

Money makes the world go round, and basic research can help

Ido Erev*

Abstract

As the adage goes, “money makes the world go round” – but which direction does it spin? This analysis considers how basic decision research can help us work out how to answer this question. It suggests that the difficulty of deriving clear predictions based on existing decision research is at least partly rooted in two restrictive conventions. The first is the focus on deviations from rational choice, and the effort to capture observed deviations by assuming subjective value functions. While it is difficult to reject the hypothesis that choice behavior reflects the weighting of subjective values, it is not clear that it advances the derivation of useful predictions. A second restrictive convention is the focus on objective hypothesis testing, which favors analyses that evaluate small refinements of the popular models. The potential benefits of relaxing these conventions are considered, with reference to recent choice prediction competitions that facilitate the exploration of distinct assumptions and model development techniques. The winners in these competitions assume very different decision processes than those assumed by the popular “subjective functions” models. The relationship of the results to the big data revolution is discussed.

Keywords: methodology, subjective value, model development, prediction competitions

1 Introduction

“Love of money,” says the New Testament, “is the root of all evil” (1 Timothy 6:10). Perhaps . . . or perhaps, as per the quotation variously attributed to Mark Twain and George Bernard Shaw, “lack of money is the root of all evil.” How is it that these statements are contradictory, yet both ring true? Or to put it another way, everybody knows that “money makes the world go round” – but what is the direction of spin?

Part of the difficulty in determining the effect of money is that this effect can be highly negative in certain settings, but positive in others. In terms of the “love of/lack of” conundrum, it appears that focusing on small sets of different environments led different thinkers to reach different conclusions. This logic highlights the importance of research that allows useful predictions of the impact of money.

Most efforts to clarify the impact of money start with different variants of the rationality assumption. Money is abstracted, under this approach, as a medium that facilitates trading and exchange (this abstraction is consistent with the behavioral interpretation of money as a secondary reinforcement, and also underlies the current analysis). The basic

rational model implies a simple effect of money: when facing a choice between two or more options, decision makers choose the one that maximizes their expected monetary gain (given their rational beliefs concerning the incentive structure and the behavior of other relevant agents). Behavioral studies demonstrate that this basic model is too simple, and highlight the value of more complex models that add subjective utilities to the basic rational processes. However, progress in this field is rather slow. Despite decades of research, the leading models still cannot clearly predict the impact of money even when the analysis is restricted to the best-studied setting: a choice between gambles in controlled experiments. The current paper aims to both illuminate and address the difficulties in advancing positive models of economic behavior. It suggests that these difficulties are partly the product of problematic working assumptions and research methods that direct behavioral decision research to a local maximum.

Figure 1 presents a topographical analogy that illustrates the main assertion of the current paper. The right-hand side abstracts the “land of problems”. Money has a negative effect in some of these problems, and a positive effect in others. The left-hand side abstracts the “land of assumptions”. Developing models, under the current analogy, amounts to selecting a point in the land of assumptions. A model provides useful predictions of behavior within its field of view. In the figure, greater elevation indicates greater descriptive and predictive value. Importantly, however, scientists do not know the topography of the land of assumptions perfectly well; nor do they have perfect knowledge of their current model’s precise field of view. They can only try different sets of assumptions, and test their value. This essay argues that the development processes and evaluation methods which underlie popular

This paper is based on Ido Erev’s EADM (European Association for Decision Making) presidential address presented at SPUDM 2019 (the biennial meeting of EADM). The author thanks Ori Plonsky and Yefim Roth for useful comments, and Meira Ben-Gad for useful comments and editorial assistance. This research was supported by grant number 535/17 from the Israel Science Foundation.

Copyright: © 2020. The author licenses this article under the terms of the Creative Commons Attribution 3.0 License.

*Faculty of Industrial Engineering and Management, Technion – Israel Institute of Technology. Email: erevido@gmail.com. ORCID 0000-0001-9889-4070.




FIGURE 1: The topographical analogy: The left-hand side abstracts the land of assumptions. The right-hand side abstracts the land of problems. Points on the left (e.g., EV) denote models, and points on the right (e.g., St. Petersburg) denote choice problems (these models and problems are described in the text below). Models provide useful predictions for the problems in their field of view.

models have often directed behavioral decision scientists to hills of low elevation.

