



Architecture of Choice and Cognitive Biases in Digital Interfaces: A Review of Empirical Evidence on Impacts on Consumer Behavior

Volodymyr Rebets

Expert in Behavioral Marketing and Branding.

Abstract

Online shopping has changed everything, but we still don't fully understand what goes on in our heads when we click "buy". This article examines choice architecture - when designers intentionally build a decision environment to push you in the right direction. We look at three mechanisms that keep cropping up in research: framing (how to present information is as important as the information itself), scarcity (constraints create urgency), and social proof (when in doubt, look at what others are doing). We looked at research from 2020 to 2023 to understand how interface design leverages these cognitive biases to drive sales. The picture is complex. The choice architecture works - the interventions show moderate effectiveness with Cohen's d around 0.43 - but the ethics are muddled. Framing effects vary across cultures unpredictably.

Scarcity tactics face diminishing returns as people get used to constant "urgent piles". Social proof mechanisms risk creating echo chambers where everyone sees distorted versions of popularity. We suggest distinguishing legitimate nudges from manipulative dark patterns based on transparency and preserving autonomy, not just outcomes. Practical recommendations include rotating scarcity campaigns every three months so consumers don't get used to it, adding "why this recommendation?" to recommendation systems, limiting the depth of discounts to maintain trust in the brand.

The main contribution: an effective architecture of choice does not require harming the well-being of consumers - but only if designers accept principled constraints that the market does not naturally reward.

Keywords: Choice Architecture, Cognitive Biases, Digital Interfaces, Behavioral Economics, UX Ethics.

INTRODUCTION

Go to any physical store - you will see products arranged in such a way that it catches the eye. Seasonal goods on the ends of the racks, check-out areas jammed with impulse purchases, sale prices written in colors that evoke urgency. Retailers have spent decades learning how spatial design affects what people buy. But transfer the same transaction online - and the game changes dramatically. A website doesn't just place products in space. It can rewrite itself for each visitor separately, A/B-test a thousand options at the same time, track exactly which psychological buttons make you click "buy now".

It's the architecture of choice in the digital age, and it's become surprisingly complex. The term itself comes from behavioral economics - it means the design of decision-making environments in such a way as to push people to make specific choices (Mertens et al., 2022). Traditional marketing tries to persuade through information or price incentives. The choice architecture works differently. It uses the mental shortcuts your brain takes when you're overwhelmed, tired, or just trying to get through the day without overthinking every decision ad nauseam.

The transition to online shopping has turned what used to be an art into something closer to a science. Physical stores

have real limitations - space on the shelves costs money, it is impossible to rearrange the entire layout for each customer, tracking who bought what requires either loyalty cards or guesswork. Digital platforms have none of these limits. They test ruthlessly. That "limited time" on the red timer? Someone experimented to find the exact shade of red that creates maximum urgency without provoking skepticism. That "87% of customers recommend this product" icon? The algorithm determined that 87% sounded more convincing than 90%, although higher numbers should logically inspire more confidence. Mathur and colleagues (2021) call this a "persuasion arms race" - platforms are racing to find ever more effective ways to manipulate your decision-making without you noticing.

And you, most likely, do not notice. This is the most inconvenient part. When you buy something during a flash sale, you're more likely to tell yourself (and anyone who asks) that you made a smart choice based on price and need. You don't think "I fell into an artificial deficit designed to keep me from thinking about a decision." When you choose a product because thousands of others have bought it before, you hardly think about whether those buying patterns reflect real quality or algorithmic manipulation. The people who design these interfaces know exactly what they are doing. People who click are mostly not.

