
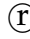


ADAPTIVE CARDINAL UTILITY¹

Arthur J. Robson²  Lorne A. Whitehead  Nikolaus Robalino³
December, 2022


ABSTRACT

This paper models the real-time adaptation of utility when individuals have a limited capacity to make fine distinctions. A simple adaptive mechanism is proposed that adapts utility to an arbitrary distribution of rewards in a way that is ultimately optimal, in terms of minimizing the probability of errors. A modified simple mechanism is approximately optimal in terms of maximizing expected utility. The model generates the so-called “hedonic treadmill” and induces a preference for rising consumption.

Keywords: Utility, cardinal, evolutionary biology, adaptation, bounded rationality.
JEL Codes: A12, D11.

1. INTRODUCTION

A strand of recent literature in economics has shown how adaptive decision utility functions might arise from a limited capacity to make fine distinctions. The first of these papers is Robson (2001) where perception is limited by a small number of thresholds. In order for the individual to minimize the probability of error, the thresholds should be at the quantiles of the distribution of rewards. The resulting

¹The order of names was determined by the AEA Author Randomization tool, as indicated by the symbol . We thank two referees and the editor for their comments. We also thank Carlo Alos-Ferrer, Andrew Caplin, Paul Glimcher, Kenway Louie, Antonio Rangel, Luis Rayo, Phil Reny, Larry Samuelson, Rava da Silveira, Jakub Steiner, Zhifeng Sun, Ryan Webb and Michael Woodford for helpful discussions. Finally, we thank audiences at numerous conferences and workshops for generous and insightful remarks. Robson thanks the Canada Research Chairs Program and the Social Sciences and Humanities Research Council of Canada for financial support.

²Corresponding author.

³Arthur J. Robson, Department of Economics, Simon Fraser University, Burnaby BC Canada V5A 1S6. Email: robson@sfu.ca.

Lorne A. Whitehead, Department of Physics and Astronomy, University of British Columbia, Vancouver BC Canada. Email: lorne.whitehead@ubc.ca.

Nikolaus Robalino, Department of Economics, Rochester Institute of Technology, Rochester, NY USA 14623. Email: ndrobalin@gmail.com.

utility function, adapted to the underlying distribution of rewards, is cardinal. In particular, utility is steepest where outcomes arise most frequently.

The second directly related paper is Netzer (2009), who investigates the Robson model using the expected fitness criterion instead of the probability of error. He shows that, whereas the Robson approach generates a limiting density of thresholds given by f , where f is the pdf of the distribution of rewards, the expected fitness approach generates a limiting density of thresholds proportional to $f^{2/3}$. This again puts fewer thresholds where f is low, but to a less dramatic extent. Intuitively, the expected fitness criterion is more concerned with low f than is the probability of error criterion, since, although the probability of error is small if f is small, the size of the error is large. Again utility must be adapted to the distribution of rewards.

Rayo and Becker (2007) address the issue of adaptation using an alternative model of bounded rationality. Individuals cannot maximize expected utility precisely, but all choices that come within some utility band are considered equivalent. The problem is to construct the optimal utility function. Under simplifying and limiting assumptions, optimal utility is a step function, jumping from 0 to 1 at maximized expected income, so concentrating all incentives at the point of greatest interest, and adapting to the distribution. They do not, however, consider any mechanism for rapid dynamic adaptation.

All these previous papers describe how adaptation would be advantageous and make predictions presuming this has occurred; the key contribution of the present paper is to suggest *how*—to provide an explicit real-time low-rationality adjustment mechanism. This paper dynamically extends Robson (2001), and Netzer (2009) by explicitly modelling the rapid adaptation of utility when an agent has a limited capacity to make fine distinctions. Our main result is that there are simple low-information adaptive rules that are ultimately optimal when the goal is to minimize the probability of error, and approximately optimal when the goal is to maximize fitness.

More specifically, we consider an individual making a binary economic choice. Two options have been drawn iid according to a cdf F . The mapping from actual

value of an outcome to the value used in the choice is a deterministic step function which generates adaptation by means of endogenously setting the location of the steps, the thresholds. Our key new contribution is to model rapid automatic adaptation, and to discuss how such adaptation is reflected in economic choices. In particular, we show that there are simple low-information mechanisms that can generate this adaptation.

Two cases are considered. In the first of these, the objective is simply to minimize the probability of error. When an outcome arrives between two thresholds, the thresholds move closer to each other, by a given increment. An irreducible Markov chain with a unique invariant distribution then describes the dynamics of the thresholds. As the increment is made smaller, the invariant distribution puts full weight on the thresholds being at the quantiles of the distribution F . Thus, the thresholds adapt to the distribution in fashion that is ultimately optimal.

In the second of these cases, the objective is to maximize the expected fitness of the chosen option. Now, when an outcome arrives between two thresholds, one of the two thresholds moves closer to the other with a probability that also depends on the distance between the thresholds. As the grid size is shrunk, the invariant distribution puts full weight on the thresholds being such that each threshold has an equal chance of moving to the right or to the left. It is not now possible to achieve exact limiting optimality with a given finite number of thresholds, given compelling restrictions on the information available. Nevertheless, such a configuration is shown to be approximately optimal for a large number of thresholds.

The model generates a tradeoff between the speed of adjustment and accuracy. It accounts immediately for the “hedonic treadmill”—that is, the reversion of average utility to its original level despite a vast shift in the distribution of rewards. The real time adaptation of utility, as reflected in the hedonic treadmill, suggests a reconsideration of national happiness measures. Finally, it produces a preference for rising consumption streams.

The present model is intended to provide a biological basis for adaptive utility. The sense in which it is “evolutionary” is worth clarifying. A fully fledged evolutionary approach would consider how optimal behavior could emerge as the winner

from a field of arbitrary possibilities. This emergence might be expressed as slow dynamic process of natural selection under which better choices generate more offspring and so come to dominate the population. Although this is certainly of substantial interest, this is *not* the approach here. Rather, the current approach cuts to the chase and considers only the long run result of natural selection. Although genetic evolution can be complicated and it is not guaranteed that the long run outcome of natural selection is the type maximizing offspring, it is reasonable to look at this case as central. Essentially, that is, we adopt the “phenotypic gambit” as described and advocated by Grafen (1984).

To emphasize, the dynamic process that is described explicitly here is not a fully fledged evolutionary approach. Rather, this dynamic process is considered part of the long run outcome of a fully fledged evolutionary approach. That is, natural selection found that the current explicit dynamic rule was advantageous in terms of long run genetic evolution. Such an approach does raise many issues, but most of them cannot be adequately addressed without a much richer understanding of the connection between genes and behavior than is currently available.

Recent neuroscience research provides a backdrop to the current investigation, a backdrop suggesting that deeper and more precise connections to economics will eventually be made. For example, economic decisions can be related to the electrical activity of neurons that release dopamine, a “neurotransmitter”. (See Stauffer *et al* (2014), for example.) Dopamine neuron activity does not, however, directly reflect the actual hedonic experience of consuming a particular option. Dopamine neurons are activated when the actual reward received exceeds the reward that was anticipated, so there is a “reward prediction error”. Intriguingly, Sharot *et al.* (2009) have shown that dopamine also signals an *expectation* of pleasure from actual consumption of an option, as relevant when an economic choice must be made.

2. THE MODEL

Consider the following model of how a choice is made. Each option has a value $y \in [0, 1]$. The individual does not directly observe this value, but anticipates

the hedonic consequences, $h(y)$, of consuming y . The function h is necessarily inaccurate, reflecting a limited ability to make fine perceptual distinctions. This generates a benefit from adaptation.

More specifically, utility $h : [0, 1] \rightarrow \{0, \delta, 2\delta, 3\delta, \dots, N\delta = 1\}$, is a non-decreasing step function characterized by thresholds x_n , at which a jump is made from one level $(n - 1)\delta$ to the next higher level $n\delta$, for $n = 1, \dots, N$. We have $0 \leq x_1 \leq \dots \leq x_N \leq 1$ where we set $x_0 = 0$ and $x_{N+1} = 1$. Adaptation of h is captured by shifts in the thresholds $x_n, n = 1, \dots, N$. The thresholds are a technical device to render the adjustment process tractable, since it involves adjusting a finite number of parameters.

For simplicity, we concentrate on the case where a binary choice is made. There are then two options, y^1 and y^2 , say, generating anticipated values, $h(y^i), i = 1, 2$. If $h(y^1) > h(y^2)$, y^1 is chosen, as is clearly optimal. If $h(y^1) = h(y^2)$, each option is chosen with probability $1/2$.

It is without much loss of generality to suppose that y represents fitness. That is, if y instead represents money, for example, which generates fitness according to an increasing concave function, only minor notational changes need to be made. The $y^i, i = 1, 2$, are assumed to be independent, distributed according to the same continuous cumulative distribution function, F , with continuous probability density function, $f > 0$ on $[0, 1]$. Although the realizations are random *ex ante*, they are realized prior to choice. The distribution F nevertheless plays an important role because the thresholds must be set in the light only of F rather than the realizations.

Errors here arise only when both $y^i, i = 1, 2$, lie in the same interval $[x_n, x_{n+1})$. Minimizing the probability of error implies that the thresholds should be equally spaced in terms of probabilities; should be then at the quantiles of the distribution.

A preferable criterion is maximization of expected fitness. The analysis is more complex in this case. Each threshold now ought to be at the *mean* of the distribution conditional on being between the two neighboring thresholds.⁴ It is then not

⁴Suppose there is only one threshold. Maximizing expected fitness is the same as minimizing the expected fitness loss from a mistake. Intuitively, placing the threshold at the mean of the distribution

possible to use purely qualitative observations to estimate the mean, as was true for the median. Nevertheless, approximately optimal rules will be derived that use the distances to the two neighboring thresholds.

Under either criterion, the thresholds should be denser where f is higher. Concentrating the the thresholds like this reduces the probability of error or the expected size of the loss.

The key contribution here is then to address the question: How could the thresholds adjust to a novel distribution, F ? We respond by showing that such adjustment can be achieved by a simple informationally undemanding mechanism. This is a “reverse engineering” approach in that the simple rule attains the optimal configuration of thresholds in the limit as the grid size shrinks, for the probability of error case. Further, a modified simple rule attains approximate optimality, in the expected fitness case, as the number of thresholds increases. We formulate a model to address these questions.

In modelling how the thresholds adjust, we make the simplifying assumption that there is a single stream of outcomes, represented as y , and abstract from the choices made. Alternatively, we could interpret the analysis as supposing that y^1 and y^2 arrive in alternate periods, with the system adapting to each of them, and a choice between them being made in every even period.

Suppose then, for simplicity, the thresholds are confined to a finite grid $\mathcal{G}_\epsilon = \{0, \epsilon, 2\epsilon, \dots, (G-1)\epsilon, 1\}$, for an integer G such that $G\epsilon = 1$. The thresholds are time dependent, given as $x_n^t \in \mathcal{G}_\epsilon$, where $0 \leq x_1^t \leq \dots \leq x_N^t \leq 1$, at time $t = 1, 2, \dots$

3. PROBABILITY OF ERROR CASE

Consider the rule of thumb for adjusting the thresholds —

Definition 3.1. Rule of Thumb. Suppose the period t outcome lands in $[x_{n-1}^t, x_n^t)$. If $x_n^t - x_{n-1}^t = \epsilon$, then nothing happens. Otherwise, if x_{n-1}^t is an interior threshold,

balances the expected fitness losses on each side of the threshold. The general result, with $N > 1$, then follows from considering the optimal placement of each threshold relative to its two neighbors.

then $x_{n-1}^{t+1} = x_{n-1}^t + \varepsilon$. If x_n^t is an interior threshold, then $x_n^{t+1} = x_n^t - \varepsilon$. No other threshold shifts when y^t lands in $[x_{n-1}^t, x_n^t)$.

