


Advancing Research on Cognitive Processes in Social and Personality Psychology: A Hierarchical Drift Diffusion Model Primer

Social Psychological and
Personality Science
1-11
© The Author(s) 2017
Reprints and permission:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/1948550617703174
journals.sagepub.com/home/spp


David J. Johnson¹, Christopher J. Hopwood¹, Joseph Cesario¹,
and Timothy J. Pleskac²

Abstract

We provide a primer on a hierarchical extension of the drift diffusion model (DDM). This formal model of decisions is frequently used in the cognitive sciences but infrequently used in social and personality research. Recent advances in model estimation have overcome issues that previously made the hierarchical DDM impractical to implement. Using examples from two paradigms, the first-person shooter task and the flash gambling task, we demonstrate that the hierarchical DDM can provide novel insights into cognitive processes underlying decisions. Finally, we compare the DDM to dual-process models of decision-making. We hope this primer will provide researchers a new tool for investigating psychological processes.

Keywords

drift diffusion model, process model, first-person shooter task, flash gambling task

Although widespread in cognitive sciences, the use of sequential sampling models to understand decision-making is rare in our field. We demonstrate how a hierarchical version of the drift diffusion model (DDM; Ratcliff, 1978) can advance research on psychological processes for social and personality psychology. Although we illustrate the model with two tasks, we stress its relevance for any task where individuals make decisions between alternatives.

The DDM

The DDM is a sequential sampling model used primarily to explain binary decisions (see Forstmann, Ratcliff, & Wagenmakers, 2016; Klauer, 2014; Ratcliff & McKoon, 2008; Ratcliff, Smith, Brown, & McKoon, 2016). The key model assumption is that people repeatedly sample decision-relevant evidence from their environment until some threshold is met, triggering a decision. The model simultaneously explains both decisions and decision speed. This dynamic element separates the DDM from other cognitive models better known to social and personality psychologists, such as process dissociation (Jacoby, 1991; Payne, 2001) or signal detection (Green & Swets, 1966), which assume a static evidence accumulation process and cannot explain decision speed.

Figure 1 depicts the DDM and Table 1 describes its parameters. The model assumes that when people make a decision between two choices they start out with a bias for one choice. They then accumulate evidence over time, and when that

evidence reaches a threshold, they make the decision. Response time is determined by reaching a threshold. Errors happen because evidence is noisy and sometimes accumulates incorrectly to the wrong threshold. Increasing the distance between thresholds decreases the chance evidence will reach the incorrect threshold but also increases response time, offering a mechanistic explanation of the speed-accuracy trade-off.

For a given relative start point β , threshold separation α , drift rate δ , and nondecision time τ , the model predicts the probability of selecting each choice and the associated response time distribution. Expressions and derivations for these parameters are well-documented (Busemeyer & Diederich, 2010; Diederich & Busemeyer, 2003; Ratcliff & Tuerlinckx, 2002; Van Zandt, 2000). Figure 1 illustrates predicted response time distribution for choosing Option A and Option B given a set of parameters. The area under each distribution reflects the predicted probability that evidence accumulation will terminate at the given threshold (the predicted choice probability). Thus, the total area under the distributions is equal to 1. Figure 2

¹ Michigan State University, East Lansing, MI, USA

² Center for Adaptive Rationality, Max Planck Institute for Human Development, Berlin, Germany

Corresponding Author:

David J. Johnson, Michigan State University, 244B Psychology, 316 Physics Road, East Lansing, MI 48824, USA.
Email: djjohnson@smcm.edu

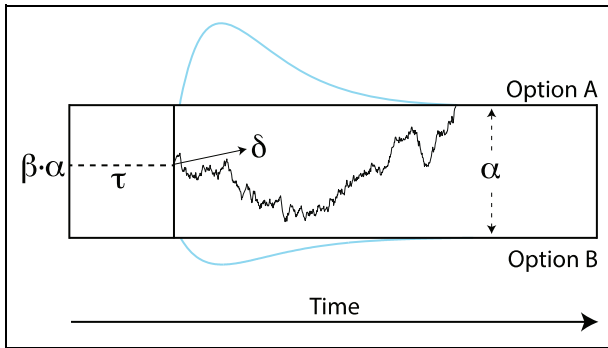


Figure 1. The diffusion model. Individuals start with an initial preference for Option A or B. This initial preference is determined by the relative start point β , which determines the relative location between the two choice thresholds. Information is then accumulated in favor of one of the two options with average strength δ . The amount of information needed to make a decision is indicated by the location of the thresholds with the bottom threshold fixed at 0 and the location of the upper threshold determined by the parameter α . The length of time for other nondecision-related processes by tau τ . Distributions (in blue) above and below the decision space indicate that the model predicts the distribution of response times for each option.

Table 1. Main Parameters of the Drift Diffusion Model.

Parameter	Interpretation
Relative start point (β)	Initial bias to favor one option at the start of the evidence accumulation process, with $0 < \beta < 1$. A value of .50 indicates no preference.
Threshold (α)	Level of evidence required to make a decision, with $0 < \alpha$. Hitting a threshold boundary triggers the relevant choice. Measures how much a person trades accuracy for speed.
Drift rate (δ)	Average quality of information extracted from a stimulus at each unit of time, with $-\infty < \delta < \infty$. Higher absolute values indicate stronger evidence, whereas values around zero indicate ambiguous evidence.
Nondecision time (τ)	Length of all response components unrelated to decision-making, with $0 < \tau$. Reflects encoding, motor response time, and other unknown contaminants. Measured in milliseconds.

shows how changes in the parameters impact predicted choice proportions and response time distributions.

