

Understanding the Hows and Whys of Decision-Making: From Expected Utility to Divisive Normalization

PAUL GLIMCHER

*Institute for the Interdisciplinary Study of Decision Making, New York University,
New York, New York 10013*

Correspondence: pg3@nyu.edu

Over the course of the last century, economists and ethologists have built detailed models from first principles of how humans and animals *should* make decisions. Over the course of the last few decades, psychologists and behavioral economists have gathered a wealth of data at variance with the predictions of these economic models. This has led to the development of highly descriptive models that can often predict *what* choices people or animals will make but without offering any insight into *why* people make the choices that they do—especially when those choices reduce a decision-maker's well-being. Over the course of the last two decades, neurobiologists working with economists and psychologists have begun to use our growing understanding of how the nervous system works to develop new models of *how* the nervous system makes decisions. The result, a growing revolution at the interdisciplinary border of neuroscience, psychology, and economics, is a new field called *Neuroeconomics*. Emerging neuroeconomic models stand to revolutionize our understanding of human and animal choice behavior by combining fundamental properties of neurobiological representation with decision-theoretic analyses. In this overview, one class of these models, based on the widely observed neural computation known as *divisive normalization*, is presented in detail. The work demonstrates not only that a discrete class of computation widely observed in the nervous system is fundamentally ubiquitous, but how that computation shapes behaviors ranging from visual perception to financial decision-making. It also offers the hope of reconciling economic analysis of what choices we should make with psychological observations of the choices we actually do make.

The postenlightenment study of human decision-making has its roots in the work of mathematicians and economists like Blaise Pascal, Daniel Bernoulli, Adam Smith, and David Ricardo. These scholars sought to explain regularities in human behavior from first principles. For these early theorists, that meant defining algorithms that maximized individual gain, profit, or well-being and that thus gave rise to observable market level regularities. These scholars took as their focus the study of “why”—why it was that people made the choices they did from a mathematical and theoretical point of view. While revising many of the key assumptions of these early scholars, economists of the early 1900s like Paul Samuelson (Samuelson 1947) and John von Neumann (Neumann and Morgenstern 1944) made huge strides in this direction by defining how decision-making ought to work in humans. They developed theories that fundamentally described how “optimal” decision-makers should behave.

In the 1960s, however, it started to become clear that humans did not always behave in exactly the ways predicted by these theories. Work by economists like Maurice Allais (Allais 1953) and psychologists like Amos Tversky and Daniel Kahneman (Kahneman and Tversky 1979) began to reveal systematic patterns of behavior that violated the predictions of the economic models. These were behaviors that appeared to reduce the well-being of

individuals and even of whole economies. The psychologist Herbert Simon (Simon 1955) responded to this observation by suggesting that the limitations of our cognitive or neurobiological architecture might be the cause of these inefficiencies. The majority of those cataloging these inefficiencies, however, devoted their energies to developing descriptive models. These were models that could predict *what* people would choose in any situation, even if the models offered little or no insight into *why* biological decision-makers produced these apparently self-destructive behaviors.

In the early 2000s, a group of neurobiologists, working with economists and psychologists, sought to reconcile this tension by studying *how* the nervous system actually makes decisions. These interdisciplinary scholars, working with like-minded economists and psychologists, rapidly gave birth to the interdisciplinary field of *Neuroeconomics*. Within Neuroeconomics, a number of models of how the nervous system encodes and processes decision-related information began to emerge in the 2000s and early 2010s. One of these models, the focus of this review, has its roots in the *Efficient Coding Hypothesis* of Horace Barlow (Barlow 1961). Barlow had hypothesized that because both building internal representations of the outside world and performing computations on those representations is metabolically costly, the representational frameworks and computational processes used by organ-

isms should naturally trade off completeness against metabolic cost. He hypothesized that an animal nervous system would be “efficient” in the sense that it would devote energy only to valuable computations and goals, thus minimizing costly redundancies in its internal representations.

In the mid-2000s, my colleagues and I began to see some parallels between these statements and Herbert Simon’s hypothesis that limitations imposed by our cognitive or neurobiological architecture might be the cause of the apparent “inefficiencies” in human choice behavior identified by psychologists. It might be, we hypothesized, that decision-making systems were trading off representational costs at the neurobiological level against benefits at the level of behavior—and that this trade-off might account for the apparent gap between the *whys* of economists and the *whats* of psychologists. If that was true, then it might even be the case that human choice behavior was efficient, once one took into account the neurobiological costs of representation and computation.

