Rational Choice in Context

Shlomi Sher[1] [ID]**, Craig R. M. McKenzie[2,3],**
Johannes Müller-Trede[4] [ID]**, and Lim Leong[3]** [ID]
[1]Department of Psychological Science, Pomona College; [2]Rady School of Management, UC San Diego;
[3]Department of Psychology, UC San Diego; and [4]IESE Business School, University of Navarra

Abstract

Human decisions are context dependent in ways that violate classical norms of rational choice. However, these norms implicitly depend on idealized descriptive assumptions that are often unrealistic. We focus on one such assumption: that information is constant across contexts. Choice contexts often supply subtle cues—which may be embedded in frames, procedures, or menus—to which human decision makers can be highly sensitive. We review recent evidence that some important context effects reflect dynamically coherent belief and preference updating, in response to ecologically valid cues. This evidence paints a more nuanced picture of human rationality in natural choice environments and opens up prospects for nonpaternalistic forms of choice architecture.

Keywords

context effects, framing effects, information leakage, nudges, rationality

Psychological research on decision making bridges two levels of analysis. At the normative level, the focus is on how ideal rational actors should make decisions. At the descriptive level, the focus is on how real human actors do make decisions. Over the past 50 years, psychologists have compiled a long catalogue of descriptive violations of compelling normative principles.

In traditional rational-actor models, coherence norms occupy a central place. These norms do not dictate the content of an agent's preferences (e.g., chocolate is better than vanilla). Rather, they require that preferences expressed in different contexts be consistent (e.g., whether chocolate is ranked above vanilla should not be affected by the other flavors on the menu). The following coherence norms are generally held to be requirements of rationality:

- **Description invariance**: Logically equivalent descriptions of a choice problem should lead to identical decisions.
- **Procedure invariance**: Different methods of eliciting a person's preferences should yield the same ordering of options.
- **Transitivity**: If a is preferred to b, and b is preferred to c, then a must be preferred to c.
- **Independence of irrelevant alternatives**: The relative ranking of any two options in a choice set should not be affected by the inclusion of other options in the set.

Violations of all four principles have been demonstrated in the psychological literature (e.g., Hsee et al., 2004; Levin et al., 1998). These findings are widely thought to deliver a bleak verdict on human rationality. "Because the assumptions of description invariance and procedure invariance are normatively unassailable but descriptively inadequate," Tversky (1996, p. 195) argued, "it is not possible to reconcile normative and descriptive accounts of individual choice."

On closer inspection, though, these coherence norms depend on two subtle assumptions, which are often left unstated and therefore untested. First, information must be constant across all choice contexts. Otherwise, different preferences in different contexts may simply reflect the different inferences those contexts trigger. Second, the normative ranking of alternatives must be *complete* (i.e., well-defined everywhere), so that there is always a well-defined optimal response for the decision maker (DM).[1] Otherwise, different responses in different contexts need not imply that any response is suboptimal (Sher et al., 2022).

In this article, we take a critical look at the first assumption—of constant information. We review research supporting three main conclusions: First, apparently incidental features of the choice context

Corresponding Author:
Shlomi Sher, Department of Psychological Science, Pomona College
Email: Shlomi.Sher@pomona.edu

often "leak" relevant information. Second, people are sensitive to these subtle contextual cues. Third, inferences drawn from contextual cues can explain some well-known violations of classical coherence norms. In light of these findings, we discuss alternative agendas for choice architecture (i.e., how choices should be presented to DMs).

Learning in the Choice Phase

Typical decision-making experiments are conceptually divided into two phases. The first is the *learning phase*: The participant reads a (usually brief) background blurb, which provides information about a (sometimes hypothetical) choice situation. The second phase, which may immediately follow the first, is the *choice phase*: The participant encounters one or more options and expresses preferences or evaluations. In the choice phase, the researcher may manipulate contextual variables, such as the menu (which options are presented), the frame (how options are described), or the procedure (how preferences are elicited). The researcher then assesses the consistency of preferences across different conditions, employing norms such as those listed above. Inconsistencies are read as indicators of irrationality.

This style of analysis depends on the implicit assumption that no learning occurs during the choice phase. This assumption is often questionable. Experiments frequently involve unfamiliar or invented attributes about which the participant has little, if any, prior knowledge. The menu of sampled options may then furnish the participant's best, or only, information about the likely real-world distribution of attributes. Furthermore, as we explain below, frames and procedures may also supply cues about the likely state of the world. Participants who are sensitive to these contextual cues may form different beliefs in different contexts—and these changing beliefs may, in turn, lead to different preferences. If rational belief is not invariant to context, neither is rational choice.[2]

To understand the rationality and psychology of context effects, then, one must study the information structure of the choice environment. The following sections review two complementary research approaches we use in our own work—studies of *ecological validity* and tests of *dynamic coherence*. In the first approach, we focus on the external world, examining the empirical relationship between contextual cues and choice-relevant features of the environment. In the second approach, we attempt to peer into the DM's internal world, asking how beliefs change in response to contextual cues and how preferences change in coordination with evolving beliefs.

