

Behavioral Economics
Peter Diamond¹

What has been happening in behavioral economics, broadly interpreted, has been very exciting. And I expect the excitement to continue.² The papers in this special issue cover a range of topics and approaches in behavioral economics, including issues of individual behavior, of the measurement of the consequences of individual behavior, and of incorporating these pictures of individual behavior into normative analysis. They are very well worth reading.

There is a wide range of behavioral economics papers generally. Some papers pursue better understanding of how individuals decide, particularly about economic transactions. Others pursue the properties of economic equilibrium given a model of how people decide (possibly with heterogeneity in individual behavior) and a model of the underlying economic environment. Others pursue policy implications. Some of these are in areas where there must be policy (for example, taxation) and ask how the insights from behavioral economics might change existing policy analyses. Others open up new areas where government intervention might be successful (for example legislating defaults in 401k plans). With the latter, one should think whether encouraging a new realm of regulation is likely to result in a successful outcome. While a similar question may arise in the former (e. g., a new type of tax expenditure), often the issue is trying to get a legislature to do better what it will do in any event, such as setting tax rates on existing tax bases.

Learning from studies in this area, as from studies in all of economics, should come from drawing inferences as a result of considering the applicability of the findings of a study beyond its theoretical model or beyond its econometric formulation and estimation. Policy advice should not come from taking formal implications literally. It seems to me, as it seemed to Alfred Marshall, that the research process involves both formal analyses and informal combining of various insights:

it [is] necessary for man with his limited powers to go step by step; breaking up a complex question, studying one bit at a time, and at last combining his partial solutions into a more or less complete solution of the whole riddle. ... The more the issue is thus narrowed, the more exactly can it be handled: but also the less closely does it correspond to real life. Each exact and firm handling of a narrow

¹ I did not attend the conference, but read the papers with interest after being invited to submit a comment for this special issue. I am grateful for comments on an earlier draft from Henry Aaron, Alan Auerbach, Nick Barr, Angus Deaton, Bengt Holmstrom, Richard Layard, George Loewenstein, Antonio Rangel, Johannes, Spinnewijn, Nick Stern, and Richard Thaler.

² It is that belief, and the hope of contributing to its development, that led me to co-organize the conference that became Diamond and Vartiainen (2007).

issue, however, helps towards treating broader issues, in which that narrow issue is contained, more exactly than would otherwise have been possible. With each step ... exact discussions can be made less abstract, realistic discussions can be made less inexact than was possible at an earlier stage. (Marshall, 1948, page 366.)

I do not think we have “a more or less complete solution of the whole riddle” that behavioral economics addresses, or pretty much any broad issue that economics addresses. But policy making, and so policy recommendations, can not wait for a complete solution. Moving ahead with what we think we know, and moving in recognition of the limitations of how much we do know is appropriate. And doing so is better than making policy while ignoring the understanding of economics we do have.

For continuing and improved success, behavioral economics needs to preserve some aspects of traditional economic analyses and avoid others. For example, with survey questionnaires, which economists have a long history of using, behavioral analyses need to continue to recognize that answers often have interesting content even though they are not necessarily accurate, exact answers to the questions the researchers (or the readers of the research report) had in mind.³ This holds for questions about facts (e. g., how much wealth do you have), questions about beliefs (e. g., what is the probability you will live to 85), and questions about condition (e. g., how happy are you).

Economists are more aware of the importance of the constraints coming from equilibrium than typical noneconomists, an awareness that matters for behavioral economics as well. Awareness of equilibrium (sometimes going from unintended consequences to undesired consequences that can be part of an evaluation) is an important part of thinking about both positive and normative dimensions of outcomes. Constraints from equilibrium come in a variety of forms. Individuals have budget constraints – there can not be a total absence of quantity responses to a single price change of a good being purchased, and some policies to lower consumption today (e. g., of positional goods) may increase their consumption tomorrow. Governments also have budget constraints – today’s deficits have real implications, for taxes and/or spending in the future. Economic agents react to changes in the economic environment (e. g., rules on disclosure) and agents react to the changes in the behavior of others. It is important to preserve awareness of equilibrium responses when considering policy changes by governments or other institutions or individuals. Behavioral analyses need to be as aware of the roles of equilibrium constraints as non-behavioral analyses.