EFFICIENT CODING AND THE THEORY OF CORTICAL REPRESENTATION

The theory of efficient coding emerged, in part, from Keffer Hartline’s observation (Ratliff and Hartline 1959) that the activity of sensory neurons in the retina encodes not just the photons falling directly on the *centers* of their receptors but is also influenced by the structure of *surrounding* visual stimuli. Barlow noted that the precise relationship between center and surround described by Hartline would have the effect of creating more neural activity when the center and surround differed and very little activity when center and surround were homogeneous. Drawing on Claude Shannons’ *Theory of Information* (Shannon and Weaver 1949), Barlow proposed that to be efficient, neurons must devote action potentials to sensory stimuli in proportion to the degree to which those stimuli are *surprising*, and he proposed that it was the differences, rather than the homogeneities, that carried surprising information.

By the early 1990s, it had become clear that the relationship between center and surround in visual cortical neurons was more complex than even Barlow had suspected. Cortical area V1 neurons sensitive to vertically oriented bars, for example, responded more strongly when horizontally oriented bars were arranged outside their “classical” receptive fields and were suppressed in a nonlinear way when vertically oriented bars were placed at those same locations. David Heeger and his colleagues (Heeger 1992; Heeger et al. 1996) responded to that observation by developing what came to be known as the *Heeger normalization model* (Fig. 1A). In that model, the firing rate of a given neuron (FR_1) is described by

$$FR_1 = \frac{Stim_1}{\sigma^2 + \sum_i Stim_i}. \quad (1)$$

Heeger’s idea was that the activity of a V1 neuron reflected a direct excitatory input from its receptive field ($Stim_1$) that was modulated by the summed activity of all of its neighbors (the set of all other $Stim_i$). This modulation was achieved by dividing the input activity by the summed activity of the neighbors, a model that decreased the FR, when the surround was similar. The model also included a constant in the denominator: σ^2 . Note that the magnitude of σ^2 controls the rate of rise of the firing rate function, as shown in Figure 1B. It should be noted that Heeger’s model also included a number of other features glossed over here.

In the late 1990s, Eero Simoncelli and others (Olshausen and Field 1997; Simoncelli 2003) began to test the hypothesis that models of this kind formally achieve Barlow’s goal: They maximize the information carried by each action potential and by each neuron. To understand their approach, consider a neural system that encodes the luminance of a single square pixel surrounded by eight other pixels (Fig. 2A). Simoncelli and his colleagues recognized that in the real visual world, the luminance of the central pixel is not independent (or, formally, *uncorrelated*) with the luminance of the surrounding pixels. In the real world, if we knew that the eight outer pixels were black, we could be almost certain that the central pixel

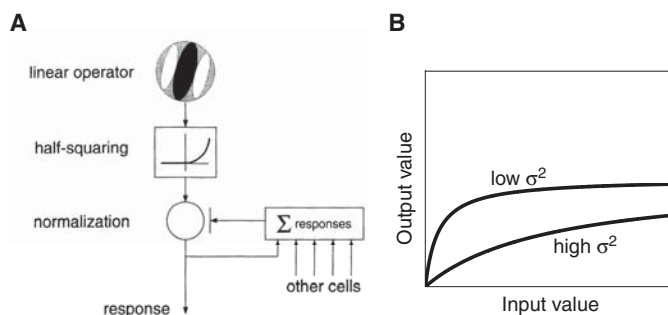


Figure 1. (A) The original divisive normalization model of Heeger and colleagues (Heeger 1992; Heeger et al. 1996). Visually sensitive neurons receive a direct excitatory input that is rectified (negative responses are removed) by half-squaring. This rectified output is then divided by the sum of the activity of similarly sensitive neighbors plus a constant, usually called σ^2 . (B) Firing rate curves for a simulated normalized neuron are shown for a high and a low value of σ^2 . A low σ^2 focuses the dynamic range of the function at low input values. A high σ^2 distributes the firing rate function over a larger range. (A, Courtesy of David Heeger.)




Figure 2. (A) An array of local pixels from a photograph. (B) The correlation between a central pixel and eight surrounding pixels is quite high in natural images.

was also black. Conditional on the outer pixels being black, there is very little information gained by learning that the central pixel is also black. An efficient neuron would know that and would waste very few action potentials to communicate this unsurprising fact.