More complex versions of the DDM exist that include trial-by-trial variability in parameters to account for slow and fast errors (Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff, Van Zandt, & McKoon, 1999), changes in information processing as attention switches between sources of information (Diederich, 1997; Diederich & Busemeyer, 2015; Krajbich, Armel, & Rangel, 2010), extra processing stages to account for confidence (Pleskac & Busemeyer, 2010), decay parameters to account for leakage of evidence (Busemeyer & Townsend, 1993; Yu, Pleskac, & Zeigenfuse, 2015), or ways to model




Figure 2. An illustration of how changing diffusion model parameters impacts decisions and response time distributions (in blue). We assume that evidence is correctly accumulated toward Option A. Top panel: higher relative start point β increases the likelihood and speed of selecting Option A by primarily increasing modal response speed. Middle panel: higher threshold α increases the likelihood of choosing Option A and decreases the speed of choosing both options by shifting the mode and lengthening the tails of responses. Bottom panel: higher drift rate δ increases the likelihood and speed of selecting Option A by shortening the tails of the responses. Nondecision time τ is not depicted as it simply shifts both distributions by a fixed amount.

responses with more than two alternatives (Diederich & Busemeyer, 2003; Smith, 2016). Each of these modifications can be compared to simpler versions to test whether the modified models and the hypotheses they formalize better account for data. Here, we focus on the standard model to illustrate its potential for social and personality research.

Benefits of the DDM

As a formal cognitive model, the DDM possesses several benefits including the ability to measure latent processes, disentangle process-level accounts of behavior, and provide explicit tests of fit (Klauer, 2014). However, the dynamic nature of the DDM provides some unique benefits that merit elaboration.

First, the DDM presents a unified model of decision-making speed and accuracy. Often in psychology, decisions and response times are analyzed separately or one is simply neglected. Even worse, when analyzing reaction time data, error trials are typically discarded based on the assumption that some unknown process has contaminated them. This problematic claim is often based on intuition rather than evidence. However, the DDM shows how *the same process can generate both correct and incorrect decisions as well as their speed*. From this perspective, ignoring nonindependence wastes data and prevents researchers from questioning whether factors that

influence accuracy and speed do so through similar or different processes.

This issue also applies to other types of cognitive models, including signal detection (Green & Swets, 1966), process dissociation (Jacoby, 1991) and quadruple models (Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005). By only taking into account decision data, they provide an incomplete picture of the dynamic decision-making process. This is evident in tasks where individuals make few errors, like the implicit association test (Greenwald, McGhee, & Schwartz, 1998). Because these models map processes onto decisions only, when there is little variability in decisions parameters cannot be estimated. In contrast, the DDM can still explain decisions as long as there is variability in decision speed. Thus, it provides a more comprehensive model of decision-making.

Another benefit is that the DDM is a formal description of the decision-making process that can be directly tested rather than relying on indirect tests through questionnaires. This is relevant to personality psychologists, who are interested in how individual differences in traits influence behavior. Typically, researchers use a correlational design where a trait is measured via self-report and then compared to some behavior. However, there are concerns about the role of self-insight and willingness to report information (McClelland, Koestner, & Weinberger, 1989). In contrast, the DDM can experimentally test the cognitive processes that underlie behavior. This method can measure processes without raising self-presentation concerns.

In summary, the DDM presents a unified model of decision-making speed and accuracy that more fully describes the decision-making process compared to other process models common to social and personality psychologists. It also provides a way to test the cognitive processes that give rise to behaviors in ways that complement research in individual differences. Despite these advantages, current implementations of the DDM have made it difficult to use with most social and personality paradigms. We now discuss how a hierarchical version of this model overcomes this difficulty.

Hierarchical DDM

There are several different methods by which the DDM can be estimated from the data (e.g., Busemeyer & Diederich, 2010; Ratcliff & Childers, 2015; Ratcliff & Tuerlinckx, 2002; Wagenmakers, Van Der Mass, & Grassman, 2007). However, the full model must be estimated from response time distributions, which requires a large amount of data (Wagenmakers, 2009). Because the standard DDM analyzes decisions at the *individual* level, this requires many trials per participant. This method is incompatible with many social and personality paradigms, which are based on few trials per participant and focus on *population*-level analyses. Researchers must make trade-offs to estimate the model at the individual level with sparse data, simplifying the model by constraining which parameters vary (e.g., Correll, Wittenbrink, Crawford, & Sadler, 2015) or increasing the number of trials (Dioux,

Brochard, Gabarrot, & Zagar, 2016; Klauer, Voss, Schmitz, Teige, & Mociemba, 2007).

In a two-step approach, researchers first estimate models at the individual level, then conduct a second analysis on the parameter estimates to test the effects of individual differences or experimental manipulations. This approach has been used with psychological phenomena such as shooting decisions (Correll et al., 2015), implicit associations (Klauer et al., 2007), weapon identification (Klauer & Voss, 2008), and stereotype priming (Dioux et al., 2016). However, this two-step approach is not without issue. First, in many experimental contexts, it may be impossible to obtain enough trials per participant to obtain reliable parameter estimates. In addition, because information about participants is isolated, this approach is less powerful than an approach that examines all data simultaneously.