Recent studies confirm that it really works. Mertens and colleagues (2022) analyzed more than 200 studies covering 440 different experiments and found that prodding produced real behavior change. The effect isn't huge -- what statisticians call "small to medium" -- but it's consistent across everything from food choices to energy savings. More importantly, the effects vary wildly depending on exactly how you implement them and who you target. Default options (when inaction leads to a pre-selected choice) hit the hardest. Framing effects (presenting the same information in different ways) work, but show a surprising cultural variation--they are stronger in Korea than in Germany, for reasons that researchers are still debating (Cheon et al., 2021). And here's the real problem: the very same techniques that can nudge you toward healthier eating or sustainable energy consumption can be twisted into what researchers now call "dark patterns"--design tricks that make you do things against your self-interest, like signing up for services you don't need or giving up privacy you'd rather keep (Gunawan et al., 2021).

This article looks at three specific cognitive triggers that keep popping up in research: framing (how the presentation of information shapes your preferences), scarcity (how perceived constraints create urgency), and social proof (how other people's behavior influences your choices). We do not see these as abstract psychological concepts. We want to understand how they are actually implemented in real interfaces, what makes them work, and where they cross the line from useful to harmful.

The analysis proceeds from an assumption that may seem obvious but is often forgotten: understanding these mechanisms serves different audiences with different purposes. Whether you design interfaces or sell products, this study offers a manual for more effective persuasion. Whether you're a consumer, a politician, or an ethicist, it reveals manipulation tactics that deserve scrutiny and possibly regulation. We think the difference between legitimate nudging and deceptive dark patterns is not just the results, but the process. Does the design respect autonomy? Does it maintain transparency? Does it avoid exploiting vulnerabilities? These questions matter, and there are no easy answers. The point here is not to condemn the entire choice architecture (which would be both impractical and naive, since every design makes choices that affect behavior). The goal is to understand where the boundaries should be.

METHODOLOGY

We needed to find out what actually happens when websites use psychological tricks to get people to buy stuff. Not theories. Not lab studies where college kids fill out questionnaires about hypothetical purchases. Real studies tracking real shoppers spending real money.

Hit up Web of Science, Google Scholar, ACM Digital Library looking for anything from 2020-2023. Threw in search terms: "choice architecture," "nudge," "behavioral economics"

mixed with "e-commerce," "online shopping," "digital interface" plus "framing," "scarcity," "social proof." Got 387 papers initially.

Started eliminating. Paper had to show actual data from experiments or field observations, not just someone's opinion about what probably works. Had to be about digital environments - didn't care about physical store research because online works completely differently. Had to cover at least one of our three things: framing, scarcity, or social proof. Had to be peer-reviewed with methodology described clearly enough that we could judge if it was any good.

Threw out public health stuff where nobody's purchasing anything. Threw out personality studies that didn't look at actual interface design. After cutting, left with 73 studies plus 15 older papers for theoretical background.

Went through everything three times. First time: extracted every design trick mentioned, categorized by type (framing/scarcity/social proof), how it was implemented (text/visual/algorithm), what got measured. Second time: tried figuring out effect sizes and what makes interventions work better or worse. When studies gave enough numbers we compared directly. When they only gave descriptions we noted which direction things went and what problems existed. Third time: looked for ethics discussions, mapped out debates and disagreements.

Key thing: we're synthesizing patterns, not doing formal meta-analysis. Studies differ too much in setup and what they measured. Cramming numbers together when they're measuring completely different things just creates fake precision. Better to say "here's the pattern" and "here's where it gets unclear."

Limitations worth mentioning. English only, so we missed stuff published in other languages. 2020-2023 window misses older work except what we cited. Publication bias definitely skews things - nobody publishes failed experiments, which means we're probably overestimating how well this manipulation works. And academic research doesn't show us what companies actually do. Amazon and Booking.com run massive A/B tests but never publish results, so we're missing huge chunks of the real picture.

RESULTS AND DISCUSSION

Digital interfaces work by exploiting how your brain actually makes decisions, not how economics textbooks say it should. When faced with a choice, your brain doesn't painstakingly weigh all the options against some intrinsic utility function. Instead, it uses mental shortcuts - heuristics - that usually work quite well with minimal effort. These labels sometimes give rise to systematic biases, predictable ways in which you deviate from purely rational choices. Choice architecture identifies these patterns and builds environments that either reinforce or suppress specific biases depending on what outcome the designer wants.

