

Instructor Rating Markets

Mithun Chakraborty, Sanmay Das, Allen Lavoie, Malik Magdon-Ismael, and
Yonatan Naamad

Rensselaer Polytechnic Institute, Troy NY 12180, USA

Abstract. We describe the design of *Instructor Rating Markets* in which students trade on the ratings that will be received by instructors, with new ratings revealed every two weeks. The markets provide useful dynamic feedback to instructors on the progress of their class, while at the same time enabling the controlled study of prediction markets where traders can affect the outcomes they are trading on. More than 200 students across the Rensselaer campus participated in markets for ten classes in the Fall 2010 semester. We show that market prices convey useful information on *future* instructor ratings and contain significantly more information than do past ratings. The bulk of useful information contained in the price of a particular class is provided by students who are in that class, showing that the markets are serving to disseminate insider information. At the same time, we find little evidence of attempted manipulation of the liquidating dividends by raters. The markets are also a laboratory for comparing different microstructures and the resulting price dynamics, and we show how they can be used to compare market making algorithms.

We present a novel application of prediction markets to instructor evaluations. Such markets have the potential to provide dynamic feedback on the progress of a class. We describe a pilot deployment of these markets at Rensselaer Polytechnic Institute in the Fall of 2010, with more than 200 students participating across 10 classes. These markets provide insights into the behavior of students in their roles as both traders and, potentially, as market manipulators (traders who are in a class directly affect the rating of that class), while also allowing us to study how market microstructure affects price formation and the information content of prices.

Each instructor-course pair is an openly traded security in the IRM. Every two weeks, each security pays a liquidating dividend derived from how students in the class rate the instructor for that two week period. Each security can be traded by anyone at the institute, but only students who are in the instructor's class may rate the instructor. A rating period opens after the first week of trading, and students who have "in class" credentials receive an email asking them to rate the instructor of their class – the rating period stays open until the end of the second week, at which point both the rating and trading windows close. If everything works well, fluctuations in the price of the "instructor security" give *real-time feedback* on how well the instructor is doing (we are not endorsing teaching to maximize "stock value", but instructor ratings exist, and it is useful to know

more about what goes into student ratings, and how they would change on a day-to-day basis if students were “polled” repeatedly). Thus, we use students, as well as their roommates and friends, as information gatherers, giving them an outlet (a fun trading game) to reveal their information. While the instructor is only rated occasionally, price movements provide continuous feedback.

There are two major differences between the IRMs and more traditional prediction markets. First, in many prediction markets, information revelation continues right up to the moment of liquidation (for example, opinion polls are released continuously during election cycles), whereas in our markets the only major information revelation is the liquidation event itself. The information revelation leading up to liquidation in IRMs is considerably more noisy (Did the instructor give a good lecture? Was there a hard homework due that week?). Second, typical large prediction markets, such as election markets, attempt to predict a much more stable statistical aggregate quantity: voting turnouts range from the tens of thousands to the tens or hundreds of millions. In contrast, the classes the IRMs ran on in our deployment had between 3 and 25 regular raters. This raises questions about the effects of insider information and potential market manipulation. The success of the markets in predicting instructor ratings is not a given.

However, we find that prices are, in fact, predictive of future instructor ratings, and significantly more predictive than are previous ratings, showing that they incorporate new information. The higher predictivity is due to the trades of insiders: our data shows that when previous and future liquidations differ, students who are enrolled in a class trade in the direction of future liquidations while others trade in the direction of the last liquidation. We also find little evidence of efforts by students to manipulate the ratings for their own benefit as traders: first, the ratings had very high correlation with the official end-of-semester student evaluations of the classes, and, second, we found few cases where students, either individually or in groups, gave surprising ratings and profited from doing so. The fact that IRM ratings are well aligned with the official end-of-semester evaluations shows that the system as a whole is relevant and useful to instructors. Combining that fact with the power of prices to predict IRM ratings is encouraging for the potential of such markets.

In addition to our primary results, we also document learning behavior along several dimensions. In particular, prices for more predictable securities become more efficient, and an early “in class” optimistic bias in traded prices disappears in later periods. The markets also have other beneficial side effects: for example, active traders are more likely to give ratings, thus providing instructors with useful feedback every two weeks. This is already an achievement over the considerably less dynamic single end-of-semester ratings typically available. Finally, we can use the IRM to study the effects of different market microstructures. In particular, we provide further validation of a Bayesian market-making algorithm, BMM, that can provide more price stability than the standard Logarithmic Market Scoring Rule (LMSR) market maker while also making more profit.