Recent advances in Bayesian modeling have overcome this issue by embedding the DDM in a hierarchical framework (Vandekerckhove, Tuerlinckx, & Lee, 2011; Wiecki, Sofer, & Frank, 2013). This method produces parameter estimates at individual and condition levels. When variability between participants is low, it increases the precision of individual estimates and shrinks them toward the group mean. This allows for precise estimates with sparse data (Kryptos, Beckers, Kindt, & Wagenmakers, 2015). This approach makes the DDM suitable for use in common tasks and increases the power of the analysis.

A depiction of the hierarchical DDM is shown in Figure 3. As the bottom of the figure shows, the joint probability of a decision and its response time are distributed according to a Weiner diffusion process. The model is structured hierarchically because the data for each subject are constrained by a higher order condition-level distribution. For example, the threshold for subject s in within-subjects condition w and between-subjects condition b is indicated by a normal distribution,

$$\alpha_{wbs} \sim \mathcal{N}(\mu_{wb}^{\alpha}, \tau_b^{\alpha}),$$

where μ_{wb}^{α} and τ_b^{α} indicate the mean and precision (inverse of the variance) of the condition-level distribution, respectively. We account for repeated measures designs by keeping the precision parameter constant across within-subjects conditions, but allowing it to vary across between-subjects conditions (Kruschke, 2014).

We estimate the posterior distributions of the hierarchical DDM in a Bayesian framework using Markov Chain Monte Carlo (MCMC) methods. These methods estimate a distribution by repeatedly drawing samples from it. They can be contrasted with maximum likelihood methods, which rely on optimization algorithms to find parameter values that maximize the likelihood of the data. These techniques become impractical as model complexity increases (Navarro & Fuss, 2009; Tuerlinckx, 2004). The DDM is already a complex model (Navarro & Fuss, 2009; Tuerlinckx, 2004) and hierarchical extensions (e.g., random effects of conditions or persons) quickly become intractable using these methods. However, MCMC methods only require the specification of a prior




Figure 3. Graphic diagram of the hierarchical drift diffusion model (DDM). Bivariate decision and response time y for subject s in within-subjects condition w and between-subjects condition b on trial t is generated by a drift diffusion process. Markers on the DDM indicate the data were censored. Subject-level decision parameters are drawn from condition-level truncated normal distributions (middle row). These distributions are given uninformative priors (top row). Each arrow represents a line of JAGS code (see Figure 5) and is indicated by the number to the left of the arrow.




Figure 4. Response times and error rates from Study 1 of Pleskac, Cesario, and Johnson (2017). Error bars are 95% confidence intervals with the standard error estimated from the mean squared error of the race by object interaction from the analysis of variance.

distribution and a tractable likelihood function to update that distribution (Gelman & Hill, 2007; Kruschke, 2014; Kruschke & Vanpaemel, 2015). This framework accommodates complex random structures with little computational burden (Vandekerckhove et al., 2011).

We use the MCMC program JAGS with the Wiener module extension (Wabersich & Vandekerckhove, 2014) to run the model. Figure 5 shows the JAGS code for the hierarchical DDM where all of the model parameters vary according to some number of between- and within-subjects manipulations. Translating the graphic diagram into JAGS code is straightforward. Each arrow corresponds to one line of code. For example,

the equation above is expressed as Line 8. The function “dnorm” indicates that α is normally distributed with mean $\mu\alpha$ and precision $prec\alpha$. The code also contains information about how distributions were truncated.¹

Condition-level parameters are given uninformative priors so the data dominate the posterior estimates. For instance, for each condition the mean threshold is uniformly distributed across all plausible values (Line 18). The precision is given a wide gamma prior ranging from zero to infinity (Line 23). When specifying the model, each parameter can be allowed to vary across all experimental conditions or be held constant when appropriate. Different priors can be investigated to understand their impact on theoretical conclusions, but typically when the data set is moderately large and priors are not severely specific, posterior estimates are robust (Kruschke, 2014).

Walkthrough With First-Person Shooter Task (FPST) Data

To demonstrate how to implement the hierarchical diffusion model, we used data from the FPST (Correll, Park, Judd, & Wittenbrink, 2002). The FPST is used to study how race influences the decision to shoot. In the task, 56 subjects saw 100 pictures of Black and White men holding guns or harmless objects (in a fully crossed design). They were instructed to shoot armed targets and not shoot unarmed targets, within a 850-ms

```

1 model {
2   #likelihood function
3   for (t in 1:nTrials) { #for each trial
4     y[t] ~ dwiener(alpha[WC[t], subject[t]], tau[WC[t], subject[t]],
5     ↪ beta[WC[t], subject[t]], delta[WC[t], subject[t]])
6   }
7   for (s in 1:nSubjects) { #for each subject
8     for (w in 1:nWithin) { #for each within-subjects condition
9       alpha[w, s] ~ dnorm(muAlpha[w, BC[s]], precAlpha[BC[s]]) T(.1, 5)
10      beta[w, s] ~ dnorm(muBeta[w, BC[s]], precBeta[BC[s]]) T(.1, .9)
11      tau[w, s] ~ dnorm(muTau[w, BC[s]], precTau[BC[s]]) T(.0001, 1)
12      delta[w, s] ~ dnorm(muDelta[w, BC[s]], precDelta[BC[s]]) T(-5, 5)
13    }
14  }
15  #priors
16  for (b in 1:nBetween) { #for each between-subjects condition
17    precAlpha[b] ~ dgamma(.001, .001)
18    precBeta[b] ~ dgamma(.001, .001)
19    precTau[b] ~ dgamma(.001, .001)
20    precDelta[b] ~ dgamma(.001, .001)
21    for (w in 1:nWithin) { #for each within-subjects condition
22      muAlpha[w, b] ~ dunif(.1, 5)
23      muBeta[w, b] ~ dunif(.1, .9)
24      muTau[w, b] ~ dunif(.0001, 1)
25      muDelta[w, b] ~ dunif(-5, 5)
26    }
27  }

```

Figure 5. JAGS code for implementing a general form of the hierarchical drift diffusion model. Each line of code that is not a loop, bracket, or comment has a corresponding arrow in the graphic diagram.

response window. This specific data set was reported in Pleskac, Cesario, and Johnson (Study 1; 2017).