While these common elements in economic analysis have served the profession well and should be continued, there are other aspects of positive analysis that need to be adapted to work well with the insights of behavioral economics. And beyond the changes that behavioral analysis brings to descriptions of the determination of equilibrium, there are new challenges and new opportunities for doing normative evaluations of equilibrium.

³ As Beshears, Choi, Laibson and Madrian (this issue) put it: “We agree that self-reports can’t be taken at face value, but we also believe that they should not be ignored completely.”

Context/Situation

In standard modeling, we assume consistent behavior across economic environments, captured in preferences that are defined only in terms of commodities acquired (absent externalities). In addition, these same preferences are generally taken to be a suitable basis for normative analysis in many settings, although this common approach has not been applied universally.⁴ One of the key messages of behavioral economics is that context (also referred to as situation) matters in ways that are not recognized in standard modeling.⁵ Purchasing context (e. g., type of store) matters for willingness to pay to acquire a particular good.⁶ Context (anticipated use by others) matters for the demand for positional goods.⁷ Context (presence of incentives that are not narrowly self-oriented) matters for the role of financial incentives.⁸ Context matters for the tradeoff between current and future concerns (e. g., trading off chocolate cake today for higher weight tomorrow is different if the cake is in front of you or merely nearby). Context, as captured in the concept of mental accounts, matters for spending.⁹ Eventually there may be enough understanding of how context matters to do well in moving inferences from one setting to another. For now, carrying estimated parameters (e. g., β from quasi-hyperbolic modeling) unchanged from one context to another is not likely to be a good estimate. That is, the description of the setting where an economic decision is being made needs to include far more information than the usual description of the setting. And behavior is likely to vary with elements of the description that play no role in the usual formulation.¹⁰

Normative Evaluations

Using behavioral insights well for policy purposes will often be hard. My message is not that these policy issues are unimportant or not worth working hard on, but that much more hard work is warranted. Where the standard model of individual choice is an adequate basis for positive analysis, it often lends itself to normative analysis since, on many issues, there is no good reason for not incorporating the individual preferences into normative analysis, either of Pareto optimality or with a social welfare function. But

⁴ The concept of merit goods (Musgrave, 1959) is a long-standing exception. So too are analyses based on presumptions of behaviors that do not succeed in being individually optimal, particularly in considering mandatory pension systems. And governments have instituted a variety of measures of consumer protection that draw on the same underlying ideas that are now making themselves felt in economic analyses, such measures as cooling-off periods after purchases from door-to-door salesmen or the banning of some consumer products. While regulation of product design can have multiple underlying sources, some of them fit this category.

⁵ Koszegi and Rabin (this issue) discuss preferences depending on the context of the choice set.

⁶ Consider beer purchase in the experiment in Thaler, 1985.

⁷ As in positional goods discussed by Frank (this issue).

⁸ As in the studies discussed by Bowles (this issue). It matters whether a circumstance is viewed as an opportunity for generosity or as a possibility for being made a sucker.

⁹ Shefrin and Thaler, 1988

¹⁰ As an example of context mattering, consider this spoof response to an actual study (Plassman et al., 2008) from ironictimes.com:

▶ **Study: Consumers Get More Pleasure From Same Wine at Higher Price**
Oil companies raise gas prices to enhance drivers' enjoyment.

when the standard model is too far from accurate, and we have empirically supported alternative models, we should consider policy recommendations based on these alternatives. As with standard analysis, going from consideration of policy to a recommendation requires a positive evaluation of what the political process will do with such a recommendation.

