Unifying Procedure-Dependent Preference Reversals Theory and Experiments

Guo, Liang

Published in:
Management Science

Published: 01/11/2024

Document Version:
Post-print, also known as Accepted Author Manuscript, Peer-reviewed or Author Final version

Published version (DOI):
10.1287/mnsc.2021.02640

Publication details:
Guo, L. (2024). Unifying Procedure-Dependent Preference Reversals: Theory and Experiments. *Management Science*, *70*(11), 8163–8186. https://doi.org/10.1287/mnsc.2021.02640

Guo, L. (2024) Unifying Procedure-Dependent Preference Reversals: Theory and Experiments. *Management Science, 70*(11), 8163-8186.

https://doi.org/10.1287/mnsc.2021.02640

Unifying Procedure-Dependent Preference Reversals: Theory and Experiments*

Liang Guo
(City University of Hong Kong)

ABSTRACT

Revealed preferences between alternatives can be systematically reversed across a variety of elicitation procedures (e.g., choice, valuation, matching, joint/separate evaluation). These puzzling findings have been usually invoked to challenge the procedure invariance principle. Yet procedure-dependent preferences can be endogenous. This paper presents a unifying theory of contextual deliberation to account for seemingly disparate phenomena of preference reversals. When attribute importance is ex ante imperfectly known, people can engage in costly information retrieval/acquisition activities (i.e., deliberation) prior to making decisions. Elicitation procedures can influence revealed preferences through affecting the incentive for deliberation. Therefore, contextual deliberation can endogenously yield procedure-dependent preference reversals, offer a common micro-foundation for extant psychological accounts (e.g., the prominence hypothesis, the evaluability hypothesis), and coherently organize apparently unrelated/inconsistent findings in the literature. We also run five experiments and document new findings that are inconsistent with extant hypotheses but can be reconciled by contextual deliberation.

Key words: deliberation, evaluation mode, evaluation scale, preference reversal, procedure invariance, prominence effect, joint evaluation

* This paper benefits from insightful comments of Wilfred Amaldoss, Sanjay Jain, and participants at 2017 Asian Meeting of the Econometric Society, 2017 Marketing Science Conference at Los Angeles, and 2018 China GBA Experimental Economics Workshop. The financial support offered by the Hong Kong RGC General Research Fund (Grant Number 16500614), the subject pool provided by the Department of Marketing at the Chinese University of Hong Kong, and excellent research assistance by Tianqi Xue and Tianjiao Wu, are gratefully acknowledged. Address for correspondence: Department of Marketing, City University of Hong Kong, Kowloon, Hong Kong, China. Email: liang.guo@cityu.edu.hk.

1 Introduction

Economic theories generally assume that preferences are exogenous and can be revealed through well-designed elicitation procedures. Procedures to elicit preferences differ in their *evaluation scale* and/or *evaluation mode* (Hsee et al. 1999). Choice over alternatives is the "golden" scale. Valuation represents the monetary equivalent of an option, which can take the form of willingness to pay (WTP) or willingness to accept (WTA). Another scale is matching, whereby respondents are asked to input the value for an option's missing attribute (e.g., price, quality) to generate equivalence to another option. In terms of evaluation mode, different options can be assessed either jointly or separately. In joint evaluation (JE), multiple options are presented and evaluated simultaneously. In separate evaluation (SE), alternatives are presented one at a time and evaluated in isolation.[1] These preference-eliciting procedures are ubiquitously and interchangeably used in academic research and in practice (e.g., Tversky et al. 1988, Carmon and Simonson 1998).

One fundamental premise of the revealed preference paradigm is the *principle of procedure independence*. That is, the same preference order should be inferred from logically equivalent preference-eliciting procedures.[2] However, the past decades have seen many persistently replicated studies (typically on two-attribute alternatives in controlled lab experiments) that demonstrate violations of the procedure independence principle. These violations suggest that preferences can be systematically varied, and even reversed in a specific direction, across elicitation procedures.

One well-known instance of scale-dependent preference reversal (PR) is between binary choice and matching. It refers to the finding that people tend to choose the option with higher value on the so-called prominent attribute (i.e., the prominent option) whereas the reverse preference order can be deduced from inputted attribute value that matches the options. For example, Tversky et al. (1988) find in a variety of categories that the proportion of respondents who prefer the prominent option is higher in choice than in matching: job candidates, public welfare programs, benefit plans, and coupons. They term this phenomenon as the *prominence effect*, which has been replicated in various settings (e.g., Nowlis and Simonson 1997, Carmon and Simonson 1998, Fischer et al. 1999, Chernev 2005). Another notable PR is between evaluation modes. The joint-separate PR is originally observed by Bazerman et al. (1992) on options of dispute resolution between neighbors. Subsequent studies document similar phenomenon in contexts such as WTP valuation (Hsee 1996), public goods (Kahneman and Ritov 1994), sportscard (List 2002), and gender bias (Bohnet et al.

[1]An evaluation scale can be in the JE or the SE mode. These two dimensions are not independent: some scales (e.g., choice, matching) are typically in the JE mode, and others (e.g., valuation) can be in either the JE or the SE mode. There are other less-used scales that are variants of choice, valuation, and matching (see Section 3.2).

[2]Procedure invariance across scales/methods of measurement is commonly presumed in other scientific fields such as physics (e.g., length and weight for objects). See Tversky et al. (1988) for more discussion.

2016). These findings provide evidence against the procedure independence principle and pose challenges to the revealed preference paradigm. They also yield profound controversies on practical applications of preference measures and techniques such as conjoint methods, demand estimation, market research, and public policy analysis (e.g., Hsee et al. 2008, Bohnet et al. 2016).

The prevalence and the importance of the procedure-dependent PRs raise a basic question that has not been adequately addressed. Why do logically equivalent procedures lead to systematic reversals? Previous research has identified domain-specific features that seem to account for a particular type of PR (e.g., the prominence hypothesis for the prominence effect, and the evaluability hypothesis for the JE-SE reversal). However, as we elaborate in this paper, the underlying mechanisms remain unclear. The extant explanations are mostly ad hoc, cannot reconcile seemingly unrelated findings within and across different types of procedure-dependent PRs, and may even lead to misleading interpretation of elicited preferences.

Our objective in this paper is to present a unifying perspective to endogenize the dependence of revealed preferences on elicitation procedures, which can rationalize extant hypotheses and organize seemingly unrelated findings. The premise of our theory is that people may not know their preferences perfectly, even for alternatives with well-defined attributes. For example, they may not be immediately sure about the relative importance of attributes. Nevertheless, they can engage in cognitive activities to retrieve or acquire relevant preference information (i.e., deliberation). To the extent that elicitation procedures (e.g., choice versus matching, joint versus separate evaluation) may influence the incentive for deliberation that in turn affects revealed preferences, procedure dependence may emerge endogenously. As a result, revealed preferences may appear to be procedure dependent, even though the underlying preference/information structure is not.

The theory of contextual deliberation can parsimoniously rationalize seemingly disparate phenomena of procedure-dependent PRs, without resorting to any nonstandard preference/belief. In particular, we present necessary and sufficient conditions for different types of choice-matching PR to emerge endogenously from contextual deliberation. An alternative can be favored more in choice than in matching, if and only if its likelihood of being preferred decreases (or increases) with deliberation but a shift to matching leads to greater (or lower) deliberation, and vice versa. Intuitively, an increasing deliberation makes the distribution of the posterior mean of an attribute's importance more dispersed, thus reducing the likelihood of preferring the alternative that enjoys advantage on that attribute under the prior belief. Importantly, this common mechanism can yield either the standard or the nonstandard choice-matching PR (i.e., the prominence effect or the reverse prominence effect), depending on which alternative is more advantageous under the prior belief and on whether optimal deliberation is lower in choice than in matching. Moreover, we show how varying attribute values may endogenously influence the relative ex ante advantageousness of the alternatives as well as optimal deliberation across the evaluation scales, thus leading to multiple non-monotonic shifts between the prominence and the reverse prominence effects.

Contextual deliberation can also give rise to the marked discrepancy between joint and separate valuations. The underlying mechanism consists of two integral parts: an increasing optimal deliberation from separate to joint valuation, and differential response of the options' mean valuations to the change in deliberation. The first part follows naturally from the different number of decisions, across the evaluation modes, that can be informed by deliberation. The second part is determined by the curvature of the utility function and the options' attributes. We identify conditions for the emergence of the standard or the nonstandard joint-separate PR, in terms of which alternative's mean valuation becomes relatively higher from the SE to the JE mode.

We conduct five experiments to document new findings that are inconsistent with extant psychological accounts but can be explained by contextual deliberation. To falsify the psychological explanations for the prominence effect (e.g., the prominence hypothesis), we investigate how revealed preferences and a proxy of cognitive activities (i.e., response time) may change when the non-prominent option's attribute value decreases. Because the relative prominence of attributes is fixed under the extant explanations, they would predict no qualitative difference in the choice-matching PR across different attribute values, while excluding the possibility of observing the reverse prominence effect. Another identifying angle is about the underlying cognitive process: the psychological hypotheses typically posit that the prominence effect is associated with lower cognition in choice than in matching. Moreover, to remove the evaluability hypothesis as an explanation for the JE-SE reversal, we deliberately select attribute values such that the option that is superior on the more-difficult-to-evaluate attribute is also relatively more advantageous. This option's evaluability would increase and hence it should be favored more from the SE to the JE mode, whereas the reverse would be predicted by the theory of contextual deliberation.

Our experiments produce results, for both incentive-based and hypothetical tasks across several categories, that are clearly inconsistent with the extant psychological hypotheses. As the non-prominent option's attribute value decreased, differences in revealed preferences across the choice-matching scales may become nonsignificant, which was not associated with a longer response time in the choice tasks. It is also observed that the prominent option was favored less in choice than in matching (i.e., the reverse prominence effect). We even observe instances where, across different attribute values, the impact of choice (vs. matching) on revealed preferences changed signs. Moreover, we find significant differences between joint and separate valuations, in the opposite (same) direction as predicted by the evaluability hypothesis (by contextual deliberation).

Starting from Aumann (1962), there is a long history in decision theory to study incomplete preferences over lotteries or acts, arising from ambiguous/uncertain beliefs (e.g., Gilboa et al. 2010, Cerreia-Vioglio et al. 2011) or indecisiveness between alternatives (e.g., Ok 2002). Ok et al. (2012) investigate how indecisiveness in beliefs and indecisiveness in the underlying tastes for the outcomes of uncertain acts can be distinguished and connected to each other. In addition, there are models that characterize selection strategies among multiple beliefs or utility functions (e.g., Gilboa and Schmeidler 1989, Cerreia-Vioglio et al. 2015). In relation to this literature, we consider a random-parameter model that leads to stochastic preference, where the realization of the preference parameter is endogenous through contextual deliberation.

Another related stream of studies are the developments on Bayesian cognition models (Gabaix 2019). The basic theme of this literature is how Bayesian posteriors under imperfect perceptions about the decision problem may yield risk aversion (Khaw et al. 2017) or nonstandard preferences such as intertemporal preference reversals (Gabaix and Laibson 2017) and inverse S-shaped probability weighting function (Enke and Graeber 2019). Conceptually, our focus is on imperfect preferences even when perceptions are perfect about parameters of the decision problem. Moreover, our model differs fundamentally in that preference uncertainty can be endogenously reduced as the decision maker invests in costly deliberation prior to decision making.

Our research is related to the literature on pre-decision information acquisition, which is normally termed as rational inattention (e.g., Sims 2003, Woodford 2009). Caplin and Dean (2015) and Matejka and McKay (2015) study revealed preferences on state-dependent stochastic choice data, and Ergin and Sarver (2010) investigate how costly contemplation may influence preferences over menus. Bartos et al. (2016) show that costly information acquisition can magnify prior beliefs and preferences to endogenously lead to discrimination against negatively stereotyped groups.

This paper adds to a growing research stream on contextual deliberation that seeks to rationalize various nonstandard preferences and beliefs, which have been typically interpreted as decision bias or errors. For instance, Guo (2016, 2022, 2023) shows that endogenous deliberation can yield nonstandard preferences such as the compromise effect, choice overload, and the anchoring of valuations. Contextual deliberation can also explain the choice-valuation reversal for lotteries (Guo 2021). Our contribution to this research stream is both theoretical and empirical. First, we present a more general model of contextual deliberation, and substantially expand the research scope to accommodate different classes of procedure-dependent PRs, including the choice-matching and the JE-SE reversals and their many variants. Importantly, we articulate how our unifying framework can rationalize various PRs, each of which has been typically viewed as an essentially different phenomenon, dominantly explained by a fundamentally divergent theory, and developed as a separate and large literature in fundamentally independent domain. As a result, our theoretical model and analysis can reconcile many seemingly unrelated findings within and across these procedure-dependent PRs. Second, we report a variety of well-replicated experimental findings across these procedure-dependent PRs. Our experiments can not only help falsify the extant psychological accounts, but also present newly-documented empirical phenomena (e.g., the reverse prominence effect, the nonstandard joint-separate PR) that are novel and important on their own.

2 Model Setup

We develop the theory of contextual deliberation to account for endogenously procedure-dependent preferences. Our model is a significant extension of Guo (2016) to a general structure of information acquisition and to endogenous procedure dependence. The building block of the theory is a two-stage sequential deliberation-then-response process that characterizes how a decision maker (DM) may optimally respond to preference-eliciting tasks under different procedures.

Let $\mathcal{S}$ be the set of feasible alternatives. An alternative $A \in \mathcal{S}$ is represented by a vector of attributes x_A (e.g., product specifications). To facilitate exposition, we concentrate on the two-alternative, two-attribute setting on which procedure-dependent PRs are typically documented (e.g., Tversky et al. 1988). The following utility function follows their specification:

$$u(\theta, x_i) = \theta v_1(x_{i_1}) + v_2(x_{i_2}), i \in \mathcal{S} = \{A, B\}, \tag{1}$$

where θ captures the relative weight for the first attribute, and $v_j(\cdot)$ is a monotone function that transforms attribute value to utility: $v_j'(\cdot) > 0$, $j \in \{1, 2\}$. Without loss of generality, the attribute values satisfy $x_{A_1} > x_{B_1}$ and $x_{B_2} > x_{A_2}$ such that no option is dominant on both attributes.

Tversky et al. (1988) assume that θ is deterministic and *directly* contingent on the elicitation method. In contrast, we assume that true preference is not immediately accessible or *ex ante* perfectly known to the DM, and that θ is uncertain and procedure independent.[3] The prior belief is that θ has distribution $F(\theta)$, density $f(\theta)$, mean $\hat{\theta}$, and support on some interval $[\underline{\theta}, \bar{\theta}]$.

The DM can engage in costly deliberation to partially improve her preference information. Deliberation can involve cognitive processes such as recall from memory, information processing, introspection, retrospection, anticipation, etc. For example, the DM can retrieve/acquire preference information by recollecting similar experiences in past decisions, thinking about the pros and cons of alternatives, contemplating over the importance of attributes, and foreseeing future usages.

[3]It is unnecessary to assume that the preference uncertainty is (only) over the attribute weight θ. For instance, we can consider preference imperfection regarding how the attribute values x_i may map to the DM's utility u in (1), which would generate conceptually similar results. Nevertheless, this alternative setting would require us to focus on specific utility functionals by parameterizing $v_j(\cdot)$. Therefore, our current setup can preserve generality in $v_j(\cdot)$, while casting direct contrast to the specification in the literature (e.g., Tversky et al. 1988).