Ecological Cues

A general problem arises in many cognitive domains: The mind must construct a model of the distal environment on the basis of ambiguous proximal input. In vision, for example, a large rod far away and a small rod nearby project the same proximal pattern onto the retina. To figure out what the distal world is probably like, the visual system must rely on subtle cues in the input. Converging lines in the retinal array, for instance, indicate likely distance, and hence may be used to estimate the rod's likely size. In this way, vision science requires an understanding of *ecological validity* (in Brunswik's, 1956, sense)—how proximal cues are statistically correlated with distal features of natural environments.

The problem of ecological validity is just as important in decision making as in vision. On the basis of a fragmentary description (the proximal input), the DM must construct an internal model of a largely unknown choice situation (the distal world). Much as an optimal visual system does, an optimal decision-making system would accomplish this by drawing on available contextual cues that, in typical environments, tend to be correlated with relevant distal facts (McKenzie et al., 2018). Table 1 provides a summary of proximal cues in standard choice contexts and the distal facts that normally correlate with them. As explained next, inferences from these proximal cues can generate rational violations of coherence norms (also summarized in Table 1).

For example, consider how a speaker chooses to frame the options. In a framing experiment, the participant may be randomly assigned to see one of two attribute frames. For example, a medical treatment may be described as having an "80% survival rate" or a "20% mortality rate." The standard finding in such experiments is a *valence-consistent shift*: Positive frames trigger more favorable evaluations than logically equivalent negative frames do (Levin et al., 1998). But whereas researchers may assign frames randomly, frame selection in natural discourse is far from random. First, speakers tend to frame descriptions in terms of attributes that are higher than the typical value. A medical treatment is more likely to be framed in terms of its mortality rate when it leads to more deaths than other treatments do (McKenzie & Nelson, 2003). Second, speakers are more likely to select positive frames when describing objects toward which they have broadly positive attitudes than when describing objects toward which they have negative attitudes (Sher & McKenzie, 2006). In the natural ecology of human communication, then, a speaker's frame selection leaks relevant information (Sher & McKenzie, 2006, 2008): A positive frame is a cue suggesting that the level of the positive attribute is unusually

Table 1. How Distal Inferences From Proximal Cues Can Generate Rational Violations of Coherence Norms

Proximal cue	Distal fact	Potential effect of inferences	Norm violation	Reference
Attribute frame	Reference point and/ or speaker's attitude	Valence-consistent shift	Description invariance	Sher & McKenzie (2006)
Default option	Choice architect's attitude	Default effect	Procedure invariance	McKenzie et al. (2006)
Choice menu	Market distribution[a]	Joint-separate reversal	Procedure invariance	Sher & McKenzie (2014)
Choice menu	Market distribution[a]	Failure of the triangle inequality[b]	Transitivity	Müller-Trede et al. (2015)
Choice menu	Market distribution[a]	Attraction and compromise effects	Independence of irrelevant alternatives	Prelec et al. (1997)

[a]"Market distribution" refers to the naturally occurring distribution of choice-relevant attributes. Different kinds of inferences about the market distribution, based on the composition of the choice menu, can lead to different effects (indicated in the last three rows). [b]The triangle inequality is a mathematical condition relating choice probabilities that holds provided that preferences (which may vary randomly across time, independently of context) are always transitive. The rational model of triangle-inequality failures assumes imperfect memory for past contexts (see Müller-Trede et al., 2015, for details).

high and the object is well regarded. Accordingly, ideally rational listeners, attuned to the statistics of human communication, would exhibit a valence-consistent shift.

The procedure is also a potential source of ecologically valid cues. Some procedures include a default option that will take effect unless the DM elects to override it. Options are more often selected when they are designated as the default, even when the effort required to reject the default is minimized (Johnson & Goldstein, 2003). Among several causes of default effects, information leakage is likely to be one important contributor. This is because choice architects (i.e., the people who design the presentation of choices) set defaults nonrandomly. Those with personal attitudes favoring organ donation, for example, are more likely to designate a donation default than are those who do not believe that people should donate organs (McKenzie et al., 2006). And contrary to one initial report (Zlatev et al., 2017), when participants are experimentally assigned a persuasive goal (e.g., to encourage a DM to choose one of two available jobs), they strategically select defaults to influence behavior in desired ways (Jung et al., 2018; McKenzie et al., 2021). Thus, as can frames, defaults can leak implicit recommendations, cuing the DM to the likely attitude of the choice architect.