One realm of behavioral analysis is the identification of circumstances where people are making “mistakes.”¹¹ This is a wonderful label since it suggests that welfare improvements, possibly even Pareto improvements, are possible and simultaneously suggests an approach to normative analysis, by assessing outcomes as individuals would if they were not making mistakes. In order to proceed with normative analysis and policy recommendations, one needs to identify mistakes and needs to evaluate alternative outcomes that the policy may mandate or induce. Consideration of outcomes needs to include changes in the behavior of others as equilibrium changes, as noted above. Support for a conclusion of mistake from analysis of the decision process is valuable and would be reassuring. Nevertheless, we should not rule out reaching a conclusion of mistake from simply considering the outcome for the individual. Sometimes people change their behavior when an alleged mistake is pointed out to them. But people can also be convinced to change their behavior in settings that might not be mistakes. So responses to advice, along with reluctance to decide, may be helpful in identifying mistakes but are not necessarily conclusive. Common sense may suggest that a mistake is a far more likely interpretation than tastes that seem unusual or peculiar in an attempt to claim a standard revealed preference interpretation. And further tests may strongly support an interpretation of mistake, e.g., experience utility measurement.¹² The historic weight given by the economics profession to revealed preference is not a good reason for rejecting all claims of mistake. General skepticism about the multiple political sources of government interventions should make one cautious about labeling mistakes. That is, perceptions of plausible political use of possible policy recommendations based on behavioral economics findings should play a role when deciding whether to push a recommendation. Sometimes it will seem likely that good use will be made of a recommendation and sometimes not.

Mistake probably can not be “proved” in many interesting settings. But then the absence of mistake can not be proved either. Contrasting models of addiction are an example where plausibility can play a role in deciding on a basis for policy recommendations. And mistake seems to me a more plausible interpretation in some addiction settings. Analysts need to conclude what interpretation is most likely, while recognizing the possibility of an erroneous interpretation. Not every finding of the presence of mistakes should lead to political recommendations.

Having recognized widespread mistakes, consideration of equilibrium responses (the full set of changes as a consequence of an intervention) and additional concerns (e. g.,

¹¹ Koszegi and Rabin (this issue) and Beshears, Choi, Laibson and Madrian (this issue) discuss how to identify mistakes and how to construct normative preferences.

¹² For example, when actions are based on memories of experiences that have systematic biases relative to contemporaneous evaluations of experiences (Kahneman, et al., 1993).

individual heterogeneity, the value of choice per se) can lead to mandates, incentives, regulations, giving advice or providing information. In the realm of retirement income, we have mandatory systems of payments from individuals when younger and to them when older; we have special tax rules for saving that is not spent before some age; we have regulation of private retirement savings plans; we have advice on how much to save; and we have information on the implications of alternative savings plans. And we fruitfully have all five at the same time. Mandatory and voluntary (tax-favored) retirement savings plans have restrictions on portfolio choice and on the form of payout with a variety of degrees of restriction, including all of the list above - mandates, tax incentives, defaults, advice and information.¹³ All of these require an evaluation of induced changes in behavior if the design is to be evaluated well. This even holds for simply providing information as Loewenstein and Ubel (this issue) note: “people who learn first about the risks of a treatment followed by its benefits make different choices than people who first learn about its benefits and then its risks. Decision aid developers have no choice but to present information in one order or another, but unfortunately the order they choose will almost inevitably affect people’s decisions.”

A second realm of normative behavioral economics arises with concerns for others, which may involve non-market interactions with others. A prime example of this is in charitable donations and the use of income and estate tax incentives to encourage donations. How to form a social welfare function from a given model of individual behavior is open to debate.¹⁴

And a third realm arises with the goal of changing preferences, as Stern, Dethier, and Rogers (2005) argue is the heart of economic development. “Changing preferences” is a term with several meanings and different uses, resulting in some confusion. I explore three examples in terms of choices. For all three examples, assume that all four of A and B at both dates cost the same in present discounted value.

