Multi-parameter utility and drift-rate functions conflate attribute weights and choice consistency

Ian Krajbich*

The Ohio State University

Department of Psychology

Department of Economics

*Correspondence to krajbich.1@osu.edu

Abstract

Standard decision models include two components: subjective-value (utility) functions and stochastic choice rules. The first establishes the relative weighting of the attributes or dimensions and the second determines how consistently the higher utility option is chosen. For a decision problem with M attributes, researchers often estimate $M-1$ utility parameters and separately estimate a choice-consistency parameter. Instead, researchers sometimes estimate M parameters in the utility function and neglect choice consistency. I argue that while these two approaches are mathematically identical, the latter conflates utility and consistency parameters, leading to ambiguous interpretations and conclusions. At the same time, behavior arises from the interaction of utility and consistency parameters, so for choice prediction they should not be considered in isolation. Overall, I advocate for a clear separation between utility functions and stochastic choice rules when modeling decision-making, and reinforce the notion that researchers should use $M-1$ parameters for M -attribute decision problems.

Introduction

In recent years, researchers have been increasingly using mathematical models to study choice behavior. Static models attempt to explain choice outcomes, while sequential sampling models such as the drift diffusion models (DDM), attempt to explain both choice outcomes and reaction times (RT). Both approaches model the tradeoffs between attributes of the choice problem, either with a utility/subjective-value function in the static case, or with a drift-rate function in the dynamic case¹.

Nearly all choice problems involve tradeoffs between attributes². What determines a person's preference is their exchange rate between the attributes, i.e. how much they weigh one attribute compared to the other(s). For example, a person may like oranges twice as much as they like apples. In that case, the simplest utility or drift-rate function to describe that person would be $v(o, a) = 2o + a$, where o denotes the number of oranges and a denotes the number of apples. Importantly, utility and drift-rate are both *ordinal* rather than *cardinal*, and are invariant to affine transformations. In other words, the same preferences could be described by $v(o, a) = 4o + 2a$, $v(o, a) = 2o + a + 1$, or more generally $v(o, a) = b(2o + a) + c$ where b and c are arbitrary constants.

Some recent papers using DDMs have used two-parameter drift-rate functions in two-attribute tasks. For example, Hutcherson et al., (2015) use a drift function $v = w_{self}(\$Self) + w_{other}(\$Other)$ to model social preferences: allocating money between oneself and another; a follow-up paper uses the same function but adds an additional term for inequality (Teoh et al., 2020). Another article used a drift function $v = s(Health) + \omega(Taste)$ to model food choice, weighing health considerations against taste (Enax et al., 2016). Yet another article used a drift function $v = \delta_A(Amount) + \delta_T(Time)$ to model intertemporal choices, weighing the amount of money against the time to receive it (Amasino et al., 2019).

¹ In the DDM, the drift rate determines how evidence accumulates over time towards one of two possible decision boundaries.

² One notable exception is learning, where the tradeoff is instead between exploring and exploiting.

Here I argue that using such formulations can lead to problematic interpretations of the parameters. The formulations above are not incorrect, but because of the way they are written, these functions entice the reader/analyst to treat a one-dimensional problem as if it was two-dimensional (or more generally an $M-1$ dimensional problem as if it was M dimensional). In other words, these tasks have a single tradeoff and yet the models have two weight parameters. Just as one cannot answer the question “how much do you value money?” without asking “relative to what?”, one cannot talk about the weights on both *self* and *other* payoffs (only *self* relative to *other*), the weights on *money* and *time* (only *money* relative to *time*), or the weights on *taste* and *health* (only *taste* relative to *health*), without something else to compare them to.³ In these M -parameter formulations, the extra parameter comes from eliminating the choice consistency parameter from the stochastic choice function (typically the inverse-temperature parameter in a logit choice function, or the drift-rate scaling factor in the DDM). Importantly, the tradeoff between attributes should be independent from the consistency with which a person chooses in line with their utility or drift-rate. In theory, one can measure the tradeoff between attributes by asking for the exchange rate directly, for example, “how many apples do you require in order to give up an orange?”

