

## Uncomfortable Decisions

Paul Bloom and L. A. Paul

A behavioral economist gets a tempting job offer and agonizes over it. She meets with a friend and tells him that she is struggling to decide. It's a very tempting offer, she says, but she's happy in her current department. The offer is at a more prestigious university, but in a less attractive city. The pay is better, but the move would put her farther away from her aging parents. She goes on and on, describing her anguish, and then her friend puts up his hand and gently tells her to stop.

He says to her:

Look, you study decision-making for a living. You teach about it, write papers about it. So, you know what to do. Write a list of the pros and cons of moving and another list of the pros and cons of staying. For each item, give it a weight corresponding to its importance. Then, just add it up: do the math, treating the pros as positive and the cons as negative.

She stares at him, getting madder and madder as he speaks, until she finally yells at him: "For God's sake, this is serious!"

\* \* \*

We like this joke because it rings true. There is something that seems wrong about making a significant decision through this type of process. It might be useful to make a list of pros and cons, and even to weigh them, but few of us would act solely on the outcome of this mathematical calculation.

Other decision processes are similarly unpalatable. We do not like to make significant decisions based on social science research. We do not like to choose randomly, even in cases where flipping a coin is plainly the wisest choice (Keren and Tiegén 2010)—as when the values of the options are impossible to distinguish and decision-making is very costly. Many of us are reluctant to defer to another person, even if we believe that the other person is wiser. We also often do not want to appeal to an artificial intelligence (AI), even if we believe it to have decision-making powers superior to those of humans, a bias dubbed "algorithm aversion" (e.g., Dietvorst, Simmons, and Massey 2015; Castelo, Bos, and Lehmann 2019; though see Logg, Minson, and Moore 2019 for cases in which this does not apply).

What do these uncomfortable decision processes share? We can characterize them as *impersonal*. It's not you who is making the decision; it is someone or something else. This is to be contrasted with personal decision-making, where, well, you just decide, perhaps after mulling it over for a long time. To put it colloquially, for normal acceptable choices, you "own" your decision. Now, at some level, this distinction is an illusion—one's personal decisions are determined by causal processes and might themselves reduce to algorithms of the sort that the friend of the behavioral economist urged her to rely on. But, nonetheless, it doesn't feel this way; a satisfying decision-making process is more than an algorithmic process. It's something that *we* do, and in a certain way.<sup>1</sup>

But what, precisely, makes certain decision-making processes seem impersonal? What do we want from our decisions?

\* \* \*

Before looking at options, we should note two subtleties about the phenomenon we're talking about.

First, our concern about impersonal decision-making processes applies to significant decisions, not so much to trivial ones. There are those who allow their partners, or sometimes even their waiters, to choose what dishes they will eat in a restaurant, and we do flip coins to resolve small dilemmas. Studies find that people are uncomfortable using a coin toss to decide which life to save but are fine with flipping a coin to decide authorship order or whether to attend the opera or the theater (Keren and Tieggen 2010). We suspect, though, that even here there is some discomfort with relying entirely on these processes. We often want to be able to *approve* the dinner recommendation, and it's a cliché that some people, finding that the coin's outcome wasn't what they wanted, might choose to flip again.

Still, we are especially interested in significant decisions here. Even people who are comfortable making trivial decisions impersonally will balk at making big ones this way. One of the authors of this chapter has argued that this distrust of such impersonal procedures is particularly salient in the case of "transformative experiences"—experiences, like choosing to become a parent or emigrating or going to war, where you must choose between new lives, or choose between having a new life or keeping your old one (Paul 2014). Notably, with these types of decisions, there's an inability to properly imagine what this new life is like.

Second, we have no qualms about using impersonal sources of information to *help* us make important decisions. We assume that every reader of this chapter regularly uses ratings—from Yelp, TripAdvisor, Amazon, and so on—when deciding where to eat, which hotel to stay at, what book to read, and so on. This is true for even the most significant of choices. Someone deciding whether to have a child is likely to be very interested in the data on whether children make you happy or give meaning to your life and very likely to want to talk to new parents and hear about their experiences. Someone who was wrestling with a decision about cancer treatment might well be very interested in data on outcomes and survival rates.