- (1) A might be chosen over B and consumed at time t , while B is chosen over A and consumed at time $t+1$, even though the situations are similar (e. g., identical relative prices). One might describe this as preferences that changed between t and $t+1$.
- (2) If A is consumed at time t , A will be chosen over B and consumed at time $t+1$. Yet, if B is consumed at time t , B will be chosen over A and consumed at time $t+1$. And if both are available at time t , A is chosen. One might describe this as consumption at time t changing preferences at time $t+1$.
- (3) At time t , A might be chosen over B for consumption at time $t+1$, while if the choice is made at time $t+1$, B would be chosen over A for consumption at time $t+1$. One might describe this as preferences that changed between t and $t+1$.

All three examples could be consistent with Arrow-Debreu preference modeling or could be the result of a mistake, with very different normative implications. Arrow-Debreu

¹³ An example in the realm of benefit form is the US requirement of notarized agreement by a spouse if a single life annuity is chosen from a pension plan.

¹⁴ See, e. g., Diamond (2006).

theory is defined in terms of the vector of consumptions at all times and states of nature. In (1), with all sequences costing the same, an Arrow-Debreu consumer might simply prefer the sequence {A at time t , B at time $t+1$ } to any of the other three sequences. This might be the result of declining marginal utility in total consumption of A or preferences that are different at different ages (as they surely are). Such preferences are different over time, but not “changing” in a sense that matters for normative analysis. Alternatively, as a behavioral interpretation, it may be that the consumer has trouble deciding and the context influenced the choices, for example, A might have been the default at time t and B at time $t+1$, and possibly A (or B) was the better choice at both times.

In (2), the choice in the first period might be made with full awareness of the implications for preferences in time $t+1$ (as in Arrow-Debreu theory). Different choices at $t+1$ might come from the absence of intertemporal additivity in preferences, while the preferences are acted upon in a time-consistent way. An individual will do more of activities that have become more enjoyable from previous experience (e. g., getting good at tennis or the clarinet) and less of those that become less interesting (or even tiresome) with more use, a form of diminishing marginal utility when viewing consumption aggregated over time. With individual awareness of such intertemporal links, there is no need for standard theory to be used any differently because of non-additivity of preferences. Alternatively, the choice at time t might be made with a myopic ignoring of future implications for preferences or a more general failure to fully appreciate what the future will be like. Indeed, some decisions, such as going to university or having a child, are known to affect future preferences, but in ways that may not be fully appreciated by the decisionmaker.

For (3), which is the prime example for quasi-hyperbolic modeling, to be consistent with Arrow-Debreu theory, some circumstance needs to change between t and $t+1$, which might be the arrival of information about which state of nature is occurring. For fully complete markets (no asymmetric information), with complete Arrow-Debreu markets one would be able to make purchases conditional on the information set at different future dates. But markets are not complete.

Part of the confusion in the term changing preferences comes from the different basic modeling of preferences in standard economics and in some behavioral economics. In Arrow-Debreu theory, preferences (the basis of choice) are defined relative to goods at different dates. In some behavioral economics, preferences (the basis of choice) are defined relative to the distance over time from the time at which a decision is being made to when the consumption actually occurs.¹⁵ This difference lies behind some confusion about the role of exponential discounting in time-consistent choice. In Arrow-Debreu theory, there is no need for exponential discounting for time-consistent behavior since the discounting applies to dates. Indeed, without intertemporal separability, the preferences may not lend themselves to the existence of something readily interpreted as a discount factor. In modeling quasi-hyperbolic preferences, discounting is applied in the same way

¹⁵ In Arrow-Debreu theory all choices are made once and for all at the start of time. But that is not how we experience purchasing.

to the “future” as measured from different decision dates, implying that we need exponential discounting in order to avoid time-inconsistent behavior.

Individual choices resulting in effects spread over time may be particularly subject to mistake. Beyond this, and at the heart of the Stern, Dethier, and Rogers view, it may be good to change individual preferences in ways the individuals do not contemplate. Governments have a range of ways of affecting choices, as discussed above, and suitable responses will be different in different settings. Similarly, norms matter for behavior and norms can be influenced by governments. Normative evaluation of government interventions in these settings requires evaluating people with different (changed) preferences as a consequence of a policy intervention. And that requires moving beyond traditional normative analyses that respect revealed preference.