In particular, the DM can acquire a signal $s \in [\underline{s}, \bar{s}]$ that potentially makes her better informed of θ over the prior. The DM sets the level of deliberation by choosing an $\alpha \in [\underline{\alpha}, \bar{\alpha}]$ that uniquely determines the signal generating process. Upon receiving a signal s, the DM updates her belief according to the Bayes' rule, yielding the posterior distribution $G(\theta|s, \alpha)$. Conditional on α and s, the posterior mean of θ is $\tilde{\theta}(s, \alpha) \equiv \int_{\underline{\theta}}^{\bar{\theta}} \theta dG(\theta|s, \alpha)$. The posterior mean $\tilde{\theta}(s, \alpha)$ is a random variable from the ex ante perspective (prior to receiving the signal s). To simplify notation, we use $\tilde{\theta}$ to denote $\tilde{\theta}(s, \alpha)$. Denote the distribution of $\tilde{\theta}$ as $H(\cdot, \alpha)$, and the density as $h(\cdot, \alpha)$.

Given that utility is linear in θ as in (1), all matters for deliberation is how it may influence the distribution of the posterior mean $\tilde{\theta}$. Our objective is to employ a flexible structure for $H(\cdot, \alpha)$ that is uniquely ranked by α. To this end, we adopt the rotation order (Johnson and Myatt 2006).

Assumption 1 (Rotation order): The family of distributions $H(\cdot, \alpha)$ are rotation ordered, i.e., there exists a rotation point $\theta^+ \in (\underline{\theta}, \bar{\theta})$ such that

$$\frac{\partial H(\tilde{\theta}, \alpha)}{\partial \alpha} > 0 \text{ if } \tilde{\theta} < \theta^+ \text{ and } \frac{\partial H(\tilde{\theta}, \alpha)}{\partial \alpha} < 0 \text{ if } \tilde{\theta} > \theta^+, \text{ for all } \alpha.$$

The rotation order implies mean-preserving spread: $\int_{\underline{\theta}}^{\tilde{\theta}} \frac{\partial H(x,\alpha)}{\partial \alpha} dx > 0$ for $\tilde{\theta} \in (\underline{\theta}, \bar{\theta})$, and $\int_{\underline{\theta}}^{\bar{\theta}} \frac{\partial H(x,\alpha)}{\partial \alpha} dx = 0$. That is, a higher α improves the DM's information about the true θ by making the distribution of the posterior mean, $H(\cdot, \alpha)$, more dispersed and converging to the prior $F(\cdot)$. The rotation order allows us to capture, in a fairly general and tractable manner, the influence of deliberation on the DM's preference information. As in the following examples, some commonly known models of information acquisition in the literature are special cases of the rotation order.

Example 1 (Gaussian learning): The DM's preference parameter θ is normally distributed with mean $\hat{\theta}$ and variance $1/\beta$: $\theta \sim N(\hat{\theta}, 1/\beta)$. A signal about θ can be acquired: $s = \theta + \varepsilon$, where ε is independent of θ and follows a normal distribution with mean 0 and variance $1/\alpha$: $\varepsilon \sim N(0, 1/\alpha)$. Conditional on s, the posterior distribution of θ is normal with mean $\tilde{\theta}(s, \alpha) = \frac{\beta\hat{\theta}+\alpha s}{\alpha+\beta}$ and variance $\frac{1}{\alpha+\beta}$. It can also be verified that the distribution of the posterior mean, $H(\tilde{\theta}, \alpha)$, is normal with mean $\hat{\theta}$ and variance $\frac{\alpha}{\beta(\alpha+\beta)}$, and is rotation ordered with the rotation point $\theta^+ = \hat{\theta}$.

Example 2 (Truth-or-noise): The DM's prior belief is that θ follows a distribution $F(\cdot)$. The DM can acquire a signal s about θ. With probability $\alpha \in [0, 1]$, $s = \theta$, and with probability $1 - \alpha$, s is independently drawn from $F(\cdot)$. Conditional on s, the posterior mean of θ is $\tilde{\theta}(s, \alpha) = \alpha s + (1-\alpha)\hat{\theta}$. The distribution of the posterior mean is then $H(\tilde{\theta}, \alpha) = F\left(\frac{\tilde{\theta}-(1-\alpha)\hat{\theta}}{\alpha}\right)$. It follows that $\frac{\partial H(\tilde{\theta},\alpha)}{\partial \alpha} = -\frac{\tilde{\theta}-\hat{\theta}}{\alpha} h(\tilde{\theta}, \alpha)$. Therefore, $H(\tilde{\theta}, \alpha)$ is rotation ordered with the rotation point $\theta^+ = \hat{\theta}$.

Deliberation is normally costly because of the engaged time and cognitive resources. It is assumed that the cost of deliberation, $k(\alpha)$, is continuous and nonnegative. In addition, $k'(\cdot) > 0$ and $k''(\cdot) > 0$ for all α, $k'(\underline{\alpha}) \to 0$, and $k'(\bar{\alpha}) \to +\infty$. This cost structure is assumed to guarantee the existence of an interior solution for optimal deliberation.

Conditional on $\tilde{\theta}$, the DM's posterior expected utility is:

$$u(\tilde{\theta}, x_i) = \tilde{\theta} v_1(x_{i_1}) + v_2(x_{i_2}), i \in \mathcal{S} = \{A, B\}. \tag{2}$$

To simplify notation, we write $u(\tilde{\theta}, x_i)$ as u_i, $i \in \mathcal{S}$. Denote the vector of posterior expected utilities for the alternatives in $\mathcal{S} = \{A, B\}$ as $u_{\mathcal{S}} = \{u_A, u_B\}$. It is also useful to write the ex ante expected utilities as $\hat{u}_i = \hat{\theta} v_1(x_{i_1}) + v_2(x_{i_2})$, $i \in \mathcal{S}$.

Depending on the elicitation procedure, the DM is asked to respond to decision tasks that are designed to reveal her preferences over $\mathcal{S}$ (e.g., choice, valuation, matching). Denote the DM's response for a representative task as d, and write the corresponding objective function as $U_E(d, u_{\mathcal{S}})$, which is to be specified subsequently according to the elicitation procedure E. The DM follows a two-stage sequential process for each task by behaving as an expected-utility maximizer: first determine how much to deliberate (i.e., α) and then, conditional on the gathered signal through deliberation (i.e., $\tilde{\theta}$), decide on the optimal response (i.e., d).

Conditional on $\tilde{\theta}$, the DM's problem for a decision task solves:

$$max_d \ U_E(d, u_{\mathcal{S}}). \tag{3}$$

Denote the optimal response as d_E^*. The DM's first-stage deliberation problem solves:

$$max_\alpha \left\{ U_E(\alpha) = \int_{\underline{\theta}}^{\bar{\theta}} U_E(d_E^*, u_{\mathcal{S}}) \ h(\tilde{\theta}, \alpha) d\tilde{\theta} - k(\alpha) \right\}. \tag{4}$$

If the optimal interim utility $U_E(d_E^*, u_{\mathcal{S}})$ is convex (concave) in $\tilde{\theta}$, an increasing α would yield a higher (lower) ex ante expected surplus $U_E(\alpha)$. Optimal deliberation is generally influenced not only by the attribute values in $\mathcal{S}$ but also by the preference elicitation procedure. That is, optimal deliberation can be written as α_E^*, capturing its dependence on the elicitation procedure. As a result, the elicited preference can appear to be influenced by the elicitation procedure, if the DM engages in endogenous deliberation. This apparent dependence of preference on procedure may arise, although the preference/information structure is presumed to be procedure invariant. As illustrated in Figure 1, thanks to the mediation of endogenous deliberation, elicitation procedures may become normatively relevant and exert systematic influences on elicited preferences.

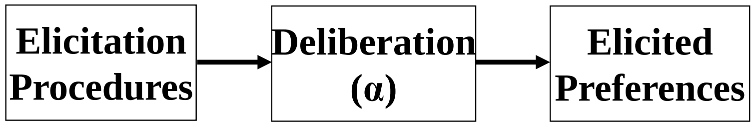

Figure 1: Deliberation Mediates the Dependence of Preferences on Elicitation Procedures

3 Endogenous Preferences across Evaluation Scales

We now investigate how elicited preferences can be reversed endogenously across evaluation scales. We focus on how contextual deliberation can yield different types of choice-matching PR (the case of choice-valuation PR for riskless alternatives is similar). We also discuss how the theory of contextual deliberation can parsimoniously reconcile extant accounts, organize other related findings in the literature, and generate new predictions.

3.1 The Choice-Matching PR

Two commonly employed scales to elicit preferences are choice and matching. In a binary choice task (i.e., $E = C$), the DM is asked to select the preferred option from the set $\mathcal{S} = \{A, B\}$. That is, the DM's problem is to determine the choice function $d_C : \mathcal{S} \to \mathcal{S}$. Preference between the options is revealed through observed choice, i.e., $A \succ B$ if and only if $d_C(\mathcal{S}) = A$.[4]

Alternatively, in a matching task (i.e., $E = M$), the DM is presented the same options except that one of the four attribute values is missing. Without loss of generality, suppose the missing value is x_{B_2}. The DM is requested to indicate a value d_M such that the generated alternative $\Delta = (x_{B_1}, d_M)$ is equivalently preferred as $A = (x_{A_1}, x_{A_2})$. Preference between the two original alternatives can be inferred by comparing the stated matching value d_M with the missing value x_{B_2}. Given that preference is increasing in each attribute, $\Delta \succ B$ if and only if $d_M > x_{B_2}$. It follows that, by $A \sim \Delta$ and transitivity, $A \succ B$ if and only if $d_M > x_{B_2}$.

Procedure invariance across choice and matching is typically tested through between-subjects experiments. Let the proportion of experiment participants who indicate a higher preference for option A under the evaluation scale $E \in \{C, M\}$ be P_E. The principle of procedure invariance implies that $P_C = P_M$. However, when option A's advantageous attribute (i.e., the first attribute, given that $x_{A1} > x_{B1}$) can be interpreted as more prominent than the other attribute, it is usually observed that $P_C > P_M$ (e.g., Tversky et al. 1988).[5] That is, the relative preference for the

[4]To simplify matters without the loss of generality, we assume away preference indifference between the options.

[5]The determination on attribute prominence is ad hoc in the literature, i.e., through experiment instructions or

prominent alternative is higher in choice than in matching, i.e., the prominence effect. We also call the prominence effect the standard choice-matching PR, and refer to the alternative outcome (i.e., $P_C < P_M$) as the reverse prominence effect or the nonstandard choice-matching PR.

We show below that both the standard and the nonstandard choice-matching PRs can arise endogenously from contextual deliberation. We start by demonstrating that the endogeneity of deliberation is necessary to generate procedure-dependent preference. Put it differently, we show that the revealed preference would be independent of the evaluation scale if deliberation were exogenous or fixed (i.e., $\alpha_C^* = \alpha_M^* = \alpha$).

Consider first the choice task. It is evident that $d_C^*(\mathcal{S}) = A$ if and only if $\tilde{\theta} > \dot{\theta} \equiv \frac{v_2(x_{B_2}) - v_2(x_{A_2})}{v_1(x_{A_1}) - v_1(x_{B_1})}$. As a result, the probability that option A is chosen is $P_C = 1 - H(\dot{\theta}, \alpha)$. Consider then the matching task under a variant of the classical BDM mechanism (Becker et al. 1964) to ensure incentive compatibility: the DM receives option A if and only if $y < d_M$, and obtains the generated option $\Lambda = (x_{B_1}, y)$ if otherwise, where y is randomly drawn from a distribution $W(y)$ on $[\underline{y}, \bar{y}]$ with density $w(y)$.[6] Let $u_\Lambda(y) = \tilde{\theta} v_1(x_{B_1}) + v_2(y)$ be the posterior expected utility for the generated option. Following (3), the optimal matching value is determined by solving:

$$max_d \left\{ U_M(d, u_{\mathcal{S}}) = \int_{\underline{y}}^{d} u_A \ w(y)dy + \int_{d}^{\bar{y}} u_\Lambda(y) \ w(y)dy \right\}. \tag{5}$$

The first-order condition is $u_A - u_\Lambda(d) = 0$, which gives rise to the optimal interior solution $d_M^* = v_2^{-1}\left(\tilde{\theta}[v_1(x_{A_1}) - v_1(x_{B_1})] + v_2(x_{A_2})\right)$.[7] It follows that $d_M^* > x_{B_2}$ if and only if $\tilde{\theta} > \dot{\theta}$. This implies that the probability of option A being preferred over B, as inferred from the optimal matching value d_M^*, is the same as that revealed by the choice task, i.e., $P_M = 1 - H(\dot{\theta}, \alpha) = P_C$. As a result, stochastic preference by itself does not lead to systematically discrepant responses between the choice and the matching tasks.

However, when deliberation is endogenous (i.e., $\alpha_C^* \neq \alpha_M^*$), procedure-dependent reversal may emerge between choice- and matching-based preferences.

---

researcher judgement (e.g., quality more prominent than price, traffic causalities more important than implementation costs for accident-reducing programs, health more important than costs of eliminating beach pollution).

[6] This mechanism can simulate the decision process for hypothetical matching tasks. For instance, consider Problem 1 in Tversky et al. (1988) where the participants are asked to specify the missing score of "Human Relations" for candidate B such that the two candidates would be equally suitable for the position of a Production Engineer. This task can be transformed (by the respondents) into a more familiar problem, which is to search sequentially and randomly in the market for a better candidate to replace current candidate A. The quality on Human Relations of potential candidates, denoted as y, is random with distribution $W(\cdot)$. The optimal solution to this search problem is to set a reservation value for the attribute such that search continues until a candidate with a score no lower than the reservation value is found. The setting of the reservation value is then equivalent to deciding on the matching value (i.e., d_M) under the modified BDM mechanism.

[7] The second-order condition is satisfied at the optimal d_M^*: $(u_A - u_\Lambda(d))w'(d) - u'_\Lambda(d)w(d) = -u'_\Lambda(d)w(d) < 0$.