The condition that if $x_n^t - x_{n-1}^t = \varepsilon$, then nothing happens, means that the process will never superimpose one threshold on top of another, and that the order of thresholds will never be reversed at some stage so that $x_{n+1}^{t+1} < x_n^{t+1}$, for example. The thresholds will then be governed by an invariant distribution over all set of all configurations with $x_1^t < \dots < x_N^t$. This condition is without loss of generality in the limit as ε goes to zero, since in this limit the thresholds can be arbitrarily close to one another.

For general N , we have

Theorem 3.1. *In the limit as $\varepsilon \rightarrow 0$, the invariant joint distribution of the thresholds, (x_1^t, \dots, x_N^t) , converges to a point mass at the vector (x_1^*, \dots, x_N^*) , where $F(x_n^*) = \frac{n}{N+1}$, $n = 1, \dots, N$.⁵ That is, in this limit, the thresholds are located optimally.*

Proof. This is a consequence of the more general Theorem 4.1 below which is proved in the Appendix.

The basic property implies that the placement of the thresholds maximizes the rate of Shannon information transfer, as in Laughlin (1981). This argument, which is for a single channel and an abstract criterion, contrasts with that for the present binary choice and minimization of the probability of error.

The property that $F(x_n^*) = n/(N+1)$, for $n = 1, \dots, N$ is that utility, U , say, adapts to the distribution, F . Indeed, in the limit as $N \rightarrow \infty$, $U = F$. If F has a typical unimodal shape, utility has an S-shape similar to that in Kahneman and Tversky (1979).

The rule of thumb illustrates that there exist low rationality mechanisms that can generate fast complete adaptation to an arbitrary distribution. This process is intended to reflect an automatic process that is readily feasible. The thresholds are a device that renders the analysis of adaptation tractable. They still allow the step function h to approximate an arbitrary increasing continuous function.

⁵For any particular $\varepsilon > 0$, the invariant distribution has full support, \mathcal{G}_ε . Only in the limit does this invariant distribution converge to the x_n^* 's.

Although the rule of thumb described above is inevitably slower than the full Bayesian approach it is vastly simpler. Full Bayesian adaptation would be less demanding if the cdf were known to come from a small parametric class—normal distributions with an unknown mean, for example. Without some such severe restriction, the complexity of Bayesian updating seems likely to preclude fast neural responses.

3.1. Class of Rules. In this section, we buttress further the present class of rules. Under some reasonable simplifying assumptions, and when there are only two thresholds, the present rule is the only rule that generates optimal long run choices.

Consider then the following class of rules. Only the interval containing each outcome is detected. To assume more is detected would contradict the spirit of the problem. The rule is Markov, in that only the current configuration of thresholds bears on the choice made next by the rule. The rule shifts any number of thresholds one step in either direction in the grid. The most obvious consequence of this one step assumption is to slow down adaptation, which could be offset by increasing the arrival rate of outcomes. Finally, make the (harmless) simplification that the rule treats thresholds symmetrically.

With two adjustable thresholds, given as $x_1, x_2 \in (0, 1)$, the issue then is: Which thresholds should be moved as a consequence of an outcome in each of the three possible intervals $I_1 = [0, x_1)$, $I_2 = [x_1, x_2]$ or $I_3 = (x_2, 1]$?

In the limit as $\varepsilon \rightarrow 0$, but where the arrival rate of outcomes is sped up in proportion to $1/\varepsilon$, it can be shown that the thresholds adapt in a way described by two linear ordinary differential equations. This is sufficient to obtain the result concerning the limit of the invariant distribution as Theorem 3.1. Indeed, we have the following

Proposition 3.1. *Within the current simple class of rules, with two adjustable thresholds, the adjustment rule from Definition 3.1 is the only one that generates optimal choice in the long-run.*

Proof. See the Appendix.

If there were more than two thresholds, the current rule may not be the only such rule. However, all rules that are optimal in the long run must share a key feature—the interval in which an outcome is observed must shrink. Suppose this were not the case. That is, suppose there is some $j \in \{0, 1, \dots, N\}$ such that the interval $[x_j, x_{j+1})$ does not shrink when an outcome arrives in that interval. It follows that there is a stable configuration with thresholds $0 = x_1 = \dots = x_j$ and $x_{j+1} = \dots = x_N = 1$, so that this rule cannot be long run optimal.

For simplicity, this discussion concerns the probability of error case. However, similar observations apply to the more general maximizing fitness case that is taken up next.

4. MAXIMIZING FITNESS

The most appealing general criterion is to maximize expected fitness. That is, individuals who successfully do this should outperform those who do not.⁶

The situation is now more complicated than it was with the criterion of minimizing the probability of error. There are no longer simple rules of thumb that implement the optimum exactly. However, there do exist simple rules of thumb that implement the optimum *approximately*, for large N . These rules of thumb involve conditioning on the arrival of a realization in the adjacent interval, as above, but also modify the probability of moving using the distance to the next threshold, in a symmetric way.

Although it is possible to accurately estimate the median of a distribution from the limited information available to such a rule of thumb, it is not possible to do this for the mean. Hence the result for the probability of error case are sharper than the results for the expected fitness case.

To see that simple rules of thumb like this cannot implement the optimum exactly, consider first the case that $N = 1$. Suppose that F has median $1/2$ but a mean that is not $1/2$. Consider a symmetric rule of thumb based on the arrival of an outcome to the left or the right of the current position of the threshold at x , say,

⁶This assumes that the risk is independent across individuals. See Robson (1996) for a treatment of this issue. Another possibility would be that fitness depends on relative payoffs.

and the distance to the ends— x or $1 - x$. This will then generate a limiting position for the threshold at $1/2$, thus failing to implement the optimum. This is also an issue for any number of thresholds, since this argument applies to the position of any threshold relative its two neighbors.

It is important that this general rule of thumb uses only the information that is available—the location of the neighboring thresholds and whether an outcome lies in the subinterval just to the right or just to the left. It would contradict the interpretation of the model here to use detailed information about the precise location of the outcome within a subinterval.

At the same time, the general rule of thumb here makes greater computational demands than does the rule of thumb for the probability of error case. The need to utilize the position of adjacent thresholds must entail a greater complexity cost.

The general rule of thumb considered here is—

Definition 4.1. General Rule of Thumb. Suppose the period t outcome lands in $[x_{n-1}^t, x_n^t)$. If $x_n^t - x_{n-1}^t = \varepsilon$, then nothing happens. Otherwise, if x_{n-1}^t is an interior threshold, then independently of any shift in x_n^t , $x_{n-1}^{t+1} = x_{n-1}^t + \varepsilon$ with probability $(x_n^t - x_{n-1}^t)^\beta < 1$, and $x_{n-1}^{t+1} = x_{n-1}^t$ with probability $1 - (x_n^t - x_{n-1}^t)^\beta > 0$. If x_n^t is an interior threshold, then independently of any shift in x_{n-1}^t , $x_n^{t+1} = x_n^t - \varepsilon$ with probability $(x_n^t - x_{n-1}^t)^\beta < 1$, and $x_n^{t+1} = x_n^t$ with probability $1 - (x_n^t - x_{n-1}^t)^\beta > 0$. No other threshold shifts when y^t lands in $[x_{n-1}^t, x_n^t)$.

If the parameter $\beta = 0$, we have the old rule of thumb. Formally, then, Theorem 3.1 follows from Theorem 4.1 below.

The Markov chain defined here is irreducible. That is, there exists a number of repetitions such that, for any initial configuration, (x_1^0, \dots, x_N^0) , say, there is positive probability of being in any final configuration, (x_1^T, \dots, x_N^T) , say. There is therefore a unique invariant distribution for this chain.

If $\beta > 0$ this will encourage the closing up of large gaps that arise where f is small, which is useful to maximize expected fitness. Consider, for example, a threshold situated so that the probability of an outcome in the adjacent interval to the left equals the probability of an outcome just to the right. Suppose, however, that the

distance to the next threshold on the right exceeds the distance to the left, because the pdf, f is lower to the right. It will then pay to move to right, since the expected fitness stakes on the right exceed those on the left. Indeed, if $\beta = 1/2$, the resulting rule will be shown to be approximately optimal for large N .

We have—

Theorem 4.1. *In the limit as $G \rightarrow \infty$ so that $\epsilon \rightarrow 0$, the invariant joint distribution of the thresholds (x_1^t, \dots, x_N^t) converges to one that assigns a point mass to the vector with components x_n^* , $n = 1, \dots, N$. These are the unique solutions to*

$$(F(x_{n+1}^*) - F(x_n^*))(x_{n+1}^* - x_n^*)^\beta = (F(x_n^*) - F(x_{n-1}^*))(x_n^* - x_{n-1}^*)^\beta,$$

for $n = 1, \dots, N$.

Proof. See the Appendix.

Theorem 4.1 straightforwardly extends Theorem 3.1. Again, the limiting position of each threshold is such that the probability of moving to the left is equal to the probability of moving to the right.

4.1. Approximate Optimality of the Rule of Thumb. For each N , there exists a unique positioning of the N interior thresholds, under the rule of thumb, in the limit as $G \rightarrow \infty$ so that $\epsilon \rightarrow 0$. We now consider the limiting properties of the rule of thumb. Consider the efficiency of the rule of thumb relative to the full information ideal, for the expected fitness criterion. (The “full information ideal” is to always choose the higher outcome.) Suppose that the expected deficit in y , for the limiting rule of thumb relative to the full information ideal, is given by $L(N)$. Consider also $L^*(N)$, the expected deficit in y , relative to the full information ideal, when the N thresholds are placed optimally.

Theorem 4.2. *Suppose f is uniformly continuous. Then, as $N \rightarrow \infty$, the limiting efficiency of the rule of thumb is characterized by*

$$(4.1) \quad (N+1)^2 L(N) \longrightarrow \frac{1}{6} \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right).$$

Moreover, as $N \rightarrow \infty$ the density of thresholds under the rule of thumb is

$$(4.2) \quad U(x) = k \cdot \int_0^x f(y)^{\frac{1}{1+\beta}} dy,$$

where k is the normalizing constant $1 / \int_0^1 f(y)^{\frac{1}{1+\beta}} dy$. That is, in the limit, the fraction of thresholds in the interval $[0, x]$ is $U(x)$. Expression 4.1 is uniquely minimized by choice of $\beta = 1/2$. Hence the rule of thumb with the best limiting efficiency satisfies

$$(N + 1)^2 L(N) \rightarrow \frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy \right)^3$$

as $N \rightarrow \infty$.

Proof. See the Appendix. The strict optimality of $\beta = 1/2$ follows from the Hölder Inequality.

The choice of $\beta = 1/2$ gives the best limiting efficiency among the simple adjustment rules considered here. The next result is that the rule of thumb with $\beta = 1/2$ gives the same limiting efficiency, moreover, as do optimally placed thresholds—

Theorem 4.3. *As $N \rightarrow \infty$ the limiting efficiency when the N thresholds are placed optimally is characterized by*

$$(4.3) \quad (N + 1)^2 L^*(N) \rightarrow \frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy \right)^3.$$

Hence the rule of thumb, when $\beta = 1/2$, is approximately optimal for large N . This approximation is additional to those already involved in i) the convergence of the Markov chain to an invariant distribution *and* ii) taking the limit of the invariant distribution as $\epsilon \rightarrow 0$.⁷

⁷The density of optimally placed thresholds is proportional to $f^{\frac{2}{3}}$ in the limit, so the optimal utility function is

$$U(x) = k \int_0^x f(y)^{2/3} dy, \text{ for } k = 1 / \int_0^1 f(y)^{2/3} dy.$$

as was shown by Netzer (2009).

5. SPEED VERSUS ACCURACY

A formal but interesting property of the model is a tradeoff between the speed and accuracy of adjustment. This is controlled by the parameter ϵ —when ϵ is large, adaptation is rapid, but imprecise; if ϵ is small, adaptation is slow but precise.

This tradeoff between speed and accuracy seems bound to be theoretically robust. That is, other models that differ in detail but still capture rapid adaptation seem bound to also produce such a tradeoff.