Look at two ways to label the exact same yogurt: "90% fat-free" or "contains 10% fat." Same yogurt. Same fat content. Math checks out either way. But people pick the first one overwhelmingly, and they'll pay more for it too. Why? Because "fat-free" sounds like a benefit you're getting, while "contains fat" sounds like a problem you're accepting. Your brain doesn't process these as equivalent statements even though logically they are.

Shan's team tested this with fake meat in seven Chinese cities (2022). Took identical products, wrote two sets of descriptions. One version emphasized what you gain: "high protein, lots of nutrients." Other version emphasized what you avoid: "low cholesterol, low fat." Same fake meat either way. The "gain" framing bumped purchase rates up 23% compared to the "loss" framing. People bought more of the exact same product just because the words changed.

Grocery store sells cereal in boxes. One design per product. Print it, ship it, stack it on shelves. Every customer sees the same damn box. That's it. Online? Completely different animal. Same product gets fifty different descriptions running at once. System tracks which wording works on which type of person, then shows each visitor whichever version is statistically most likely to make them buy. You and the person next door could be looking at the exact same blender and see totally different copy because the algorithm decided you respond better to "professional-grade power" while they respond to "easy cleanup."

Kim and Choi ran tests in 2021, found weird results with discount wording. Expensive items: "Don't miss 30% off" beat "Save 30%." Cheap items: opposite. Gain framing won. Makes no logical sense - 30% off equals saving 30%, right? But brains don't work logically. Optimal manipulation apparently depends on price bracket, what you're selling, how much the buyer cares. Brick-and-mortar can't adjust for this stuff. Algorithm does it automatically per visitor.

Ethics gets murky fast. Emphasizing your product's strengths with positive language - fine, probably. Marketing did this forever. Actively concealing problems or warping truth to fool people? Not fine.

Subscription companies mastered the sleazy version. Go cancel Netflix or Spotify. Cancellation button won't say "stop getting charged fifteen bucks monthly." Will say "lose premium content" or "end ad-free listening." Same exact thing happens either way - subscription stops, money stops leaving account. But wording switches what your brain focuses on. First version makes you think about money you keep. Second makes you think about stuff disappearing. Exploits loss aversion to keep you subscribed to something you're actively attempting to cancel. Can't call that persuasion with straight face. It's overriding what you said you want.

Scarcity manipulation hooks into evolutionary wiring. Our ancestors lived in environments where limited resources

meant grab now or starve later. Food scarce? Take what's available immediately instead of wondering if better options exist tomorrow. Survival logic. Applies way less to buying sneakers on the internet but the ancient brain circuitry fires anyway.

Online shopping industrialized scarcity. Countdown clocks on every page. "Only 3 in stock!" that might be total lies. "Limited time sale!" running continuously for months. Ladeira analyzed 89 studies in 2023 - scarcity messaging does boost purchase intent, effect size $r=0.31$, moderate. Effectiveness varies wildly though based on scarcity type claimed.

Two scarcity types function differently. Supply: "only 5 left." Time: "sale ends 4 hours." Gierl and Huettl ran tests in 2010 finding supply scarcity works best on luxury goods where owning rare stuff shows status. Time scarcity works best on everyday junk where urgency matters more than exclusivity. Sites combine them: "Only 2 rooms at this price" stacks quantity constraint with price-change threat.

Problem though: constant scarcity messaging stops working. Every website yelling urgency teaches shoppers to ignore it. Fades into background. Researchers call this "persuasion knowledge" - recognizing manipulation weakens it. Zong and Guo measured brain activity in 2022. Chronic scarcity exposure quiets reward-processing areas while activating skepticism and distrust regions.