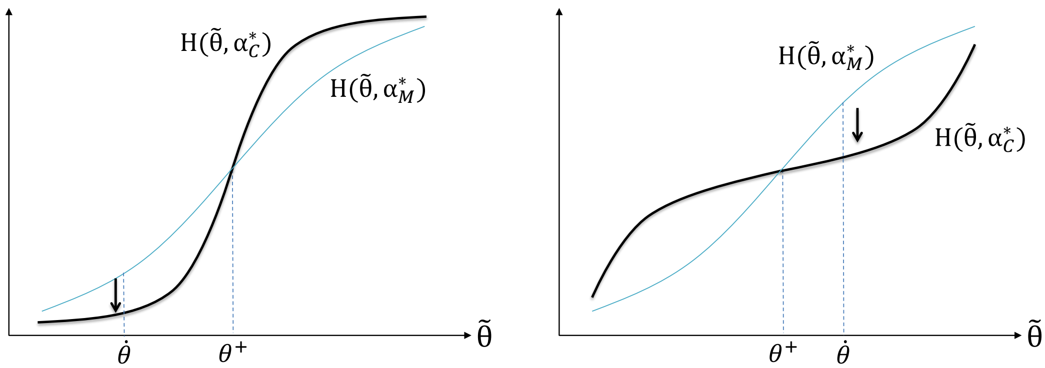

(a) $\dot{\theta} < \theta^+$ and $\alpha_C^* < \alpha_M^*$ (b) $\dot{\theta} > \theta^+$ and $\alpha_C^* > \alpha_M^*$

Figure 2: The Standard Choice-Matching PR (The Prominence Effect): $P_C > P_M$

Theorem 1 *$P_C > P_M$ if and only if one of the following conditions is satisfied: (i) $\dot{\theta} < \theta^+$ and $\alpha_C^* < \alpha_M^*$; or (ii) $\dot{\theta} > \theta^+$ and $\alpha_C^* > \alpha_M^*$.*

This theorem presents the necessary and sufficient conditions for the relative preference of the prominent option A to be greater in choice than that in matching (i.e., the prominence effect), which are illustrated in Figure 2. The conditions hinge on the comparison of the preference cut-off point (i.e., $\dot{\theta}$) with the rotation point (i.e., θ^+), and on the comparison of the optimal deliberation levels across the evaluation scales. When the preference cut-off point is lower than the rotation point, a higher deliberation decreases the probability that option A is preferred, i.e., $\frac{\partial H(\dot{\theta},\alpha)}{\partial \alpha} > 0$ for $\dot{\theta} < \theta^+$ according to Assumption 1. As a result, if optimal deliberation is lower in choice than in matching, as in Theorem 1(i) and Figure 2(a), the probability that option A is favored would loom larger in the choice task. Alternatively, if the preference cut-off value is high relative to the rotation point, option A would be more likely to be preferred as deliberation increases, i.e., $\frac{\partial H(\dot{\theta},\alpha)}{\partial \alpha} < 0$ if $\dot{\theta} > \theta^+$. For $P_C > P_M$ to happen, it would then be necessary and sufficient that optimal deliberation under choice is higher than that under matching, which is represented by Theorem 1(ii) and Figure 2(b).

Importantly, as implied by Theorem 1, we can also generate the necessary and sufficient conditions for $P_C < P_M$ (i.e., the reverse prominence effect). In particular, if the preference cut-off point is lower than the rotation point and optimal deliberation is higher in choice than in matching, as in Figure 3(a), the probability that option A is preferred in choice would be lower than that implied by the matching task. In addition, in case the preference cut-off point is higher than the rotation point, we would have $P_C < P_M$ if optimal deliberation is lower under choice than that under matching, which is illustrated in Figure 3(b).

Next, we examine how parameters of the decision tasks (i.e., attribute values) may influence the

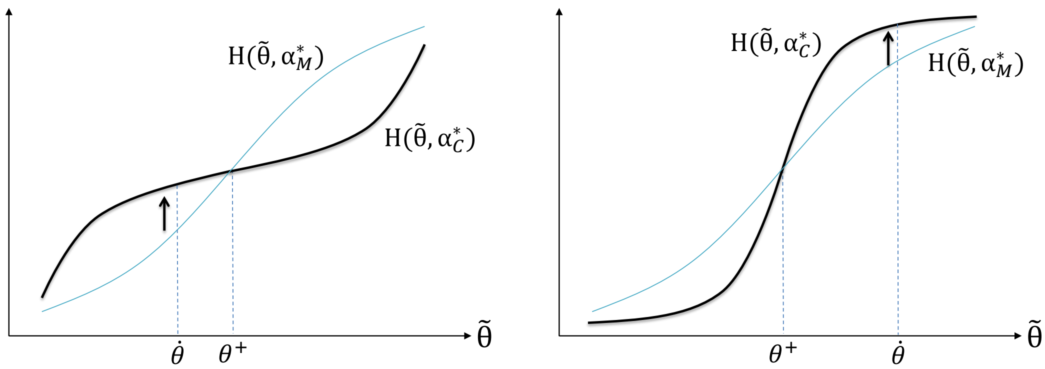

(a) $\dot{\theta} < \theta^{+}$ and $\alpha_C^* > \alpha_M^*$

(b) $\dot{\theta} > \theta^{+}$ and $\alpha_C^* < \alpha_M^*$

Figure 3: The Nonstandard Choice-Matching PR (The Reverse Prominence Effect): $P_C < P_M$

revealed preferences across the elicitation scales. To this end, we need to investigate how optimal deliberation levels may vary with the elicitation scales. Given the optimal responses in the choice and the matching tasks we have examined, their interim utilities can be written as, respectively,

$$U_C(d_C^*, u_{\mathcal{S}}) = \max\{u_A, u_B\},$$

$$U_M(d_M^*, u_{\mathcal{S}}) = \int_{\underline{y}}^{d_M^*} u_A \; w(y)dy + \int_{d_M^*}^{\bar{y}} u_\Lambda(y) \; w(y)dy.$$

Optimal deliberation can then be derived by substituting $U_E(d_E^*, u_{\mathcal{S}})$ into (4), for $E \in \{C, M\}$.

Theorem 2 *Optimal deliberation under choice (i.e., α_C^*) first increases and then decreases with x_{B_2}. In addition, there exist $z^1 < z^2$ such that optimal deliberation under choice is higher than that under matching (i.e., $\alpha_C^* > \alpha_M^*$) if and only if $x_{B_2} \in (z^1, z^2)$.*

This theorem characterizes the impact of the attribute value x_{B_2} on optimal deliberation under choice relative to that under matching. It demonstrates that an increase in x_{B_2} can influence α_C^* in an inverted-U manner.[8] When the value of x_{B_2} is sufficiently low, option A would be the no-brainer choice. This is because the potential utility loss of not choosing option B, in case the true preference parameter θ turns out to be relatively small, is negligible. Analogously, when x_{B_2} becomes sufficiently high, the DM would not lose much by choosing option B. As a result, in either case the

[8] Similarly, α_C^* first increases and then decreases with x_{A_2}. The impact of x_{A_1} or x_{B_1} on α_C^* is not completely unambiguous for the whole range of attribute values. Nevertheless, it can be shown that α_C^* increases with x_{A_1} for relatively low x_{A_1}, decreases with x_{B_1} for relatively high x_{B_1}, and converges to the minimum level $\underline{\alpha}$ for extreme values of either attribute.

information value of deliberation is endogenously low and the DM does not have much incentive to improve her information about the relative weight of the attributes. In comparison, for intermediate x_{B_2}, it is not ex ante clear which option should be more preferred and thus the DM would gain more from being better informed of her preference. This contrasts with the corresponding matching task, in which the DM's optimal level of deliberation is strictly above the lower bound $\underline{\alpha}$ and invariant of the (missing) attribute value x_{B_2}. Moreover, as demonstrated in Theorem 2, optimal deliberation is higher in choice than that in matching, if and only if x_{B_2} is in an intermediate range.

We are then ready to determine how x_{B_2} may influence the emergence of the standard and the nonstandard choice-matching PRs in Figures 2 and 3. When x_{B_2} is sufficiently small, the preference cut-off point $\dot{\theta}$ would be lower than the rotation point θ^+ and optimal deliberation under choice would be lower than that under matching. It implies that the prominence effect would be observed as in Figure 2(a). Nevertheless, when x_{B_2} increases and falls into the intermediate rage (z^1, z^2) such that $\alpha_C^* > \alpha_M^*$ (see Theorem 2), we may obtain either the scenario in Figure 3(a) or that in Figure 2(b), depending on whether the preference cut-off point $\dot{\theta}$ is still lower or becomes higher than θ^+, respectively. That is, a moderately small x_{B_2} would generate the reverse prominence effect, whereas a moderately large x_{B_2} would lead to the re-emergence of the prominence effect. Finally, when x_{B_2} becomes sufficiently high such that $\dot{\theta} > \theta^+$ but $\alpha_C^* < \alpha_M^*$, we would expect to observe again the reverse prominence effect as in Figure 3(b). To summarize, we show that the impact of varying the attribute value x_{B_2} on the occurrence of the standard and the nonstandard choice-matching PRs may involve multiple non-monotonic shifts.

3.2 Theoretical and Empirical Implications

We present theoretical and empirical implications of accounting for scale-dependent reversals by contextual deliberation. First, contextual deliberation can readily reconcile previously-proposed explanations for the observed choice-matching PR (e.g., the prominence effect). Second, contextual deliberation can coherently organize an array of seemingly unrelated findings in the literature.

3.2.1 Relation to Extant Explanations of the Standard Choice-Matching PR

The prominence hypothesis proposed by Tversky et al. (1988) is by far the most widely recognized account for the standard choice-matching PR (i.e., the prominence effect).[9] It is construed as an application of a more general *compatibility principle*, according to which the weight of any input of the decision task (e.g., attribute characteristics) is enhanced by its compatibility with

[9]We (ought to) distinguish a psychological/theoretical account from the phenomenon it intends to explain.

the decision output (e.g., preference responses). The relative prominence between attributes is a qualitative (or ordinal) task feature, although their specific levels are quantitative arguments for the decision task. It then follows from the principle of compatibility that the qualitative consideration of attribute prominence looms larger in the ordinal procedure of choice than in the cardinal procedure of matching, whereas the specific attribute values are considered more in matching than in choice, i.e., the prominence hypothesis. Another major account for the prominence effect is the task-goal hypothesis (Fischer et al. 1999). It posits that the prominent attribute is weighted more heavily in tasks where the primary goal is to differentiate between alternatives (e.g., choice) than in tasks where the primary goal is to equate alternatives (e.g., matching).

The prominence and the task-goal hypotheses are not fundamentally different from each other. Both propose that normatively irrelevant task features may influence the weighting of different attributes. As acknowledged by Fischer et al. (1999), these two hypotheses can be viewed as alternative interpretations of the distinction between qualitative and quantitative tasks, focusing respectively on the scale or the purpose of the response. Moreover, both accounts share the same underlying premise that cognitively simple decision rules can be applied to resolve conflicts for indecisive choices but not for matching tasks. To support the heuristic of selecting options by attribute prominence, Tversky et al. (1988, p. 372) argue that "This procedure...does not require the decision maker to assess the trade-off between the attributes, thereby reducing mental effort and cognitive strain." In addition, their rationale for the compatibility principle is that "noncompatibility (in content, scale, or display) between the input and the output requires additional mental transformations, which increase effort and error, and reduce confidence and impact" (p. 376). Fischer et al. (1999, p. 1060) present similar reasoning for the task-goal hypothesis: "giving greater weight to the prominent attribute makes it easier to arrive at a clear choice."

The theory of contextual deliberation provides a micro-foundation for the choice-matching PR by endogenizing the dependence of revealed preferences on the elicitation scales. It also explicates why, consistent with the psychological accounts, the prominence effect may be associated with a less thoughtful decision process under choice than under matching (see Theorem 1(i) and Figure 2(a)). In this regard, contextual deliberation can be viewed as a deeper-level explanation for the psychological hypotheses, about when and why their described situation may or may not arise.

However, there are fundamental differences across the different perspectives, in terms of both the underlying premises as well as testable implications. According to either the prominence or the task-goal hypothesis, the relative prominence of attributes is determined in an ad hoc way, the prominent attribute necessarily weighs more in choice than in matching, and it is always desirable to simplify the decision process and to avoid extensive attribute tradeoff (which is feasible for ordinal comparisons between alternatives but not for cardinal assessments). In contrast, under the theory of contextual deliberation, we do not need to make judgement about the attribute prominence, the DM's posterior beliefs about the relative weight of attributes are mean-preserving and need not always favor a particular attribute across the evaluation scales, and optimal deliberation can be either lower or higher for choice tasks. Essentially, proponents of the prominence or the task-goal hypothesis interpret cognitive efforts only as costly activities that should be overcome, while ignoring their instrumental value of informing the DM's decision. It is the consideration of this tradeoff, under the contextual deliberation theory, that underscores the characterization of the endogenous emergence of both the standard and the nonstandard choice-matching PRs.

Consequently, we can empirically differentiate the alternative accounts by testing their predictions. The prominence effect should be observed irrespective of the value of x_{B_2}, if it is due to the prominence hypothesis or the task-goal hypothesis. Nevertheless, the reverse prominence effect is inconsistent with either of the two hypotheses. Moreover, the psychological accounts predict that the preference-relevant cognition level for choice should be lower when the respondents exhibit the prominence effect than that when they do not. In contrast, the theory of contextual deliberation predicts that changing the attribute value x_{B_2} may influence optimal deliberation and revealed preferences, to generate either the standard or the nonstandard choice-matching PR, in systematic but non-monotonic manners (see Figures 2 and 3).

3.2.2 Empirical Evidence from Previous Findings

Previous research has documented a number of findings in controlled experiments, besides the prominence effect itself, that may appear elusive but can be organized by the theory of contextual deliberation. Fischer et al. (1999) compare the inferred preferences, over several pairs of options A and B, among choice, matching, and several variants of choice. In a *difference comparison* task, the subjects are asked whether they would prefer to improve from an alternative $C = (x_{B_1}, x_{A_2})$ to $A = (x_{A_1}, x_{A_2})$ or to $B = (x_{B_1}, x_{B_2})$, where $x_{A_1} > x_{B_1}$ and $x_{A_2} < x_{B_2}$. In a *high-low* task, the subjects are asked whether x_{B_2} would be too high or too low to make A and B equally attractive, where the corresponding attribute value for option B is initially not displayed. In addition, the *choice-based matching* variant involves a sequence of adaptive choices between $A = (x_{A_1}, x_{A_2})$ and $B^t = (x_{B_1}, x^t_{B_2})$: the value of $x^t_{B_2}$ in each sub-task t is adjusted up (or down) if A (or B^{t-1}) is selected in the previous sub-task $t-1$, and the process continues until $|x^t_{B_2} - x^{t-1}_{B_2}|$ is smaller than some given threshold. Moreover, two variants of the high-low task are considered by adding a standard matching question, either immediately (*high-low-match-now*) or distantly (*high-low-match-later*), after the corresponding high-low task for the same pair of alternatives A and B.

The findings in Fischer et al. (1999) can be remarkably explained by contextual deliberation (i.e., the scenario in Figure 2(a)): the inferred preferences for the prominent option as elicited by a procedure are close to those by the choice or by the matching task, depending on their relative *informational similarity*. In particular, both the difference comparison and the high-low tasks are logically and informationally equivalent to choices, and hence preferences for the prominent option under these three elicitation procedures resemble each other and are all significantly and substantially higher than those implied by matching responses. Conversely, the percentage of prominent preferences evoked by choice-based matching is only slightly higher than that by the standard matching task, but considerably lower than those by the choice and the difference comparison tasks. This pattern is also consistent with contextual deliberation: even though each sub-task under choice-based matching is essentially a choice and may involve limited deliberation, preference information can be accumulated from previous sub-tasks to the final converging sub-task within each sequence. In other words, choice-based matching is logically similar to standard choice, but informationally close to matching. Following this reasoning, if information spiller-over is shut down within a sequence of choice-based matching responses, the elicited prominent preferences would be reverted to those under choices. This expectation is indeed confirmed: when different sequences of choice-based matching responses are interlaced such that the sub-tasks within a sequence are temporally separated, the prominent preferences resemble those in choice, but are significantly and sizably higher than those under matching and than those when choice-based matching sub-tasks are presented consecutively within a sequence.