The problem that can arise from the M -parameter utility function, as opposed to the $M-1$ parameter version, is that all M parameters then reflect choice consistency. When comparing parameter values across individuals or groups, researchers may come to misleading conclusions. For example, a researcher studying age differences in health vs. taste considerations might find larger weights on health in older individuals and conclude that they are more health conscious, when in reality they are simply more consistent in what they eat (reflected by larger weights on taste as well). To examine the importance of health, the researcher would instead need to examine the ratio of the health weight to the taste weight.

³ This issue has also appeared in the literature on model-based (MB) vs. model-free (MF) reinforcement learning. In early years, this tradeoff was modeled as $w * Q_{MB} + (1 - w) * Q_{MF}$, but recently some researchers have switched to modeling it as $\beta_{MB} * Q_{MB} + \beta_{MF} * Q_{MF}$ (Decker et al., 2016). Here Q refers to the learned value of an alternative, based either on the MB or MF system.

This is not an idle concern. For example, Teoh et al. (2020) report that a larger weight on others' outcomes negatively correlates with selfish gaze biases. However, because they used an M -parameter utility function, the weight on others' outcomes also reflects choice consistency. Thus, an alternative possibility is that individuals with larger selfish gaze biases are less consistent in their choices. Similarly, the authors also conclude that time pressure increases the weights on one's own outcomes and on fairness, but an alternative possibility is that time pressure increases choice consistency. The use of M -parameter utility functions can clearly complicate the interpretation of modeling results.

Below I outline the basic ambiguity involved in simultaneously estimating preferences and choice consistency, and I propose a simple rule when using such approaches. In particular, I argue that one should use $M-1$ parameters for an M -attribute decision problem. At the same time, researchers need to keep in mind that utility and drift-rate function parameters interact with choice consistency parameters to determine choice outcomes and RTs. Thus, the decision of whether to analyze the decision parameters in isolation, or their interaction, depends on whether the researcher is interested in the underlying preferences or the resulting behavior.

Utility function + Logit

Despite the power gained by leveraging RT data to enhance our understanding of behavior, it is important to keep in mind that many of the limitations of the utility-function approach also apply to the DDM. Utility functions allow decision analysts to estimate how people trade off different attributes against each other; they establish an "exchange rate" between the attributes in play⁴. The practical implication of this fact is that there is a fixed number of parameters that a utility function can have. For a decision problem with M attributes, the utility function should have at most $M-1$ free parameters⁵. For each additional attribute beyond the first, we need exactly 1 parameter to establish that attribute's exchange rate with the first.

⁴ Here I am assuming a linear utility function; for more complex functions the problem is more nuanced.

⁵ This assumes no intercept in the utility function. If there is an intercept, that adds one additional parameter. However, this should not affect any arguments in the paper.

Now, when we model real choices, we know that individuals do not always choose in line with their utility function. If they did, the arguments in this paper would be trivial. If an individual chose perfectly consistently, it would be impossible to identify more than $M-1$ parameters. Instead, researchers often use a logit function to model how consistently individuals choose in line with their utility function.

Take a simple example from the social preference literature, where we are interested in how much a person weighs their own payoff relative to another person's payoff. We can write their utility function as:

$$U(x_i, x_j) = x_i + \varphi x_j$$

where x_i denotes the decision-maker's payoff, x_j denotes the other's payoff, and φ is the exchange rate between self and other⁶.

We can then write their logit choice function for accepting a particular allocation of money between self and other as:

$$P(\text{accept}) = \frac{1}{1 + e^{-\lambda(x_i + \varphi x_j)}}$$

The logit function itself has an "inverse temperature" parameter (λ) that captures the decision maker's consistency (i.e. how often they choose in line with their utility function), as a function of the utility.⁷ However, we can use simple algebra to distribute λ , yielding:

⁶ Note that the choice of where to put φ is arbitrary; we could just as easily have put it in front of x_i .

