Present Bias: Lessons Learned and To Be Learned[†]

By Ted O'Donoghue and Matthew Rabin*

Present bias is an old idea. The notion that people are susceptible to the over-pursuit of immediate gratification dates (at least) to the ancient Greeks. In academic research, psychologists working with animals in the 1960s and 1970s proposed "hyperbolic discounting"—a functional form of discounting that generates present bias—as a natural way to represent how animals respond to time delays, and later research in psychology extended this idea to humans.[1] Also in the 1960s, economists investigating general implications of time-inconsistent preferences used as an example the now-popular β, δ functional form—which also generates present bias.[2] But present bias really took hold in economics following David Laibson's dissertation (Laibson 1994).

The literature has blossomed in the past 20 years. Research has led to a much better theoretical understanding of present bias, when and how to apply it, and which ancillary assumptions are appropriate in different contexts. Empirical analyses have demonstrated how present bias can improve our understanding of behavior in various economic field contexts. While the model is clearly not "correct"—no model is—for many contexts it is proving a useful, tractable, and (importantly) disciplined improvement in economic analysis. Nonetheless, there is still much to learn.

I. A Brief Overview of Present Bias

Like exponential discounting, present bias is a model of discounting. One indication of the success of present bias is that, much as for exponential discounting, most readers do not need a review of the structure of the model. Hence, here we present only a brief summary.

Suppose that intertemporal preferences from the perspective of period t can be represented by $U^t = \sum_{\tau=t}^{T} D(\tau - t)u_\tau$, where u_τ is instantaneous utility experienced in period τ and $D(x)$ reflects the discounting associated with a delay $x \in \{0, 1, 2, \ldots\}$. While more general variants of present bias exist, the β, δ functional form typically used assumes that

$$D(x) = \begin{cases} 1 & \text{if } x = 0 \\ \beta \times \delta^x & \text{if } x > 0. \end{cases}$$

With this functional form, $\beta = 1$ corresponds to exponential discounting, while $\beta \in (0, 1)$ reflects present bias. When $\beta < 1$, the model's predictions may depend on an ancillary assumption about whether one is aware of how preferences change over time (sophisticated), unaware of how preferences change over time (naïve), or something in between (partially naïve).

II. Some Lessons Learned

This section summarizes some lessons learned over the past 20 years. While these lessons have been learned well among those immersed in the literature, we highlight these lessons for others.

Lesson #1: Present Bias Operates on Utility.—All discounting models—exponential, present bias, or other—operate on the timing of utility. Importantly, they do *not* operate on the timing of purchases or on the timing of monetary payments. Indeed, perhaps the most basic lesson from standard saving-consumption models is that, in the absence of liquidity constraints and

*O'Donoghue: Department of Economics, Cornell University, 482 Uris Hall, Ithaca, NY 14853 (e-mail: edo1@cornell.edu); Rabin: Department of Economics, Harvard University, Littauer Center M-8, Cambridge, MA 02138 (e-mail: matthewrabin@fas.harvard.edu). For helpful comments, we thank Dan Benjamin, Nava Ashraf, and other participants in our AEA Session.
[†] Go to http://dx.doi.org/10.1257/aer.p20151085 to visit the article page for additional materials and author disclosure statement(s).

[1] See Ainslie (1992) for an overview.
[2] See in particular Phelps and Pollak (1968) and Pollak (1968). These papers build on an earlier paper by Strotz (1956) that investigates general time-inconsistent preferences without special focus on the case of present bias.

uncertainty, choices between different streams of monetary payments are driven entirely by maximizing the present discounted value of wealth, and preferences are irrelevant.

The idea that discounting models, present-biased or not, operate on utility is not so much a lesson learned over time as a fundamental lesson inherited from standard economics. However, we emphasize this lesson because it is sometimes forgotten by those new to present bias.

To highlight the importance of distinguishing utility flows from money flows, we note two natural intuitions for how present bias can lead people to *delay* money flows—the opposite of what one would predict if one applied present bias directly to money flows. First, Laibson (1997) demonstrates how sophisticated present bias can lead people to choose to constrain their future liquidity, and one way to do so is to defer money flows until the more distant future. Second, O'Donoghue and Rabin (1999b) demonstrate how naïve present bias applied to the effort required to optimize one's retirement saving can lead people to delay the accrual of money flows rather than move it forward.

