Bridging between Hypothetical and Incentivized Choice *

Arash Laghaie †
Thomas Otter ‡

Current Version: February 28, 2025

Abstract

The hypothetical nature of choices collected in typical discrete choice experiments (DCEs) for market research has been a source of concern for both researchers in academia and industry. Recent studies in marketing indeed demonstrate that choices change qualitatively when respondents are properly incentivized. However, studies that try to model these changes in the framework of random utility models show that incentives make choices *less* consistent reflected in a decrease of the scale of preference estimates, which is widely considered as a measure of choice quality in the literature. This result is concerning for market researchers that consider incentive-alignment and seemingly contradicts the large body of behavioral literature that suggests incentives increase decision effort and performance. In this paper, we propose a model that sheds light on the mechanism underlying the differences between hypothetical (HYP) and incentivized (ICA) multi-attribute choice. The model is in the class of sequential sampling models of choice and measures the differential cognitive processing between respondents in HYP and ICA choice settings. The empirical application of the model shows that respondents in the ICA group attend to larger sets of attributes on average and therefore resolve more trade-offs than those in the HYP group. Consequently, ICA respondents process more information and across more diverse sources of utility when making inside choices and make *less* consistent choices overall that nevertheless are of higher quality than HYP choices in that they are more aligned with respondents' deep preferences. Insights from this paper can help market researchers when deciding to invest in incentive-alignment mechanisms to collect better data.

*This research was supported by the German Science Foundation (DFG) grant OT 447/1-1.
†Nova School of Business and Economics. Email: arash.laghaie@novasbe.pt
‡Goethe University Frankfurt & Nova SBE. Email: otter@marketing.uni-frankfurt.de

1 Introduction

Discrete choice experiments (DCEs) are widely adopted in economics, marketing and other disciplines to elicit so called stated preferences (see McFadden, 1986 for the theory and e.g., Roy et al., 1996 and Louviere and Lancsar, 2009 for applications). Research in experimental economics documents that preferences elicited in such inconsequential, hypothetical setups may generalize poorly to prevailing preferences when choice consequences are real (e.g., Smith, 1982; Camerer et al., 1999). A number of recent studies in marketing demonstrate empirically that data collected in hypothetical (HYP) settings results in inferences that do not extrapolate well to choices that have real consequences (e.g., Ding et al., 2005; Ding, 2007; Ding et al., 2009; Dong et al., 2010; Wertenbroch and Skiera, 2002; Miller et al., 2011). These studies suggest different mechanisms making choices incentive-aligned (ICA), linking choices in a DCE to real consequences, in order to incentivize respondents to behave as when revealing their preferences in a market transaction.

However choice in these studies is modeled under the standard random utility assumption that the decision maker (DM) acts on fully observed utilities of alternatives. As a result, HYP and ICA choices inform two sets of preference parameters that fit the differences between the two incentive schemes. The HYP preference parameters have to be discarded and prediction is done using the ICA preferences. This implies that there is no useful information in the HYP data even though marketing decisions in practice have relied on HYP DCEs for decades.

Attempts to generalize across choices under different incentive settings and stable preferences have modeled the impact of incentives as changing the scale parameter in a Scale Multinomial logit model (Scale MNL). This model suggests that incentive-alignment does not affect preferences but changes choice consistency. A recent study by Hauser et al. (2019) shows that incentive-alignment, all else equal, results in *less* consistent choices reflected in a decrease in the preference scale parameter estimate.[1] These results are concerning for practitioners and academics who consider employing incentive-aligment in

[1] Hauser et al. (2019) also show that the scale estimate in text-based DCE choices are higher than that in image-based DCEs and argue that this might be because text-based choices are "easier for respondents to respond" than image-based choices. However, within the text-based DCEs they conduct, which are the standard practice in market research and the context of this paper, ICA choices have lower scale than HYP choices. The same pattern can be seen across HYP and ICA DCE choices in the data in Miller et al. (2011).

their market research as the scale parameter in Random Utility Models (RUM) is widely considered as a measure of accuracy and choice quality (Toubia et al., 2004; Evgeniou et al., 2005; Howell et al., 2021).

In contrast, a large body of literature across behavioral economics, psychology, and neuroscience suggests that incentives increase attention and effort and consequently improve task performance. In economics, Camerer et al. (1999), reviewing 74 experiments, show that incentives increase attention and effort when it is efficacious, which in turn improves performance. In the experimental literature, studies that use proximal measures of decision effort such as response time, pupillary dilation, brain activity, or eye-traces point to increased decision effort when choice is incentivized (e.g., Kahneman and Peavler, 1969; Kahneman, 1973; Wilcox, 1993; Yang et al., 2018; Camerer and Mobbs, 2017). In neuroscience, studies that focus on neural correlates of cognitive effort (e.g., Shenhav et al., 2013), or monitor cognitive task performance (e.g., Padmala and Pessoa, 2011) show that incentives enhance effort and performance.

The goal of this paper is to reconcile the above results by studying the mechanism through which incentives translate into qualitatively different choice behavior. To achieve this, we develop a model that estimates preferences and effort from choice data. Applying the model to a large-scale experiment consisting of DCEs conducted at different levels of incentive-alignment we find that the source of inconsistency brought about by incentive-alignment and reflected in the scale in the RUM framework, is not due to lack of attention and cognitive effort. To the contrary, DMs at higher incentives attend to more attributes and process more information when making inside choices. Our model explains how attending to more attributes can result in less consistent choices. Moreover, we show that despite being less consistent, choices at higher incentives have higher quality in that they are more in line with the DM's deep preferences. Insights from our study can help market researchers decide whether to invest in incentive-alignment of DCEs to improve the responses.

Our main assumption is that HYP choices recruit different amount of decision effort than ICA choices, but under the same set of deep preferences. Therefore in our model the information set, i.e., what attributes are used at all, and the level of processing of attributes and alternatives change across the HYP and the ICA choice setting. The underlying, deep interpretation and source of valuation of attributes however is assumed to be invariant. We believe this to be a reasonable assumption in the context of a familiar collection of attributes under study, which we test in our empirical application. In other words, we conceive of the decision process as constructing valuations of alternatives in a choice context by integrating attribute information, based on a set of invariant, deep parameters underlying this preference construction. In addition, our model features parameters that capture the effective information set and the level of cognitive involvement with the information set. These parameters affect if and to what extent the DM integrates the value of different product attributes in the choice process[2], and they change across contexts that induce more (ICA) and less (HYP) effort.

We adopt a sequential sampling model (SSM), and specifically the recently proposed dependent Poisson race model (DPRM) (Ruan et al., 2008) as our modeling framework. Sequential sampling models that encompass diffusion and race models (see Ratcliff and Smith, 2004, for an overview) originate in mathematical psychology, and motivate choice from a latent process of evidence accumulation. The idea is that the DM encodes characteristics of the alternatives in the choice set, and by processing them produces evidence that accumulate in alternative-specific counters. In this way, the DM constructs preferences for alternatives in the choice set (see Bettman et al., 1990, for a general discussion of constructive aspects of preferences). Once the accumulated evidence or preference supporting one alternative reaches a preset threshold, the process stops and the choice is made. If two or more alternatives reach the threshold at the same time, which is only possible when counters are dependent, the tie is broken at random. Processing alternatives' characteristics, i.e., evaluating attributes in the context of a particular choice set and integrating these evaluations to an overall assessment of which is the best alternative is effortful. The model that we develop both measures which characteristics are processed and how much processing occurs (in expectation) before making a choice. As a consequence, the model can predict choices under different levels of effort induced by HYP and ICA choice settings.

[2]In contrast to search models, what is costly in our model is not finding out the value of attributes but translating that value into utility and integrate it into product valuation. This aspect of the model usefully fits in with the nature of DCEs where information about all the attributes is equally readily available to the DM.

The evidence accumulation process posited in the DPRM builds on the evaluation of attributes in the context of multi-attribute alternatives that define a choice set. The DM produces and integrates attribute-specific pieces of evidence to determine which alternative should be chosen. When the DM processes information to figure out which is the best alternative in a set, the overall process is structured as a super-position of simpler processes associated with individual attribute levels. For example, the color red, as a level of the attribute color, could be associated with a simple process that generates positive evidence at a rate proportional to the deep, underlying attractiveness of that color. As we will show, the DPRM models these simple processes as Poisson such that process increments are discrete units (so called 'hits') and the states are aggregate counts of past hits. Whenever the simple process associated with the color red generates a hit, this hit will accrue to all alternatives in the set that share the color red, and thereby increase alternative-specific evidence counters. Therefore the constructed preferences that lead to a choice are stochastic manifestations of the DM's deep preferences, which are the partworth parameters in the model. As a consequence of building up evidence for alternatives from attribute-level processing, the model implies similarity effects. This way it structurally deviates from the widely criticized independence of irrelevant alternatives (IIA) property.

A basic premise of the DPRM, and other models in this class, is that the processing that translates into valuations is effortful. As we will show, structuring the overall evidence for an alternative as coming from attribute-level specific valuation processes implies that decisions between more comparable alternatives (sharing similar attribute levels) endogenously require less time and processing effort, and are generally more deterministic (consistent). In contrast, choices among less comparable alternatives in the sense of requiring (more) trade-offs between attributes endogenously require more time and processing effort and are generally less deterministic, more volatile. These implications of the DPRM are consistent with well documented findings in experimental psychology (Zhang and Markman, 2001; Bettman et al., 1990; Tversky, 1969; Tversky and Russo, 1969). Similarly, to the extent that attending to more attributes results in more trade-offs, *ignoring* attributes could result in *less* information processing and *more* consistency. This is in analogy to numerical solutions to lower (higher) dimensional integrals that require less (more) computational effort to achieve the same level of accuracy.

Higher evidence thresholds (required for a decision) result in effortfully accumulating more evidence and choosing the best alternative more consistently, everything else equal. Therefore, the DM acts under a threshold level that reflects a general trade-off between decision effort and decisions accuracy (Fehr and Rangel, 2011). To set the threshold, she weighs the prior expected gain from a perfectly accurate choice, against the opportunity cost of time and processing effort. The model captures this trade-off in reduced form. As we will show, the reduced form is likelihood-identified from choice data. The evidence threshold is fixed under rational expectations about the prior expected gain and invariant time and effort cost functions. In contrast to, e.g., court rulings or decisions about scientific theories, the goal is not to maintain a high level of accuracy per se but to balance benefits from choice with time and effort costs. In this respect the choice process in the DPRM falls in the class of satisficing algorithms (Simon, 1955). These algorithms, as Gigerenzer and Goldstein (1996) note, need to meet more criteria than just accuracy. Before seeing the alternatives, based on the incentive level and prior knowledge about the possible attribute levels, the DM gauges the importance of the decision relative to the opportunity cost of time and processing, and sets the threshold accordingly. Just as in Simon's satisficing rule, what is important for the DM is for the choice to reach a threshold rather than some relative ranking (Stüttgen et al., 2012). However, the DPRM can accommodate a range of behaviors from attribute based satisficing to (essentially) deterministic choice of the utility maximizing alternative.

As the incentive level increases, the DM exerts more decision effort by expanding the information set (set of attributes) processed. Expanding the information set and attending to more attributes poses more trade-offs and information integration across non-comparable attributes in choice. Conditional on a fixed threshold, increasing the set of attributes processed leads to more information processing before a choice occurs. This is because increasing the dimensionality of the source of utility decreases the comparability of alternatives in a choice set and hence increases the number of trade-offs facing a DM. Therefore, to the extent that the DM has limited cognitive resources, the decision about the information set and the decision about the threshold level are not independent.

Even though more attribute attendance and a higher threshold level both lead to more information processing, they have different implications for choice behavior. Higher threshold results in more evidence accumluation and more accurate choices based on the DM's deep preferences. The probability of tie decreases in the threshold value. Consequently, at high thresholds, the effect of attribute-level similarity vanishes as simultaneous evidence no longer create tied races (assuming at least some minimal differentiation between alternatives). On the other hand, attending to more attributes increases the amount of attribute trade-offs and makes choices less consistent.

These properties imply that, conditional on a fixed set of deep preferences, the DPRM can capture various qualitatively different choice behavior depending on a combination of the threshold level and number of attributes attended. The different behaviors induced by the model generalize beyond more or less random choices and, as we will show, can potentially lead to preference reversals. We exploit this feature of the DPRM to capture reactions to incentive change.

The model learns about preferences and the threshold from choice. Variation in attribute levels allows for the identification of preferences from repeated choices. The threshold is identified from overall choice consistency and the degree of attribute-based (vs. alternative-based) processing in choice, that result in different substitution patterns. Having parameters associated with preferences and decision effort makes the DPRM a plausible model for studying the mechanisms by which incentives change choice behavior. The separation of threshold and rate parameters is an established feature of SSM models. The mathematical psychology literature shows that manipulating accuracy or speed goals in choice *only* changes the threshold (e.g. Ratcliff and Rouder, 1998; Starns and Ratcliff, 2014) and manipulating the information quality *only* changes the information acquisition rate (e.g. White et al., 2010; van Ravenzwaaij et al., 2012). However, note that information sets in these studies, in contrast to our multi-attribute setting, only require lower level integration of information, such as e.g., recognizing a background letter in a pixelated image (Ratcliff et al., 2009).

In the DPRM too, the rate of evidence accumulation and the threshold are separate parameters reflecting deep preferences and DM's tradeoff between cognitive effort and accuracy, respectively. Ruan et al. (2008) show that the rates and the threshold can be jointly identified from choice data alone. The intuition for the joint identification of preference parameters and the threshold is that each of them affect the likelihood equation in a different nonlinear way. This contrasts with preference parameters and the error variance in RUM models (e.g., logit). These two parameters enter the likelihood as linear arguments to the utility function. Therefore making choices more random by increasing the error standard deviation cannot be distinguished from a corresponding multiplicative change of preference parameters. This is why attempts to identify randomness in choice separately from preferences in RUM essentially structure the distribution of heterogeneity (Fiebig et al., 2010).

The DPRM as developed in Ruan et al. (2008) only accomodates binary choice data with ordinal attributes, which constitute only a small proportion of the DCEs conducted in market research. We develop feasible Bayesian inference for the DPRM in larger choice sets that can contain categorical attributes. The algorithm leverages recent results on the convergence of Metropolis-Hastings (MH) sampling using simulated likelihoods (Andrieu and Roberts, 2009). We use GPU-based parallel computing (Nvidia, 2011) for a practicable implementation of the algorithm. The estimation method we develop and its implementation procedure can be applied to other models of decision making with intractable likelihoods.

For the empirical application we conduct a coffee DCE with different incentive groups. We document empirical evidence supporting the assumption of stable underlying (deep) preferences across incentive groups. We show that a Scale MNL that conditions on the same set of preferences but allows for different error scales across different groups replicates the counter-intuitive result that choices become *less* consistent in higher incentive groups (e.g., Hauser et al., 2019). In line with Hauser et al. (2019), the scale MNL estimates for our data suggest that choices become *less* consistent, in higher incentive groups. If we take the scale parameter as a measure of choice quality (Toubia et al., 2004; Evgeniou et al., 2005; Howell et al., 2021), this result suggests that investing in incentive-aligned DCEs is not beneficial, and seems to contradict the prior studies that show increased attention and information processing at high incentives.

