

The Case for a Common Currency of Value as a Basis for Decision-Making and Intersystem Competition

Chandra Sripada
University of Michigan

Abstract

Inspired by economic theory, neuroeconomics contends that value serves as a common currency for decision. This *valuational* view has amassed impressive support over the last two decades, but critics have recently raised serious concerns. They argue that there are important pathways to action that bypass value, highly dissimilar psychological systems do not share a common value currency, and the view cannot accommodate well documented kinds of “irrationality”, such as context effects. This essay offers a systematic response. Key refinements are made to the valuational view allowing for multiple formats for value information and two levels of valuation (options-level and systems-level). These refinements help address the criticisms and yield a more nuanced understanding of the role of value in human decision-making.

Introduction

What explains why people do the things they do? Many influential answers have been given [1]. Behaviorists proposed learned associations [2]; ethologists appealed to intrapsychic hydraulic forces [3]; religious thinkers invoked an immaterial self that is the ultimate chooser.

Contemporary computational cognitive science, specifically the field of *neuroeconomics*, offers a distinctive answer inspired by economic theory: People make decisions about what to do based on information about the expected value of their options, and they choose the option with highest value [4–7]. Value is a scalar variable that represents the subjective goodness or badness of an option, similar in some ways to the role played by utility in economic theories.

Two Core Theses of Neuroeconomics

While neuroeconomics encompasses a number of claims over which there are disagreements, this idea—that decisions are made on the basis of value—is one core thesis that unites the research paradigm.

A second core thesis widely endorsed in the field is that value plays the role of a “common currency” for choice [8–10]. People make decisions that trade off very different kinds of goods, for example food, money, safety, leisure, and social gains, among other things. Neuroeconomics claims that value serves as a common medium of exchange that allows for selection among dissimilar alternatives.

The preceding two theses, which we can call the *valuational view of choice*, are at the heart of the neuroeconomic research program. Moreover, over the last two decades, the valual view has gained substantial empirical support. For example, there is strong evidence from behavioral, chronometric, computational, and imaging methods for value representations as the basis for a wide range of economic and social choices [11–14]. In addition, mechanistic models that are aligned with the valual view are prominent across multiple psychological domains including attention allocation [15–18], habitual action [19], working memory [20,21], and executive control [22,23], among others.

Criticisms of the Neuroeconomic Valual View of Choice

Nonetheless, the neuroeconomic research program, and the valual view specifically, has come in for a number of lines of criticism. These arguments are scattered about the literature [24–30], though a recent article by Benjamin Hayden & Yael Niv [31] assembles several of the most important critiques.

One argument looks closely at algorithms from the reinforcement learning (RL) field [32]. These algorithms are highly important for the neuroeconomics program because they provide computationally feasible and biologically plausible accounts to explain how agents can, using ongoing feedback from the environment, acquire information about the value of their options. There are, however, certain algorithms in the RL framework that are said to work by calculating the best actions “directly”, without the need for value estimation, for example so-called policy gradient algorithms [33]. These “value-free” algorithms, it is argued, demonstrate that there are alternative pathways for decision and action selection that bypass calculation of value, seemingly at odds with the first core thesis of neuroeconomics [25,31].

Another argument says that representing the value of one’s options, while it certainly happens in some cases, is in fact relatively rare because it is much too computationally demanding. Instead, people use *heuristics* to rapidly select among a set of options without having to exhaustively calculate the value of the options [34,35]. It is now well established

that heuristics are employed ubiquitously in human decision-making, and this fact is said to challenge the neuroeconomic valuational view of decision [31].

A third argument takes issue with the neuroeconomics claim that value plays a common currency role. The human brain is outfitted with numerous mechanisms and systems that contribute to choice and action guidance, many of which operate using distinctive computational principles [36,37]. For example, some systems rely on a model of the world, other systems use “model-free” methods, still others use Pavlovian principles, and so forth. These observations raise doubts about whether there is a common medium of value that is shared across all these disparate systems allowing for commensuration across them

A fourth argument says the neuroeconomic valuational view, inspired by economic theory, has trouble accounting for certain forms of “irrationality.” People display a number of well-documented context effects—their preferences change depending on seemingly irrelevant features of the context in which they are elicited. Critics argue the valuational view of choice cannot account for these phenomena [24,27,30].

A Defense of the Valuational View

The aim of this essay is to systematically defend the neuroeconomic valuational view against these criticisms while acknowledging that the critics’ arguments involve key insights and yield important lessons. Two clarifications play key roles throughout the essay.

The first clarification concerns different kinds of information available to an agent about the value of an option. It has been largely assumed in the neuroeconomics literature, among supporters and critics alike, that value information is conveyed in a magnitude-type format, analogous to cardinal (i.e., real-valued) utilities in economics. It will be argued, however, that certain important algorithms and psychological processes provide value information in alternative, non-magnitude formats, and this fact has important implications for how to understand the valuational view.

The second clarification distinguishes the role of value as a common currency across dissimilar options (“options-level valuation”) versus across competing processes or systems (“systems-level valuation”). This distinction has been underappreciated in the field. Moreover, systems-level valuation fills in a key gap: It helps to explain how the common currency thesis can still be maintained even when there is no shared medium of value in which the options themselves can be directly compared.

The essay is divided into five parts. Part 1 delves deeper into the idea that value information can be supplied in multiple formats. The next four parts take up the four arguments against the valuational view put forward by critics. It is argued that when key clarifications are made and certain misconceptions are addressed, the observations made

by critics of the valuational view are seen to no longer count against it. Indeed, in many cases these observations end up providing additional strong support for the two core theses of the neuroeconomics program.

1. Multiple Formats for Value Information

The valuationist view of decision can be broken up into two stages. In stage 1, the agent assembles information about the value of the options. In stage 2, the agent selects the option with highest value.

This two-stage picture, with a separate information assembly stage, helps to highlight a key distinction between *value* itself and *information about value*. The neuroeconomic valuational view is committed to the idea that value itself is a scalar variable. It represents the expected subjective goodness or badness of an option in a continuous “cardinal” scale, similar to utility in economic theory.