As we noted earlier, the menu of options provides a further relevant cue. When DMs have limited prior knowledge of the distribution of product attributes, the options sampled in the menu may lead them to update their beliefs about the market (Kamenica, 2008). For example, a DM who is unfamiliar with total harmonic distortion as an attribute of sound systems (cf. Hsee,

1996) may use the mean value in the choice menu to estimate the likely mean value in the market. Such sample-based inferences may in turn affect the DM's preferences, as we discuss next.

Dynamic Coherence

When contexts are informative, rational choices will not satisfy classical coherence norms. Instead, a richer normative framework, illustrated in Figure 1, is required. In this framework, learning and decision making are intimately intertwined. Exposed to a cue in the choice phase, DMs update their models of the world and then revise their preferences in light of these new beliefs. This two-stage process can generate rational violations of description invariance, procedure invariance, independence of irrelevant alternatives, and even (as Müller-Trede et al., 2015, showed) transitivity.

The enriched normative framework supplies a more nuanced empirical test of rationality. Rather than screening for static invariance, we test for dynamic coherence: Do the observed effects of context on preference coincide with the combined effects of (a) context on belief and (b) belief on preference?

There is ample evidence that people draw inferences from contextual cues, including defaults (McKenzie et al., 2006) and frames (Sher & McKenzie, 2006). An empirical test of dynamic coherence goes further, examining how these inferences are coordinated with changes in preference. To do so, we measure belief and preference updating in tandem in a yoked *modeler-recipient design*: Modeler participants are exposed to one of several choice contexts and then report their
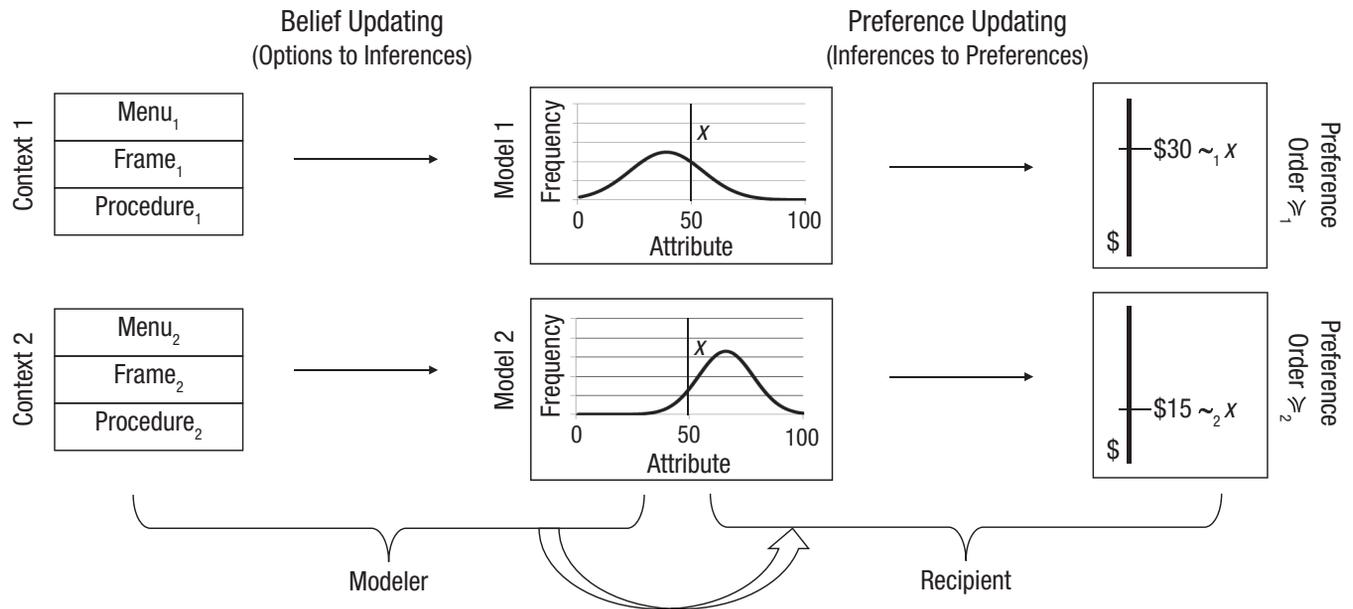
Fig. 1. Dynamically coherent belief and preference updating. Two different experimental contexts (frames, procedures, and/or menus), shown at left, result in different inferences about the distal world. The middle panel illustrates models of the distal environment (here, the market distribution of a choice-relevant attribute) inferred from different proximal cues. In this example, Context 2 leads the decision maker (DM) to infer an upward-shifted attribute distribution. For each inferred model, product x, with an attribute value of 50, is shown. As shown in the right panel, these different inferred models result in different preference orders, symbolized as $\succcurlyeq_1$ and $\succcurlyeq_2$ in Contexts 1 and 2, respectively. Here, the upward-shifted model inferred from Context 2 indicates that x's attribute value of 50 is relatively low. This results in a lower monetary equivalence point for x in preference order $\succcurlyeq_2$: Whereas the DM exposed to Context 1 (who believes that x is above average on the attribute) is indifferent between x and \$30 (i.e., $x \sim_1 \$30$), the cues in Context 2 result in a lower indifference point ($x \sim_2 \$15$). Dynamic coherence can be empirically tested in a *modeler-recipient design*, in which the context-based inferences of modeler participants (middle panel) are provided as background information to yoked recipient participants, who express their preferences in an otherwise constant context (right panel).