2 The impact of money and the value of rational models

Simple analyses of social interactions illustrate how rationality-based abstractions can highlight the impact of money. Analyses of this type point to both negative and positive effects. Two examples are presented below.

2.1 Example 1: The tragedy of the commons and the value of taxation

A clear demonstration of the negative effect of money is found in the well-known “tragedy of the commons” – Hardin’s (1968, based on Lloyd’s, 1833) analysis of the impact of incentives on shared resources. Hardin describes a community of N herdsmen who share a pasture. The pasture is large enough to support all the herdsmen with their current herds, and the current herds provide all the milk, meat and wool the herdsmen want to consume. A herdsman who adds an animal can later sell it (to people outside the community of herdsmen), accruing an expected gain of 1 monetary unit. However, each animal above the maximum supported by the pasture reduces the grazing area available to all the others, meaning that all the herdsmen together share an expected loss of more than 1. When the total expected loss is less than N , money-loving (rational) herdsmen are predicted to expand their herds – behavior that eventually will lead to collapse of the shared resource, potentially driving the community to extinction. Activity that pollutes a shared resource while benefiting the individual is another example of such behavior.

However, Hardin also suggests that while love of money triggers the tragedy of the commons, love of money is also the key to an easy solution. For example, the community can agree on a taxation system which reduces the benefits from behaviors that impair social welfare (e.g., through overuse or polluting a shared resource).

2.2 Example 2: Land protection

For a clear example of the positive potential of money, consider the following problems:

Problem 1: Two farmers living in a small valley divided by a stream consider joining together to build a temporary dam so as to protect their land during the approaching rainy season. If they do not cooperate, Farmer Left (who owns the land on the left bank of the stream) can expect to protect 60 units of land, and Farmer Right can expect to protect 40 units. If they do cooperate, they must choose between three possible locations for the dam. The three locations differ with respect to the number of land units protected for each farmer, as follows (protected for Left, protected for Right):

A:(50, 250) B:(180, 20) C: (70, 50).

Assume that crossing the stream is difficult, and the farmers cannot trade land. What is their joint decision?

Problem 2: Now consider a variant of Problem 1 with the same options, but with the understanding that each land unit is worth one monetary unit, and the agreement can include monetary transfer. What then is the farmers’ joint decision?

The difference between Problem 1 and 2 demonstrates how money can improve social efficiency. It increases the joint payoff from 120 to 300. The logic is simple: Without money (Problem 1), Left rejects (50,250) because the number of units Left would protect is below his/her outside option, and Right rejects (180, 20) for the same reason. Thus, the agreement is (70,50) and the number of protected units is 120. With money (Problem 2), bargaining moves the agents towards the (50, 250) option, with an agreement to transfer money (typically 100) from Right to Left. Thus, money (allowing compensation for the loss of land) increases efficiency (more land is protected) and can enhance fairness.

3 Deviations from maximization, subjective functions, and the impact of experience

The derivation of the impact of money in the examples presented above is based on the most basic rational model. It assumes that people maximize the expected value (EV) of an act or agreement. However, while this basic rational

model provides useful insights in these examples, it is not always accurate. The best-known violation of this model is the St. Petersburg paradox (Bernoulli, 1738/1954), presented below.

The St. Petersburg paradox: A fair coin will be flipped until it comes up heads. The number of flips is denoted by the letter k . The casino promises to pay each winning gambler 2^k monetary units. What is the maximum you are willing to pay in order to play this game one time?

While the EV from playing this game is infinite,¹ most people are not willing to pay more than 10 monetary units. Bernoulli proposed expected utility theory (EUT) to capture this finding. Under EUT people seek to maximize the expected utility of an act, and the utility of money is a concave function. This assumption implies that the utilities (subjective values of the possible outcomes) are weighted by their objective probabilities. As commonly understood, EUT adds a risk aversion parameter to the EV rule. This addition can be described as a move from point EV to point EUT in Figure 1’s land of assumptions. Follow-up research (von Neumann & Morgenstern, 1944; Savage, 1954) clarifies this move by explicitly presenting the assumptions that underlie EUT.