The FPST hypothesis is that correct responses to stereotype congruent targets (armed Blacks, unarmed Whites) will be faster than correct responses to stereotype incongruent targets (unarmed Blacks, armed Whites). Figure 4 shows that subjects were faster to correctly shoot armed Black targets than armed White targets, $t(55) = -6.50, p < .001$, and slower to correctly not shoot unarmed Black targets than unarmed White targets, $t(55) = 5.97, p < .001$. That is, participants showed evidence of racial bias. However, this bias was not evident in error rates; subjects made fewer errors for unarmed Black targets than for unarmed White targets, $t(55) = -3.25, p = .002$, and there were no significant differences for armed targets.

While these analyses reveal that race impacts shooting decisions, they do not illustrate *what part* of the decision-making process is influenced by race. Correll, Park, Judd, and Wittenbrink (2002) argued race might influence decisions by biasing people to shoot, changing interpretation the object, or reducing decision-making certainty. Ultimately, they conclude that stereotypes “may theoretically affect any or all of these processes, and it is difficult to disentangle them theoretically, let alone empirically” (p. 1326). While testing these hypotheses is difficult when looking at decisions alone, the DDM can tease apart these possibilities.

Model Specification and Estimation

To examine the FPST data, we embedded the DDM within a hierarchical framework and specified the model according to the guidelines from Pleskac, Cesario, and Johnson (2017). All parameters were allowed to vary as a function of race, but only

Table 2. Sample of the First-Person Shooter Task Data Analyzed With the Hierarchical Drift Diffusion Model.

Subject	Race	Object	Shoot	RT
I	Black	Gun	Yes	464
I	Black	Nongun	No	658
I	Black	Gun	Yes	776
I	White	Gun	Yes	646
I	White	Nongun	No	624

Note. Each trial is represented by five columns: one indicating subject number, two indicating the experimental conditions: race of the target and whether they were holding a weapon, and two indicating the decision and the speed at which it was made.

drift rate and nondecision time were allowed to vary as a function of object. This design reflects a balance between exploring the different possibilities of how race might influence shooting decisions and model parsimony, as measured by comparing alternative specifications, their posterior distributions, and model fit. A graphic diagram of the model is provided in the Supplemental Materials.

We estimated the model using JAGS 4.20 (Plummer, 2003) for MCMC sampling and R (R Core Team, 2016) to send data to JAGS and analyze the predicted values. JAGS code, R code, and data to replicate these analyses are provided in the Supplemental Materials. A sample of the data set is given in Table 2, and each trial represents a single row.

Hypothesis Testing in a Bayesian Framework

We estimate the posterior distribution of parameters by generating a large sample from it using MCMC methods. Each sample in the MCMC chain provides one credible combination of parameter values given the data and prior distribution. We repeated this process with four independent chains 20,000 times for a total of 80,000 samples. We used the diagnostic methods recommended by Kruschke (2014) to ensure the representativeness and accuracy of the samples; these are detailed in the Supplemental Materials.

We describe the posterior distributions for each parameter (all 80,000 values) by their modal value and 95% highest density interval (HDI). The modal posterior value has the highest probability density, making it the most credible estimate. Values in the 95% HDI have a higher probability density than values outside the interval and so are more credible (Kruschke, 2014). We use an estimation approach to hypothesis testing (see Chapter 12, Kruschke, 2014). We ask if the credible values in the 95% HDI of a parameter contain some sort of null value (e.g., 0). If the null value is not among the credible values, we reject it.

Testing for differences between conditions is simple within this framework. We take the posterior distributions for two parameters of interest such as the threshold for Black targets and White targets. We then calculate the difference between the parameter values in each sample. This produces a




Figure 6. Posterior modes (dots) and 95% highest density interval (lines) for the condition-level parameter estimates of the hierarchical drift diffusion model on first-person shooter task data as reported in Pleskac, Cesario, and Johnson (2017). W = White; B = Black.

distribution of credible mean differences. If the 95% HDI for this distribution does not contain zero, the condition effect is credible.

Results and Discussion

The full results of this analysis are reported in Pleskac, Cesario, and Johnson (2017). We summarize them here in Figure 6 to illustrate the inferences one can make with the DDM. Starting with the drift rate, while race did not influence the strength of evidence (drift rates) for nongun objects ($M = 0.07$, 95% HDI $[-0.25, 0.44]$), it did influence the strength of evidence for gun objects ($M = 0.64$, 95% HDI $[0.31, 0.98]$). Evidence for the shoot decision was stronger when Black men held guns compared to White men. This analysis is consistent with the account that race influences identification of guns because of stereotypic associations between Black men and violence.