This lesson is also important when using quasi-linear preferences. In such models, if a person purchases some good at a price p, the person is assumed to experience a utility gain from consuming the good, and a utility cost of $-p$ interpreted as forgone consumption of other goods. When applying present bias, one clearly must be explicit about whether the utility from the purchased item is experienced immediately or over time. But when does that forgone consumption occur? In most applications, the most natural assumption is that, regardless of when the monetary payment is made, it is forgone *future* consumption.

Finally, we note that, related to this lesson, economists are finally appreciating that the use of money alone is not appropriate for experiments that investigate time preferences. When choosing between time-dated monetary payoffs, subjects' time preferences are irrelevant unless they are liquidity constrained. For most experiments, it seems highly unlikely that subjects are liquidity constrained in a relevant way.[3] In

recognition of this issue, recent experiments have asked subjects to choose instead between time-dated utility flows, such as when to exert real (unpleasant) effort.[4]

Lesson #2: Present Bias is About Now.— Psychologists suggested a hyperbolic functional form for discounting—i.e., $D(x) = 1/(1 + kx)$. Economists instead adopted the β, δ functional form in large part for tractability, and this model quickly became the workhorse for the literature. In the early years, we worried about whether using the β, δ functional form was restrictive. As time passed and our intuitions developed, we became less worried. Indeed, many of us now believe that, in fact, the β, δ functional form better captures the underlying psychology—that the vast majority of the action (relative to time consistency) is biased toward *now*.

To be fair, there is little direct evidence that compares different functional forms. Psychologists primarily compare the hyperbolic versus exponential functional forms, and do not consider alternative functional forms. At the same time, economists have primarily limited ourselves to studying the β, δ form.

On a related note, researchers have pointed out, correctly, two confounds in time-preference experiments: (i) payoffs received now might be viewed as certain while payoffs to be received in the future might be viewed as uncertain, and (ii) payoffs to be received in the future might involve higher transactions costs. Some have suggested eliminating these confounds by attaching a front-end delay to all payoffs. If present bias is about now, however, this approach also eliminates present bias as an influence on behavior. Indeed, when experiments with a front-end delay find little evidence of deviations from exponential discounting, such experiments provide support for the β, δ functional form relative to the hyperbolic functional form.

Lesson #3: Any Noticeable Short-Term Discounting is Evidence of Present Bias.—Most early evidence on present bias emphasized time inconsistency as the smoking gun for present bias. Researchers are now more comfortable

---

[3] Yet in many experiments that use monetary payoffs, many subjects seem not to behave in a wealth-maximizing way. It is an open question exactly what those subjects are doing.

[4] See, e.g., Augenblick, Niederle, and Sprenger (2013), who directly compare the use of money versus real effort, and find evidence of present bias only for the latter.

with a simpler argument based on calibration. Exponential discounting does not permit any noticeable discounting over short horizons because such discounting would compound to predict counterfactually severe discounting over longer horizons. Present bias, in contrast, is all about noticeable short-term discounting.[5]

To illustrate, suppose the only thing we observed about Johnny is that he cares 1 percent more about his utility today than tomorrow. If Johnny were an exponential discounter, this observation would imply a yearly discount factor of $(0.99)^{365} = 0.026$. This is clearly counterfactual: nobody cares 40 times more about now than one year from now. Hence, Johnny could not be an exponential discounter.

What about smaller short-term discounting? Under exponential discounting, caring just 0.1 percent more about today than tomorrow implies a yearly discount factor of $(0.999)^{365} = 0.694$, or caring 1 percent more about today than next week implies a yearly discount factor of $(0.99)^{52} = 0.593$. At first glance, these numbers might appear reasonable, because estimates suggest that discounting one year from *now* by 0.69 might be realistic. However, exponential discounting implies that the same 0.69 applies between any two years—e.g., people would need to care only 0.69 as much about eight years from now as seven years from now. This is again counterfactual. Observing more than infinitesimal short-term impatience is sufficient to reject exponential discounting, and quite consonant with present bias.

