

Schnytzer, Adi; Šušteršič, Janez

Working Paper

The regression tournament: A novel approach to prediction model assessment

Working Paper, No. 2011-10

Provided in Cooperation with:

Department of Economics, Bar-Ilan University

Suggested Citation: Schnytzer, Adi; Šušteršič, Janez (2011) : The regression tournament: A novel approach to prediction model assessment, Working Paper, No. 2011-10, Bar-Ilan University, Department of Economics, Ramat-Gan

This Version is available at:

<https://hdl.handle.net/10419/96027>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Regression Tournament: a Novel Approach to Prediction Model Assessment

By

Adi Schnytzer¹ and Janez Šušteršič²

Abstract

Standard methods to assess the statistical quality of econometric models implicitly assume there is only one person in the world, namely the forecaster with her model(s), and that there exists an objective and independent reality to which the model predictions may be compared. However, on many occasions, the reality with which we compare our predictions and in which we take our actions is co-determined and changed constantly by actions taken by other actors based on their own models. We propose a new method, called a regression tournament, to assess the utility of forecasting models and taking these interactions into account. We present an empirical case of betting on Australian Rules Football matches where the most accurate predictive model does not yield the highest betting return, or, in our terms, does not win a regression tournament.

February 2011

¹ Departments of Economics, Bar Ilan University, Israel.

² Faculty of Management, University of Primorska, Slovenia.

1. Introduction

In this paper, we propose a novel way to assess the relative quality of econometric models. We demonstrate empirically that the models with higher predictive accuracy, as measured by goodness of fit statistics, are not necessarily those that provide the most useful guidance for real life decisions. There is a need for an alternative assessment method that takes into account the usefulness of models for making real-life choices.

The basic idea comes from gambling studies. In a gambling setting, such as betting on a horse race or football match winners, what matters is not the goodness of fit of your predictive model *per se*, but rather the return it gives you when used for making bets at the odds set by a bookmaker. That is, the utility of your model is not assessed by comparing it to the real world outcomes (i.e. the actual match winners), but rather by interacting it with reality *as perceived* by your bookmaker and expressed by the odds offered to you.

Schnytzer (2010) has shown empirically that the two assessment criteria – the traditional “goodness of fit” and the gambling return made by betting on your regressions – do not always choose the same model as the best one. He developed a set of models to predict winners of Australian Rules Football matches, and simulated the returns one would achieve by using the model predictions to bet at an internet sport betting site, taking the last available fixed odds offered. The surprising result was that the model which would give him the highest betting return was not the one that had the highest goodness of fit in purely statistical terms.

In the following, we develop a simple taxonomy of forecasting to see clearly the issues involved, and then propose a method to assess the quality of econometric models by interacting them with each other, rather than just simply comparing each one separately with real world outcomes. We also provide an empirical illustration of the proposed method in a betting context.

2. Assessing predictions: statistical quality vs. utility of models

There are in general two ways of assessing quality of regression models. One is to use the model to estimate (or predict) values of the dependent variable and then compare these estimates to the actual real world values. The comparison is made by calculating some forecast accuracy measures,

such as mean standard error or goodness of fit statistics.³ This is a standard procedure; henceforth we refer to it as assessing the “statistical quality” of a model.

The other way is to take some action based on your regression model and then evaluate whether your action has achieved intended results. Betting on football matches is one example. Trading in financial markets based on your predictions of stock movements is another. An example from the public choice domain is fine-tuning a country’s budget to the most recent macroeconomic forecast. In all these cases, what matters is not the accuracy of the model itself, but the utility one gets by acting on model predictions. We call this procedure assessing the “utility” of a model.

It is usually taken for granted that there is a positive correlation between the statistical quality of models and their utility, i.e. that the more accurate model would also prove more useful if used for real-life decisions. But it is not straightforward to validate this assumption, and, as we show in this paper, it may actually be false.

In a real world setting, such as trading stocks or betting on sports, at least some of the actors will produce their own private forecasts that will not be publicly available and hence not observable to an outsider. Now assume we find that actor B achieves higher market returns than does actor A. This superior performance may be due to B having a more accurate forecasting model, or to B having a better ability to respond to a given forecast in a profitable manner, or both. As outside observers, we cannot tell which is the case, because we cannot assess the statistical quality of actors’ private forecasting models. We can only apply the “utility assessment”; that is, we may compare performance (outcomes, profits) of competing actors, but we can’t tell whether variation in performance comes from variation in forecasts or from variation in actions taken in response to forecasts.

There are two special cases to this general logic. One is a situation with no private models; that is, all decision-makers base their separate actions on the same set of publicly available forecasts. As an example, one can think of a set of publicly available macroeconomic forecasts and a number of economic agents (enterprises, investors, consumers, etc.) taking separate and diverging actions achieving different results. Or think of publicly known predictions of oil-reserves depletion and actions, if any, by actors that would probably be affected. Or think simply of a public weather

³ There are, of course, many more complicated approaches to measuring the quality of a forecast (including combinations of forecasts) but the essential point remains that the forecast is compared against reality. See Ramanathan (2002), 5th Ed, chapter 11.

forecast and people deciding whether to carry an umbrella or not. In such a situation when all forecasts are public, the variation in results (real-world returns of individual actors) must come from variation in responses to these forecasts. In this case then, the “utility assessment” strictly compares only the ability of decision-makers to act on the given publicly available information.⁴

The other special case is when actors have different private forecasting models, but all use the same decision rules for acting on their divergent forecasts. In such a situation, we would be safe to say that all variations in returns come from variations in the statistical quality of individual models, and not from differences in individual actions. In this case of homogenous responses to predictions, we could use the observable differences in returns as an (alternative) measure of statistical quality of the underlying models.