In sum, analyzing how to make good use of positive behavioral economics for policy purposes is a major and important ongoing topic.

Happiness¹⁶

Three of the papers in this issue relate to happiness, including analysis of survey questions about happiness in general (Layard, Mayraz and Nickell¹⁷), analysis of happiness moment by moment (Krueger and Schkade¹⁸) and discussion of how to use the content of these sources of data for normative analysis (Loewenstein and Ubel). My beliefs about the different literatures echo what I said at the start – studies like these contain real information, but even after adjusting for measurement issues, as the papers do, the answers to survey questions should not be taken literally, but used for plausible inferences.

How should we interpret answers to the question “How happy are you these days?” Or answers to questions asking about life satisfaction? (These may tap somewhat different underlying perspectives, as one might be satisfied with a life devoted to an activity that does not generate happiness; probably these should not be treated as interchangeable questions.) The papers consider sources of noise and error in attempts to answer the question the investigators have in mind. But we should consider that the respondents may have something related but somewhat different in mind. If people answer whether they are satisfied with their lives in terms of their perceived relative position in happiness, that does not necessarily mean that happiness is based on relative position, rather than that the question being answered by the respondent is a relative happiness question. And if it is a relative comparison that at least partially lies behind the answer, then different people may be using different bases for comparison – their own past, their

¹⁶ I have been only a casual follower of the happiness literatures. My willingness to write this comment, nevertheless, came from seeing the strong similarity to my basic methodology in Loewenstein and Ubel (this issue).

¹⁷ This paper focuses on answers to the questions “Taking all things into account, how happy are you these days?” and “All things considered, how satisfied are you with your life as a whole these days?”

¹⁸ This paper considers affect measures collected from a diary method. It also includes a life satisfaction question in its surveys.

perception of some set of others, possibly a particular peer group, or people locally, nationally, somewhat internationally, or their own prior expectations. And this may lie behind the lack of consistency in time-series and cross-section relations of happiness answers relative to income.¹⁹ Consider asking people how tall or how old they are on a scale from zero to ten. Would cross sections and time series of repeated surveys match respondents' measured heights and ages consistently? I doubt it. I suspect that for some purposes we would want to use their answers (e. g., behavior that relates to self-image) and for some their measurements (e.g., the effects of nutrition on height and mortality). Some exploration has been done of the impact on reported happiness of the growth of incomes and the incomes of neighbors. But such studies may not shed light on the question of how much well-being depends on one's relative standing and how much the respondent looks to relative standing in order to answer the survey question.

A happiness answer reflects some period of time – this instant in some types of studies, an intent to have a wider time frame in others. And different respondents may be implicitly reflecting different time frames, perhaps because they are considering different elements. For example, having just cut consumption to take advantage of a particular retirement savings opportunity, one might answer a happiness question based on current reduced consumption or on the satisfaction from having provided better for the future. If one were to try to use happiness measures for retirement savings policy, one would need to know what is underlying an answer at a particular point in time. And, as always, correlation does not necessarily imply causation. Given a large individual (personality?) component in happiness, one can wonder how much happier people earn more or save more (for higher income later) as well as vice versa.

Experience utility studies have an obvious role in looking for mistakes.²⁰ And happiness studies may play a role in normative evaluations given the ordinal nature of standard preferences.²¹ Even if one were to accept some particular answer relating happiness to income, that still leaves several questions.²² Is measured happiness the appropriate basis

¹⁹ That time series and cross-section relations of answers relative to income do not match is a familiar pattern in economics, most visibly in studies of savings.

²⁰ As noted above, in some settings memories of the utility of an experience have systematic biases relative to evaluations of utility collected during the experience (Kahneman, et al., 1993).