Turning to the relative start point, all conditions showed a bias to favor the shoot decision ($\beta > .5$). This bias is expected given the payoff structure of the FPST, where participants are on average rewarded more for shoot decisions. It also speaks to the validity of the hierarchical DDM. Past work using a

nonhierarchical DDM (Correll et al., 2015) has shown a bias to favor *not* shooting, which does not accurately reflect the payoff structure of the task.

Participants also showed a bias to favor *not* shooting Black targets relative to White targets ($M = .04$, 95% HDI $[-.01, .08]$) and required *more* evidence ($M = .09$, 95% HDI $[-.00, .18]$) when making decisions for Black targets than White targets. These counterstereotypic biases partially but not completely offset shooting errors caused by biased evidence accumulation and may reflect motivated strategies to reduce racial bias in shooting decisions. Past work using nonhierarchical versions of the DDM has not found such race-based differences (Correll et al., 2015). This may be partially due to the simplifications researchers make to estimate the model as well as decreased power from analyzing data in isolation. These findings demonstrate the hierarchical DDM may provide a richer account of shooting decisions because the race of an individual may have different and opposing effects on the decision process.

These findings also explain why race bias shifts from errors to response times when individuals have longer to respond (Correll et al., 2002). The first factor is that drift rates are stronger for the shoot decision for Black men than White men. All

things equal, subjects should be more likely and faster to shoot Black men. In practice, whether bias manifests in decisions or response times depends on whether enough information is collected to avoid errors. As the threshold increases, error bias should decrease but response time differences should increase. This is the pattern we observed in the current data set. However, when response windows are shorter (e.g., 630 ms), thresholds decrease. This process-level difference parallels a behavioral increase in all errors and reveals decision bias (Pleskac, Cesario, & Johnson, 2017). Thus, the model parsimoniously explains how bias in errors and decision speed results from different factors impacting the *same process*.

Finally, an examination of nondecision time revealed a main effect of object. Nondecision times were faster for guns than nonguns ($M = 47$, 95% HDI [35, 60]). This may be partially due to greater variation among nongun stimuli (e.g., wallets, cell phones) than gun stimuli. There was also an effect of race on nondecision times for gun objects ($M = 21$, 95% HDI [2, 38]), but not nongun objects ($M = -14$, 95% HDI [-34, 4]).

In summary, FPST analyses that focus only on decisions or response times are not well equipped to explain *how* race impacts different stages of the decision to shoot. Because of this, several questions have remained elusive, such as whether race bias results from prior biases or changes in object interpretation, and why bias shifts from errors to response times when subjects are given longer to respond. The hierarchical DDM provides parsimonious answers to these questions while providing a more coherent and richer account of the decision to shoot than nonhierarchical versions.

Having provided a framework for how to estimate and interpret results from the hierarchical DDM with a common social psychology task, we turn to a more theoretical discussion of adapting the DDM to study individual differences.

The Flash Gambling Task (FGT)

Personality psychologists have established associations between impulsivity (Acton, 2003) and risky behavior such as gambling or substance abuse. But, what are the processes that lead impulsive people to engage in risky behavior? One facet of impulsivity is *deliberation*, or how much people think carefully before acting. People scoring low on this facet of deliberation are typically described as hasty, impulsive, careless, and impatient (Barratt, 1993; Whiteside & Lynam, 2001). From this perspective, the DDM provides a formal description of the deliberation process and thus provides a means to better understand the processes underlying impulsivity.

Pleskac, Yu, Hopwood, and Liu (2017) used the FGT (Zeigenfuss, Pleskac, & Liu, 2014) to examine this connection within risky decision-making. In the FGT, participants make repeated choices between a certain and uncertain payoff. A typical FGT trial is shown in Figure 7. Each option in the FGT is represented by an array of randomly positioned dots. Each dot corresponds to a small amount of money and the total number of dots in each array represents a sampled payoff. The *certain*




Figure 7. Diagram of flash gambling task stimulus. Participants choose between an option that provides a certain payoff and an option that provides an uncertain payoff. The certain option is represented by a fixed number of dots, where the total number of dots signifies the payoff. The uncertain option is also represented by a field of dots, but the number of dots is determined by random draws updated every 50 ms from a normal distribution with some mean and variance. Participants observe sequential samples from the option to form an impression about it. Thus, the sequential sampling process assumed by the diffusion model is made observable. When the participant makes a choice, he or she receives the next draw from the option as a payoff.

option has a fixed number of dots in its array (e.g., 130) whereas the *uncertain* option has a changing number of dots that updates when a new sample is drawn (every 50 ms). Participants use these payoffs to determine which option they prefer and earn the next payoff from their chosen option. Importantly, this makes the sequential sampling process observable, whereas applications like the FPST assume a latent sequential sampling process. In other words, for a given trial the stochastic path of evidence accumulation in Figure 1 can be observed. This allows a better understanding of how people use sequential samples of information during deliberation (see Brown, Steyvers, & Wagenmakers, 2009; Irwin, Smith, & Mayfield, 1956; Vickers, Burt, Smith, & Brown, 1985).

Several studies have found that people are risk seeking, preferring the uncertain option over the certain option even when the certain option is equal to or greater in expected value. For example, Pleskac, Yu, et al. (2017) had participants choose between a certain option with 130 dots and an uncertain option with a mean of 115, 130, or 145 dots. Even when the uncertain option was worse than the certain option, participants typically chose the uncertain option. Moreover, Pleskac, Yu, et al. (2017) went beyond fitting the hierarchical DDM to choice behavior and analyzed how participants used the stream of payoff information from the uncertain option. Results were consistent with the DDM; participants accumulated the stream of information as evidence and used it to determine their choice. Thus, the DDM appears to accurately capture the deliberation process during risky decision-making in the FGT.