We are similarly happy to get information from more "internal" procedures such as list making and priority rankings. It would kill the joke if her friend suggested that she

write down the pros and cons of each option and they talk them through together—Who would get upset at that? Indeed, Charles Darwin famously wrote a list of pros and cons in his journal when deciding whether to marry his cousin Emma Wedgwood—“a nice soft wife on a sofa with good fire, & books & music perhaps,” but then again, “perhaps quarrelling”—but it’s clear that he’s just spitballing ideas to himself and the decision is all his own (see Quammen 2007). Other internal procedures include imagining that you made a decision and assessing how you feel about it or figuring out what one would recommend to a third party (Galef 2021 enumerates several processes of this sort). We are comfortable with all of these—so long as the final decision is left up to us.

\* \* \*

What is the source of our reluctance?

One possibility is that we don’t trust the reliability of certain impersonal decision-making processes. This concern is clearest enough when it comes to survey data. Perhaps this is because we believe we are special. This can be rational in certain cases. One of the authors dislikes cheese, and if he went to a restaurant with the best reviews, he might end up at one that specializes in cheese—most people do like cheese—and this would be terrible.

Taken to the extreme, though, this skepticism is unreasonable. Nobody should refuse a vaccine because, though it works for others, they believe (without any evidence) that their body is different. The same holds for preferences. One can imagine, along the lines of philosophical thought experiments, Perverse Man—a person whose tastes are the opposite of most everyone else’s. A hotel that gets 9.8/10 on TripAdvisor—“best hotel in Cancun,” “a perfect paradise”—would be misery for Perverse Man, while a hotel ranked 1.4—“disgusting,” “horrible”—would be perfect for him. But there is no Perverse Man. While everyone is unique and nobody should choose just by the numbers (a hotel can be ranked highly because of its golf course, but what if you don’t golf?), surely, in the absence of other information, everyone should prefer a highly ranked hotel to one that everyone else hates.

Our suspicion is that the source of the reluctance to rely on certain decision-making processes isn’t actually concern about their utility. For one thing, as mentioned before, most people believe that ratings, anecdotes, and so on *are* useful. They’re just not sufficient. The behavioral economist in the joke might believe, sincerely, that ranking the pros and cons of a decision is a useful way to proceed, but nonetheless not want to rely solely on such a process.

We suggest instead that this aversion has deeper roots, and we are interested in two of them here—one concerning authenticity, the other autonomy. (These terms are used in many ways, and we don’t have settled views on how to capture all their ordinary or even technical connotations, so we will work with a rough distinction.)

A certain type (not the only type) of authentic decision-making, as we’ll understand it, is decision-making in which an agent intentionally and knowledgeably chooses an option in a way that is “true to themselves.” This way of choosing requires a certain kind of understanding of one’s options, which we’ll expand upon in the following

paragraphs. Such decision-making can be important in contexts where one is making a life-changing decision of great import, such as the choice to emigrate, start a family, or embark on a major career change. A paradigm case of making such an authentic decision concerning oneself requires this type of understanding, which is (arguably) gained through reflection, careful assessment, and imaginatively evolving the world forward under different possible actions in order to knowledgeably and volitionally choose the outcome that will lead to the most happiness or greatest life satisfaction for oneself (and relevant others).<sup>2</sup>

As one of us, a fan of authenticity, has put it:

You use your memories from the past, your beliefs about the present, and your anticipations about possible subjective futures to formulate your current and evolving [preferences], as well as to develop temporally extended, forward-looking, subjective projections about what will happen. . . . You use this reflection on what you think these events will be like, that is, what you think your lived experience will be like, to authentically determine your preferences about your future. (Paul 2014: 106–7)

Autonomy, on the other hand, involves control. It involves making the decision yourself, and not off-loading the process to ratings, polls, an algorithm, an oracle, and so on, though, again, you might be receptive to information from these other sources.

There's a sense of which this prizing of autonomy might be perverse. Suppose a website is better at choosing books for you to read than you are—and you know it. The books it chooses make you happier, more engaged, you learn more, and so on than the books you choose for yourself. If your goal is to be happier, you should just let the website choose, even if it's not the choice that you want to make at the time. But autonomy considerations would push against this.

Note critically that autonomous choices can be made on any grounds at all. One can make an autonomous choice based on a gut feeling or even based a criterion you might acknowledge as foolish (“I want to go to this hotel because its name reminds me of the hotel I went to on my honeymoon.”). There's certainly no requirement that one imaginatively reconstruct the outcome of a decision-making process. Many authentic processes are autonomous ones. But what we are interested in here is in exploring the discomfort involved with certain kinds of autonomous decisions that are not, in the way we are defining it here, authentic.

