Prediction oriented behavioral research and its relationship to classical decision research

Ori Plonsky and Ido Erev

Technion – Israel Institute of Technology

Abstract

This paper argues that two of the common methods used in behavioral and social sciences to reduce the chances that models overfit the available data, namely heavy reliance on benchmark models and rigorous parameter estimation techniques, can slow the advancement of these sciences. An examination of classical decision research highlights how applying these methods shaped the field but have also led to limited success. As an alternative, the paper proposes a prediction-oriented approach to the development of behavioral models. Evaluating and comparing models based on their predictive power inherently guards against overfitting and also facilitates accumulation of knowledge. The paper reviews research employing the prediction-oriented approach in behavioral decision research and demonstrates that, in contrast to a common misconception, the focus on predictions can also facilitate better understanding of the underlying processes.

Keywords: prediction; overfitting; behavioral models; decision making; choice prediction competitions.

In their revolutionary book “Theory of Games and Economic Behavior” von Neumann and Morgenstern (1947) suggest that the slow progress of the social and behavioral sciences, relative to the natural sciences, is a result of the use of inappropriate mathematics. They propose game theory to facilitate the development of quantitative models that facilitate accumulation of knowledge. While game theory has triggered the development of quantitative models of choice behavior that advanced the social and behavioral sciences toward the natural sciences, the gap is still wide.¹ The current paper considers the impact of another potential contributor for this gap: the use of restrictive practices in model development. We argue that the common methods used by social and behavioral scientists to reduce the risk of overfitting the data can slow the accumulation of knowledge. We demonstrate how a prediction-oriented approach to model development can help mitigate these problems.

The classical method to avoid overfitting, and the possibility of a local maximum

Overfitting occurs when models erroneously attribute sample-specific noise as signal. Although all models are wrong, models that overfit the available data are also unlikely to be useful. Overfitting leads to incorrect conclusions concerning the underlying behavioral phenomena and impairs prediction of future behavior. Traditionally, to avoid overfitting, developers of quantitative behavioral models followed two reasonable practices: reliance on strong benchmarks, and hypothesis testing.

Minimal refinements of strong benchmark models

The natural way to avoid overfitting involves reliance on well-established theoretical benchmarks. Under this method, new models are built on other published models, but make the minimal changes necessary to capture new data of interest. Assuming that previous research has made good progress towards better understanding of the relevant behavior and its underlying processes, this practice makes a lot of sense and reduces the chance of overfitting. Previous models were probably developed based on previous data so building on these models and adjusting them to form improved ones increases the chance that the model captures more than the limited data available to the researcher. In particular, when changes from the well-established benchmark are not radical, it should be relatively easy to verify that

¹ One demonstration of the significance of this gap comes from comparison of the GRE test used to evaluate potential graduate students in different disciplines. While most of the questions in physics require the examinee to predict the results of different experiments, most of the questions in psychology require the examinee to explain the meaning of specific terms (Erev & Livne-Tarandach, 2005).

the main set of empirical phenomena this benchmark was meant to capture is also captured with the modified version. Moreover, this practice aligns with researchers' incentives. Adjusting a known model is more easily justifiable to editors, reviewers, and readers, and probably requires less modeling effort than suggesting a completely new model.

Table 1 summarizes three of the successes of this “reliance on strong benchmarks” method in behavioral decision research. The first is Bernoulli's (1738/1954) solution to the St. Petersburg paradox, the observation that people are willing to pay relatively little for the chance to play a game with infinite expected value (EV). This observation is a violation of the prominent benchmark model that assumes people choose options that maximize their EV. To solve the paradox, Bernoulli changed the assumption of EV maximization to one of expected utility maximization (where utility is a subjective function of wealth). In practice, this implied extending the benchmark model with a single parameter that can be interpreted as an abstraction of risk aversion. In other words, Bernoulli added a new mechanism that can help the existing benchmark capture new data of interest. Two centuries later, von Neumann and Morgenstern (1947) showed Bernoulli's ideas can be derived from reasonable assumptions concerning human preferences, and used it to justify the cardinal utilities assumption in game theory. Their elegant analysis made expected utility theory (EUT) a natural benchmark for the study of human decision making. Indeed, countless studies since have demonstrated deviations from the benchmark and suggested modifications of EUT that can capture these deviations (e.g., Friedman & Savage, 1948; Markowitz, 1952; Preston & Baratta, 1948).