This result leads to questions about whether aspects of the deliberation process, such as initial biases and the amount of information required to make decisions, might relate to the connection between trait impulsivity and risky behaviors like substance abuse. Pleskac, Yu, et al. (2017) correlated the hierarchical DDM parameters from the FGT with measures

of risky behavior and trait impulsivity. Interestingly, the DDM parameters measuring deliberation during the FGT were not related to self-reported measures of risk attitudes and impulsivity. Nevertheless, aspects of the deliberation process as measured by the DDM did provide incremental validity over an impulsivity questionnaire for predicting substance use. Specifically, individuals with an initial bias toward the uncertain option as well as individuals less sensitive to the samples of information (a drift rate proxy) were more likely to have abused substances. Correlations between the FGT and risky behavior were modest, like other cognitive/behavioral tasks (Cyders & Coskunpinar, 2011; Reynolds, Ortengren, Richards, & de Wit, 2006). However, this result suggests that the FGT as parameterized by the hierarchical DDM may be useful in assessing behavioral risk.

It is interesting that the DDM parameters during the FGT were related to substance use behavior even though they were not related to questionnaire-assessed impulsivity. Bornstein (2012) reviewed similar patterns in assessment and found that while self-report and performance-based measures of dependency are uncorrelated, they nevertheless are incrementally informative about dependent behavior. Based on such findings, Bornstein argued for a process-dissociation approach to psychological assessment that integrates self-report and behavioral assessment. Rather than assuming that measures targeting similar constructs should be correlated, it may be more useful to consider how they capture unique aspects of a complex process.

For instance, an individual answering a questionnaire may be able to recall aspects of their behavior and reliably report on the likelihood of their own impulsive behavior, yet their responses may not capture features of that behavior common to all people. In contrast, tasks such as the FGT may capture the underlying cognitive processes associated with risky decision-making, independent of risk-taking history. When predicting risky decision-making, it might be useful to know both something about an individual's general tendency for impulsive behavior, as assessed by an impulsivity questionnaire, and the underlying cognitive processes that make such behavior more likely, as assessed by measures like the FGT.

In summary, within the FGT, the DDM provides a nuanced understanding of how people weigh information in real time under conditions of subjective preference and ultimately helps determine whether someone looks hasty, careless, and impatient. Initial work suggests that aspects of the deliberation process measured during the FGT are associated with the likelihood of substance abuse but are independent of questionnaire-assessed impulsivity. These results highlight the importance of using self-report and performance-based measures to understand traits, as well as implications that follow such as assessing risk-taking propensity.

Discussion

These examples demonstrate how the hierarchical DDM provides novel insights into social and personality processes. For

social psychologists, the DDM can disentangle different theoretical accounts of how race influences the decision to shoot. For personality psychologists, the model provides a formal description of deliberation, a facet of impulsivity. In both cases, the model illuminates aspects of the decision process unobservable from analyses of behavior alone.

Comparisons to Dual-Process Models

The DDM is built to test a theory of decision-making based on sequential sampling, whereas dual process theories provide a different perspective. Formal dual-process models are well equipped to divide decisions into associative and deliberative components but do not describe the dynamic nature of this process. The benefits gained from these models may not outweigh the costs associated with ignoring that decisions are made over time.

One misconception about the DDM is that it merely quantifies the “automatic” process underlying decisions and does not speak to the more “controlled” process. Therefore, the dual-process distinction is maintained and the DDM is used to clarify the automatic process. However, this reflects an erroneous assumption that fast decisions reflect the influence of purely automatic processes (Payne, 2001). Similarly, the DDM has been applied to decisions lasting over a second (Krajbich et al., 2010; Pleskac & Busemeyer, 2010), which most psychologists would not consider automatic. The point is that the DDM is a qualitatively different decision model than dual-process models. It does not focus on whether its parameters are controlled or automatic, and each parameter “probably represents a complex mixture of controlled and automatic components” (Klauer, 2014, p. 148).

Interpreting DDM Parameters

A benefit of the DDM is the potential to link its model parameters to psychological constructs (Klauer, 2014). For example, the threshold parameter measures the amount of information required to make a decision and is often interpreted as indexing caution. However, moving from parameter to psychological construct requires validating that interpretation. For example, a manipulation of caution should only influence threshold (convergent validity) and not other parameters (discriminant validity). Although some work has validated DDM parameters in the FPST and FGT, more work is needed. This same logic applies to *any* task where researchers want to interpret parameters in terms of psychological mechanisms.

The benefits of linking DDM parameters with psychological constructs should not be understated. For example, in accordance with Bornstein's (2012) process dissociation approach, a DDM decomposition of the FGT may capture cognitive processes underlying risky decision-making, independent of risk-taking history. These results may reveal computational phenotypes for risk-taking behaviors. For example, changes in relative start point as a function of drug use may reflect sensitivity to labels indicating whether an option is certain or

uncertain. This could explain why substance abuse is sometimes understood as a cue-induced urge (Bonson et al., 2002; Ehrman, Robbins, Childress, & O'Brien, 1992).