In reality, it is of course not possible to expect all individuals to follow the same decision rule. But it can be done on paper. This is what Schnytzer (2010) did in the above mentioned article. He applied the same betting rule to winning probabilities calculated using different forecasting models. Using the same rule for action, one would naturally expect that any differences in performance (betting returns) would strictly reflect differences in the statistical quality of forecasting models. However, and this is quite puzzling, the statistically superior model did not produce the highest profit. In this paper, we further explore this puzzling finding.

3. The puzzle

In Schnytzer (2010), forecasts produced by different models were used to simulate betting at the actual odds set by real-life bookmakers. In this way, the models were indirectly compared to each other by interacting their winning predictions with the odds set by the unknown bookmaker’s model. An obvious shortcoming of such procedure is that the model used to set the odds against which the simulated bets were waged was not observable. That is, we know the odds, but not the model (statistical or intuitive) used by the bookie to estimate winning probabilities and transform them into betting odds.

In this paper, we try to generalize the idea of testing the quality of models by interacting them to each other, without the need to have a third party (a bookie) offering odds independently. For an

⁴ Note that “acting on forecasts” also involves deciding on how much faith and weight one wishes to put on different publicly available forecasts. In deciding what to do, actors also rely on their expert judgements, informal rules of thumb or simply their gut feeling, which makes it hard to express their behaviour by a simple formal »decision rule«.

empirical illustration, we take the same models as used by Schnytzer (2010) to estimate the winning probabilities for Australian Rules Football matches. However, instead of using the actual odds put forward by the bookmakers, we take one of our own models to play the role of the bookie.

We use two types of models to estimate winning probabilities of Australian Football league matches: a linear probability model (LPM) and a conditional logit model (CLOG) (McFadden 1974). Four different specifications were used with each model. The first one (LPM1 and CLOG1) uses only player-level variables. The second pair of regressions (LPM2 and CLOG2) add two dummy variables, the first indicating whether or not the home team has an a priori home ground advantage and the second indicating whether or not the stadium in which the current game is being played is a neutral ground, offering no a priori advantage to either side.⁵ Regressions LPM 3 and CLOG3 add to the extant explanatory variables two team-level dummy variables which indicate whether or not the team has clinched a place in the finals or whether the team has definitely been eliminated from the finals race immediately prior to the game to be played, respectively. Finally, regressions LPM 4 and CLOG4 add a further team-level variable which measures the proportion of wins accumulated by the team so far in the current season prior to the current game.

For purposes of prediction, the regressions are run on the data subset containing all observations from the first round of 1998 through the 2000 Grand Final. These regressions are used to predict the winning probabilities of the teams in round 1 of 2001 by substituting the mean values of the player-level explanatory variables for the 1998-2000 period into the obtained regression results. For each player in the team, each regression predicts a probability which may be interpreted as that player's predicted contribution to the team's winning probability. In the case of the conditional logit regressions, these probabilities sum to 1 for each game. Thus, summing them across players in any given team yields the predicted winning probability for that team. The linear probability model requires an extra step since probabilities do not generally sum to 1 for each game.⁶ Accordingly, these predictions are normalized over each game and the resultant sums per team taken as the predicted winning probabilities for the relevant team.

⁵ For a thorough analysis of the subtleties of home and neutral grounds in the AFL, see Schnytzer and Weinberg (2008). Note also that "neutral" does not appear in any CLOGIT regression because it must always, by definition, receive the same value for both teams in a game and the conditional logit regression conducts its estimation by distinguishing between the two teams in a game exclusively.

⁶ It is interesting to note, however, that the predicted probabilities per game always fall between 0 and 1 in our data set.

The models are now re-estimated adding round one of 2001 data and used to predict winners in round two of 2001. This procedure is repeated for all rounds in all seasons up to the 2007 Grand Final. The simulated betting is on those teams for which the predicted winning probability exceeds 0.5 (i.e. the predicted favorites in the game) and the amount bet is in proportion to the predicted winning probability. This betting system is adopted as it is the method adopted by many Australian professional punters.⁷

Before introducing the notion of a regression tournament, we reproduce the essential results from Schnytzer (2000). Table 1 provides the basic goodness of fit statistics for the 8 model specifications, and the simulated average annual rates of return achieved by using the respective models for betting. Since returns in the first three years of simulation (2001-2004) were negative for all models, we calculate average returns also for the sub-period 2004-2007.

Table 1: Goodness of fit and betting returns

Model	LPM1	LPM2	LPM3	LPM4	CLOG1	CLOG2	CLOG3	CLOG4
Adjusted / pseudo R ² *	0.0180	0.0477	0.0581	0.0821	0.0147	0.0371	0.0530	0.0852
Rank **	4	3	2	1	4	3	2	1
Average rate of return 2001-07	-0.1767	-0.0350	-0.0030	-0.0096	-0.1414	-0.0347	-0.0154	-0.0430
Rank	8	5	1	2	7	4	3	6
Average rate of return 2004-07	-0.1687	0.0139	0.0395	0.0113	-0.1337	0.0123	0.0214	-0.0236
Rank	8	3	1	5	7	4	2	6

Source: adapted from Schnytzer (2000). Notes: * Ranks reported among models of the same regression type (LPM, CLOG), as goodness of fit statistics are not directly comparable. ** For the relevant regression run on the entire data set.

The table clearly demonstrates the puzzle. Among the LPM models, LPM4 is the best in statistical terms, but betting with the help of LPM3 or even LPM2 would clearly yield higher returns. Similarly, specification no. 4 is statistically the most accurate one among CLOG models, but CLOG3 and CLOG2

⁷ The authors thank Terry Pattinson (formally Australian sports betting bookmaker and currently Head of In-Play Development for William Hill PLC) for this insight.

clearly outperform it in terms of betting returns. The highest returns overall are realized by LPM3. Second place goes to LPM2 or, for the sub-period, to CLOG2.