By far the most successful modification of EUT is prospect theory (Kahneman & Tversky, 1979), now considered by many to be a truly descriptive model for decision making, one that “tell[s] us how people actually make [risky] choices” (Thaler, 2016, p. 1591). The bottom two panels in Table 1 demonstrate violations of EUT that can be captured with prospect theory, and its derivatives (Tversky & Kahneman, 1992; Wakker, 2010). Specifically, prospect theory modifies EUT by adding “probability weighting” parameters to capture the experiment described in the second panel, and a “loss aversion” parameter to capture the experiment described in the third panel.

Table 1: Three classical demonstrations of deviations from benchmark models, and the leading explanations.

Simple demonstrations of deviations from the benchmark model	Implications and explanation	Empirical phenomena
<p>St. Petersburg paradox/risk aversion (Bernoulli, 1738/1954): A fair coin will be flipped until it comes up heads. The number of flips will be denoted by the letter k. The casino pays a gambler 2^k. What is the maximum amount of money that you are willing to pay for playing this game?</p> <p>Modal response: less than 8</p>	<p>The results demonstrate that people do not maximize EV (which in this game is infinite). They can be captured by assuming maximization of concave subjective utility function that implies risk aversion.</p>	<p>Buying insurance.</p>
<p>Kahneman & Tversky (1979): Choose one of the options in each of the following pairs:</p> <p>S1. Win 5 for sure R1 0.1% to win 5000, 0 otherwise.</p> <p>S2. Lose 5 for sure R2 0.1% to lose 5000, 0 otherwise.</p> <p>Main results: Most participants chose R1 and S2.</p>	<p>The results violate the rational explanation of expected utility theory assuming fixed risk attitude. They can be captured by assuming a bias that implies oversensitivity to rare events.</p>	<p>Many people buy both lottery tickets, and insurance policies.</p>
<p>Redelmeier & Tversky (1992), following Samuelson (1963): Imagine that you have the opportunity to play a gamble that offers a 50% chance to win \$2000 and a 50% chance to lose \$500. Would you play the gamble?</p> <p>Main results: Most participant rejected the gamble.</p>	<p>The results violate the common abstraction of risk aversion under expected utility theory. They can be explained by assuming a bias that implies oversensitivity to losses.</p>	<p>Equity premium puzzle: return for risky equity is far too high relative to risk-free assets, implying people demand too high premium for potential losses.</p>

Hypothesis testing

A second common practice to avoid overfitting is to examine if the contribution of each proposed mechanism, parameter, or variable is statistically significant. Because including additional mechanisms or parameters in the model always improves in-sample fit, simply adding these mechanisms or parameters endangers overfitting. It then makes perfect sense to only include a mechanism if rigorous hypothesis testing shows there is evidence that it sufficiently contributes to model fit. And, again, this practice aligns with researchers' incentives. For example, reviewers may raise questions as to why a certain variable or

mechanism is *not* included in the model. Showing that adding it does not “significantly” improve the fit is an excellent answer. Similarly, readers and reviewers may appreciate a “formal” test of the underlying assumptions (although often the data does not adhere to the assumptions of this formal test) and may also feel more comfortable interpreting and reading into parameters that have gone through such formal testing.

For example, one of the contributors to the popularity of prospect theory is the fact that its derivative, cumulative prospect theory (Tversky & Kahneman, 1992) captures the distinct observed deviations from the EUT benchmark with specific parameters. Then, with the publication of methods and tutorials for careful estimation of the model’s parameters (e.g., Harless & Camerer, 1994; Nilsson et al., 2011), it became easy to test of the significance of these parameters and shed light on the significance of the assumed biases.

Obstacles to the accumulation of knowledge

Both extending/correcting previous models and using careful elegant parameter estimation techniques to determine the contribution of different parts of the models can indeed lower the risk of overfitting.² While there is nothing inherently wrong with either of these practices, a problem can arise because following these common practices direct research to a focus on a limited space of quantitative models and experiments. When building on and making relatively small changes in previous models, the new models tend to share at least some of the main assumptions made in previous modeling efforts. Models that do not adhere to these assumptions are not even considered. Similarly, to be able to perform rigorous parameter estimation under standard statistical techniques, there is a clear incentive to prefer simpler models to which there is a known framework for parameter estimation. More complex models in which elegant estimation of parameters is less straightforward will likely be ignored. The two common practices also feed off each other. Published models are more likely to have clear parameter estimation methods, and researchers are then likely to propose relatively small corrections to these models, again reinforcing the same tendency to propose similar models. Finally, to facilitate comparison of new models to the known models, it is often necessary to focus on tasks for which the known models have clear predictions. Limiting the space of tasks that are considered in this manner further restricts the set of potential models that should be considered because models that have larger scope than necessary are likely to be needlessly complex.