Similar reasoning can be applied to the comparisons with the two high-low variants. The subjects would be induced to deliberate more under the high-low-match-now task than under the high-low or the choice task, because the acquired preference information can guide the immediately following matching response. Nevertheless, the information value would be mitigated under the high-low-match-later procedure, because the matching question is presented with temporal delay. The observed prominent preferences among the elicitation procedures are indeed ranked as expected according to the following significantly descending order: choice and high-low, high-low-match-later, high-low-match-now, and matching (see Tversky et al. 1988 on similar results).

Supportive evidence for contextual deliberation can be seen in other variants of the choice-matching experiment. As mentioned in Carmon and Simonson (1998, p. 342), respondents are significantly more likely to choose the prominent option in stand-alone choices than in choices that are preceded by matching questions on the same pairs of options. This finding is akin to the above-mentioned comparison between choice and choice-based matching (Fischer et al. 1999), thanks to the carry-over effect (on the focal choice) of previously acquired information from the preceding matching question or from the within-sequence choice sub-tasks, respectively. In addition, Attema and Brouwer (2013) demonstrate that preferences are significantly more consistent within choice-based matching than within standard matching tasks, and that no systematic PR is found between these two types of tasks. This echoes the results in Fischer et al. (1999), and is compatible with the interpretation that, even though less deliberation under each choice sub-task generates higher internal response consistency, within-sequence information spill-over under choice-based matching can lead to inferred preferences that are close to those evoked by matching.

3.3 The Choice-Valuation PR

The theory of contextual deliberation can also account for procedure-dependent preferences between other evaluation scales. Guo (2021) develops a model of contextual deliberation for risk preferences, which can explain not only the asymmetric PR between choices and certainty equivalents but also its mirror image between choices and probability equivalents. Therefore, contextual deliberation can accommodate this class of procedure-dependent reversals that have been traditionally explained by the *scale compatibility principle* (e.g., Tversky et al. 1988, Slovic et al. 1990, Tversky et al. 1990). Moreover, Guo (2021) reports experimental findings, on the choice-valuation PR between lotteries, that are consistent with the premises and predictions of contextual deliberation but are not predicted by either the scale compatibility or the prominence hypothesis.

The choice-valuation PR for riskless alternatives is recently documented by O'Donnell and Evers (2019). They show that people tend to choose hedonic goods (e.g., ice cream) over utilitarian ones (e.g., trash bags), but are more likely to indicate a higher WTP for the latter. They posit that choices are made under gut feelings and affect-based heuristics, whereas valuations require the consideration and integration of various information, which is consistent with the intuitive versus reasoned (and fast versus slow) two-system dichotomy of decision-making processes (Kahneman 2003).[10] This premise is spiritually similar to that invoked by the prominence and the task-goal hypotheses to explain the prominence effect. Nevertheless, the theory of contextual deliberation can explain why and when people may endogenously choose to deliberate more in valuation tasks than in choices and why and how their preferences may be reversed across the evaluation scales, as it does for the prominence effect (the scenario in Figure 2(a)).[11] Moreover, the deliberation explantation can account for the moderating results in O'Donnell and Evers (2019): instructing the subjects to write down the benefits of products before indicating their preferences can mitigate the choice-valuation PR by undermining their tendency to choose the hedonic product (without modifying the relative valuations). Note that choices can become better informed by this experimental procedure,

[10] To explain the choice-valuation PR in their experiments, O'Donnell and Evers (2019) need to further cite previous findings to assume that people prefer hedonic over utilitarian goods when affect-based heuristics are employed.

[11] Similarly, Miller et al. (2011) find in a large-scale field study that price sensitivity inferred from valuation-based measures is higher than from dichotomous choices.

as by the preceding matching questions in Carmon and Simonson (1998), or as in the choice-based matching tasks in Fischer et al. (1999). Across these seemingly unrelated settings, choices are affected in a systematic manner that is consistent with contextual deliberation.

4 Endogenous Preferences across Evaluation Modes

Elicited preferences can be reversed across the evaluation modes as well. We apply the theory of contextual deliberation to articulate the rationale for the endogenous emergence of the JE-SE reversal. We also discuss how contextual deliberation can parsimoniously reconcile extant hypotheses for the mode-dependent and the scale-dependent PRs. Moreover, we present findings in the literature that are consistent with the contextual-deliberation interpretation.

4.1 The PR between Joint and Separate Evaluations

Procedure invariance across the evaluation modes is typically tested through experiments in which preferences are elicited by some cardinal scale. For example, Hsee (1996) elicits WTP valuations by randomly assigning participants to three conditions: one for the JE mode, and the other two for the SE mode. Let d_{Ei} denote the representative DM's stated valuation (WTP or WTA), and $E[d_{Ei}]$ denote the mean valuation, for option $i \in \mathcal{S} = \{A, B\}$ under the evaluation mode $E \in \{J, S\}$.

The principle of procedure invariance requires that if $E[d_{SA}] > E[d_{SB}]$, then $E[d_{JA}] > E[d_{JB}]$, and vice versa. However, the preference ranking of the options, as measured by comparing the mean valuations, can be reversed across the evaluation modes (e.g., Hsee 1996). In particular, the joint-separate PR normally occurs in the direction of favoring the option in the JE mode whose superior attribute is harder to evaluate than the other attribute.[12] It is without loss of generality in our setting to let attribute 2 be more difficult to evaluate than attribute 1. The standard joint-separate PR then means $E[d_{SA}] > E[d_{SB}]$ but $E[d_{JA}] < E[d_{JB}]$. In addition, we refer to the alternative scenario, $E[d_{SA}] < E[d_{SB}]$ but $E[d_{JA}] > E[d_{JB}]$, as the nonstandard joint-separate PR.

Consider the BDM instrument in a WTP valuation: the DM pays a monetary amount equal to y to receive option $i \in \mathcal{S}$ if and only if $y \leq d_{Ei}$, and obtains nothing if otherwise, where y is randomly drawn from a distribution $W(y)$.[13] For simplicity, we concentrate on the case where the distribution $W(\cdot)$ is uniform. The DM's objective function for a WTP valuation task under the

[12] The evaluability of an attribute can be enhanced if, for instance, the participants are informed of the meaning of the attribute, the range of the attribute values is known, and/or the attribute values are dichotomous or categorial.

[13] The BDM device resembles real market mechanism where y represents the random market price. It is similar to the auction design adopted by List (2002) in a field experiment to document the standard joint-separate PR.

evaluation mode $E \in \{J, S\}$ is given by

$$max_d \left\{ U_E(d, u_i) = \int_{\underline{y}}^{d} [u_i - v(y)]/(\bar{y} - \underline{y}) dy \right\}, i \in \mathcal{S} = \{A, B\}, \tag{6}$$

where $v(\cdot)$ denotes the DM's indirect utility function of wealth that assumes the common features: $v'(\cdot) > 0$ and $v''(\cdot) \neq 0$. That is, the utility is strictly increasing and nonlinear in wealth.

The first-order condition for the WTP valuation is:

$$u_i - v(d) = 0, i \in \mathcal{S},$$

yielding the optimal interior solutions $d^*_{Ei} = v^{-1}(u_i)$ for either evaluation mode.[14] The first-stage problem for optimal deliberation can then be analyzed by inputting d^*_{Ei} and substituting the optimal interim utilities into (4), for the JE or the SE mode, respectively:

$$max_\alpha \left\{ \int_{\underline{\theta}}^{\bar{\theta}} U_J(d^*_{JA}, u_A)\ h(\tilde{\theta}, \alpha) d\tilde{\theta} + \int_{\underline{\theta}}^{\bar{\theta}} U_J(d^*_{JB}, u_B)\ h(\tilde{\theta}, \alpha) d\tilde{\theta} - k(\alpha) \right\}, \tag{7}$$

$$max_{\alpha_i} \left\{ \int_{\underline{\theta}}^{\bar{\theta}} U_S(d^*_{Si}, u_i)\ h(\tilde{\theta}, \alpha_i) d\tilde{\theta} - k(\alpha_i) \right\}, i \in \mathcal{S}. \tag{8}$$

Theorem 3 $\alpha^*_J > \max\{\alpha^*_{SA}, \alpha^*_{SB}\}$.

This theorem suggests that optimal deliberation is higher under the JE mode than that under either of the SE tasks. This general result simply reflects the fact that both options are evaluated simultaneously under the JE mode but separately under the SE mode. Intuitively, the total marginal value of information is higher under the JE mode, even though the marginal information value for each individual valuation task remains the same across the evaluation modes.

The optimal valuation d^*_{Ei} can be approximated by taking the second-order Taylor expansion at the ex ante expected utility $\hat{u}_i$:

$$d^*_{Ei} = v^{-1}(u_i) \approx v^{-1}(\hat{u}_i) + (v^{-1})'(\hat{u}_i)(u_i - \hat{u}_i) + \frac{1}{2}(v^{-1})''(\hat{u}_i)(u_i - \hat{u}_i)^2. \tag{9}$$

It follows that the mean valuation $E[d^*_{Ei}]$ can be approximated as:

$$E[d^*_{Ei}] \approx v^{-1}(\hat{u}_i) + \frac{1}{2}(v^{-1})''(\hat{u}_i) v_1(x_{i_1})^2 \sigma^2_{Ei}, \tag{10}$$

[14]The second-order condition is satisfied: $-v'(d)/(\bar{y} - \underline{y}) < 0$.

where σ_{Ei}^2 is the variance of $\tilde{\theta}$ when option $i \in \mathcal{S}$ is evaluated under the evaluation mode $E \in \{J, S\}$. Theorem 3 implies that $\sigma_{JA}^2 = \sigma_{JB}^2 > \max\{\sigma_{SA}^2, \sigma_{SB}^2\}$. To facilitate exposition, denote $\hat{\nu}_i = v^{-1}(\hat{u}_i)$ and $\hat{\eta}_i = \frac{1}{2}(v^{-1})''(\hat{u}_i)v_1(x_{i_1})^2 \neq 0$, $i \in \mathcal{S}$. Let us then examine how the comparison between the mean valuations may vary with the evaluation modes.

Theorem 4 *$E[d_{SA}^*] > E[d_{SB}^*]$ and $E[d_{JA}^*] < E[d_{JB}^*]$ if the following conditions are satisfied: $\hat{\nu}_A - \hat{\nu}_B > \hat{\eta}_B\sigma_{SB}^2 - \hat{\eta}_A\sigma_{SA}^2$, $\hat{\eta}_A - \hat{\eta}_B < 0$, and $\sigma_{Ji}^2 > \dot{\sigma} = \frac{\hat{\nu}_A - \hat{\nu}_B}{\hat{\eta}_B - \hat{\eta}_A}$. Alternatively, $E[d_{SA}^*] < E[d_{SB}^*]$ and $E[d_{JA}^*] > E[d_{JB}^*]$ if the following conditions are satisfied: $\hat{\nu}_A - \hat{\nu}_B < \hat{\eta}_B\sigma_{SB}^2 - \hat{\eta}_A\sigma_{SA}^2$, $\hat{\eta}_A - \hat{\eta}_B > 0$, and $\sigma_{Ji}^2 > \dot{\sigma}$.*

This theorem presents conditions under which the relative ranking of the mean valuations is reversed across the evaluation modes. The conditions for the emergence of the standard joint-separate PR are illustrated in Figure 4(a), where the ex ante expected utility of option A relative to that of option B (i.e., $\hat{\nu}_A - \hat{\nu}_B$) is not too low but an increase in the deliberation leads to a relatively higher mean valuation for option B than for option A (i.e., $\hat{\eta}_A - \hat{\eta}_B < 0$).[15] It means that neither option can dominate the other and that the order of their mean valuations is moderated by the level of deliberation. Furthermore, for the order of the mean valuations to be reversed across the evaluation modes (i.e., $E[d_{JA}^*] < E[d_{JB}^*]$), optimal deliberation under the JE mode (i.e., σ_{Ji}^2) needs to be sufficiently high.

The conditions for the nonstandard joint-separate PR are analogous. As illustrated in Figure 4(b), in this case the ex ante expected utility of option A is not too high relative to that of option B whereas the disadvantage would be mitigated by a higher deliberation level (i.e., $\hat{\eta}_A - \hat{\eta}_B > 0$). Analogously, the mean valuation of option A under the JE mode would become higher than that of option B (i.e., $E[d_{JA}^*] > E[d_{JB}^*]$), if the JE-mode optimal deliberation (i.e., σ_{Ji}^2) is sufficiently high.

4.2 Theoretical and Empirical Implications

This section is devoted to discussing the insights of interpreting the JE-SE reversal from the lens of contextual deliberation. We examine how extant explanations for the (standard) joint-separate PR can be rationalized, and reconciled with those for the scale-dependent PRs, by the theory of

[15]Note that, for the mean valuation of option A under the SE mode to be higher than that of option B (i.e., $E[d_{SA}^*] > E[d_{SB}^*]$), unless $\hat{\eta}_A < 0 < \hat{\eta}_B$, it is not always necessary to have $\hat{\nu}_A - \hat{\nu}_B > 0$. In addition, $\hat{\eta}_i$ can be either positive or negative, which implies that the utility function $v(\cdot)$ can be (locally or globally) concave or convex. Moreover, as can be seen from Figure 4(a), optimal deliberation under the SE mode can be higher for option A or for option B (i.e., σ_{SA}^2 can be higher or lower than σ_{SB}^2). Analogously, for the nonstandard joint-separate PR, we do not need $\hat{\nu}_A - \hat{\nu}_B < 0$, $\hat{\eta}_i$ can be either positive or negative, and σ_{SA}^2 can be higher or lower than σ_{SB}^2.

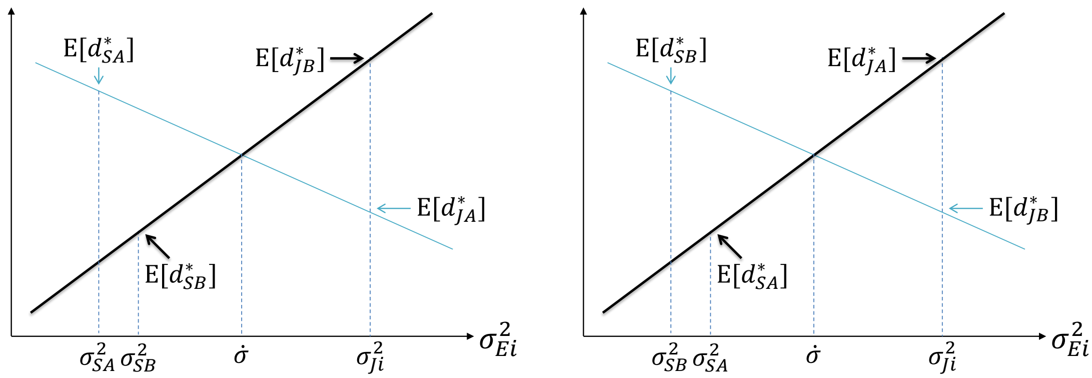

(a) The Standard Joint-Separate PR (b) The Nonstandard Joint-Separate PR

Figure 4: Endogenous Deliberation Yields the JE-SE Preference Reversal

contextual deliberation. We also present findings on the JE-SE reversal that are consistent with the deliberation interpretation.