beliefs about a relevant aspect of the choice environment. Each modeler's reported belief is then embedded in background information provided to a yoked recipient participant, who makes a decision in an otherwise constant context.

We (Sher & McKenzie, 2014) employed a modeler-recipient design to test the dynamic coherence of *joint-separate reversals* (JSRs). A typical JSR involves two options, one of which is better than the other on an unfamiliar attribute. For example, in one problem (Hsee, 1996), two applicants for a position involving a special programming language differ in programming experience in that language (an unfamiliar attribute) and GPA (a familiar attribute). Candidate A has written 10 programs, whereas Candidate B (who has a lower GPA) has written 70 programs. Some participants evaluate the two options jointly, assigning salaries to both, but others evaluate a single option in isolation. The general finding is that the option that is superior on the unfamiliar attribute (Candidate B) receives the higher evaluation when the options are considered jointly, but the lower evaluation when the options are considered separately. These effects are generally regarded as counternormative violations of procedure invariance (Hsee et al., 2004).

Results in our experiment (Sher & McKenzie, 2014) demonstrated the standard JSR (Fig. 2, left panel). But as Figure 2 shows, the effect turns out to be dynamically coherent. Modelers presented with the two applicants, those presented with only Candidate A, and those presented with only Candidate B drew markedly different inferences about the distribution of typical programming experience (Fig. 2, middle panel). Critically, when modelers' inferred distributions were provided as background information to yoked recipients (who always evaluated a single candidate), recipients' evaluations reproduced the JSR pattern (Fig. 2, right panel). Thus, inferences drawn from option samples suffice to account for the JSR.

Using a related paradigm, Prelec et al. (1997) showed that sample-based inferences can at least partly explain well-known context effects (*attraction* and *compromise* effects) that violate independence of irrelevant alternatives. And in a modeler-recipient study, we (Leong et al., 2017) established the dynamic coherence of a typical attribute framing effect. Framing a basketball player's performance in terms of "shots made" (rather than "shots missed") led modelers to infer a lower average shooting percentage in the general population of players. These inferred averages were provided as background statistics
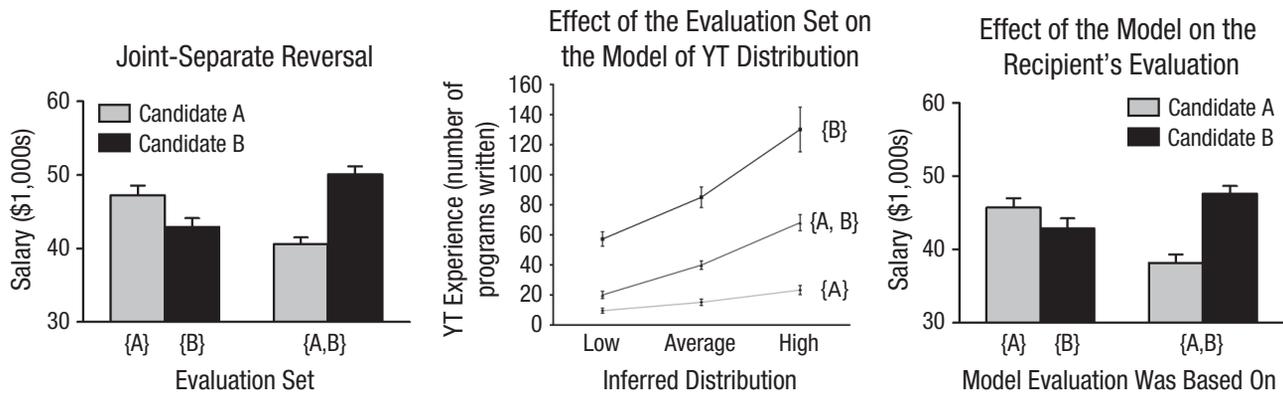
Joint-Separate Reversal