² Rabin (2013) provides further reasons to follow the practice of extending an existing (economic) model with a (behavioral) mechanism embedded with a new parameter.

To clarify the possible problems, it is convenient to think about the development of descriptive models as a hike in the “land of assumptions” in an attempt to find a set of assumptions (a hill) with a wide field of view on the “land of phenomena” (Erev, 2020). In this analogy, a point from which a phenomenon can be seen represents a model that captures that phenomenon. Good models are then points of high elevation that provide good view over the land of phenomena and thus good explanations for behavior. This analogy highlights two ways in which overreliance on benchmarks and hypothesis testing might impair the accumulation of knowledge. The first involve convergence to a local maximum. When the benchmark model is a point on a hill of low elevation (i.e., the model that provides the best approximation of the underlying process is very different than the commonly accepted benchmark), the focus on small modification of this benchmark (a hill climbing progress) will reduce the chance of leaving this hill, and mask the possibility of finding hills with much better points of view. The second problem involves the risk of hiking across a saddle to a nearby hill in an attempt to view certain phenomena of interest, but in the process losing sight of other, previously visible, phenomena.

The evolution of descriptive models of human decision making clarifies the significance of these problems. One indication for the negative impact of the tendency to stay on the benchmark hill, is illustrated by the fact that EUT was designed to address decisions among fully described payoff distributions. To test the quantitative predictions of this theory it is necessary to estimate the decision makers’ utility and assume that the participants in the experiments read, understand, and believe these descriptions. The most influential descriptive models (including prospect theory) continued the study this relatively narrow set of situations. This line of research seems to have (implicitly) assumed that the understanding of these situations will also shed light on natural decisions that are made based on incomplete description and past experience.

The right-hand column in Table 1 presents examples of natural phenomena that can be – and often are – explained with the respective deviations from the benchmark model. In particular, prospect theory can explain these and additional important economic phenomena from the real world (Barberis, 2013; Camerer, 2000; DellaVigna, 2009). Yet, some of the best examples of such phenomena come from settings in which people make many repeated decisions among partially known prospects, gain experience, and get feedback on their decisions (e.g., in financial markets). The tendency to apply prospect theory to address natural decisions from experience was not reduced by experimental demonstrations that experience can reverse the direction of the deviation from maximizations that prospect theory

was originally designed to capture (Erev et al., 2017; Hertwig & Erev, 2009).³ The focus on the narrow set of situations for which EUT applies (in an attempt to find modifications that improve the explanatory power of the model in this set) thus masked the fact while some important real-world phenomena that emerge when people rely on experience *can* be explained by prospect theory, experimental evidence shows that in such situations prospect theory is unlikely to even apply.

An example of the tendency to cross a saddle to a nearby hill is illustrated by the progress from EUT to prospect theory described in Table 1. Prospect theory and its derivatives helped explain the phenomena illustrated in the bottom two panels. However, as it turned out (Blavatskyy, 2005), cumulative prospect theory cannot capture, with the same set of parameters, both these phenomena and the St. Petersburg paradox (top panel of Table 1) that motivated the development of its ancestor. While it may be possible to develop a modification of prospect theory that addresses this challenge, it may be hard to account for many other phenomena that strictly violate the basic assumptions of benchmark model (e.g., Bernheim & Sprenger, 2020; Birnbaum, 2008; Cohen et al., 2020; Payne, 2005; Thaler & Johnson, 1990). Moving to a completely different “mountain ridge” may be more useful.

Avoiding overfitting by focusing on prediction

When researchers must rely on small data sets, it is difficult to find good alternatives to the popular practices criticized above. However, when more data is available, behavioral scientists can learn from the recent success of data scientists and artificial intelligence researchers, and use direct methods designed to avoid overfitting. Specifically, we suggest quantitative models will be evaluated (also) based on their ability to predict new data. Evaluation based on predictive performance inherently limits the danger of overfitting to in-sample data, because models are optimized for prediction of out-of-sample data. With evaluation based on prediction, quantitative behavioral scientists could relax their reliance on

³ How is it then that these natural phenomena are consistent with the original deviations? To answer this, it is important to consider the possibility that when researchers actively search for empirical phenomena that are consistent with some model—indeed almost any model—they will inevitably find some that do. This point was figuratively clarified by Economist Douglas Bernheim in a 2020 conference discussion on the predictive value of prospect theory (Plonsky, 2020). Bernheim suggested the hypothetical “Itzhak Theory” which posits that all parents name their son Itzhak. Then, to show that the theory is valid, he gives examples of ten people named Itzhak, allegedly corroborating a prediction the theory makes. Indeed, other phenomena in the same domains appear to contradict the assumptions of prospect theory. For example, many financial investors hold risky and under-diversified portfolios (Goetzmann & Kumar, 2008), which appears to be inconsistent with loss aversion.

these methods and allow themselves more degrees of freedom when developing quantitative behavioral models.