The *evaluability hypothesis* (Hsee 1996, Hsee et al. 1999) is the most cited to account for the typical JE-SE reversal for same-category alternatives with well-defined attributes.[16] It posits that the relative importance of an attribute is positively related to its evaluability, and that an attribute may become easier to evaluate from the SE to the JE mode.[17] Therefore, if the alternatives involve one easy-to-evaluate and one difficult-to-evaluate attribute, the alternative superior on the difficult-to-evaluate attribute would be favored more under joint evaluation than under separate evaluation. Another interpretation for the joint-separate evaluation reversal is the *want/should hypothesis* (Bazerman et al. 1998). It proposes that the attributes can be classified along the want/should dimension, and that the importance of the should attribute relative to that of the want attribute looms larger in the JE mode than in the SE mode.

To apply these two hypotheses to account for the JE-SE reversal, one must be willing and able to determine *a priori* the attribute features: low versus high evaluability, or what people want versus what they should have. The central question is then about why attribute features/importance change systematically across the evaluation modes. This is presumed in the want/should hypothesis. The answer offered by the evaluability hypothesis is that the presence of both options under joint

[16]Preferences can also be reversed between alternatives from different categories that are not directly comparable. For example, people can be more willing to pay in the SE mode to deal with threats to Australian mammal species, but indicate a higher preference in the JE mode to solve the problem of skin cancer for farm workers (Kahneman and Ritov 1994). The *norm theory* proposed by Kahneman and Miller (1986) can be applied to explain this type of JE-SE reversal: an alternative is evaluated independently relative to its category norm under the SE mode, whereas the overall importance of a category would be more salient in across-category assessments under the JE mode.

[17]A similar account called *attribute ambiguity* is proposed by Loewenstein et al. (1994).

evaluation provides more information about the difficult-to-evaluate attribute (Hsee et al. 1999, Bazerman et al. 1999, Bohnet et al. 2016). However, more questions remain to be addressed, e.g., what is the uncertainty/information about, and why does more information lead to higher attribute importance? Nevertheless, the two extant accounts share similar view about how the evaluation modes are associated with different cognitive processes. They both propose that decision making under separate evaluation is more intuitive and automatic and that under joint evaluation is more reasoned and reflective (Bazerman et al. 1999, Zhang 2015, Bohnet et al. 2016).[18]

The theory of contextual deliberation provides a micro-foundation to account for the (standard) joint-separate PR, with subtle but fundamental differences from extant explanations. It posits that, what the DM is uncertain about is not attribute values *per se*, but the mapping from attribute values to utility (e.g., attribute importance). The mere presence of more attribute values from jointly evaluated alternatives need not provide more information about preferences. Rather, it is the additional decision task under joint valuations that increases the incentive to acquire information (Theorem 3). Thus, contextual deliberation can explicate why joint valuations should be cognitively more thoughtful than separate valuations, as supposed by the evaluability or the want/should hypothesis. In addition, the endogenously improved information increases the dispersion in the posterior mean of the attribute importance, without making an attribute more important. That is, the influence of information on preferences is neutral. It is because the optimally stated valuation is nonlinear in the posterior attribute importance that enhanced deliberation may reverse the alternatives' relative valuations across the evaluation modes (Theorem 4).

Although reversals across evaluation modes are generally distinguished as disparate phenomena from those across evaluation scales (e.g., choice, matching, valuation), as discussed above, their extant explanations are based on the same framework of two cognitive systems (e.g., Kahneman 2003). However, these explanations do not seem to be coherently driven by a general principle about when and why a cognitive system should be employed. Recall that, according to the prominence and the task-goal hypotheses (and the perspective offered by O'Donnell and Evers on the choice-valuation PR), the intuitive System 1 is the default and should be naturally evoked to save cognitive efforts (e.g., choice), and it would be necessary to resort to the effortful System 2 only when heuristic decision rules fail to work (e.g., matching, valuation). In comparison, the evaluability hypothesis appears to suggest that effortful attribute tradeoffs should be made in the decision process (e.g., joint evaluation), unless it is infeasible to do so when comparative information is unavailable (e.g., separate evaluation). After all, if System 1 were the default, choices would not be reversed across the joint/separate modes, which would contradict the findings in Gonzalez-Vallejo and Moran (2001).

[18]One might be somewhat surprised by the conflicting stance of the evaluability hypothesis: joint evaluation makes an attribute easer to evaluate although the decision process becomes cognitively more engaged.

They show that the choice between two alternatives can differ systematically from the revealed preference when each task involves an accept/reject decision for a separately presented alternative. Similar pattern of reversals can be seen in Bohnet et al. (2016) between jointly (A and B and an outside option) and separately presented choices (A/B and an outside option).

To further illustrate the inherent inconsistency across the extant hypotheses for the procedure-dependent PRs, the following thought experiment can be considered. Let us compare choice and separately-presented valuations, for two alternatives that are defined on one easy-to-evaluate and one difficult-to-evaluate attribute. According to the prominence and the task-goal hypotheses, the choice task represents an ordinal assessment and should evoke System-1 cognition, and System 2 should be activated for the cardinal valuations. The reverse would be predicted though from the perspective of evaluability, because the choice is in the JE mode and contains more information on attribute values. Moreover, no consistent prediction can be made across the extant hypotheses about how the relative attribute importance may change between these two evaluation procedures. The easy-to-evaluate attribute is prominent (because of its higher evaluability) and hence should loom larger in choice than in valuation according to the prominence/task-goal hypotheses, but the difficult-to-evaluate attribute would become more evaluable and hence more important in the JE mode (i.e., choice) than in the SE mode (i.e., valuation), a contradiction.

These inconsistencies can be reconciled by the theory of contextual deliberation. Unlike extant explanations, it does not need to pre-specify or fix the (desirable) level of cognitive efforts for a particular elicitation procedure. For instance, it does not presume that indecisive choices necessarily involve low-cognition heuristics. Instead, optimal deliberation is endogenously determined by gauging the benefits and the costs of information acquisition. As a result, choices can lead to varying levels of deliberation and reversals across the JE and the SE modes (e.g., Bohnet et al. 2016). Moreover, optimal deliberation need not increase with the availability of alternatives in the stimuli. This means that, in contrast to the evaluability hypothesis, choices in the JE mode need not be cognitively more effortful than separate valuations. Similarly, depending on the attribute values, optimal deliberation can be either higher or lower for ordinal tasks than for cardinal tasks (Theorem 2). Therefore, the problems in the above-mentioned thought experiment do not exist under the theory of contextual deliberation.

There are other findings that are supportive of contextual deliberation. Bazerman et al. (1992) show that preferences continue to be reversed between joint and separate evaluations, even when the respondents were given explicit pre-response information about possible attribute range.[19] This

[19] Bazerman et al. (1992) find that people tend to choose higher but unequal payoffs (e.g., \$600 to self and \$800 to other), but are more likely to give higher ratings for payoffs that are lower but equal (e.g., \$500 to self and \$500 to other). Their expectation is that cognitive efforts are lower under choices (JE mode) than under ratings (SE mode), and that the Self-Payoff attribute becomes less decisive under the reflective rating tasks than the Equal/Unequal

implies that the observed PR is not caused by the increasing availability of alternatives in the JE mode, which is against the evaluability hypothesis but consistent with the contextual deliberation theory. Moreover, List (2002) finds that experienced dealers are less likely to exhibit the joint-separate PR than inexperienced consumers. One plausible explanation is that lower deliberation costs for the dealers reduce the difference of optimal deliberation across the evaluation modes.

5 Experiments

We conduct five experiments to examine the role of different explanations in procedure-dependent preferences. We focus on documenting findings that are inconsistent with predictions of extant psychological explanations but can only be explained by contextual deliberation. The first four experiments involve incentive-based tasks on the choice-matching and the JE-SE procedures, respectively. The fifth experiment is a hybrid study that elicits hypothetical responses.

5.1 Studies 1 and 2 (The Choice-Matching PR)

The first two experiments involve a between-subjects design (choice vs. matching) that is typically implemented in the literature to examine the prominence effect. The experimental stimuli include two options in each of four categories. As shown in Table 1, each option is characterized by two attributes. Option A has relatively higher value on the first attribute, whereas option B is relatively advantageous on the second attribute. We consider two levels for the superior attribute of option B in each category, which are nonetheless missing for the matching task. We call these two variants the high-value option B (HB) and the low-value option B (LB), respectively. We intentionally make the HB-LB variation a within-subjects factor for the choice tasks, because it is necessarily within subjects for matching (i.e., only one matching task in each category). The experiment includes eight binary choices and four matching tasks, leading to a total of sixteen preference measures. Note that two preference measures can be inferred from each matching task, by comparing the stated matching value (i.e., d_M) with each of the two levels of option B's missing attribute.

The first experiment was run on participants of a data-based course (one master and two undergraduate sessions). In each session, the students were briefed about choice and attribute matching as two alternative methods for preference elicitation. They were then invited to take part in the study on a voluntary basis (about 74% of the enrolled students completed the experiment). The partici-

attribute. This is akin to our explanation for the prominence effect in Section 3.1 (i.e., the scenario in Figure 2(a)), and is different from the conventional interpretation that the harder-to-evaluate Self-Payoff attribute (relative to the Equal/Unequal attribute) becomes more evaluable under choices than under ratings (e.g., Hsee et al. 1999).

Table 1. Options and Attribute Values

	Attributes	Option A	Option B
Lottery	Probability of Winning	95%	50%
	Amount of Winning	$90	$230/$160
Face Mask	ASTM Level	Level 3	Level 2
	Pieces (Individual Pack)	40	70/50
Delayed Payoff	When to Receive	14 Days from Today	56 Days from Today
	Amount to Receive	$110	$150/$115
Chocolate	Product	GODIVA 72% Dark Chocolate Tablet	FERRERO Collection
	Pieces	10	32/18

Note: There are two variants for option B in each category, each taking a different value for the second attribute. The other attribute values are fixed. Both options for the Face Mask category have the following common features: Watson 3-PLY Medical Face Masks; Made in Hong Kong; Individual Pack; Extra Soft Earloop; Blue Color.

pants were provided a hyperlink to access the Qualtrics-based study. The experiment instructions were presented both on the webpage and by the experimenter (i.e., the course instructor).

The subjects were told that they would be randomly assigned to respond to either the choice or the matching tasks. They were instructed that, at the end of the class, three persons would be randomly selected. As a result, nine people were selected across the three sessions. For each selected person, one task would be randomly drawn for real implementation. We emphasized that the probability to draw each subject/task was the same. In each choice task, the subjects were asked to indicate the option they prefer. For each matching task, the subjects were asked to input a matching value (an integer between 0 and 300), on the missing attribute of option B, that makes the two options equally attractive. The subjects were instructed that, to implement a matching task, a random offer value between 0 and 300 would be generated. The selected subject would receive option B (with the offer value on the missing attribute) if the generated offer value was greater than or equal to the inputted matching value, and would be given option A if otherwise. It was emphasized that it was of the subjects' best interest to respond honestly to all tasks.

After the instructions were explained, the subjects were directed to respond to two practice questions (one on lotteries and the other on delayed payoffs) that were designed to improve their understanding of the incentive scheme for the matching tasks. They were then prompted to ask questions before being given a passcode to start the experiment. After responding to all tasks, the subjects were asked to fill up their demographic information. The payoff scheme was then executed based on the subjects' indicated preferences, by using Excel to draw random numbers.

Each task was presented on one single page. No page could be returned after the response was finalized, or skipped before moving to the next page. The response time (RT) for each task, defined as the duration (in seconds) from when the page was loaded to when the response on that page was finalized, was automatically recorded. This measure would be used to approximate preference-relevant cognitive efforts during the decision making process.

According to previous research, the first attribute in our setting can be interpreted as the more prominent attribute in each category, making option A the prominent option. For example, for Lottery the probability of wining is normally treated as the prominent attribute relative to the amount of winning (e.g., Slovic et al. 1990). Similarly, time or quality is the primary attribute while the amount to receive is secondary (e.g., Tversky et al. 1988, Carmon and Simonson 1998, Fischer et al. 1999). Moreover, as we have discussed extensively (Section 3.2.1), both the prominence and the task-goal hypotheses predict that the prominence effect should be observed irrespective of the attribute values, and that its occurrence is driven by a cognitively simpler process for choice than for matching. This leads us to test the following predictions.

H1 For both the HB and the LB variants, the proportion of preferring option A over option B is higher for choice than for matching.

H2 Across the HB and the LB variants, the RT under choice is negatively associated with the occurrence of the prominence effect.

We note that RT is a noisy measure of preference-relevant cognition: the subjects may spend a substantial proportion of time on other activities during the response process (e.g., semantic information processing, typing), which can be naturally shorter for choice than for matching. Therefore, to examine the preference-relevant cognition between choice and matching, it is not sensible to directly compare their RT levels. Instead, as in H2, we compare the choice RT between the HB and the LB variants, given that they correspond to the same matching task for each category.

The results on the revealed preferences between choice and matching are presented in Table 2.[20] We start with testing the null hypothesis, for each of the HB scenarios, that the proportion of subjects preferring option A over option B was the same across the two elicitation scales. The subjects were more likely to prefer option A in choice than in matching for three categories (two-tailed independent-samples t-test assuming unequal variances: $p < 0.010$ for Lottery and Face Mask, $p = 0.064$ for Chocolate). The preference for the category Delayed Payoff exhibited the opposite direction, but the null hypothesis could not be rejected. Similar pattern can be seen if the

[20]We report in the Appendix the summary statistics of the responses and the RT for matching in the first two experiments (Tables A1-A2), and those of the stated valuations in the JE-SE experiments (Tables A3-A4).

Table 2. Proportions of Preferring Option A between Choice and Matching in Study 1

	High Value for Option B				Low Value for Option B			
	Choice	Matching	Independent-Samples t	Mann-Whitney-U z	Choice	Matching	Independent-Samples t	Mann-Whitney-U z
Lottery	0.55	0.14	3.903	-3.375	0.75	0.75	0.000	0.000
			(<0.001)	(<0.001)			(1.000)	(1.000)
Face Mask	0.68	0.36	2.674	-2.571	0.83	0.71	1.043	-1.076
			(0.010)	(0.010)			(0.302)	(0.282)
Delayed Payoff	0.68	0.75	-0.669	-0.663	0.90	0.89	0.093	-0.095
			(0.506)	(0.507)			(0.926)	(0.925)
Chocolate	0.43	0.21	1.885	-1.794	0.68	0.64	0.270	-0.274
			(0.064)	(0.073)			(0.788)	(0.784)

Note: Tests of the equality, between choice and matching responses, of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=40 for choice and n=28 for matching. P values in the parenthesis.

non-parametric Mann-Whitney-U (hereafter MWU) test was used instead.

However, under the LB variant where option B's superior attribute value decreased, the null hypothesis can no longer be rejected for any category. That is, the differences of the proportions of subjects preferring option A between the elicitation scales became nonsignificant. As a result, there were clear discrepancies, in the impact of the elicitation scales on the revealed preferences, across the two variants of option B. The prominence effect observed for the HB variant was not sustained for the LB variant, leading us to reject H1 that is predicted by the prominence or the task-goal account. This result is unlikely to be driven by the choice tasks' within-subjects design increasing the salience to the subjects, because the differences in the revealed preferences between the HB and the LB variants were smaller for choice than for matching.