\* \* \*

Now let's go back to our questions about comfortable and uncomfortable decisions. If we are choosing whether to embark upon a new life, we can do it comfortably by deliberating about our options, accurately assessing the value of each possibility through imagining what it would be like to be those possible new selves leading those new lives, and personally choosing the best one. This approach hits all the “like” buttons. It is both authentic and autonomous in the way we've been framing these features. It involves making the choice ourselves, in an informed way, where we

undertake an appropriately informed process of evaluating our options and comparing them in order to make the best choice.

In the case of the job offer, if our behavioral economist were to authentically engage in making this decision, she'd imaginatively put herself in each possible job, and assess (or maybe create) the value of that option. Then she'd compare these values to form a preference, either preferring Job A over Job B (if A's value is higher than B's value) or vice versa. And then she would act on this preference. Autonomy plus authenticity—what's not to like?

But then there are decisions that are neither autonomous nor authentic. Indeed, if the relevant kind of authenticity requires an imaginative recreation, a simulation, of the life one would end up taking, then perhaps transformative experiences are simply not amenable to this kind of authentic decision-making process. We can't decide in a satisfying way to become a parent, say, because someone without children cannot fully imagine what it is to have children (Paul 2014). Perhaps the same is true for a decision like taking a new job.

Consider now a less welcome alternative. Perhaps we are not far away from the day when our behavioral economist would be able to hire a benevolent AI to do the calculation for her. She might engage the Google Deepmind Concierge to have a bespoke algorithm provide an analysis based on masses of data collected on her (and the rest of humanity) since the birth of the internet. The AI uses the algorithm to assess what jobs that others relevantly similar to her have liked most, scales these results (somehow) to make comparisons meaningful, and chooses the highest-scoring job. For a small added fee, it informs her prospective employers and politely declines the alternative position on her behalf. Our professor is highly likely (let's assume) to get a result that, afterwards, she will testify to as making her happier than other options would have. But her decision is neither authentic nor autonomous.

There are similar decision processes that also don't meet these two criteria. She might flip a coin. Or be paralyzed with indecision, so that the offer goes away. Or consider decisions that are made through involving certain forms of nudging: perhaps someone else sets up the world for her in ways that unobtrusively guides her to the choice that is her best option. Or, perhaps someone else simply makes the choice for her, acting in her best interests, someone like the AI, only made of flesh and blood. None of these alternatives are autonomous or authentic.

But now, as a distinct class of cases, consider choices that are autonomous but not authentic. Imagine, in search of doing something a little more satisfying, our professor rejects the services of the AI, opting instead to have Google package all the relevant numbers, including the assigned, scaled values, and deliver it to her in easy-to-read, bite-sized chunks. Given this information, she maps out the decision tree, does the math—and perhaps sleeps on it, indulges in silent prayer—and then decides on which job to take. This is likely to get the same result as using the benevolent AI. (Maybe it's still worth it to hire the AI to decline the offer she doesn't want.) But it's not authentic in the sense we are exploring, because it doesn't require simulating the different alternative lives and making the comparisons based on this knowledge.

\* \* \*

The two authors of this chapter agree on the importance of autonomy. People are happier when they themselves are making the decision, rather than when it's offloaded to a coin toss, an AI, and so on.<sup>3</sup> The authors also agree that for non-trivial decisions, the assumed efficacy of the procedure matters. Our behavioral economist is unlikely to ask her four-year-old niece for advice or bring out an Ouija board.

We disagree about the authenticity requirement, however. One of us (Paul Bloom) thinks that people would be comfortable appealing to the AI-facilitated process described earlier, so long as the ultimate choice is in their hands; the other (L. A. Paul) is skeptical. To make this sort of decision authentically, she skeptically argues, one has to get a real sense of what each new life will bring, and to do so, one has to be able to accurately imagine themselves, or perhaps a transformed version of themselves, in the new situation.

The problem, she thinks, is that it simply isn't acceptable to rely on the happiness utilities assigned by an expert AI to make this sort of choice. There just seems to be something wrong, at an intuitively deep level, with using numerical quantities to capture what matters. For her, this is why the joke works. Officially, experts in behavioral economics and other social sciences endorse rigorous, mathematically based treatments of decision-making and evidentially based policy formation. Officially, the right way for our economist to make the decision between job offers is to assign numerical values, calculate expected utilities, and choose in a way that maximizes expected value. But, as the joke illustrates, this feels wrong, at a gut level—it's not at all how we feel we should make the decision.