To further explore these seemingly contradictory results, we calibrate our DPRM-based framework with the same data. The DPRM estimates imply that, based on a fixed set of preferences, respondents in higher incentive groups attend to more attributes on average which results in more trade-offs and less consistent choices. Given the larger information set to process and hence expecting more effortful choices, these respondents set slightly *lower* thresholds compared to HYP DMs.

We estimate the effective amount of information processing for the respondents, and show that the measure of information processing in our model is a significant predictor of decision time. Moreover, we show that the combination of the increased information set processed and the threshold level chosen by the incentivized respondents results in higher (lower) overall information processing compared to the hypothetical respondents when choosing among more attractive (collectively less attractive) inside options. Therefore our model can explain both the increased attention levels and less consistent choices as incentives increase.

Finally, we show that the decreased consistency resulted from higher incentives does not lead to lower choice quality. We compute the valuation of alternatives in respondents' choice sets based on their estimated deep preferences, and compare the rate at which respondents' choices are in line with their preferences. We find that respondents in higher incentive groups are more likely to choose the alternative with the highest valuation in their choice sets.

The out of sample prediction performance of the DPRM is similar to that of the Scale MNL model and both of them are very close to that of HB logit models trained on the incentive groups separately. Comparing the out of sample prediction performance of the separate HB logit models with the MNL scale and the DPRM model, which pool the data across the incentive groups, strongly supports our assumption that (deep) preferences remain stable as incentives change.

The remainder of the paper is organized as follows. In section 2 the related streams of literature are reviewed. Section 3 explains the DPRM in greater detail and shows how it provides a measure of the effective amount of information processing. In Section 4 we discuss Bayesian inference and data requirements for this model, and lay out how we combine data from a HYP-DCE and an ICA-DCE to study their differences. In Section 5, we show that applying the Scale MNL model to a DCE dataset we collected that has data from different incentive groups replicates that choices become less consistent as incentives increase. We then apply our proposed model to the same dataset to explain the seemingly contradictory results, and show that the inconsistency in incentive-aligned choices measured by Scale MNL is in fact a result of increased decision effort and should not be mistaken for lower quality as compared to hypothetical choices. Finally, we discuss our contributions and avenues for further research in Section 6.

2 Related literature

This paper relates to four streams of literature.

First, at a conceptual level, our approach is related to Swait and Louviere (1993), Louviere et al. (1999), and Louviere and Eagle (2006) who propose that revealed and stated preferences can be traced back to a common set of preference parameters, after accounting for different error scales in a multinomial logit model. However, our operationalization of effort extends beyond the notion of more or less stochastic choice behavior as implied by different error scales, and predicts systematic changes in the preferences for choice alternatives as a function of effort. This is particularly relevant because unless the mechanism underlying larger error scale estimates is not studied, these can be interpreted as a sign of low data quality (Dellaert et al., 1999; Toubia et al., 2004; Evgeniou et al., 2005; Howell et al., 2021), and can mislead market researchers in their decision about investing in incentive-alignment of DCEs. Moreover, practitioners that use consistency as a way to measure choice quality, acknowledge low quality respondents that use simplifying strategies (Orme, 2019), however these cannot be identified in RUM framework due to the inherent confoundedness of preferences and stochasticity. Through modeling information processing and decision effort separately, our framework provides evidence of the employment of these strategies at low incentives.

Second, in terms of capturing decision effort across choices in different incentive settings our framework shares the goal pursued in Cao and Zhang (2020). These authors propose a structural model where the DM chooses her decision effort level after observing the price and realization probability of the choice. The main differences between their model and the framework we propose is that, first, their approach is confined to binary choice between one inside option described by an alternative-specific constant and the outside good, whereas we model multi-attribute, multi-alternative choice. As a consequence, and second, their model does not structure information processing as a function of attributes.

However, in contrast to our approach their model assumes that processing of the price information is automatic and deterministically accurate. Third, just as when estimating separate logit models for HYP and ICA data, their model implies that HYP data does not contain useful information.

Third, is the literature that develops psychological process models, particularly SSM. SSM models have recently become popular in economics and marketing because they explain suboptimal decisions (Krajbich et al., 2014) and well-documented context effects such as similarity (Otter et al., 2008), they have shown to have a better fit than the commonly used RUM models such as multinomial logit (Ruan et al., 2008; Clithero and Rangel, 2013; Chen et al., 2019), and unlike RUM models, give an interpretation of choice stochasticity that goes beyond analytical convenience (Woodford, 2014). They have been used to model experimental data from e.g. decisions on firms' stock performance (Frydman and Nave, 2017), choice of snacks (Baldassi et al., 2020) or food (Krajbich et al., 2010; Fudenberg et al., 2018), to model stochasticity in repeated choice tasks (Agranov and Ortoleva, 2017), or responses to mobile advertisements (Chiong et al., 2024). In spite of their recent popularity, the vast majority of applications have been limited to simple binary choices (cf. Otter et al., 2008)[3]. This limitation precludes them from being used in many interesting real-world economic choice settings. Furthermore, in published empirical applications that are comparable to typical DCEs used in market research, the evidence accumulation process lacks structure beyond the functional form assumptions defining the respective SSM. We argue that structuring the process of evidence accumulation as a function of attributes is important. Such structure provides a natural link between unobserved information acquisition and observed aspects of a choice task and thereby accounts for potential context dependency of the effective amount of information processing before choice under a given threshold. We extend the DPRM, which is in the class of sequential sampling models, to a multi-alternative, multi-attribute setting. This extension makes our model applicable to multi-attribute discrete choice experiments with choice sets with more than two alternatives that are common in economics and marketing.

Finally, the MH algorithm we devise to estimate the multi-alternative DPRM and its implementation using parallel computing contributes to the methodological literature on

[3]In addition, attempts to extend SSM models beyond binary choice assume a pairwise comparison process (e.g. Noguchi and Stewart, 2018; Baldassi et al., 2020).

the estimation of models with intractable likelihood (e.g. Turner and Sederberg, 2014; Holmes, 2015; Holmes and Trueblood, 2018).

3 Model

The DPRM proposed by Ruan et al. (2008) is a race model (e.g., Townsend and Ashby, 1983). Similar to other SSM models, a race model assumes that the cognitive processes that eventually lead to choice can be conceptualized as producing evidence in favor of choice alternatives over time.

The process interpretation of choice makes DPRM flexible in that it can explain different choice behavior from a DM with fixed deep preferences depending on the amount of information she chooses to process in expectation, given prior beliefs about alternatives and attribute configurations in choice sets. Whereas in RUM models alternatives are assumed to be readily evaluated based on the DM's deep preferences and the observed decision maximizes utility that is only partially observable by the analyst, in the DPRM choice is a result of costly information acquisition subject to limited cognitive resources and opportunity cost of time. Therefore, the DM chooses based on a noisy representation of the alternatives' true values. This representation depends on the amount of information processed. To construct this representation, the DM sets a decision threshold and starts accumulating evidence in favor of the alternatives. As soon as evidence in favor of an alternative in the set reaches the threshold, that alternative is chosen and the process stops. With a sufficiently high threshold, the DM accumulates enough evidence to consistently choose the best alternative according to her deep preferences. But generating evidence is costly. An intuitive way to think about the cost of evidence accumulation, is through decision time[4]. The threshold optimizes the trade-off between cognitive effort and accuracy given expectations about what the DM may see in a choice set. Hence, the threshold is *not* endogenous to what is offered in a *specific* choice set, but it is a satisficing criterion depending on the incentive level and the cost of time and cognitive processing. However, as we will show, the effective amount of information processed before a decision occurs is endogenous to the characteristics of alternatives in the choice set.

[4]In the empirical application in Section 5 we show that the amount of information processing measured by the model significantly predicts decision time.

In the absence of cognitive costs, the DM would accumulate an inifinite amount of evidence and deterministically choose the best alternative as a function of deep preferences for attributes. As such, the DPRM does not equate observed randomness in choice with unobserved aspects of utility but motivates randomness from imperfect information processing. This change in perspective seems inherently attractive in the context of DCEs, where unobserved aspects of utility are hard to motivate given that the DM and the analyst ex ante share the same information set.

When we think about the process of building up the valuation of a multi-attribute product, a natural starting point is the valuation of the attributes. In fact, the reason that the DM cannot immediately act on her true preferences in the DPRM is the assumption that the evaluation and interpretation of the attributes in the context of a particular choice set require cognitive effort. When an attribute level is processed, an individual piece of evidence accrues according to a Poisson process in the DPRM. This process generates discrete units of evidence called 'hits'. The rates of the Poisson processes are an exponential function of the deep valuation of the attributes by the DM. The more valuable an attribute is to the DM, the higher the attribute's rate and the more often hits are produced in favor of alternatives that possess that attribute. Equation 1 shows how rates are computed for attribute j in a binary choice set[5] where x_{ij} encodes attribute j's level for alternative i and β_j is the attribute's partworth.

$$\rho_{1j} = e^{x_{1j}\beta_j} - 1, \qquad \rho_{2j} = e^{x_{2j}\beta_j} - 1 \tag{1}$$

Subtracting 1 in Equation 1 accomplishes that attribute levels assigned a value of zero do not generate hits at all. The influence of these baseline attribute levels comes through a shared rate that we discuss further below.

An important feature of the choice process in the DPRM is that the DM allocates hits produced by thinking about a particular characteristic to all alternatives in a set that share this characteristic. In other words what is actually effortful for the DM is not looking at different alternatives and realizing what attributes they have, but evaluating an attribute based on her deep preferences in the context of a choice set. Therefore, when an

[5] For ease of exposition and without loss of generality we explain the characteristics of the model in a binary choice setting. The model extends to larger choice sets in conceptually the same way.

attribute is processed, a "shared hit" is produced in favor of all the alternatives that have the same level of that attribute. In a binary choice, what the alternatives have in common in an attribute contributes to a shared Poisson process that generates simultaneous hits in favor of both of the alternatives, and the part of the attribute they uniquely possess contributes to their unique Poisson processes. This is shown formally in Equation 2, where ρ_{sj} ("s" stands for shared) is the rate of the shared Poisson process of attribute j, and ρ_{u1j} and ρ_{u2j} ("u" stands for unique) are respectively the rates of alternative 1 and alternative 2's unique processes.

ρ_{sj} is the minimum of the rates of the attribute levels of the alternatives. In case of an ordinal attribute, it is the component of the attribute that both alternatives possess. The unique component of the attribute, that is possessed by the alternative with the higher attribute level, will contribute to that alternative's unique rate and produce hits in favor of that alternative only. In case of a categorical attribute, if the alternatives share the same level of the attribute, they share the hits generated from processing it, and if they possess different levels of that attribute, $\rho_{sj} = 0$ and processing those attribute levels generates unique hits in favor of the respective alternatives.

$$\rho_{sj} = min(\rho_{1j}, \rho_{2j}), \qquad \rho_{u1j} = \rho_{1j} - \rho_{sj}, \qquad \rho_{u2j} = \rho_{2j} - \rho_{sj} \tag{2}$$

As a result of the shared processes, the DPRM captures attribute-level similarity between the alternatives both along a priori ordered, and hence more comparable, and inherently categorical attributes that are less comparable. Therefore choices in the DPRM do not exhibit the IIA property. According to IIA, the ratio of choice probabilities for two alternatives in the choice set is independent of other alternatives included in the choice set. Consequently, a new product entering the market takes market share from other products proportional to their overall utility levels and regardless of their similarity to the new product. IIA is criticized in the literature as being unintuitive and assuming IIA in choice may lead to bias e.g. in market simulation (Debreu, 1960; Hausman and Wise, 1978; Zeithammer and Lenk, 2006; Dotson et al., 2018).

The concept of shared processes to account for similar aspects of the alternatives contrasts with modeling relative evidence as the sole driver of choice in e.g. drift diffusion models. Whereas discounting the shared evidence might seem reasonable as it suggests that the DM rationally ignores attributes that equally affect all the alternatives (in this example price), it cannot differentiate between choice sets with equally good and equally suboptimal inside options, when there is an outside option. However, to the extent that the DM could realize the relative value of the common aspects of the inside options relative to the value of the outside option, she may arbitrarily choose one of the inside options in the former case, but opt out in the latter. The implications of a higher shared rate of inside goods for choice also aligns with the quick resolution of the implied approach-approach conflicts (Lewin, 2013) between inside alternatives.

The DM constructs valuations of alternatives by integrating the incremental valuation of attributes. The model thus combines attribute-specific Poisson processes to yield the alternative-specific processes. More precisely, the alternative-specific Poisson processes are superpositions of the attribute-specific processes and their rates are computed by summing over the attribute-specific rates that pertain to their respective alternatives (Equation 3). Note that there is one shared process for each combination of alternatives in the choice set. Hence in a multi-alternative choice set, there will be more than one shared rate. The overall shared rate (here ρ_s) has an intercept parameter β_0 that captures common characteristics that are not shown in the choice set and are assumed to be similar in all the alternatives. The intercept also accounts for the shared aspect of the worst levels of all the attributes (cf. Equation 1).

$$\rho_s = e^{\beta_0} + \sum_{j=1}^{J} \rho_{sj} \qquad \rho_{u1} = \sum_{j=1}^{J} \rho_{u1j} \qquad \rho_{u2} = \sum_{j=1}^{J} \rho_{u2j} \tag{3}$$

To demonstrate information processing in the DPRM and the implications of unique and shared processes we provide a stylized example of a binary choice set. Table 1 shows a choice set of laptops based on brand and price. The values in the table correspond to the values of the attributes weighted by the DM's preferences.