3.1 Prospect theory (PT)

An effort to capture choice behavior with EUT reveals several violations of this model as it is commonly understood. The clearest examples are the Allais paradox (Allais, 1953) and the coexistence of gambling and insurance (Friedman & Savage, 1948). Kahneman and Tversky (1979) replicated these violations by studying the problems presented in Table 1, and proposed a generalization of EUT known as prospect theory (PT) that can capture them (see also Tversky & Kahneman, 1992; Wakker, 2010). The key assumption behind the generalization implied by PT states that people select the option with the highest weighted subjective value, and the subjective values are weighted by a subjective function of their probabilities. The other assumptions focus on the shape of the relevant subjective functions: the “value function” and the “weighting function”. The value function assumes sensitivity to a reference point; losses relative to the reference point loom larger than gains. The weighting function assumes oversensitivity to low-probability extreme outcomes.

3.2 The impact of experience

While PT provides effective predictions for one-shot decisions under risk, efforts to use it to capture behavior in repeated settings reveal mixed results. Several studies have shown that experience can reverse the deviations from EUT

¹The EV = $2(1/2) + 4(1/4) + 8(1/8) + 16(1/16) + \dots = 1 + 1 + 1 + 1 + \dots \rightarrow \infty$

TABLE 1: Two of the violations of EUT replicated by Kahneman & Tversky (1979).

Phenomenon	Problems	R-rate
1. The Allais paradox	R1: 4000, .8; 0 otherwise	.20
	S1: 3000 with certainty	
	R2: 4000, .20; 0 otherwise	.65
	S2: 3000, .25; 0 otherwise	
2. Preference for lotteries and insurance (overweighting of rare events)	R3: 5000, .001; 0 otherwise	.72
	S3: 5 with certainty	
	R4: -5000, .001; 0 otherwise	.17
	S4: -5 with certainty	

Note: Data from Kahneman and Tversky (1979). Participants were asked to choose once between the two prospects in each problem, based on complete descriptions of the prospects. The payoffs were hypothetical. The modal choice pattern in the Allais problems (S1 and R2) violates EUT, as this theory predicts R2 if and only if R1 is selected. A preference for both gambling and insurance violates the prevailing interpretation of EUT, which assumes a fixed risk-aversion attitude.

captured by PT (Barron & Erev, 2003; Hertwig et al., 2004). Indeed, feedback has been found to reverse deviations from rational choice even when it does not add information. For example, consider Erev et al.’s (2017) experimental study of the four conditions summarized in Table 2. The results reveal an initial tendency to overweight rare events (in accordance with the predictions of PT), and a reversal of this pattern (i.e., underweighting of rare events) after several trials with feedback. The difference between the initial reaction to description and the reaction to experience is known as the description–experience gap (Hertwig & Erev, 2009).

To clarify the significance of the description–experience gap, consider the effort to avert the tragedy of the commons, described above, by taxing animals beyond the number supported by the pasture. Since collecting tax is costly, the most effective taxation under PT involves large fines for a relatively small (but not too small) proportion of the additional animals; for example, under the median parameters estimated by Tversky and Kahneman (1992), a tax of 100 on 1% of the additional beasts should suffice. The observation that experience reduces the weighting of low-probability events suggests that this solution might work initially, even when it is likely to fail in the long term.

TABLE 2: Two of the effects of feedback on decisions under risk documented by Erev et al. (2017).

Phenomenon	Problems	R-rate				
		Block: 1 (no FB)	2	3	4	5
The Allais paradox and its elimination with experience						
	R5: 4, .8; 0 otherwise					
	S5: 3 with certainty	.42	.57	.57	.60	.65
	R6: 4, .20; 0 otherwise					
	S6: 3, .25; 0 otherwise	.61	.62	.62	.64	.62
Over- and underweighting of rare events						
	R7: 101, .01; 1 otherwise	.55	.45	.43	.42	.42
	S7: 2 with certainty					
	R8: -20, .05; 0 otherwise	.48	.63	.62	.62	.64
	S8: -1 with certainty					

Note: The participants faced each problem for 25 trials (5 blocks of 5 trials). They were presented with complete descriptions of the prospects, and received feedback starting in the sixth trial. The final payoffs (in Israeli shekels, 4 shekels = 1 euro) were determined by a randomly selected trial.

It is constructive to distinguish between two approaches to capturing the impact of experience on choice behavior. The first tries to improve the abstraction of the underlying subjective functions. For example, the impact of experience on decisions under risk, demonstrated in Table 2, can be captured by hypothesizing that experience affects the parameters of the subjective functions assumed by this theory (see related analyses in Abdellaoui, L'Haridon & Paraschiv, 2011; Glöckner, Hilbig, Henninger & Fiedler, 2016).