Before we detail how we think a prediction-oriented approach to development of models should manifest, it is constructive to address a common critique of this approach. Some behavioral scientists object the focus on prediction on the grounds that the main goal of the models in behavioral science is to provide an understanding of the phenomena and its underlying processes. This is usually interpreted as discovery of the true causal structure of these processes. Avoiding overfitting by alluding to predictive models may be problematic for two reasons. First, such models are often frowned upon for being uninterpretable “black boxes” from which it is impossible to learn general principles. Yet, the fact that some highly predictive models are difficult to interpret does not mean that there are no predictive models that are simple enough for humans to understand. Second, to be predictive, a model need not uncover any causal relations. Rather, it is sufficient if it identifies a structure that is highly correlated with the data. Indeed, it is known that it is possible to derive misspecified models that predict better than correctly-specified ones (Shmueli, 2010; Yarkoni & Westfall, 2017). To this we answer: It is true correlation does not imply causation, but models with high correlation with out-of-sample outcomes are at least good candidates for development and testing of causal theories, and surely better candidates than models that have poor out-of-sample predictions.

A prediction-oriented approach

Table 2 lists four steps we propose researchers take when using a more prediction-oriented approach to development of behavioral models. It starts with an explicit definition of the space to which the model should apply and the known empirical phenomena it should capture in this space. When fitting and evaluating models in-sample using the available data, the “space” in which the models are evaluated is implicit even if it is not clearly defined: It is simply the space of the problems that are used in the fitting exercise. However, when models should be evaluated out-of-sample, it is necessary to explicitly define the space of tasks or problems because at time of model development the members of this space that will serve as the test set are unknown (if they are known, the model can be developed to overfit the test set, and hence overfit in general). Explicitly defining the boundaries of the space will allow generating the test data independent of the model development. To facilitate accumulation of knowledge, it is important, of course, that the space will be wide enough to cover the phenomena explained by the leading benchmark models.

Table 2. A prediction-oriented approach

No	Step	Comment
1	Define a space to which the model applies	The space should be wide enough to allow replications of the puzzles that inspired the leading benchmark models.
2	Collect train data	This may include a demonstration of new empirical phenomena, but should include large, diverse set of training examples from the space
3	Develop a new model	The model should capture the main empirical phenomena from step 1, but optimized for maximal prediction scores.
4 (optional)	Test the model on new data	Generate a new set of examples that the model has predictions for, similarly to how training examples were selected. Publish the model and its predictions for this new set in a pre-print server. Collect the new data pertaining to the new set and examine the purely out-of-sample predictive accuracy of the model.

After the space and the empirical phenomena in that space are defined, the next step is to collect data that will be used for model development and training. When collecting data, it is important to remember that a good predictive model for a space covers the whole space, not only isolated pockets of the space in which “interesting” phenomena emerge. Specifically, under the prediction-oriented approach, the model should be designed to capture the gaps that are “in-between” known empirical phenomena: What happens when several phenomena may emerge at once? What happens when neither should explicitly emerge? Hence, while the collected data may include demonstrations of new empirical phenomena, it must be sufficiently large and diverse to cover the whole space. Often, this can be done by randomly sampling (many) problems from the space.

In the final steps, the model is developed and evaluated with a focus on prediction. This focus allows more degrees of freedom in model development while guarding against overfitting. It is highly encouraged that the model’s predictive power (i.e., lack of overfitting) will be verified with independent data collected only after the model has been finalized.

To clarify the potential of this approach we review research that used it to extend the study of decision making under risk and uncertainty considered in Table 1. Part on the importance of this field is a result of the fact that the development of useful game theoretic models is built on the way people react to risk and uncertainty. The following section describes how we implemented the prediction-oriented approach from Table 2 to develop a predictive model in this domain. We also provide some alternatives and tools that can help

researchers perform similar prediction-oriented research. Finally, we show this approach can lead to new theoretical insights that can be generalized and promote understanding.