Parthworths	Brand1	Brand2	Price
Laptop1	0.15	-	0.3
Laptop2	-	0.1	0.3

Table 1: A binary choice of laptops based on brand and price

Both laptops have the same price and Laptop1 has a slightly more attractive brand. The identical price contributes to the shared process[6] that has a rate of $\rho_s = e^{0.3} - 1$. The two brands contribute to the unique processes with rates $\rho_{u1} = e^{0.15} - 1$, and $\rho_{u2} = e^{0.1} - 1$ generating hits in favor of the respective laptops. Realizations of the choice process can be simulated by using the normalized rates $P_s = \frac{\rho_s}{\rho_s+\rho_{u_1}+\rho_{u_2}} = \frac{e^{0.3}-1}{e^{0.15}+e^{0.1}+e^{0.3}-3}$, $P_1 = \frac{\rho_{u_1}}{\rho_s+\rho_{u_1}+\rho_{u_2}} = \frac{e^{0.15}-1}{e^{0.15}+e^{0.1}+e^{0.3}-3}$, $P_2 = \frac{\rho_{u_2}}{\rho_s+\rho_{u_1}+\rho_{u_2}} = \frac{e^{0.1}-1}{e^{0.15}+e^{0.1}+e^{0.3}-3}$ as probabilities of a multinomial distribution to generate hits. In this example alternative-specific processes are composed of a single attribute ($P_s = P_{Price}$, $P_1 = P_{B1}$, and $P_2 = P_{B2}$).[7]

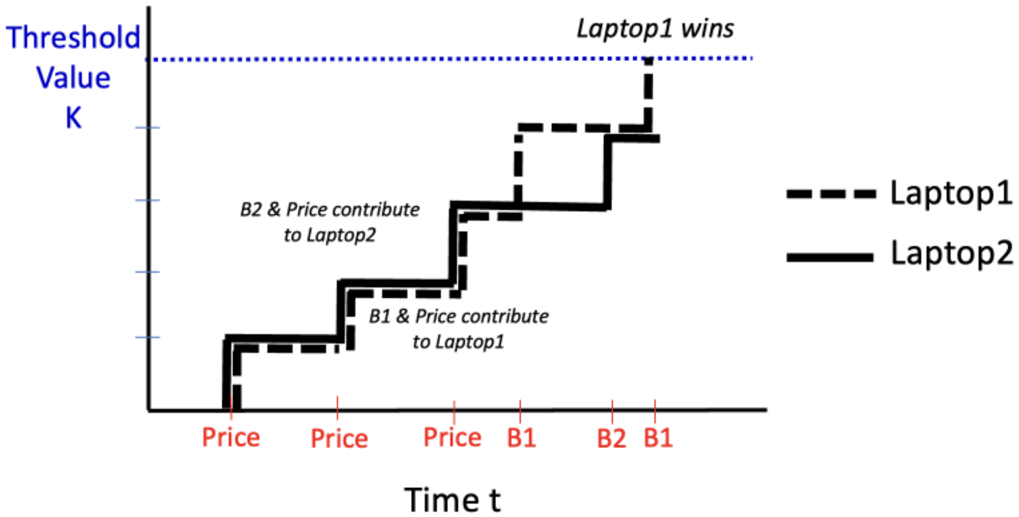

Figure 1: An illustration of a DPRM decision process simulated from the choice set in Table 1

The decision threshold commits the DM to processing a total number of hits in a certain range as explained next. For instance by setting K = 5, she commits to processing at least 5 hits (unique or shared), and at most 9 hits (4 unique hits in favor of each alternative and the 5th hit in favor of the winner). Figure 1 shows a realization of a choice process by a DM with deep valuations as shown in Table 1 and a threshold value of K = 5. The dashed and the solid traces show how the states of the hit-counters for laptops 1 and 2 in this example develop over time [8] The horizontal dotted line in the figure indicates the

[6] e^{β_0} is assumed to be negligible in this example.

[7] If an alternative-specific or a shared process is the superposition of multiple attribute-specific processes, we can additively decompose the hit-generating probabilities into their attribute-specific components and generate hits as coming from a specific attribute in the simulation.

[8] As we discuss in detail in Section 4.4, we simulate the hit realizations by draws from a multinomial distribution indexed by the normalized alternative-specific rates. Equivalently, they can be simulated by assuming that the realized hit at each stage of the race comes from the process with the smallest "arrival" time, where arrival time of hits is exponentially distributed with rates equal to the rates of the Poisson processes (see e.g. Bradlow and Park, 2007, for an application of this relationship in the context

evidence threshold an alternative needs to reach for a choice to occur. Accumulation of hits continues until the evidence in favor of one of the options, in this realization Laptop1, reaches the threshold $K = 5$ first and this alternative is chosen.

Threshold	Average Hits	Tie Probability	Accuracy
1	1	0.5661	0.547
5	6	0.1927	0.636
10	11	0.1138	0.6868
15	18	0.0876	0.727
20	24	0.0693	0.7678
25	30	0.0598	0.7875
30	36	0.0494	0.8104
60	72	0.0245	0.8926
100	120	0.0105	0.9432
150	181	0.0059	0.9727
200	241	0.0026	0.987
300	362	3e-04	0.9964
500	603	0	0.9998
1000	1205	0	1

Table 2: Choice behavior as a function of threshold simulated for the choice set in Table 1

Table 2 characterizes the DPRM choice process as a function of the threshold parameter for the choice set in Table 1. Table 2 shows that when the threshold increases, the average number of hits accumulated increases. Based on the DM's limited cognitive capacity, hits are assumed to be processed sequentially (Anderson et al., 2004). Therefore decision time is expected to increase in the number of hits. Accumulating more hits makes choices more accurate and makes ties less probable. Consequently choices become deterministic at high thresholds and attribute-level similarity, which may result in ties at lower thresholds will matter less at high thresholds. Intuitively, in a less careful decision, the DM would perceive both laptops equally attractive, or "good enough", with a higher probability.

Ties in the accumulation to bound process in the DPRM highlight the satisficing properties of the model. Satisficing, as modelled by Stüttgen et al. (2012), is a process that alternates between searching for information about the alternatives and evaluating them as satisfactory or unsatisfactory. The DPRM accommodates a sequence of information processing that stops once an alternative is identified as satisfactory. The threshold sets a deterministic satisficing rule in line with Simon's original theory.

of modeling bidding behavior). However the sum of these times does *not* uniquely correspond to the DM's response time because of the presence of shared hits and random tie-breaking.

Even though setting the threshold exogenously specifies the range of processed hits before the decision, attribute (non-)attendance and the structure of the choice set play a systematic role in determining the number of hits processed. This is an important consequence of shared hits, which makes information processing endogenous to the specifics of the choice set and the effective information set processed by the DM. Processing similar alternatives results in shared hits. As a consequence, the total number of hits processed in a choice between similar alternatives is smaller than between dissimilar alternatives. Under the usual design rules in DCEs, attending to more attributes translates into less similar and directly comparable alternatives and hence increases the number of processed hits before a choice in expectation. In sum, decisions between comparable (vs non-comparable) alternatives, need less information processing, are faster, and are more consistent in the DPRM under the same rates and threshold. The same applies to decisions made based on fewer attributes (vs more attributes).

To illustrate this, consider Table 3, which shows partworths in a choice between laptops based on brand, color, and price under two different information acquisition scenarios. Price is reverse-coded so that higher values correspond to lower prices[9], therefore the DM prefers laptop2 based on higher overall preference values.

In scenario "a", the DM attends to all the attributes and hence integrates evidence for alternatives over a higher dimensional space (constructs alternative-specific utilities from a higher dimensional space). In scenario "b" she ignores brand and color and chooses solely based on price.

a)

Partworths	Brand1	Brand2	Silver	Black	Price
Laptop1	5	-	1.5	-	2
Laptop2	-	5.2	-	2	1.5

b)

Partworths	Price
Laptop1	2
Laptop2	1.5

Table 3: A binary laptop choice under 2 scenarios: a)full attribute attendance, b)attribute non-attendance

We simulate 20,000 choices in each of these choice scenarios, for a DM with a fixed threshold parameter of K = 15. Table 4 shows the rounded average number of hits

[9] In general, for the DPRM, ordinal attributes are coded such that higher values in the design correspond to better attribute levels. The worst level of an ordinal attribute will be the baseline and = 0.

accumulated before the choices, the variance in choices, and the accuracy of choices (rate of choosing the better alternative based on the DM's preferences, here Laptop2) in each scenario.

The first insight from the results in this table is that attending to more attributes, and hence integrating utility over a higher dimensional space and hence less comparable alternatives, results in more information processing (as reflected by the number of hits processed), and less consistency in choices (as reflected by higher variance in choices). Second, even though ignoring attributes results in higher consistency, it can lead to very low quality choices. In scenario b, by ignoring all the attributes except price, the DM is faced with simple choices between directly comparable alternatives based on price and consistently chooses the cheaper alternative, which is the worse alternative overall, based on her preferences (in Table 3).

Of course, in each of these cases, choice becomes more difficult when alternatives are closer to each other in value. This, all else equal, would be reflected in lower consistency in both cases and higher number of total hits processed, except for special cases.

Choice scenario	a) Full attendance	b) Non-attendance
Avg. hits	25	15
	(0.04)	(0.03)
Choice variance	0.21	0
Accuracy	0.69	0

Table 4: Average (standard errors) hits processed, variance (consistency) of choices and accuracy (quality) across different attribute attendance scenarios presented in Table 3

These choice patterns are in line with well documented behavior in the experimental literature. For example Tversky and Russo (1969) and Zhang and Markman (2001) show that subjects can distinguish more easily between comparable than non-comparable alternatives and Tversky (1969) and Bettman et al. (1990) argue that tradeoffs in choice make choices complex and resolving them needs more cognitive effort by the DM.

The above examples highlight the structure of information acquisition in the DPRM. This structure plays a key role in changing decision behavior as the DM processes more hits. The DM generates hits by evaluating the attributes. Shared hits stem from processing levels of attributes that two or more alternatives have in common, and capture attribute-level similarity. The hits generated by attribute evaluation accumulate to that of alternatives.

Conditional on attending to a fixed set of attributes, when the DM processes more hits, first, decision converges to a deterministic choice of the best alternative based on deep preferences; second, similarity effects vanish (without reverting to IIA), because the less attractive alternatives will be chosen progressively less frequently regardless of their similarity to more attractive ones; and third, choices become more alternative-based (vs. attribute-based), because the DM accumulates hits across more attributes. These properties make the total hits generated in the DPRM process a plausible measure of the exerted effort that causes meaningful changes in choice probability, conditional on the deep preferences of the DM.

To illustrate, Table 5, presents a ternary choice set of laptops based on two ordinal attributes: memory(RAM) and processor speed(CPU). Similar to the previous examples, the values in the table correspond to the values of the attributes weighted by the DM's preferences. Laptop2 has the highest total value, but Laptop1 and Laptop3 have respectively the highest RAM and the fastest CPU. Figure 2 shows the evolution of choice probabilities as a function of theshold value (K) for this choice set.

Partworths	RAM	CPU
Laptop1	3.2	0
Laptop2	2	1.7
Laptop3	0	2.7

Table 5: A ternary choice of laptops based on two ordinal attributes memory (RAM) and processor speed (CPU)

In spite of its highest overall value, the second laptop has the lowest choice probability in low effort levels. The reason is that choice at low thresholds is a result of processing attributes in isolation. At K=1, for instance, the DM processes one hit, which is a piece of evidence associated with one of the two attributes. If the hit is generated processing RAM, Laptop1 will be chosen with a high probability, and if it is caused by processing CPU, Laptop3 is chosen with a high probability. Laptop2 is chosen only as a result of a random tie-break in case of a shared hit at either of the above scenarios.

But as the threshold increases and the DM exerts more effort, she integrates information across the two attributes and realizes the high overall attractiveness of the second laptop. Therefore the choice probability of Laptop2 increases in K, and at K equal and higher than 5, Laptop2 is chosen with the highest probability.

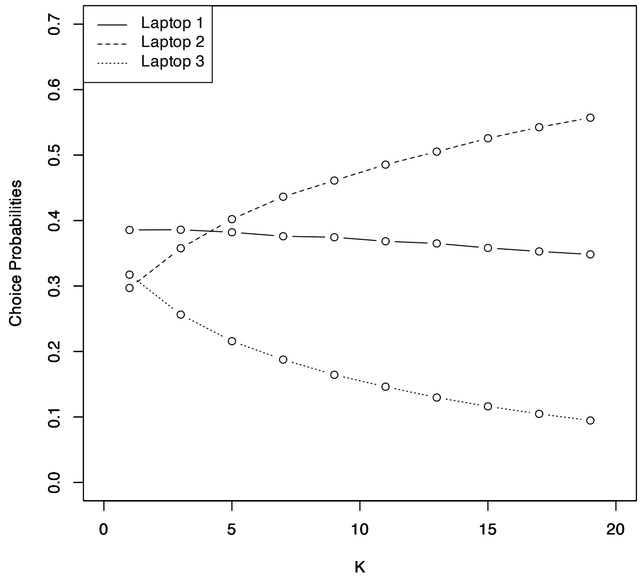

Figure 2: DPRM choice probabilities as a function of threshold value for the choice task in Table 5

The evolution of choice probabilities in Figure 2 clearly suggests more deterministic choices (note the increasing gap between the choice probabilities) and the transition from attribute-based to alternative-based processing as more hits are processed. In fact the order of choice probabilities according to Elimination by Aspects (EBA) (Tversky, 1972), which assumes a purely attribute-based choice rule, is P(Laptop1) > P(Laptop3) > P(Laptop2), consistent with DPRM choice at low effort. The choice probability ordering of multinomial logit, which assumes alternative-based processing is P(Laptop2) > P(Laptop1) > P(Laptop3), and in line with the DPRM choice at high effort. Attribute-based processing at low and holistic processing at high effort levels in the DPRM echo the experimental literature results that show that subjects choose the former strategy to minimize effort and the latter to maximize accuracy (e.g. Russo and Dosher, 1983; Lohse and Johnson, 1996). The process of integrating hits from processing attributes to evaluate alternatives is consistent with the distinction between attribute-based and alternative-based processing in behavioral decision theory (Bouyssou and Pirlot, 2004; French, 1986). In other words, the DPRM captures choice strategies across different effort levels that potentially give rise to different preference orderings, even under invariant deep preferences.

4 Estimation

Figure 3 shows the DPRM's data generating process in the form of a directed acyclic graph. Individual i's choice (y_i) in the DPRM is generated by a set of attribute preferences β_i, which determine the evidence accumulation rates, and the evidence threshold K_i. These parameters follow hierarchical prior distributions indexed by their respective hyperparameters $\bar{\beta}, \Sigma_\beta$, λ.

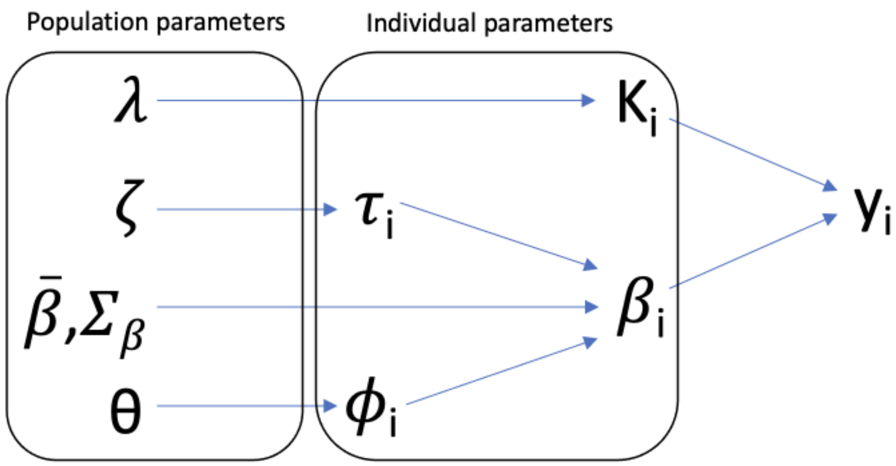

Figure 3: Directed acyclic graph of the DPRM

The two additional parameters τ_i and ϕ_i operating on β_i respectively capture attribute attendance and the least attractive level of each categorical attribute. The former allows for the possibility that respondents only integrate information over a subset of atttributes, and the latter enables inference about the DM's preference for categorical attributes—both are features we add to the original DPRM developed by Ruan et al. (2008). In the following two sections we elaborate on the role of these two sets of parameters in the model.