In terms of the FPST, insofar as the threshold can be interpreted as a measure of caution, the finding of higher thresholds for Black men than White men suggests that participants may be motivated to avoid biased responding. This could be tested by measuring individual differences in the motivation to avoid prejudice or by manipulating this motivation through instruction. Similarly, insofar as stronger drift rates toward the shoot decision represent a bias to see objects held by Black men as threatening, training to focus on the object (rather than race) may decrease this bias. This could explain why officers show less bias in the FPST than civilians (Correll et al., 2007; Sim, Correll, & Sadler, 2013). Both possibilities demonstrate how the DDM can further our understanding of this important decision and even help create training programs.

Conclusion

The DDM is a dynamic cognitive model based on sequential sampling. Two applications of this model were presented that highlight its potential to provide novel insights into cognitive processes. In particular, the model can reveal differences in the decision-making process that are difficult to discern when only observing behavior. Recent advances in Bayesian modeling make it feasible to implement a hierarchical version of the DDM to social and personality psychology tasks. Although we focused on two examples, the model can in principle be used to explain any task where individuals make quick decisions. Given the ubiquity of these tasks, we hope this primer equips researchers with a new tool for investigating psychological processes.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was partially supported by a grant from the National Institutes of Health (R03DA033455) to the second and fourth author, a grant from the National Science Foundation (1230281) to the third author, and a grant from the National Science Foundation (0955410) to the fourth author.

Supplemental Material

The online data supplements are available at <http://journals.sagepub.com/doi/suppl/10.1177/1948550617703174>.

Note

1. Truncation facilitates model estimation. Parameter boundaries were chosen to be far from the observed parameter estimates, and parameter estimates are robust to boundary changes.

References

- Acton, G. S. (2003). Measurement of impulsivity in a hierarchical model of personality traits: Implications for substance use. *Substance Use and Misuse, 38*, 67–83.
- Barratt, E. S. (1993). Impulsivity: Integrating cognitive, behavioral, biological, and environmental data. In W. G. McCown, J. L. Johnson, & M. B. Shure (Eds.), *The impulsive client: Theory, research, and treatment* (pp. 39–56). Washington, DC: APA.
- Bonson, K. R., Grant, S., Contoreggi, M. D., Links, J. M., Metcalfe, J., Weyl, L., . . . London, E. D. (2002). Neural systems and cue-induced cocaine craving. *Neuropsychopharmacology, 26*, 376–386.
- Bornstein, R. F. (2012). From dysfunction to adaptation: An interactionist model of dependency. *Annual Review of Clinical Psychology, 8*, 291–316.
- Brown, S., Steyvers, M., & Wagenmakers, E.-J. (2009). Observing evidence accumulation during multi-alternative decisions. *Journal of Mathematical Psychology, 53*, 453–462.
- Busemeyer, J. R., & Diederich, A. (2010). *Cognitive modeling*. Los Angeles, CA: Sage.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review, 100*, 432–459.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology, 89*, 469–487.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer's dilemma: Using ethnicity to disambiguate potentially threatening individuals. *Journal of Personality and Social Psychology, 83*, 1314–1329.
- Correll, J., Park, B., Judd, C. M., Wittenbrink, B., Sadler, M. S., & Keesee, T. (2007). Across the thin blue line: Police officers and racial bias in the decision to shoot. *Journal of Personality and Social Psychology, 92*, 1006–1023.
- Correll, J., Wittenbrink, B., Crawford, M. T., & Sadler, M. S. (2015). Stereotypic vision: How stereotypes disambiguate visual stimuli. *Journal of Personality and Social Psychology, 108*, 219–233.
- Cyders, M. A., & Coskunpinar, A. (2011). Measurement of constructs using self-report and behavioral lab tasks: Is there overlap in nomothetic span and construct representation for impulsivity? *Clinical Psychology Review, 31*, 965–982.
- Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology, 41*, 260.
- Diederich, A., & Busemeyer, J. R. (2003). Simple matrix methods for analyzing diffusion models of choice probability, choice response time, and simple response time. *Journal of Mathematical Psychology, 47*, 304–322.
- Diederich, A., & Busemeyer, J. R. (2015). *Multi-stage sequential sampling model of multi-attribute decision making* (Tech. Rep.). Bremen, Germany: Jacobs University.
- Dioux, V., Brochard, R., Gabarrot, F., & Zagar, D. (2016). The young and the reckless: Does priming age stereotype in young adults affect speed or cautiousness? *Social Cognition, 34*, 324–342.