Lesson #4: Naïvete Makes Sense, and Doesn't Always Lead to "Crazy" Behavior.—Early work on present bias in economics focused on the assumption of complete sophistication, in part based on a belief that naïvete would lead to "crazy" behavior. In our own work, we highlighted how complete or partial naïvete are valid assumptions in the sense that they can be modeled in a rigorous fashion (see O'Donoghue and Rabin 1999a, 2001, and subsequent research).[6]

Moreover, it is not the case that naïvete in general predicts "crazy" behavior, at least not in the sense of predicting unrealistic behavior. In some situations, the predictions of naïvete are identical to the predictions of sophistication. In other situations, naïvete predicts very damaging behavior, but of the sort that, unfortunately, we observe way too often in the world. For instance, while it is inconsistent with exponential discounting or sophisticated present bias for a person to predict hundreds of times that she'll start a diet, quit smoking, or write a referee report tomorrow when she won't, these seem to be the types of behaviors that we observe. More and more research is suggesting that models that incorporate naïvete (at least to some degree) seem to better explain behavior.

Lesson #5: There is a Natural Intuition for How to Identify the Parameters of Present Bias.—Economists often estimate parameter values in structural models. Over time, researchers have developed a good intuition for how to estimate the parameters of present bias (i.e., β and δ). Specifically, one needs data on multiple types of choices, some which involve trade-offs between immediate utility and future utility, and others which involve trade-offs between future utility and further-future utility—i.e., some decisions which are heavily influenced by present bias (β), and other decisions which are primarily influenced by longer-term discounting (δ). Two applications illustrate this intuition nicely. Angeletos et al. (2001) study present bias in the context of savings-consumption decisions, and describe how identification can come from a combination of a household's credit-card borrowing to finance current consumption (heavily influenced by present bias) and a household's savings accumulation for retirement (primarily influenced by longer-term discounting). DellaVigna and Paserman (2005) study present bias in the context of job search, and describe how identification can come from a combination of a person's effort to search for a job (heavily influenced by present bias) and a person's reservation wage applied to job offers (primarily influenced by longer-term discounting).

---

[5] This argument is analogous to calibration arguments for how any noticeable risk aversion over modest stakes is evidence against expected utility.

[6] Prior to our work, research on naïvete was scarce. Strotz (1956) introduced the distinction between sophistication and naïvete, but his formal analysis considered only sophistication. Pollak (1968) formally analyzed naïvete, but

only as a methodological approach to prove a result about sophistication. Akerlof (1991) was the first to seriously consider implications of naïvete, although his analysis was not framed in terms of present bias.

Researchers have been less successful at generating well identified parameter estimates for naïvete. In principle, one needs data on choices that reflect (perhaps indirectly) people's predictions for their own future behavior—e.g., purchases now intended to be consumed in the future, or decisions now that impact the future prices that will be faced. In practice, clean data that permit precise identification of the degree of naïvete has proven hard to find. More often, we see either evidence that indicates at least some sophistication (e.g., commitments) or evidence that indicates at least some naïvete (e.g., commitments that don't work as intended or clearly inefficient procrastination).

Lesson #6: Welfare Analysis is Doable.— When preferences are time-inconsistent, welfare analysis becomes tricky because there are multiple preferences that one could use. While there have been growing pains—some economists initially suggested that, given this issue, we shouldn't do welfare analysis at all—economists seem to have accepted that welfare analysis is doable. In particular, the most important lesson on this dimension is that one should be rigorous and precise in exactly the way that economists usually are—be very clear about what assumptions one is making about how to assess welfare, and assess whether one's welfare conclusions are robust to other assumptions about how to assess welfare.

While there is no agreed-upon welfare criterion, we have argued for the use of "long-run utility," wherein we use the intertemporal utility function U^t evaluated from a prior (or long-run) perspective, which is equivalent to using $\beta = 1$. The early literature suggested instead using a Pareto criterion in which intertemporal utility from all perspectives is taken into account. It turns out that these two approaches frequently yield the same conclusions. Based on such results, we conjecture that long-run utility will in the end be seen as best single criterion.[7]

III. Lessons To Be Learned

Despite the progress of the past 20 years, there is still much to be learned. We next discuss some important open questions.