Is this just some sort of residual anti-scientific bias, some fear of being reduced to mere mathematics, a last gasp of religiosity that needs to be put to rest in the modern world? Perhaps we simply need to bring our gut feeling to heel and impose, as careful thinkers, a more rational approach. But our pro-authenticity author argues otherwise. She thinks, rather, that our gut intuition reflects our knowledge of a real fact: numerical quantities, while useful in many contexts, are inadequate for representing the value of the nature and character of the kinds of lives we are contemplating. (In fact, she'd argue, doing the work of imaginative evaluation can be itself seen as a rational process.)

The need for a deep, intuitive understanding of how each job choice would change one's life is why the imaginative evaluation is so important. The point isn't that you can't compare the values on this approach. You can: you can compare them, and have a defined preference, and understand which option it is that you'd prefer. The point is that there's a richness to the way we want to assign values to these life choices that isn't being captured by the numbers. There's something about the experience of imaginatively understanding these different lives and the different ways that you would respond to each job choice that allows you to make a more informed assessment about the pros and cons, and thus know how to assign and represent their values in a way that will let you make a fully informed, and thus authentic, choice. However it is that we want to represent and assign values to our life options, it isn't captured by the official approach.

And, she thinks, as the joke shows, we will not be fully satisfied with a process that lacks such authenticity. Knowing, say, that becoming a parent made 65% of people

happier, or that declining a job most likely will maximize one's priorities (the pros will outweigh the cons) just isn't enough in this context.

To illustrate the kind of details that could be obscured by simply looking at the numbers, consider the following point about these kinds of transformative decisions.

Making a major job change does more than change what you are doing with your life. It changes who you are. This is most obvious when we think about choosing between very different careers. If you become a dentist, who you are and what you care about will be very different from who you are and what you care about if you become a pianist. In this sense, the choice between careers is a deeply personal transformative choice about whom you want to become, about what to give up and what to embrace, and about choosing the life you want to live. As we can put it, it's a choice between your future selves. Which self do you want to become? This is the problem that the person must solve.

This, L. A. Paul argues, suggests that in such cases, the decision-making process matters. The process is internal in a special way because your choice determines what you want. If you choose to become a dentist and it makes you happy, it's the process of becoming a dentist that makes you happy. You are glad, at that point, that you hadn't decided to become a pianist! On the other hand, if you choose to become a pianist, it's the process of becoming a pianist that makes you happy about being a pianist. You're glad, at that point, that you hadn't decided to become a dentist!

To put it differently, such a choice is partially endogenous. And this endogeneity creates a problem with our interpretation of the results of the choice. What if your dentist self would testify to a +4 happiness, saying that your wealthy, comfortable life is fantastic, while the pianist would testify to a +2 happiness, but claim that your life, devoted as it is to music, is so much more satisfying? Given that the future self you choose will replace the self that you are now, how are you to know which of them is the best, most natural extension of who you are right now? How are those selves comparable to who you are now, as you choose, and by extension, how are they to be compared with each other (Paul and Healy 2018; Paul 2020)?

The job choice faced by our behavioral economist is like this—the choice isn't quite as stark, but if making a major job change can endogenously affect you in the way that choosing a career can, then the structure reappears. If the behavioral economist chooses the more prestigious university, as she embeds herself into her new life, it will change her in ways that, in the end, will make her value her new life over her old one. If she stays put, she'll strengthen her current values and be glad she chose being close to her parents over a fancy intellectual climate.

With so much to lose (or gain) and so many unknowns to face, the stakes are high. Which life is the better one for her to choose? Which one of these lives is most authentically the extension of the life she is living right now, as she makes her choice? It's tempting to think the question could be resolved if she could just imagine herself into the lives of her different possible selves and compare them. Then, at least, she could decide which possible self better captures who she is and what she really wants from life.

Unfortunately, this is precisely what she can't do. And, in the case where you are supposed to appeal to an AI-facilitated process, you can't either. You can't, and

don't, simulate these new lives and compare them yourself. Instead, you turn to an impersonal process of assessment and scaling performed by an AI. And it simply isn't clear whether the AI has the resources to assess the options and thus determine your preferences in the right way. If all that matters to the AI is maximizing happiness, other things that seem important could be left out. In particular, the choice that the AI makes for you could fail to be a legitimate extension of your current preferences. Rather, it could be a choice that involves a replacement of your current preferences. Say the AI tells you to choose the dentist option: after all, it brings greater happiness (+4 over +2). In this scenario, you could be choosing to replace your life with a new life that, while it is +4 happier, isn't really an extension of what you care about now. In other words, it is not a choice that truly captures who you are, at this moment of choosing. Rather, the replacement life brings with it a kind of psychological rearrangement, making you glad, as a result of the process, that the replacement has occurred. If you rely solely on the AI, there is no way of knowing whether your current preferences will be satisfied, as opposed to merely replaced.