The second approach calls for a larger change in the abstraction of the underlying decision process by relaxing the assumption that deviations from maximization reflect the shape of subjective functions. To clarify this approach, note that under the current topographical analogy (Figure 1), EUT, PT and similar analyses aim to find the best model by focusing on an area of the land of assumptions that can be described as the “subjective functions” hill. The hypothesis that the observed impact of experience calls for a larger change in the underlying assumptions suggests that it may be possible to find hills of higher elevation that allow more useful predictions of behavior. This suggestion is implicit in B. F. Skinner’s (1985) critique of early behavioral decision research. Skinner refers to the outcomes of past experiences with similar decisions as “contingencies of reinforcement”, and writes: “Neglected contingencies of reinforcement can be subtle. Kahneman and Tversky (1984) have reported that people say they would be less likely to buy a second ticket to

the theatre if a first was lost than to buy a ticket after losing the money that was set aside for that purpose. The difference is said to be due to a difference in categorization. A difference in relevant contingencies should not be overlooked. A boy who usually washes his hands before sitting down to dinner quite justly protests when told to wash them if he has already done so” (Skinner, 1985, p. 297).

To clarify the difference between the models on the subjective functions hill, and the hills suggested by the contingencies of reinforcement approach, let us return to the St. Petersburg paradox. The models on the subjective functions hill assume that decision makers have in mind a hypothetical casino that can pay lucky winners with certainty. That is, these models ignore the fact that the prize on offer is potentially higher than could be paid by any real casino (see related argument in Tversky & Bar-Hillel, 1983), and they assume that the low valuation of the bet (typically below 10) reflects the shape of the subjective value function. In contrast, the natural contingencies of reinforcement explanation states that the description of the bet reminds decision makers of past experiences in which they were offered preposterous deals (e.g., an email beginning “Good news: You’ve won \$25,000,000!!”). In light of such past experiences, decision makers facing the St. Petersburg scenario choose not to waste time trying to understand the description, and offer low valuations.

4 Hypothesis testing vs. choice prediction competitions

Classical decision research advances by objective and systematic hypothesis testing. In the context of model development, this method implies the study of one assumption at a time. Thus, it is not easy to justify large changes in the underlying assumptions needed to explore alternative hills (“alternative paradigms” under Kuhn’s [1962] terminology). In order to facilitate exploration of this type, my co-authors and I organized several choice prediction competitions (e.g., Erev et al., 2010a, 2010b, 2017). To assist exploration, competition participants (model developers) are not asked to justify their choice of assumptions, and they are encouraged to use large data sets collected by the organizers.

The 2015 choice prediction competition (Erev et al., 2017; CPC15) focused on decisions under risk and ambiguity, with and without experience. To set up the competition, we first identified 14 robust choice phenomena (including the four phenomena described in Tables 1 and 2, and a finite variant of the St. Petersburg paradox). We then described a 12-dimensional space of choice problems that can give rise to all 14 phenomena, and ran an experiment that examined 30 problems in an effort to replicate these phenomena. The results showed that all the phenomena are replicable, but that some of the initial tendencies are reversed by the availability of feedback.

In the next step, we ran another experiment which examined behavior in 60 additional problems randomly selected from this space, and presented a baseline model that was shown to capture the 14 phenomena and the aggregated choice rates in all 90 problems. Notably, this baseline model – dubbed the Best Estimate And Simulation Tools (BEAST) – is different from PT in many ways, and cannot be described as a point on the subjective functions hill. BEAST does not assume that deviations from the EV rule reflect the impact of subjective functions. Rather, it assumes that, while decision makers are sensitive to the best estimate of the two prospects’ EVs based on the problem description, they also doubt this description, and behave as if they weight it based on their experience with similar problems encountered in the past. It predicts quick learning to maximize the expected return only when the maximizing option also minimizes the probability of regret.

For the choice prediction competition proper, we (the organizers) published the results and our baseline model (BEAST), and challenged other researchers to propose better models. Interested researchers were asked to submit models implemented in computer programs that would read the (12) parameters of each choice problem as inputs, and provide the predicted choice rates as outputs. The submitted models were then compared in a third “target” experiment using 60 new randomly selected choice problems. The results revealed that the 12 best submissions (out of 25) were

variants of BEAST. None of the leading models relied on the subjective functions assumption.