⁷ In keeping with the idea that decision makers always maximize utility, λ is meant to capture the weight on the attributes in the modeled utility function, relative to the attributes that are not modeled in the utility function. This means that the value of λ depends both on the units of x_i , but also on which attributes appear in the utility function. The more predictive the modeled attributes, the higher λ should be. In the limit where the utility function perfectly explains all of a subject's choices, $\lambda = \infty$. In this scenario one can see a clear advantage for the

$$P(\text{accept}) = \frac{1}{1 + e^{-\lambda x_i - \lambda \varphi x_j}}$$

We can define $\delta = \lambda \varphi$, leaving us with what looks like a two-parameter utility function:

$$U(x_i, x_j) = \lambda x_i + \delta x_j$$

This is merely an illusion. To give our utility function two parameters, we have sacrificed our ability to estimate a separate inverse temperature parameter. In other words, λ is not really a weight on x_i ; it is just the inverse temperature parameter in disguise. Moreover, the weight on x_j now also contains the inverse temperature parameter.

DDM analogue

When we move to the DDM, we are in a very similar situation. The analogy to the logit example is particularly compelling when one considers that the DDM yields a logit choice function (Webb, 2018).

As with a utility function + logit implementation, a drift rate function has two components: the relative weights of the attributes and the choice consistency (i.e. the rate at which the person chooses in the direction of the drift).

Consider the example above. To implement a DDM for this type of decision we would typically fix the within-trial noise parameter at $\sigma = 0.1$, estimate the boundary separation a , and then also estimate a drift function:

$$v = \gamma * (x_i + \varphi x_j)$$

one-parameter utility function over the two-parameter utility function, since the former would still allow us to estimate φ , while the latter would yield $\lambda = \delta = \infty$, telling us nothing about the exchange rate between x_i and x_j .

where γ is a scaling factor that determines how quickly/consistently the decision variable hits the correct boundary for a given decision problem ($x_i + \varphi x_j$) and boundary separation a .⁸

So, what we are left with is a drift function that looks almost exactly like what we had in the logit equation above. Just like in that case, we can distribute the γ parameter through the equation and we're left with

$$v = \gamma x_i + \gamma \varphi x_j$$

Similar to before, we can define $\beta = \gamma \varphi$ which leaves us with what looks like a two-parameter drift function:

$$v = \gamma x_i + \beta x_j$$

And as before, this is merely an illusion. To give our drift function two parameters, we sacrificed our ability to estimate the scaling parameter. γ here is not a weight on x_i it is just the scaling parameter in disguise. Similarly, the weight on x_j contains the scaling parameter too.

One can unveil this illusion even more clearly by noting that the DDM is fundamentally not fully identifiable. That is, the units of “evidence” in the model are arbitrary, which means that we are forced/free to set one of the parameters $\{\gamma, a, \sigma\}$ equal to a constant. In this case we chose to follow the convention of setting $\sigma = 0.1$ (Ratcliff, 1978), but we could just as easily have tried to fit σ . Had we tried to fit $\{a, \sigma\}$ as well as our two-parameter utility function with $\{\gamma, \beta\}$, we would be in a serious bind; there would be an infinite set of best-fitting parameters. For example, a model with $a = 1$, $\sigma = 1$, $\gamma = 1$, and $\beta = 0.5$ is identical to a model with $a = 2$, $\sigma = 2$, $\gamma = 2$, and $\beta = 1$. If we did not want to set a or σ to a constant, we would be forced to abandon our two-parameter drift function in favor of a single-parameter one.

⁸ Note that the estimation of the boundary separation can be done independently, based on the RT distribution when the drift rate is zero, i.e. when $x_i = -\varphi x_j$.

In other words, an arbitrary modeling choice on the part of the analyst can make it impossible to estimate an M parameter drift function for a decision problem with M attributes. As with the utility + logit approach, one can really only estimate $M-1$ drift-function parameters for an M attribute choice problem.