Simoncelli and his colleagues (Schwartz and Simoncelli 2001; Wainwright et al. 2002) thus took two approaches to revising the Heeger equation. At a purely mathematical level, they set out to prove that there exists some variant of the Heeger model that maximizes information per action potential and per neuron for the real visual world. They were successful in this regard. At an empirical level, they set about a massive exercise analyzing real world photographs to see just exactly how adjacent pixels of images are related. What this second exercise yielded, and this is a critical insight, was a matrix that shows the degree of relatedness between adjacent pixels in the real world. The numbers in each box in Figure 2B reveal how much knowing the state of that pixel tells you about the central pixel. This is, in the language of linear algebra, a *correlation matrix*. Now with that insight in hand, we can easily make sense of the model Simoncelli and his colleagues produced:

$$FR_1 = \frac{Stim_1}{\sigma^2 + \sum_i \omega_i Stim_i} \quad (2)$$

The numerator remains the direct input to the neuron, but now the denominator takes all of the activity in nearby neighbors and multiplies each by the degree to which it allows one to predict the value of $Stim_1$. The sum $\omega_i Stim_i$ aggregates the total knowledge one can derive from the activity of the neighbors about our pixel, and then the equation pulls that information out of the firing rate by dividing it away. The result is a maximally efficient signal, “cleaned” of redundant information.

What followed this demonstration were a host of papers showing that divisive normalization models accurately explained the firing rates of neurons in cortical sensory areas ranging from V1 to V2 and even in auditory and olfactory areas. (For a review of this literature, see Reynolds and Heeger 2009; Carandini and Heeger 2012.)

NORMALIZATION MODELS AND DECISION-MAKING

The observation that many cortical and subcortical systems appear to use a normalization-based encoding scheme led Kenway Louie and I to test the hypothesis in the late 2000s that cortical circuits important for decision-making might also use this encoding strategy (Louie et al. 2011). Over the preceding decade my colleagues and I had conducted detailed studies of the activity in the lateral intraparietal area (LIP) of the monkey parietal cortex. We had shown that the activity of these neurons often appears to encode decision variable–like signals (Platt and Glimcher 1999; Glimcher 2013). Each neuron in LIP, we found, encodes the desirability of looking at a particular point in extrapersonal space. We had hypothesized that LIP firing rates could be thought of as encoding a “priority map.” The higher the firing rate, the more the animal wants to look at the location encoded by that neuron.

Recall that the *why* theories of economists like von Neumann tell us what an idealized decision-maker *should* do when she makes a choice: As the magnitude of a reward increases, a chooser should be more likely to pick that reward (it should have higher priority), and as the probability of being rewarded declines, a chooser should be less likely to pick that reward. These theories hypothesize that choosers behave exactly *as if* somewhere in their brains a numerical value called a decision variable is computed and represented for each option and the process of choosing is selecting the option having the highest decision variable. Working from those theories in the late 1990s, Michael Platt and I (Platt and Glimcher 1999) had shown that neurons in LIP do follow the general predictions of these *why* models, and we hypothesized that neurons like those in area LIP could be viewed as physically instantiating these decision variables, which my colleagues and I came to call *subjective values* (Fig. 3).

Of course we knew that the *why* theories of economists are imperfect predictors of behavior. We had always argued that the *why* models we used to explain cortical firing rates should only be considered approximations. It was with that caveat in mind that Louie and I formally examined the firing rates of LIP neurons with a cortical normalization model. To do that, we used a simplified form of the Heeger–Simoncelli model that John