Our authenticity-friendly author doesn't think that people would be worried about this particular problem. Rather, she thinks they would be worried that, in some more general sense, simply using AI-generated numerical values won't capture the experiential complexity and meaningfulness of each option. In this situation, discomfort stems from the fact that the numbers might not represent all the factors that matter to you, since these factors go beyond what they can capture. If you simply rely on the AI, you must choose without thinking it through for yourself. All you have are the outcome-numbers gathered by the AI. And this is uncomfortable.

\* \* \*

The other author, less persuaded about the importance of authenticity, agrees with many of the discussed points. He very much agrees that people want more than happiness—we also want to live meaningful lives, to be moral, to have purpose, and other things as well. (Indeed, he has just published a book making exactly this argument: Bloom 2021.) Our motivational pluralism makes decision-making difficult, as these values often clash. And it suggests that we would be most interested in a decision-making process that draws upon multiple sources of information about how we will end up after the decision has been made. If the AI only took happiness into account, it would be inadequate.

Bloom agrees as well that the choices one makes influence how one evaluates the outcome. This is one of the more robust findings in social psychology, a classic example of the phenomenon of cognitive dissonance (see Cushman 2020 for review). Get people to choose between two things that are valued roughly equally, and, later, they will tend to like the chosen one more and the unchosen one less (e.g., Egan, Santos, and Bloom 2007). This effect occurs even when the choice is blind, where they don't know what they're choosing (Sharot, Velasquez, and Dolan 2010).

Finally, he agrees that people often don't like to think about important decisions in terms of numerical rankings. Among other things, it's rather unseemly. It would be a poor marriage proposal indeed if the suitor, on bended knee, proudly states to

his beloved the precise estimated values of “married” versus “not married” (“Married wins!” he sings to her.). Better to say instead, “I couldn’t imagine any other life than one with you. It was no choice at all.” Indeed, there may be some intrapersonal benefit to thinking of important and difficult-to-reverse decisions (getting married, having children, religious conversions) as no-brainers, the only conceivable thing to do. It certainly reduces regret. Perhaps this is one function of cognitive dissonance.

Decision-making has to happen, though. To choose A over other options is to come to believe that A ranks above the rest, and this requires putting the options on a common scale. Yes, if you make the decisions through simulating alternative possibilities, you will have access to information that is rich and qualitative, and yes, assigning the values is a complex matter. Perhaps the values you assign are represented as numbers, or perhaps they are in some non-numerical representational format, but in the end, comparing values is what decision-making comes down to. If the suitor insists that his own decision-making process is exempt from this, then he is deluding himself.

More generally, both authors agree that none of these points—our motivational pluralism, our distaste for numerical comparison, and the fact that making a decision influences what we think about the outcome—entail that we want authenticity in our decision-making process. If you concede that non-authentic processes can tell you that some choice is likely to make you happier, then it suggests that this process can also tell you that this choice will make you more satisfied with your life in some more abstract sense. To take a real example, while there is a lot of debate over the interpretation of the data (e.g., Nelson, Kushlev, and Lyubomirsky 2014), there is some evidence that parents report having less pleasurable experiences than non-parents—but also report more meaningful lives (see Bloom 2021 for review and discussion). Pluralistic decision-making does not require authenticity.

Similarly, both agree that cognitive dissonance is a phenomenon that is general to decision-making, applying both to major life experiences and to totally mundane choices like deciding which of two identical cookies to pick up. Thus, the partially endogenous nature of decision-making does not in itself provide evidence for a desire for authentic decision-making processes, nor does the fact that people often don’t like to think in terms of quantitative comparisons. This is again compatible with all sorts of processes that are not authentic.

In the end, the core disagreement between the authors might be about how people construe decisions, particularly high-stakes decisions. The author who is sympathetic to authenticity concerns sees such decisions in terms of changes of self, and perhaps the replacement of one self by another—and she believes that, at some level, when making high-stakes decisions, non-philosophers recognize that such choices bring major changes in the kind of person one is. And so authenticity becomes important; people want to know how the choice will change them. The author who is skeptical about the importance of authenticity rejects this view of common-sense decision-making. Someone deciding whether to be a parent, say, will want to make the choice that will lead to a life that has the most happiness, purpose, and meaning—won’t naturally construe it as involving a change in the sort of person they are. And so authenticity, as defined here, doesn’t matter—if they trust the data and can choose autonomously, that’s all they need.