Information about value, on the other hand, is information or evidence that the agent has regarding the values of its options. Importantly, while the value of an option is itself a scalar variable, the information that the agent has about the value of its options can involve a variety of representational formats, including formats that convey non-scalar information.

These ideas can be illustrated with the example of *price*. The price of a good or bundle of goods is a scalar variable, a continuous quantity that indicates how much must be paid at checkout. In contrast, *information about price* can be supplied in a variety of formats. For example, consider apples and bananas, the only two things sold by a certain store. One can say: an apple is 25 cents, and a banana is 40 cents; a banana is more expensive than an apple; a banana is more expensive than an apple by 15 cents; or an apple has the lowest price. The first statement presents information about price in a magnitude-type format, but the others do not. For example, the statement that an apple has the lowest price assigns a nominal property to the apple: being the least expensive.




Figure 1: Multiple Formats for Value Information. The valuational view distinguishes information gathering at stage 1 and selection at stage 2. Additionally, the model allows that even though value is itself a scalar variable, the information about value gathered at stage 1 can involve multiple representational formats. EV = expected value.

Figure 1 presents an analogous situation with respect to value. Though value itself is a scalar variable, information about value can be supplied in multiple different formats, including magnitude-type format, ordinal format, relative value format, or maximal value format, among others.

The fact that information about value can be supplied in a variety of formats becomes especially important when we additionally consider that, on the valuational view, the agent selects the option with the highest value. For this selection task, having magnitude-type information is not always necessary.

To see why, suppose you are trying to identify which of two shopping carts of goods, **A** or **B**, has the lowest price. And suppose further you know the carts contain identical items except that cart **B** has an organic apple and cart **A** has a cheaper regular apple—a piece of price information in an ordinal format. That is sufficient to infer that **A** has the lowest price, even if you don't know the magnitudes of the respective prices of the carts themselves.

A similar point applies to identifying which option, **A** or **B**, has the highest value. Suppose you know that **A** beats **B** on the only attribute on which they differ—a piece of value information in an ordinal format. That is sufficient to infer that **A** has the highest expected

value, even if you don't know the magnitudes of the respective expected values of the options themselves.

We can call processes like these—processes in which non-magnitude-type value information is used to draw conclusions about scalar expected value—*alternative pathway processes*. Importantly, these processes are fully compatible with the valuational view. This view is only committed to the idea that value itself is, like price, a scalar variable that allows comparison across dissimilar items. It is agnostic about the format of the information that is utilized to draw conclusions about which option has highest value, and this information might be in variety of non-magnitude-type formats.

Critics of the neuroeconomics program appear to have assumed that the field is committed to the view that information about value only occurs in a magnitude-type representational format. And theorists associated with the program themselves have often presented their views in ways that support this impression and have done little to correct it. Thus, the fact that value information can be supplied in a variety of formats, including non-magnitude-type formats, has not typically been made salient.

In the sections that follow, this key clarification will be applied in the context of criticisms of the neuroeconomic valuational view. Several algorithms and processes are discussed that have been taken by critics to be counterexamples to the valuational view. It is argued that these algorithms and processes have been misunderstood. They do not in fact bypass value, as the critics allege, but rather, they are alternative pathway processes. They supply information about value in non-magnitude type formats, and they are in fact fully consistent with the valuational view.

2. Reinforcement Learning and Policy Gradients: Do Some Learning Algorithms Bypass Value?

The reinforcement learning (RL) framework addresses the problem of how agents learn to behave in complex environments in ways that maximize achievement of their aims [32]. There are a number of importantly different kinds of algorithms within the RL framework. Many of these algorithms supply decision-making processes with magnitude-type information about the value of the available options. For example, in the well-known temporal difference reinforcement learning algorithm [32,38], the agent maintains estimates of the scalar expected values of their actional options for each environmental state, and these estimates are iteratively adjusted based on ongoing environmental feedback. In many settings, with sufficient experience, the algorithm eventually converges on accurate scalar value estimates to guide the agent's decisions.

There is substantial evidence that these kinds of scalar value-based RL algorithms do in fact play wide-ranging roles across multiple areas of cognition, including attention allocation [15–18], habitual action [19], working memory [20,21], and executive control [22,23]. These observations, in turn, provide compelling evidence for the neuroeconomic valuational view. This because the view says that agents use value information to select the option with highest value; this was the first core thesis of the neuroeconomics program. Scalar value-based RL algorithms provide one important explanation, though certainly not the only one, for where this value information comes from.

But there are some RL algorithms that are said to avoid estimates of value. For example, so-called policy gradient methods are said to arrive at an action policy “directly” without calculating values [33]. Policy gradient algorithms, or closely related algorithms such as actor-critic methods [39], also have biological plausibility as algorithms implemented in the human brain, especially as a basis of routinized action control [25,39]. This fact potentially presents a serious problem for the neuroeconomic valuational view: Policy gradient methods point to the presence of important pathways to action that bypass value.

However, a closer look at the workings of policy gradient algorithms calls this line of criticism into doubt.

In policy gradient algorithms, the agent assigns so-called action weights to its actional options for each environmental state, where these weights encode the probability of performing the associated action (the collection of action weights assigned across all states is called a “policy”). These action weights can initially be set randomly. The agent next observes rewards that arise during action sequences and uses these observations to, via a gradient-based method, nudge the action weights up or down based on their observed rewardingness. That is, actions that are observed to yield higher rewards are nudged up and these actions are, going forward, performed with higher probability; actions observed to yield lower rewards are correspondingly nudged down. It can be shown that across many settings, repeated application of this gradient-based process converges on an optimal policy for action [33].

Hayden and Niv [31] argue that that an agent selecting its actions through policy gradient methods does not choose their actions based on value. A policy gradient agent acts on the basis of action weights—these weights determine the probability of performing the respective actions. These weights, they argue, do not convey information about the value of the corresponding actions. Instead, they are more like direct recommendations for how to act.

For example, suppose a policy gradient algorithm is used to learn action weights for the children’s game rock/papers/scissors, and eventually learns an optimal policy: 1/3 weight assigned to each action. These action weights do not represent the value of the respective actions. Rather, they are recommendations that one should perform each action 1/3 of the time. Haden and Niv thus conclude that policy gradient algorithms present a

counterexample for the neuroeconomic view that decisions for how to act are based on value.