A prediction-oriented study of choice under risk and uncertainty

To develop a new model of choice under risk and uncertainty, we (with our colleague Eyal Ert) implemented the prediction-oriented approach described in Table 1, using a choice prediction competition methodology (Erev et al., 2017). As the competition was held in 2015, we hereafter refer to it as CPC15.

The first step of this approach includes explicit definition of the space of prediction tasks that will be used to compare models. This seems like a natural step for any model development process. For example, Kahneman and Tversky (1979) defined prospect theory as a model that applies to simple prospects with at most two non-zero monetary outcomes and with stated probabilities. Note the definition implies both fixed dimensions that define the paradigm (monetary outcomes, stated probabilities) and variable dimensions that define the volume of the space (up to two non-zero outcomes). Under the prediction-oriented approach, it is important to define the space in such a way that it would be possible to generate tasks from the space. In this respect, the definition of prospect theory's space is incomplete. For example, it is unclear how values of the monetary outcomes are chosen. In CPC15, we defined an 11-dimensional space of tasks and an algorithm that generates tasks from within the space. In addition, with a choice of an experimental paradigm, some additional dimensions of the space are implicitly stated. In CPC15, each choice task included 25 repeated decisions within each task, the first five without feedback and the rest with full feedback. This essentially defined additional fixed dimensions to which CPC15 models apply.

To facilitate accumulation of knowledge, it is constructive to consider a space wide enough so that it includes tasks and phenomena examined in interesting previous research. Ideally, one should also replicate the phenomena using the paradigm and within the space defined. This provides support to the claims that the model that will be developed can capture the phenomena of interest. For example, the space examined in CPC15 is wide enough to demonstrate the emergence of 14 distinct behavioral decision making phenomena, including some that emerge when choosing between simple fully described prospects (a-la Kahneman and Tversky); others that emerge when repeatedly choosing between unmarked alternatives that provide feedback (i.e. "decisions from experience"); another (the Ellsberg paradox) that emerges when probabilities of one of the alternatives are unknown (i.e. the alternative is

ambiguous); and another (St. Petersburg paradox) that emerges when one option provides many possible outcomes. We therefore defined a space in which options can have just a few or many outcomes, and with “feedback” and “ambiguity” dimensions. We then replicated all 14 phenomena within this space.

Note that an explicit definition of the dimensions and the boundaries of the space to which the model applies does not necessarily mean that the model would not or cannot apply outside the defined space. It only reflects uncertainty with respect to the degree that it would. While researchers are unlikely to consider and define boundaries using every possible factor that could potentially affect the response of interest, even *not* explicitly considering some factor can be informative for readers and potential users of the model. Empirical phenomena that emerge outside the space that is explicitly defined may involve an additional factor that was ignored, and this makes any potential issues with generalization more salient. That way, it is far more difficult to cherry-pick empirical phenomena that are consistent with the model. If a phenomenon emerges within the specified space, then the model should capture it. If the phenomenon robustly emerges only outside the space, it cannot be used as supporting evidence for the model’s validity without explicit qualification. For example, any claims that the models developed in CPC15 generalize to tasks in which, for example, decision makers get partial feedback or get feedback from the very first trial should be taken with caution until empirically verified.

As mentioned above, in the (first) data collection step, it is important to collect a large and diverse set of examples that cover as much of the space as possible. That is, the training data should include many different combinations of the varying dimensions that define the space. Ideally, choice of training data examples will be done by random sampling of values for the dimensions that define the space. However, in many cases certain combinations are technically or logically implausible (or even impossible) and in other cases too many combinations lead to tasks in which behavior will be trivial to predict. In such cases, it may be a good idea to bias the task selection algorithm such that it would produce less (but some) trivial or implausible tasks. Of course, introducing such bias changes the distribution of tasks within the space, and should be done with great caution. In CPC15, the problem selection algorithm we used limited the likely difference between the expected values of the two alternatives. Moreover, because problems were played by participants for real money, alternatives with possible very high, or, more so, very low (negative) outcomes were removed.

The next step is the actual model development. As in any model development exercise, this is the hardest part and there is no recipe for how this should be done. During development, we propose that any potential model would be evaluated based on two criteria. First, it should capture the empirical phenomena described in previous steps. Second, it should be able to provide useful predictions for new data. For the latter, one very useful concept is cross-validation. In cross-validation, the researcher splits the available data to k (where k is between 2 and the number of training examples, but 5 is often a good choice) mutually exclusive “folds” (parts), of which $k-1$ folds are used for training and one is used for testing. That is, the model is trained on $k-1$ folds and provides predictions for the test data. This process is usually repeated k times with each fold used as a test set once. The estimate for out-of-sample predictive performance of the model is then the average performance over all folds. The practice is widely used in data science and gaining grounds in behavioral sciences. The interested reader should look elsewhere for more details. Here we only mention that it is of extreme importance that no knowledge that can be gained from the held-out fold should be available to the model at the time of training (Domingos, 2012). For example, if the model predicts the behavior of an individual over two distinct time points, it is wrong to separate these two correlated data points to two different folds.