To examine the cognitive process underlying the revealed preferences, we compare the RT for each pair of the HB and the LB choice tasks. As shown in Table 3, across all categories the subjects spent significantly more time choosing between option A and the high-value variant of option B than that for the low-value variant. The results were remarkably similar for either the two-tailed paired-samples t-test or the Wilcoxon signed-rank (hereafter WSR) test. H2 is not supported, because the occurrence of the prominence effect was positively associated with the choice RT across the HB and the LB variants.

The robustness of the results is examined in the second experiment that involve some enhanced features: sample size, expected stake, task order, and experiment environment. In particular, we

Table 3. RT between Choice Tasks with High-Value and Low-Value Option B in Study 1

	High Value Option B		Low Value Option B			
	Mean	SD	Mean	SD	Paired-Samples t	Wilcoxon Singed-Rank z
Lottery	15.75	9.19	10.88	8.08	2.740 (0.009)	-3.024 (0.002)
Face Mask	25.81	19.23	6.75	2.90	6.431 (<0.001)	-5.444 (<0.001)
Delayed Payoff	10.66	8.26	6.30	6.34	2.595 (0.013)	-3.952 (<0.001)
Chocolate	13.43	12.05	6.03	2.45	3.908 (<0.001)	-4.597 (<0.001)

Note: Tests of the equality of RT (in seconds), between the choice tasks with high versus low value for option B's second attribute, for each category. Two-tailed paired-samples t tests. n=40. P values in the parenthesis.

increased the number of participants to 123. They were recruited from a subject pool and granted credit for a marketing course. In addition, thirty percent of the subjects were randomly drawn to receive response-based compensation, which increased the per-participant expected payment to about HK$83 (from about HK$47 in the first experiment). The task order was counterbalanced between the HB and the LB choices (the order across the categories was still fixed). The experiment was run in a lab setting where the subjects were seated at remotely separated desktop computers. It consisted of 7 sessions, each lasting about 30 minutes. The other components of the experiment design and procedure resembled those in experiment 1.

The results of experiment 2 are reported in Tables 4 and 5. Similar to experiment 1, we continue to find that there are three categories (i.e., Lottery, Face Mask, and Chocolate) in which the subjects were more likely to prefer option A in choice (than in matching) under the HB variant but not under the LB variant. Nevertheless, some new findings are notable. First, we now observe the reverse prominence effect for Delayed Payoff under the HB variant: the negative effect of the choice scale (vs. matching) on the prominent option A became clearly significant. Second, for Chocolate, changing option B's superior attribute value modified the sign of the impact of the elicitation scales on the revealed preferences. That is, under the HB variant the probability to favor option A was significantly larger in choice than in matching (the prominence effect), but became marginally significantly smaller under the LB variant (the reverse prominence effect). Third, across all categories the differences in the choice RT were not significantly different between the HB and the LB variants (Table 5). Again, neither H1 nor H2 can be supported.

To summarize, across the first two experiments, we find that the prominence effect (if any) was not sustainable across the two variants of option B, the reverse prominence effect could arise, and the choice RT when the prominence effect was observed was higher or comparable to that when it was not observed. Therefore, the prominence effect we observe (under the HB variant of three categories)

Table 4. Proportions of Preferring Option A between Choice and Matching in Study 2

	High Value for Option B				Low Value for Option B			
	Choice	Matching	Independent-Samples t	Mann-Whitney-U z	Choice	Matching	Independent-Samples t	Mann-Whitney-U z
Lottery	0.25	0.08	2.651	-2.634	0.80	0.69	1.385	-1.372
			(0.009)	(0.008)			(0.169)	(0.170)
Face Mask	0.46	0.22	2.857	-2.796	0.69	0.58	1.346	-1.338
			(0.005)	(0.005)			(0.181)	(0.181)
Delayed Payoff	0.59	0.81	-2.704	-2.659	0.92	0.95	-0.837	-0.847
			(0.008)	(0.008)			(0.404)	(0.397)
Chocolate	0.42	0.23	2.254	-2.230	0.53	0.69	-1.846	-1.833
			(0.026)	(0.026)			(0.067)	(0.067)

Note: Tests of the equality, between choice and matching responses, of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=59 for choice and n=64 for matching. P values in the parenthesis.

is unlikely to be explained by the prominence or the task-goal hypothesis, and the emergence of the reverse prominence effect (experiment 2) is utterly at odds with the psychological accounts. If it were the relative attribute prominence that led the prominent option A to be chosen comparatively more often, we should have observed qualitatively similar results for the LB scenarios. Moreover, if the subjects had followed the cognitively-saving heuristics to choose the so-called prominent option more frequently, they should have taken less time in the HB than in LB choices. However, our findings clearly run counter to these predictions of the psychological hypotheses.

Then how can the theory of contextual deliberation account for these findings? It is likely that the observed prominence effect under the three HB scenarios (i.e., Lottery, Face Mask, and Chocolate) conforms to the case in Figure 2(b): the subjects may have an ex ante inclination to prefer option B because its superior attribute value x_{B_2} is moderately high, but are induced to deliberate more and hence become relatively more likely to favor the other option A in choice than in matching. As x_{B_2} decreases, they may tend instead to favor option A instinctively, whereas this propensity can be mitigated by a higher deliberation. The reverse prominence effect we document under the LB variant for Chocolate (in experiment 2) can then be interpreted as being consistent with the scenario in Figure 3(a).[21] Similarly, it is possible that, given the attribute values for Delayed Payoff, the subjects have an ex ante intuitive preference for option A even under the HB variant so as to yield the reverse prominence effect (as in Figure 3(a)) in experiment 2.

[21] The other LB scenarios without significant choice-matching PR may correspond to a scenario that is somewhere between Figure 2(a) and 3(a).

Table 5. RT between Choice Tasks with High-Value and Low-Value Option B in Study 2

	High Value Option B		Low Value Option B			
	Mean	SD	Mean	SD	Paired-Samples t	Wilcoxon Singed-Rank z
Lottery	12.87	9.94	13.18	9.75	-0.192 (0.848)	-0.257 (0.797)
Face Mask	13.72	9.27	18.08	21.22	-1.454 (0.151)	-1.427 (0.154)
Delayed Payoff	8.56	4.85	7.81	4.00	0.951 (0.346)	-0.823 (0.411)
Chocolate	9.12	8.16	11.88	16.96	-1.121 (0.267)	-0.785 (0.432)

Note: Tests of the equality of RT (in seconds), between the choice tasks with high versus low value for option B's second attribute, for each category. Two-tailed paired-samples t tests. n=59. P values in the parenthesis.

5.2 Studies 3 and 4 (The Joint-Separate PR)

The third experiment consists of two between-subjects conditions (joint vs. separate evaluation), each involving the same set of eight valuation tasks. The stimuli were similar to those in the first two studies. In particular, we employed option A and the high-value option B in each category, because contextual deliberation would generate opposite predictions from those made by the currently dominant explanation for the JE-SE reversal. According to the evaluability hypothesis (e.g., Hsee et al. 1999, Bazerman et al. 1999, Zhang 2015), attributes with known range or involving dichotomous values or categorial variations are easier to evaluate than those with unknown range or incremental variations. This implies that, in each category of our setting, the first attribute is easier to evaluate than the other attribute (Table 1), which yields the following prediction.

H3 The relative preference for option A over option B is lower in JE than in SE.

In contrast, from the perspective of the contextual deliberation theory, we expect that option A is relatively more favorable in joint than in separate valuations. It is because, as can be seen from our choice-matching experiments (Tables 2 and 4), the subjects may have a natural propensity to favor the high-value option B (except for Delayed Payoff), whose ex ante advantage would nonetheless be mitigated as optimal deliberation increases from SE to JE. In other words, as illustrated in Figure 4(b), we expect to observe the nonstandard joint-separate PR.

The experimental procedure was similar to the first experiment, except some necessary modifications. Both experiments were run on the same pool of class participants, with a five-week interval between them. The subjects were asked to state an integer value between 0 and 300 that was equally attractive as each presented item. They were instructed that an offer price between 0 and 300 would be randomly generated. It was made clear that a selected subject would receive the generated offer price if it was greater than or equal to the stated value, and would be given the item if otherwise (i.e., the BDM scheme). As part of the instructions, two specific examples were provided to demonstrate that the expected welfare would be lowered if true valuations were either over- or under-stated. The subjects were randomly assigned to either the JE condition to value both options of a category jointly in the same page, or to the SE condition that involved one task per page. Accordingly, RT was recorded for each page but not necessarily for each task.[22]

We used two measures to capture the relative preference between the options. The first one was the difference in stated valuations (i.e., $d_A - d_B$), and the second was whether option A was valued higher than option B. As can be seen from Table 6, the mean valuation differences were negative across all categories and both evaluation modes, and the proportions of preferring option A were mostly less than half. This is basically consistent with our expectation (as well as the results in our choice-matching experiments) that option B had an ex ante advantage over option A. Importantly, the comparisons of the preference measures between the evaluation modes were all directionally opposite to what would be predicted by the evaluability hypothesis. In particular, the mean valuation differences were larger in JE than in SE, which were clearly significant for at least two categories (i.e., Face Mask and Chocolate) in both the independent-samples t and the MWU tests. This pattern became even more salient for the other preference measure: the increases in the likelihood of valuing option A more highly, from the SE to the JE mode, were clearly significant. These results were evidently against the evaluability hypothesis (i.e., H3), but were in support of the theory of contextual deliberation.

We run another experiment (study 4) to examine the robustness of the results. The basic design and procedure in experiment 3 were retained, except the following enhancements/modifications that were similar to those in experiment 2. We nearly doubled the sample size (to 135). All participants received credit for a marketing course. We randomly selected thirty percent of the subjects to receive response-based compensation, increasing the per-subject expected payment (from about HK$15 in the third experiment) to about HK$47. We counterbalanced the task order between the two options in separate valuations (while fixing the order across the categories). The experiment included seven 30-minutes lab sessions on remotely separated desktop computers.

The results of experiment 4 are presented in Table 7, which are remarkably similar to those in Table 6. The general pattern is inconsistent with the prediction of the evaluability hypothesis that the relative preference for option A should be smaller in JE than in SE (i.e., H3). Instead, we find that, moving from the SE to the JE mode, the mean differences in stated valuations and the

[22]We compare the RT of joint valuation with that of either separate valuation in each category. We find that the subjects spent significantly more time in the JE page than for each SE task. Nevertheless, we do not formally report this result which is not our main focus, because it can be predicted by both the theory of contextual deliberation and the evaluability hypothesis.

Table 6. Preferences for Option A between JE and SE in Study 3

	Mean Valuation Difference (A-B)				Proportion of Preferring Option A			
	JE	SE	Independent-Samples t	Mann-Whitney-U z	JE	SE	Independent-Samples t	Mann-Whitney-U z
Lottery	-22.14	-38.21	1.716 (0.091)	-1.488 (0.137)	0.26	0.06	2.322 (0.024)	-2.233 (0.026)
Face Mask	-4.00	-42.18	3.358 (0.001)	-3.501 (<0.001)	0.51	0.14	3.478 (<0.001)	-3.212 (<0.001)
Delayed Payoff	-1.14	-10.38	1.106 (0.273)	-1.868 (0.062)	0.57	0.26	2.680 (0.009)	-2.562 (0.010)
Chocolate	-6.03	-38.56	2.451 (0.017)	-2.307 (0.021)	0.46	0.18	2.595 (0.012)	-2.483 (0.013)

Note: Tests of the equality, between JE and SE responses, of mean valuation differences and of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=35 for JE and n=34 for SE. P values in the parenthesis.

proportions of valuing option A more than option B were increased, especially for Face Mask and Chocolate. These results provide further support to the theory of contextual deliberation.

5.3 A Hybrid Study with Hypothetical Tasks

We aim to investigate the robustness of our results, when the incentive for deliberation is supposedly low and heuristics are expected to be employed more often in the decision-making process. To this end, we follow previous studies in the literature to use hypothetical tasks in the fifth experiment. We integrate both the choice-matching and the JE-SE tasks as two separate parts in the experiment. We seek to enhance independence between these two major parts of the questionnaire by inserting between them a set of forty questions on the Big Five Personality Traits. The preference-elicitation tasks were the same as in experiments 1-4, except that no monetary incentive was offered. Accordingly, the instructions were simplified and no practice questions were asked. We run the experiment as an online survey by recruiting participants from a subject pool at a public university. They voluntarily participated in the survey in exchange for course credit. After registering through a cloud-based research system, they were provided a hyperlink to access the online study and completed the responses at their own pace. The subjects were randomly assigned to either condition of the evaluation scale/mode within either part of the survey.

The results on the choice-matching tasks are presented in Tables 8 and 9, for the relative preferences and the choice RT, respectively. The overall pattern is remarkably similar to that in

Table 7. Preferences for Option A between JE and SE in Study 4

	Mean Valuation Difference (A-B)				Proportion of Preferring Option A			
	JE	SE	Independent-Samples t	Mann-Whitney-U z	JE	SE	Independent-Samples t	Mann-Whitney-U z
Lottery	-28.54	-41.15	1.563	-1.963	0.12	0.07	0.898	-0.900
			(0.120)	(0.050)			(0.371)	(0.368)
Face Mask	-6.13	-19.70	2.527	-3.396	0.42	0.25	2.085	-2.062
			(0.013)	(<0.001)			(0.039)	(0.039)
Delayed Payoff	-14.88	-19.76	0.802	-0.863	0.40	0.29	1.326	-1.323
			(0.424)	(0.388)			(0.187)	(0.186)
Chocolate	-13.29	-39.68	3.823	-3.260	0.31	0.10	3.091	-3.005
			(<0.001)	(0.001)			(0.003)	(0.003)

Note: Tests of the equality, between JE and SE responses, of mean valuation differences and of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=67 for JE and n=68 for SE. P values in the parenthesis.

the first two experiments (Tables 2-5). For Lottery and Face Mask, the subjects were significantly more likely to favor option A over the high-value option B in choice than in matching (i.e., the prominence effect), whereas the differences were not significant between option A and the low-value option B. The results on Delayed Payoff are even harder to be explained by the psychological hypotheses. The preference for option A over either variant of option B decreased from matching to choice. That is, the probability of choosing the prominent option was persistently smaller than that implied by matching (i.e., the reverse prominence effect). The most challenging case emerged for Chocolate. The relative preference of option A over the HB variant was significantly higher, but significantly lower for that over the LB variant, in choice than in matching. Consequently, no matter how the relative prominence between the attributes would be classified, the co-existence of the standard and the nonstandard choice-matching PRs is necessarily incompatible with the prominence or the task-goal hypothesis. Moreover, across all categories, the time the subjects spent responding to the HB choice task was higher or comparable to that for the corresponding LB choice task, contradicting H2. Nevertheless, as we have discussed in Section 5.1, the theory of contextual deliberation can be applied to account for the findings in Tables 8 and 9, including the intriguing emergence of the choice-matching PR in both directions for the same category (Chocolate).