Effect of the Evaluation Set on the Model of YT Distribution

Effect of the Model on the Recipient's Evaluation

Fig. 2. Dynamic coherence of a typical joint-separate reversal. The problem (adapted from Hsee, 1996) involved job applicants who had differing experience with an unfamiliar programming language (YT). Candidate A had less experience with YT (10 programs) than Candidate B did (70 programs) but had a higher grade point average. Modelers evaluated one candidate in isolation—set {A} or set {B}—or both candidates jointly—set {A, B}—and then estimated the distribution of YT experience. The graph on the left shows the mean salary assigned to each candidate in each condition and demonstrates a standard joint-separate reversal (JSR). The graph in the middle shows the modelers' mean beliefs about the distribution of low, average, and high values of YT experience in the population of job candidates. Modelers drew very different inferences from different option sets. The graph on the right shows the mean salary assigned to each candidate by yoked recipients who were provided modelers' estimates as background information and who always evaluated a single candidate in isolation. The inferences from the option sets sufficed to reproduce the JSR pattern. Error bars represent +1 or ±1 standard error of the mean. From "Options as Information: Rational Reversals of Evaluation and Preference," by S. Sher and C. R. M. McKenzie, 2014, *Journal of Experimental Psychology: General, 143*(3), p. 1132 (https://doi.org/10.1037/a0035128). Copyright 2013 by the American Psychological Association. Reproduced and adapted with permission.

to yoked recipients, who evaluated a neutrally framed player. Frame-based inferences reproduced the usual valence-consistent shift: Recipients rated the player more favorably when the background information they were given matched the beliefs of a modeler who was exposed to the positive frame, compared with when the background information they received matched the beliefs of a modeler who was exposed to the negative frame.

Contextual Cues in Choice Architecture

Human decisions, as we have discussed, are heavily context dependent. This raises a question for policymakers and other choice architects: When eliciting potentially context-dependent preferences, what context should the architect employ?

Thaler and Sunstein (2008) proposed an influential answer to this question, grounded in a gloomy view of human rationality: Because behavior is bound to be arbitrarily pushed around by irrelevant contextual factors, the architect might as well engineer these factors to "nudge" the DM toward the best choice. This approach is known as *libertarian paternalism*. It preserves freedom of choice, but presumes that the architect often knows, better than the DM does, what is in the DM's own interests (Sugden, 2017).

The research reviewed here paints a less gloomy picture of context-dependent decision making: DMs are not just passive playthings of blind contextual forces.

They are architects in their own right, actively constructing a model of the environment in response to contextual cues, and constructing their preferences to sensibly align with that model. This view opens up an alternative role for the choice architect—not as nudging parent, but as cooperative communicator, crafting contexts that effectively convey valid and useful information to DMs. This approach to choice architecture requires an understanding of the signals that cues convey (information leakage) and the inferences they trigger (information absorption). (For a review of information leakage and absorption in public-policy contexts, see Krijnen et al., 2017.) The architect-as-communicator may also draw on insights from opinion research, through which survey designers have learned that the answers they receive from respondents reflect, in part, the subtle cues their questions transmit to respondents (Schwarz, 1999).