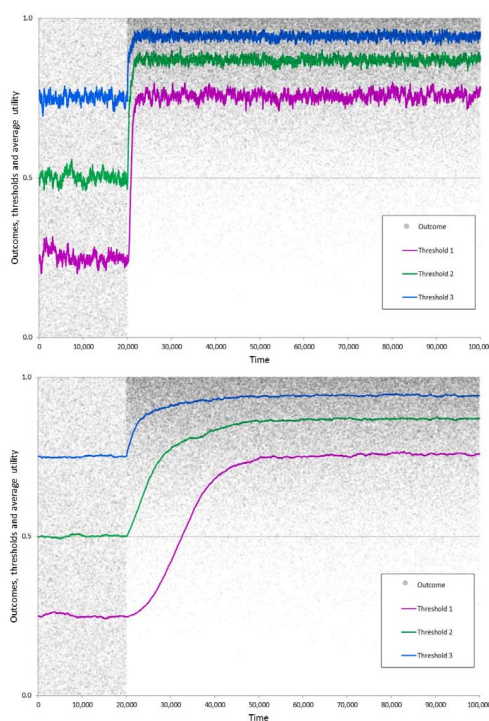


FIGURE 1. Speed versus Accuracy

Figure 1 illustrates these observations in the probability of error case. It depicts the evolution of the three thresholds over time, now contrasting two different values of the grid size ϵ ; namely 0.002, and 0.000125, top and bottom, respectively. It is

evident here that a smaller value of ϵ slows down the speed of adjustment but improves the precision of the ultimate allocation of thresholds.

This discussion could be sharpened by assuming that the underlying cdf, F , was subject to occasional change. Suppose, to be more precise, that there is a (finite, say) set of cdf's $\{F_j\}$. With a Poisson arrival rate, the current cdf from this set is switched to a new one, drawn at random from this set. It is intuitively compelling that there should then be an optimal $\epsilon > 0$ and that this should vary with the rate of introduction of novelty, in particular.⁸

Adaptation should be slow when circumstances change infrequently; but fast when circumstances change frequently. (This would consider the parameter ϵ as endogenous, tailored to the circumstances.) This is consistent with adaptation to living in a new locale taking several years; but adaptation to playing a game of penny ante poker being much faster.

6. HEDONIC TREADMILL

The model has immediate implications for the hedonic treadmill. Schkade and Kahneman (1998) formulate a well-known version of this treadmill that concerned students at the University of Michigan and at UCLA. Students in the two locations reported similar degrees of life satisfaction, but Michigan students projected that UCLA students would be significantly happier.

Schkade and Kahneman describe this as a conflict between “decision utility”—which is applied when deciding whether to move from Michigan, and which is based on an expected substantial increase in life satisfaction in California—and “experienced utility”—which reflects the more modest increase ultimately obtained once there. Schkade and Kahneman argue then that “decision utility” is defective.

The model here entails the adaptation of utility, which implies such a distinction between decision and experienced utilities. There is no sense, however, in which either decision or experienced utility is defective, in contrast to Schkade

⁸In a similar spirit, the number of thresholds might be allowed to vary with the problem at hand. That is, if a problem has particularly high stakes, N might be allowed to increase, but at a cost.

and Kahneman.⁹ Either utility function is *optimal* for the corresponding circumstances, in the light of the inability to make arbitrarily fine distinctions.

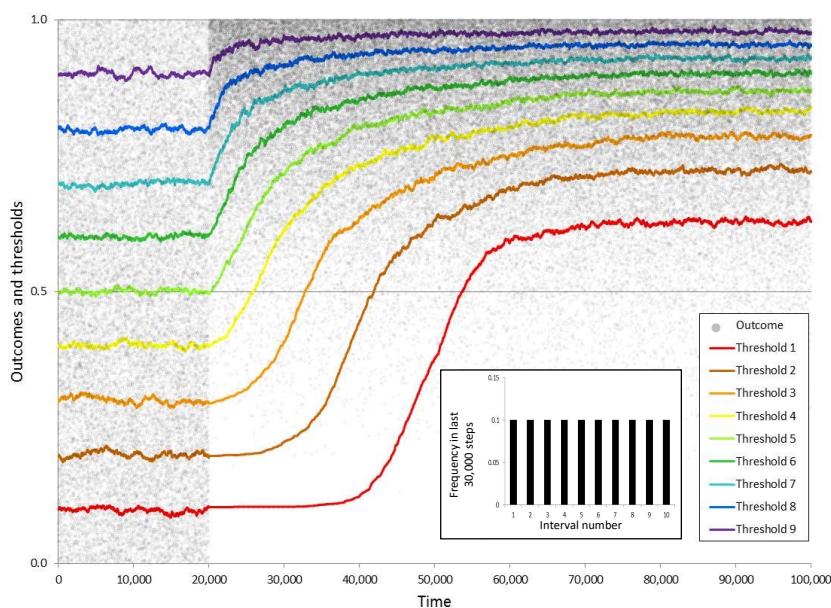


FIGURE 2. Adaptation of the Thresholds to a Novel Distribution.

The results of simulating this version of the model are presented in Figure 2. Consider the class of cdf's given by $F(x) = x^\gamma$ with pdf's $f(x) = \gamma x^{\gamma-1}$, with $\gamma > 0$, for all $x \in [0, 1]$. Suppose $\epsilon = 0.0005$. Consider the probability of error case, for example, so that $\beta = 0$, with nine thresholds, so that these thresholds will be optimally positioned at the deciles of the distribution. Take 100,000 periods, where $\gamma = 1$ for the first 20,000 periods and $\gamma = 5$ thereafter, so that probability mass is shifted to the upper end of the interval $[0, 1]$. Suppose the thresholds are placed initially at 0.1, 0.2, ..., 0.9—that is, at the deciles of the distribution for $\gamma = 1$. This is essentially equivalent to supposing that the $\gamma = 1$ regime has been in effect for

⁹Robson and Samuelson (2011) addressed these issues using the Rayo and Becker (2007) model of a limited ability to make fine distinctions. However, this model, in common with other papers in this small literature, lacks an explicit account of real time adaptation.

a long time. (All the simulations here were done using Excel.) The distribution of thresholds quickly puts most mass near the deciles, as shown by the restoration of the uniform empirical frequency of outcomes in each interval.

Figure 2 illustrates the adaptation induced by the current rule of thumb. An outcome near 0.5, for example, once generated utility near 0.5; after the shift to a substantially more favorable distribution, it generates a rock-bottom level of utility. But the distribution of thresholds was optimal for the original distribution and became optimal again for the new distribution.

For $\beta = 0$, as in Figure 2, average utility reverts completely to its original level, after a shift in the cdf F . This is a pure form of the “hedonic treadmill”. For $\beta = 1/2$, however, reversion is generally incomplete or can overshoot. Figure 3 illustrates partial reversion, presenting a rolling average of utility, where utility is defined so that average utility for $\gamma = 1$ is normalized to 0.5.¹⁰




FIGURE 3. Rule of Thumb with $\beta = 1/2$. Modified Hedonic Treadmill.

¹⁰Average utility needs to be smoothed to be meaningful. We use a rolling average of the last 1,000 periods.

Indeed, long run average expected utility generally depends on the distribution, when $\beta > 0$.¹¹ To illustrate overshooting, consider before and after distributions F and G , respectively, say, such that the average utility under G exceeds that under F . Suppose now that G is scaled down, by halving all outcomes, perhaps, so that F first-order stochastically dominates G . Such scaling ensures that average utility will plunge with the advent of G . However, in the long run, average utility is independent of this scaling-down, no matter how dramatic this might be, and so average utility will eventually climb back past its old level.

Importantly, Figure 2 also serves to demonstrate the robustness of the limiting results Theorems 3.1 and 4.1 which concern limits of invariant distributions as the grid size ϵ tends to 0. That is, these results hold approximately for finite time and reasonable positive grid sizes. Theorems 4.2 and 4.3 rely on taking the additional limit as $N \rightarrow \infty$, then showing that $\beta = 1/2$ yields a rule of thumb that is approximately optimal. Figure 3 then shows $\beta = 1/2$ is approximately optimal even for small values of N . These results are also then robust. Further, although there is a gain from $\beta > 0$, this gain is not overwhelming, relative to $\beta = 0$. The additional complexity cost of rules of thumb with $\beta > 0$ might then outweigh the gain over the rule with $\beta = 0$.

From a welfare point of view, not merely is average utility an unreliable guide as to the extent of the change in the distribution, it may reverse the apparent direction.

The explanation the current model provides for the hedonic treadmill can be applied to the criterion of average national happiness, as has been recently suggested as an alternative to GNP. Clark *et al.* (2018) go beyond the present *positive* approach to utility to embrace a *normative* view. Adaptation raises various issues for the positive view that is adopted here, but it is more problematic as a normative approach. If all monetary outcomes are doubled, for example, this causes an immediate increase in average utility. Nevertheless, average utility is unaffected in the long run. This does not seem a desirable normative property of a welfare criterion.

¹¹Expected utility also depends on the distribution under the *optimal* allocation of a finite number of thresholds.

In the probability of error case, there is *complete* long run adaptation, in general, because $\int F(x)dF(x) = 1/2$ for all cdf's F . This is because the mechanism described here generates purely *relative* valuations.

The expected fitness criterion differs on this score. That is, in the limiting large N case, if the criterion is expected fitness, it follows that $U'(x) = kf(x)^{2/3}$ where k is such that $U(1) = 1$ (as in Netzer, 2009). Although it is still true that $\int U(x)dF(x)$ is fully invariant in the long run to rescaling the cdf F , it now depends on the non-scaleable properties of the distribution.¹² Indeed, average utility can be made arbitrarily close to 1—the absolute maximum value—by taking distributions converging to a point mass at 1 in a particular fashion.

The problem is to maximize $\int_0^1 U(x)f(x)dx$ subject to $F'(x) = f(x)$ $U'(x) = kf(x)^{2/3}$, where $F(0) = U(0) = 0$ and $F(1) = 1$, and where k is such that $U(1) = 1$.

This can be translated to the equivalent problem

$$\max_f \frac{V(1)}{W(1)}$$

subject to

$$F'(x) = f(x); W'(x) = f(x)^{2/3}; V'(x) = W(x)f(x),$$

and $F(0) = W(0) = V(0) = 0$, and $F(1) = 1$. This “Problem of Mayer” is awkward in the sense that, although first-order necessary conditions are straightforward, it is difficult to obtain local or global sufficient conditions. (See Hestenes, 1966, Ch. 7.) This issue can be finessed here since it can be shown directly that an unbeatable payoff can be obtained in a limiting sense. It must be that $V(1)/W(1) \leq 1$, since $V(1)$ is the expectation of $W(x)$ which has maximum value 1. Moreover—

Lemma 6.1. *For the problem described above, there is a sequence of f_n , tending to a point mass at 1, such that $V_n(1)/W_n(1) \rightarrow 1$.*

¹²Scale invariance in this expected fitness case holds in the following sense. Consider a reference cdf F and pdf f with support $[0, 1]$, where $f > 0$ on $(0, 1)$. Consider now the scaled cdf $F^\lambda(x) = F(\lambda x)$ and pdf $f^\lambda(x) = \lambda f(\lambda x)$, with support $[0, \frac{1}{\lambda}]$ for $\lambda > 0$. Scaled utility is then $U^\lambda(x) = k\lambda^{2/3} \int_0^x f(\lambda x')^{2/3} dx'$ so that $U^\lambda(\frac{y}{\lambda}) = \frac{k}{\lambda^{1/3}} \int_0^y f(y')^{2/3} dy'$ where $k = \frac{\lambda^{1/3}}{\int_0^1 f^{2/3}(y') dy'}$ so that $U^\lambda(\frac{1}{\lambda}) = 1$. It follows immediately that $U^\lambda(\frac{y}{\lambda}) = U(y)$ so that $\int_0^{\frac{1}{\lambda}} f^\lambda(x) U^\lambda(x) dx = \int_0^1 f(y) U(y) dy$.