²¹ It is common to assert that choice under uncertainty gives some degree of cardinality (up to an affine transformation). But is the “utility” in expected utility the same as the utility we seek for normative purposes? I think this is a semantic sleight of hand. Consistent with maximizing expected utility is maximizing any increasing function of expected utility. But converting that formulation into choice under certainty gives a preference known up to an increasing transform not restricted to being an affine transform. That is, drawing the common inference involves an assumption that “utilities” in different states of nature are the utilities we are concerned with, and not some transform of them. Similarly, if lifetime utility of consumption today and consumption tomorrow for person n is $u_1[x_1] + \delta_n u_2[x_2]$, is this the function that should be added across people by a utilitarian or, the ordinally equivalent, $u_1[x_1]/\delta_n + u_2[x_2]$, or something else?

²² Layard, Mayraz and Nickell focus on estimating a single CRRA parameter. They find that log-linear fits fairly well except at the ends of the income distribution, which may or may not reflect measurement issues. For some policies it is the ends that are particularly important (poverty relief, taxation of very high incomes).

for normative evaluations? If happiness is measured on a scale of zero to ten, what transform of these answers is the right measurement of happiness for a utilitarian to aggregate? Since happiness measures for an individual will vary over time, how should these be aggregated to obtain lifetime happiness? And there remains the issue of the extent to which concerns about inequality in happiness should enter a social welfare function (if at all). That is, would we want to simply add up happiness measures across the population or consider a concave function of them?

Using happiness measures for normative purposes will involve other assumptions. Layard, Mayraz and Nickell note that they make the assumption that reported happiness depends on other variables only via the dependence of true utility on those variables. While some assumption is needed, it is not apparent that this is the only assumption worth considering. As shown in some experiments, answers are influenced by recent events and current circumstances that are not particularly relevant for a life satisfaction measure. One could explore empirically the influence of recent events (e. g., income growth, promotion, sunshine on the day interviewed) on reported happiness and consider the difference in implications as fractions of that effect are attributed to a difference between reported and true utility. Layard, Mayraz and Nickell also note the tensions across different elicitation methods, a tension that needs resolution for more confidence in using any one of them.

Loewenstein and Ubel (this issue) question the identification of happiness with the base for normative evaluation, a questioning with which I fully agree:

The main problem with experience utility is its failure to incorporate non-hedonic aspects of experience, such as meaning and capabilities (even if such capabilities are not used) that are important to people but have little impact on their happiness.

Behavioral research is showing us very interesting things about human behavior. In addition to continuing to look for interesting insights, we need to learn more about how to use what we are learning. I conclude as I began – what has been happening in behavioral economics is very exciting and the papers in this special issue are very well worth reading.

References

Diamond, Peter. 2006. "Optimal Tax Treatment of Private Contributions for Public Goods with and without Warm Glow Preferences", *Journal of Public Economics*, vol. 90, no. 4-5, May, pp. 897-919.

Diamond, Peter, and Hannu Vartiainen. 2007. *Behavioral Economics and Its Applications*, Princeton: Princeton University Press.

Kahneman, D., Fredrickson, D. L., Schreiber, C. A., & Redelmeier, D. A. 1993. When more pain is preferred to less: Adding a better end. *Psychological Science*, 4, 401-405.

Musgrave, Richard A. 1959. *The Theory of Public Finance: A Study in Political Economy*. New York: McGraw-Hill.

Marshall, Alfred. 1948. *Principles of Economics*, 8th ed. New York: Macmillan.

Plassman, Hilke, John O'Doherty, Baba Shiv, and Antonio Rangel. 2008. Marketing actions can modulate neural representations of experienced pleasantness. *Proceedings of the National Academy of Sciences*, 105, 3, (January 22) 1050-1054.

Shefrin, Hersh, M., and Richard H. Thaler. 1988. "The behavioral life-cycle hypothesis." *Economic Inquiry*, 26, 609-643.

Stern, Nicholas, Jean-Jacques Dethier, and F. Halsey Rogers. 2005. *Growth and Empowerment: Making Development Happen*, Cambridge: MIT Press.

Thaler, Richard H. 1985. "Mental Accounting and Consumer Choice." *Marketing Science*, Vol. 4, No. 3 (Summer), pp. 199-214.