The above two criteria may sometime conflict. A model that provides the best predictive performance according to a cross-validation analysis may not capture all empirical phenomena of interest. When this happens, the researcher should make a judgement call that depends on, among other things, the level of confidence the researcher has that the empirical phenomenon not captured is robust, the size of training data, and the flexibility of the model. When confidence is high, the training data is small and the model is not very flexible, it may be better to favor a version of the model that sacrifices a bit of predictive power to capture the phenomenon (and vice versa). Of course, it is also possible to try to modify the model.

At this point, a model which was developed with prediction in mind has been finalized (this includes fixing values for any free parameters) and has an estimate for its predictive performance in the space. A highly encouraged final step is to show that it is indeed usefully predictive out-of-sample. Cross validation measures can be too optimistic (Malik, 2020), so to get an unbiased estimate of the model’s predictive power, the researcher should now generate a new set of points from the space to which the final model can provide predictions. These predictions should be made public (e.g., on a pre-registration or pre-print platform) and only then the data for these points should be collected. This provides a clear

out-of-sample measure for how well the model predicts new data, the epitome of the prediction-oriented approach.

Additional tools and considerations

Common estimates for the predictive power of a model, like mean-squared-error (MSE), can be hard to interpret because its value is highly dependent on the domain. One way to clarify the meaning of a model's predictive power is to report an interpretable measure that spans across domains. ENO (Equivalent Number of Observations; Erev et al., 2007) is an example of such measure. The ENO of a model is an estimation of the number of new measurements that have to be collected until the average measurement provides a better prediction for a new data point than the prediction of the model. Thus, it gives an estimate for how much we can save on data collection by using the prediction of the model. For example, the model we suggested in CPC15 (called Best Estimate and Sampling Tools, or BEAST) achieved $ENO > 12$. This means that it provides a better prediction for the behavior of a participant in an unseen choice task drawn from the space than predicting the participant's behavior based on the average behavior of 12 participants in that task. The ENO is a function of the pooled estimated variance S^2 (pooled over all the prediction tasks), and the models' MSE (in predicting each of the observations). Specifically, $ENO = S^2 / (MSE - S^2)$

Another way to report a more intelligible estimate of a model's predictive power is to compare its performance to that of other models. Yet, researchers have clear incentives to compare their new models to some models and not to others (e.g., to those that happen to perform worse than their new model). Holding a prediction tournament (as was did in CPC15) helps guard against this problem. Here, before collecting the final dataset (and, often, before even generating a new set of tasks from the space), researchers make their suggested model and its results on the previously collected data public. They then challenge other researchers to submit better models that are evaluated based on their predictive performance in the new set. The advantage of such tournaments is the ability to compare many different models on a common dataset, another widely used practice in data science.

When comparing the predictive performance of two models in a space, it is important to verify that the relative ranking of the models is robust to changes of the (random) selection of test cases drawn from the space. One way to do this is with a bootstrap analysis in which many test sets are created from the original test set, and for each, the MSE (or ENO or another indicator of prediction error) of both models is computed. Then, the differences between the MSEs across tasks are computed and averaged. By truncating the maximal 2.5%

and the minimal 2.5% of average differences, a 95% confidence interval for the difference in MSEs can be retrieved. For example, in CPC15, the winning model achieved ENO of 13.5 but using this bootstrap analysis it was revealed its difference from the ENO of BEAST was not significant.

Collecting data sufficiently large to cover a wide space of tasks, not to mention organizing a prediction tournament for that space, requires considerable resources. One way to reduce the burden on individual research teams is to organize them using large scale collaborations between many research labs. This also has the advantage of providing different perspectives when deciding on the exact boundaries of a space of interest and the important empirical phenomena that it includes.

Once a large and diverse dataset has been collected, and a predictive model has been developed for the space, and assuming these have been made publicly available, they can serve individual researchers in model development. Importantly, new models need not necessarily be similar to the existing benchmark model that exists for the space. The actual benchmark to use is the model’s predictive performance. It is however, important to note that once the test set and the benchmark’s performance for it has been published, it is possible to overfit the test set (Dwork et al., 2015). Claims that a new model predicts better an already published test set than an original model should be taken with caution.