There are two problems with interpreting the results in table 1. First, the goodness of fit statistics for LPM and CLOG regressions may not be directly compared. Second, the differences between the statistical quality and the real-life betting utility of models may depend on the specifics of their interactions with the model (explicit or implicit) used by the bookmaker to set the odds, which of course is unobservable to us. To circumvent these problems, we now proceed to what we call a regression tournament. This procedure allows us to compare the relative utility of our models without having to use the odds set by an exogenous bookmaker, and without computing any goodness of fit statistics.

4. A regression tournament

The idea is simple. We use our models for betting not against the odds set exogenously, but against the odds set by one of our other models. The odds are set for each match according to the predicted winning probabilities of the two teams. When converting the probabilities into bookmakers' prices, an "over-round" is added that corresponded to that used by actual bookmakers. For example, if the sum of the bookies' odds (expressed as prices) for a given match was 1.03, the predicted probabilities from our model were multiplied by 1.03.

The procedure for assessing the quality of our models is now the following. We use one of the models, e.g. LPM1, as our bookmaking model and set the odds for each game played since the first round of 2001. Then we use all other models to wage bets, according to winning probabilities estimated by each model, and calculate the betting return that would be made by using each model separately over the entire 2001 to 2007 period. Once this calculation is completed, we take another model (LPM2) to act as a bookie, and again use all other models to wage bets and calculate profits. We repeat the procedure until all models have taken the role of a bookie. To facilitate a fair comparison of models, the regressions were run over the identical sample of matches with no data missing. This required eliminating round one and the finals matches from each season in the forecast period (2001-2007).

In this way, we use the same decision rule with all betting models and hence, as explained in Section 2, the differences in betting returns should reflect only the differences in statistical quality of the models. On the hypothesis of correspondence between statistical quality and utility of models, we

would expect the models with better goodness of fit (of their predictions) to yield higher returns when used for betting. However, as results of our exercise in Table 2 demonstrate, this is clearly not the case. This means that our method – a regression tournament – provides an alternative way of assessing models' quality that does not necessarily give same results as traditional methods.

Table 2

Rates of return* made by models betting at the odds set by the bookmaking model

Bookmaking model	<u>Betting (predicting) model</u>							
	LPM1	LPM2	LPM3	LPM4	CL1	CL2	CL3	CL4
LPM1		0.0327	0.0262	0.0027	0	0.0324	-0.0103	-0.0652
LPM2	-0.1329		-0.0406	-0.0400	-0.1190	0.0324	-0.0103	-0.0652
LPM3	-0.1265	-0.0027		-0.0443	-0.1149	0.0319	-0.0103	-0.0652
LPM4	-0.0668	0.1518	0.1625		-0.0397	0.0353	-0.0107	-0.0652
CL1	0	0.0327	0.0291	-0.0060		0.0324	-0.0103	-0.0652
CL2	-0.1329	0	-0.0718	-0.0671	-0.1190		-0.0593	-0.0652
CL3	-0.1039	0.0116	0.2024	-0.6200	-0.0897	-0.0264		-0.0710
CL4	-0.0522	0.1043	0.1176	0.1768	-0.0250	0.0958	0.0842	
Negative returns	6	1	2	5	6	1	6	7
Positive returns	0	5	5	2	0	6	1	0
No bets	1	1	0	0	1	0	0	0
Adjusted R ² / pseudo R ² **	0.0187	0.0530	0.0689	0.0920	0.0155	0.0414	0.0643	0.0945

Notes: * A zero indicates that the odds set by the bookmaking model were always too low to bet, so no bets were made. ** For the relevant regression run on the entire data set. The figures differ slightly from the ones in table 1 due to some omitted matches as explained in the text.

Among the LPM models, specifications 2 and 3 clearly outperform the other two. Betting by help of LPM2 or LPM3 yields positive returns when used against odds set by five out of the seven other models. They only give negative or zero returns when bet against each other and against CLOG2. If we compare returns of LPM2 and LPM3 when used against the same bookmaking model, LPM2 yields higher returns when betting against specification 1 and 2 (LPM1, CLOG1, CLOG2), whereas LPM3 is better in betting against specification 3 and 4 (CLOG3, LPM4, CLOG4). The statistically superior specification LPM4 is clearly inferior for betting purposes, as it yields positive returns only twice.

Among the CLOG models, specification 2 stands out clearly. It yields positive returns in all cases except when used for betting against CLOG3. All other specifications, including the statistically superior CLOG4, yield negative returns in all or almost all cases.

Another advantage of our procedure is that it allows us to compare LPM and CLOG models directly. It is not possible to compare their statistical merits (goodness of fit), as one cannot directly compare the size of adjusted R^2 (LPM) to pseudo R^2 (CLOG), for example, but our procedure allows us to directly compare their utility (in terms of betting returns). In our case, it is interesting to compare the best LPM specifications (LPM2 and LPM3) to the best CLOG specification (CLOG2). When compared to specifications 1 and 4, there is always one of the LPM specifications that yields the highest return, with CLOG2 being ranked in second or third place. This would suggest that LPM models are generally more useful for this kind of betting than are CLOG models. However, when CLOG2 is used for betting directly against odds set by LPM2 or LPM3, it yields positive returns, while LPM2 and LPM3 produce a loss or a zero when used against CLOG2. This prohibits a clear conclusion as to the relative merits of LPM and CL models.

Table 3 offers a different look at our results, a sort of a league table for our regressions. Instead of counting how many times a model gives a positive return when used for betting, as in Table 2, we now look at pair wise comparisons and think of them as matches between two regressions. If the return is positive, the betting model wins. If the return is negative, the bookmaking model wins. If the return is zero, we declare it a draw (D). Table 2 lists winners of all such matches and counts how many times each model emerges as a winner. Now, specification CLOG2 emerges as a tournament winner with 12 wins, while LPM2 and LPM3 end up second with 11 points each. Interestingly, LPM4 and CLOG3 models, with 6 points each, now emerge as better than the rest (LPM1, CL1, CL4), which in Table 1 was not the case.