There is an important problem, however, for Hayden and Niv's analysis. Notice that policy gradient algorithms are designed to acquire an *optimal* policy. Optimality in the RL framework means that an agent acting on the action weights specified in the policy achieves *maximum scalar expected value* (relative to the case in which actions are selected on the basis of alternative action weight assignments, up to the limit of ties). So, in the rock/paper/scissors example, the optimal policy that is learned effectively conveys that doing each action 1/3 of the time will maximize expected value. It follows that policy gradient methods clearly *do* convey information about scalar value: They convey that performing certain actions, the ones specified in the learned optimal policy, maximizes expected value.

We can frame the preceding point more sharply by leveraging the ideas discussed in Part 2 that value information can be conveyed in different formats. Policy gradient algorithms, which are designed to learn an optimal policy, supply information about scalar value in a maximum value-type format. In the terminology from Part 2, they represent "alternative pathway processes." Like the information that an apple has the lowest price, an optimal policy conveys that the specified actions (or probabilistic combination of actions) have highest expected value.

The information about value supplied by policy gradient methods is, of course, different than that supplied by algorithms like the temporal difference algorithm, which supply information about value in a magnitude-type format. Policy gradient methods in effect throw away magnitude information during the learning process, retaining just a specification of the actions (or probabilistic combination of actions) that yield maximum expected value. A key advantage of this approach is that it obviates the need at the time of decision for comparison of magnitude-type information about the values of the actional alternatives to identify the one with highest value, increasing the speed with which responses can be generated. Thus, policy gradient methods are frequently deployed for problems in which speed is at a premium, such as control of routinized movement.

A more general conclusion is that while some authors have spoken loosely that some RL algorithms learn value information while other algorithms do not [32,33], this way of putting things is not exactly correct. *All* RL algorithms supply information about value to guide decisions about how to act. What these authors likely meant is that many RL algorithms, such as the temporal difference algorithm, supply magnitude-type value information, but other RL algorithms, such as policy gradient algorithms, specify maximum value-type information.

Policy gradient RL algorithms thus do not present a problem for the neuroeconomic valuational view. Instead, the situation is somewhat the opposite. It was noted earlier that RL algorithms of various kinds are ubiquitously observed across multiple psychological

domains, including attention allocation, habitual action, working memory, and executive control. Furthermore, all RL algorithms supply information about value in some format or other for decision and action guidance. These two observations jointly provide strong evidence *for* the neuroeconomic valuational view.

3. Heuristics for Complex Decisions: Are Decision Shortcuts Incompatible with Neuroeconomic Theory?

Heuristic Decision-Making

People often need to make choices among a large number of options that differ across multiple attributes. Canonical normative models recommend a comprehensive aggregation strategy: value information is aggregated across the attributes of each option and this information is aggregated across the full set of options [40]—a computationally demanding endeavor. There is a sizable literature in judgment and decision-making research documenting that people often instead use heuristics, or decision shortcuts, instead of comprehensive aggregation [34,35,41,42].

For example, in the *priority* heuristic [43], a decision-maker compares their options along a single, important attribute dimension. If one option is found to be sufficiently better along this dimension, the option is selected. If not, another important dimension is identified and the process repeats.

In the *elimination by aspects* heuristic [44], the decision-maker identifies some important feature that an option should have. The option set is then culled of any options that lack this critical feature. If multiple options remain, another important feature is identified, and another round of culling ensues. The process is repeated until only one option is left, which is the option that is selected.

Heuristics and the Valuational View

The use of these sorts of heuristics might be seen to present a problem for the neuroeconomic valuational view of choice. This view, inspired by economic theory, says agents make choices on the basis of the value of their options, selecting the option with highest value. But an agent using heuristics ostensibly seems to be doing something quite different; they don't seem to be calculating the value of the options at all.

Hayden and Niv take up this line of critique [31]. They argue that whereas neuroeconomic theories take value to be essential for choice, heuristic approaches solve the problem of how to choose “without the need for computation of values.” They add, “Importantly for our argument, the success of heuristic approaches demonstrates that value calculations are not essential for neuroeconomic theories of choice. In particular, heuristic theories

can solve many problems for which value is proposed to be needed”, including for “comparison of multi-dimensional goods and bundles, and for choice across dissimilar goods.”

Heuristics are indeed widespread in decision-making. What is up for debate, however, is the further claim, captured in the quote from Hayden and Niv, that heuristic decision-making represents a problem or threat to the neuroeconomic valuational view. A key point that goes missing in Hayden and Niv’s discussion is that heuristics in fact rely on value information at multiple places. Once we see how information about value is essential for the operation of heuristics, it becomes clear that the use of heuristics is not incompatible with the valuational view after all.

The Role of Valuation in Heuristic Decision-Making

One place that value enters into heuristic decision-making is in selecting to use a heuristic in the first place. Consider an agent presented with a set of options, for example a choice to purchase one of several kinds of cars. A number of strategies are available for this decision-making problem. They could undertake an exhaustive multiattribute aggregation approach. Or they could take a heuristic approach, and there are a number of importantly different heuristic strategies available. How do agents select among all these available strategies?

This question has been the focus of decades of research in the judgment and decision-making field, and a pair of key observations are well established. First, people use a wide variety of decision-making strategies across different problem contexts [45]. For example, sometimes they focus on important dimensions, other times they make pairwise comparisons, and still other times they opt for exhaustive aggregation.

Second, people tend to employ strategies that are well-suited to the specific problem context at hand [45,46]. For example, in “non-compensatory” environments, in which looking at a small set of attributes is sufficient for optimal choice, people tend to employ heuristic strategies that home in on a small set of difference-making attributes. But in “compensatory” environments, in which a large number of attributes must be considered for optimal choice, people switch to strategies that more closely resemble exhaustive multi-attribute aggregation.