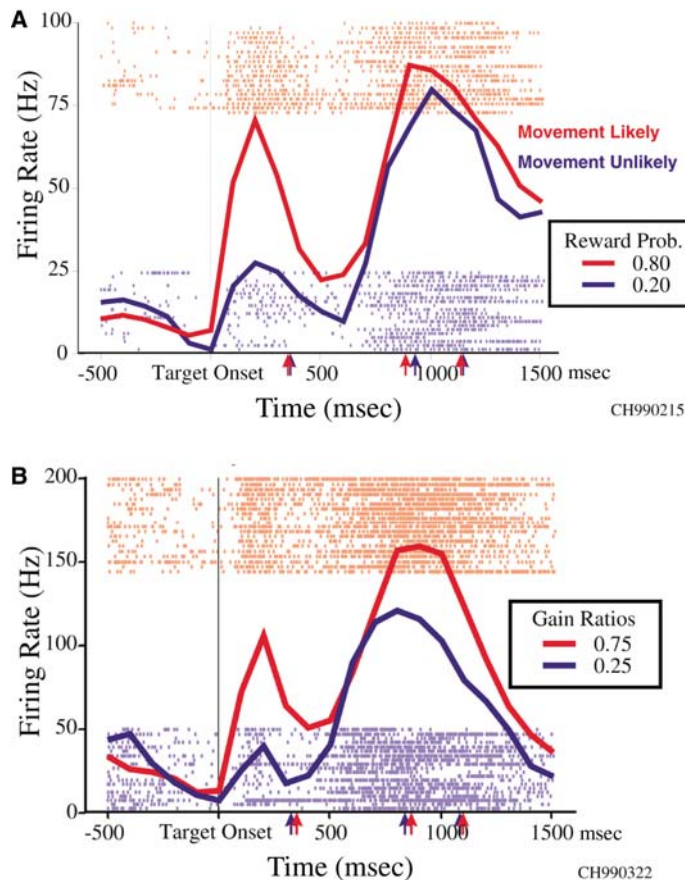


Figure 3. Firing rates for two neurons in the lateral intraparietal area (LIP). (A) Solid lines indicate average firing rates for a neuron as a function of time under two conditions. The red line indicates average response on trials in which there was a high probability of reward. The blue line indicates firing rate on trials with a low probability of reward. (B) A similar figure for varying reward magnitude. Red identifies trials on which the monkey expects a large reward, and blue identifies trials on which the monkey expects a small reward. (Adapted from Platt and Glimcher 1999.)

Reynolds and Heeger (Reynolds and Heeger 2009) had recently proposed:

$$FR_1 = \frac{\text{Baseline} + V_1}{\sigma^2 + \sum_i \omega_i V_i}, \quad (3)$$

which included a baseline firing rate and took as inputs the objective values of the rewards that the monkey (and his neurons) were evaluating. Recall that a core prediction of the normalization models is that stimuli in the surround influence firing rate through the denominator. One can see this influence as a process of decorrelation, achieved by multiplying the activity in each of the other neurons by a covariance matrix (ω). Entirely equivalently, one can see this as a process of “contextual suppression,” which says the same thing in a less mathematical way. But how should one think about contextual suppression in a decision-making system? One might imagine that the other options currently under consideration by a chooser (like the other stimuli presented near a visual receptive field) form a kind of “context” that might modulate the firing rate of the neuron under study.

Consider a situation in which the option represented by neuron FR_1 yields a much larger reward (say 1 mL of

water) than two other options being represented by two other neurons. Under those conditions, our neuron might be expected to have a very high firing rate, correctly signaling that the reward under consideration is large compared with the other two options under consideration. But under conditions in which the two other options are of much higher value than the 1 mL reward encoded by FR_1 , then we might expect FR_1 to show a much reduced firing rate—even though the absolute value of the reward it encodes is the same. The key idea here is that the FR_1 neuron, if it follows the normalization model, should report the value of an option modulated by current context just as V1 neurons report the orientation of a stimulus modulated by local orientation context. (And, of course, both of these contextual modulations are a form of decorrelation.)

To test the hypothesis that LIP neurons do use a normalization-like encoding scheme, Louie and I examined the delay-period activity of LIP neurons while monkeys were presented with three options of variable magnitude (Louie et al. 2011). What we found was that the neurons were very much influenced by local context. And that this local context effect could be best described using a normalization model. Figure 4 shows this for a database of neurons and compares the ability of three models to ac-




Figure 4. Comparison of predicted and observed firing rates for three models: the original relative value model of Platt and Glimcher (1999); the “difference model” underlying firing rates in the drift diffusion model; and the normalization model.

count for those delay-period firing rates. Firing rates are much better accounted for by the normalization model than by two competing models that have been widely used in the past.