4.1 Attribute attendance

The parameter vector τ_i, with hierarchical prior indexed by ζ, is a vector of indicators specifying the size of the information set—the subset of attributes that the DM attends to in the choice process. In this model, each individual $i = 1, \ldots, N$ is characterized by a logical vector τ_i that points to the active variables in individual i's likelihood function, i.e., to the attributes that determine this individual's rates of evidence accumulation. Thus, the vector τ_i reflects individual i's information set. The collection of logical vectors across all individuals $\{\tau_i\}_{i=1}^N$ informs the parameters in the hierarchical prior distribution of information sets, ζ. In the simplest case, ζ indexes a vector of independent Bernoulli distributions (see Gilbride et al., 2006). The elements in τ_i determine the variable specific (parameter specific) hierarchical prior setting for individual specific parameters β_i. For the p-th element τ_{ip} equal to one, the marginal hierarchical prior for β_{ip} is then $log(\beta_{ip}) \sim N(\bar{\beta}_p, \Sigma_{pp})$, where this distribution conditions on all the other elements in the preference vector (β_{i-p}), and for τ_{ip} equal to zero this marginal prior becomes $log(\beta_{ip}) \sim N(\bar{\beta}_p - c, c^{-1}\Sigma_{pp})$, where this distriution conditions on β_{i-p}, and c is a large (relative to the likely support of $N(\bar{\beta}_p, \Sigma_{pp})$) positive constant, say, 1000, such that $\int_{-\infty}^{+\infty} \beta_{ip} N(log(\beta_{ip})|\bar{\beta}_p - c, c^{-1}\Sigma_{pp}) dlog(\beta_{ip}) \approx 0$, and $P(\beta_{ip} \approx 0|\tau_{ip} = 0) >> P(\beta_{ip} \approx 0|\tau_{ip} = 1)$. For $\beta_{ip} = 0$ the corresponding attribute level does not contribute to the race at all, reflecting that this level has been ignored by the respondent when making the observed choices.[10] We conceptualize attribute attendance as paying attention to all the levels of an attribute. Thus in case the DM ignores a categorical attribute, all the elements in τ_i corresponding to the levels of that attribute equal zero.

4.2 Categorical attributes

Because the DPRM is a race model, all rates need to be positive (e.g., Ruan et al., 2008) and thus the rates from Equation 1 are computed as functions of attributes, their values, and β_i, which has a log-normal distribution and is therefore non-negative. This is in contrast to e.g. logit that has unconstrained preference parameters. In the logit, the baseline level of categorical attributes can be selected arbitrarily by the analyst. Preference for the other categories are then measured relative to the baseline category. A category less attractive than the baseline gets a negative partworth. In the DPRM, to implement the rates' non-negativity constraint for ordinal attributes, the respective lowest attribute levels are normalized to zero. For categorical attributes, however, setting the baseline is not as straightforward as in logit. Baseline levels must be the least attractive levels but since categorical attributes do not have a preferred direction, the least attractive categories

[10]It is true that this formulation cannot per se distinguish between attribute levels of no intrinsic, deep value to the respondent and attributes a respondent just did not attend to. However, under the assumption of invariant deep valuations of attributes, differences in the distribution of τ_i between HYP-DCE and ICA-DCE, i.e., signed differences between ζ^{HYP} and ζ^{ICA} are strongly consistent with the idea that attributes that could be valued are simply ignored, depending on the incentive level.

cannot be determined a priori. Therefore we infer the least attractive level of each categorical attribute from the data. For each individual $i = 1, \ldots, N$, $\phi_i = \{\phi_{i1}, ..., \phi_{iP_{cat}}\}$, where P_{cat} is the number of categorical attributes. That is, ϕ_i is a vector of elements each indicating the baseline level of a categorical attribute. The collection of these sets across all individuals $\{\phi_i\}_{i=1}^{N}$ informs the parameters in the hierarchical prior distribution of the baseline indicators, θ.

Similar to τ_i, ϕ_i affects β_i through the hierarchical prior. The j-th element of ϕ_{ip} will be equal to one if the j-th category of the p-th categorical attribute is the least attractive, and hence baseline, level of that attribute. Consequently, the marginal prior of the baseline category of the attribute becomes $N(\bar{\beta}_j - c, c^{-1}\Sigma_{jj})$, where this distribution conditions on all the other elements in the preference vector β_{-j}, and $c = 1000$.

4.3 Bayesian inference

We use a Bayesian hierarchical specification similar to Ruan et al. (2008). These authors explicitly compute the individual level DPRM likelihood, conditional on a threshold value K and normalized Poisson rates. For example, in a binary choice we have P_1, P_2, and P_s representing normalized rates for the unique counters 1 and 2, and the shared counter, respectively. Choice probabilities are obtained by summing over all possible combinations of hits, i.e., all possible evidence accumulation paths (races), such as that in Figure 1, leading to the choice of the observed alternative conditional on the threshold value K, weighted by the probability of that race occuring. The probability of choosing the first alternative in a binary choice $P(y = 1)$ is computed in Equations 4, 5, and 6.

$$P(tie) = \sum_{i=1}^{K} \frac{(2K-i-1)!}{(i-1)!(K-i)!(K-i)!} P_s^i P_1^{K-i} P_2^{K-i} \tag{4}$$

$$\begin{gathered} P(1) = \sum_{k=0}^{K-1} \frac{(K+k-1)!}{k!(K-1)!} P_1^K P_2^k + \\ \sum_{i=1}^{K-1} \sum_{k=0}^{K-i-1} \left[\frac{(K+k-1)!}{(i-1)!(K-i)!k!} + \frac{(K+k-1)!}{i!(K-i-1)!k!} \right] P_s^i P_1^{K-i} P_2^k \end{gathered} \tag{5}$$

$$P(y = 1) = P(1) + 0.5P(tie) \tag{6}$$

$P(y=2)$ is computed in the same way. Equation 7 exhibits the likelihood of the model for individual i's choices in choice sets $\{L_t\}_{t=1}^{T}$, where Θ is the set of the individual parameters and there are T observations per individual.

$$\ell(Y_i|\Theta_i) = \prod_{t=1}^{T} \prod_{l \in L_t} P(y = l|\Theta_i)^{I(y_{it}=l)} \tag{7}$$

For inference we assume a log normal distribution and a shifted Poisson distribution respectively for preferences (β) and the threshold parameter (K). The population parameters of preferences and the threshold follow a normal-inverse Wishart, and lognormal subjective prior, respectively. Equation 8 lays out the distributional assumptions of all the model parameters.

$$\begin{gathered} K \sim Poisson(\lambda) + 1, \quad log(\beta) \sim MVN(\bar{\beta}, \Sigma_\beta), \\ \tau \sim B(\zeta), \; and \quad \phi \sim Mult(\theta) \\ where \;\; \bar{\beta}|\Sigma_\beta \sim N(B, A^{-1}\Sigma_\beta), \;\; \Sigma_\beta \sim IW(\nu, V) \\ \zeta \sim Beta(a, b), \quad \theta \sim Dirichlet(D), \\ and \;\; ln(\lambda) \sim N(\overline{\lambda}, \sigma_\lambda) \end{gathered} \tag{8}$$

In case of binary and ternary choice sets, computing the likelihood as per Equation 7, inference can be done straightforwardly (see Ruan et al., 2008). However, as the number of alternatives in the choice set increases, the number of Poisson counters and thus the number of possible races that need to be summed over for likelihood evaluation explodes such that computing the exact likelihood for choice sets with more than three alternatives becomes prohibitive. To adapt the DPRM to larger choice sets, we circumvent the exact computation of the likelihood by developing an algorithm based on the Pseudo-Marginal MH sampler originally proposed by Beaumont (2003). Instead of calculating the likelihood at each iteration exactly, this method relies on a Monte Carlo approximation to the likelihood based on a finite number of simulated races.

At each MCMC iteration, given the proposed values of the individual parameters, the algorithm simulates M choices from the DPRM to approximate the likelihood by taking the expectation over the simulated choices as shown in Equation 9. The collection of $\bar{P}(y|\Theta)$ for the repeated measurements are then used instead of their exact counterparts

$P(y|\Theta)$ to approximate the likelihood in Equation 7, which in turn is used by the MH algorithm to sample from the individual level posterior.

$$\bar{P}(y|\Theta) = M^{-1} \sum_{m=1}^{M} P(y|z_m) \tag{9}$$

where Θ is the proposed set of individual parameters, $z_m \sim p(z|\Theta)$ obtained by forward simulating races in parallel. The following example demonstrates the forward-simulation of the DPRM's choice process to generate races $\{z_m\}$.

4.4 Forward-simulation of a DPRM race

Table 6 shows a ternary choice set with three equally attractive brands, where two of them have the same color. Processing Color2 generates simultaneous hits in favor of Alt.2 and Alt.3, while processing Color1 generates hits uniquely in favor of Alt.1. Similarly, processing each brand generates hits in favor of its respective alternative. The frequency of the hits can be computed by the attributes' normalized rates. For example, the relative frequency of the hits generated from processing Color2 is $P_{Color2} = \frac{\rho_{Color2}}{\rho_{Brand1}+\rho_{Brand2}+\rho_{Brand3}+\rho_{Color1}+\rho_{Color2}} = \frac{e^{0.2}}{e^{0.2}+e^{0.2}+e^{0.2}+e^{0.1}+e^{0.2}}$.

Attributes	Brand1	Brand2	Brand3	Color1	Color2
Alt.1	0.2	0	0	0.1	0
Alt.2	0	0.2	0	0	0.2
Alt.3	0	0	0.2	0	0.2

Table 6: A ternary choice based on brand and color attributes

Given a threshold value K, a DPRM race can then be simulated by drawing from the multinomial distribution indexed by these attribute-specific probabilities:
$Mult(P_{Brand1}, P_{Brand2}, P_{Brand3}, P_{Color1}, P_{Color2})$. Table 7 shows consecutive hit draws from the multinomial distribution indexed by the attributes-specific relative rates and the allocation of hits to their respective alternatives. Figure 4 depicts the race implied by the hits in Table 7, when the DM's threshold is K=5.

In this race, the DM initially receives a hit in favor of alternative 2 by processing its brand. She continues by processing color2, whereupon both alternatives 2 and 3 receive a

Alternatives↓ Hits→	B2	C2	B3	B3	B1	B2	B2	B3	C2	Total
Alt.1	0	0	0	0	1	0	0	0	0	1
Alt.2	1	1	0	0	0	1	1	0	1	5
Alt.3	0	1	1	1	0	0	0	1	1	5

Table 7: Forward simulation of a hit sequence resulted from processing the choice set in Table 6

hit. She then processes Brand1, Brand3, and the processing goes on until upon processing Color2 for the third time, both alternatives 2 and 3 reach the threshold simultanously and the race finishes in a tie. The choice between the alternatives that tie is random.

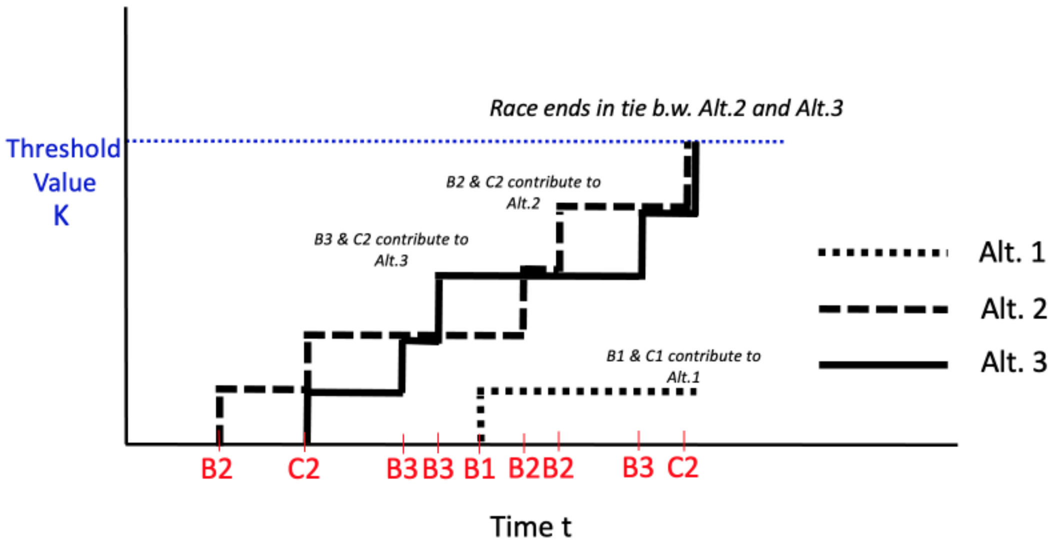

Figure 4: An Illustration of the race associated with the hit sequence in Table 7

Andrieu and Roberts (2009) demonstrate that using the approximated likelihood $\bar{P}_{z_m}(y|\Theta)$ instead of $P(y|\Theta)$ in the MH algorithm enables exact inference from the posterior, provided that the current likelihood at each iteration is approximated using the z_m's "recycled" from the previous iteration. This sampling strategy has been coined "Pseudo-Marginal" by Andrieu and Roberts (2009) because unobserved races are integrated out, but not fully, at each iteration of the sampler.

Even though implementing this algorithm in `Rcpp` greatly reduces the runtime, to get a likelihood approximation smooth enough that facilitates faster convergence of the MCMC we need a very high number of monte-carlo simulations (M) at each iteration. This makes the algorithm computationally intensive. In order to achieve convergence in a reasonable time, we simulate the races at each iteration in parallel. We use `CUDA C` to utilize the graphics processing unit(GPU) for parallelized simulation of DPRM races (see e.g. Lin et al., 2019, for a use of Cuda parallelization for approximating densities). The parallel algorithm takes advantage of the fact that the parameter update within a MCMC iteration is done independently across individuals. Therefore the algorithm simulates $N \times T \times M$ races in parallel to update the individual parameters at each iteration. Comparing the posteriors obtained from this method and those from the MCMC that evaluates the exact likelihood[11] shows that with $M = 300$ Monte-Carlo approximations, the Pseudo-Marginal method converges in a reasonable number of iterations. We use the same value of M for estimation in this paper. That is, for a dataset with $N = 1000$ respondents, and $T = 10$ choice observations per respondent, the algorithm simulates $3,000,000$ races at every iteration in parallel to update the individual parameters (for more details on conditional updating of the individual parameters, see Appendix A). Parallelization has made the estimation faster roughly by a factor of 100. Given the high number of simulations per iteration, parallel implementation is an essential part of a feasible estimation of the model.