These observations raise a further question: How exactly do people know which strategy to use for which problem context? Building on prior work, Leider and Griffiths [47] model strategy selection in terms of a value learning problem. They propose that agents learn a predictive model that maps each problem type they encounter into strategies that maximize the value of computation (VOC) for that problem type. VOC in turn reflects a balance between the benefits of using a strategy, assumed to be primarily accuracy, and the costs, assumed to be primarily time spent executing the strategy.

The agent learns the VOC of available strategies for each problem-type through an RL-type iterative value learning algorithm, and when they encounter a new decision problem, they select the strategy with highest VOC for that problem type. In short, then, one place that value enters into the use of heuristics is at the very start: The choice to employ a heuristic strategy is itself a value-based decision aimed at maximizing accuracy and speed.

Another place that value enters into heuristic decision-making is in the application of a heuristic. Suppose the decision-maker selects the priority heuristic for the specific decision problem at hand. Value representations are required at a number of points to apply this procedure.

For example, the priority heuristic requires that the agent must select the most “important” dimension, and importance here is defined in valuational terms. That is, the most important dimension is the one with the highest “predictive validity”, that is the one that most distinguishes the options in terms of their respective values. For example, in deciding between a Jeep versus a minivan, for many people, safety emerges as an important dimension while fuel efficiency does not. This is because safety distinguishes the options substantially in terms of value (Jeeps are uniquely prone to rollovers) but fuel efficiency does not (the vehicles get similar mileage).

Having selected an important dimension, the agent must next assess which option is better along that dimension and by how much. For example, minivans have a slightly lower tendency to roll over than Jeeps. Does that make a minivan moderately better, substantially better, or decisively better than a Jeep? That depends on how much the agent values this increment in safety.

Value also appears in another critical place in the operation of heuristics: in the product of the heuristic.

Consider an agent who initially knows nothing about which of two options, **A** or **B**, is better. They then use the priority heuristic and establish that **A** is much better than **B** with respect to some highly important attribute. Clearly, the agent gains important evidence that **A** is better than **B** *overall*. That is, though they have looked at just one dimension, the agent nonetheless gains evidence about the overall value of **A** relative to **B**. The strength of this evidence will of course be higher in some contexts than others (e.g., higher in non-compensatory environments). Furthermore, the evidence is defeasible—the value of **B** might in fact be higher than **A** when comprehensively aggregating across all attributes. The key point, however, is that heuristics clearly do provide evidence, indeed often very solid evidence, about value.

This point is strengthened by considering the finding, discussed earlier, that agents flexibly employ specific decision-making strategies that are well-suited to the specific problem-types they encounter [45–47]. That is, agents select heuristics that tend to produce reasonably accurate information about which option is overall better while expending

much less effort than required for comprehensive aggregation. For example, in certain situations, the priority heuristic strategy, i.e., comparing options **A** and **B** on one specific highly important dimension, tends to yield reasonably accurate information about which option is better overall, and agents tend to employ the priority heuristic in these situations. On the other hand, where heuristics perform poorly and do not deliver accurate information about which option is overall better, approaches that rely on more exhaustive multiattribute aggregation are instead employed.

These two observations help to underscore that heuristics are not attempts to bypass value, which is how it initially seemed and which is what Hayden and Niv had argued. Instead, heuristics are attempts to acquire reasonably accurate information about value with greater ease. An agent using the priority heuristic can figure out that one option is very likely to be overall higher in value much more quickly and with much less effort. Value is not bypassed; it is instead ascertained much more cheaply.

It bears emphasis that heuristics typically do not provide information about value in a magnitude-type format. For example, depending on the details of exactly how the heuristic is applied, the priority heuristic supplies information about which option is likely to be *better* or *best* overall, and thus it supplies information about value in a relative value or maximum value-type format [26,27]. If a critic assumes (incorrectly) that the valuational view is committed to the claim that all value information is supplied in a magnitude-type format, then widespread use of the priority heuristic would indeed present a problem. But as was pointed out in Part 2, even though value itself is a scalar variable, information about value can be supplied in a variety of formats, any one of which can enable the agent to select the option with the highest value. Heuristics for complex decisions, then, are best understood as examples of alternative pathway processes.

In sum, detailed examination of the priority heuristic reveals a number of places in which value plays a pivotal role: In the selection of the heuristic, in the application of the heuristic, and in the product of the heuristic. Nothing about these claims is unique to the priority heuristic; the claims apply to heuristic decision-making more broadly. It follows that heuristic decision-making does not bypass value after all, it instead richly employs it. Heuristic decision-making is thus not a problem for the neuroeconomic valuational view. Instead, the ubiquity of heuristic decision-making provides further support for the view.

4. Are Dissimilar Alternatives Compared in a Common Currency of Value?

The neuroeconomics valuational view claims that value serves as common currency for decision-making [8,9]. But it is useful to distinguish two importantly different common currency roles: options-level and systems-level. As we shall see, the first kind of common currency encounters limitations that the second is well positioned to address.

Options-Level Common Currency

People make decisions that trade off an astonishing diversity of goods, including food, money, safety, social gains, morally-valenced outcomes, and many others. A question arises as to how choices are made involving such highly disparate items.

Neuroeconomic theory provides a compelling answer: Disparate options are compared in a common medium of value, and the option with the highest value is selected. This approach treats value as an options-level common currency because value attaches directly to the options themselves (i.e., goods to be consumed, outcomes to be brought about, actions to be executed, and so forth). In contrast, in a systems-level common currency, which is discussed in the next section, value does not attach to the options directly.

Multiple lines of evidence support the view that value serves as an options-level common currency in human decision-making.

One line of evidence comes from neuroimaging studies in which people are presented rewards of various kinds and their brain activation patterns are quantified. These studies have identified several neural regions, most prominently ventral striatum and ventral medial prefrontal cortex, that exhibit a trio of features. First, their activation scales with the expected value of rewards such that larger expected values, for example, larger expected monetary gains, are associated with stronger activation [14,48–50]. Second, they activate to a very broad range of reward types, including food, attractive faces, monetary rewards, consumer items, social gains, and moral outcomes, among others [14,51–53]. Third, activation in these regions is predictive of exchange rates or actual choice patterns among these rewards [8,54,55]. These observations have led some theorists to conjecture that these regions encode a common currency of value that enables choice over disparate option types [8,9,56].