THE “NORMALIZATION MODEL” CONSIDERED AS A NETWORK

A key idea of the normalization model is that neurons do not just represent their direct inputs, but that these inputs are modulated by the activity of their neighbors, a *contextual sensitivity* that maximizes their efficiency. Barlow, Heeger, and Simoncelli had considered this in a static sense: How should steady state firing rates incorporate context to maximize the information carried by each action potential and by each neuron? Simoncelli largely solved that problem, proving that a variant of the original model (one incorporating a covariance matrix) maximizes information. As Tom LoFaro, Louie, and I worked on this problem, however, we began to wonder how real neural circuits accomplish this decorrelation (LoFaro et al. 2014). Figure 5 shows the simplest possible network for computing this kind of normalization model. The values of two options, V_1 and V_2 , serve as excitatory inputs to encoding neurons R_1 and R_2 . As in Figure 1, the R neurons receive an inhibitory input (assumed to be divisive although that is not critical) that pools together the outputs of all of the R neurons. (In this network there are only two




Figure 5. A simple network for divisive normalization. Value inputs (V) enter excitatory neurons R , the activity of which is reduced by a divisive inhibitory interneuron, G . For simplicity, the model shown here includes only two pairs of neurons and thus supports pairwise choice. Extension of the model to n pairs yields results identical to the ones reported here for two-pair models.

R neurons, though we have shown that our findings generalize to large numbers of R – G neuron pairs.)

The excitatory neurons of the network are its backbone. They get excitatory inputs that we think of as “value” signals, but which in the visual cortex one might think of as the contrasts of pixels. A network of inhibitory neurons receives inputs from the excitatory neurons. One can think of the strengths of the synapses connecting the excitatory neurons to the inhibitory neurons as implementing the covariance matrix of Equation 2, although for simplicity here we will just treat those synapses as having a fixed unit strength. That is all there is to this network. To examine how such a simple network would behave dynamically we need to represent how firing rates change when inputs change with simple differential equations:

$$R \text{ neuron : } \tau \frac{dR}{dt} = -R + \frac{V_1}{1 + G_1}, \quad (4)$$

$$G \text{ neuron : } \tau \frac{dG}{dt} = -G + \sum_i \omega_i R_i. \quad (5)$$

Notice that these equations describe how neuronal firing rates change from one moment to the next, they report the first derivative of firing rate with respect to time. The first equation represents the first derivative of the firing rate for an R neuron, and the second equation represents it for a G neuron. τ is a time constant for these neurons.¹ Notice that the R neuron responds with an increase in firing rate when the strength of the input V increases. It responds with a (divisive) decrease in firing rate when the inhibitory neuron G increases its firing rate.

To better understand the behavior of such a network, we can select a value for V_1 and then solve the equations to see how the firing rate of each neuron would evolve over time for any possible preexisting firing rate in the R and the G neuron. Figure 6 plots a graph of that mathematical “experiment.” The arrows at each point on the graph show how firing rates evolve from that point when $V_1 = 10$. On such a plot, called a phase plane, we can look for all of the places where (for a given input V_1) the first derivative of the R neuron or the G neuron goes to zero. These two lines, called *null clines*, show where each neuron is stable. Wherever these lines intersect, where neither neuron changes, the network is in a stable *equilibrium* state. Interestingly, we were able to prove in a quite general way that these dynamic networks always have one (and only one) stable equilibrium. Once the input is set, the network evolves toward a unique equilibrium (regardless of how many R – G neuron pairs the network contains).

To understand how important this fact is, we need to rewrite the equations above in terms of actual firing rates for R neurons. Equation 6 shows such a restating of these

¹Here we use a single time constant for both kinds of neurons, but our findings are robust to the selection of time constant. We often set the time constant for the R neurons to 10 msec and for the G neurons to 5 msec to mirror the membrane time constants of excitatory and inhibitory cortical neurons.




Figure 6. A phase plane trajectory for a single pair network. The red line indicates the firing rate path when the value input is abruptly stepped to a positive value. Dashed lines identify the nullclines for the two neurons. The intersection of the dashed lines marks the unique equilibrium state for the network: *divisive normalization*.

equations, but rather than being plotted in continuous time we plot it here in discrete time (this discrete form representation is equivalent at equilibrium to the continuous time form in Equations 4 and 5):

$$R_t = \frac{V_1}{1 + \sum_{i=0}^{-\infty} \sum_i \delta^i \omega_i R_i}. \quad (6)$$

What you can see is that the firing rate of an R neuron at time t reflects the current excitatory input (V_1) divided by a denominator that may look complicated, but is not. In the denominator is a double sum. The right-hand part of the sum is familiar. It is simply the activity of the inhibitory neuron that computes a (covariance matrix) weighted sum of the R neurons (rather than the V inputs) in the network: $\sum_i \omega_i R_i$. To capture the dynamics, though, we have added one additional feature, a sum over time. The outer sum and the δ (which is a number < 1) cause the equation to weight current value-related activity heavily (when t is small) and allows older value-related activity to decay away (as t grows into the past). The function of this “temporal summation” is simply to exponentially decrease the effects of old inputs on the network—to make old inputs decay away. We can think of the size of δ as representing the time constant of the network.