Ethical problems intensify when claims of scarcity are found to be false. Regulatory agencies have begun to sanction companies that display false scarcity -- countdown timers that reset after reaching zero, inventory indicators that show arbitrary numbers, "limited time" offers that are available forever. This goes from legitimately creating urgency to outright fraud. Even authentic scarcity raises questions when designed to prevent deliberation. Ultra-short countdown timers or flash sales that force instant decisions may boost conversion rates, but hurt consumers who don't have time to compare or reconsider.

Social proof taps into something very basic in us. When we're not sure what to do, we look around and copy other people. That's why websites are packed with reviews, star ratings, "bestseller" labels, and those little pop-ups saying "Someone in Chicago just bought this." It works because, for most of human history, paying attention to what others were doing was a survival skill. Online platforms simply scaled this instinct to an industrial level.

Whether social proof actually helps or just manipulates depends a lot on context - and researchers are still sorting that out. A study by Mrkva and colleagues in 2021 showed that social proof can sometimes be genuinely useful. When people feel lost, they use others' choices as a shortcut to expertise. Think about buying insurance: you don't fully understand deductibles, coverage options are confusing, and the language is opaque. Seeing what similar people chose

can provide real guidance, not just pressure to buy. In those cases, social proof adds information rather than noise.

There is an obvious ugly side to all of this. Platforms don't just display social signals - they decide which ones you're allowed to see. They own the entire chain. Reviews can be fake. Ratings and likes are bought in bulk. Algorithms don't care whether something is good; they care whether it makes money. As a result, genuine social proof - real people making real choices - gets pushed out of sight. What surfaces instead is a carefully manufactured consensus that looks spontaneous but isn't. Mathur and colleagues showed this back in 2021, documenting how platforms visually downplay negative reviews and amplify positive ones through layout and formatting tricks. It works for a simple reason: most users take the first information they see and move on.

Algorithmic curation makes the situation even murkier. Labels like "trending" or "customers also bought" sound informative, but they're opaque. A product might actually be popular. Or it might simply be more profitable to promote. From the user's perspective, there's no way to tell. Both scenarios look exactly the same on the screen. That lack of transparency makes social proof unreliable - it's impossible to know whether it reflects collective behavior or just an algorithm nudging people toward what benefits the platform.

Researchers try to draw a line between ethical nudges and manipulative design, but that line is blurry. Ahuja and Kumar argued in 2022 that ethical interventions respect user autonomy - the ability to choose in line with one's own values - while dark patterns undermine it. On paper, that distinction makes sense. In practice, it's messy. Autonomy isn't something you either have or don't have; it varies by person and situation. A reminder about an abandoned cart can be genuinely helpful for someone who forgot, and deeply annoying for someone who didn't. Context matters. Intent matters. But neither gives designers or regulators a clear rule to enforce.

Empirical work shows how widespread the problem is. Gunawan and colleagues found in 2021 that nearly nine out of ten e-commerce platforms rely on at least one manipulative tactic. The most common ones weren't subtle: making subscriptions hard to cancel, hiding extra fees until late in the process, or repeatedly interrupting users to push a preferred option. These tactics all follow the same logic. Choices the company wants become easy and smooth. Choices it doesn't want become slow, confusing, and exhausting. Users are technically free to refuse - but doing so costs time, attention, and mental energy. Resistance is allowed, just deliberately made painful.

Ethical choice architecture probably isn't about whether designs change behavior - of course they do, every design does. Question is whether designs play fair. Transparent nudges let you see the mechanism instead of hiding it.

Choice-preserving designs keep actual alternatives open without making you pay for choosing them. Well-being-aligned interfaces prioritize what helps users, not just what maximizes revenue. Doesn't answer every situation but at least gives criteria for judging.

Research suggests that persuasion doesn't have to cross ethical lines. In theory, it's possible to design influence without manipulation. In practice, markets pull in the opposite direction. Competitive pressure doesn't reward restraint - it rewards whatever pushes people to click faster.