4.5 The bridge framework

The DPRM can separately measure preferences and decision effort. We exploit this feature to develop a procedure that uses a combination of HYP and ICA data for inference. Figure 5 provides a summary of the procedure for a combination of HYP-DCE data from 1,...,V respondents and ICA-DCE data from 1,...,W respondents. The procedure could also be used as a cost-effective improvement of an ICA-DCE with W respondents, as it increases the efficiency of the estimation by utilizing cheap HYP data. At the individual level, we infer deep preferences for product characteristics β, evidence thresholds K, and finally indicator vectors ϕ and τ that respectively point to the baseline level of categorical attributes and the information set encoded when making a choice (see the column 'Individual Level Parameters' in Figure 5).

The deep preferences of individuals in both groups are assumed to come from the same distribution, therefore both the data from the HYP-DCE and the ICA-DCE will be used to inform *one* population distribution of deep valuations of attributes, indexed by $\bar{\beta}, \Sigma_\beta$ in Figure 5. It follows that the ϕ_i vectors indicating the least attractive level of the categorical attributes also have the same distibution under both incentive schemes. However, we allow for different threshold levels in the HYP and the ICA settings through parameters λ^{HYP}

[11]Note that for binary choice, exact likelihood computation is feasible. Hence we could use it as a benchmark to test the performance of the Pseudo-Marginal method at different values of M.

and λ^{ICA} indexing the respective hierarchical prior distributions of evidence thresholds K and different distributions of information sets under HYP and ICA; ζ^{ICA} and ζ^{HYP} index hierarchical prior distributions that specify probabilities for attributes entering the information set. Furthermore, as respondents perceive the no-choice option (NC) differently under HYP than under ICA, the preference for no-choice $\beta_{i_{NC}}$, and the baseline brand indicator $\phi_{i_{BR}}$[12] of the HYP respondents inform $\overline{\beta}_{NC}^{HYP}$ and $\overline{\theta}_{BR}^{HYP}$, and those of the ICA respondents inform $\overline{\beta}_{NC}^{ICA}$ and $\overline{\theta}_{BR}^{ICA}$.

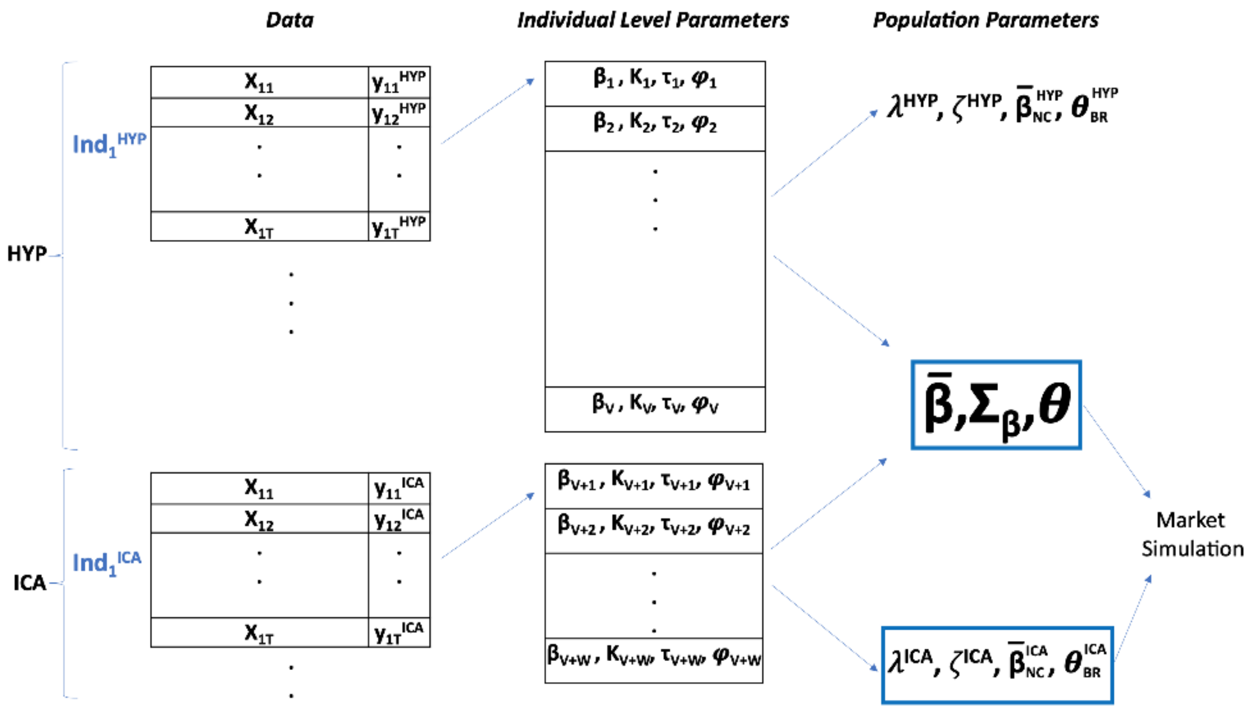

Figure 5: The bridge framework

Using the marginal posterior of the hierarchical prior

$$p(\lambda^{ICA}, \zeta^{ICA}, \overline{\beta}_{NC}^{ICA}, \overline{\theta}_{BR}^{ICA}, \bar{\beta}, \Sigma_{\beta}, \theta | Y^{HYP}, Y^{ICA}) \tag{10}$$

it will be possible to compute conterfactual choice distributions, i.e., market shares under the effort level induced by ICA, but harnessing the information from both the HYP-DCE and the ICA-DCE.

[12]We code NC as another brand in the data.

5 Empirical application: specialty coffee

We conduct a choice experiment, where 1277 respondents were randomized into 5 different treatment groups: The incentive-alignment mechanism was to make respondents "live with their choices" (Ding et al., 2005) by making them buy one of their choices if incentive-alignment is realized. The realization probability was different for each group.

Conducting an incentive-aligned choice experiment at this scale is challenging due to several reasons. This makes a framework that can pool across hypothetical and incentivized choices empirically useful. First, there are considerable additional monetary and non-monetary costs due to incentive-alignment:

Design constraints: As opposed to hypothetical experiments, the per respondent cost of the experiment is not the participation fee, but the cost of the product. Therefore, the cost of DCE's for expensive products can become prohibitive for large experiments[13], unless the product is given to the respondents based on a lottery (see e.g. Ding, 2007); a mechanism that decreases the realism of the experiment. In addition to the cost, product availability needs to be ensured. Every product shown in the experiment should be available for purchase at the prices presented to the respondents. This can be difficult to achieve in practice due to stockouts and price volatility. Suppliers often do not accept to hold inventory or to commit to a certain price to sell products at a later date. Moreover, product availability affects experiment design. Constraints on the combination of attributes can limit variation in the available choice sets for large experiments.

Implementation constraints: Market research companies are reluctant to implement large-scale incentive-aligned experiments because running such experiments requires operations that goes beyond the scope of their usual operation. For instance, the company needs to handle the communications with the supplier for product purchases, shipping the products to the respondents, and the monetary transactions. Moreover, it has to handle the data privacy issues that arise when the respondents are asked to provide their addresses for product shipment.

[13]We are grateful to German Science Foundation, for the grant that covered the costs associated with our experiment.

5.1 Experimental design

Specialty coffee[14] was a desirable product category for our experiment because it would allow us to generate a large number of product profiles with different attributes, we could find coffee suppliers that offered a wide variety of products, it could be shipped to the respondents at a reasonable cost, and it had a price range not too high to make the experiment infeasible and not too low to make the choice trivial. The experiment was conducted online in Germany, where respondents were recruited from a panel of a market research company. Respondents were initially screened based on their coffee consumption habits. Screening questions were designed to ensure that the respondents were coffee drinkers, and they purchase coffee to make at home. A number of filler questions (e.g. questions about meat consumption) were included so that the purpose of the screening questions is not revealed. The screened respondents were randomized into the following treatment groups: HYP, ICA1, ICA33, ICA66, and ICA100, each of them with a sample size of roughly N=255. In each group, we asked respondents to make 12 choices each consisting of four different specialty coffee sorts and one outside option. The coffee sorts were described along 6 attributes: roast, body, process, aroma, quality score, and price per 250gr package. All these attributes were explained to the respondents before they would go through the choice tasks. The attributes and their levels are shown in Table 8.

All respondents were given 11 euros. In the HYP group, respondents would receive the whole amount as their participation fee regardless of choice outcomes. Respondents in ICA1, ICA33, ICA66, and ICA100 were told that with a certain probability, one of their choices, drawn randomly, would be realized and they would have to buy it with this amount. That is, they would receive their chosen coffee plus the difference between the price of the coffee and the 11 euros. The realization probabilities were 1%, 33%, 66%, and 100% for ICA1, ICA33, ICA66, and ICA100 respectively, and were communicated to the respondents. Respondents were informed that if in the randomly drawn choice set, they had chosen the outside option, they would receive 11 euros without having to buy any product. To prevent respondents in the incentive-aligned groups from excessively choosing the outside option, which would result in low variation in the data, we asked for

[14] Specialty coffee is a term used to describe high grade coffee that is grown and processed usually in a single geographic area with a special micro-climate. The term is used as the opposite of commodity coffee, which is grown in large quantities and is usually a blend of beans from different regions.

prices that were between 40 to 50% lower than in the market and communicated this to respondents before they engaged with the choice tasks.

Attributes	**Levels**	**Number of Att. Levels**
Roast	Espresso, Filter	2
Body	Full, Medium, Light	3
Process	Washed, Natural	2
Aroma	Chocolaty, Nutty, Fruity	3
Quality score	86, 87, 88, 89	4
Price	EUR 7.4, EUR 8.3, EUR 9.2, EUR 10.1, EUR 11	5

Table 8: Product attributes and attribute levels.

5.2 Pre-treatment variables

Before starting the choice tasks, respondents were asked questions about their coffee drinking and purchasing habits. To ensure that random assignment of respondents into incentive groups was successful, we compared these pre-treatment variables across the groups. Table 9 shows the means and standard deviations of the pre-treatment variables across the groups. Respondents appear to be similar in terms of their coffee consumption habits: regardless of the incentive group, most of the respondents drink 1-2 cups daily, the favorite preparation method is filter coffee, and the highest percentage of respondents in each group consumes commercial coffee and spends less than 10 euros on coffee per week.

The only exception is coffee expenditure in ICA33, which has a higher percentage of respondents (7% on average) with lower coffee expenditure compared to other groups.

5.3 Self-reported and model-free measures

Self-reported attribute rankings *After* completing the choice tasks, respondents were asked to rank the coffee attributes, which appeared in the experiment in order of importance. Figure 6 shows the distribution of the attribute rankings for different incentive groups. Each line plot contains the relative frequency of the six different ranks respondents used to rank the attributes by, for a certain incentive group/attribute pair. The shades around

Variable	HYP	ICA1	ICA33	ICA66	ICA100
Coffee Consumption per Day					
% 1-2 cups	0.402	0.422	0.407	0.415	0.376
% 2-3 cups	0.324	0.31	0.343	0.285	0.384
% >3 cups	0.273	0.267	0.25	0.3	0.239
Coffee Preparation					
% French Press	0.027	0.031	0.036	0.05	0.035
% Filter coffee	0.477	0.434	0.435	0.385	0.459
% Espresso machine	0.066	0.085	0.089	0.088	0.075
% Fully automatic	0.309	0.353	0.298	0.346	0.329
% Other	0.121	0.097	0.141	0.131	0.102
Coffee Type					
% Specialty	0.242	0.217	0.21	0.238	0.216
% Commercial	0.758	0.783	0.79	0.762	0.784
Coffee Expenditure (in eur)					
% < 10	0.594	0.574	0.653	0.596	0.584
% 10-20	0.34	0.36	0.274	0.346	0.314
% 20-30	0.051	0.05	0.052	0.058	0.075
% > 30	0.016	0.016	0.02	0	0.027

Table 9: Pre-treatment characteristics by condition group

the line plots represent the standard errors around the relative frequencies. The overlap of the uncertainty shades is consistent with preference invariance across incentive groups. [15]

The most important attributes for all the groups are roast and aroma. Respondents within all the groups appear to be similarly divided with respect to the importance of price and process, and body and quality score are consistently the least important attributes.

Self-reported attribute attendance We also asked respondents to indicate whether they ignored any attributes when making their choices. Figure 7 shows the percentage of respondents who reported that they have not ignored any attributes throughout the experiment. There are more respondents who reported full attendance in the ICA groups compared to the HYP group.

Outside option value Despite the discount in coffee prices in the experiment, outside choices increase in the incentive level. Table 10 compares the respondents' average outside

[15]Later in this section we provide further model based evidence for the invariance of deep preference parameters underlying observed choices across incentive groups.

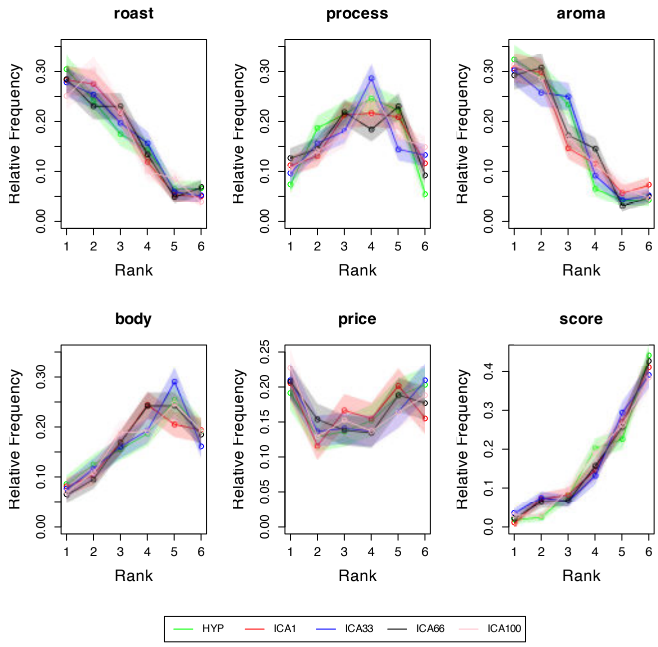

Figure 6: Attribute ranking distribution for respondents in different incentive groups (1= most important)

choice count across different incentive groups. We see that in the ICA groups, the average number of outside choices is higher in groups where choices are more likely to be realized. Outside choice of respondents in the HYP group, however, do not follow this pattern. This can be due to the fact that regardless of the respondents' choices in the HYP group, they would receive the same amount of money. However in the ICA groups, choosing outside increases the probability of receiving the full 11 euros, instead of buying the product. These results support having group-specific outside option preference parameters in our proposed framework.

	HYP	ICA1	ICA33	ICA66	ICA100
Outside choices	1.91	1.55	1.66	2.02	2.17
	(0.22)	(0.20)	(0.21)	(0.22)	(0.24)

Table 10: Average (standard errors) of outside choices across DMs in different incentive groups.