Proof. See Appendix.

Surveys of happiness vary to a remarkably small extent with such obvious factors as income. Although Denmark is much richer than Bhutan, for example, it is only somewhat happier. This approximate constancy may reflect adaptation. It would be interesting to explain the variation that is left in terms of non-scaleable effects of the distribution.

The discussion of happiness economics largely involves the static properties of the model. The last application reconsiders its new explicitly dynamic properties.

7. BEHAVIORAL CONTRAST

There are a number of puzzling phenomena in economics that relate to the psychological phenomenon of *behavioral contrast*.¹³ For example, Vestergaard and Schultz (2015) consider bidding behavior in second-price auctions. Lesser-valued options involving the growth of value were shown to be systematically preferred to better options with declining time profiles. Similarly, Khaw *et al* (2017) show that current observed “BDM” bids in an auction are higher if the average value of recent items is lower. Lesser-valued options arising in a context involving the growth of value were shown to be systematically preferred to better options with declining time profiles.

In a labor market context, Dustmann and Meghir (2005) provide related evidence of the returns to tenure. These returns seem to represent a direct market response to a preference by workers for rising wages. Skilled workers enjoy a substantial initial wage gain with experience, gains from staying with the same firm, and some gain from staying in the same sector. It is particularly relevant that unskilled workers’ wages grow noticeably with firm tenure, but grow with experience only for a short period, and not at all with sector tenure.

The current model generates results consistent with such observations. Consider the probability of error case, for simplicity. Suppose rewards are from an arbitrary

¹³The classic experiments of Crespi (1942) concerned rats. Rats ran faster towards larger rewards than towards smaller ones. However, rats trained on small rewards ran still faster towards large rewards than did rats exposed to large rewards all along.

cdf F that shifts to the right with time in a linear fashion, at rate α . If the adaptive rule shifts thresholds a distance $\varepsilon > 0$, as in Definition 3.1, then the cdf at time t is $F(x, t) = F(x - \varepsilon\alpha t)$. (The shift in the cdf's is $\varepsilon\alpha$ because otherwise the rule of thumb, shifting thresholds by ε , could not keep up with the improving distributions.) The grid containing the x_n thresholds is now $\mathcal{G}_\varepsilon = \{0, \varepsilon, 2\varepsilon, \dots, m\varepsilon, \dots\}$.

Consider then the random variables $\hat{x}_n = x_n - \varepsilon\alpha t, n = 1, \dots, N$. In the long run these are governed by an invariant distribution, say $\hat{\pi}_{\varepsilon\alpha}$. Theorem 3.1 can readily be adapted to show—

Lemma 7.1. *Suppose $\alpha < \frac{2}{N(N+1)}$ in the modified model with growth of rewards.¹⁴ Then, in the limit as $\varepsilon \rightarrow 0$, the invariant joint distribution of the modified thresholds, $\hat{\pi}_{\varepsilon\alpha}$, converges to a point mass at the vector $(\hat{x}_1^*, \dots, \hat{x}_N^*)$, where the components are characterized by*

$$(7.1) \quad (F(\hat{x}_{n+1}^*) - F(\hat{x}_n^*)) = (F(\hat{x}_n^*) - F(\hat{x}_{n-1}^*)) + \alpha, n = 1, \dots, N.$$

Proof. See Appendix.

Thus, if $\alpha > 0$, the long run effect of such increase in rewards is to shift all thresholds down, with each threshold being shifted down relative to the one before. Hence average utility is unambiguously shifted up, with this effect being more pronounced for larger α . If $\alpha < 0$, the effect is reversed, with utility being shifted down. It follows that a rising sequence of rewards may generate higher average utility than a decreasing sequence, even if the decreasing sequence is unambiguously better.

¹⁴When the condition $\alpha < \frac{2}{N(N+1)}$ is not satisfied there will not be an interior solution for equation (7.1).

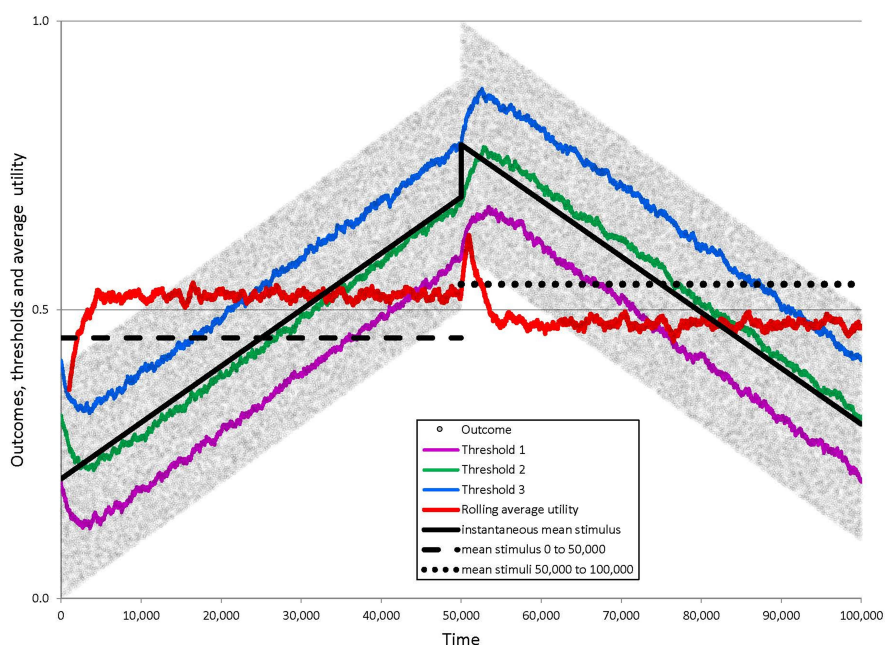


FIGURE 4. Behavioral Adaptation

This analysis is illustrated in Figure 4, where the cdf is the uniform distribution, and there are three thresholds. This cdf shifts up in a linear fashion over the first 50,000 periods. During this phase, the thresholds lag behind the quartiles of the current cdf, as the above argument implies. This raises average utility.

At period 50,000, there is an abrupt jump up in the cdf, which creates a spike in average utility. After period 50,000, the cdf declines linearly, at the same rate that it rose before, the thresholds are above the quartiles, and average utility falls, despite the higher average rewards in this second phase.

Consider then a choice between rising consumption—as in the first phase in Figure 4—and falling consumption—as in the second phase. After a temporal reversal, the rewards in the second phase vector dominate those in the first, but, except for the immediate aftermath of the jump in rewards at the midpoint, which is an artifact of the precise sequence here, utility levels in the first phase similarly vector

dominate those in the second. This provides a basis for predicting a preference for the first phase.¹⁵

8. CONCLUSIONS

We present a simple model where utility shifts in real time in response to changing circumstances. This adaptation acts to reduce the error caused by a limited ability to make fine distinctions, and is ultimately evolutionarily optimal. This model sheds light on the hedonic treadmill, and happiness economics and behavioral adaptation.

There is no deep conflict with economics, in the sense that limited discrimination is the only reason there are mistakes at all; as this ability improves, behavior converges to that implied by standard economics.

9. APPENDIX: PROOFS

9.1. Proof of Proposition 3.1. Make the innocuous simplification that the cdf F is uniform on $[0, 1]$. The optimal placement of the thresholds is then $(1/3, 2/3)$. In the limit as $\varepsilon \rightarrow 0$, but where the arrival rate of outcomes is sped up in proportion to $1/\varepsilon$, it can be shown that the thresholds adapt in a way described by two linear ordinary differential equations.

In general, the differential equations that hold in the limit are

$$\dot{x}_1 = \Delta_1^1 x_1 + \Delta_2^1 (x_2 - x_1) + \Delta_3^1 (1 - x_2),$$

$$\dot{x}_2 = \Delta_1^2 x_1 + \Delta_2^2 (x_2 - x_1) + \Delta_3^2 (1 - x_2),$$

where $\Delta_i^n \in \{0, 1, -1\}$ represents how an outcome in $I_i, i = 1, 2, 3$, moves the threshold $x_n, n = 1, 2$. For $(1/3, 2/3)$ to be an equilibrium of this dynamic, it must be that $\sum_{i=1}^3 \Delta_i^n = 0, n = 1, 2$, so that nontrivial Δ_i^n are permutations of $\{0, 1, -1\}$ for $n = 1, 2$. Symmetry means that the process can be reversed from left to right,

¹⁵This explanation obviates the need to invoke “recency bias”, which could also predict a preference for the first phase over the second.

so that¹⁶

$$\Delta_1^1 = -\Delta_3^2, \Delta_2^1 = -\Delta_2^2, \text{ and } \Delta_3^1 = -\Delta_1^2.$$

The current rule is that

$$\Delta_1^1 = -\Delta_3^2 = -1, \Delta_2^1 = -\Delta_2^2 = 1, \text{ and } \Delta_3^1 = -\Delta_1^2 = 0$$

and the limiting differential equations for x_1 and x_2 are then

$$\dot{x}_1 = (x_2 - x_1) - x_1 = x_2 - 2x_1 \text{ and } \dot{x}_2 = (1 - x_2) - (x_2 - x_1) = 1 - 2x_2 + x_1.$$

This linear system is globally asymptotically stable with $(x_1, x_2) \rightarrow (1/3, 2/3)$

because the matrix $\begin{bmatrix} -2 & 1 \\ 1 & -2 \end{bmatrix}$ has two strictly negative eigenvalues.

One alternative rule would simply reverse the current rule. Not surprisingly, this rule is unstable because the relevant matrix $\begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}$ has two strictly positive eigenvalues. The Markov process cannot satisfy Theorem 3.1.

The remaining rules can be sorted by choice of $i = 1, 2, 3$ such that $\Delta_i^1 = 0$. One possibility is

$$\Delta_1^1 = -\Delta_3^2 = -1, \Delta_2^1 = -\Delta_2^2 = 0, \text{ and } \Delta_3^1 = -\Delta_1^2 = 1.$$

The problem with such a rule is evident—it pays no attention to outcomes in I_2 which means it cannot adapt appropriately to the probability of I_2 . The differential equations become

$$\dot{x}_1 = -x_1 + (1 - x_2) = \dot{x}_2,$$

so that $x_1 - x_2$ does not change. The relevant matrix $\begin{bmatrix} -1 & -1 \\ -1 & -1 \end{bmatrix}$ has one negative eigenvalue and one that is 0. The solution to the differential equations tends to a limit $(x_1^*, 1 - x_1^*)$, dependent on the initial values, for $x_1^* \in [0, 1]$. This limit is generally distinct from $(1/3, 2/3)$. The invariant distribution of the Markov process tends to concentrate at such limits but displays no tendency to move to $(1/3, 2/3)$ in particular.

¹⁶Symmetry is for the sake of exposition. It can be shown that asymmetry does not generate additional viable rules.

Another option would reverse the signs of the Δ_i^n . It still follows that $\dot{x}_1 = \dot{x}_2$. The matrix $\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$ now has one positive eigenvalue and one that is 0, so that points of the form $(x_1^*, 1 - x_1^*)$ are not stable limits of the differential equations. The Markov process does not satisfy Theorem 3.1.

The final pair of possibilities involve $\Delta_1^1 = 0$. The first is

$$\Delta_1^1 = -\Delta_3^2 = 0, \Delta_2^1 = -\Delta_2^2 = -1, \text{ and } \Delta_3^1 = -\Delta_1^2 = 1.$$

The differential equations are then

$$\dot{x}_1 = (1 - x_2) - (x_2 - x_1) = 1 + x_1 - 2x_2 \text{ and } \dot{x}_2 = (x_2 - x_1) - x_1 = x_2 - 2x_1.$$

This system is not then globally asymptotically stable because the matrix $\begin{bmatrix} 1 & -2 \\ -2 & 1 \end{bmatrix}$ has one strictly positive eigenvalue and one strictly negative. The only remaining possibility reverses the previous rule. The resulting differential equation system is not asymptotically stable, because the relevant matrix $\begin{bmatrix} -1 & 2 \\ 2 & -1 \end{bmatrix}$ again has one strictly positive eigenvalue and one strictly negative. The Markov process for either of these two rules cannot then satisfy Theorem 3.1.