Cooperative communication also requires an understanding of the dynamics of attention. As do other forms of constructive cognition, belief and preference updating must conform to general capacity limits of attention and working memory (Marois & Ivanoff, 2005). Thus, along with the inference-based context effects reviewed here (in which contexts provide distinct cues), there are attention-based context effects (in which contexts differentially call attention to a given cue).[3] To effectively communicate to the DM, the architect must calibrate salience (how much attention does a cue attract?) to relevance (how much attention does a cue deserve?).
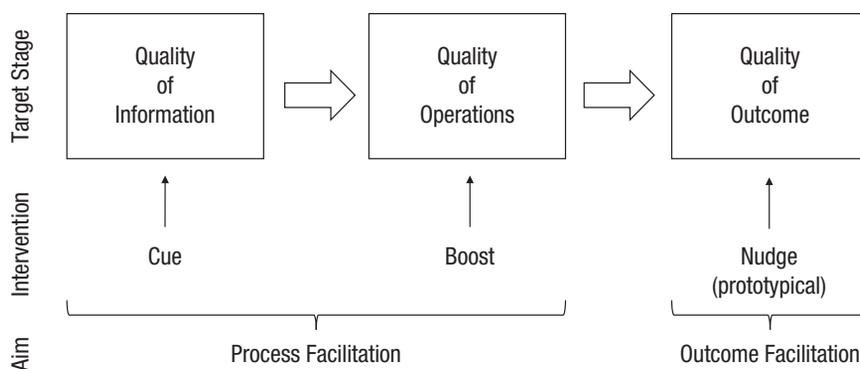
Fig. 3. Two aims and three intervention targets in choice architecture. A choice architect may aim to improve the choice outcome or the decision process leading to it. A libertarian-paternalist choice architect aims to facilitate outcomes, "nudging" decision makers to make the choice that (in the choice architect's view) is best for them. Although such an architect may intervene at any of the target stages, in the prototypical nudge, the architect attempts to improve the decision outcome without improving the decision process. In process facilitation, by contrast, the choice architect seeks to aid the decision process, without prejudging the outcome. The architect may try to enhance computational operations (e.g., through "boosts") and/or the informational inputs to those operations (e.g., through the cues discussed in the text). From "Constructed Preferences, Rationality, and Choice Architecture," by C. R. M. McKenzie, S. Sher, L. M. Leong, and J. Müller-Trede, 2018, *Review of Behavioral Economics*, *5*(3–4), p. 355 (https://doi.org/10.1561/105.00000091). Copyright 2018 by the authors. The original article is available under the Creative Commons CC-BY license.

These two views of context-dependent preference—the gloomy view of blind forces and the brighter view of perceptive inferences—present a striking contrast. Yet in the final analysis, they are not mutually exclusive. Human cognition is a multifaceted, multisystem affair, and context effects can arise in myriad ways. Some effects, such as the dynamically coherent updating depicted in Figure 1, may reflect forms of rational preference construction (for further examples, see Sher et al., 2022). But other context effects may defy a rational analysis. This checkered picture of human rationality confronts the architect with a choice: The architect can engineer a range of contextual factors; some leverage, whereas others bypass, the DM's potential for rationality. Which factors should the architect manipulate, and in what ways?

The answer depends on the architect's aims. Figure 3 distinguishes two broad objectives for choice architecture: outcome facilitation and process facilitation (McKenzie et al., 2018). The former is the guiding star of libertarian paternalists, who seek to nudge DMs toward the best decision outcome. Alternatively, the architect may strive to aid the decision process, without making paternalistic assumptions about which outcome is best for the DM. An architect may aid the decision process in either of two ways: by enhancing the relevance, salience, and ecological validity of the informational inputs (as outlined above) or by trying to improve the computational operations that are applied to those inputs (as in the "boosts" proposed by Hertwig & Grüne-Yanoff, 2017).

Though outcome facilitation and process facilitation are related goals, they should not be confused. Outcome-oriented architects may work to improve the decision process if a better process yields the best outcome (Sunstein's, 2018, "educative" nudges), but they are equally willing to reroute cognitive biases so they happen to point in the right direction (the prototypical nudge). At times, they may even contrive frames that, although not literally false, serve to subtly mislead the DM (e.g., to overestimate the frequency of rare desirable behaviors; Demarque et al., 2015). If the DM makes the right choice for the wrong reasons, the outcome-oriented architect is satisfied. The process-oriented architect is not.

What, then, are the proper ends of choice architecture? In light of the great diversity of choices and architects, no blanket prescription is possible (cf. Schmidt & Engelen, 2020). Yet we believe that in many settings, process facilitation, despite its more modest scope, has much to recommend it. By capitalizing on DMs' potential for rationality, it respects their dignity. As a consequence, it may also be better poised to preserve their trust (Arad & Rubinstein, 2018; McKenzie et al., 2018).

Ideals and Idealizations

Psychologists draw a sharp line between normative ideals and descriptive facts. But classical statements of normative ideals (invariance, independence, and transitivity) implicitly depend on subtle descriptive assumptions about the DM and the environment (completeness

of preference and constancy of information). These implicit descriptive idealizations are often dubious. As a result, behavioral violations of coherence norms are often ambiguous. They may represent genuine failures of rationality or, instead, mere failures of the normative theory's idealized background assumptions (see also Einhorn & Hogarth, 1981).

Because these presuppositions are often overlooked, researchers are quick to classify context effects as failures of rationality. The result tends to be a picture of the human DM as a hapless puppet of irrational forces—and correspondingly, of the choice architect as paternalistic puppeteer, harnessing irrational means to utilitarian ends. Yet when these idealizations are critically examined, context dependence emerges in a new light. Some (though not all) norm violations reflect learning in the choice phase—that is, dynamically coherent belief and preference updating in response to ecologically valid cues. This perspective opens up non-paternalistic approaches to choice architecture, along with a richer view of rational choice in context.