Another line of evidence comes from computational models of multiattribute decisions, such as decision field theory [57,58] and other sequential sampling-based models of preferential choice [59–61]. These models are often applied to decision problems involving options with highly dissimilar attributes, for example vehicles that differ on cost, safety, and style. To enable comparisons across dissimilar dimensions, these models make a basic assumption that the agent can directly make pairwise comparisons of how much better an option is versus another option along a specific attribute dimension, i.e., the agent can make attribute-wise relative value assessments. Moreover, it is assumed value information arising from these comparisons can be summed up across comparisons involving different attributes. In short, then, these models make the assumption that value serves as a common medium of comparison across dissimilar attributes and across options.

Importantly, over the last two decades, decision field theory and related sequential sampling models have enjoyed substantial empirical success in capturing a range of behavioral phenomena including patterns of choice and reaction times [59,61]. In addition, they provide the leading explanations for well-documented context effects such as the similarity effect, compromise effect, and decoy effect (discussed in Part 5). The success of these computational models of multiattribute decision-making provides strong evidence that their key assumption—that value serves as a common currency across dissimilar attributes and options—is in fact correct.

There is, in addition, a venerable theoretical argument, which might be dubbed the “How else?” argument, for why we observe widespread use of value as an options-level common currency in human decision-making (see [1,8,56,62,63] for somewhat different versions of this general argument).

Value provides a highly efficient and parsimonious medium in which to encode countless tradeoffs across dissimilar goods, and to adjust these tradeoffs based on ongoing experience via value learning algorithms (for example, RL algorithms). Without the use of value for this purpose, encoding tradeoffs between even just two goods becomes prohibitively complex (see [63] for a detailed argument and [1] for a compelling illustration). Moreover, there is the further problem of how to adjust these tradeoffs if the goods change in decision-relevant ways (for example, if a certain food item no longer yields certain anticipated benefits).

In short, one likely reason we observe value ubiquitously as an options-level common currency in human decision-making is that alternative methods to manage complex tradeoffs across dissimilar goods are simply too inefficient to be realistically workable.

Systems-level Common Currency

Even if value is widely used as an options-level common currency, there are principled reasons to expect limitations on this approach. There is widespread agreement that there are multiple systems in the human brain that contribute to choice and action guidance [36,37,64,65]. Examples include model-based RL systems, model-free RL systems, Pavlovian systems, and stimulus-response learning systems, among others.

The presence of multiple action guidance systems in the human brain raises the question of what happens if more than one system is concurrently engaged, and they produce divergent outputs. This question becomes especially pressing when the systems utilize very different algorithms, and their outputs are in very different formats.

Consider an example, similar to one suggested by Hayden and Niv [31], in which there is competition between a model-based system and policy gradient-based system. The model-based system uses a scalar value-based RL algorithm that favors selecting option **A**, while the policy gradient algorithm uses a gradient-based method to learn an optimal policy, and it favors selecting option **B**. These systems operate by very different

computational principles and the value information that they supply, magnitude-type information for the former and maximum value-type information for the latter, are in different formats. As such, there does not appear to be a single options-level common currency in which the options can be compared, and the one with the highest value selected.

Noncomparability of this type presents a potentially serious problem for the neuroeconomic valuational view, which is committed to the idea that value serves as a common currency for decision-making.

Addressing this challenging requires distinguishing two different kinds of valuation, options-level and systems-level, which yield two correspondingly different kinds of common currencies. In options-level valuation, which is the type of valuation that was assumed in the previous subsection, the agent assigns value directly to the bearers of value, for example goods to be consumed, outcomes to be brought about, or actions to be executed, and so forth.

In systems-level valuation, the agent assigns value to competing systems, not as direct bearers of value themselves, but rather as *recommendation sources*. More specifically, the agent assembles information about the expected value of relying on the recommendations of the respective systems, using indicators such as their observed prior accuracy, speed in delivering recommendations, cognitive demands, and so forth. This information is aggregated and the system with overall higher expected value as a recommendation source is selected, and its recommendation is executed. In short, instead of asking a first-order question of which option is better, the agent asks a second-order question of which system is it better to follow.

There is considerable evidence for systems-level valuation in human decision-making, which will be reviewed shortly. But a key point, which appears to have been underappreciated, is that this type of valuation allows for an elegant solution for the noncomparability problem.

Recall the earlier example of competition between a model-based RL system and a policy gradient-based RL system. Because the respective systems operate with different computational principles and supply value information in different formats, an options-level common currency was not available. Importantly, however, the presence of this kind of first-order noncomparability only hampers using value to answer the first-order question of which option is better. It is no bar to asking the second-order question of which system is it better to follow. The agent can answer this higher-level question leveraging indicators such as the respective systems' prior accuracy, speed, cognitive demands, etc. In short, even in the absence of a first-order common currency, value can still serve as a second-order common currency allowing adjudication among competing systems in a common medium of value.

The preceding solution for the noncomparability problem is not merely theoretical. There is extensive evidence that systems-level valuation is in fact widespread across multiple domains of human cognition.

A number of researchers investigated competition between model-based and model-free reinforcement learning, identifying a set of factors that tend to promote use of each of the respective systems [66–69]. These factors have been integrated in a cost/benefit model [70] in which model-based control is favored when its benefits in terms of accuracy outweigh its costs in terms of slower speed and opportunity cost in consuming limited cognitive resources.

A similar model has been put forward for explaining competition between controlled and automatic processing. In the *expected value of control* framework from Shenhav, Cohen, and colleagues [22,23,71], controlled processing is favored by cues indicating higher expected value for exercising control, which include errors on previous trials, surprising outcomes, the presence of response conflict, and incentives for accurate performance.

Another place where we see systems-level valuation is, arguably, in the strategy selection phase of heuristic decision-making, which was discussed in Part 3. In the model from Leider and Griffiths [47], agents select among candidate decision strategies by learning the expected value of these strategies for each problem context. If we interpret “system” expansively as any process that delivers recommendations on what to do, heuristic decision strategies, such as the priority heuristic and elimination by aspects heuristic, might be thought of as instances of “systems” that are selected on the basis of systems-level valuation.