We suggest that when a large dataset and a predictive model already exists, researchers who develop new models can contribute in at least four ways. One way is by demonstrating new empirical phenomena that emerge in the space and showing that the new model both captures these new phenomena and predicts similarly or better than the current benchmark in the space. Second, researchers can use the existing dataset for training both their new model and the existing benchmark model, derive predictions of both trained models on a new test set and then compare the models on the new set.

Third, researchers can suggest a new dimension to the space. That is, expand the space and then augment the published dataset with new data. The latter approach was taken in another prediction tournament (CPC18) that we organized (Plonsky et al., 2019). In that tournament, we expanded the space from CPC15 such that both alternatives could have included up to 10 different outcomes. In addition, we used the full dataset of CPC15 for training and augmented it with new data for training and for testing.

Finally, researchers can show that their new model predicts similarly or better than the current benchmark, and that in addition it provides more intuitive predictions and is more

understandable by potential users. A more understandable and intuitive model is likely to be more useful in practice.

Prediction can promote understanding

As a final note, we would like to address a common source of skepticism behavioral scientists have with prediction-oriented research. Namely, the goal of science is to understand, often interpreted as revealing the true causal structure that underlies phenomena, but prediction concerns finding good correlative structure that can foreshadow observed outcomes. This is true. Yet, when we have a good predictive model, we should explore what is it about this model that makes it predictive in this space. The answer can lead to insights and better understanding. This can be done in several ways. For example, Fudenberg and Liang (2019) derived a predictive model of initial play in 2-player economic games and then carefully examined the tasks in which their model predicted much better than theoretical benchmarks. They then discovered that these tasks have regularity which led them to add a new mechanism to the theoretical model.

In another example, Plonsky et al. (2017) developed a machine learning model that outperforms BEAST in CPC15 data. Their model used sets of features based on BEAST as input in a random forest algorithm. They then systematically tested how removing different sets of features that were derived based on different mechanisms of BEAST changes the behavior of the model. This led them to conclude that features that mimic reliance on small samples and EV maximization are most important (Barron & Erev, 2003; Erev & Roth, 2014).

Finally, the success of models that assume reliance on small samples in a series of tournaments for prediction of choice behavior in repeated choice with feedback (Erev et al., 2017; Erev, Ert, & Roth, 2010; Erev, Ert, Roth, et al., 2010; Plonsky et al., 2019) has strengthened our confidence in the usefulness of this assumption to describe behavior in such tasks. Yet, it has also pushed us to theorize on what underlies this assumption and what it is about these tasks that makes this assumption so useful. Our analysis led us to consider the idea that reliance on small samples is driven by a sophisticated attempt to follow sequential patterns, even though these do not exist (Plonsky et al., 2015; Plonsky & Erev, 2017). This idea led to new testable predictions, specifically that people will exhibit a “wavy recency” response to rare events, and that when the task does involve sequential patterns, most people will get them (Gaissmaier & Schooler, 2008; Plonsky & Teodorescu, 2020). It is worth stressing this last point. The reliance on small samples assumption was useful for prediction

in the different tournaments because these tournaments focused on spaces without sequential patterns, and the implications of the reliance on small samples assumption are correlated with the behavior of people in such tasks. Had we considered a space that does involve patterns, this assumption will not have worked. That is, the best models provided a useful approximation of the underlying process, but they were also wrong. Nevertheless, carefully considering why they provided a useful approximation has helped us advance towards better understanding.

References

- Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, 27(1), 173–196.
- 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.
- Bernheim, B. D., & Sprenger, C. (2020). On the empirical validity of cumulative prospect theory: Experimental evidence of rank-independent probability weighting. *Econometrica*, 88(4), 1363–1409.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk (original 1738). *Econometrica*, 22(1), 23–36.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, 115(2), 463–501. <https://doi.org/10.1037/0033-295X.115.2.463>
- Blavatskyy, P. R. (2005). Back to the St. Petersburg Paradox? *Management Science*, 51(4), 677–678. <https://doi.org/10.1287/mnsc.1040.0352>
- Camerer, C. F. (2000). Prospect theory in the wild: Evidence from the field. In C. F. Camerer, G. F. Loewenstein, & M. Rabin (Eds.), *Advances in behavioral economics* (pp. 148–161). Princeton University Press.
- Cohen, D., Plonsky, O., & Erev, I. (2020). On the impact of experience on probability weighting in decisions under risk. *Decision*, 7(2), 153–162.
- DellaVigna, S. (2009). Psychology and economics: Evidence from the field. *Journal of Economic Literature*, 47(2), 315–372.
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87.
- Dwork, C., Feldman, V., Hardt, M., Pitassi, T., Reingold, O., & Roth, A. (2015). The reusable holdout: Preserving validity in adaptive data analysis. *Science*, 349(6248), 636–638.
- Erev, I. (2020). Money makes the world go round, and basic research can help. *Judgment & Decision Making*, 15(3), 304–310.
- 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. <https://doi.org/10.1037/rev0000062>
- Erev, I., Ert, E., & Roth, A. E. (2010). A choice prediction competition for market entry games: An introduction. *Games*, 1(2), 117–136. <https://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. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, 23(1), 15–47. <https://doi.org/10.1002/bdm.683>
- Erev, I., & Livne-Tarandach, R. (2005). Experiment-based exams and the difference between the behavioral and the natural sciences. In R. Zwick & A. Rapoport (Eds.), *Experimental business research* (pp. 297–308). Springer.