Recommended Reading

Krijnen, J. M., Tannenbaum, D., & Fox, C. R. (2017). (See References). A review of information leakage and absorption in choice architecture, highlighting the important point that decision makers are not just receivers, but may also be strategic senders, of signals.

McKenzie, C. R. M., Sher, S., Leong, L M.., & Müller-Trede, J. (2018). (See References). A discussion of cues in typical choice contexts and their implications for choice architecture.

Sher, S., & McKenzie, C. R. M. (2011). Levels of information: A framing hierarchy. In G. Keren (Ed.), *Perspectives on framing* (pp. 35–64). Psychology Press/Taylor & Francis Group. An analysis of multiple ways in which frames can be equivalent (e.g., logical equivalence, formal economic equivalence, information equivalence) and of circumstances in which these different kinds of equivalence come apart.

Sher, S., & McKenzie, C. R. M. (2014). (See References). A report on an application of a modeler-recipient design to test the dynamic coherence of joint-separate reversals.

Transparency

Action Editor: Piotr Winkielman
Editor: Robert L. Goldstone
Declaration of Conflicting Interests
   The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.
Funding

ORCID iDs

Shlomi Sher (iD) https://orcid.org/0000-0002-4978-4534
Johannes Müller-Trede (iD) https://orcid.org/0000-0002-3669-3078
Lim Leong (iD) https://orcid.org/0000-0002-9098-9718

Notes

1. Formally, the completeness axiom states that, for any pair *a*, *b* of options, either *a* is definitely preferred to *b*, *b* is definitely preferred to *a*, or the DM is precisely indifferent between *a* and *b*. Despite its mathematical convenience, a number of economists and philosophers have argued that this axiom is not a plausible requirement of rationality (e.g., Mandler, 2001).
2. Our analysis focuses on classic experimental paradigms in which choice options are explicitly described to participants and researchers assume that no learning occurs during the choice phase. These experiments contrast with recent studies of *decision from experience*, in which participants learn about options through making repeated choices and observing outcomes. In these studies (reviewed by Lejarraga & Hertwig, 2021), choice and learning are deliberately conjoined.
3. In fact, a single manipulation can have effects via both pathways, as DMs may draw inferences from the fact that the architect chose to make a particular contextual item salient.

References

Arad, A., & Rubinstein, A. (2018). The people's perspective on libertarian-paternalistic policies. *The Journal of Law and Economics*, *61*(2), 311–333. https://doi.org/10.1086/698608

Brunswik, E. (1956). *Perception and the representative design of psychology experiments* (2nd ed.). University of California Press.

Demarque, C., Charalambides, L., Hilton, D. J., & Waroquier, L. (2015). Nudging sustainable consumption: The use of descriptive norms to promote a minority behavior in a realistic online shopping environment. *Journal of Environmental Psychology*, *43*, 166–174. https://doi.org/10.1016/j.jenvp.2015.06.008

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgement and choice. *Annual Review of Psychology*, *32*, 53–88. https://doi.org/10.1146/annurev.ps.32.020181.000413

Hertwig, R., & Grüne-Yanoff, T. (2017). Nudging and boosting: Steering or empowering good decisions. *Perspectives on Psychological Science*, *12*(6), 973–986. https://doi.org/10.1177/1745691617702496

Hsee, C. K. (1996). The evaluability hypothesis: An explanation for preference reversals between joint and separate evaluations of alternatives. *Organizational Behavior and Human Decision Processes*, *67*(3), 247–257. https://doi.org/10.1006/obhd.1996.0077

Hsee, C. K., Zhang, J., & Chen, J. (2004). Internal and substantive inconsistencies in decision making. In D. J. Koehler & N. Harvey (Eds.), *Blackwell handbook of judgment and decision making* (pp. 360–378). Blackwell.