Table 3

A league table for a regression tournament: winners of games between pairs of regression

Bookmaking model	<u>Betting (predicting) model</u>							
	LPM1	LPM2	LPM3	LPM4	CL1	CL2	CL3	CL4
LPM1		LPM2	LPM3	LPM4	D	CL2	LPM1	LPM1
LPM2	LPM2		LPM2	LPM2	LPM2	CL2	LPM2	LPM2
LPM3	LPM3	LPM3		LPM3	LPM3	CL2	LPM3	LPM3
LPM4	LPM4	LPM2	LPM3		LPM4	CL2	LPM4	LPM4
CL1	D	LPM2	LPM3	CL1		CL2	CL1	CL1
CL2	CL2	D	CL2	CL2	CL2		CL2	CL2
CL3	CL3	LPM2	LPM3	CL3	CL3	CL3		CL3
CL4	CL4	LPM2	LPM3	LPM4	CL4	CL2	CL3	
Number of wins	2	11	11	6	3	12	6	2

In table 3, each match between two models is treated separately. But since all pairs meet twice (interchanging the roles of bookmaking and betting models), it is interesting to establish the combined winner of both games. We may have three situations. If the same model wins both matches, that is, if it gains a positive return as a bettor and inflicts a negative return on the other model when setting the odds, than clearly this model is the combined winner. The second situation is when both matches are won by the bookmaking models. In this case, the model which inflicts a higher loss on the other one, is the winner. A third possible situation is that both matches are won by the betting models, and of course the one that achieves a higher rate of return would be the match winner. Interestingly, this situation never appears in our simulations. Note also that when one of the two games is a draw, the combined winner is the model that won the other game.

Table 4 present results of the regression tournament when both matches between a pair of models are considered together and a combined winner is established. The basic result of table 3, that specifications CL2 is the most useful, is not changed. But there are no more shared positions: LPM2 now has a point more than LPM3, and CL3 a point more than LPM4.

Table 4

A league table for a regression tournament: winners of pairs of games between pairs of regression

Bookmaking model	<u>Betting (predicting) model</u>							
	LPM1	LPM2	LPM3	LPM4	CL1	CL2	CL3	CL4
LPM1		LPM2	LPM3	LPM4	D	CL2	CL3	LPM1
LPM2			LPM2	LPM2	LPM2	CL2	LPM2	LPM2
LPM3				LPM3	LPM3	CL2	LPM3	LPM3
LPM4					LPM4	CL2	CL3	LPM4
CL1						CL2	CL3	CL1
CL2							CL2	CL2
CL3								CL3
CL4								
Number of wins	1	6	5	3	1	7	4	0

Table 5 summarizes our results. It ranks the predictive models with respect to their simulated rates of return when betting against bookmaker's odds, their goodness of fit in predicting the match winners, and their performance in all three types of regression tournaments introduced above (tables 2 to 4). The problem this table tries to address can be described as follows. Assume you have several models available that you could potentially use for betting. Assume further that you have no data available on past odds, so you cannot determine the utility of models by simply simulating betting returns for past matches. You are thus left with two options for choosing the best model for your purpose: one is to rely on the statistical quality of the models in terms of predicting past match results. The other is to run a regression tournament, by which you interact models with each other and simulate betting returns that would be achieved if odds were set by your own models. This second procedure, proposed in this paper, is closer to what we termed the real-life utility of models.

Table 5

Comparing different methods for choosing the best betting model

<i>Rank of the models according to:</i>						
Model	rates of return, 2004-07	rates of return, 2001-07	pseudo / adj. R ² *	regression tournament in:		
				table 2	table 3	table 4
LPM3	1	1	3.5	2.5	2.5	3
CLOG3	2	3	3.5	5	4.5	4
LPM2	3	5	5.5	2.5	2.5	2
CLOG2	4	4	5.5	1	1	1
LPM4	5	2	1.5	4	4.5	5
CLOG4	6	6	1.5	7	7.5	8
CLOG1	7	7	7.5	7	6	6.5
LPM1	8	8	7.5	7	7.5	6.5
<i>Pearson's correlation between ranks:</i>						
for rates of return, 2001-07	0.83		0.63	0.66	0.64	0.54
for rates of return, 2004-07		0.83	0.39	0.71	0.74	0.71

Notes: * The model were ranked within their regression type (LPM and CLOG). To ranks were then multiplied by 1.5 to make them comparable with other rankings (ie. the first ranked model within the LPM group was assigned the rank of 1.5, the second ranked 3.5 etc). This is effectively the same to assuming that the best LPM and the best CLOG share the first place, the second bests share third place, etc.

Table 5 shows that none of the procedures – calculating statistical goodness of fit or running a regression tournament – successfully predicts which of the models will produce the highest betting returns. However, there are several indications that the regression tournament may represent a useful approach. First, by using the regression tournament winner (CLOG4) one would achieve higher returns than by using the goodness of fit favorites (LPM4 or CLOG4). Secondly, the model that gives the highest returns (LPM3) is ranked higher in all tournaments than it is by goodness of fit statistics. Finally, for the subperiod 2004-07, the rank correlations with the order of model based on simulated rates of return are significantly higher for the regression tournament rankings than for the goodness of fit rankings.

4. Conclusions

We have presented the idea of a regression tournament as a new procedure for comparative assessment of forecasting models. The basic idea is to take several models that forecast the same variable, impose the same decision rule on all of them, interact them with each other and then compare the relative performance of models in such pair wise interactions.

In our case, the variable forecast was the winning probability in a sports event; the decision rule was how to set odds / wage bets according to the estimated probabilities; the performance compared was the rate of betting profits made. We explored the empirical case of Australian Rules football, where the regression tournament procedure, compared to standard goodness of fit statistics, gave us better information on which model would yield higher real-life betting returns.

