Current Directions in Psychological Science, in press

Rational Choice in Context

Shlomi Sher[1], Craig R. M. McKenzie[2, 3], Johannes Müller-Trede[4], and Lim Leong[3]

[1] Department of Psychological Science, Pomona College

[2] Rady School of Management and [3] Department of Psychology, UC San Diego

[4] IESE Business School, University of Navarra

Word count: 3278 (abstract, main text, and endnotes)

Author Note

This research was supported by a James S. McDonnell Foundation Scholar Award to S.S. and by a National Science Foundation grant (SES-2049935) to C.R.M.M.

Correspondence concerning this article should be addressed to Shlomi Sher, Department of Psychological Science, Pomona College, 647 N. College Way, Claremont, CA 91711. Email: Shlomi.Sher@pomona.edu.

Abstract

Human decisions are context-dependent in ways that violate classical norms of rational choice. However, these norms implicitly depend on idealized descriptive assumptions which 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 non-paternalistic forms of choice architecture.

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

Rational Choice in Context

Psychological research on decision making bridges two levels of analysis. At the normative level, we ask how ideal rational actors should make decisions. At the descriptive level, we ask how real human actors do make decisions. Over the last 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 shouldn't 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 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, Zhang, & Chen, 2004; Levin, Schneider, & Gaeth, 1988). 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 needn't imply that any response is suboptimal (Sher, Müller-Trede, & McKenzie, 2022).

This article takes 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 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.

Learning in the Choice Phase

Typical decision making experiments are conceptually divided in 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 like 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, we must study the information structure of the choice environment. The following sections review two complementary research approaches – studies of *ecological validity* and tests of *dynamic coherence*. The first approach looks at the external world, examining the empirical relationship between contextual cues and choice-relevant features of the environment. The second approach seeks 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, based on 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. Based on a fragmentary description (the proximal input), the DM must construct an internal model of a largely unknown choice situation (the distal world). To do so, an optimal decision making system, like an optimal visual system, would draw on available contextual cues that, in typical environments, tend to be correlated with relevant distal facts (McKenzie, Sher, Leong, & Müller-Trede, 2018). Table 1 provides a summary of proximal cues in standard choice contexts and the distal facts that normally correlate with them. As explained below, inferences from these proximal cues can generate rational violations of coherence norms.

Table 1

Proximal Cues, Distal Inferences, and Resulting Violations of Coherence Norms

Proximal cue	Distal fact	Inferences may result in:	Violation of:	Reference
Attribute frame	Reference point and/or speaker's attitude	Valence-consistent shift	Description invariance	Sher & McKenzie, 2006
Default	Policy maker's attitude	Default effect	Procedure invariance	McKenzie et al., 2006
Composition of choice menu	Market distribution of choice-relevant attribute	Joint-separate reversal	Procedure invariance	Sher & McKenzie, 2014
		Failure of triangle inequality*	Transitivity	Müller-Trede et al., 2015
		Attraction and compromise effects	Independence of irrelevant alternatives	Prelec et al., 1997

Note. The triangle inequality (*) is a formal condition that holds whenever preferences (which may randomly vary across time) are always transitive. The rational model of triangle inequality failures assumes imperfect memory for past contexts.

For example, consider the speaker's choice of frame. In a framing experiment, the participant may be randomly assigned to see one of two attribute frames – e.g., 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 (Levin et al., 1988). But while 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 (McKenzie & Nelson, 2003). Second, speakers are more likely to select positive frames when describing objects to which they have broadly positive 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 positive attribute is unusually abundant 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, like the frame, is also a potential source of ecologically valid cues. Some procedures include a "default" option, which 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 for default rejection 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" set defaults non-randomly. Those with personal attitudes favoring organ donation, for example, are more likely to designate a donation default (McKenzie, Liersch, & Finkelstein, 2006). And contrary to one initial report (Zlatev, Daniels, Kim, & Neale, 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, Sun, & Nelson, 2018; McKenzie, Leong, & Sher, 2021). Like frames, defaults can thus leak "implicit recommendations," cueing 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, if the DM is unfamiliar with the Total Harmonic Distortion attribute for sound systems (cf. Hsee, 1996), they may use the mean value in the choice menu to estimate the likely mean value in the market. These sample-based inferences may in turn affect the DM's preferences, as we discuss next.