5.4 Model calibration

To have a benchmark against the DPRM, we estimated the Scale MNL (Fiebig et al., 2010) that is the state of the art model for generalizing choices across environments that induce different accuracy. Similar to our framework, this model assumes an invariant preference distribution $N(\bar{\beta}, \Sigma_\beta)$ in the population under different incentive groups, however tries

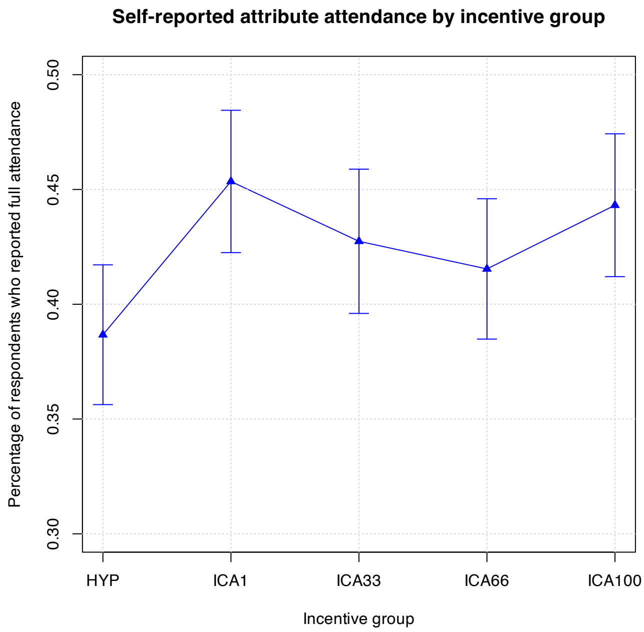

Figure 7: Percentage of respondents who reported that they have attended to all attributes across different incentive groups

to capture differences caused by incentive change across groups by allowing for different preference scales(c) across the groups. The scale factor of HYP respondents is fixed to 1, and the scales in other groups are estimated relative to that of the HYP respodents Therefore the likelihood of choice j in a choice set L for individual i in the ICA group is:

$$P(y_j|\beta_i, c) = \frac{exp(c \cdot x'_j\beta_i)}{\sum_{l\in L} exp(c \cdot x'_l\beta_i)} \tag{11}$$

Table 11 shows the population posterior moments of the scale parameter in different groups. The scale is decreasing in the incentive level that suggests choices are more consistent at lower incentive levels. Even though these results are consistent with the literature (Hauser et al., 2019), this seems at odds with the conjecture that cognitive effort increases in incentives.

	HYP	ICA1	ICA33	ICA66	ICA100
Scale(c)	1	0.92	0.86	0.9	0.89
	(0.00)	(0.04)	(0.05)	(0.05)	(0.04)

Table 11: Average (standard errors) of scale estimates across different incentive groups. Scale of HYP preferences is fixed to 1 for identification.

Estimating the DPRM can shed more light on the decision making process as it accounts for the effect of incentive through different channels: first, through the decision threshold, second, through attribute attendance, and third, varying outside option preferences.

Table 12 shows the population posterior moments of the parameters of the DPRM that vary with changing incentives. In line with the Scale MNL results, the threshold is slightly decreasing in the incentive level, which is an indication of *more* consistency in choices at *lower* incentive levels.

However, looking at the attribute attendance estimates, we see that the average number of attributes attended is increasing in the incentive level. Moreover, and consistent with the outside choice count across different incentive levels, the preference for the outside option is increasing in the incentive level.[16]

Next we use the model estimates to explore the impact of incentive-alignment on the amount of information processed in choice.

	HYP	ICA1	ICA33	ICA66	ICA100
Threshold(k)	66.98	64.54	57.33	58.74	60.79
	(1.42)	(2.56)	(2.04)	(2.05)	(1.88)
Attribute attendance	4.9	5.04	4.64	5.3	5.3
	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Outside option preference (in log scale)	-1.02	-0.86	-1.2	-0.86	-0.74
	(0.12)	(0.14)	(0.20)	(0.13)	(0.14)

Table 12: Averages (standard errors) of the DPRM estimates affected by incentive change. Threshold is the population threshold parameter, attribute attendance is the average number of attributes attended, and outside preference is the population preference parameter for the outside option in log scale.

Effective information processing Given the DPRM parameter estimates for each individual we can simulate M races for each choice set that the individual faced in the experiment (see Section 4.4) and average the total number of realized hits across the M races to estimate the amount of information processed by the individual in that choice set.

Response time is commonly used as a proxy measure for information processing in the literature (Guo, 2022). To corroborate our model-based measure of information processing we use it to explain the individual decision-time in the choice tasks. A linear regression (with robust standard errors) of the log-transformed decision time of each individual on

[16] ICA33 has an unusually small value in all three measures and does not follow the pattern seen in the other ICA groups. Respondents in this group have reported low coffee expenditure which might suggest their low preference for specialty coffee a priori, see Section 5.2.

the total (simulated) hits processed by that individual in the experiment yields a positive and highly significant coefficient (Table 13).

	Estimate	Std. Error	t value	Pr(>|t|)
(Intercept)	4.1626	0.1856	22.43	0.0000
hits	0.0009	0.0002	4.32	0.0000

Table 13: Regression of log-transformed DCE decision time on the total hits processed for every individuals in the experiment. Robust standard errors are reported.

Conditional on the same information set (i.e. attributes), high thresholds lead to higher amount of information processing as measured by a larger number of realized hits in expectation (see Table 2). However, this relationship might not hold in the presence of attribute non-attendance and varying outside option preferences.

In Section 3 we showed that processing more attributes in the choice set can lead to more hits processed, because the DM needs to compare alternatives along non-comparable alternatives and resolve more trade-offs. On the other hand, high preference for the outside option can lead to fewer number of hits processed, because the DM does not spend effort to compare and resolve the trade-offs between the inside options.

To demonstrate these two opposing forces in the context of our empirical application, we simulate choices based on our model estimates in two scenarios.

	Score	Price(eur)	Roast	Body	Process	Aroma
Alt.1	89	8.3	Filter	Full	Washed	Chocolaty
Alt.2	88	7.4	Filter	Full	Natural	Nutty
Alt.3	87	8.3	Filter	Medium	Natural	Nutty
Alt.4	87	7.4	Filter	Medium	Washed	Chocolaty
Alt.5			Non of the above			

	Score	Price(eur)	Roast	Body	Process	Aroma
Alt.1	86	11	Espresso	Medium	Natural	Nutty
Alt.2	87	10.1	Espresso	Medium	Washed	Fruity
Alt.3	86	11	Espresso	Light	Washed	Fruity
Alt.4	87	10.1	Espresso	Light	Natural	Nutty
Alt.5			Non of the above			

Table 14: Two different choice sets from the coffee DCE design. A choice set where inside options are attractive (top) and one where inside options are unattractive. Attractiveness of the categorical attribute levels are gauged based on the population preference estimates of the DPRM calibrated using the coffee DCE data.

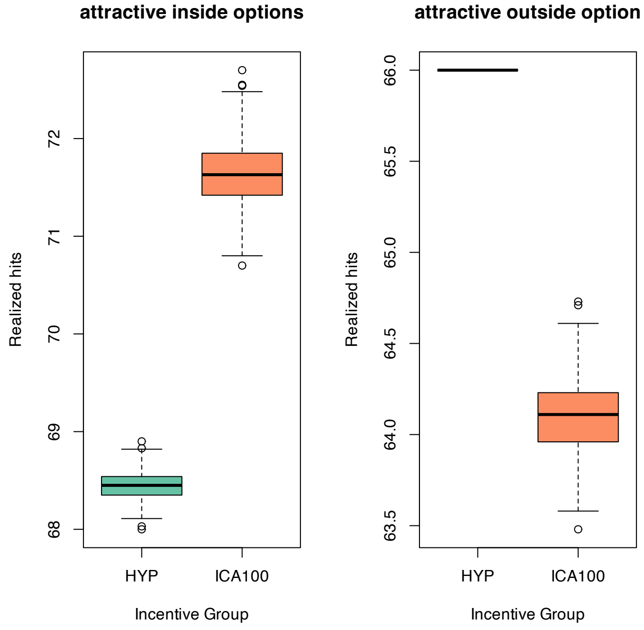

Figure 8: Distribution of realized hits from simulating 1000 choices by representative DMs in the HYP and ICA100 groups, in each of the choice sets in Table 14.

Table 14 shows two design matrices, the top one represents a choice set where the inside options are attractive and the bottom choice set represents one where the inside options possess unattractive attribute levels.[17]

Figure 8 compares the average realized hits from simulating 1000 choices of a representative DM from the HYP and ICA100 groups in each of the choice sets in Table 14. Representative DMs are constructed using the population posterior averages of the DPRM parameters. The average number of processed hits by the ICA100 DM is larger than that by the HYP DM in the first choice set, where the inside options are more attractive, but lower in the second choice set where they are unattractive. The reason is that, when the inside options are more attractive, the ICA100 DM faces more tradeoffs as she attends to more attributes, and this leads to higher effective information processing, higher realized hits. In the second choice set, both DMs mainly process hits from the outside option, and the lower threshold of the ICA100 as well as higher outside option rate leads to smaller number of realized hits, faster decisions.

[17]Alternatives in the choice set on the top have the lowest price levels, higher quality scores and the more attractive levels for aroma, process, body, and roast based on the population preference estimates. In contrast, alternatives in the choice set at the bottom are drawn from the less attractive levels in each attribute.

In fact we see the same pattern when comparing augmented hits conditional on the individual posterior draws and actual choices. ICA100 DMs processed more hits on average ($mean = 78.51, se = .18$) compared to HYP DMs ($mean = 74.86, se = .23$) when they chose an inside option, but processed fewer hits on average ($mean = 72.2, se = .56$) compared to HYP DMs ($mean = 77.51, se = .53$) when they chose outside.

To explore how the above changes in information processing affect the quality of choices, we computed the extent to which choices of respondents in each incentive group are aligned with their deep preferences.

We compare the rate of choosing the best alternatives in the choice set by the DMs in different incentive groups. The best alternative is defined as the alternative with the highest utility level in the choice set, as implied by measured (deep) preferences for attributes.[18] More formally, we evaluate to what extent the observed choice corresponds to the utility maximizing choice in the counterfactual world where respondents process all attributes. The counterfactual preference distribution is observed in the fully incentivized condition and applies to all incentive groups. As a consequence, we can impute preferences for an attribute that is ignored by a respondent at a lower incentive level from the joint posterior distribution of preferences given this respondent's observed (through the lens of the model) preferences for attended attributes. In fact, this imputation is automatically performed as part of calibrating the hierarchical prior distribution of preferences which does not depend on the incentive level. Thus, we obtain what would have been the utility-maximizing choice by setting all elements of $\tau_i = 1$ in the counterfactual calculation in Equation 12:

$$p(y_{obs} = y_{best}) = \int I(y_{obs} = y_{best}(\beta_i, \tau_i = 1))p(\beta_i|Y)d\beta_i \tag{12}$$

where $p(\beta_i|Y)$ is the marginal posterior distribution of respondent i's deep preferences.

We assume additively separable utilities and compute the utility levels based on the individual level preference estimates. Rates are obtained by taking averages across 1000 draws from the posterior of individual preferences to account for posterior uncertainty. Figure 9 shows that the rate of choosing the best alternatives in choice sets increases

[18]Recall that deep preferences are model parameters and realized preferences correspond to the number of Poisson hits associated with a particular attribute level leading up to a choice.

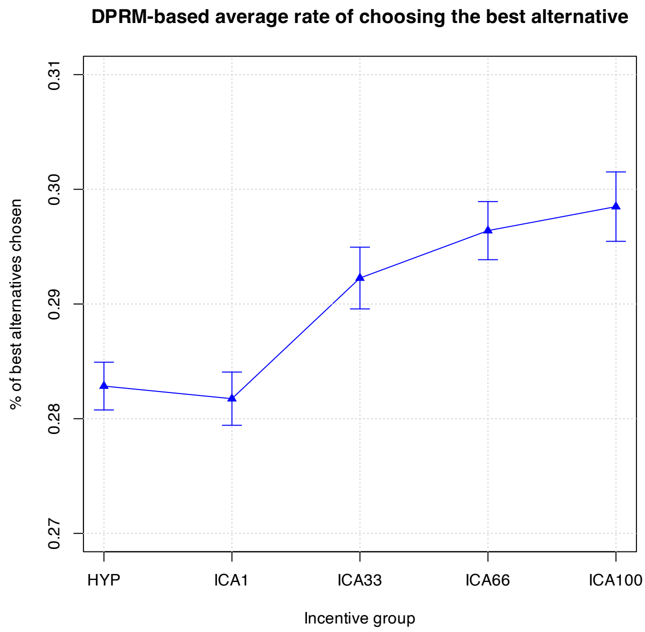

Figure 9: Average rate (standard error) of choosing best choices in a choice set by alternatives across different incentive groups. Best choices are alternatives with the highest utility level in their choice sets. Alternative utility is computed based on the individual level preferences accounting for posterior uncertainty.

in incentives and choices at higher incentive settings are more aligned with DMs deep preferences. This is in line with prior literature that shows incentives improves average performance (Camerer et al., 1999).

Model comparison To evaluate the relative performance of the DPRM, we compare its log marginal density (Newton and Raftery, 1994), predictive log likelihood, and out of sample hit rates with those of the Hierarchical Bayes (HB) logit (Allenby and Rossi, 1998), and MNL Scale (see Table 15). Both the DPRM and the MNL Scale pool across different incentive groups assuming fixed preferences, while the HB logit's predictions in each incentive group are made based on calibrating the model with the respective group's data separately.

To obtain HB logit predictive likelihood we compute $\prod_{i=1}^{N}\left[R^{-1}\sum_{r=1}^{R}p(y_i|x_i,\beta_{ir})\right]$ in the log scale. That is, we take the expectation over all the individual level posterior draws instead of using the posterior mean.

Hit rates are computed based on predicting one holdout choice per respondent. For the HB logit, we predicted individual i's holdout choice by finding the choice that maximizes: $R^{-1}\sum_{r=1}^{R}p(y_i|\beta_{ir})$ that is computed by taking the expectation of likelihood evaluations over R draws from the individual's preference posterior.

The predictive measures of MNL Scale in each incentive group are computed analogous to those of the HB logit, however preferences drawn from the posterior are scaled by the scale parameter (c) that is drawn from the posterior of the respective group.

The DPRM predictive likelihoods and hit rates are computed in a similar way. The difference is that we do not have the DPRM's likelihood $p(y_i|\beta_i, k_i)$ in closed-form expression and we approximate it by taking the expectation over simulated choices given posterior draws. Hence the computation of the predictive likelihood of holdout choices y_{hi} is done as in Equation 13.

$$
\begin{aligned}
p(y_{hi}|Y) &= \int p(y_{hi}|\beta_i, K_i)p(\beta_i, K_i|Y)d\{\beta_i, K_i\} \\
&= \int\int p(y_{hi}|z_{hi})p(z_{hi}|\beta_i, K_i)dz_{hi}p(\beta_i, K_i|Y)d\{\beta_i, K_i\} \\
&= R^{-1}M^{-1}\sum_{r=1}^{R}\sum_{m=1}^{M} p(y_{hi}|z_{hirm})
\end{aligned} \tag{13}
$$

where $z_{hirm} \sim \int p(z_{hi}|\beta_i, K_i)p(\beta_i, K_i|Y)d\{\beta_i, K_i\}$ and obtained by forward simulating races in parallel, and $p(\beta_i, K_i|Y)$ is the individual level posterior distribution.