In sum, there are thus at least two different levels at which a common currency can operate, and both are widely used in human decision-making. At the first level, dissimilar options (i.e., entities that are the direct bearers of value) compete in an options-level common currency of value to resolve the question of which option is best. At the second level, dissimilar systems compete in a systems-level common currency of value to resolve the question of which system is it best to follow. In both cases, the same underlying strategy is deployed: a common medium of value is used to resolve competition among dissimilar alternatives.

5. Context Effects and Preference Reversals: Can the Valuation Model Accommodate “Irrational” Choice?

The neuroeconomic valuation model says that people make decisions based on information about the value of the options, selecting the option with highest value. Some theorists suggest that models of this type, which are inspired by economic theory, cannot

accommodate observations of “irrational” decision-making such as context effects and preference reversals [24,26,72,73].

For example, consider a decision problem that produces the so-called “similarity effect” [74]. Suppose job candidate **A** is strong on personality and weak on work ethic while **B** is the reverse, and people generally prefer each roughly equally. Adding a third job candidate **S** that is very similar to **A** to the option set (slightly better than **A** on personality, slightly worse than **A** on work ethic) tends to decrease the proportion of choices for **A** [74].

Some theorists contend that the similarity effect and other preference reversals are a problem for valuational models that posit a common currency of value. For example, Walsek and Brown argue that, “Context-induced changes in attribute weightings are ... hard to understand if there is a single common currency of value—because the existence of such a currency mandates a fixed trade-off between attribute values.” They add, “Such inconsistencies are typically taken as problematic, because utility functions cannot be easily constructed to accommodate systematic preference reversals.”

Walsek and Brown’s critique, however, appears to rest on a misconception because neuroeconomic valuational models can indeed accommodate context effects, and, moreover, they provide detailed mechanistic explanations for how these effects arise.

Neuroeconomic valuational models say that people use information about value to select the option with highest value. But, critically, this value information is not somehow already present and ready at the time it is needed. Rather, it must be actively *assembled*: Individual informational elements need to be accessed, and in many cases, they must be tabulated, summed, or inferentially elaborated. It is only after this active assembly stage is complete that the relevant value information is ready for the purposes of informing choice (see “Stage 1” in Figure 1).

In some cases, this assembly stage for value information unfolds rapidly and fairly directly, as happens with cached value representations furnished by model-free RL algorithms [66]. But in other cases, such as sequential sampling multiattribute models (see Part 4) and heuristic decision-making (see Part 3), assembling value information unfolds over time and relies on factors such as attention, memory retrieval, pairwise comparisons, noisy evidence accumulation, probability estimation, and so forth. Importantly, all of these factors can influence the ongoing assembly of value information in context sensitive ways, which in turn can influence subsequent choices.

Multi-attribute decision field theory (MDFT) provides a computationally precise illustration of this phenomenon [58], though these points apply generally to a number of related sequential sampling-type models [59–61]. In the MDFT model, the agent focuses, via a stochastic attention process, on one of the attributes of the options and compares each pair of options along this attribute dimension. In the job candidate example discussed earlier, if the attribute currently in focus were work ethic, **B** would strongly beat **A** and **S**,

while if the attribute in focus were personality, **A** and **S** would strongly beat **B**. As attention wanders from attribute to attribute, value information arising from these pairwise comparisons accumulates for each option over time under conditions of noise. When a prespecified confidence threshold is reached, the agent selects the associated option.

Importantly, the MDFT model provides a natural explanation for the similarity effect in terms of division of attentional advantages. When attention focuses more on the work ethic dimension, **B** uniquely gains an advantage, but when attention focuses more on the personality dimension, *both* **A** and **S** gain an advantage, roughly splitting the “wins” between them. Thus, the addition of **S** to the option set tends to leave the proportion of choices of **B** unchanged while choices for **A** are halved—precisely the contextual effect that is observed in experimental settings.

For our purposes, the key point is that MDFT explains context effects in terms of the specific factors that operate during the assembly stage for value information, for example, stochastic attention, pairwise value comparisons, noisy evidence accumulation, etc. Agents who choose their options based on this type of value assembly process are necessarily susceptible to context effects. Thus, contrary to what the critics suggested, context effects are not after all incompatible with neuroeconomic valuational models, and, instead, these effects are in fact predicted by and explained by these models.

Stepping back a bit, the critics appear to have directed their charge against the wrong target. Their target is not neuroeconomic valuational models, but rather standard expected utility models of the kind associated with John von Neumann & Oskar Morgenstern [75] and Leonard Savage [76], among others. Standard expected utility theory treats value as a means to represent patterns of preferences over options. That is, the theory starts with an agent who displays a certain pattern of preferences over pairwise options and, if these preferences satisfy a set of consistency conditions specified in the axioms of the theory, then a utility function is next constructed to represent this preference structure. Context effects are a decisive problem for standard expected utility theory because inconsistent preference patterns violate these axiomatic consistency conditions and cannot be represented by a utility function, just as the critics allege.

In neuroeconomic valuational models, in contrast, value does not play this post hoc representational role. Instead, value information serves as an intermediate variable that is constructed during the decision-making process, and it is in turn causally implicated in choice. That is, agents actively assemble value information as a basis to choose; value is not merely ascribed to them to reflect their choices. This active, assembly process relies on factors such as attention, memory retrieval, pairwise comparisons, noisy evidence accumulation, etc., all of which can contribute to contextually sensitive modulation of value information, which in turn can lead to contextually sensitive choices.

Neuroeconomic valuational models are thus compatible with the presence of context effects and preference reversals, and indeed are ideally positioned to provide detailed mechanistic explanations for how these effects arise.

Conclusion

Neuroeconomics is an influential research program in cognitive science inspired by economic theory that puts forward mathematically precise mechanistic models of decision and action guidance. These models are supported by behavioral, chronometric, computational, and neural methods, and they have found application across diverse domains of cognition.