Dynamic Coherence

When contexts are informative, rational choices won't 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, the DM updates their model of the world, and then revises 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, Sher, and McKenzie (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 (1) context on belief, and (2) belief on preference?

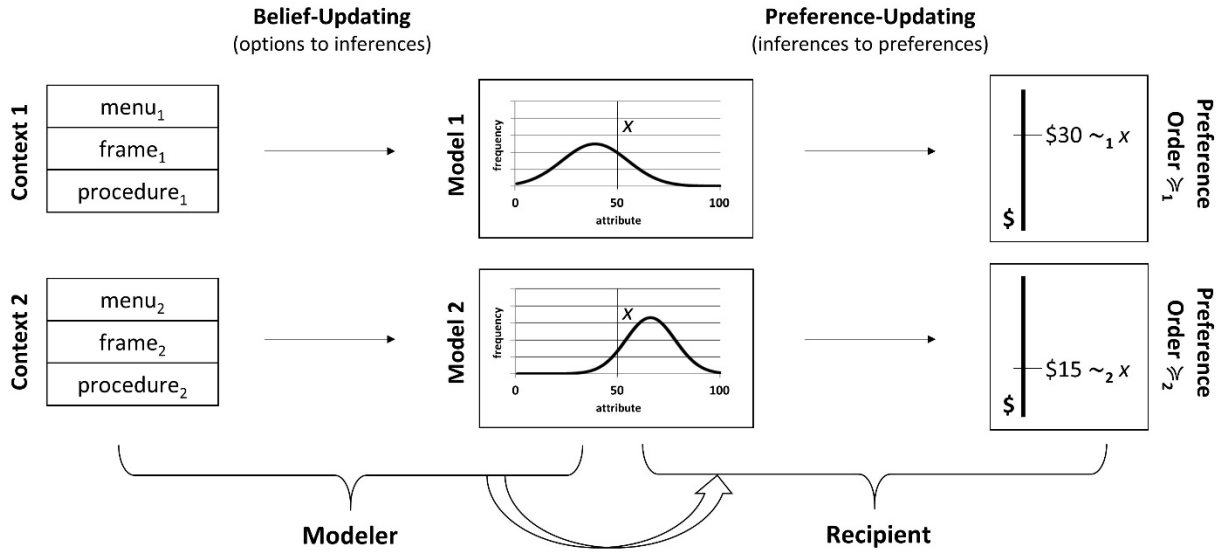

Figure 1. Dynamically coherent belief- and preference-updating. Two different experimental contexts (frames, procedures, and/or menus), shown at left, result in different inferences about the distal world. The middle panel illustrates "models" of the distal environment (here, the market distribution of a choice-relevant attribute) inferred from different proximal cues. In this example, Context 2 leads the DM to infer an upward-shifted attribute distribution. Relative to both inferred models, a 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, $\succcurlyeq_1$ and $\succcurlyeq_2$. 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 $\succcurlyeq_2$: While 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).

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 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.

Sher and McKenzie (2014) employed a modeler-recipient design to test the dynamic coherence of "joint-separate reversals" (JSRs; Hsee, 1996). A typical JSR features two options, one of which is better on an unfamiliar attribute. For example, in applications for a position involving a special programming language, Candidate A may have written 10 programs, while Candidate B (who has a lower GPA) has written 70 programs. Some participants evaluate both options jointly, while others see and evaluate one option in isolation. The general finding is that the option that excels on the unfamiliar attribute (Candidate B, in this case) receives higher joint evaluations but lower separate evaluations. These effects are generally regarded as counter-normative violations of procedure invariance (Hsee et al., 2004).