Within the current simple class of Markov rules, with two adjustable thresholds, the current rule is then the only one that generates optimal choice in the long run, as claimed.

9.2. Proof of Theorem 4.1. Fix N , the number of interior thresholds. Consider a sequence of distributions $\{F_\varepsilon\}$ converging to F , as ε tends to zero, where, furthermore, for each $\varepsilon > 0$, the equilibrium interior thresholds \mathbf{x}_ε^* corresponding to F_ε (as described in Theorem 4.1) all lie on the grid $\{\mathcal{G}_\varepsilon\}^N$. The vector $\mathbf{x}_\varepsilon(t)$ describes the placement of the thresholds when the distribution over outcomes is F_ε and the placement of the cutoffs evolves according to the rule of thumb from Definition 4.1.

Given the rule of thumb, to each cdf F_ε there correspond transition probabilities, $P_\varepsilon\{\mathbf{x}_\varepsilon(\tau) \mid \mathbf{x}_\varepsilon(t)\}$ —i.e., the probabilities that the state transitions from $\mathbf{x}_\varepsilon(\tau)$ in period τ to $\mathbf{x}_\varepsilon(t)$ in period $t > \tau$. Since the Markov chain is irreducible, these one

step transition probabilities generate π_ε , the unique invariant distribution of the thresholds under F_ε . In the following, E_ε denotes the expectation with respect to P_ε .

Suppose π^* is the distribution that assigns point mass to the vector of thresholds $\mathbf{x}^* = (x_0^*, \dots, x_{N+1}^*)$, where these are as described in Theorem 4.1, for F . We prove now that π_ε converges to π^* as ε tends to zero, proving Theorem 4.1. Consider first the following result.

Lemma 9.1. *Given F_ε as described above define the “Manhattan” distance of $\mathbf{x}_\varepsilon(t)$ from \mathbf{x}_ε^* —*

$$\Delta_t = \sum_{n=1}^N |x_n^t - x_{\varepsilon n}^*|.$$

Then, for each t and $\tau \geq t$,

$$(9.1) \quad E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t)) - \Delta_t \leq \varepsilon \cdot N.$$

Proof of Lemma 9.1. Fix F_ε throughout. In the following, we suppress the ε subscript on $x_{\varepsilon n}^*$. Consider threshold configurations in $\{\mathcal{G}_\varepsilon\}^{N+2}$ such that at least one interior threshold x_n , $n = 1, \dots, N$, is at its equilibrium location. The set of these is $X^* \subseteq \{\mathcal{G}_\varepsilon\}^{N+2}$, i.e.,

$$X^* = \{\mathbf{x} \in \{\mathcal{G}_\varepsilon\}^{N+2} : \min_{n=1, \dots, N} |x_n - x_n^*| = 0\}.$$

Then, let $X_r^* \subseteq \{\mathcal{G}_\varepsilon\}^{N+2}$ denote configurations with exactly r thresholds at their equilibrium locations. Write $I_n(t) = [x_{n-1}^t, x_n^t]$, and $I_n^* = [x_{n-1}^*, x_n^*]$, for $n = 1, \dots, N+1$.

Lemma 9.1 follows from Lemmas 9.2, 9.3, and 9.4, which are stated and proved below.

Lemma 9.2. *Suppose $\mathbf{x}_\varepsilon(t) \notin X^*$, then $E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t)) - \Delta_t \leq 0$.*

Proof. Define, for each t , sets of threshold indices—

$$\mathcal{I}(t) \subset \{1, \dots, N+1\} \text{ is such that } n \in \mathcal{I}(t) \text{ if and only if } I_n(t) \subset I_n^*,$$

$$\mathcal{O}(t) \subset \{1, \dots, N+1\} \text{ is such that } n \in \mathcal{O}(t) \text{ if and only if } I_n^* \subset I_n(t),$$

$$\mathcal{N}(t) = \{1, \dots, N+1\} \setminus \{\mathcal{I}(t) \cup \mathcal{O}(t)\}.$$

Suppose, as in the statement of the lemma, that $\mathbf{x}_\varepsilon(t) \notin X^*$. Consider $n \in \mathcal{N}(t)$. Given that $\mathbf{x}_\varepsilon(t) \notin X^*$ it follows that $n \notin \{1, N+1\}$. Therefore, when y_t lands in $I_n(t)$, either x_{n-1}^t or x_n^t , or both, can shift. One shift decreases Δ_t by ε while the other shift increases Δ_t by ε . Since it is equally likely that either x_{n-1}^t or x_n^t shifts, conditional on $y_t \in I_n(t)$, it follows that $n \notin \mathcal{N}(t)$ implies

$$E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t), y_t \in I_n(t)) - \Delta_t = 0.$$

Thus—

$$\begin{aligned} E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t)) - \Delta_t &= \sum_{n \in \mathcal{S}(t)} (E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t), y_t \in I_n(t)) - \Delta_t) - \\ &\quad \sum_{n \in \mathcal{O}(t)} (E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t), y_t \in I_n(t)) - \Delta_t). \end{aligned}$$

Next consider $n \in \mathcal{O}(t)$. If the outcome y_t lands in $I_n(t)$, then any threshold shift can only bring the configuration closer to \mathbf{x}_ε^* . In particular, if $j \in \{1, 2\}$ thresholds are moved, then $\Delta_{t+1} = \Delta_t - j\varepsilon$. It follows that $n \in \mathcal{O}(t)$ implies

$$E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t), y_t \in I_n(t)) - \Delta_t = -\varepsilon(\phi_{n-1} + \phi_n)(x_n^t - x_{n-1}^t)^\beta,$$

where $\phi_m = 0$ if $m \in \{0, N+1\}$, and $\phi_m = 1$ otherwise. Consider the last case, $n \in \mathcal{S}(t)$. When y_t lands in $I_n(t)$, if $j \in \{1, 2\}$ thresholds shift, then $\Delta_{t+1} = \Delta_t + j\varepsilon$. Thus, $n \in \mathcal{S}(t)$ implies

$$E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t), y_t \in I_n(t)) - \Delta_t = \varepsilon(\phi_{n-1} + \phi_n)(x_n^t - x_{n-1}^t)^\beta.$$

The foregoing implies

$$\begin{aligned} E_\varepsilon(\Delta_{t+1} | \mathbf{x}_\varepsilon(t)) - \Delta_t &= \\ (9.2) \quad &\sum_{n \in \mathcal{S}(t)} \varepsilon(\phi_{n-1} + \phi_n)H_n(t) - \sum_{n \in \mathcal{O}(t)} \varepsilon(\phi_{n-1} + \phi_n)H_n(t), \end{aligned}$$

where

$$H_n(t) = (F_\varepsilon(x_n^t) - F_\varepsilon(x_{n-1}^t))(x_n^t - x_{n-1}^t)^\beta.$$

By the definition of the equilibrium configuration for F_ε , if $n \in \mathcal{O}(t)$, and $m \in \mathcal{S}(t)$, then $H_n(t) \geq H_m(t)$. Therefore, in view of Eq (9.2), to complete the proof, it

suffices to show that

$$(9.3) \quad \sum_{n \in \mathcal{S}(t)} (\phi_{n-1} + \phi_n) \leq \sum_{m \in \mathcal{O}(t)} (\phi_{m-1} + \phi_m),$$

whenever $\mathbf{x}_\varepsilon(t) \notin X^*$.

To show that Eq (9.3) is true we first show that in Eq (9.3) each ϕ_n , $n < N + 1$, under the $\mathcal{S}(t)$ sum is offset by a ϕ_{m-1} , $m > n$, under the $\mathcal{O}(t)$ sum. To do this, note that $n \in \mathcal{S}(t)$ implies $x_n^t < x_n^*$. Therefore, if $n < N + 1$, then $n + 1 \in \mathcal{N}(t) \cup \mathcal{O}(t)$. That is, $x_{n+1}^t > x_{n+1}^*$ implies $n + 1 \in \mathcal{O}(t)$, and $x_{n+1}^t < x_{n+1}^*$ implies $n + 1 \in \mathcal{N}(t)$, and that these cases exhaust the possibilities, since $\mathbf{x}_\varepsilon(t) \notin X^*$ implies $x_{n+1}^t \neq x_{n+1}^*$. Now, if $n + 1 \in \mathcal{O}(t)$, then ϕ_n appears under both the $\mathcal{S}(t)$ and $\mathcal{O}(t)$ sums in Eq (9.3), and thus these terms cancel each other in that expression. If instead $n + 1 \in \mathcal{N}(t)$, then similarly $n + 2 \in \mathcal{N}(t) \cup \mathcal{O}(t)$. If $n + 2 \in \mathcal{O}(t)$ then ϕ_{n+1} under the $\mathcal{O}(t)$ sum cancels ϕ_n under the $\mathcal{S}(t)$ sum. Otherwise, proceed by induction to eventually find an $m > n$ with $m \in \mathcal{O}(t)$ and $k \in \mathcal{N}(t)$ for each $k = n + 1, \dots, m - 1$. Note that $N + 1 \notin \mathcal{N}(t)$. This ϕ_m under the $\mathcal{O}(t)$ sum then cancels ϕ_n under the $\mathcal{S}(t)$ sum. In a similar fashion it can be shown that each ϕ_{n-1} , $n > 0$, under the $\mathcal{S}(t)$ sum is offset by a ϕ_m , $m < n$, under the $\mathcal{O}(t)$ sum. Eq (9.3) thus follows. In view of Eq (9.2), this completes the proof of Lemma 9.2.

We will also need the following result to prove Lemma 9.1.

Lemma 9.3. *Suppose $\mathbf{x}_\varepsilon(t) \in X_r^*$. Let σ be the first date, following t , a threshold is shifted away from its equilibrium placement. Then, for each $\tau > t$,*

$$(9.4) \quad E_\varepsilon(\Delta_\tau \mid \mathbf{x}_\varepsilon(t), \{\tau < \sigma\}) - \Delta_t \leq 0.$$

Proof. By the definition of σ there are r thresholds $\tilde{X} = \{x_{n_1}^t, \dots, x_{n_r}^t\}$ where, conditional on $\tau < \sigma$, for each $s = 1, \dots, r$, $x_{n_s}^t = x_{n_s}^*$, in periods t, \dots, τ . The thresholds in \tilde{X} can then be treated as fixed endpoints of their adjacent intervals when computing the expectation in Eq (9.4). The argument used in the proof of Lemma 9.2 then applies for the thresholds inside these fixed intervals.

Lemma 9.1 now follows immediately from the next result.

