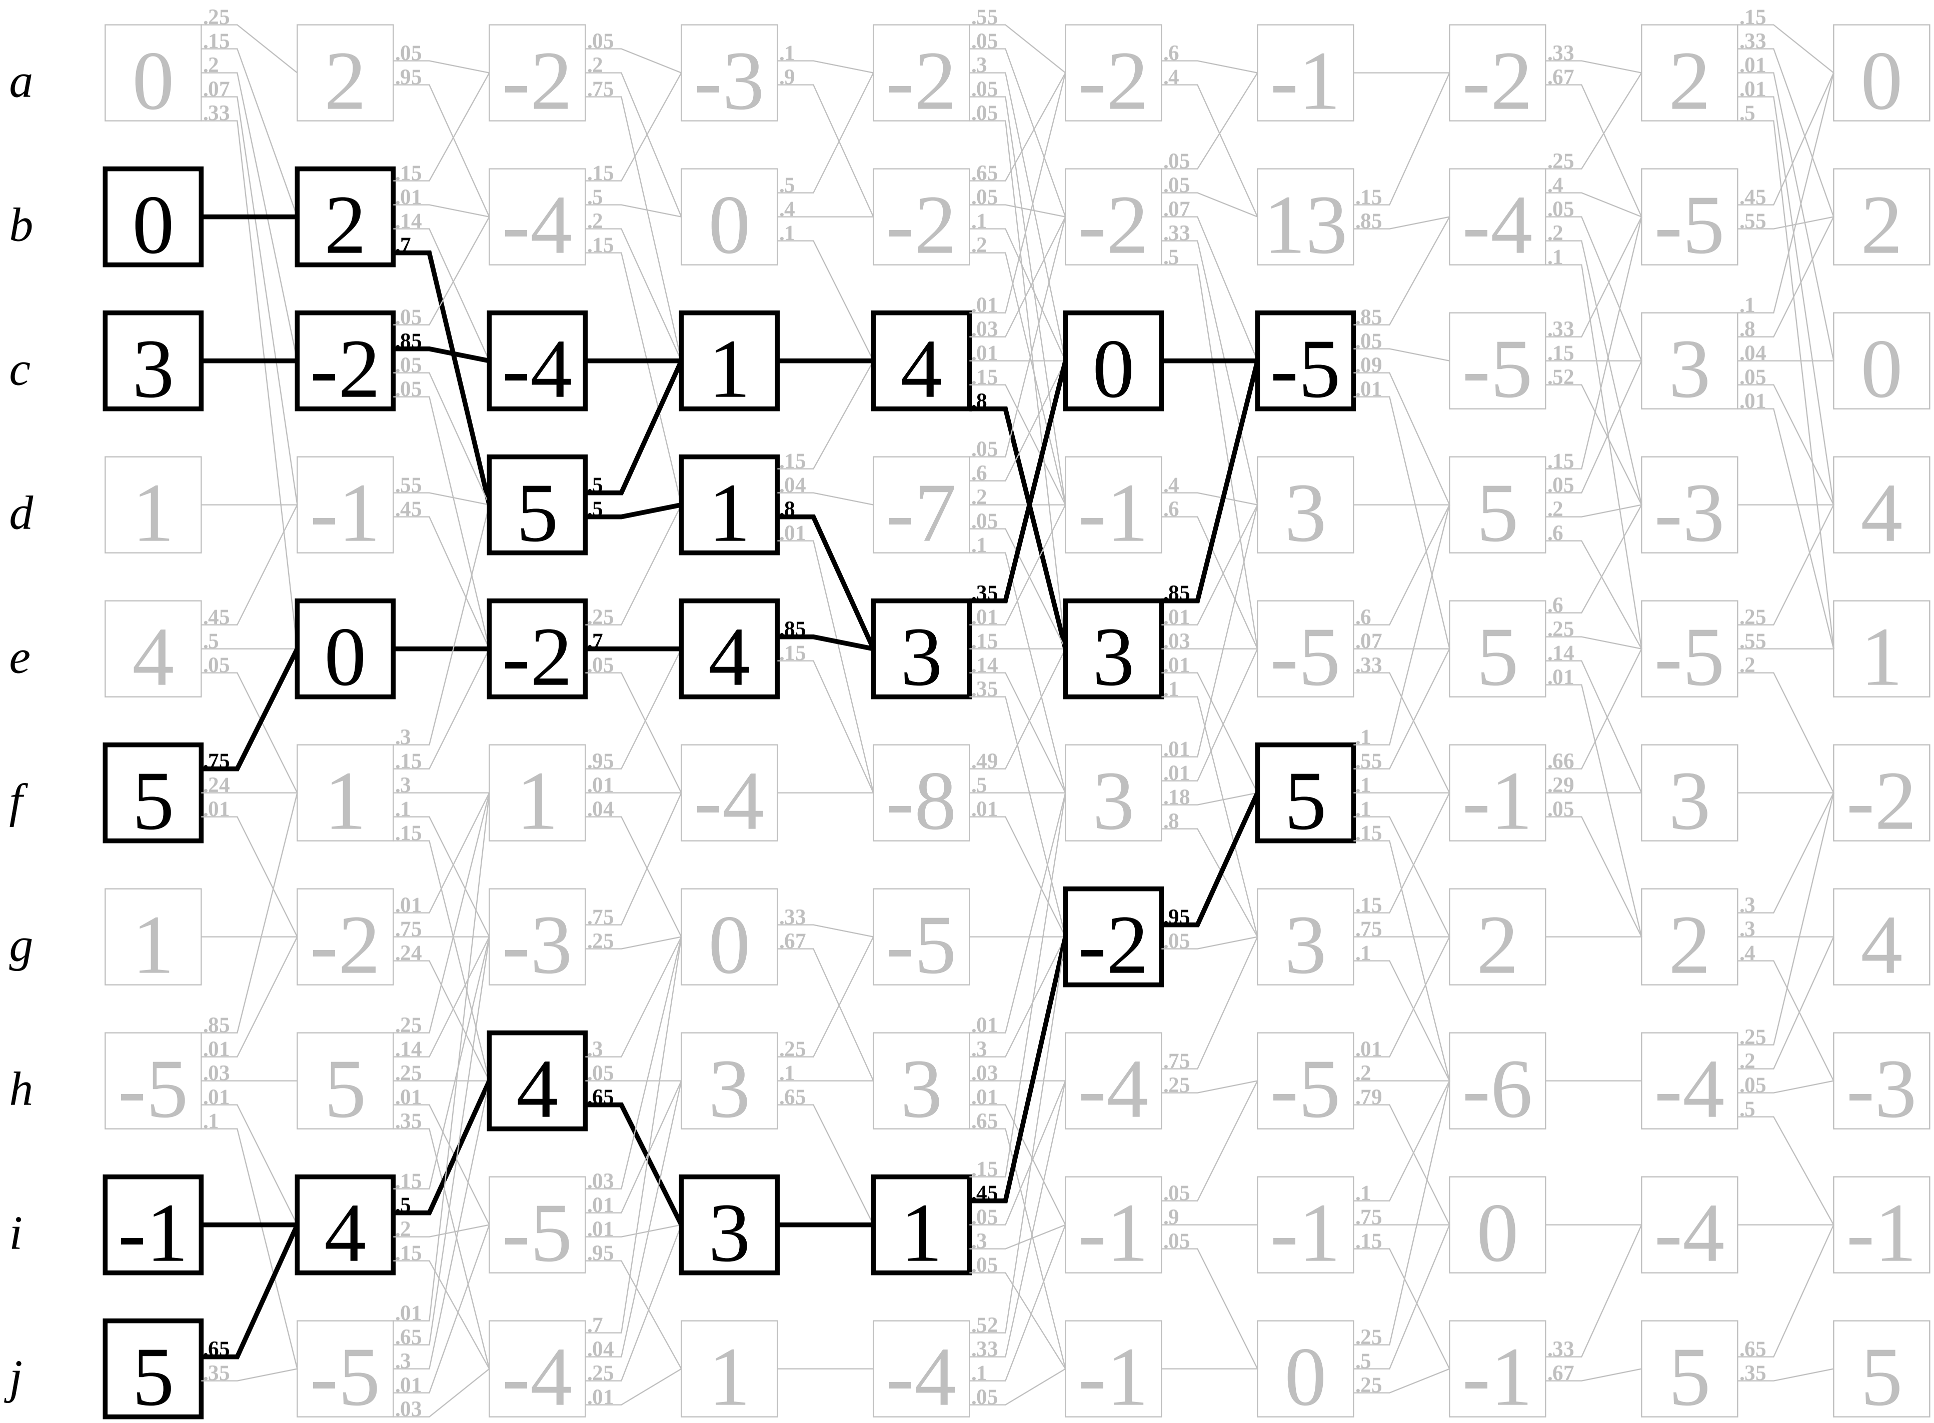
a
b
c
d
e
f
g
h
i
j