To take our understanding one step further, let us algebraically separate the denominator into two parts: The current activity of the excitatory network of R cells and the past activity of those same cells. The next equation shows that

$$R_t = \frac{V_1}{[1 + \sum_{i=-1}^{-\infty} \sum_i \delta^i \omega_i R_i] + \sum_i \omega_i R_i}, \quad (7)$$

$$R = \frac{V_1}{[\sigma^2] + \sum_i \omega_i V_i}. \quad (8)$$

Now compare the dynamic equation (Equation 7) with the original static Heeger–Simoncelli shown in Equation 8. The two equations are almost identical—the only big difference is that the constant σ^2 has turned into something much more complicated and dynamic, as shown in

the bracketed parts of Equation 7. That in itself is an important point; when you build a divisive network of almost any kind you have something that generates as an output much like the original normalization models. To put that another way, normalization models are what divisive (and many subtractive) networks naturally do.²

Now consider the term that has replaced σ^2 . This is the term that captures all of the temporal dynamics of the system. When the firing rate is changing but the inputs are constant, it is because this part of the equation is changing as time passes. What is critical about that is: *Whenever the stuff in the bracket is not changing—when the stuff in the bracket is a constant and the inputs are fixed—the system is at equilibrium and becomes a Heeger–Simoncelli-type model.* The Heeger–Simoncelli model is, it turns out, a special case of a divisive network. It is what a divisive network does when it comes to rest.

Recall from Figure 1A that the value of σ^2 controls how steeply the firing rate function rises. A small constant here causes the function to rise quickly, and a large constant causes it to rise slowly. So what is this constant in the general form? It is simply an exponentially weighted sum of past activity derived from past input values—what we would call past “options” in decision-making. In more formal terms, it is an “adaptive expectation.” *It is the expected average value of the upcoming inputs, based on an analysis of recent history.* What the “constant” does in this dynamic network is to adjust the firing rate function so it has the greatest possible resolution near the current expectation. What that means is that these classes of networks are much more powerful and general than one might have expected. At equilibrium they compute a form of divisive normalization. Their denominators contain a “tuning constant” that maximizes the firing rate function’s sensitivity around the network’s expected input.

First, that tells us that the static equations of the Heeger and Simoncelli forms in fact describe the unique equilibrium state of *nearly any network of this class*. If you have an excitatory input to your excitatory neuron and a pooling inhibitory network, you wind up with a network that at equilibrium computes divisive normalization.

NORMALIZATION NETWORKS AND NEURONAL DYNAMICS

Interestingly, these dynamic networks make an additional prediction. Recall that R neurons receive excitatory value input (V s) and inhibitory input from the G neurons, G , but the G input is delayed, reflecting the time it takes for the inhibitory network to become active. One can see this in Figure 6, which shows the R neuron overshooting before being pulled down by the G neuron.

²Note, however, that the denominator in the Heeger form (Equation 8) has V s in it. The denominator here (Equation 7) has R s in it. We have shown that this distinction does not make much of a difference at equilibrium. It does make a significant difference with regard to network dynamics. Louie, LoFaro, Webb, and I have shown that the actual dynamics of LIP neurons are much more closely approximated by this form, a point discussed next and in Louie et al. (2014).