The DPRM has a higher log marginal density than both the separate HB logit and MNL scale models. Moreover, both models that generalize across different incentive groups have better marginal densities than the joint marginal density of HB logit models calibrated to data from different incentive groups in isolation. This supports the argument that some aspects of choice remain invariant across different incentive levels. If incentives changed preferences in addition to information processing, both MNL scale and the DPRM would underperform compared to the (separate) HB logit.

	logMargDen	Predictive Likelihood					Hit Rate				
Model		HYP	ICA1	ICA33	ICA66	ICA100	HYP	ICA1	ICA33	ICA66	ICA100
MNL	-11115	-243	-254	-235	-250	-232	0.62	0.59	0.59	0.63	0.62
MNL Scale	-10429	-243	-256	-232	-249	-228	0.62	0.57	0.60	0.62	0.63
DPRM	-10390	-248	-253	-233	-253	-224	0.58	0.60	0.60	0.62	0.64

Table 15: Comparing the Log Marginal Density (Newton and Raftery, 1994) as well as predictive log likelihood and hit rates on the holdout data for the MNL, MNL Scale, and the DPRM. MNL is the HB logit with sign constraint trained on each incentive group dataset separately. MNL Scale is the HB logit with different scales of the error term for each incentive group, and DPRM represents the proposed bridging framework allowing for attribute nonattendance and different preferences for outside option.

6 Discussion

Several researchers indicated that hypothetical choices in discrete choice experiments, as common in applied market research, may generalize poorly to market choices because of their hypothetical nature (e.g. Ding et al., 2005; Ding, 2007; Ding et al., 2009). This conclusion is based on comparing choices across standard hypothetical settings and so called incentive-aligned settings. A common way to align incentives (between the respondent and the researcher) is to endow respondents with a budget under the contract that the respondent's choice in a randomly chosen choice set will become an actual transaction. Based on the different inference from standard choice models fitted to data from hypothetical and incentivized experiments, the current consensus is that incentives change respondents' preferences. As a consequence, only preferences inferred from incentivized experiments should be used for market simulation.

In this paper we propose instead that the major difference between hypothetical and incentive-aligned choices is not in the underlying preference structure, but in the cognitive effort respondents invest when making their choices. At a conceptual level, our proposal is related to Louviere's classic idea that seemingly different preferences across experimental settings or across stated and revealed preferences may be spurious consequences of different error distributions (Swait and Louviere, 1993; Louviere et al., 1999; Louviere and Eagle, 2006). However, we model differences in decision effort in the richer, and theoretically grounded framework of sequential sampling models of choice (e.g. Townsend and Ashby, 1983; Ratcliff and Smith, 2004). Supported by a body of experimental work, the proposed model identifies preferences independently from parameters that reflect a respondent's trade-off between accuracy and decision effort (Ratcliff and Rouder, 1998; White et al., 2010; Eidels et al., 2010; van Ravenzwaaij et al., 2012). The model structures the process that eventually leads to choice as accumulating information about the valuation of attributes in a choice context, and thus motivates more attribute-based and strictly alternative-based decision making from different effort levels. This mechanism can explain preference reversals from different effort levels under the same set of underlying, deep preference parameters. In this way, the model can bridge between data from hypothetical and incentivized experiments extracting information about invariant, deep preferences from both data sources while accounting for the difference in cognitive effort.

An alternative approach to modeling how incentives affect choice was recently proposed by Cao and Zhang (2020). These authors suggest that decision makers update the prior valuation of a product as a function of the expected net-gains from choosing based on effortfully constructed preferences. While some may view the structural formulation of the model in Cao and Zhang (2020) as an advantage, the model literally implies that hypothetical choices do not convey useful information. This implication is hard to reconcile with standard industry practice and raises the question what respondents actually do when participating in hypothetical choice experiments. The sequential sampling model we propose in this paper, instead, suggests that respondents act under the same deep preferences but exert more decision effort as choices are made to be consequential in an incentive-aligned experiment. Preferences for alternatives are constructed based on effortful valuations of attributes and levels of alternatives in a choice set, where low effort leads to more attribute-based decision making and high effort results in alternative based decision making that effectively integrates well deliberated valuations across all attributes describing alternatives. By formulating how valuations of alternatives in a choice develop, the proposed model applies straightforwardly to multi-attribute, multi-alternative decisions. In contrast, Cao and Zhang (2020) only address the binary choice between one inside good and the outside good. In addition, these authors assume that the processing of price information is automatic and perfectly accurate.

In contrast to the prevailing interpretation of the error term in random utility models as aspects of utility only observed by the decision maker, the proposed model explicitly motivates randomness in choices, after conditioning on attributes and levels defining a choice set, as the consequence of a trade-off between accuracy and decision costs by the decision maker. This trade-off governs the threshold parameter in our model under rational expectations about attributes and alternatives in a choice set and knowledge about the incentive structure in a particular experiment. Given this threshold and a set of deep, invariant preferences, the effective amount of processing before a choice occurs is endogenous to the specifics of alternatives in a choice task. For example, the decision maker needs to process more information (in expectation) to resolve a balanced trade-off as compared to when choosing among similar alternatives or among alternatives that involve trade-offs but one alternative is clearly better. This is consistent with observations of proximal measures of information processing made by other researchers (e.g. Yang et al., 2018).

As a technical contribution, we show how to leverage recent theoretical results on the convergence of MCMC-samplers when likelihoods needs to be simulated (Andrieu and Roberts, 2009), and develop a massively parallel algorithm to forward simulate from our model on a GPU (Nvidia, 2011). We are working on making the algorithm publicly available as an R-package.

Empirically, we find that the model captures the differences among data collected at different incentive levels choices under a common set of preferences equally well as different HB-logit models fit separately to data from each incentive level. This supports our model's preference invariance assumption. If preferences were to change across incentive levels, we would expect the model that pools across different incentives holding preferences fixed to perform worse than separate HB-logit models. Our model reconciles the results from the literature that shows incentive-alignment in DCEs lead to less consistent choices as measured by the scale of the error term in the MNL scale model (Hauser et al., 2019), with the vast literature that suggests incentives increase attention and cognitive effort in choice (e.g. Camerer et al., 1999; Wilcox, 1993; Kahneman and Peavler, 1969). We show that the incentivized respondents attend to more attributes on average and have a slightly lower decision threshold than hypothetical respondents. This leads to less consistent choices compared to hypothetical respondents. However, even at a lower decision threshold, as the incentivized respondents process a larger attribute set and have to resolve more trade-offs, they effectively process more information as measured by the model. We further corroborate the measure of information processing (hits) in the DPRM by showing that the estimated average number of hits processed for the respondets in our experiment significantly correlates with their decision time. Finally, we show that the quality of choices improves in incentives and choices become more consistent with the respondents' deep preferences. Our findings help market researchers when deciding about the design of their DCEs to get higher quality data.

A topic of future research is modeling definite evidence against alternatives in a choice set. The DPRM models lack of evidence in favor of attributes by setting the rate of that attribute to zero. However, capturing definite evidence *against* an alternative requires dependence between the *attribute*-specific processes, not just dependence between alternatives. Such a structure could capture non-compensatory choice behavior (Gilbride and Allenby, 2004). Further, similarity in categorical attribute levels is only measured on an "all-or-none" basis, where e.g. alternatives with identical brands share all the hits generated by that brand and those with different brands share none. Another possible extension of the model is capturing the similarity between different levels of a categorical attribute by defining a measure of similarity for that attribute.

Finally, another possible avenue for research is incorporating process measures such as eye-tracking and decision time as dependent variables in the model. Decision time in race models can be modelled by recording exponentially distributed interarrival time of hits (Otter et al., 2008). Visual fixation has been modelled e.g. in drift diffusion models as a paramater modulating the accumulation rate (Krajbich et al., 2010). Integrating these measures into the model could further corroborate the underlying process that the DPRM puts forward.

References

Agranov, M. and P. Ortoleva (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy 125*(1), 40–68.

Allenby, G. M. and P. E. Rossi (1998). Marketing models of consumer heterogeneity. *Journal of Econometrics 89*(1), 57–78.

Anderson, J. R., D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin (2004). An integrated theory of the mind. *Psychological review 111*(4), 1036.

Andrieu, C. and G. O. Roberts (2009). The pseudo-marginal approach for efficient monte carlo computations. *The Annals of Statistics*, 697–725.

Baldassi, C., S. Cerreia-Vioglio, F. Maccheroni, M. Marinacci, and M. Pirazzini (2020). A behavioral characterization of the drift diffusion model and its multialternative extension for choice under time pressure. *Management Science*.

Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics 164*(3), 1139–1160.

Bettman, J. R., E. J. Johnson, and J. W. Payne (1990). A componential analysis of cognitive effort in choice. *Organizational behavior and human decision processes 45*(1), 111–139.

Bouyssou, D. and M. Pirlot (2004). Preferences for multi-attributed alternatives: Traces, dominance, and numerical representations. *Journal of Mathematical Psychology 48*(3), 167–185.

Bradlow, E. T. and Y.-H. Park (2007). Bayesian estimation of bid sequences in internet auctions using a generalized record-breaking model. *Marketing Science 26*(2), 218–229.

Camerer, C. and D. Mobbs (2017). Differences in behavior and brain activity during hypothetical and real choices. *Trends in cognitive sciences 21*(1), 46–56.

Camerer, C. F., R. M. Hogarth, D. V. Budescu, and C. Eckel (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. In *Elicitation of Preferences*, pp. 7–48. Springer.

Cao, X. and J. Zhang (2020). Preference learning and demand forecast. *Marketing Science*.

Chen, N., J. A. Clithero, and M. Hsu (2019). Demand estimation and forecasting using neuroeconomic models of consumer choice. *Available at SSRN 3397895*.

Chiong, K., M. Shum, R. Webb, and R. Chen (2024). Combining choice and response time data: A drift-diffusion model of mobile advertisements. *Management Science 70*(2), 1238–1257.

Clithero, J. A. and A. Rangel (2013). Combining response times and choice data using a neuroeconomic model of the decision process improves out-of-sample predictions. *unpublished, California Institute of Technology*.

Debreu, G. (1960). Review of individual choice behavior by rd luce. *American Economic Review 50*(1), 186–188.

Dellaert, B. G., J. D. Brazell, and J. J. Louviere (1999). The effect of attribute variation on consumer choice consistency. *Marketing Letters 10*(2), 139–147.

Ding, M. (2007). An incentive-aligned mechanism for conjoint analysis. *Journal of Marketing Research 44*(2), 214–223.

Ding, M., R. Grewal, and J. Liechty (2005). Incentive-aligned conjoint analysis. *Journal of marketing research 42*(1), 67–82.

Ding, M., Y.-H. Park, and E. T. Bradlow (2009). Barter markets for conjoint analysis. *Management Science 55*(6), 1003–1017.

Dong, S., M. Ding, and J. Huber (2010). A simple mechanism to incentive-align conjoint experiments. *International Journal of Research in Marketing 27*(1), 25–32.

Dotson, J. P., J. R. Howell, J. D. Brazell, T. Otter, P. J. Lenk, S. Maceachern, and G. M. Allenby (2018). A probit model with structured covariance for similarity effects and source of volume calculations. *Journal of Marketing Research*, jmr–13.

Eidels, A., C. Donkin, S. D. Brown, and A. Heathcote (2010). Converging measures of workload capacity. *Psychonomic bulletin & review 17*(6), 763–771.

Evgeniou, T., C. Boussios, and G. Zacharia (2005). Generalized robust conjoint estimation. *Marketing Science 24*(3), 415–429.

Fehr, E. and A. Rangel (2011). Neuroeconomic foundations of economic choice–recent advances. *Journal of Economic Perspectives 25*(4), 3–30.

Fiebig, D. G., M. P. Keane, J. Louviere, and N. Wasi (2010). The generalized multinomial logit model: accounting for scale and coefficient heterogeneity. *Marketing Science 29*(3), 393–421.

French, S. (1986). *Decision theory: an introduction to the mathematics of rationality.* Halsted Press.

Frydman, C. and G. Nave (2017). Extrapolative beliefs in perceptual and economic decisions: Evidence of a common mechanism. *Management Science 63*(7), 2340–2352.

Fudenberg, D., P. Strack, and T. Strzalecki (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review 108*(12), 3651–84.

Gigerenzer, G. and D. G. Goldstein (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological review 103*(4), 650.

Gilbride, T. J. and G. M. Allenby (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science 23*(3), 391–406.

Gilbride, T. J., G. M. Allenby, and J. D. Brazell (2006). Models for heterogeneous variable selection. *Journal of Marketing Research 43*(3), 420–430.

Guo, L. (2022). Testing the role of contextual deliberation in the compromise effect. *Management Science 68*(6), 4326–4355.

Hauser, J. R., F. Eggers, and M. Selove (2019). The strategic implications of scale in choice-based conjoint analysis. *Marketing Science 38*(6), 1059–1081.

Hausman, J. A. and D. A. Wise (1978). A conditional probit model for qualitative choice: Discrete decisions recognizing interdependence and heterogeneous preferences. *Econometrica: Journal of the econometric society*, 403–426.

Holmes, W. R. (2015). A practical guide to the probability density approximation (pda) with improved implementation and error characterization. *Journal of Mathematical Psychology 68*, 13–24.

Holmes, W. R. and J. S. Trueblood (2018). Bayesian analysis of the piecewise diffusion decision model. *Behavior research methods 50*(2), 730–743.

Howell, J. R., P. Ebbes, and J. C. Liechty (2021). Gremlins in the data: Identifying the information content of research subjects. *Journal of Marketing Research 58*(1), 74–94.

Kahneman, D. (1973). *Attention and effort*, Volume 1063. Prentice-Hall Englewood Cliffs, NJ.

Kahneman, D. and W. S. Peavler (1969). Incentive effects and pupillary changes in association learning. *Journal of Experimental Psychology 79*(2p1), 312.

Krajbich, I., C. Armel, and A. Rangel (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience 13*(10), 1292–1298.

Krajbich, I., B. Oud, and E. Fehr (2014). Benefits of neuroeconomic modeling: new policy interventions and predictors of preference. *American Economic Review 104*(5), 501–06.

Lewin, K. (2013). *A dynamic theory of personality-selected papers*. Read Books Ltd.

Lin, Y.-S., A. Heathcote, and W. R. Holmes (2019). Parallel probability density approximation. *Behavior research methods 51*(6), 2777–2799.

Lohse, G. L. and E. J. Johnson (1996). A comparison of two process tracing methods for choice tasks. *Organizational Behavior and Human Decision Processes 68*(1), 28–43.