Question #1: How Can We Improve the Predictions of Present Bias?—Typically our models explain only some of the variation in the data, and it is natural to seek ways to improve our models. A popular approach among researchers is to enrich the model of present bias. Two potentially important ways to enrich the model have been discussed. First, one might incorporate heterogeneity in present bias, and thereby explain some of the variation in behavior across individuals. Indeed, more and more research is finding correlations between measures of present bias—e.g., from a survey—and field behaviors. Second, one might incorporate ways in which the magnitude of present bias depends on context, and thereby explain some of the variation in behavior across contexts. Quantitative estimates of discounting do tend to vary across contexts, and some models (e.g., dual-process models) explicitly incorporate context-specific discounting.[8]

While such enrichments are surely useful, we worry that researchers are perhaps excessively focused on the details of present bias, and not focused enough on other details. Present bias makes no predictions about behavior independent of (i) utility functions (what people like and don't like), (ii) the timing of decisions, and (iii) constraints and transactions costs. In seeking to improve our models, we must not forget to be careful in accounting for these other factors that standard economic theory deems relevant.

We worry, for example, about attempts to explain heterogeneity in behavior primarily due to heterogeneity in present bias. Heterogeneity in cigarette consumption, for instance, is far more likely due to heterogeneity in tastes for cigarettes, or in prices and extent of peer pressure toward cigarettes experienced in one's youth. We also think there are reasons to resist overexcitement about variation of present

---

[7] There also exist other exotic welfare criteria which (we believe) are less in the spirit of traditional economics. But even for these, we conjecture that the ancillary assumptions needed to fit economists' intuitions about welfare will make those models line up with the long-run utility criterion.

[8] One might also search for improvements to the functional form of present bias. We conjecture, however, that such improvements are unlikely to be important.

bias across contexts. Such variation is often confounded with variation in the utility function and constraints across contexts—indeed, we conjecture that the latter is likely far more pronounced.

Question #2: How Important is Temporal Aggregation?—Data come in different frequencies—e.g., consumption data might come at a monthly or quarterly frequency. More importantly, data often come at a frequency that arguably reflects the net effects of a series of underlying decisions. In such cases, empirical analyses typically develop a model at the same frequency as the data—e.g., if the data come at the monthly frequency, then a period in the model is assumed to be one month. Such analyses ignore the underlying temporal aggregation.

Under exponential discounting, we suspect this issue doesn't matter much. Under present bias, in contrast, it could be quite important. For instance, suppose data come at the monthly frequency, and reflect the net behavior of 30 daily decisions each impacted by a small present bias. If one uses this data to estimate a model in which a period is a month, estimated impatience will be very large. Moreover, that estimated impatience would not predict well decisions on simple trade-offs between utility now versus utility one month from now. The right way to approach such data would be to explicitly model how a series of underlying decisions aggregate into predictions at the frequency of the data.

Question #3: How to Assess the Impact of Present Bias Against Other Phenomena?— Present bias is being incorporated into more and more analyses. However, the success of present bias has perhaps been to the detriment of other potential improvements to economic models of intertemporal choice. Indeed, economists are sometimes prone to misattribute behaviors to present bias that more likely are due to other shortcomings of the classical economics model.

Four intertemporal phenomena seem particularly relevant. First, there is the old idea of habit formation wherein one's utility from consumption depends on one's own past consumption. Second, there is projection bias wherein one's decisions are distorted by mis-

predictions of future tastes (Loewenstein, O'Donoghue, and Rabin 2003). Third, there is anticipatory utility wherein one experiences utility now from anticipating future consumption (Loewenstein 1987). Fourth, there is intertemporal "news" utility wherein one experiences gain-loss utility whenever one's behavior deviates from one's expectations (Kőszegi and Rabin 2009). In each case, we think there are examples where researchers have confused these phenomena with either present bias or refutations of present bias.[9] It is important to develop ways to tease these phenomena apart, and more generally to develop better intuitions for the types of environments in which each is likely to be important.

Question #4: How to Assess Whether Commitments Are Due to Present Bias?—A prominent prediction of sophisticated present bias is commitment. Economists understand this prediction well, and frequently point to observed commitments as indicative of present bias. We worry, however, that researchers are sometimes too quick to attribute any observed commitment to sophisticated present bias. There are reasons to be careful.

First, there are other reasons why people make commitments—e.g., models of belief-based utility, such as anticipatory utility or news utility, generate a motive to commit so as to influence beliefs and thereby utility. Second, some behaviors that appear to be commitments might in fact not be, but rather reflect incorrect beliefs about future behavior. For instance, suppose people purchase snacks frequently in small packages when they could have saved money by purchasing large packages. At first glance, this behavior might be seen as evidence of sophisticated present bias, where people purchase small packages so as to prevent overconsumption at home.[10] However, this behavior is equally consistent with people repeatedly purchasing small packages because they repeatedly predict that they won't want to consume much at home—e.g., due to naïvete about present bias or mispredicting the impact of hunger.