It is important to note that the results of CPC15 do not imply that it is possible to find a single model that can capture all areas in the land of problems. They only demonstrate that relaxing the subjective functions assumption can improve our ability to predict the impact of economic incentives.

5 Three connections with the big data revolution

To clarify the implications of the current analysis, it is convenient to discuss them in light of the big data revolution. Consider first the suggestion that behavioral decision researchers should relax their working assumptions, and explore wider classes of models (alternative hills in the land of assumptions) even when the exploration cannot be justified based on systematic hypothesis testing. This suggestion is perfectly consistent with the approach that has led to the most important advances in the data sciences. Indeed, the main difference between the contemporary data sciences and traditional statistics is the replacement of the focus on hypothesis testing with a focus on predictions.

The second implication involves the analysis of relatively small data sets. In CPC15, the competitors could develop and estimate their model based on 90 experiments with a total of 214,500 observations. While this data set is much larger than the sets considered in typical decision research, the results show that it is not large enough for pure (“theory-free”) machine learning analysis. The most effective predictions are derived by basing machine learning algorithms on theory-based features. One demonstration of the value of this approach is presented by Plonsky et al.’s (2017) analysis of CPC15 (Erev et al., 2017). While their theory-free machine learning submission did not perform well in CPC15, Plonsky et al. (2017) show that it is possible to improve upon the winning models with a machine learning algorithm that uses BEAST (the baseline model) as a feature. This observation was supported in a more recent choice prediction competition (CPC18), which was won by just such a machine learning algorithm (Plonsky et al., 2019). Importantly, however, the improvement over the best variant of BEAST was neither large nor significant.

The third, and most important, implication of the present analysis involves the observation that new big data technology allows the redesign of many incentive systems so as to address difficult social problems. Specifically, the new technology can facilitate dynamic pricing and provide more feedback. For example, ride sharing applications like Uber try to encourage drivers to work more during rush hour by increasing their earnings using surge pricing, and by providing feedback concerning the options they have missed. These and similar developments increase the importance of models

that allow useful predictions of the impact of incentives and experience on human behavior.

6 Alternative uses of rational models

The analysis presented above treats simple rational models, like the EV rule, as tools that can help predict the impact of money. The best models in CPC15 clarify the conditions under which these tools are likely to be effective. In contrast, most previous behavioral decision studies use rational models as benchmarks: typical experimental studies clarify deviations from these benchmarks, and the leading theoretical analyses examine which assumptions best capture these deviations.

The use of rational models as benchmarks has been highly effective in facilitating communication between economists and psychologists. It has helped clarify the value of experimental study of choice behavior. Yet it may also limit behavioral research, as the set of environments in which rational models provide clear predictions is relatively narrow. For example, when decision makers do not receive a complete description of the incentive structure surrounding a choice (or cannot trust the accuracy of the description), many behaviors can be justified as rational under certain prior beliefs. This fact has led mainstream experimental decision research to focus on decisions from description, in situations in which the decision makers are likely to Read, Understand, and Believe (RUB) the instructions. While this “RUB convention” allows clear tests of the rationality benchmark, it sheds limited light on behavior in natural settings in which decision makers may not receive clear descriptions, and do not always read, understand, and believe the information they receive.

7 Summary

The current analysis suggests that basic decision research has high potential, but its present impact is much lower than it could be. The potential is high because money does indeed make the world go round, but the direction of spin is sensitive to the exact incentive structure. Thus, an effective means to predict the impact of distinct incentive structures on human decisions can help address existing social problems and prevent new ones. Moreover, the importance of basic decision research to the design of effective incentive structures is enhanced by the big data revolution, which allows the redesign of many social environments.

The limited present impact of extant decision research may be the product of convergence to a local maximum – that is, convergence to an area of low elevation in the land of assumptions illustrated in Figure 1. The current analysis highlights the importance of two conventions that appear to direct previous research to a hill of low elevation. The first is the use

of restrictive working assumptions. Mainstream decision research assumes that the main deviations from maximization reflect the impact of subjective functions. This assumption clarifies the differences between the proposed models and EUT, but it also leads basic research to a bounded area in the land of assumptions. The second restrictive convention involves the reliance on systematic hypothesis testing methods which direct research to focus on small refinements of the popular models, and ignore the possibility that their basic assumptions should be re-considered.