A simple solution for existing results with M -parameter functions

If a researcher has already run their analysis, or are using, for example, a logistic regression, they may be stuck with parameters on all the attributes. However, this does not mean that the results need to be analyzed or written up that way. Suppose the following model has been run:

$$v = \gamma x_i + \beta x_j$$

All one needs to do is to divide through by γ or β and this will separate out the scaling/temperature parameter from the preference component⁹:

$$v = \gamma(x_i + \frac{\beta}{\gamma} x_j)$$

Then one can report the preference parameter $\frac{\beta}{\gamma} = \varphi$, and separately, the scaling/temperature parameter γ . Note that one needs to know both parameters in order to quantify how much an increase in x_j or x_i will change behavior.

Conclusions

When researchers use utility functions or drift functions with M parameters for M -attribute problems, they are not wrong, but they have to recognize that those parameters reflect a mix of relative weights between attributes and overall choice consistency. Thus, those parameters

⁹ One might wonder whether it matters whether you divide through by γ or β . The answer is that globally it does not matter, but locally one does need to be consistent about it within a study. An easy way to see this is to again think of the preference parameter as an exchange rate: think of x_i as being money in USD and x_j as being money in Euros. Because drift rates and exponents are unitless this means that γ is in units of USD⁻¹ and β is in units of Euros⁻¹. If you factor out γ then you are using USD as your units within the parentheses and your preference parameter $\frac{\beta}{\gamma}$ is in units of USD per Euro. On the other hand, if you factor out β then you are using Euros as your units within the parentheses and your preference parameter $\frac{\gamma}{\beta}$ is in units of Euros per USD. Similarly, in the preference domain, one can talk about tradeoffs between *money* and *time* as \$ per day or days per \$, or between *taste* and *health* as lip-licks per calorie or calories per lip-lick.

cannot simply be thought of as weights on attributes. So, to reduce ambiguity, I recommend sticking to the simple rule: when you have M attributes, only fit $M-1$ utility or drift-rate function parameters. If a researcher is more interested in behavior than in the underlying preferences, they can always additionally consider the interaction between the utility/drift parameters and the choice consistency parameter.

Acknowledgments

Thanks to Dianna Amasino, Scott Huettel, Nicolette Sullivan for multiple rounds of comments on earlier drafts, and Cendri Hutcherson, Antonio Rangel, Bernd Weber, Todd Hare, Fadong Chen, and Arkady Konovalov for comments and conversations. Thanks also to the Cattell Sabbatical Fund and the National Science Foundation Career Award 1554837 .

Appendix

Relationship between inverse temperature and DDM parameters

In the DDM, there is a fundamental identification problem where we can only fit two out of three of the parameters: drift multiplier, boundary separation, and within-trial noise. In the standard formulation of the DDM, people assume the within-trial noise parameter is constant across conditions. After accounting for non-decision time, the boundary separation parameter can be estimated purely from RT, by looking at the RT distribution for trials where drift rate is zero (i.e. where the subject is indifferent). The drift multiplier can then be estimated based on the accuracy data, conditional on the estimated boundary separation. So, in the standard DDM formulation, the drift multiplier (multiplied by boundary separation) is analogous to the inverse temperature parameter in the logit. One can derive similar relationships for other formulations. For example, if we set the boundary separation constant across conditions (as in the aDDM; Krajbich, Armel, & Rangel, 2010), then it is the within-trial noise parameter that can be determined from the RTs (in trials where the drift rate is zero), and the drift multiplier (divided by the noise) again is analogous to the inverse temperature in the logit. Note that all of this assumes no across-trial variability in starting point, drift rate, or non-decision time; parameters that are often also fit in DDM studies.

Models with attributes that enter at different times (i.e. latency models)

Recently some papers have introduced the idea that attributes enter the drift rate at different points in time (Maier et al., 2020). For instance, Sullivan et al. (2014) argue that in food decisions, tastiness is considered earlier than healthiness. Similarly, Amasino et al. (2019) argue that in temporal discounting decisions, patient subjects consider amounts earlier than delays. The question is whether such models are subject to the same criticism presented in this paper. The answer is yes.