Scarcity is a good example. It collapses the moment it's used too often. When a site shouts "LAST CHANCE" or "ONLY 3 LEFT" day after day, users stop reacting. The signal burns out. Over time it becomes background noise. The issue isn't scarcity itself, but repetition without variation. Rotating cues - countdowns for a period, then low-stock messages, then seasonal availability - keeps attention from going numb. Still, rotation alone isn't enough. Scarcity only works when it reflects something real. Holiday shipping cutoffs exist. Warehouses do run out of stock. Some products are genuinely seasonal. These constraints create urgency because people can check them against reality. Artificial urgency, repeated endlessly, does the opposite: it teaches users not to believe anything.

Framing works the same way - it depends on context. With expensive products, absolute numbers feel real. "\$300 off" lands harder than "15% off" when the price tag is high. With cheap items, percentages do the work. "50% off" sounds dramatic; "save \$3" doesn't. Loss framing follows a similar pattern. When money is on the line and the decision feels risky, messages like "Don't miss this" can be effective. For small, routine purchases, pressure isn't necessary. People just want to know what they're getting.

Discounting habits quietly reveal whether a brand's positioning is honest. A company that calls itself premium but runs 40-50% discounts all the time sends a clear message: full price is for suckers. Over time, customers learn to wait. The premium story falls apart. Brands that actually want to stay premium have to limit discounts - not eliminate them, but keep them restrained. Smaller, predictable discounts look deliberate. Budget brands play a different game. Deep sales fit their identity, and customers expect them. In both cases, consistency matters more than how big the discount is. When prices jump around without a pattern, people lose track of what the product is even worth.

The same logic applies to scarcity. Vague phrases like "Limited time offer" don't mean much on their own. Limited why? For how long? When urgency is tied to operations - shipping schedules, restocking cycles, production limits - it feels credible. "Order by 5 p.m. for next-day shipping" makes sense. "Next restock on February 12" gives people something solid. Urgency grounded in logistics is trusted. Urgency that exists only as a sales performance is not.

All of this assumes companies are trying to balance profit with ethics. Most aren't. Conversion metrics reward whatever works in the short term, even if it slowly erodes trust. Regulation lags behind practice. The EU's Digital Services Act gestures toward transparency and bans some dark patterns, but enforcement is uneven and limited. Industry self-regulation exists mostly as aspiration - UX ethics guidelines sound nice, but no one is punished for ignoring them. The distance between what research identifies as ethical design and what competitive markets actually incentivize remains wide.

CONCLUSION

Choice architecture exists where persuasion, autonomy, profit all collide. Sites use cognitive tricks - framing, fake scarcity, phony social proof - pushing purchases. Works initially. People adapt though. Become immune. Culture changes outcomes. Shoppers recognize scams eventually. Nudging clearly changes behavior across different areas but results depend massively on execution and circumstances.

Ethical issue isn't eliminating choice architecture - impossible, design always influences. Issue: separating help from harm. Three criteria keep appearing: transparency (hiding tactics or showing?), autonomy (genuine choice or fake?), welfare (serving customers or maximizing extraction?). Interfaces revealing persuasion methods, maintaining real alternatives, prioritizing customer welfare over profits - legitimate assistance. Everything else: exploitation masquerading as help.

Unresolved tensions pile up. Personalization enabling choice architecture simultaneously creates filter bubbles, narrows visible options. Behavioral insights powering helpful nudges become dark patterns with minimal technical modification. Competitive conversion optimization pressure drives exploiting vulnerabilities despite ethical objections.

Digital commerce exploiting cognitive shortcuts continues inevitably given competitive dynamics, technical capabilities. Critical question: exploitation creating mutual value or extracting customer money for shareholders? Ethically-constrained choice architecture legitimately helps navigating overwhelming decisions. Conversion-optimized choice architecture treats shoppers as marks, not humans. Encouraging first, limiting second: defining ethical struggle for online commerce this decade.