The main result of our paper is that standard goodness of fit statistic are not appropriate for choosing the best model for making real life decisions. Our proposed regression tournament is an interesting, although clearly not perfect alternative assessment method. These findings open up a vast area for further exploration. Why exactly do statistical properties of models not reveal their real-life utility? The obvious, but facile answer, that our models suffer from missing variables bias does not help because in reality this will always be true to some unknown (and unknowable) extent or other. How should regression tournaments be designed to yield useful new information on models' usefulness for real life choices? Could a regression tournament be used for comparing regression types (i.e. different logit and probit type models for which no comparable goodness of fit statistics exist)? Could it be used for more complex predictions, such as predicting the final position for all horses competing in a race?

The importance of our findings extends far behind the sports betting contexts. There are many other real world examples where the utility of one's model depends on interaction of his actions with actions taken by others according to their own forecasting models. Financial markets are an obvious example. It is not only my prediction of stock movements and my strategy that matters; the profits I can achieve depend also on actions taken by other traders based on their own models. Another example would be trying to protect the real value of my assets against inflation. It is not only my inflation expectations and my actions that matter; it is also the expectations and actions of others (price setters, governments, central banks etc.) that will interact with my decisions and co-determine the final outcome.

Standard methods to assess the statistical quality of models neglect such interactions. They implicitly assume there is only one person in the world, namely the forecaster with her model(s), and that there exists an objective and independent reality to which the model predictions may be compared. However, on many occasions, the reality to which we compare our predictions and in which we take our actions is codetermined and constantly changed by actions taken by other actors based on their own models. We propose that a regression tournament may be an interesting way of taking these interactions into account when assessing the real utility of our models.

References

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Eds.), *Frontiers in econometrics: Economic theory and mathematical economics* (pp. 105-142). New York: Academic Press.

Ramanathan, R.,(2002), *Introductory Econometrics with Applications* 5th ed., South-Western.

Schnytzer, A. and G. Weinberg (2008), "Testing for Home Team and Favorite Biases in the Australian Rules Football Fixed Odds and Point Spread Betting Markets", *Journal of Sports Economics*, Volume 9, No. 2, 2008, 173-190.

Schnytzer, A. (forthcoming), "The Prediction Market for the Australian Football League", in L. Vaughn Williams, *Prediction Markets*, Routledge.

Bar-Ilan University
Department of Economics
WORKING PAPERS

- 1-01 **The Optimal Size for a Minority**
Hillel Rapoport and Avi Weiss, January 2001.
- 2-01 **An Application of a Switching Regimes Regression to the Study of Urban Structure**
Gershon Alperovich and Joseph Deutsch, January 2001.
- 3-01 **The Kuznets Curve and the Impact of Various Income Sources on the Link Between Inequality and Development**
Joseph Deutsch and Jacques Silber, February 2001.
- 4-01 **International Asset Allocation: A New Perspective**
Abraham Lioui and Patrice Poncet, February 2001.
- 5-01 **מודל המועדון והקהילה החרדית**
יעקב רוזנברג, פברואר 2001.
- 6-01 **Multi-Generation Model of Immigrant Earnings: Theory and Application**
Gil S. Epstein and Tikva Lecker, February 2001.
- 7-01 **Shattered Rails, Ruined Credit: Financial Fragility and Railroad Operations in the Great Depression**
Daniel A. Schiffman, February 2001.
- 8-01 **Cooperation and Competition in a Duopoly R&D Market**
Damiano Bruno Silipo and Avi Weiss, March 2001.
- 9-01 **A Theory of Immigration Amnesties**
Gil S. Epstein and Avi Weiss, April 2001.
- 10-01 **Dynamic Asset Pricing With Non-Redundant Forwards**
Abraham Lioui and Patrice Poncet, May 2001.
- 11-01 **Macroeconomic and Labor Market Impact of Russian Immigration in Israel**
Sarit Cohen and Chang-Tai Hsieh, May 2001.

Electronic versions of the papers are available at
http://www.biu.ac.il/soc/ec/wp/working_papers.html

- 12-01 **Network Topology and the Efficiency of Equilibrium**
Igal Milchtaich, June 2001.
- 13-01 **General Equilibrium Pricing of Trading Strategy Risk**
Abraham Lioui and Patrice Poncet, July 2001.
- 14-01 **Social Conformity and Child Labor**
Shirit Katav-Herz, July 2001.
- 15-01 **Determinants of Railroad Capital Structure, 1830–1885**
Daniel A. Schiffman, July 2001.
- 16-01 **Political-Legal Institutions and the Railroad Financing Mix, 1885–1929**
Daniel A. Schiffman, September 2001.
- 17-01 **Macroeconomic Instability, Migration, and the Option Value of Education**
Eliakim Katz and Hillel Rapoport, October 2001.
- 18-01 **Property Rights, Theft, and Efficiency: The Biblical Waiver of Fines in the Case of Confessed Theft**
Eliakim Katz and Jacob Rosenberg, November 2001.
- 19-01 **Ethnic Discrimination and the Migration of Skilled Labor**
Frédéric Docquier and Hillel Rapoport, December 2001.
- 1-02 **Can Vocational Education Improve the Wages of Minorities and Disadvantaged Groups? The Case of Israel**
Shoshana Neuman and Adrian Ziderman, February 2002.
- 2-02 **What Can the Price Gap between Branded and Private Label Products Tell Us about Markups?**
Robert Barsky, Mark Bergen, Shantanu Dutta, and Daniel Levy, March 2002.
- 3-02 **Holiday Price Rigidity and Cost of Price Adjustment**
Daniel Levy, Georg Müller, Shantanu Dutta, and Mark Bergen, March 2002.
- 4-02 **Computation of Completely Mixed Equilibrium Payoffs**
Igal Milchtaich, March 2002.
- 5-02 **Coordination and Critical Mass in a Network Market – An Experimental Evaluation**
Amir Etziony and Avi Weiss, March 2002.