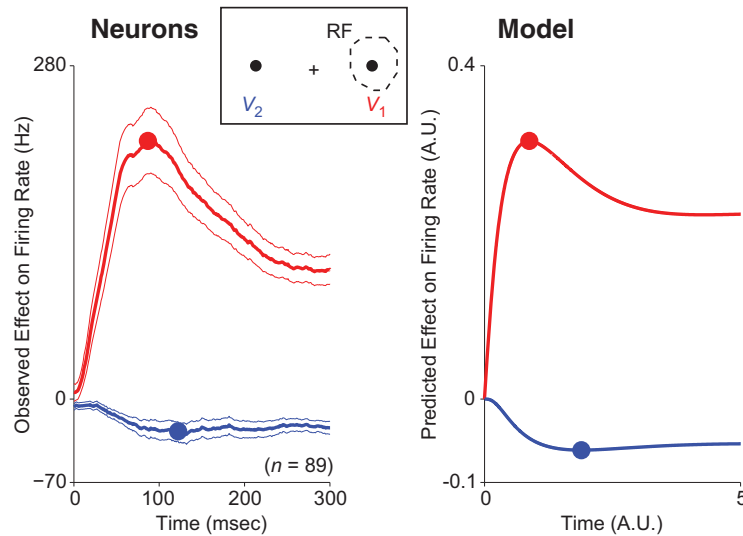


Figure 7. Excitatory and inhibitory contributions to population firing rates in area LIP. (*Left*) Observed data from an LIP population. (*Right*) The prediction of the network model that achieves divisive normalization at equilibrium.

Thus, R neurons respond to a step change in V with a transient burst-like increase before they settle down to stable “delay-period” activity. In real life, area LIP neurons have long been known to show what is often called a “visual burst” before sustained delay-period activity. What the model reveals is that for divisive (or subtractive) networks of this kind there is nothing “visual” about the “visual burst.” The early rise in firing rate is the natural network response when its input goes up. The “burst” is the off-equilibrium response of the network before it settles to a sustained equilibrium.

Louie and I (Fig. 7; Louie et al. 2014) examined the dynamics of LIP neurons to see how closely their “visual burst” could be predicted by these kinds of model. What we found was that the magnitude and timing of the burst and the structure of its later suppression were exactly what would be expected of such a network.

NORMALIZATION NETWORKS AND CHOICE BEHAVIOR

Normalization networks capture the idea that neurons have limited capacity and that they distribute this capacity to maximize their efficiency. But their limited capacity necessarily means that they perform imperfectly as the amount of information they encode grows. Although we have not focused on this issue yet, the source of that capacity limitation is neuronal noise. If firing rates could encode input values to 20 decimal places, then they would have no real capacity limit. The fact that firing rate variance (roughly) scales with mean rate, however, means that these neurons are quite limited in their precision.

Consider a network representing two high-valued options, each of almost equal value and a third low-value (and undesirable) option. Figure 8A shows a firing rate *distribution* for three R neurons encoding the values of these three options at equilibrium. Each of the colored


curves represents the distribution of observed firing rates over time. What we can see is that the network does a good job of separating (decorrelating) the two highest-valued options, even given that the firing rates fluctuate randomly within the colored distributions. Consider, however, what happens as we increase the value of the undesirable option (Fig. 8B). The firing rate functions begin to overlap—the decorrelation begins to fail. And thus it gets less and less likely that the network will uniquely identify the highest-valued option as option number grows.³

To see whether choice behavior degrades in exactly this way, Ryan Webb, LoFaro, Louie, and I (Louie et al. 2013; Webb et al. 2014) developed a series of human and animal choice experiments in which we searched for anomalies like the one described above. Perhaps surprisingly, we found that each of the predictions our models made accurately described an anomalous choice behavior found in humans and in monkeys that defied traditional economic analysis. This is a key point: Our model predicts that the predictions of traditional economic theory are wrong because traditional theory neglects the cost of representation. In any real system, where neurons are costly and capacity is limited, choice accuracy will have to trade off representational costs against benefits. What this suggests is that the gap between the *whys* of the economists and the *whats* of the psychologists really reflects a misspecification of the costs of neural computation.

CONCLUSION

Neuroeconomic models of human and animal decision-making emerged as an effort to reconcile economic theory and psychological observation. Economists had

³Louie and I showed that this is true for essentially any noise regime. That is a somewhat counterintuitive fact covered in Louie et al. (2011).



As u_3 increases, the distance between u_1 and u_2 shrinks

Figure 8. Firing rate distributions for normalization models show a sensitivity to the value of irrelevant options. See text.

long argued that choice was a process of optimization and had developed lovely theories of how that optimization should work. But what the economists had failed to take into account was the fact that neural computation is costly and neural representational capacity is limited. Psychologists had correctly noted the failures of traditional economic theory but offered little deep explanation for why these anomalies arise. What has become clear in the last decade is that models that incorporate real constraints faced by the nervous system can close much of this gap. These models recognize that perfect choice behavior requires an infinitely costly nervous system. Real animals operating under evolutionary pressure accept errors whenever correcting those errors would be more costly (in terms of a larger or more precise nervous system) than beneficial. The study of human decision-making stands poised to use neuroscience to develop a next-level understanding of the *whys*, *whats*, and *hows* of human behavior.