But as Figure 2 shows, the effects turn out to be dynamically coherent. In the programmer problem, when modelers were exposed to different evaluation sets ({A}, {B}, or {A, B}), they exhibited a standard JSR (Figure 2A). They also drew markedly different inferences about the distribution of typical programming experience (Figure 2B). 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 (Figure 2C). Inferences from option-samples suffice to account for the JSR.

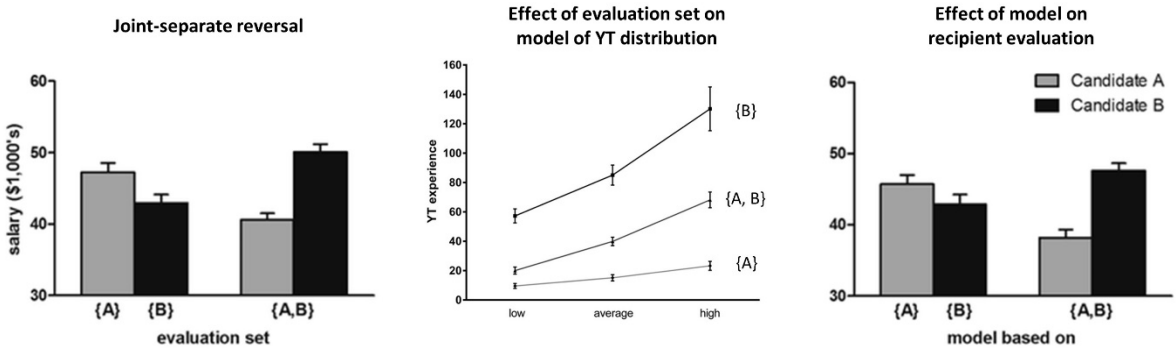

Figure 2. Dynamic coherence of a typical joint-separate reversal (JSR). In a problem adapted from Hsee (1996), involving experience with an unfamiliar programming language (YT), modelers saw and evaluated one candidate in isolation ({A} or {B}) or both candidates jointly ({A, B}), where Candidate B (70 programs) has more experience but a lower GPA than Candidate A (10 programs). *Left panel:* Modelers exhibited a standard JSR, assigning higher mean salaries to A in separate evaluation and to B in joint evaluation. *Middle panel:* Modelers also drew different inferences from different evaluation sets. Exposed to different option-samples, modelers formed different beliefs about the distribution of programming experience (low, average, and high values in the population of job candidates). *Right panel:* When modeler estimates were provided as background information to yoked recipients (who always evaluated a single candidate in isolation), the standard JSR pattern was reproduced. Recipients assigned a high salary to A when their beliefs about the YT distribution were matched to a modeler who had seen A alone, but they assigned a low salary to A when their beliefs were matched to a modeler who had seen both A and B. Adapted from Sher & McKenzie, 2014.

Using a related paradigm, Prelec, Wernerfelt, and Zettelmeyer (1997) showed that sample-based inferences can at least partly explain well-known context effects (so-called "attraction" and "compromise effects") that violate independence of irrelevant alternatives. And in a modeler-recipient study, Leong, Sher, Müller-Trede, and McKenzie (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 to yoked recipients, who evaluated a neutrally framed player.

Frame-based inferences reproduced the usual valence-consistent shift: When their background information matched the beliefs of a modeler who was exposed to the positive frame, recipients rated the player more favorably.

Contextual Cues in Choice Architecture

Human decisions, as we have seen, are heavily context-dependent. This raises a question for policymakers and other "choice architects," who present people with choices: When eliciting potentially context-dependent preferences, what context should the architect instate?