- 6-02 **Inviting Competition to Achieve Critical Mass**
Amir Etziony and Avi Weiss, April 2002.
- 7-02 **Credibility, Pre-Production and Inviting Competition in a Network Market**
Amir Etziony and Avi Weiss, April 2002.
- 8-02 **Brain Drain and LDCs' Growth: Winners and Losers**
Michel Beine, Frédéric Docquier, and Hillel Rapoport, April 2002.
- 9-02 **Heterogeneity in Price Rigidity: Evidence from a Case Study Using Micro-Level Data**
Daniel Levy, Shantanu Dutta, and Mark Bergen, April 2002.
- 10-02 **Price Flexibility in Channels of Distribution: Evidence from Scanner Data**
Shantanu Dutta, Mark Bergen, and Daniel Levy, April 2002.
- 11-02 **Acquired Cooperation in Finite-Horizon Dynamic Games**
Igal Milchtaich and Avi Weiss, April 2002.
- 12-02 **Cointegration in Frequency Domain**
Daniel Levy, May 2002.
- 13-02 **Which Voting Rules Elicit Informative Voting?**
Ruth Ben-Yashar and Igal Milchtaich, May 2002.
- 14-02 **Fertility, Non-Altruism and Economic Growth: Industrialization in the Nineteenth Century**
Elise S. Brezis, October 2002.
- 15-02 **Changes in the Recruitment and Education of the Power Elites in Twentieth Century Western Democracies**
Elise S. Brezis and François Crouzet, November 2002.
- 16-02 **On the Typical Spectral Shape of an Economic Variable**
Daniel Levy and Hashem Dezhbakhsh, December 2002.
- 17-02 **International Evidence on Output Fluctuation and Shock Persistence**
Daniel Levy and Hashem Dezhbakhsh, December 2002.
- 1-03 **Topological Conditions for Uniqueness of Equilibrium in Networks**
Igal Milchtaich, March 2003.
- 2-03 **Is the Feldstein-Horioka Puzzle Really a Puzzle?**
Daniel Levy, June 2003.

- 3-03 **Growth and Convergence across the US: Evidence from County-Level Data**
Matthew Higgins, Daniel Levy, and Andrew Young, June 2003.
- 4-03 **Economic Growth and Endogenous Intergenerational Altruism**
Hillel Rapoport and Jean-Pierre Vidal, June 2003.
- 5-03 **Remittances and Inequality: A Dynamic Migration Model**
Frédéric Docquier and Hillel Rapoport, June 2003.
- 6-03 **Sigma Convergence Versus Beta Convergence: Evidence from U.S. County-Level Data**
Andrew T. Young, Matthew J. Higgins, and Daniel Levy, September 2003.
- 7-03 **Managerial and Customer Costs of Price Adjustment: Direct Evidence from Industrial Markets**
Mark J. Zbaracki, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, September 2003.
- 8-03 **First and Second Best Voting Rules in Committees**
Ruth Ben-Yashar and Igal Milchtaich, October 2003.
- 9-03 **Shattering the Myth of Costless Price Changes: Emerging Perspectives on Dynamic Pricing**
Mark Bergen, Shantanu Dutta, Daniel Levy, Mark Ritson, and Mark J. Zbaracki, November 2003.
- 1-04 **Heterogeneity in Convergence Rates and Income Determination across U.S. States: Evidence from County-Level Data**
Andrew T. Young, Matthew J. Higgins, and Daniel Levy, January 2004.
- 2-04 **“The Real Thing:” Nominal Price Rigidity of the Nickel Coke, 1886-1959**
Daniel Levy and Andrew T. Young, February 2004.
- 3-04 **Network Effects and the Dynamics of Migration and Inequality: Theory and Evidence from Mexico**
David Mckenzie and Hillel Rapoport, March 2004.
- 4-04 **Migration Selectivity and the Evolution of Spatial Inequality**
Ravi Kanbur and Hillel Rapoport, March 2004.
- 5-04 **Many Types of Human Capital and Many Roles in U.S. Growth: Evidence from County-Level Educational Attainment Data**
Andrew T. Young, Daniel Levy and Matthew J. Higgins, March 2004.

- 6-04 **When Little Things Mean a Lot: On the Inefficiency of Item Pricing Laws**
Mark Bergen, Daniel Levy, Sourav Ray, Paul H. Rubin and Benjamin Zeliger,
May 2004.
- 7-04 **Comparative Statics of Altruism and Spite**
Igal Milchtaich, June 2004.
- 8-04 **Asymmetric Price Adjustment in the Small: An Implication of Rational Inattention**
Daniel Levy, Haipeng (Allan) Chen, Sourav Ray and Mark Bergen, July 2004.
- 1-05 **Private Label Price Rigidity during Holiday Periods**
Georg Müller, Mark Bergen, Shantanu Dutta and Daniel Levy, March 2005.
- 2-05 **Asymmetric Wholesale Pricing: Theory and Evidence**
Sourav Ray, Haipeng (Allan) Chen, Mark Bergen and Daniel Levy,
March 2005.
- 3-05 **Beyond the Cost of Price Adjustment: Investments in Pricing Capital**
Mark Zbaracki, Mark Bergen, Shantanu Dutta, Daniel Levy and Mark Ritson,
May 2005.
- 4-05 **Explicit Evidence on an Implicit Contract**
Andrew T. Young and Daniel Levy, June 2005.
- 5-05 **Popular Perceptions and Political Economy in the Contrived World of Harry Potter**
Avichai Snir and Daniel Levy, September 2005.
- 6-05 **Growth and Convergence across the US: Evidence from County-Level Data (revised version)**
Matthew J. Higgins, Daniel Levy, and Andrew T. Young , September 2005.
- 1-06 **Sigma Convergence Versus Beta Convergence: Evidence from U.S. County-Level Data (revised version)**
Andrew T. Young, Matthew J. Higgins, and Daniel Levy, June 2006.
- 2-06 **Price Rigidity and Flexibility: Recent Theoretical Developments**
Daniel Levy, September 2006.
- 3-06 **The Anatomy of a Price Cut: Discovering Organizational Sources of the Costs of Price Adjustment**
Mark J. Zbaracki, Mark Bergen, and Daniel Levy, September 2006.