Next, we report the results on the JE-SE tasks in Table 10, which are notably similar to those in our incentive-based studies (Tables 6-7). The mean differences in stated valuations between options A and B increased from SE to JE, which were significant or marginally significant except for Delayed Payoff. The mean valuation differences for Face Mask and Chocolate even changed

Table 8. Proportions of Preferring Option A between Choice and Matching in Study 5

	High Value for Option B				Low Value for Option B			
	Choice	Matching	Independent-Samples t	Mann-Whitney-U z	Choice	Matching	Independent-Samples t	Mann-Whitney-U z
Lottery	0.74	0.22	8.252 (<0.001)	-7.044 (<0.001)	0.87	0.80	1.165 (0.246)	-1.163 (0.245)
Face Mask	0.71	0.34	5.345 (<0.001)	-4.976 (<0.001)	0.84	0.87	-0.452 (0.652)	-0.453 (0.651)
Delayed Payoff	0.51	0.71	-2.852 (0.005)	-2.799 (0.005)	0.82	0.92	-2.045 (0.042)	-2.031 (0.042)
Chocolate	0.52	0.28	3.494 (<0.001)	-3.392 (<0.001)	0.64	0.81	-2.585 (0.011)	-2.548 (0.011)

Note: Tests of the equality, between choice and matching responses, of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=90 for choice and n=91 for matching. P values in the parenthesis.

Table 9. RT between Choice Tasks with High-Value and Low-Value Option B in Study 5

	High Value Option B		Low Value Option B			
	Mean	SD	Mean	SD	Paired-Samples t	Wilcoxon Singed-Rank z
Lottery	11.42	11.29	8.07	5.49	2.691 (0.009)	-3.159 (0.002)
Face Mask	11.32	22.33	9.18	19.76	1.238 (0.219)	-1.121 (0.262)
Delayed Payoff	11.75	33.84	4.47	2.20	2.043 (0.044)	-4.484 (<0.001)
Chocolate	21.98	158.04	6.06	6.45	0.962 (0.339)	-0.195 (0.845)

Note: Tests of the equality of RT (in seconds), between the choice tasks with high versus low value for option B's second attribute, for each category. Two-tailed paired-samples t tests. n=90. P values in the parenthesis.

signs across the evaluation modes. Similarly, the positive impact of JE on the probability of valuing option A more highly was unambiguously significant in the same three categories. These findings provide additional evidence to reject the prediction of the evaluability hypothesis (i.e., H3) but support that of contextual deliberation.

6 Summary and Discussions

Revealed preferences over alternatives can be systematically reversed, depending on how they are measured. The PR phenomena have been observed across a variety of elicitation procedures (e.g.,

Table 10. Preferences for Option A between JE and SE in Study 5

	Mean Valuation Difference (A-B)				Proportion of Preferring Option A			
	JE	SE	Independent-Samples t	Mann-Whitney-U z	JE	SE	Independent-Samples t	Mann-Whitney-U z
Lottery	-12.02	-32.45	1.893	-1.782	0.37	0.20	2.580	-2.563
			(0.060)	(0.075)			(0.011)	(0.010)
Face Mask	11.84	-21.95	3.541	-3.704	0.59	0.33	3.706	-3.590
			(<0.001)	(<0.001)			(<0.001)	(<0.001)
Delayed Payoff	-6.02	-11.62	0.637	-0.522	0.48	0.47	0.041	-0.041
			(0.525)	(0.602)			(0.967)	(0.967)
Chocolate	15.16	-19.79	3.186	-3.598	0.57	0.33	3.369	-3.284
			(0.002)	(<0.001)			(<0.001)	(0.001)

Note: Tests of the equality, between JE and SE responses, of mean valuation differences and of the proportions of subjects who preferred option A over the corresponding option B, for each category. Two-tailed independent-samples t tests (assuming unequal variances). n=86 for JE and n=95 for SE. P values in the parenthesis.

choice, matching, valuation, joint or separate evaluation), and may lead to unintended consequences. They cast direct challenges to the fundamental principle of procedure invariance, an implicit and critical axiom that has been universally assumed in academic research and in practice. Extant accounts for these puzzles center on ad hoc hypotheses about how domain-specific characteristics (e.g., attribute prominence/evaluability) may relate to a particular type of PR. This paper offers a synthesizing theory of contextual deliberation to rationalize procedure-dependent PRs. The primitives of the theoretical framework are completely procedure independent, whereas it is endogenous deliberation that may lead to the emergence of procedure-dependent preferences. Therefore, this theory provides a common micro-foundation for these seemingly disparate phenomena, reconciles extant hypotheses, and parsimoniously organizes apparently unrelated/inconsistent empirical findings within and across different types of procedure-dependent PRs.

We report five experiments that are designed to identify the potential explanations of procedure-dependent PRs. We find that the influence of elicitation procedures on revealed preferences can diverge systematically, across tasks with varying attribute values, in ways that are unexpected in the eyes of extant psychological perspectives. In addition, the change in the cognition level underlying decision making (as measured by response time) is inconsistent with the prediction of previous explanations of the prominence effect. We even observe the reverse prominence effect and the nonstandard joint-separate PR, which are opposite to what would be predicted by the extant accounts. Nevertheless, contextual deliberation can reconcile these seemingly puzzling results. Our experiments also generate new and novel findings that have not been documented in the past.

Our experiments are designed to directly falsify extant psychological explanations (e.g., the prominence hypothesis, the evaluability hypothesis), but contextual deliberation *per se* is falsifiable only in the experiments on the joint-separate PR (not in those on the choice-matching PR). The issue is that, as we have shown in Section 3.1 (e.g., Theorems 1 and 2), the impact of varying the attribute value x_{B_2} on the observables (revealed preferences and response time) is quite complicated and may change signs multiple times. Therefore, as the current experiment design involves only two variants for x_{B_2}, definite predictions cannot be made by the theory of contextual deliberation. However, this does not mean that the contextual deliberation theory is unfalsifiable (as an explanation for the choice-matching PR). For example, it predicts that the dependence of preferences on elicitation procedures can be non-monotonically moderated by the cost/need of deliberation. Intuitively, when the deliberation cost becomes too low or too high, optimal deliberation would be basically invariant of elicitation procedures. This means that scale- and mode-dependent PRs would be mitigated when deliberation becomes either costless or extremely costly. Similar falsifiable prediction can be made for manipulations that exogenously modify the respondents' preference information. These predictions can be tested in future research by examining antecedents or proxies of the cost/need of deliberation, e.g., cognitive skills, time constraint, depletion of cognitive resources, sample provision, preference articulation, emotions, and task order.

At a more general level, the theory of contextual deliberation offers a normative perspective on the elicitation of preferences. The basic point is that, conditional on information, inherent preferences are well defined and can be represented by expected utilities. Nevertheless, revealed preferences can vary systematically, depending on how preference information is retrieved/acquired via optimal deliberation. Essentially, a person's preference structure over a set of alternatives may involve multiple stochastic orders, each one corresponding to a particular level of information/delibeation. Each instance of responding to a preference-eliciting task evokes one of the stochastic orders, as determined by optimal deliberation that may vary endogenously with seemingly irrelevant factors such as context and/or elicitation procedures. In other words, apparently context- and procedure-dependent preferences may simply be different manifestations of the same preference structure that is inherently stable and context/procedure free. In short, our view is that preferences can indeed be constructed, but preference construction *per se* can be endogenous and stable.

References

Attema, Arthur E., and Werner B. F. Brouwer (2013), "In Search of A Preferred Preference Elicitation Method: A Test of the Internal Consistency of Choice and Matching Tasks," *Journal of Economic Psychology*, 39, 126-140.

Aumann, Robert J. (1962), "Utility Theory without the Completeness Axiom," *Econometrica*, 30 (3), 445462.

Bartos, Vojtech, Michal Bauer, Julie Chytilova, and Filip Matejka (2016), "Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition," *American Economic Review*, 106 (6), 1437-1475.

Bazerman, Max H., G. F. Loewenstein, and S. B. White (1992), "Reversals of Preference in Allocation Decisions: Judging an Alternative versus Choosing among Alternatives," *Administrative Science Quarterly*, 37, 220-240.

Bazerman, Max H., Don A. Moore, Ann E. Tenbrunsel, Kimberly A. Wade-Benzoni, and Sally Blount (1999), "Explaining How Preferences Change across Joint versus Separate Evaluation," *Journal of Economic Behavior and Organization*, 39, 41-58.

Bazerman, Max H., H. Schroth, P. P. Shah, K. A. Diekmann, A. E. Tenbrunsel (1994), "The Inconsistent Role of Comparison Others and Procedural Justice in Reactions to Hypothetical Job Descriptions: Implications for Job Acceptance Decisions," *Organizational Behavior and Human Decision Processes*, 60, 326-352.

Bazerman, Max. H., Ann. E. Tenbrunsel, and Kimberly A. Wade-Benzoni (1998), "Negotiating with Yourself and Losing: Understanding and Managing Competing Internal Preferences," *Academy of Management Review*, 23, 225-241.

Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak (1964), "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 9(3), 226-232.

Bohnet, Iris, Alexandra van Geen, and Max Bazerman (2016), "When Performance Trumps Gender Bias: Joint vs. Separate Evaluation," *Management Science*, 62(5), 1225-1234.

Caplin, Andrew, and Mark Dean (2015), "Revealed Preference, Rational Inattention, and Costly Information Acquisition," *American Economic Review*, 105 (7), 2183-2203.

Carmon, Ziv, and Itamar Simonson (1998), "Price-Quality Trade-offs in Choice Versus Matching: New Insights Into the Prominence Effect," *Journal of Consumer Psychology*, 7 (4), 323-343.

Cerreia-Vioglioy, Simone, David Dillenbergerz, and Pietro Ortoleva (2015), "Cautious Expected Utility and the Certainty Effect," *Econometrica*, 83 (2), 693728.

Cerreia-Vioglio, S., F. Maccheroni, M. Marinacci, and L. Montrucchio (2011), "Uncertainty Averse Preferences," *Journal of Economic Theory*, 146 (4), 12751330.

Chernev, Alexander (2005), "Feature Complementarity and Assortment in Choice," *Journal of Consumer Research*, 31 (4), 748-759.

Enke, Benjamin, and Thomas Graeber (2019), "Cognitive Uncertainty," working paper.

Ergin, Haluk, and Todd Sarver (2010), "A Unique Costly Contemplation Representation," *Econo-*

metrica, 78 (4), 1285-1339.

Fischer, Gregory W., Ziv Carmon, Dan Ariely, and Gal Zauberman (1999), "Goal-based Construction of Preferences: Task Goals and the Prominence Effect," *Management Science*, 45 (8), 1057-1075.

Gabaix, Xavier (2019), "Behavioral Inattention," in Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., Handbook of Behavioral Economics-Foundations and Applications 2, Elsevier, 261-343.

Gabaix, Xavier, and David Laibson (2017), "Myopia and Discounting," working paper.

Gilboa, I., F. Maccheroni, M. Marinacci, and D. Schmeidler (2010), "Objective and Subjective Rationality in A Multiple Prior Model," *Econometrica*, 78 (2), 755770.

Gilboa, I., and D. Schmeidler (1989), "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18 (2), 141-153.

Gonzalez-Vallejo, Claudia, and Elizabeth Moran (2001), "The Evaluability Hypothesis Revisited: Joint and Separate Evaluation Preference Reversal as a Function of Attribute Importance," *Organizational Behavior and Human Decision Processes*, 86 (2), 216-233.

Guo, Liang (2016), "Contextual Deliberation and Preference Construction," *Management Science*, 62 (10), 2977-2993.

Guo, Liang (2021), "Contextual Deliberation and the Choice-Valuation Preference Reversal," *Journal of Economic Theory*, 195 (July), 105285.

Guo, Liang (2022), "Testing the Role of Contextual Deliberation in the Compromise Effect," *Management Science*, 68 (6), 4326-4355.

Guo, Liang (2023), "Recovering the Anchoring of Economic Valuations," *American Economic Journal: Microeconomics*, forthcoming.

Hsee, Christopher K. (1996), "The Evaluability Hypothesis: An Explanation of Preference Reversals between Joint and Separate Evaluations of Alternatives," *Organizational Behavior and Human Decision Processes*, 46, 247-257.

Hsee, Christopher K., Jean-Pierre Dube, and Yan Zhang (2008), "The Prominence Effect in Shanghai Apartment Prices," *Journal of Marketing Research*, April, 133-144.

Hsee, Christopher K., George F. Loewenstein, Sally Blount, and Max H. Bazerman (1999), "Preference Reversals between Joint and Separate Evaluations of Options: A Review and Theoretical Analysis," *Psychological Bulletin*, 125(5), 576-590.

Johnson, Justin P., and David P. Myatt (2006), "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96(3), 756-784.

Kahneman, Daniel (2003), "A Perspective on Judgment and Choice: Mapping Bounded Rationality," *American Psychologist*, 58 (9), 697.

Kahneman, Daniel, and Miller, D. T. (1986), "Norm Theory: Comparing Reality with Its Alternatives," *Psychological Review*, 93, 136-153.

Kahneman, Daniel, and I. Ritov (1994), "Determinants of Stated Willingness to Pay for Public

Goods: A Study in the Headline Method," *Journal of Risk and Uncertainty*, 9, 5-38.

Khaw, Mel Win, Ziang Li, and Michael Woodford (2017), "Risk Aversion as A Perceptual Bias," working paper.

List, John A. (2002), "Preference Reversals of A Different Kind: The 'More Is Less' Phenomenon," *American Economic Review*, 92(5), 1636-1643.

Loewenstein, G. F., Blount, S., and Bazerman, M. H. (1994), "Reversals of Preference Between Independent and Simultaneous Evaluation of Alternatives," unpublished manuscript, Carnegie Mellon University.

Matejka, Filip, and Alisdair McKay (2015), "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model," *American Economic Review*, 105 (1), 272-298.

Miller, Klaus M., Reto Hofstetter, Harley Krohmer, and Z. John 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.

Nowlis, Stephen M., and Itamar Simonson (1997), "Attribute-Task Compatibility as a Determinant of Consumer Preference Reversals," *Journal of Marketing Research*, 34 (May), 205-218.

O'Donnell, Michael, and Ellen R. K. Evers (2019), "Preference Reversals in Willingness to Pay and Choice," *Journal of Consumer Research*, 45 (6), 1315-1330.

Ok, E. (2002), "Utility Representation of An Incomplete Preference Relation," *Journal of Economic Theory*, 104 (2), 429449.

Ok, E., P. Ortoleva, and G. Riella (2012), "Incomplete Preferences under Uncertainty: Indecisiveness in Beliefs vs. Tastes," *Econometrica*, 80 (4), 17911808.

Sims, Christopher A. (2003), "Implications of Rational Inattention," *Journal of Monetary Economics*, 50 (3), 665-690.

Slovic, Paul, Dale Griffin, and Amos Tversky (1990), "Compatibility Effects in Judgment and Choice," in R. Hogarth (Ed.) *Insights in decision making: Theory and applications*, University of Chicago Press.

Tversky, Amos, Shmuel Sattath, and Paul Slovic (1988), "Contingent Weighting in Judgment and Choice," *Psychological Review*, 95 (3), 371-384.

Tversky, Amos, Paul Slovic, and Daniel Kahneman (1990), "The Causes of Preference Reversal," *American Economic Review*, 80 (March), 204-217.