Louviere, J. J. and T. Eagle (2006). Confound it! that pesky little scale constant messes up our convenient assumptions. In *Sawtooth Software Conference*. Sawtooth Software Inc.

Louviere, J. J. and E. Lancsar (2009). Choice experiments in health: the good, the bad, the ugly and toward a brighter future. *Health Economics, Policy and Law 4*(4), 527–546.

Louviere, J. J., R. J. Meyer, D. S. Bunch, R. Carson, B. Dellaert, W. M. Hanemann, D. Hensher, and J. Irwin (1999). Combining sources of preference data for modeling complex decision processes. *Marketing Letters 10*(3), 205–217.

McFadden, D. (1986). The choice theory approach to market research. *Marketing science 5*(4), 275–297.

Miller, K. M., R. Hofstetter, H. Krohmer, and Z. J. Zhang (2011). How should consumers' willingness to pay be measured? an empirical comparison of state-of-the-art approaches. *Journal of Marketing Research 48*(1), 172–184.

Newton, M. A. and A. E. Raftery (1994). Approximate bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society Series B: Statistical Methodology 56*(1), 3–26.

Noguchi, T. and N. Stewart (2018). Multialternative decision by sampling: A model of decision making constrained by process data. *Psychological review 125*(4), 512.

Nvidia, C. (2011). Nvidia cuda c programming guide. *Nvidia Corporation 120*(18), 8.

Orme, B. (2019). Consistency cutoffs to identify" bad" respondents in cbc, acbc, and maxdiff. *Sawtooth Software Research Paper Series*.

Otter, T., G. M. Allenby, and T. Van Zandt (2008). An integrated model of discrete choice and response time. *Journal of Marketing Research 45*(5), 593–607.

Otter, T., J. Johnson, J. Rieskamp, G. M. Allenby, J. D. Brazell, A. Diederich, J. W. Hutchinson, S. MacEachern, S. Ruan, and J. Townsend (2008). Sequential sampling models of choice: Some recent advances. *Marketing letters 19*(3), 255–267.

Padmala, S. and L. Pessoa (2011). Reward reduces conflict by enhancing attentional control and biasing visual cortical processing. *Journal of cognitive neuroscience 23*(11), 3419–3432.

Ratcliff, R., M. G. Philiastides, and P. Sajda (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the eeg. *Proceedings of the National Academy of Sciences 106*(16), 6539–6544.

Ratcliff, R. and J. N. Rouder (1998). Modeling response times for two-choice decisions. *Psychological science 9*(5), 347–356.

Ratcliff, R. and P. L. Smith (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological review 111*(2), 333.

Rossi, P. E., G. M. Allenby, and R. McCulloch (2012). *Bayesian statistics and marketing.* John Wiley & Sons.

Roy, R., P. K. Chintagunta, and S. Haldar (1996). A framework for investigating habits,"the hand of the past," and heterogeneity in dynamic brand choice. *Marketing science 15*(3), 280–299.

Ruan, S., S. N. MacEachern, T. Otter, and A. M. Dean (2008). The dependent poisson race model and modeling dependence in conjoint choice experiments. *Psychometrika 73*(2), 261–288.

Russo, J. E. and B. A. Dosher (1983). Strategies for multiattribute binary choice. *Journal of Experimental Psychology: Learning, Memory, and Cognition 9*(4), 676.

Shenhav, A., M. M. Botvinick, and J. D. Cohen (2013). The expected value of control: an integrative theory of anterior cingulate cortex function. *Neuron 79*(2), 217–240.

Simon, H. A. (1955). A behavioral model of rational choice. *The quarterly journal of economics 69*(1), 99–118.

Smith, V. L. (1982). Microeconomic systems as an experimental science. *The American Economic Review 72*(5), 923–955.

Starns, J. J. and R. Ratcliff (2014). Validating the unequal-variance assumption in recognition memory using response time distributions instead of roc functions: A diffusion model analysis. *Journal of memory and language 70*, 36–52.

Stüttgen, P., P. Boatwright, and R. T. Monroe (2012). A satisficing choice model. *Marketing Science 31*(6), 878–899.

Swait, J. and J. Louviere (1993). The role of the scale parameter in the estimation and comparison of multinomial logit models. *Journal of marketing research 30*(3), 305–314.

Toubia, O., J. R. Hauser, and D. I. Simester (2004). Polyhedral methods for adaptive choice-based conjoint analysis. *Journal of Marketing Research 41*(1), 116–131.

Townsend, J. T. and F. G. Ashby (1983). *Stochastic modeling of elementary psychological processes*. CUP Archive.

Turner, B. M. and P. B. Sederberg (2014). A generalized, likelihood-free method for posterior estimation. *Psychonomic bulletin & review 21*(2), 227–250.

Tversky, A. (1969). Intransitivity of preferences. *Psychological review 76*(1), 31.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological review 79*(4), 281.

Tversky, A. and J. E. Russo (1969). Substitutability and similarity in binary choices. *Journal of Mathematical psychology 6*(1), 1–12.

van Ravenzwaaij, D., G. Dutilh, and E.-J. Wagenmakers (2012). A diffusion model decomposition of the effects of alcohol on perceptual decision making. *Psychopharmacology 219*(4), 1017–1025.

Wertenbroch, K. and B. Skiera (2002). Measuring consumers' willingness to pay at the point of purchase. *Journal of marketing research 39*(2), 228–241.

White, C. N., R. Ratcliff, M. W. Vasey, and G. McKoon (2010). Using diffusion models to understand clinical disorders. *Journal of Mathematical Psychology 54*(1), 39–52.

Wilcox, N. T. (1993). Lottery choice: Incentives, complexity and decision time. *The Economic Journal 103*(421), 1397–1417.

Woodford, M. (2014). Stochastic choice: An optimizing neuroeconomic model. *American Economic Review 104*(5), 495–500.

Yang, L., O. Toubia, and M. G. de Jong (2018). Attention, information processing, and choice in incentive-aligned choice experiments. *Journal of Marketing Research 55*(6), 783–800.

Zeithammer, R. and P. Lenk (2006). Bayesian estimation of multivariate-normal models when dimensions are absent. *Quantitative Marketing and Economics 4*(3), 241–265.

Zhang, S. and A. B. Markman (2001). Processing product unique features: Alignability and involvement in preference construction. *Journal of Consumer Psychology 11*(1), 13–27.

A Appendix

A.1 Bayesian estimation of the DPRM

$$P(\gamma, K, \bar{\gamma}, \Sigma_\gamma, \tau, \phi, \lambda, \zeta, \theta | Y) \propto$$
$$P(Y|\gamma, K, \bar{\gamma}, \Sigma_\gamma, \tau, \phi, \lambda, \zeta, \theta) P(\bar{\gamma}, \Sigma_\gamma) P(\tau|\zeta) P(\zeta) P(\phi|\theta) P(\theta) P(K|\lambda) P(\lambda) P(\epsilon)$$

The above equation shows the joint posterior of the model parameters up to the normalizing constant. γ is the log-transformed vector of attribute weights ($\gamma = log(\beta)$), and K is the threshold parameter. τ is the indicator vector of attribute selection and ϕ is the indicator vector for categorical attributes' baseline levels.
For an alternative with 3 attributes, the first of which is ordinal and the other two categorical with 2 and 3 categories, preferences of the DM are coded in a vector $\gamma = \{\gamma_0, \gamma_1, ..., \gamma_6\}$, where the elements denote the intercept, the ordinal attribute, two levels of the first categorical attribute and three levels of the second categorical attribute. The DM's attribute attendance is indicated by the vector $\tau = \{\tau_0, \tau_1, \tau_2, \tau_3\}$, where τ_p is a dummy indicating if attribute p has been attended to (τ_0 corresponds to the intercept and is =1). The baseline (least attractive) level of the categorical attributes is indicated by $\phi = \{\phi_1, \phi_2\} = \{\phi_{11}, \phi_{12}, \phi_{21}, \phi_{22}, \phi_{23}\}$, where $\phi_{pj} = 1$ if level j of the $p-$th categorical attribute is the baseline.

The MCMC algorithm is as follows:

1. Draw individual parameters with a series of MH steps:

i) For each attribute $p = 1, ..., P$ Set $\tau_p = 1$, with probability ζ or 0 otherwise.

ii) For each categorical attribute $p = 1, ..., P_{cat}$ set ϕ_p equal to a draw from a multinomial distribution with probability θ_p.

iii) Draw γ^*, a candidate value for γ , from an importance sampling Pseudo-Marginal MH chain (see Appendix A.2 for the details), such that:

$$\gamma^* \sim MVN(C + \bar{\gamma}^{(n-1)}, s\Sigma_\gamma^{(n-1)})$$

where each element of γ^* is drawn conditional on the other elements. $C = \{C_1, C_2, ..., C_P\}$, where C_p is a scalar and $= 1000(\tau_p - 1)$ if attribute p is ordinal and is a vector and

$= 1000(\tau_p^* - 1)(\phi_p)$ if attribute p is categorical. s is a small positive number (we took s $= 0.4$ in our examples). The acceptance probability is computed as

$$\alpha_{(\gamma,\{z_m\}_{m=1}^M)} = min\left(\frac{\overline{P}(y|\gamma^*, K^{(n-1)})}{\overline{P}(y|\gamma^{(n-1)}, K^{(n-1)})}, 1\right)$$

Note that the proposal distribution here is equal to the prior. Therefore the priors and proposals have canceled each other out in the acceptance probabilty.

iv) Draw K^*, a candidate value for K, from a random walk Pseudo-Marginal MH chain, such that:

$$K^* \sim K^{(n-1)} \pm 1$$

The acceptance probability is computed as

$$\alpha_{(K,\{z_m\}_{m=1}^M)} = min\left(\frac{\overline{P}(y|\gamma^{(n)}, K^*, \epsilon)P(K^*|\lambda^{(n-1)})}{\overline{P}(y|\gamma^{(n)}, K^{(n-1)}, \epsilon)P(K^{(n-1)}|\lambda^{(n-1)})}, 1\right)$$

2. Draw γ's hyperparameters with a Gibbs step (see Rossi et al., 2012):

   i) First draw $\Sigma_\gamma^{(n)}$, from:

$$\Sigma_\gamma^{(n)} \sim IW(\nu + N, V + S)$$
$$S = E'E,$$
$$E = Y - X\tilde{B} + (\tilde{B} - \bar{B})'A(\tilde{B} - \bar{B}),$$
$$\tilde{B} = (X'X + A)^{-1}(X'Y + A\bar{B}).$$

   ii) Draw $\overline{\gamma}$ from:

$$\text{vec}(\overline{\gamma}^{(n)})|\Sigma_\gamma^{(n)} \sim N(\text{vec}(\bar{B}), \Sigma_\gamma^{(n)} \otimes (X'X + A^{-1}))$$

   N is the total number of respondents. The subjective priors are set to $\nu = 20$, $V = 14I_P$, $\bar{B} = 0$, and $A = 0.1$.

3. Draw $\tau = \{\tau_P\}_{p=1}^P$'s hyperparameter from a conjugate Beta distribution :

   i) For HYP respondents draw $\zeta^{HYP} = \{\zeta_p^{HYP}\}_{p=1}^P$ such that:

$$\zeta_p^{HYP} \sim Beta(a + \sum_{i=1}^{hyp} \tau_{ip}, b + \sum_{i=1}^{hyp} \tau_{ip})$$

ii) For ICA respondents draw $\zeta^{ICA} = \{\zeta_p^{ICA}\}_{p=1}^{P}$ such that:

$$\zeta_p^{ICA} \sim Beta(a + \sum_{i=1}^{ica} \tau_{ip}, b + \sum_{i=1}^{ica} \tau_{ip})$$

τ_{ip} is the element p of the vector τ_i for respondent i. hyp and ica are the number of HYP and ICA respondents. a and b are the subjective prior parameters and in our case are both set = 2.

4. Draw $\phi = \{\phi_p\}_{p=1}^{P_{cat}}$'s hyperparameter from a conjugate Dirichlet distribution :

   i) Draw $\theta = \{\theta_p\}_{p=1}^{P_{cat}}$, such that:

$$\theta_p \sim Dirichlet(D + \sum_{i=1}^{N} \phi_{ip})$$

P_{cat} is the number of categorical attributes, ϕ_{ip} is the phi vector for categorical attribute p for respondent i, N is the total number of respondents, D is a vector of subjective prior parameter with all its elements set = 2 in our case.

5. Draw K's hyperparameter with a random walk MH step:

   Draw λ^*, a candidate value for λ, such that:

$$ln(\lambda^*) \sim N(\lambda^{(n-1)}, \sigma)$$

In our examples we set $\sigma = .1$. The acceptance probability is computed as

$$\alpha_\lambda = min\left(\frac{P(K-1|\lambda^*)P(\lambda^*|\overline{\lambda},\sigma_\lambda)}{P(K-1|\lambda^{(n-1)})P(\lambda^{(n-1)}|\overline{\lambda},\sigma_\lambda)}, 1 \right)$$

The subjective priors are $\overline{\lambda} = 0$, and $\sigma_\lambda = 1$.

A.2 The Pseudo-Marginal Metropolis-Hastings algorithm

In general, inference about the posterior $p(\delta|\Delta)$ of an individual level parameter δ conditional on the rest of the individual parameters Δ, based on the approximated likelihood function $\bar{p}(y|\delta, \Delta)$, can be done using "$Iter$" number of iterations of the Pseudo-Marginal MH algorithm as follows:

i) Set the initial value: $\delta^{(n-1)}$ for $n = 1$

ii) Sample $m = 1, ..., M$ independent draws of races (z_m) from $p(z|\delta_0, \Delta)$ and compute the approximated likelihood as follows:

$$\bar{p}_{(n-1)}(y|\delta, \Delta) = \int p(y|z)p(z|\delta, \Delta)dz = M^{-1} \sum_{m=1}^{M} p(y|z_m)$$

where $z_m \sim p(z|\delta, \Delta)$ and obtained by forward simulating races.

iii) Draw δ^* from the proposal distribution $q(.)$

iv) Sample M independent draws of paths (z) from $p(z|\delta^*, \Delta)$ and compute $\bar{p^*}(y|\delta^*, \Delta)$

v) Accept δ^* and $\bar{p^*}(y|\delta)$ with probability:

$$\alpha_{(\delta, \{z_m\}_{m=1}^{M})} = min\left(\frac{\bar{p^*}(y|\delta^*, \Delta,)p(\delta^*)q(\delta^{(n-1)}))}{\bar{p}_{(n-1)}(y|\delta^{(n-1)}, \Delta)p(\delta^{(n-1)})q(\delta^*))}, 1\right)$$

otherwise, $\delta^{(n)} = \delta^{(n-1)}$ and $\bar{p}_{(n)} = \bar{p}_{(n-1)}$

vi) If $n = Iter$ stop, otherwise set $n = n + 1$ and go to v

The fundemental difference between this algorithm and a standard MH is that the current likelihood value ($\overline{p}(.)$) *must* be preserved at each MCMC iteration. This is because the likelihood value is a function of the probabilistic z draws, which implicitly become a part of the state space.