Lemma 9.4. *Suppose $\mathbf{x}_\varepsilon(t) \in X_r^*$. Then, for each $\tau \geq t$, $E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t)) - \Delta_t \leq r\varepsilon$.*

Proof. Suppose $\mathbf{x}_\varepsilon(t) \in X_r^*$. Consider stopping times $\mathcal{D}_\tau^j = \{\sigma : t \leq \sigma \leq \tau\}$, $j = 1, 2$, where at each date in \mathcal{D}_τ^j exactly j thresholds are shifted onto their equilibrium locations. Similarly, consider stopping times $\mathcal{U}_\tau^j = \{\sigma : t \leq \sigma \leq \tau\}$, $j = 1, 2$, where at each date in \mathcal{U}_τ^j exactly j thresholds are displaced from their equilibrium locations. Write $\mathcal{D}_\tau = \mathcal{D}_\tau^1 \cup \mathcal{D}_\tau^2$, and $\mathcal{U}_\tau = \mathcal{U}_\tau^1 \cup \mathcal{U}_\tau^2$. Now,

$$\begin{aligned} E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t)) - \Delta_t &= E_\varepsilon \left(\sum_{s=t}^{\tau-1} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right) \\ &= E_\varepsilon \left(\sum_{s \in \mathcal{U}_\tau} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right) \\ &\quad + E_\varepsilon \left(\sum_{s \in \mathcal{D}_\tau} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right) \\ &\quad + E_\varepsilon \left(\sum_{s \notin \mathcal{U}_\tau \cup \mathcal{D}_\tau} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right). \end{aligned}$$

Lemma 9.3 implies the sum in the last line is bounded above by zero. Therefore

$$\begin{aligned} E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t)) - \Delta_t &\leq E_\varepsilon \left(\sum_{s \in \mathcal{U}_\tau} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right) \\ &\quad + E_\varepsilon \left(\sum_{s \in \mathcal{D}_\tau} E_\varepsilon(\Delta_{s+1} | \mathbf{x}_\varepsilon(s)) - \Delta_s \mid \mathbf{x}_\varepsilon(t) \right) \\ &= E_\varepsilon \left(\sum_{s \in \mathcal{U}_\tau^1} \varepsilon - \sum_{s \in \mathcal{D}_\tau^1} \varepsilon \mid \mathbf{x}_\varepsilon(t) \right) + 2E_\varepsilon \left(\sum_{s \in \mathcal{U}_\tau^2} \varepsilon - \sum_{s \in \mathcal{D}_\tau^2} \varepsilon \mid \mathbf{x}_\varepsilon(t) \right). \end{aligned}$$

The lemma therefore follows if

$$E_\varepsilon(|\mathcal{U}_\tau^1| + 2|\mathcal{U}_\tau^2| - |\mathcal{D}_\tau^1| - 2|\mathcal{D}_\tau^2| | \mathbf{x}_\varepsilon(t)) \leq r.$$

Suppose $|\mathcal{U}_\tau^1| + 2|\mathcal{U}_\tau^2| = r + k$, where $k > 0$. Since $\mathbf{x}_\varepsilon(t) \in X_r^*$, by assumption, it follows that for each of the additional k shifts where a threshold was displaced from its equilibrium location, there must have been an offsetting earlier shift where a threshold landed on its equilibrium location. In particular, we must then have $|\mathcal{D}_\tau^1| + 2|\mathcal{D}_\tau^2| \geq k$. This completes the proof of the lemma.

Lemma 9.1 follows immediately from Lemmas 9.2 and 9.4. We can now prove the following, which establishes that π_ε converges to π^* as ε tends to zero, hence proving Theorem 9.1.

Lemma 9.5. *Consider the process $\mathbf{x}_\varepsilon(t)$. Suppose Δ_t is the distance of $\mathbf{x}_\varepsilon(t)$ from \mathbf{x}_ε^* . Then, there is a date t such that for all $\tau \geq t$,*

$$P_\varepsilon\{\Delta_\tau \geq \sqrt{\varepsilon N}\} \leq 2\sqrt{\varepsilon N}.$$

Proof. Lemma 9.1 implies

$$\sup_{\tau \geq t} E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t)) \leq \Delta_t + \varepsilon N.$$

Hence,

$$\sup_{\tau \geq t} E_\varepsilon(\Delta_\tau | \mathbf{x}_\varepsilon(t), \Delta_t \leq \varepsilon N) \leq 2\varepsilon N.$$

Markov's inequality then gives

$$\sup_{\tau \geq t} P_\varepsilon\{\Delta_\tau \geq \sqrt{\varepsilon N} | \Delta_t \leq \varepsilon N\} \leq 2\sqrt{\varepsilon N}.$$

The Markov chain is irreducible. Hence, with probability one, the event $\Delta_t \leq \varepsilon N$ happens infinitely often, and thus

$$\limsup_{\tau} P_\varepsilon\{\Delta_\tau \geq \sqrt{\varepsilon N}\} \leq 2\sqrt{\varepsilon N}$$

as claimed.

9.3. Proof of Theorem 4.2. Fix N throughout. Recall that $(x_0^*, \dots, x_{N+1}^*)$ describes the equilibrium placement of thresholds, in the limit as $\varepsilon \rightarrow 0$, for the rule of thumb, adapted to the cdf F , when there are N interior thresholds. Each N then defines λ_N , where

$$\lambda_N = (x_n^* - x_{n-1}^*)^\beta \cdot (F(x_n^*) - F(x_{n-1}^*)), \quad n = 1, \dots, N + 1.$$

Define

$$\alpha_n = \frac{F(x_n^*) - F(x_{n-1}^*)}{x_n^* - x_{n-1}^*}.$$

Then

$$(9.5) \quad \lambda_N = (x_n^* - x_{n-1}^*)^{1+\beta} \cdot \alpha_n, \quad n = 1, \dots, N + 1.$$

Raising each side to the power of $\frac{1}{1+\beta}$ and then summing over $n = 1, \dots, N + 1$ gives

$$(9.6) \quad \lambda_N^{\frac{1}{1+\beta}} = \frac{1}{N+1} \sum_{n=1}^{N+1} \alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*).$$

We will use, moreover, the following, which is implied by Eqs (9.5) and (9.6)—

$$(9.7) \quad \begin{aligned} x_n^* - x_{n-1}^* &= \left(\frac{\lambda_N}{\alpha_n} \right)^{\frac{1}{1+\beta}} \\ &= \left(\frac{1}{\alpha_n} \right)^{\frac{1}{1+\beta}} \frac{1}{N+1} \sum_{m=1}^{N+1} \alpha_m^{\frac{1}{1+\beta}} (x_m^* - x_{m-1}^*), n = 1, \dots, N + 1. \end{aligned}$$

Recall that $L(N)$ is the expected loss, relative to the full information ideal, when the thresholds are placed according to $(x_0^*, \dots, x_{N+1}^*)$. In particular,

$$L(N) = \sum_{n=1}^{N+1} L_n(N) (F(x_n^*) - F(x_{n-1}^*))^2,$$

where $L_n(N)$ is the expected loss, relative to the full information ideal, conditional on the outcomes on both arms arriving in $[x_{n-1}^*, x_n^*]$. More precisely,

$$L_n(N) = \left(\frac{1}{F(x_n^*) - F(x_{n-1}^*)} \right)^2 \times \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} \frac{\max\{y, y'\} - \min\{y, y'\}}{2} f(y) f(y') dy dy'.$$

since conditional on y, y' landing in $[x_{n-1}^*, x_n^*]$ the payoff is $\max\{y, y'\}$ under the full information ideal, while under limited information the payoff is $\max\{y, y'\}$ with probability $1/2$, and $\min\{y, y'\}$ with probability $1/2$. Consider now an approximation of $L(N)$, where the expected loss relative to the full information ideal is integrated against a step function approximation of f . The steps in the approximation occur at the placement of the equilibrium cutoffs for F , given that there are N cutoffs.

Definition 9.1. For each N consider the step function f_N , where $f_N(1) = f(1)$, and

$$f_N(y) = \frac{F(x_n^*) - F(x_{n-1}^*)}{x_n^* - x_{n-1}^*} \text{ for all } y \in [x_{n-1}^*, x_n^*), n = 1, \dots, N + 1.$$

Then define

$$\hat{L}_n(N) = \left(\frac{1}{F(x_n^*) - F(x_{n-1}^*)} \right)^2 \times \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} \frac{\max\{y, y'\} - \min\{y, y'\}}{2} f_N(y) f_N(y') dy dy',$$

and

$$\hat{L}(N) = \sum_{n=1}^{N+1} \hat{L}_n(N) (F(x_n^*) - F(x_{n-1}^*))^2.$$

We will use Lemmas and 9.6, 9.7, and 9.8 to characterize the limiting efficiency of the rule of thumb as in Theorem 4.2.

Lemma 9.6. *Since f is uniformly continuous, and bounded away from zero, f_N converges uniformly to f as $N \rightarrow \infty$.*

Proof. Since f is uniformly continuous, it follows that for each $\eta > 0$ there is a $\delta > 0$ such that $|y - y'| < \delta$ implies $|f(y) - f(y')| < \eta$. The mean value theorem implies that for each $n = 1, \dots, N + 1$ there is a $y_n \in [x_{n-1}^*, x_n^*]$ such that

$$f(y_n) = \frac{F(x_n^*) - F(x_{n-1}^*)}{x_n^* - x_{n-1}^*}.$$

Thus

$$\max_{y \in [x_{n-1}^*, x_n^*]} \left| f(y) - \frac{F(x_n^*) - F(x_{n-1}^*)}{x_n^* - x_{n-1}^*} \right| < \eta, n = 1, \dots, N + 1,$$

whenever $|x_n^* - x_{n-1}^*| < \delta$. In order to establish the uniform convergence of f_N to f we then need only show that $\max_{n=1, \dots, N} |x_n^* - x_{n-1}^*|$ converges to zero as N tends to infinity. Recall, given N , the constant

$$\lambda_N = (x_n^* - x_{n-1}^*)^\beta \cdot (F(x_n^*) - F(x_{n-1}^*)), \quad n = 1, \dots, N + 1.$$

We thus have

$$x_n^* - x_{n-1}^* = \left(\frac{\lambda_N}{f(y_n)} \right)^{\frac{1}{1+\beta}}.$$

The pdf, f , is bounded away from zero, by assumption. Thus

$$x_n^* - x_{n-1}^* \leq c \lambda_N^{\frac{1}{1+\beta}}, \quad n = 1, \dots, N + 1,$$

for some constant, c , that does not depend on $n = 1, \dots, N + 1$, or N . From Eq (9.6), it follows that $\max_{n=1, \dots, N} |x_n^* - x_{n-1}^*|$ converges to zero as N tends to infinity. This completes the proof of the lemma.

Lemma 9.7. *Since f is uniformly continuous, and bounded away from zero, $(N + 1)^2 \times |L(N) - \hat{L}(N)| \rightarrow 0$, as $N \rightarrow \infty$.*

Proof. We begin with

$$\begin{aligned} |L(N) - \hat{L}(N)| &= \left| \sum_{n=1}^{N+1} (L_n(N) - \hat{L}_n(N)) (F(x_n^*) - F(x_{n-1}^*))^2 \right| \\ &\leq \sum_{n=1}^{N+1} \left| (L_n(N) - \hat{L}_n(N)) (F(x_n^*) - F(x_{n-1}^*))^2 \right|. \end{aligned}$$

Now

$$\begin{aligned} (L_n(N) - \hat{L}_n(N)) (F(x_n^*) - F(x_{n-1}^*))^2 &= \\ \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} \frac{\max\{y, y'\} - \min\{y, y'\}}{2} (f(y)f(y') - f_N(y)f_N(y')) dy dy'. \end{aligned}$$

Therefore,

(9.8)

$$\begin{aligned} |(L_n(N) - \hat{L}_n(N)) (F(x_n^*) - F(x_{n-1}^*))^2| &\leq \\ \frac{1}{2} \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} |\max\{y, y'\} - \min\{y, y'\}| \cdot |f(y)f(y') - f_N(y)f_N(y')| dy dy' &\leq \\ \frac{1}{2} \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} (x_n^* - x_{n-1}^*) |f(y)f(y') - f_N(y)f_N(y')| dy dy' &\leq \\ \frac{1}{2} \left(\max_{y, y' \in [x_{n-1}^*, x_n^*]} |f(y)f(y') - f_N(y)f_N(y')| \right) \int_{x_{n-1}^*}^{x_n^*} \int_{x_{n-1}^*}^{x_n^*} (x_n^* - x_{n-1}^*) dy dy' &= \\ \frac{1}{2} \left(\max_{y, y' \in [x_{n-1}^*, x_n^*]} |f(y)f(y') - f_N(y)f_N(y')| \right) (x_n^* - x_{n-1}^*)^3 &\leq \psi_N \cdot (x_n^* - x_{n-1}^*)^3, \end{aligned}$$

where

$$\psi_N = \max_n \left\{ \max_{y, y' \in [x_{n-1}^*, x_n^*]} |f(y)f(y') - f_N(y)f_N(y')| \right\}.$$