---

[9]For instance, intertemporal news utility can generate behavior that looks like present bias.

[10]Wertenbroch (1998) provides evidence on package sizes, and discusses the intuition of sophistication.

Third, some commonly discussed commitments actually have a negative price, and thus people might make such commitments even if they don't value the commitment. Indeed, Laibson (1997) and subsequent work emphasize how sophisticated present bias can lead people to invest heavily in illiquid assets as a commitment device. In real-world markets, however, illiquid assets pay a premium, and thus, for instance, even people with naïve present bias might choose to invest heavily—because they expect not to want to consume their savings for quite some time.

IV. Concluding Thoughts

We hope we've reached a point where present bias will be treated like other mature hypotheses—like basic discounting per se, or risk aversion, or a preference for variety. For any particular application, researchers should use their best judgment for whether present bias is potentially at play, and assess scientifically the extent to which it impacts economic outcomes.

Some might worry that present bias is more complicated than exponential discounting, and thus lead to more complicated analyses. But the longer run "general-equilibrium" effect of using present bias may be to simplify economics.[11] To illustrate, suppose present bias is in fact influencing outcomes across a range of applications. For any particular application, researchers who assume exponential discounting can almost surely account for the observed impact of present bias by adding some extra assumptions about, e.g., liquidity constraints or transactions costs. When realistic, such assumptions enhance our models. But when contrived solely to avoid present bias, these "patches" will leave economists with a more complicated and less coherent view of behavior than if we incorporate realistic levels of present bias.

---

[11] The logic here parallels closely the logic from Camerer (2000) with regard to prospect theory.

REFERENCES

Ainslie, George W. 1992. *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person.* Cambridge, UK: Cambridge University Press.

Akerlof, George A. 1991. "Procrastination and Obedience." *American Economic Review* 81 (2): 1–19.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation." *Journal of Economic Perspectives* 15 (3): 47–68.

Augenblick, Ned, Muriel Niederle, and Charles Sprenger. 2013. "Working Over Time: Dynamic Inconsistency in Real Effort Tasks." National Bureau of Economic Research Working Papers 18734.

Camerer, Colin. 2000. "Prospect theory in the wild: Evidence from the field." In *Choices, Values, and Frames*, edited by Daniel Kahneman and Amos Tversky, 288–300. New York: Russell Sage Foundation.

DellaVigna, Stefano, and M. Daniele Paserman. 2005. "Job Search and Impatience." *Journal of Labor Economics* 23 (3): 527–88.

Kőszegi, Botond, and Matthew Rabin. 2009. "Reference-Dependent Consumption Plans." *American Economic Review* 99 (3): 909–36.

Laibson, David. 1994. "Hyperbolic Discounting and Consumption." PhD diss. Massachusetts Institute of Technology.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–77.

Loewenstein, George. 1987. "Anticipation and the Valuation of Delayed Consumption." *The Economic Journal* 97 (387): 666–84.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118 (4): 1209–48.

O'Donoghue, Ted, and Matthew Rabin. 1999a. "Doing It Now or Later." *American Economic Review* 89 (1): 103–24.

O'Donoghue, Ted, and Matthew Rabin. 1999b. "Procrastination in Preparing for Retirement." In *Behavioral Dimensions of Retirement Economics*, edited by Henry J. Aaron, 125–56. Washington, DC: Brookings Institution Press.

O'Donoghue, Ted, and Matthew Rabin. 2001. "Choice and Procrastination." *Quarterly Journal of Economics* 116 (1): 121–60.

O'Donoghue, Ted, and Matthew Rabin. 2006. "Optimal Sin Taxes." *Journal of Public Economics* 90 (10–11): 1825–49.

Phelps, E. S., and Robert A. Pollak. 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35 (2): 185–99.

Pollak, Robert A. 1968. "Consistent Planning." *Review of Economic Studies* 35 (2): 201–8.

Strotz, R. H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23 (3): 165–80.

Wertenbroch, Klaus. 1998. "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice." *Marketing Science* 17 (4): 317–37.