- 4-06 **Holiday Non-Price Rigidity and Cost of Adjustment**
Georg Müller, Mark Bergen, Shantanu Dutta, and Daniel Levy.
September 2006.
- 2008-01 **Weighted Congestion Games With Separable Preferences**
Igal Milchtaich, October 2008.
- 2008-02 **Federal, State, and Local Governments: Evaluating their Separate Roles in US Growth**
Andrew T. Young, Daniel Levy, and Matthew J. Higgins, December 2008.
- 2008-03 **Political Profit and the Invention of Modern Currency**
Dror Goldberg, December 2008.
- 2008-04 **Static Stability in Games**
Igal Milchtaich, December 2008.
- 2008-05 **Comparative Statics of Altruism and Spite**
Igal Milchtaich, December 2008.
- 2008-06 **Abortion and Human Capital Accumulation: A Contribution to the Understanding of the Gender Gap in Education**
Leonid V. Azarnert, December 2008.
- 2008-07 **Involuntary Integration in Public Education, Fertility and Human Capital**
Leonid V. Azarnert, December 2008.
- 2009-01 **Inter-Ethnic Redistribution and Human Capital Investments**
Leonid V. Azarnert, January 2009.
- 2009-02 **Group Specific Public Goods, Orchestration of Interest Groups and Free Riding**
Gil S. Epstein and Yosef Mealem, January 2009.
- 2009-03 **Holiday Price Rigidity and Cost of Price Adjustment**
Daniel Levy, Haipeng Chen, Georg Müller, Shantanu Dutta, and Mark Bergen,
February 2009.
- 2009-04 **Legal Tender**
Dror Goldberg, April 2009.
- 2009-05 **The Tax-Foundation Theory of Fiat Money**
Dror Goldberg, April 2009.

- 2009-06 **The Inventions and Diffusion of Hyperinflatable Currency**
Dror Goldberg, April 2009.
- 2009-07 **The Rise and Fall of America's First Bank**
Dror Goldberg, April 2009.
- 2009-08 **Judicial Independence and the Validity of Controverted Elections**
Raphaël Franck, April 2009.
- 2009-09 **A General Index of Inherent Risk**
Adi Schnytzer and Sara Westreich, April 2009.
- 2009-10 **Measuring the Extent of Inside Trading in Horse Betting Markets**
Adi Schnytzer, Martien Lamers and Vasiliki Makropoulou, April 2009.
- 2009-11 **The Impact of Insider Trading on Forecasting in a Bookmakers' Horse Betting Market**
Adi Schnytzer, Martien Lamers and Vasiliki Makropoulou, April 2009.
- 2009-12 **Foreign Aid, Fertility and Population Growth: Evidence from Africa**
Leonid V. Azarnert, April 2009.
- 2009-13 **A Reevaluation of the Role of Family in Immigrants' Labor Market Activity: Evidence from a Comparison of Single and Married Immigrants**
Sarit Cohen-Goldner, Chemi Gotlibovski and Nava Kahana, May 2009.
- 2009-14 **The Efficient and Fair Approval of "Multiple-Cost–Single-Benefit" Projects Under Unilateral Information**
Nava Kahanaa, Yosef Mealem and Shmuel Nitzan, May 2009.
- 2009-15 **Après nous le Déluge: Fertility and the Intensity of Struggle against Immigration**
Leonid V. Azarnert, June 2009.
- 2009-16 **Is Specialization Desirable in Committee Decision Making?**
Ruth Ben-Yashar, Winston T.H. Koh and Shmuel Nitzan, June 2009.
- 2009-17 **Framing-Based Choice: A Model of Decision-Making Under Risk**
Kobi Kriesler and Shmuel Nitzan, June 2009.
- 2009-18 **Demystifying the 'Metric Approach to Social Compromise with the Unanimity Criterion'**
Shmuel Nitzan, June 2009.

- 2009-19 **On the Robustness of Brain Gain Estimates**
Michel Beine, Frédéric Docquier and Hillel Rapoport, July 2009.
- 2009-20 **Wage Mobility in Israel: The Effect of Sectoral Concentration**
Ana Rute Cardoso, Shoshana Neuman and Adrian Ziderman, July 2009.
- 2009-21 **Intermittent Employment: Work Histories of Israeli Men and Women, 1983–1995**
Shoshana Neuman and Adrian Ziderman, July 2009.
- 2009-22 **National Aggregates and Individual Disaffiliation: An International Study**
Pablo Brañas-Garza, Teresa García-Muñoz and Shoshana Neuman, July 2009.
- 2009-23 **The Big Carrot: High-Stakes Incentives Revisited**
Pablo Brañas-Garza, Teresa García-Muñoz and Shoshana Neuman, July 2009.
- 2009-24 **The Why, When and How of Immigration Amnesties**
Gil S. Epstein and Avi Weiss, September 2009.
- 2009-25 **Documenting the Brain Drain of «la Crème de la Crème»: Three Case-Studies on International Migration at the Upper Tail of the Education Distribution**
Frédéric Docquier and Hillel Rapoport, October 2009.
- 2009-26 **Remittances and the Brain Drain Revisited: The Microdata Show That More Educated Migrants Remit More**
Albert Bollard, David McKenzie, Melanie Morten and Hillel Rapoport, October 2009.
- 2009-27 **Implementability of Correlated and Communication Equilibrium Outcomes in Incomplete Information Games**
Igal Milchtaich, November 2009.
- 2010-01 **The Ultimatum Game and Expected Utility Maximization – In View of Attachment Theory**
Shaul Almakias and Avi Weiss, January 2010.
- 2010-02 **A Model of Fault Allocation in Contract Law – Moving From Dividing Liability to Dividing Costs**
Osnat Jacobi and Avi Weiss, January 2010.