In this essay, neuroeconomic valuational models, defined in terms of two core theses, were defended from several influential criticisms. The essence of the response was that the criticisms are focused on a relatively simple model of valuation that admittedly encounters significant problems. Over the course of this essay, a more sophisticated model of valuation was elaborated and defended. This revised view retains the two core theses that are the hallmarks of the neuroeconomic research program, but it clarifies and amends these theses in important ways.

The first core thesis is that people make decisions about what to do based on information about the expected value of their options, and they choose the option with highest value. But in the more sophisticated view of valuation defended here, it was clarified that information about value can be supplied in a variety of formats, including non-magnitude type formats. This clarification is critical to understanding how certain algorithms and processes, including policy gradient RL algorithms and heuristic decision-making strategies, are fully accounted for within neuroeconomic theory.

The second core thesis is that value serves as a common currency for choice and action guidance. But a distinction was made between an options-level common currency, in which the targets of valuation are dissimilar options, and a systems-level common currency, in which the targets of valuation are systems as recommendation sources. This clarification addressed critical “noncomparability” cases in which an options-level common currency is not available; competition between dissimilar systems can nonetheless still be resolved through a systems-level common currency.

Overall, in adopting a more nuanced picture that includes multiple formats for value information and two levels of valuation, we achieve a more powerful and psychologically realistic version of the neuroeconomic valuational view.

A final point is that the arguments of this essay further underscore the generality of neuroeconomic valuational view—this view applies very broadly across a vast array settings and domains.

For example, the neuroeconomic valuational view applies to choices in which value information is supplied by diverse kinds of RL algorithms, including policy gradient algorithms which are sometimes described as “value free”, and there is strong evidence

that diverse RL algorithms are operative across multiple psychological domains. The valuational view additionally applies to various kinds of deliberative decision-making, which includes exhaustive multiattribute aggregation but also various kinds of heuristic decision-making strategies. Furthermore, the valuational view applies to cases in which disparate options are directly traded off. But it also applies to more complex situations when dissimilar systems compete—such competition is also resolved in a common currency of systems-level value. Finally, the valuational model extends to “irrational” choices involving contextual effects and preference reversals.

Overall, the more nuanced neuroeconomic valuational view put forward and defended in this essay has grand aspirations to be a truly general approach to understanding human decision-making across diverse psychological domains.

References

1. Berridge, K.C. (2004) Motivation concepts in behavioral neuroscience. *Physiol. Behav.* 81, 179–209
2. (1963) Operant behavior. *Am. Psychol.* 18, 503
3. Hinde, R.A. (1956) Ethological models and the concept of ‘drive.’ *Br. J. Philos. Sci.* 6, 321–331
4. Montague, P.R. and Berns, G.S. (2002) Neural economics and the biological substrates of valuation. *Neuron* 36, 265–284
5. Glimcher, P.W. and Rustichini, A. (2004) Neuroeconomics: The Consilience of Brain and Decision. *Science* 306, 447–452
6. Sanfey, A.G. *et al.* (2006) Neuroeconomics: cross-currents in research on decision-making. *Trends Cogn. Sci.* 10, 108–116
7. Glimcher, P.W. and Fehr, E. (2013) *Neuroeconomics: Decision making and the brain*, Academic Press
8. Levy, D.J. and Glimcher, P.W. (2011) Comparing apples and oranges: using reward-specific and reward-general subjective value representation in the brain. *J. Neurosci.* 31, 14693–14707
9. Levy, D.J. and Glimcher, P.W. (2012) The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* 22, 1027–1038
10. Sugrue, L.P. *et al.* (2005) Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* 6, 363–375
11. Rangel, A. *et al.* (2008) A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci* 9, 545–56
12. Montague, P.R. *et al.* (2006) Imaging Valuation Models in Human Choice. *Annu. Rev. Neurosci.* 29, 417–448
13. Fehr, E. and Camerer, C.F. (2007) Social neuroeconomics: the neural circuitry of social preferences. *Trends Cogn. Sci.* 11, 419–427
14. Bartra, O. *et al.* (2013) The valuation system: a coordinate-based meta-analysis of BOLD fMRI experiments examining neural correlates of subjective value. *Neuroimage* 76, 412–427
15. Theeuwes, J. *et al.* (2022) What to expect where and when: how statistical learning drives visual selection. *Trends Cogn. Sci.*
16. Anderson, B.A. (2019) Neurobiology of value-driven attention. *Curr. Opin. Psychol.* 29, 27–33
17. Failing, M. and Theeuwes, J. (2018) Selection history: How reward modulates selectivity of visual attention. *Psychon. Bull. Rev.* 25, 514–538
18. Anderson, B.A. and Kim, H. (2018) Mechanisms of value-learning in the guidance of spatial attention. *Cognition* 178, 26–36
19. Dolan, R.J. and Dayan, P. (2013) Goals and Habits in the Brain. *Neuron* 80, 312–325
20. D’Ardenne, K. *et al.* (2012) Role of prefrontal cortex and the midbrain dopamine system in working memory updating. *Proc. Natl. Acad. Sci.* 109, 19900–19909