The foregoing calculation shows that

$$|L(N) - \hat{L}(N)| \leq \psi_N \sum_{n=1}^{N+1} (x_n^* - x_{n-1}^*)^3.$$

The uniform convergence of f_N to f (Lemma 9.6) implies ψ_N can be made arbitrarily small by choice of N . In order to complete the proof it therefore suffices to show that

$$(N+1)^2 \sum_{n=1}^{N+1} (x_n^* - x_{n-1}^*)^3$$

is finite in the limit. Eq (9.7) implies

$$(9.9) \quad (x_n^* - x_{n-1}^*)^2 = \left(\frac{1}{\alpha_n}\right)^{\frac{2}{1+\beta}} \left[\frac{1}{N+1} \sum_{m=1}^{N+1} \alpha_m^{\frac{1}{1+\beta}} (x_m^* - x_{m-1}^*) \right]^2.$$

The term in the square brackets does not depend on n . Therefore,

$$\begin{aligned} \sum_{n=1}^{N+1} (x_n^* - x_{n-1}^*)^3 &= \left(\sum_{n=1}^{N+1} \left(\frac{1}{\alpha_n}\right)^{\frac{2}{1+\beta}} (x_n^* - x_{n-1}^*) \right) \times \\ &\quad \left(\frac{1}{N+1} \sum_{n=1}^{N+1} \alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*) \right)^2 \\ &= \left(\frac{1}{N+1}\right)^2 \left(\int \left(\frac{1}{f_N(y)}\right)^{\frac{2}{1+\beta}} dy \right) \left(\int f_N(y)^{\frac{1}{1+\beta}} dy \right)^2. \end{aligned}$$

The assumption that f bounded away from zero implies $1/(f(y))^{\frac{2}{1+\beta}}$ is integrable. Since f_N converges to f uniformly, and since the domain of integration is finite, there is clearly some integrable g such that $g \geq f_N$ for each N . The same is true for negative powers of f . The dominated convergence theorem thus implies

$$\begin{aligned} \int \left(\frac{1}{f_N(y)}\right)^{\frac{2}{1+\beta}} dy &\longrightarrow \int \left(\frac{1}{f(y)}\right)^{\frac{2}{1+\beta}} dy, \text{ and} \\ \int f_N(y)^{\frac{1}{1+\beta}} dy &\longrightarrow \int f(y)^{\frac{1}{1+\beta}} dy, \end{aligned}$$

as $N \rightarrow \infty$. Hence,

$$(N+1)^2 \sum_{n=1}^{N+1} (x_n^* - x_{n-1}^*)^3 \longrightarrow \left(\int \left(\frac{1}{f(y)}\right)^{\frac{2}{1+\beta}} dy \right) \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2.$$

This completes the proof of the lemma.

Next we have—

Lemma 9.8. For each $n = 1, \dots, N+1$,

$$\hat{L}_n(N) = \frac{x_n^* - x_{n-1}^*}{6}.$$

Proof. $\hat{L}_n(N)$ is equal to the loss, relative to the full information ideal, from choosing at random when outcomes are uniformly distributed on $[x_{n-1}^*, x_n^*]$. Thus consider a uniform distribution with pdf $1/s$ on the interval $[0, s]$. The loss from choosing at random relative to the full information ideal is $s/6$. Specifically, the expected payoff from choosing randomly between the two arms is clearly $s/2$. The expected payoff from choosing the higher of the two arms, on the other hand, as would be the full information ideal, is $2s/3$. To see this, suppose

$$K(y) = \Pr\{\max\{y', y''\} < y\} = \Pr\{y' \& y'' < y\} = (y/s)^2.$$

Hence $\int_0^s y dK(y) = 2s/3$. It follows that the expected loss from choosing at random is $s/6$, proving Lemma 9.8.

Theorem 4.2 will follow from Lemmas 9.9, 9.10, and 9.11, which are given next. The next result characterizes the limiting efficiency of the rule of thumb.

Lemma 9.9. *Since f is uniformly continuous, and bounded away from zero, the limiting efficiency of the rule of thumb is characterized by*

$$(N+1)^2 L(N) \longrightarrow \frac{1}{6} \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right),$$

as N tends to infinity.

Proof. In view of Lemma 9.7, it suffices to show that, for the approximation, $\hat{L}(N)$,

$$(N+1)^2 \hat{L}(N) \longrightarrow \frac{1}{6} \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right).$$

Lemma 9.8 gives

$$(9.10) \quad \hat{L}(N) = \frac{1}{6} \sum_{n=1}^{N+1} (x_n^* - x_{n-1}^*)^3 \alpha_n^2.$$

Equation 9.7 gives

$$(x_n^* - x_{n-1}^*)^2 = \left(\frac{1}{\alpha_n} \right)^{\frac{2}{1+\beta}} \left(\frac{1}{N+1} \sum_{m=1}^{N+1} \alpha_m^{\frac{1}{1+\beta}} (x_m^* - x_{m-1}^*) \right)^2.$$

Therefore,

$$\begin{aligned}\hat{L}(N) &= \frac{1}{6} \left(\frac{1}{N+1} \right)^2 \left(\sum_{n=1}^{N+1} \alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*) \right)^2 \left(\sum_{n=1}^{N+1} \alpha_n^{\frac{2\beta}{1+\beta}} (x_n^* - x_{n-1}^*) \right) \\ &= \frac{1}{6} \left(\frac{1}{N+1} \right)^2 \left(\int f_N(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f_N(y)^{\frac{2\beta}{1+\beta}} dy \right).\end{aligned}$$

Recall that f_N converges uniformly to f . Clearly then

$$(N+1)^2 \hat{L}(N) \longrightarrow \frac{1}{6} \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right).$$

This completes the proof of Lemma 9.9.

Lemma 9.10. *In the limit as N tends to infinity the distribution of thresholds under the adaptive rule of thumb is*

$$(9.11) \quad U(x) = \frac{\int_0^x f(y)^{\frac{1}{1+\beta}} dy}{\int_0^1 f(y)^{\frac{1}{1+\beta}} dy}.$$

That is, in the limit under the rule of thumb the fraction of thresholds in the interval $[0, x]$ is $U(x)$.

Proof. Equation 9.7 implies

$$\alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*) = \frac{1}{N+1} \sum_{m=1}^{N+1} \alpha_m^{\frac{1}{1+\beta}} (x_m^* - x_{m-1}^*), \quad n = 1, \dots, N+1.$$

Therefore, for each $m = 1, \dots, N+1$ we have

$$\frac{\sum_{n=1}^m \alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*)}{\sum_{n=1}^{N+1} \alpha_n^{\frac{1}{1+\beta}} (x_n^* - x_{n-1}^*)} = \frac{m}{N+1},$$

which is the fraction of thresholds in $[0, x_m^*]$. This is just

$$\frac{\int_0^{x_m^*} f_N(y)^{\frac{1}{1+\beta}} dy}{\int_0^1 f_N(y)^{\frac{1}{1+\beta}} dy} = \frac{m}{N+1}.$$

Define $U_N(x)$ such that $U_N(1) = 1$, and for each $x \in [x_{m-1}^*, x_m^*)$, $m = 1, \dots, N + 1$,

$$(9.12) \quad U_N(x) = \frac{\int_0^{x_m^*} f_N(y)^{\frac{1}{1+\beta}} dy}{\int_0^1 f_N(y)^{\frac{1}{1+\beta}} dy}.$$

Notice that $U_N(x_m^*) = m/(N + 1)$ for each $m = 1, \dots, N + 1$. Recall that f_N converges to f . The dominated convergence theorem gives that U_N converges pointwise to U , where U is the cdf in the statement of the lemma. This completes the proof.

Lemma 9.11. *The expression*

$$\left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^2 \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right).$$

is minimized uniquely by choice of $\beta = 1/2$.

Proof. This follows from the Hölder Inequality (Royden, 1988, p. 119) since

$$\begin{aligned} \int f(y)^{\frac{2}{3}} dy &= \int f(y)^{\frac{2}{3} \cdot \frac{\beta}{1+\beta}} f(y)^{\frac{2}{3} \cdot \frac{1}{1+\beta}} dy \\ &\leq \left(\int f(y)^{\frac{2\beta}{1+\beta}} dy \right)^{\frac{1}{3}} \left(\int f(y)^{\frac{1}{1+\beta}} dy \right)^{\frac{2}{3}}. \end{aligned}$$

Furthermore, equality can only hold here if $f(y)^{2\beta/(1+\beta)} = f(y)^{1/(1+\beta)}$; that is, only if $\beta = 1/2$.

It follows that, when $\beta = 1/2$,

$$(N + 1)^2 L(N) \longrightarrow \frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy \right)^3.$$

This completes the proof of Theorem 4.2.

9.4. Proof of Theorem 4.3. We now prove that the rule of thumb with $\beta = 1/2$ has the same limiting efficiency as the optimal placement of the N thresholds. Suppose (x_0, \dots, x_{N+1}) describes the optimal placement of thresholds for F given that there are N interior thresholds. Let α_n now be $\frac{F(x_n) - F(x_{n-1})}{x_n - x_{n-1}}$. Let $L^*(N)$ denote the expected loss in y from the optimally placed thresholds relative to the full information ideal. As in the previous section consider an approximation $\bar{L}(N)$ of

$L^*(N)$ where the loss is integrated against $g_N(y)$ where $g_N(1) = f(1)$, and for each $n = 1, \dots, N+1$, $g_N(y) = \alpha_n$, for all $y \in [x_{n-1}, x_n]$. Moreover, let $L_n^*(N)$ denote the expected loss from the optimally placed thresholds, relative to the full information ideal, given that both outcomes land in $[x_{n-1}, x_n]$, and let $\bar{L}_n(N)$ denote the approximation of $L_n^*(N)$ where the loss is integrated against g_N , instead of f .

Lemma 9.8 applies here, and this gives $\bar{L}_n(N) = (x_n - x_{n-1})/6$. Thus

$$\bar{L}(N) = \frac{1}{6} \sum_{n=1}^{N+1} (x_n - x_{n-1})(F(x_n) - F(x_{n-1}))^2,$$

and hence

$$L^*(N) = \frac{1}{6} \sum_{n=1}^{N+1} (x_n - x_{n-1})^3 \alpha_n^2 + L^*(N) - \bar{L}(N).$$

The calculations leading to Eq (9.8) in the proof of Lemma 9.7 can be repeated here to show that

$$L^*(N) - \bar{L}(N) \geq -(\psi_N/6) \sum_{n=1}^{N+1} (x_n - x_{n-1})^3,$$

where

$$(\psi_N/6) = \max_n \left\{ \max_{y, y' \in [x_{n-1}, x_n]} |f(y)f(y') - g_N(y)g_N(y')| \right\}$$

has been redefined for convenience. Thus

$$L^*(N) \geq \frac{1}{6} \sum_{n=1}^{N+1} (x_n - x_{n-1})^3 (\alpha_n^2 - \psi_N).$$

The proof of Lemma 9.6 can be repeated here to show that g_N converges uniformly to f . (In particular, as in the proof of Lemma 9.6, $\max_n |x_n - x_{n-1}| \rightarrow 0$, as $N \rightarrow \infty$, yields the uniform convergence of g_N to f .) Hence, $\psi_N \rightarrow 0$. The term $\alpha_n^2 - \psi_N$ is then positive for each sufficiently large N , since α_n is bounded away from zero, by assumption. The Hölder Inequality then gives, for sufficiently large N ,

$$\begin{aligned} \sum_{n=1}^{N+1} (x_n - x_{n-1}) (\alpha_n^2 - \psi_N)^{\frac{1}{3}} &\leq \left(\sum_{n=1}^{N+1} (x_n - x_{n-1})^3 (\alpha_n^2 - \psi_N) \right)^{\frac{1}{3}} \left(\sum_{n=1}^{N+1} (1) \right)^{\frac{2}{3}} \\ &= \left(\sum_{n=1}^{N+1} (x_n - x_{n-1})^3 (\alpha_n^2 - \psi_N) \right)^{\frac{1}{3}} (N+1)^{\frac{2}{3}}. \end{aligned}$$

Hence,

$$\begin{aligned} \sum_{n=1}^{N+1} (x_n - x_{n-1})^3 (\alpha_n^2 - \psi_N) &\geq \left(\frac{1}{N+1}\right)^2 \left(\sum_{n=1}^{N+1} (x_n - x_{n-1}) (\alpha_n^2 - \psi_N)^{\frac{1}{3}}\right)^3 \\ &= \left(\frac{1}{N+1}\right)^2 \left(\int (g_N(y)^2 - \psi_N)^{\frac{1}{3}} dy\right)^3. \end{aligned}$$

It then follows that

$$(9.13) \quad L^*(N) \geq \frac{1}{6} \left(\frac{1}{N+1}\right)^2 \left(\int (g_N(y)^2 - \psi_N)^{\frac{1}{3}} dy\right)^3.$$

Since $g_N(y)^2 - \psi_N \leq \bar{f}^2$ converges pointwise to $f(y)^2$,

$$\liminf\{(N+1)^2 L^*(N)\} \geq \frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy\right)^3.$$

The rule of thumb with $\beta = 1/2$ has limiting efficiency that is exactly equal to $\frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy\right)^3$. It then follows from the assumed optimality of the configurations (x_0, \dots, x_{N+1}) that

$$\lim\{(N+1)^2 L^*(N)\} = \frac{1}{6} \left(\int f(y)^{\frac{2}{3}} dy\right)^3,$$

as well, since otherwise the rule of thumb with $\beta = 1/2$ would be better. This completes the proof of Theorem 4.3.