Thaler and Sunstein (2008) proposed an influential answer to this question, grounded in a gloomy view of human rationality: Since behavior is bound to be arbitrarily pushed around by irrelevant contextual factors, the architect might as well engineer these factors to "nudge" the DM towards 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 (re)constructing their preferences to sensibly align with it. 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, Tannenbaum, and Fox (2017). The architect-as-communicator may also draw on insights from opinion research, where 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. Belief- and preference-updating, like other forms of constructive cognition, must conform to general capacity limits of attention and working memory (Marois & Ivanoff, 2005). Thus, along with the inference-based context effects reviewed here (where contexts provide distinct cues), there are attention-based context effects (where 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?).

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 multi-faceted, multi-system affair, and context effects can arise in myriad ways. Some effects may reflect forms of rational preference construction, such as the dynamically coherent updating depicted in Figure 1 (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, while 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 towards 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. They may

aid the decision process in either of two ways – first, by enhancing the relevance, salience, and ecological validity of its informational inputs (as outlined above); or second, by trying to improve the computational operations that are applied to those inputs (as in the "boosts" proposed by Hertwig & Grüne-Yanoff, 2017).

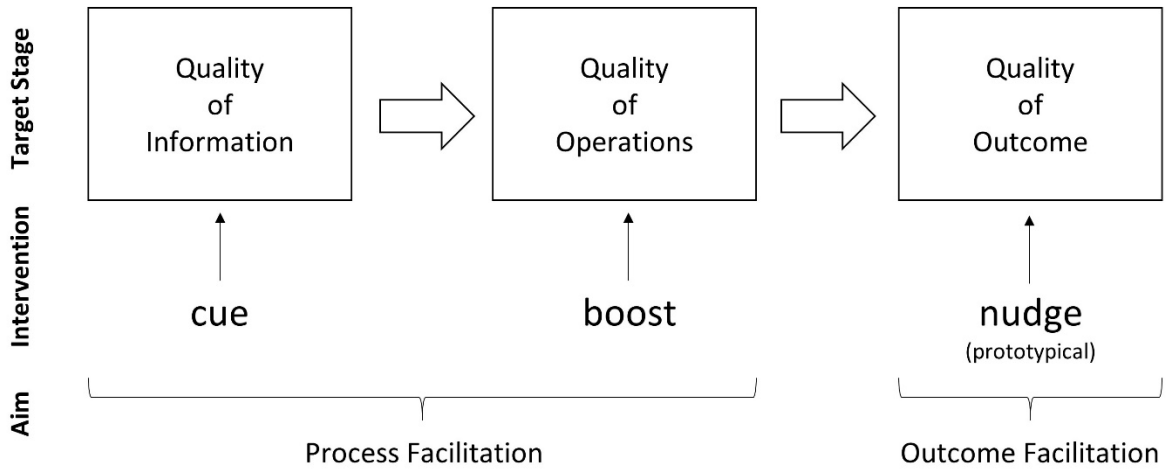

Figure 3. Two aims and three intervention targets in choice architecture. A choice architect may aim to improve the choice *outcome*, or the decision *process* leading to it. Libertarian paternalists aim to facilitate outcomes, "nudging" DMs to make the choice that (in the choice architect's view) is best for the DM. Process facilitation, by contrast, seeks to aid the decision process, without prejudging the outcome. Its interventions may try to enhance computational operations (as in "boosts") and/or the informational inputs to those operations (as in the cues discussed in the text). Reprinted from McKenzie et al. (2018).

Though process and outcome aims are related, 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 re-route cognitive biases so they happen to point in the right direction (the prototypical nudge). At times, they may even contrive frames which, while not literally false, serve to subtly mislead the DM

(e.g., to overestimate the frequency of rare desirable behaviors; Demarque, Charalambides, Hilton, & Waroquier, 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? Considering the great diversity of choices and architects, no blanket prescription is possible (cf. Schmidt & Engelen, 2020). Yet in many settings, we believe that process facilitation, despite its more modest scope, has much to recommend it. By capitalizing on the DM's 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 we critically examine these idealizations, context-dependence emerges in a new light. We find that some (though not all) norm violations reflect learning in the choice phase – i.e., 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.) 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., & Müller-Trede, J. (2018). (See references). 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 Keren, G. (Ed.), *Perspectives on framing.* Psychology Press - Taylor & Francis Group, 35-64. 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.) Application of modeler-recipient design to test the dynamic coherence of joint-separate reversals.