21. Hazy, T.E. et al. (2006) Banishing the homunculus: Making working memory work. *Neuroscience* 139, 105–118
22. Shenhav, A. et al. (2013) The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron* 79, 217–240
23. Shenhav, A. et al. (2016) Dorsal anterior cingulate cortex and the value of control. *Nat. Neurosci.* 19, 1286
24. Walasek, L. and Brown, G.D. (2023) Incomparability and incommensurability in choice: no common currency of value? *Perspect. Psychol. Sci.*
25. Bennett, D. et al. (2021) Value-free reinforcement learning: policy optimization as a minimal model of operant behavior. *Curr. Opin. Behav. Sci.* 41, 114–121
26. Vlaev, I. (2018) Local choices: rationality and the contextuality of decision-making. *Brain Sci.* 8, 8
27. Vlaev, I. et al. (2011) Does the brain calculate value? *Trends Cogn. Sci.* 15, 546–554
28. Spurrett, D. (2009) Need there be a common currency for decision-making? *South Afr. J. Philos.* 28, 210–221
29. Lichtenstein, S. and Slovic, P. (2006) *The Construction of Preference*, Cambridge University Press
30. Slovic, P. (1995) The construction of preference. *Am. Psychol.* 50, 364
31. Hayden, B.Y. and Niv, Y. (2021) The case against economic values in the orbitofrontal cortex (or anywhere else in the brain). *Behav. Neurosci.* 135, 192
32. Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning: An Introduction*, A Bradford Book
33. Sutton, R.S. et al. (1999) Policy gradient methods for reinforcement learning with function approximation. *Adv. Neural Inf. Process. Syst.* 12
34. Tversky, A. and Kahneman, D. (1974) Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 1124–1131
35. Gigerenzer, G. and Todd, P.M. (1999) Fast and frugal heuristics: The adaptive toolbox. In *Simple heuristics that make us smart*, pp. 3–34, Oxford University Press
36. Sanfey, A.G. and Chang, L.J. (2008) Multiple systems in decision making. *Ann. N. Y. Acad. Sci.* 1128, 53–62
37. Van Der Meer, M. et al. (2012) Information processing in decision-making systems. *The Neuroscientist* 18, 342–359
38. Schultz, W. et al. (1997) A neural substrate of prediction and reward. *Science* 275, 1593–9
39. Joel, D. et al. (2002) Actor–critic models of the basal ganglia: New anatomical and computational perspectives. *Neural Netw.* 15, 535–547
40. Raiffa, H. (2006) Preferences for multi-attributed alternatives. *J. Multi-Criteria Decis. Anal.* 14
41. Gilovich, T. et al. (2002) *Heuristics and Biases: The Psychology of Intuitive Judgment*, Cambridge University Press
42. Kahneman, D. (2003) A perspective on judgment and choice: Mapping bounded rationality. *Am. Psychol.* 58, 697–720
43. Brandstätter, E. et al. (2006) The priority heuristic: making choices without trade-offs. *Psychol. Rev.* 113, 409

44. Tversky, A. (1972) Elimination by aspects: A theory of choice. *Psychol. Rev.* 79, 281–299
45. Payne, J.W. (1993) *The adaptive decision maker*, Cambridge University Press
46. Payne, J.W. et al. (1988) Adaptive strategy selection in decision making. *J. Exp. Psychol. Learn. Mem. Cogn.* 14, 534
47. Lieder, F. and Griffiths, T.L. (2017) Strategy selection as rational metareasoning. *Psychol. Rev.* 124, 762
48. Breiter, H.C. et al. (2001) Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30, 619–639
49. Knutson, B. et al. (2005) Distributed neural representation of expected value. *J. Neurosci.* 25, 4806–4812
50. Kable, J.W. and Glimcher, P.W. (2007) The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* 10, 1625–1633
51. FitzGerald, T.H. et al. (2009) The role of human orbitofrontal cortex in value comparison for incommensurable objects. *J. Neurosci.* 29, 8388–8395
52. Kim, H. et al. (2011) Overlapping responses for the expectation of juice and money rewards in human ventromedial prefrontal cortex. *Cereb. Cortex* 21, 769–776
53. Lin, A. et al. (2012) Social and monetary reward learning engage overlapping neural substrates. *Soc. Cogn. Affect. Neurosci.* 7, 274–281
54. Plassmann, H. et al. (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J. Neurosci.* 27, 9984–9988
55. Smith, D.V. et al. (2010) Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *J. Neurosci.* 30, 2490–2495
56. Kable, J.W. and Glimcher, P.W. (2009) The neurobiology of decision: consensus and controversy. *Neuron* 63, 733–745
57. Busemeyer, J.R. and Townsend, J.T. (1993) Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100, 432
58. Roe, R.M. et al. (2001) Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychol. Rev.* 108, 370
59. Trueblood, J.S. et al. (2014) The multiattribute linear ballistic accumulator model of context effects in multialternative choice. *Psychol. Rev.* 121, 179
60. Howes, A. et al. (2016) Why contextual preference reversals maximize expected value. *Psychol. Rev.* 123, 368
61. Busemeyer, J.R. et al. (2019) Cognitive and neural bases of multi-attribute, multi-alternative, value-based decisions. *Trends Cogn. Sci.* 23, 251–263
62. McFarland, D.J. and Sibly, R.M. (1975) The behavioural final common path. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 270, 265–293
63. Spurrett, D. (2021) The descent of preferences. *Br. J. Philos. Sci.*
64. Frank, M.J. et al. (2009) Multiple systems in decision making: A neurocomputational perspective. *Curr. Dir. Psychol. Sci.* 18, 73–77
65. Daw, N.D. and O’Doherty, J.P. (2014) Multiple systems for value learning. In *Neuroeconomics*, pp. 393–410, Elsevier
66. Daw, N.D. et al. (2005) Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704–1711

67. Keramati, M. *et al.* (2011) Speed/accuracy trade-off between the habitual and the goal-directed processes. *PLoS Comput. Biol.* 7, e1002055
68. Lee, S.W. *et al.* (2014) Neural computations underlying arbitration between model-based and model-free learning. *Neuron* 81, 687–699
69. Patzelt, E.H. *et al.* (2019) Incentives boost model-based control across a range of severity on several psychiatric constructs. *Biol. Psychiatry* 85, 425–433
70. Kool, W. *et al.* (2017) Cost-benefit arbitration between multiple reinforcement-learning systems. *Psychol. Sci.* 28, 1321–1333
71. Kool, W. *et al.* (2017) Cognitive Control as Cost-Benefit Decision Making. In *The Wiley Handbook of Cognitive Control*, pp. 167–189, Wiley-Blackwell
72. Chater, N. and Vlaev, I. (2011) The instability of value. *Atten. Perform.* 23, 81
73. Tversky, A. and Simonson, I. (1993) Context-dependent preferences. *Manag. Sci.* 39, 1179–1189
74. Huber, J. *et al.* (1982) Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *J. Consum. Res.* 9, 90–98
75. Von Neumann, J. and Morgenstern, O. (1947) *Theory of games and economic behavior*, 2nd rev, Princeton University Press
76. Savage, L.J. (1972) *The foundations of statistics*, Courier Corporation