REFERENCE

1. Mertens, S., Herberz, M., Hahnel, U. J. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1), e2107346118. <https://doi.org/10.1073/pnas.2107346118>
2. Lindstrom, K. N., Tucker, J. A., & McVay, M. (2023). Nudges and choice architecture to promote healthy food purchases in adults: A systematized review. *Psychology of Addictive Behaviors*, 37(1), 87-103. <https://doi.org/10.1037/adb0000892>
3. Mrkva, K., Posner, N. A., Reeck, C., & Johnson, E. J. (2021). Do nudges reduce disparities? Choice architecture compensates for low consumer knowledge. *Journal of Marketing Research*, 58(4), 801-806. <https://doi.org/10.1177/0022242921993186>
4. Cheon, J. E., Nam, Y., Kim, K. J., Lee, H. I., Park, H. G., & Kim, Y. (2021). Cultural variability in the attribute framing effect. *Frontiers in Psychology*, 12, 754265. <https://doi.org/10.3389/fpsyg.2021.754265>
5. Shan, L., Jiao, X., Wu, L., Shao, Y., & Xu, L. (2022). Influence of framing effect on consumers' purchase intention of artificial meat-Based on empirical analysis of consumers in seven cities. *Frontiers in Psychology*, 13, 911462. <https://doi.org/10.3389/fpsyg.2022.911462>
6. Kim, J., & Choi, S. (2021). The differential impact of loss- versus gain-framed price-promotion messages on purchase intentions. *Journal of Retailing and Consumer Services*, 58, 102282. <https://doi.org/10.1016/j.jretconser.2020.102282>
7. Ladeira, W. J., Santini, F. de O., Araujo, C. F., & Sampaio, C. H. (2023). A meta-analysis on the effects of product scarcity. *Psychology & Marketing*, 40(6), 1139-1156. <https://doi.org/10.1002/mar.21816>
8. Mathur, A., Kshirsagar, M., & Mayer, J. (2021). What makes a dark pattern... dark? Design attributes, normative considerations, and measurement methods. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1-18. <https://doi.org/10.1145/3411764.3445610>
9. Gunawan, J., Choffnes, D., Hartzog, W., & Wilson, C. (2021). A comparative study of dark patterns across mobile and web modalities. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), Article 377. <https://doi.org/10.1145/3479521>
10. Ahuja, S., & Kumar, J. (2022). Conceptualizations of user autonomy within the normative evaluation of dark patterns. *Ethics and Information Technology*, 24(4), 52. <https://doi.org/10.1007/s10676-022-09672-9>
11. Furnham, A., & Boo, H. C. (2011). A literature review of the anchoring effect. *The Journal of Socio-Economics*, 40(1), 35-42. <https://doi.org/10.1016/j.socjec.2010.10.008>
12. Zong, Y., & Guo, X. (2022). An experimental study on anchoring effect of consumers' price judgment based on consumers' experiencing scenes. *Frontiers in Psychology*, 13, 794135. <https://doi.org/10.3389/fpsyg.2022.794135>

13. Gierl, H., & Huettl, V. (2010). Are scarce products always more attractive? The interaction of different types of scarcity signals with products' suitability for conspicuous consumption. *International Journal of Research in Marketing*, 27(3), 225-235. <https://doi.org/10.1016/j.ijresmar.2010.02.002>
14. Dowling, K., Guhl, D., Klapper, D., Spann, M., Stich, L., & Yegoryan, N. (2020). Behavioral biases in marketing. *Journal of the Academy of Marketing Science*, 48(3), 449-477. <https://doi.org/10.1007/s11747-019-00699-x>
15. Acciarini, C., Brunetta, F., & Boccardelli, P. (2021). Cognitive biases and decision-making strategies in times of change: A systematic literature review. *Management Decision*, 59(3), 638-652. <https://doi.org/10.1108/MD-07-2019-1006>

Citation: Volodymyr Rebets, "Architecture of Choice and Cognitive Biases in Digital Interfaces: A Review of Empirical Evidence on Impacts on Consumer Behavior", Universal Library of Business and Economics, 2024; 1(2): 78-83. DOI: <https://doi.org/10.70315/uloap.ulbec.2024.0102015>.

Copyright: © 2024 The Author(s). This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.