REFERENCES

- Allais M. 1953. Behavior of the rational man before risk—Criticism of American school postulates and axioms. *Econometrica* **21**: 503–546.
- Barlow HB. 1961. Possible principles underlying the transformation of sensory messages. In *Sensory communication* (ed. Rosenblith WA), MIT Press, Cambridge, MA.
- Carandini M, Heeger DJ. 2012. Normalization as a canonical neural computation. *Nat Rev Neurosci* **13**: 51–62.
- Glimcher PW. 2013. Value-based decision-making. In *Neuroeconomics: Decision-making and the brain* (ed. Glimcher PW, Fehr E), Academic, New York.
- Heeger DJ. 1992. Normalization of cell responses in cat striate cortex. *Vis Neurosci* **9**: 181–197.
- Heeger DJ, Simoncelli EP, Movshon JA. 1996. Computational models of cortical visual processing. *Proc Natl Acad Sci* **93**: 623–627.
- Kahneman D, Tversky A. 1979. Prospect theory—Analysis of decision under risk. *Econometrica* **47**: 263–291.
- LoFaro T, Louie KL, Webb R, Glimcher PW. 2014. The temporal dynamics of cortical normalization models in decision-making tasks. *Lett Biomath* **1**: 209–220.
- Louie K, Gratton LE, Glimcher PW. 2011. Reward value-based gain control: Divisive normalization in parietal cortex. *J Neurosci* **31**: 10627–10639.
- Louie K, Khaw MW, Glimcher PW. 2013. Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci* **110**: 6139–6144.
- Louie K, LoFaro T, Webb R, Glimcher PW. 2014. Dynamic divisive normalization predicts time-varying value coding in decision-related circuits. *J Neurosci* **34**: 16046–16047.
- Neumann J, Morgenstern O. 1944. *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.
- Olshausen BA, Field DJ. 1997. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vis Res* **37**: 3311–3325.
- Platt ML, Glimcher PW. 1999. Neural correlates of decision variables in parietal cortex. *Nature* **400**: 233–238.
- Ratliff F, Hartline HK. 1959. The responses of *Limulus* optic nerve fibers to patterns of illumination on the receptor mosaic. *J Gen Physiol* **42**: 1241–1255.
- Reynolds JH, Heeger DJ. 2009. The normalization model of attention. *Neuron* **61**: 168–185.
- Samuelson PA. 1947. *Foundations of economic analysis*. Harvard University Press, Cambridge, MA.
- Schwartz O, Simoncelli EP. 2001. Natural signal statistics and sensory gain control. *Nat Neurosci* **4**: 819–825.
- Shannon CE, Weaver W. 1949. *The mathematical theory of communication*. University of Illinois Press, Urbana, IL.
- Simon HA. 1955. A behavioral model of rational choice. *Q J Econ* **69**: 99–118.
- Simoncelli EP. 2003. Vision and the statistics of the visual environment. *Curr Opin Neurobiol* **13**: 144–149.
- Wainwright MJ, Schwartz O, Simoncelli EP. 2002. Natural image statistics and divisive normalization: Modeling nonlinearity and adaptation in cortical neurons. In *Probabilistic models of the brain: Perception and neural function* (ed. Rao RPN, Olshausen BA, Lewicki MS), pp. 203–222. MIT Press, Cambridge, MA.
- Webb R, Glimcher PW, Louie KL. 2014. *Rationalizing context-dependent preferences: Divisive normalization and neurobiological constraints on choice*. SSRN Electronic Library 2462895.



Cold Spring Harbor Symposia on Quantitative Biology

Understanding the Hows and Whys of Decision-Making: From Expected Utility to Divisive Normalization

Paul Glimcher

Cold Spring Harb Symp Quant Biol published online January 30, 2015
Access the most recent version at doi:[10.1101/sqb.2014.79.024778](https://doi.org/10.1101/sqb.2014.79.024778)

P<P Published online January 30, 2015 in advance of the print volume.

Email alerting service Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#)

Advance online articles have not yet appeared in the print volume. Citations to Advance online articles must include the digital object identifier (DOI) and date of initial publication.

To subscribe to *Cold Spring Harbor Symposia on Quantitative Biology* go to:
<http://symposium.cshlp.org/subscriptions>