Johnson, E. J., & Goldstein, D. (2003). Do defaults save lives? *Science*, *302*(5649), 1338–1339. https://doi.org/10.1126/science.1091721

Jung, M. H., Sun, C., & Nelson, L. D. (2018). People can recognize, learn, and apply default effects in social influence. *Proceedings of the National Academy of Sciences, USA*, *115*(35), E8105–E8106. https://doi.org/10.1073/pnas.1810986115

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review*, *98*(5), 2127–2149. https://doi.org/10.1257/aer.98.5.2127

Krijnen, J. M. T., Tannenbaum, D., & Fox, C. R. (2017). Choice architecture 2.0: Behavioral policy as an implicit social interaction. *Behavioral Science & Policy*, *3*(2), 1–18. https://doi.org/10.1353/bsp.2017.0010

Lejarraga, T., & Hertwig, R. (2021). How experimental methods shaped views on human competence and rationality. *Psychological Bulletin*, *147*(6), 535–564. https://doi.org/10.1037/bul0000324

Leong, L. M., McKenzie, C. R. M., Sher, S., & Müller-Trede, J. (2017). The role of inference in attribute framing effects. *Journal of Behavioral Decision Making*, *30*(5), 1147–1156. https://doi.org/10.1002/bdm.2030

Levin, I. P., Schneider, S. L., & Gaeth, G. J. (1998). All frames are not created equal: A typology and critical analysis of framing effects. *Organizational Behavior and Human Decision Processes*, *76*(2), 149–188. https://doi.org/10.1006/obhd.1998.2804

Mandler, M. (2001). A difficult choice in preference theory: Rationality implies completeness or transitivity but not both. In E. Milgram (Ed.), *Varieties of practical reasoning* (pp. 373–402). MIT Press.

Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, *9*(6), 296–305. https://doi.org/10.1016/j.tics.2005.04.010

McKenzie, C. R. M., Leong, L. M., & Sher, S. (2021). Default sensitivity in attempts at social influence. *Psychonomic Bulletin & Review*, *28*(2), 695–702. https://doi.org/10.3758/s13423-020-01834-4

McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science*, *17*(5), 414–420. https://doi.org/10.1111/j.1467-9280.2006.01721.x

McKenzie, C. R. M., & Nelson, J. D. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin & Review*, *10*(3), 596–602. https://doi.org/10.3758/BF03196520

McKenzie, C. R. M., Sher, S., Leong, L. M., & Müller-Trede, J. (2018). Constructed preferences, rationality, and choice architecture. *Review of Behavioral Economics*, *5*(3–4), 337–360. https://doi.org/10.1561/105.00000091

Müller-Trede, J., Sher, S., & McKenzie, C. R. M. (2015). Transitivity in context: A rational analysis of intransitive choice and context-sensitive preference. *Decision*, *2*(4), 280–305. https://doi.org/10.1037/dec0000037

Prelec, D., Wernerfelt, B., & Zettelmeyer, F. (1997). The role of inference in context effects: Inferring what you want from what is available. *Journal of Consumer Research*, *24*(1), 118–125. https://doi.org/10.1086/209498

Schmidt, A. T., & Engelen, B. (2020). The ethics of nudging: An overview. *Philosophy Compass*, *15*(4), Article e12658. https://doi.org/10.1111/phc3.12658

Schwarz, N. (1999). Self-reports: How the questions shape the answers. *American Psychologist*, *54*(2), 93–105. https://doi.org/10.1037/0003-066X.54.2.93

Sher, S., & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, *101*(3), 467–494. https://doi.org/10.1016/j.cognition.2005.11.001

Sher, S., & McKenzie, C. R. M. (2008). Framing effects and rationality. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 79–96). Oxford University Press.

Sher, S., & McKenzie, C. R. M. (2014). Options as information: Rational reversals of evaluation and preference. *Journal of Experimental Psychology: General*, *143*(3), 1127–1143. https://doi.org/10.1037/a0035128

Sher, S., Müller-Trede, J., & McKenzie, C. R. M. (2022). *Choices without preferences: Principles of rational arbitrariness* [Manuscript in preparation]. Department of Psychological Science, Pomona College.

Sugden, R. (2017). Do people really want to be nudged towards healthy lifestyles? *International Review of Economics*, *64*(2), 113–123. https://doi.org/10.1007/s12232-016-0264-1

Sunstein, C. R. (2018). Misconceptions about nudges. *Journal of Behavioral Economics for Policy*, *2*(1), 61–67. https://ideas.repec.org/a/beh/jbepv1/v2y2018i1p61-67.html

Thaler, R. H., & Sunstein, C. R. (2008). *Nudge: Improving decisions about health, wealth, and happiness*. Yale University Press.

Tversky, A. (1996). Rational theory and constructive choice. In K. J. Arrow, E. Colombatto, M. Perlman, & C. Schmidt (Eds.), *The rational foundations of economic behavior* (pp. 185–197). Macmillan.

Zlatev, J. J., Daniels, D. P., Kim, H., & Neale, M. A. (2017). Default neglect in attempts at social influence. *Proceedings of the National Academy of Sciences, USA*, *114*(52), 13643–13648. https://doi.org/10.1073/pnas.1712757114