- Erev, I., & Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111(Supplement 3), 10818–10825.
- Erev, I., Roth, A. E., Slonim, R. L., & Barron, G. (2007). Learning and equilibrium as useful approximations: Accuracy of prediction on randomly selected constant sum games. *Economic Theory*, 33(1), 29–51. <https://doi.org/10.1007/s00199-007-0214-y>
- Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *Journal of Political Economy*, 56(4), 279–304.
- Fudenberg, D., & Liang, A. (2019). Predicting and understanding initial play. *American Economic Review*, 109(12), 4112–4141.
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109(3), 416–422. <https://doi.org/10.1016/j.cognition.2008.09.007>
- Goetzmann, W. N., & Kumar, A. (2008). Equity portfolio diversification. *Review of Finance*, 12(3), 433–463.
- Harless, D. W., & Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica: Journal of the Econometric Society*, 62(6), 1251–1289.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, 13(12), 517–523. <https://doi.org/10.1016/j.tics.2009.09.004>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–292.
- Malik, M. M. (2020). A hierarchy of limitations in machine learning. *ArXiv Preprint ArXiv:2002.05193*.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy*, 60(2), 151–158.
- Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, 55(1), 84–93.
- Payne, J. W. (2005). It is whether you win or lose: The importance of the overall probabilities of winning or losing in risky choice. *Journal of Risk and Uncertainty*, 30(1), 5–19. <https://doi.org/10.1007/s11166-005-5831-x>
- Plonsky, O. (2020). *Round table with Prof. Daniel Kahneman titled “Does prospect theory predict well?”* YouTube. <https://www.youtube.com/watch?v=Klia9OrJ7PI>
- 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. (2017). Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, 93, 18–43. <https://doi.org/10.1016/j.cogpsych.2017.01.002>
- Plonsky, O., Erev, I., Hazan, T., & Tennenholtz, M. (2017). Psychological forest: Predicting human behavior. *31st AAAI Conference on Artificial Intelligence, AAAI 2017*.
- Plonsky, O., & Teodorescu, K. (2020). Perceived patterns in decisions from experience and

- their influence on choice variability and policy diversification: A response to Ashby, Konstantinidis, & Yechiam, 2017. *Acta Psychologica*, 202.
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, 122(4), 621–647.
- Preston, M. G., & Baratta, P. (1948). An experimental study of the auction-value of an uncertain outcome. *The American Journal of Psychology*, 61(2), 183–193.
- Rabin, M. (2013). An approach to incorporating psychology into economics. *American Economic Review*, 103(3), 617–622.
- Redelmeier, D. A., & Tversky, A. (1992). On the framing of multiple prospects. *Psychological Science*, 3(3), 191–193.
- Samuelson, P. A. (1963). Risk and uncertainty: A fallacy of large numbers. *Scientia*, 98, 108–113.
- Shmueli, G. (2010). To explain or to predict? *Statistical Science*, 25(3), 289–310.
- Thaler, R. H. (2016). Behavioral economics: Past, present, and future. *American Economic Review*, 106(7), 1577–1600.
- Thaler, R. H., & Johnson, E. J. (1990). Gambling with the house money and trying to break even: The effects of prior outcomes on risky choice. *Management Science*, 36(6), 643–660.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323. <https://doi.org/10.1007/BF00122574>
- von Neumann, J., & Morgenstern, O. (1947). *Theory of Games and Economic Behavior* (2nd ed.). Princeton university press.
- Wakker, P. P. (2010). *Prospect Theory: For Risk and Ambiguity*. Cambridge University Press.
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100–1122.