9.5. Proof of Lemma 6.1. Choose any strictly increasing and continuous $f : [0, 1] \rightarrow (0, \infty)$ where $f(x) \rightarrow \infty$ as $x \rightarrow 1$ and $\int_0^1 f(x) dx = 1 - m$ for $m \in (0, 1)$. Define $h(\delta)$ as the unique strictly decreasing solution of $\int_0^{1-\delta} f(x) dx + h(\delta)\delta = 1$, and consider the pdf given by f on $[0, 1 - \delta)$ and the constant $h(\delta)$ on $[1 - \delta, 1]$. It follows that $W(x) = W(1 - \delta) + h^{2/3}(x - 1 + \delta)$ if $x > 1 - \delta$ so that $V(1) = \int_0^{1-\delta} W(x)f(x)dx + h \int_{1-\delta}^1 W(x)dx$. After some algebra, it follows that $V(1) = V(1 - \delta) + h\delta W(1 - \delta) + \frac{h^{2/3}\delta^2}{2}$. As $\delta \rightarrow 0$, have $h\delta \rightarrow m$ so that $h^{2/3}\delta \rightarrow 0$. It follows that $W(1 - \delta) \rightarrow W(1)$. Hence $\frac{V(1)}{W(1)} \geq h\delta \frac{W(1-\delta)}{W(1)} \rightarrow m$, as $\delta \rightarrow 0$. For any sequence of $f_n \rightarrow 0$ such that $m_n \rightarrow 1$, it follows that the associated V_n and W_n satisfy

$$\frac{V_n(1)}{W_n(1)} \rightarrow 1,$$

since $\frac{V_n(1)}{W_n(1)} \leq 1$.¹⁷

Proof of Lemma 7.1.

As in Section 9.2 consider distributions F_ε converging to F as ε tends to zero. For each ε , let $\mathbf{x}_\varepsilon^* = (x_{\varepsilon 1}^*, \dots, x_{\varepsilon N}^*)$ denote the equilibrium vector of the modified thresholds from Lemma 7.1, but for the cdf F_ε . Consider the process $\hat{\mathbf{x}}_\varepsilon(t) = \mathbf{x}_\varepsilon(t) - \varepsilon \cdot \alpha \cdot t$, where $\mathbf{x}_\varepsilon(t)$ describes the placement of thresholds evolving according to the basic rule of thumb, given F_ε .

We give a direct proof of the result for the case $N = 1$. For general N the proof is more involved but proceeds along similar lines. To ease the notation we drop the $n = 1$ and ε subscripts in the remainder when referring to the one threshold, x^t , and the modified threshold, \hat{x}^t . Notice first that \hat{x}^t must visit the interval $[0, 1]$ infinitely often in the limit. This follows because $\alpha < 1$ and hence while the threshold x^t lies outside the support of the improving distribution of outcomes, it is drawn toward the support at a rate greater than $\varepsilon\alpha$. In the remainder then consider a subsequence $\hat{x}^{tr}, r = 1, 2, \dots$ where each term of the sequence lies in $[0, 1]$. The proof of Lemma 7.1 follows upon showing that Lemma 9.1 holds for this sub-sequence, with a suitable adjustment to the right hand side of equation (9.1) there. (For terms not in the sequence, Lemma 9.1 holds trivially since the modified threshold cannot be drawn farther away from its equilibrium placement when it lies outside $[0, 1]$.) With this in mind, notice that the equilibrium threshold \hat{x}^* in this case is such that

¹⁷The necessary conditions for the problem of Mayer described before Lemma 6.1 are as follows (Hestenes, 1966, Theorem 4.1, p. 315). It is necessary to impose a finite upper bound, \bar{f} , say, on the density f , since otherwise existence may not hold. Define then the Hamiltonian

$$\mathcal{H} = \psi_V W f + \psi_F f + \bar{\psi}_W f^{2/3},$$

where ψ_V, ψ_F and ψ_W are the costate variables corresponding to the state variables V, F and W , respectively. Hence $\psi'_V = \psi'_F = 0$ and $\bar{\psi}'_W = -\frac{\partial \mathcal{H}}{\partial W} = -\psi_V f$, so that $\bar{\psi}_W = \psi_W - \psi_V F$, for a constant ψ_W . The objective is to minimize $-\frac{V(1)}{W(1)}$. It is necessary that \mathcal{H} is minimized over $f \in [0, \bar{f}]$ with transversality conditions $\psi_V(1) = \frac{\psi_0}{W(1)}$ and $\bar{\psi}_W(1) = \psi_W - \psi_V = -\frac{\psi_0 V(1)}{W(1)^2}$, for some constant $\psi_0 \geq 0$. It can be shown (eventually) that these necessary conditions admit a unique closed form solution on $[0, 1]$. This solution involves a strictly increasing $f \in (0, \bar{f})$ initially but then a second and final phase where $f = \bar{f}$. If a well-behaved solution exists, this must be it. As the upper bound $\bar{f} \rightarrow \infty$, it follows that $f \rightarrow 0$ on the initial phase and $\frac{V(1)}{W(1)} \rightarrow 1$. The essential features of this construction are used in the argument given above. But that argument is direct, and finesses issues of existence or sufficiency.

$F_\varepsilon(\hat{x}^*) = 1 - F_\varepsilon(\hat{x}^*) - \alpha$, that is, such that $1 - 2 \cdot F_\varepsilon(\hat{x}^*) = \alpha$. Now, suppose the threshold, \hat{x}^t , is such that $\hat{x}^t \geq \hat{x}^* + \varepsilon$. Recall that, by assumption, $\alpha \leq 1$, then

$$E(|\hat{x}^{t+1} - \hat{x}^*| | \hat{x}^t) = |\hat{x}^t - \hat{x}^*| - \alpha \cdot \varepsilon + \varepsilon \cdot (1 - 2 \cdot F(\hat{x}^t)).$$

It follows that $E(|\hat{x}^{t+1} - \hat{x}^*| | \hat{x}^t) < |\hat{x}^t - \hat{x}^*|$, since $\hat{x}^t > \hat{x}^*$ by assumption, and hence $-\alpha \cdot \varepsilon + \varepsilon \cdot (1 - 2 \cdot F(\hat{x}^t)) < 0$. Now suppose $\hat{x}^t \leq \hat{x}^* - \varepsilon$. Then,

$$E(|\hat{x}^{t+1} - \hat{x}^*| | \hat{x}^t) = |\hat{x}^t - \hat{x}^*| + \alpha \cdot \varepsilon - \varepsilon \cdot (1 - 2 \cdot F(\hat{x}^t)).$$

Again, clearly $E(|\hat{x}^{t+1} - \hat{x}^*| | \hat{x}^t) < |\hat{x}^t - \hat{x}^*|$. Next suppose $\hat{x}^t \in (\hat{x}^* - \varepsilon, \hat{x}^* + \varepsilon)$. In this case, clearly $|\hat{x}^{t+1} - \hat{x}^*| \leq |\hat{x}^t - \hat{x}^*| + \varepsilon \cdot (1 + \alpha)$. We have then

$$E(|\hat{x}^{t+1} - \hat{x}^*| | \hat{x}^t) - |\hat{x}^t - \hat{x}^*| \leq \varepsilon \cdot (1 + \alpha).$$

This completes the proof that Lemma 9.1 holds for \hat{x}^t , when the constant on the right hand side of equation (9.1) is multiplied by $1 + \alpha$. This factor makes no difference when using Lemma 9.1 in Lemma 9.5 to verify the limiting invariant distribution of Theorem 4.1. Hence Lemma 7.1 now follows immediately.

REFERENCES

- [1] Andrew E. Clark, Sarah Flèche, Richard Layard, Nattavudh Powdthavee, and George Ward. *The Origins of Happiness: The Science of Well-Being Over the Life Course*. Princeton University Press, Princeton and Oxford, 2018.
- [2] Leo P. Crespi. Quantitative variation of incentive and performance in the white rat. *The American Journal of Psychology*, 55(4):467–517, 1942.
- [3] Christian Dustmann and Costas Meghir. Wages, experience and seniority. *Review of Economic Studies*, 72:77–108, 2005.
- [4] Alan Grafen. Natural selection, kin selection and group selection. In J.R. Krebs and N.B. Davies, editors, *Behavioural Ecology*, pages 62–84. Blackwell Scientific Publications, Oxford, 1984.
- [5] M. R. Hestenes. *Calculus of Variations and Optimal Control Theory*. Wiley, New York, 1966.
- [6] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.
- [7] Mel W. Khaw, Paul W. Glimcher, and Kenway Louie. Normalized value coding explains dynamic adaptation in the human valuation process. *Proceedings of the National Academy of Sciences of the USA*, 114(48):12696–12701, 2017.
- [8] Simon Laughlin. A simple coding procedure enhances a neuron’s information capacity. *Zeitschrift für Naturforschung C*, 36:910–912, 1981.

- [9] Nick Netzer. Evolution of time preferences and attitudes towards risk. *American Economic Review*, 99(3):937–55, 2009.
- [10] Luis Rayo and Gary Becker. Evolutionary efficiency and happiness. *Journal of Political Economy*, 115(2):302–337, 2007.
- [11] Arthur J. Robson. A biological basis for expected and non-expected utility. *Journal of Economic Theory*, 68(2):397–424, 1996.
- [12] Arthur J. Robson. The biological basis of economic behavior. *Journal of Economic Literature*, 39(1):11–33, 2001.
- [13] Arthur J. Robson and Larry Samuelson. The evolutionary optimality of decision and experienced utility. *Theoretical Economics*, 6:311–339, 2011.
- [14] H. L. Royden. *Real Analysis*. Prentice Hall, Englewood Cliffs, NJ, 1988.
- [15] David A. Schkade and Daniel Kahneman. Does living in California make people happy? A focusing illusion in judgments of life satisfaction. *Psychological Science*, 9(5):340–346, 1998.
- [16] Tali Sharot, Tamara Shiner, Annemarie C. Brown, Judy Fan, and Raymond J. Dolan. Dopamine enhances expectation of pleasure in humans. *Current Biology*, 19(24):2077–2080, 2009.
- [17] William R. Stauffer, Armin Lak, and Wolfram Schultz. Dopamine reward prediction error responses reflect marginal utility. *Current Biology*, 90:927–39, 2014.
- [18] Martin Vestergaard and Wolfram Schultz. Choice mechanisms for past, temporally extended outcomes. *Proceedings of the Royal Society: Series B*, 282:20141766, 2015.