First consider the simpler modeling exercise using just the utility + logit approach discussed earlier. For decisions that terminate before the latency for the second attribute, one would estimate the following logistic regression:

$$P(\text{accept}) = \frac{1}{1 + e^{-\lambda(x_i)}}$$

Here λ would represent how often the subject chooses “accept” when x_i is positive. Essentially, λ here captures the alignment between choice and attribute x_i when not considering the other attribute(s).

Separately, on decisions that terminate after the latency for the second attribute, one would estimate a different logistic regression:

$$P(\text{accept}) = \frac{1}{1 + e^{-\lambda(x_i + \varphi x_j)}}$$

Note that λ here would take on a different value than in the first logit, unless x_j is non-predictive of choice. In other words, incorporating additional information (i.e. the second attribute) will change how consistently the decision-maker chooses. Clearly, it would still be misleading to think of either λ as an absolute weight on x_i .

The DDM case is analogous. Up until the latency for the second attribute, one would model the drift rate as:

$$v = \gamma * (x_i)$$

Here γ would represent how quickly the decision variable approaches the choice threshold for a given value of x_i , ignoring the other attribute(s).

After the latency period, one would then estimate a different drift rate:

$$v = \gamma * (x_i + \varphi x_j)$$

As in the logit case, here γ would take on a different value than in the first case. Just like with the logit, incorporating more information should change the drift rates. As above, it would still be misleading to think of either γ value as an absolute weight on x_i .

When estimating these parameters in a DDM, Amasino et al. (2019) estimate a single γ under the implicit assumption that γ in the first case is equal to γ in the second case. This is a restriction in the model, which amounts to saying that considering additional information does not change choice consistency. It is possible that this kind of restriction may lead to inaccurate parameter estimates.

References

- Amasino, D. R., Sullivan, N. J., Kranton, R. E., & Huettel, S. A. (2019). Amount and time exert independent influences on intertemporal choice. *Nature Human Behaviour*, 3(4), 383. <https://doi.org/10.1038/s41562-019-0537-2>
- Decker, J. H., Otto, A. R., Daw, N. D., & Hartley, C. A. (2016). From Creatures of Habit to Goal-Directed Learners: Tracking the Developmental Emergence of Model-Based Reinforcement Learning. *Psychological Science*, 27(6), 848–858. <https://doi.org/10.1177/0956797616639301>
- Enax, L., Krajbich, I., & Weber, B. (2016). Salient nutrition labels increase the integration of health attributes in food decision-making. *Judgment and Decision Making*, 11(5), 460–471.
- Hutcherson, C. A., Bushong, B., & Rangel, A. (2015). A Neurocomputational Model of Altruistic Choice and Its Implications. *Neuron*, 87(2), 451–462. <https://doi.org/10.1016/j.neuron.2015.06.031>
- Maier, S. U., Raja Beharelle, A., Polanía, R., Ruff, C. C., & Hare, T. A. (2020). Dissociable mechanisms govern when and how strongly reward attributes affect decisions. *Nature Human Behaviour*, 4(9), 949–963. <https://doi.org/10.1038/s41562-020-0893-y>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59–108.
- Sullivan, N., Hutcherson, C., Harris, A., & Rangel, A. (2014). Dietary Self-Control Is Related to the Speed With Which Attributes of Healthfulness and Tastiness Are Processed. *Psychological Science*, 0956797614559543.

Teoh, Y. Y., Yao, Z., Cunningham, W. A., & Hutcherson, C. A. (2020). Attentional priorities drive effects of time pressure on altruistic choice. *Nature Communications*, 11(1), 3534.

<https://doi.org/10.1038/s41467-020-17326-x>

Webb, R. (2018). The (Neural) Dynamics of Stochastic Choice. *Management Science*, 65(1), 230–255. <https://doi.org/10.1287/mnsc.2017.2931>