Endnotes

1. Formally, completeness 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 the completeness 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 repeated choice with feedback. 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*, 311-333.

Brunswik, E. (1956). *Perception and the representative design of psychology experiments* (2nd ed.). Berkeley: 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.

Einhorn, H. J., & Hogarth, R. M. (1981). Behavioral decision theory: Processes of judgement and choice. *Annual Review of Psychology*, *32*, 53-88.

Hertwig, R., & Grüne-Yanoff, T. (2017). Nudging and boosting: Steering or empowering good decisions. *Perspectives on Psychological Science*, *12*, 973-986.

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*, 247-257.

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). Oxford: Blackwell.

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*, 201810986.

Johnson E. J., & Goldstein D. G. (2003). Do defaults save lives? *Science*, *302*, 1338–39

Kamenica, E. (2008). Contextual inference in markets: On the informational content of product lines. *American Economic Review*, *98*, 2127-49.

Krijnen, J. M., Tannenbaum, D., & Fox, C. R. (2017). Choice architecture 2.0: Behavioral policy as an implicit social interaction. *Behavioral Science & Policy*, *3*, 1-18.

Lejarraga, T., & Hertwig, R. (2021). How experimental methods shaped views on human competence and rationality. *Psychological Bulletin*, *147*, 535-564.

Leong, L. M., McKenzie, C. R., Sher, S., & Müller-Trede, J. (2017). The role of inference in attribute framing effects. *Journal of Behavioral Decision Making*, *30*, 1147-1156.

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*, 149-188.

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

Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, *9*, 296-305.

McKenzie, C. R., Leong, L. M., & Sher, S. (2021). Default sensitivity in attempts at social influence. *Psychonomic Bulletin & Review*, *28*, 695-702.

McKenzie, C. R., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations implicit in policy defaults. *Psychological Science*, *17*, 414-420.

McKenzie, C. R., & Nelson, J. D. (2003). What a speaker's choice of frame reveals: Reference points, frame selection, and framing effects. *Psychonomic Bulletin & Review*, *10*, 596-602.

McKenzie, C. R. M., Sher, S., Leong, L. M., & Müller-Trede, J. (2018). Constructed preferences, rationality, and choice architecture. *Review of Behavioral Economics*, *5*, 337-360.

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*, 280-305.

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*, 118-125.

Schmidt, A. T., & Engelen, B. (2020). The ethics of nudging: An overview. *Philosophy Compass*, *15*, e12658.

Schwarz, N. (1999). Self-reports: How the questions shape the answers. *American Psychologist*, *54*, 93-105.

Sher, S. & McKenzie, C. R. M. (2006). Information leakage from logically equivalent frames. *Cognition*, *101*, 467-494.

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

Sher, S. & McKenzie, C. R. M. (2014). Options as information: Rational reversals of evaluation and preference. *Journal of Experimental Psychology: General*, *143*, 1127-1143.

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

Sugden, R. (2017). Do people really want to be nudged towards healthy lifestyles? *International Review of Economics*, *64*, 113-123.

Sunstein, C. R. (2018). Misconceptions about nudges. *Journal of Behavioral Economics for Policy*, *2*, 61-67.

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

Tversky, A. (1996). Rational theory and constructive choice. In Arrow, K. J., Colombatto, E., Perlman, M., & Schmidt, C. (Eds.), *The rational foundations of economic behavior* (pp. 185-197). London: 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*, *114*, 13643-13648.