- Ehrman, R., Robbins, S. J., Childress, A. R., & O'Brien, C. P. (1992). Conditioned responses to cocaine-related stimuli in cocaine abuse patients. *Psychopharmacology*, *107*, 523–529.
- Forstmann, B. U., Ratcliff, R., & Wagenmakers, E. J. (2016). Sequential sampling models in cognitive neuroscience: Advantages, applications, and extensions. *Annual Review of Psychology*, *67*, 641–666.
- Gelman, A., & Hill, J. (2007). *Data analysis using regression and multi-level/hierarchical models*. New York, NY: Cambridge University.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Oxford, England: Wiley.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464–1480.
- Irwin, F. W., Smith, W. A. S., & Mayfield, J. F. (1956). Tests of two theories of decision in an “expanded judgment” situation. *Journal of Experimental Psychology*, *51*, 261–268.
- Jacoby, L. L. (1991). A process dissociation framework: Separating automatic from intentional uses of memory. *Journal of Memory and Language*, *30*, 513–541.
- Klauer, K. C. (2014). Random-walk and diffusion models. In J. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual process theories of the social mind* (pp. 139–152). New York, NY: Guilford.
- Klauer, K. C., & Voss, A. (2008). Effects of race on responses and response latencies in the weapon identification task: A test of six models. *Personality and Social Psychology Bulletin*, *34*, 1124–1140.
- Klauer, K. C., Voss, A., Schmitz, F., & Teige-Mocigemba, S. (2007). Process components of the Implicit Association Test: A diffusion-model analysis. *Journal of Personality and Social Psychology*, *93*, 353–368.
- Krajbich, I., Armel, C., & Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience*, *13*, 1292–1298.
- Kruschke, J. K. (2014). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). New York, NY: Academic/Elsevier.
- Kruschke, J. K., & Vanpaemel, W. (2015). Bayesian estimation in hierarchical models. In J. R. Busemeyer, Z. Wang, J. T. Townsend, & A. Eidels (Eds.), *The Oxford handbook of computational and mathematical psychology* (pp. 279–299). Oxford, England: Oxford University.
- Kryptos, A.-M., Beckers, T., Kindt, M., & Wagenmakers, E.-J. (2015). A Bayesian hierarchical diffusion model decomposition of performance in Approach-Avoidance Tasks. *Cognition and Emotion*, *29*, 1424–1444.
- McClelland, D. C., Koestner, R., & Weinberger, J. (1989). How do self-attributed and implicit motives differ? *Psychological Review*, *96*, 690–702.
- Navarro, D. J., & Fuss, I. (2009). Fast and accurate calculations for first-passage times in Wiener diffusion models. *Journal for Mathematical Psychology*, *53*, 222–230.
- Payne, B. K. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*, 181–192.
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, *117*, 864–901.
- Pleskac, T. J., Cesario, J., & Johnson, D. J., (2017). *How race affects evidence accumulation during the decision to shoot*. Manuscript submitted for publication.
- Pleskac, T. J., Yu, S., Hopwood, C., & Liu, T. (2017). *Characterizing deliberation during preferential choice*. Manuscript submitted for publication.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*.
- R Core Team. (2016). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Ratcliff, R., & Childers, R. (2015). Individual differences and fitting methods for the two-choice diffusion model of decision making. *Decision*, *2*, 237–279.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, *20*, 873–922.
- Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*, 347–356.
- Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Science*, *20*, 260–281.
- Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, *9*, 438–481.
- Ratcliff, R., Van Zandt, T., & McKoon, G. (1999). Connectionist and diffusion models of reaction time. *Psychological Review*, *106*, 261–300.
- Reynolds, B., Ortengren, A., Richards, J. B., & de Wit, H. (2006). Dimensions of impulsive behavior: Personality and behavioral measures. *Personality and Individual Differences*, *40*, 305–315.
- Sim, J. J., Correll, J., & Sadler, M. S. (2013). Understanding police and expert performance: When training attenuates (vs. exacerbates) stereotypic bias in the decision to shoot. *Personality and Social Psychology Bulletin*, *39*, 291–304.
- Smith, P. L. (2016). Diffusion theory of decision making in continuous report. *Psychological Review*, *123*, 425–451.
- Tuerlinckx, F. (2004). The efficient computation of the distribution function of the diffusion process. *Behavior Research Methods, Instruments, & Computers*, *36*, 702–716.
- Van Zandt, T. (2000). How to fit a response time distribution. *Psychonomic Bulletin & Review*, *7*, 424–465.
- Vandekerckhove, J., Tuerlinckx, F., & Lee, M. D. (2011). Hierarchical diffusion models for two-choice response times. *Psychological Methods*, *16*, 44–62.
- Vickers, D., Burt, J., Smith, P., & Brown, M. (1985). Experimental paradigms emphasizing state or process limitations: I. Effects on speed accuracy tradeoffs. *Acta Psychologica*, *59*, 129–161.
- Wabersich, D., & Vandekerckhove, J. (2014). Extending JAGS: A tutorial on adding custom distributions to JAGS (with a diffusion model example). *Behavior Research Methods*, *46*, 15–28.

- Wagenmakers, E.-J. (2009). Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology, 21*, 641–671.
- Wagenmakers, E. J., Van Der Maas, H. L., & Grasman, R. P. (2007). An EZ-diffusion model for response time and accuracy. *Psychonomic Bulletin & Review, 14*, 3–22.
- Whiteside, S. P., & Lynam, D. R. (2001). The five factor model and impulsivity: Using a structural model of personality to understand impulsivity. *Personality & Individual Differences, 30*, 669–689.
- Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical Bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics, 7*, 1–10.
- Yu, S., Pleskac, T. J., & Zeigenfuse, M. D. (2015). Dynamics of post-decisional processing of confidence. *Journal of Experimental Psychology: General, 144*, 489–510.
- Zeigenfuse, M. D., Pleskac, T. J., & Liu, T. S. (2014). Rapid decisions from experience. *Cognition, 131*, 181–194.

Author Biographies

David J. Johnson is a graduate student at Michigan State University. His research uses cognitive modeling to understand the processes underlying decision-making.

Chris J. Hopwood is an Associate Professor of Psychology at Michigan State University. His research focuses on interpersonal processes, personality assessment, and personality disorders.

Joseph Cesario is Associate Professor of Psychology at Michigan State University. He studies social cognitive processes from an evolutionary approach.

Timothy J. Pleskac studies how people make judgments and decisions; how these processes shape behavior at the individual, group, and organizational level; and how we can help people make better judgments and decisions. He investigate these questions with computational modeling and methods from the behavioral and cognitive sciences.

Handling Editor: Gregory Webster