While the authors disagree about the actual importance of authenticity and about whether people value it, we agree that at least the second question is an empirical one. What *do* people think of authentic and non-authentic decision-making processes? And so we will end with a clichéd proposition—but one that we believe really applies here: more research is needed.<sup>4</sup>

## Notes

- 1 Indeed, the same psychological forces that make us want to “own” decisions we make in the world might also lead us to discomfort about theories of the mind that deny a singular agentic self. As part of a critique of massive modularity, Jerry Fodor (1998) expresses this desire for decision-making “ownership” with characteristic verve: “If, in short, there is a community of computers living in my head, there had also better be somebody who is in charge; and, by God, it had better be me.”
- 2 “True to themselves” and “knowledgeably and volitionally” admit of many interpretations, but we will work with these intuitive glosses in hopes of making incremental progress.
- 3 This is a claim about how people like to make decisions when we have to; it is not meant to imply that we always enjoy making decisions. Often, we do not (e.g., Schwartz 2004).
- 4 For some research on this question, which came out as this chapter went to press, see “Deciding to be Authentic: Intuition is Favored over Deliberation When Authenticity Matters,” by Kerem Oktar and Tania Lombrozo, forthcoming in *Cognition*.

## Bibliography

- Bloom, P. (2021). *The Sweet Spot: The Pleasures of Suffering and the Search for Meaning*. New York: Ecco.
- Castelo, N., M. W. Bos and D. R. Lehmann (2019). “Task-Dependent Algorithm Aversion.” *Journal of Marketing Research*, 56 (5): 809–25.
- Cushman, F. (2020). “Rationalization is Rational.” *Behavioral and Brain Sciences*, 43: 1–16.
- Dietvorst, B. J., J. P. Simmons and C. Massey (2015). “Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err.” *Journal of Experimental Psychology: General*, 144 (1): 114.
- Egan, L. C., L. R. Santos and P. Bloom (2007). “The Origins of Cognitive Dissonance: Evidence from Children and Monkeys.” *Psychological Science*, 18 (11): 978–83.
- Fodor, J. (1998). “The Trouble with Psychological Darwinism.” *London Review of Books*, 20 (2): 11–13.
- Galef, J. (2021). *The Scout Mindset: Why Some People See Things Clearly and Others Don't*. New York: Penguin.
- Keren, G. and K. H. Teigen (2010). “Decisions by Coin Toss: Inappropriate but Fair.” *Judgment and Decision Making*, 5 (2): 83.
- Logg, J. M., J. A. Minson and D. A. Moore (2019). “Algorithm Appreciation: People Prefer Algorithmic to Human Judgment.” *Organizational Behavior and Human Decision Processes*, 151: 90–103.

- Nelson, S. K., K. Kushlev and S. Lyubomirsky (2014). "The Pains and Pleasures of Parenting: When, Why, and How is Parenthood Associated with More or Less Well-Being?" *Psychological Bulletin*, 140 (3): 846–95.
- Oktar, K. and T. Lombrozo (forthcoming). "Deciding To Be Authentic: Intuition is Favored Over Deliberation When Authenticity Matters." *Cognition*.
- Paul, L. A. (2014). *Transformative Experience*. Oxford: Oxford Univ. Press.
- Paul, L. A. (2020). "Who Will I Become?" In J. Schwenkler and E. Lambert (eds.), *Becoming Someone New: Essays on Transformative Experience, Choice, and Change*, 16–36. Oxford: Oxford Univ. Press.
- Paul, L. A. and K. Healy (2018). "Transformative Treatments." *Noûs*, 52: 320–35.
- Quammen, D. (2007). *The Reluctant Mr. Darwin: An Intimate Portrait of Charles Darwin and the Making of his Theory of Evolution (Great Discoveries)*. New York: WW Norton & Company.
- Schwartz, B. (2004). *The Paradox of Choice: Why More is Less*. New York: Ecco.
- Sharot, T., C. M. Velasquez and R. J. Dolan (2010). "Do Decisions Shape Preference? Evidence from Blind Choice." *Psychological Science*, 21 (9): 1231–5.