I propose that these difficulties can be addressed through choice prediction competitions that facilitate the exploration of wider sets of assumptions, and reduce the risk of convergence to a local maximum in the land of assumptions. CPC15 demonstrates the potential of competitions of this type. It shows that prediction of decisions under risk and ambiguity can be improved by replacing efforts to incrementally refine the popular models with models that assume very different underlying processes. While the popular models assume that decisions reflect the impact of subjective functions which bias the final outcomes, the best predictions in CPC15 were provided by models in which the expected payoffs are weighted with the outcomes obtained in small samples of past experiences.

The key difference between the current analysis and mainstream behavioral decision research involves the use of rational models. While mainstream research uses rational models as benchmarks, the current analysis uses them as tools that can help predict the impact of money. For example, BEAST assumes some sensitivity to EV, and also predicts when behavior is likely to converge to EV maximization.

In order to clarify the wider implications of the present analysis, it is important to recall that people have more important goals than earning money. I believe that the current analysis can help us advance our understanding of these more important goals – which are also typically more difficult to study – in two ways. First, the approach supported here, namely facilitating the exploration of large changes in our working assumptions by using prediction competitions, can help clarify the impact of these other goals. Second, in certain settings the understanding of the impact of money can help advance other goals and increase social welfare.

References

- Abdellaoui, M., L’Haridon, O., & Paraschiv, C. (2011). Experienced vs. described uncertainty: Do we need two prospect theory specifications? *Management Science*, 57(10), 1879–1895.
- Allais, M. (1953). Le comportement de l’homme rationnel devant le risque: critique des postulats et axiomes de l’école Américaine. *Econometrica*, 21(4), 503–546. <http://doi.org/10.2307/1907921>.

- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16(3), 215–233.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk (original 1738). *Econometrica*, 22(1), 22–36. Retrieved from <http://www.jstor.org/stable/1909829>
- Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, 124(4), 369–409.
- Erev, I., Ert, E., & Roth, A. E. (2010a). A choice prediction competition for market entry games: An introduction. *Games*, 1(2), 117–136. <http://doi.org/10.3390/g1020117>
- Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., Hertwig, R., Stewart, T., West, R. & Lebiere, C. (2010b). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47. <http://doi.org/10.1002/bdm.683>.
- Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of political Economy*, 56(4), 279–304.
- Glöckner, A., Hilbig, B. E., Henninger, F., & Fiedler, S. (2016). The reversed description-experience gap: Disentangling sources of presentation format effects in risky choice. *Journal of Experimental Psychology: General*, 145(4), 486.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.
- Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, 15(8), 534–539.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292. JOUR. <http://doi.org/10.2307/1914185>.
- Kahneman, D., & Tversky, A. (1984). Choices, values, and frames. *American Psychologist*, 39, 341–50.
- Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago, IL: The Univ. of Chicago Press.
- Lloyd, W. F. (1833). Two lectures on the checks to population: Delivered before the University of Oxford, in Michaelmas Term 1832. Oxford: J. H. Parker.
- Plonsky, O., Apel, R., Ert, E., Tennenholtz, M., Bourgin, D., Peterson, J.C., Reichman, D., Griffiths, T.L., Russell, S.J., Carter, E.C., Cavanagh, J.F., & Erev, I. (2019). Predicting human decisions with behavioral theories and machine learning. arXiv preprint arXiv:1904.06866
- Plonsky, O., Erev, I., Hazan, T., & Tennenholtz, M. (2017, February). Psychological forest: Predicting human behavior. In Thirty-First AAAI Conference on Artificial Intelligence.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York, NY: John Wiley and Sons.
- Skinner, B. F. (1985). Cognitive science and behaviourism. *British Journal of Psychology*, 76, 291–301. Retrieved from <http://onlinelibrary.wiley.com/doi/10.1111/j.2044-8295.1985.tb01953.x/abstract>
- Tversky, A., & Bar-Hillel, M. (1983). Risk: The long and the short. *Journal of Experimental Psychology Learning Memory and Cognition*, 9, 713–717
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323. <http://doi.org/10.1007/BF00122574>.
- von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.
- Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity*. Cambridge University Press.