Woodford, Michael (2009), "Information Constrained State Dependent Pricing," *Journal of Monetary Economics*, 56, S100-S124.

Zhang, Jiao (2015), "Joint vs. Separate Modes of Evaluation: Theory and Practice," *The Wiley Blackwell Handbook of Judgment and Decision Making*, Gideon Keren and George Wu eds., West Sussex, UK: Wiley Blackwell, 213-238.

Appendix

Proof of Theorem 2: The ex ante expected surplus for the choice task is $U_C(\alpha) = \int_{\underline{\theta}}^{\dot{\theta}} u_B\ h(\tilde{\theta},\alpha)d\tilde{\theta} + \int_{\dot{\theta}}^{\bar{\theta}} u_A\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha) = \hat{\theta}v_1(x_{A_1}) + v_2(x_{A_2}) + \int_{\underline{\theta}}^{\dot{\theta}} [\tilde{\theta}(v_1(x_{B_1}) - v_1(x_{A_1})) + v_2(x_{B_2}) - v_2(x_{A_2})]\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$, where $\dot{\theta} = \frac{v_2(x_{B_2}) - v_2(x_{A_2})}{v_1(x_{A_1}) - v_1(x_{B_1})}$. Integrating by parts yields $U_C(\alpha) = \hat{\theta}v_1(x_{A_1}) + v_2(x_{A_2}) + \int_{\underline{\theta}}^{\dot{\theta}} (v_1(x_{A_1}) - v_1(x_{B_1}))\ H(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$. It follows that the first-order condition for the deliberation choice is:

$$\frac{\partial U_C(\alpha)}{\partial \alpha} = \int_{\underline{\theta}}^{\dot{\theta}} [v_1(x_{A_1}) - v_1(x_{B_1})]\ \frac{\partial H(\tilde{\theta},\alpha)}{\partial \alpha} d\tilde{\theta} - k'(\alpha).$$

The marginal impact of x_{B_2} on the above first-order condition is:

$$\frac{\partial^2 U_C(\alpha)}{\partial \alpha \partial x_{B_2}} = [v_1(x_{A_1}) - v_1(x_{B_1})] \frac{\partial H(\dot{\theta},\alpha)}{\partial \alpha} \frac{\partial \dot{\theta}}{\partial x_{B_2}}.$$

Note that $\dot{\theta} = \frac{v_2(x_{B_2}) - v_2(x_{A_2})}{v_1(x_{A_1}) - v_1(x_{B_1})}$ is increasing in x_{B_2}. When x_{B_2} is relatively low such that $\dot{\theta} < \theta^+$, we would have $\frac{\partial^2 U_C(\alpha)}{\partial \alpha \partial x_{B_2}} > 0$. This is because $v_1(x_{A_1}) > v_1(x_{B_1})$ and, according to Assumption 1, $\frac{\partial H(\dot{\theta},\alpha)}{\partial \alpha} > 0$ for $\dot{\theta} < \theta^+$. This implies that optimal deliberation under choice would increase with x_{B_2} for relatively low x_{B_2}. Conversely, when x_{B_2} is relatively high such that $\dot{\theta} > \theta^+$, $\frac{\partial^2 U_C(\alpha)}{\partial \alpha \partial x_{B_2}}$ would be negative, which implies that α_C^* would decrease with x_{B_2}.

In addition, when x_{B_2} is sufficiently low such that $\dot{\theta} \to \underline{\theta}$, or when x_{B_2} is sufficiently high such that $\dot{\theta} \to \bar{\theta}$, the first part of the above first-order condition (i.e., the marginal value of deliberation) would converge to zero. It follows that optimal deliberation would then converge to the lower bound, i.e., $\alpha_C^* \to \underline{\alpha}$.

Consider then the matching task. The interim utility as a function of $\tilde{\theta}$ is $U_M(\tilde{\theta}) = U_M(d_M^*, u_{\mathcal{S}}) = \int_{\underline{y}}^{d_M^*} u_A\ w(y)dy + \int_{d_M^*}^{\bar{y}} u_\Lambda(y)\ w(y)dy = \int_{\underline{y}}^{d_M^*} [\tilde{\theta}v_1(x_{A_1}) + v_2(x_{A_2})]\ w(y)dy + \int_{d_M^*}^{\bar{y}} [\tilde{\theta}v_1(x_{B_1}) + v_2(y)]\ w(y)dy$, where $d_M^* = v_2^{-1}\left(\tilde{\theta}[v_1(x_{A_1}) - v_1(x_{B_1})] + v_2(x_{A_2})\right)$. The first-order derivative with respect to $\tilde{\theta}$ is $U_M'(\tilde{\theta}) = \int_{\underline{y}}^{d_M^*} v_1(x_{A_1})\ w(y)dy + \int_{d_M^*}^{\bar{y}} v_1(x_{B_1})\ w(y)dy$. The second-order derivative is then:

$$U_M''(\tilde{\theta}) = [v_1(x_{A_1}) - v_1(x_{B_1})] w(d_M^*) \frac{\partial d_M^*}{\partial \tilde{\theta}}.$$

Note that $\frac{\partial d_M^*}{\partial \tilde{\theta}} > 0$, which means that $U_M''(\tilde{\theta}) > 0$ for all $\tilde{\theta}$. Given that the interim utility is convex in $\tilde{\theta}$, the ex ante expected surplus $U_M(\alpha) = \int_{\underline{\theta}}^{\bar{\theta}} U_M(\tilde{\theta})\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$ is strictly increasing in α at $\alpha \to \underline{\alpha}$. It follows that $\alpha_M^* > \underline{\alpha}$ for any x_{B_2}.

Next, we compare optimal deliberation across the elicitation scales for intermediate x_{B_2}. Integrating by parts, we can rewrite the ex ante expected surplus for the matching task as $U_M(\alpha) = \int_{\underline{\theta}}^{\bar{\theta}} U_M(\tilde{\theta})\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha) = U_M(\bar{\theta}) - \int_{\underline{\theta}}^{\bar{\theta}} U_M'(\tilde{\theta}) H(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$. Similarly, the ex ante expected surplus for the choice task is $U_C(\alpha) = \int_{\underline{\theta}}^{\bar{\theta}} U_C(\tilde{\theta})\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha) = U_C(\bar{\theta}) - \int_{\underline{\theta}}^{\dot{\theta}} U_C'(\tilde{\theta}) H(\tilde{\theta},\alpha)d\tilde{\theta} -$

$\int_{\underline{\theta}}^{\bar{\theta}} U'_C(\tilde{\theta})H(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$, where $U_C(\tilde{\theta}) = \max\{u_A(\tilde{\theta}), u_B(\tilde{\theta})\}$ is the optimal interim utility under the choice problem. The comparison between the marginal impacts of deliberation on the ex ante expected surpluses across the elicitation methods leads to:

$$\frac{\partial U_C(\alpha)}{\partial \alpha} - \frac{\partial U_M(\alpha)}{\partial \alpha} = \int_{\underline{\theta}}^{\dot{\theta}} [U'_M(\tilde{\theta}) - U'_C(\tilde{\theta})]\frac{\partial H(\tilde{\theta},\alpha)}{\partial \alpha}d\tilde{\theta} + \int_{\dot{\theta}}^{\bar{\theta}} [U'_M(\tilde{\theta}) - U'_C(\tilde{\theta})]\frac{\partial H(\tilde{\theta},\alpha)}{\partial \alpha}d\tilde{\theta}.$$

Recall that $U'_M(\tilde{\theta}) = \int_{\underline{y}}^{d^*_M} v_1(x_{A_1})\ w(y)dy + \int_{d^*_M}^{\bar{y}} v_1(x_{B_1})\ w(y)dy$, which is less than $v_1(x_{A_1})$ and higher than $v_1(x_{B_1})$. In comparison, $U'_C(\tilde{\theta}) = v_1(x_{B_1})$ for $\tilde{\theta} < \dot{\theta}$, and $U'_C(\tilde{\theta}) = v_1(x_{A_1})$ for $\tilde{\theta} > \dot{\theta}$. Therefore, $\frac{\partial U_C(\alpha)}{\partial \alpha} - \frac{\partial U_V(\alpha)}{\partial \alpha} > 0$ as $\dot{\theta} \to \theta^+$. This implies that, when x_{B_2} is such that $\dot{\theta} = \theta^+$, we must have $\alpha^*_C > \alpha^*_M$. The theorem follows from recalling that α^*_C first increases and then decrease in x_{B_2}, α^*_M is constant in x_{B_2}, and $\alpha^*_C < \alpha^*_M$ for sufficiently low x_{B_2} or sufficiently high x_{B_2}.

Proof of Theorem 3: Consider first the SE mode. The interim utility as a function of $\tilde{\theta}$ is $U_{Si}(\tilde{\theta}) = U_S(d^*_{Si}, u_i) = \int_{\underline{y}}^{d^*_{Si}}[u_i - v(y)]/(\bar{y}-\underline{y})dy = \int_{\underline{y}}^{d^*_{Si}}[\tilde{\theta}v_1(x_{i_1}) + v_2(x_{i_2}) - v(y)]/(\bar{y}-\underline{y})dy$, where $d^*_{Si} = v^{-1}(u_i)$, $i \in \mathcal{S} = \{A, B\}$. Let $U_{Si}(\alpha) = \int_{\underline{\theta}}^{\bar{\theta}} U_{Si}(\tilde{\theta})\ h(\tilde{\theta},\alpha)d\tilde{\theta} - k(\alpha)$ be the ex ante expected surplus when option i is evaluated under the SE mode. Note that the ex ante expected surplus under the JE mode is $U_J(\alpha) = U_{SA}(\alpha) + U_{SB}(\alpha) + k(\alpha)$. This implies that $\frac{\partial U_J(\alpha^*_{Si})}{\partial \alpha} > \frac{\partial U_{Si}(\alpha^*_{Si})}{\partial \alpha} = 0$. As a result, $\alpha^*_J > \alpha^*_{Si}$, $i \in \mathcal{S} = \{A, B\}$.

Next, we compare optimal deliberation between the options under the SE mode. The first-order derivative of $U_{Si}(\tilde{\theta})$ with respect to $\tilde{\theta}$ is $U'_{Si}(\tilde{\theta}) = \int_{\underline{y}}^{d^*_{Si}} v_1(x_{i_1})/(\bar{y} - \underline{y})dy = v_1(x_{i_1})(d^*_{Si} - \underline{y})/(\bar{y} - \underline{y})$. The second-order derivative with respect to $\tilde{\theta}$ is then:

$$U''_{Si}(\tilde{\theta}) = v_1(x_{i_1})/(\bar{y} - \underline{y})\frac{\partial d^*_{Si}}{\partial \tilde{\theta}} = v_1(x_{i_1})^2 \left(v^{-1}\right)'(u_i)/(\bar{y} - \underline{y}).$$

In addition, the difference of the marginal information value between the alternatives is:

$$\frac{\partial U_{SA}(\alpha)}{\partial \alpha} - \frac{\partial U_{SB}(\alpha)}{\partial \alpha} = \int_{\underline{\theta}}^{\bar{\theta}} [U''_{SA}(\tilde{\theta}) - U''_{SB}(\tilde{\theta})]\ \int_{\underline{\theta}}^{\tilde{\theta}} \frac{\partial H(x,\alpha)}{\partial \alpha}dxd\tilde{\theta}.$$

As a result, $\alpha^*_{SA} > \alpha^*_{SB}$ if $U''_{SA}(\tilde{\theta}) - U''_{SB}(\tilde{\theta}) > 0$, i.e., $v_1(x_{A_1})^2\,(v^{-1})'(u_A) > v_1(x_{B_1})^2\,(v^{-1})'(u_B)$ for all $\tilde{\theta} \in [\underline{\theta}, \bar{\theta}]$, and vice versa.

Table A1. Summary Statistics for Matching in Study 1

	Mean	Median	SD	Min	Max
			Matching Responses		
Lottery	173.79	177.50	69.33	10	300
Face Mask	65.14	60.00	36.99	20	200
Delayed Payoff	217.29	200.00	83.16	0	300
Chocolate	33.93	20.00	55.38	3	300
			RT for Matching		
Lottery	69.31	23.48	230.82	6.92	1244.56
Face Mask	30.17	23.49	24.36	9.51	128.42
Delayed Payoff	27.70	20.19	26.65	7.58	146.37
Chocolate	47.68	16.33	132.79	8.71	720.79

Note: n=28.

Table A2. Summary Statistics for Matching in Study 2

	Mean	Median	SD	Min	Max
			Matching Responses		
Lottery	169.84	180.00	53.51	42	300
Face Mask	63.38	60.00	43.85	1	250
Delayed Payoff	240.25	265.00	65.19	90	300
Chocolate	37.09	24.00	51.86	2	300
			RT for Matching		
Lottery	27.65	22.47	18.63	8.33	91.50
Face Mask	31.48	26.74	18.65	6.22	114.22
Delayed Payoff	25.24	20.05	17.01	5.45	77.29
Chocolate	32.03	27.26	19.61	4.97	97.95

Note: n=64.

Table A3. Summary Statistics for Stated Valuations in Study 3

	Mean	Median	SD	Min	Max
			Joint Valuations		
Lottery A	74.29	80.00	19.12	10	100
Lottery B	96.43	100.00	47.09	5	200
Face Mask A	78.43	80.00	42.94	0	200
Face Mask B	82.43	70.00	45.10	0	200
Delayed Payoff A	104.00	100.00	36.64	50	270
Delayed Payoff B	105.14	100.00	51.37	20	300
Chocolate A	82.23	70.00	68.19	5	250
Chocolate B	88.26	65.00	64.68	0	300
			Separate Valuations		
Lottery A	77.15	81.50	19.43	3	95
Lottery B	115.35	115.00	45.49	5	240
Face Mask A	67.53	55.00	52.61	3	200
Face Mask B	109.71	80.00	86.05	5	300
Delayed Payoff A	96.12	100.00	49.08	5	300
Delayed Payoff B	106.50	110.00	50.61	8	300
Chocolate A	38.32	30.00	32.74	0	150
Chocolate B	76.88	56.50	66.54	0	300

Note: n=35 for JE and n=34 for SE.

Table A4. Summary Statistics for Stated Valuations in Study 4

	Mean	Median	SD	Min	Max
			Joint Valuations		
Lottery A	86.83	85.00	52.82	10	300
Lottery B	115.37	115.00	51.72	20	300
Face Mask A	68.81	45.00	55.28	0	300
Face Mask B	74.93	60.00	54.77	0	300
Delayed Payoff A	99.39	100.00	34.79	8	300
Delayed Payoff B	114.27	120.00	54.36	3	300
Chocolate A	64.07	40.00	63.60	10	300
Chocolate B	77.36	70.00	48.94	5	260
			Separate Valuations		
Lottery A	80.05	85.00	26.67	10	250
Lottery B	121.21	115.00	46.76	20	300
Face Mask A	53.39	47.50	37.45	1.5	200
Face Mask B	73.09	67.50	47.65	1	210
Delayed Payoff A	93.99	100.00	25.73	5	200
Delayed Payoff B	113.75	110.00	50.45	2	300
Chocolate A	40.28	32.00	29.03	8	180
Chocolate B	79.96	70.00	49.87	10	256

Note: n=67 for JE and n=68 for SE.