- 2010-03 **Coordination and Critical Mass in a Network Market: An Experimental Investigation**
Bradley J. Ruffle, Avi Weiss and Amir Etziony, February 2010.
- 2010-04 **Immigration, fertility and human capital: A model of economic decline of the West**
Leonid V. Azarnert, April 2010.
- 2010-05 **Is Skilled Immigration Always Good for Growth in the Receiving Economy?**
Leonid V. Azarnert, April 2010.
- 2010-06 **The Effect of Limited Search Ability on the Quality of Competitive Rent-Seeking Clubs**
Shmuel Nitzan and Kobi Kriesler, April 2010.
- 2010-07 **Condorcet vs. Borda in Light of a Dual Majoritarian Approach**
Eyal Baharad and Shmuel Nitzan, April 2010.
- 2010-08 **Prize Sharing in Collective Contests**
Shmuel Nitzan and Kaoru Ueda, April 2010.
- 2010-09 **Network Topology and Equilibrium Existence in Weighted Network Congestion Games**
Igal Milchtaich, May 2010.
- 2010-10 **The Evolution of Secularization: Cultural Transmission, Religion and Fertility Theory, Simulations and Evidence**
Ronen Bar-El, Teresa García-Muñoz, Shoshana Neuman and Yossef Tobol, June 2010.
- 2010-11 **The Economics of Collective Brands**
Arthur Fishman, Israel Finkelstein, Avi Simhon and Nira Yacouel, July 2010.
- 2010-12 **Interactions Between Local and Migrant Workers at the Workplace**
Gil S. Epstein and Yosef Mealem, August 2010.
- 2010-13 **A Political Economy of the Immigrant Assimilation: Internal Dynamics**
Gil S. Epstein and Ira N. Gang, August 2010.
- 2010-14 **Attitudes to Risk and Roulette**
Adi Schnytzer and Sara Westreich, August 2010.

- 2010-15 **Life Satisfaction and Income Inequality**
Paolo Verme, August 2010.
- 2010-16 **The Poverty Reduction Capacity of Private and Public Transfers in Transition**
Paolo Verme, August 2010.
- 2010-17 **Migration and Culture**
Gil S. Epstein and Ira N. Gang, August 2010.
- 2010-18 **Political Culture and Discrimination in Contests**
Gil S. Epstein, Yosef Mealem and Shmuel Nitzan, October 2010.
- 2010-19 **Governing Interest Groups and Rent Dissipation**
Gil S. Epstein and Yosef Mealem, November 2010.
- 2010-20 **Beyond Condorcet: Optimal Aggregation Rules Using Voting Records**
Eyal Baharad, Jacob Goldberger, Moshe Koppel and Shmuel Nitzan, December 2010.
- 2010-21 **Price Points and Price Rigidity**
Daniel Levy, Dongwon Lee, Haipeng (Allan) Chen, Robert J. Kauffman and Mark Bergen, December 2010.
- 2010-22 **Price Setting and Price Adjustment in Some European Union Countries: Introduction to the Special Issue**
Daniel Levy and Frank Smets, December 2010.
- 2011-01 **Business as Usual: A Consumer Search Theory of Sticky Prices and Asymmetric Price Adjustment**
Luís Cabral and Arthur Fishman, January 2011.
- 2011-02 **Emigration and democracy**
Frédéric Docquier, Elisabetta Lodigiani, Hillel Rapoport and Maurice Schiff, January 2011.
- 2011-03 **Shrinking Goods and Sticky Prices: Theory and Evidence**
Avichai Snir and Daniel Levy, March 2011.
- 2011-04 **Search Costs and Risky Investment in Quality**
Arthur Fishman and Nadav Levy, March 2011.

- 2011-05 **To What Extent do Investors in a Financial Market Anchor Their Judgments? Evidence from the Hong Kong Horserace Betting Market**
Johnnie E.V. Johnson, Shuang Liu and Adi Schnytzer, March 2011.
- 2011-06 **Attitudes to Risk and Roulette**
Adi Schnytzer and Sara Westreich, March 2011.
- 2011-07 **False Consciousness in Financial Markets: Or is it in Ivory Towers?**
Adi Schnytzer and Sara Westreich, March 2011.
- 2011-08 **Herding in Imperfect Betting Markets with Inside Traders**
Adi Schnytzer and Avichai Snir, March 2011.
- 2011-09 **Painful Regret and Elation at the Track**
Adi Schnytzer and Barbara Luppi, March 2011.
- 2011-10 **The Regression Tournament: A Novel Approach to Prediction Model Assessment**
Adi Schnytzer and Janez Šušteršič, March 2011.
- 2011-11 **Shorting the Bear: A test of anecdotal evidence of insider trading in early stages of the sub-prime market crisis**
Les Coleman and Adi Schnytzer, March 2011.
- 2011-12 **SP Betting as a Self-Enforcing Implicit Cartel**
Adi Schnytzer and Avichai Snir, March 2011.
- 2011-13 **Testing for Home Team and Favorite Biases in the Australian Rules Football Fixed Odds and Point Spread Betting Markets**
Adi Schnytzer and Guy Weinberg, March 2011.
- 2011-14 **The Impact of Insider Trading on Forecasting in a Bookmakers' Horse Betting Market**
Adi Schnytzer, Martien Lamers and Vasiliki Makropoulou, March 2011.
- 2011-15 **The Prediction Market for the Australian Football League**
Adi Schnytzer, March 2011.

2011-16 **Information and Attitudes to Risk at the Track**

Adi Schnytzer and Sara Westreich, March 2011.

2011-17 **Explicit Evidence on an Implicit Contract**

Andrew T. Young and Daniel Levy, March 2011.

2011-18 **Globalization, Brain Drain and Development**

Frédéric Docquier and Hillel Rapoport, March 2011.