

Conjunction and Disjunction Fallacies in Prediction Markets

Michael D. Lee (mdlee@uci.edu)

Emily Grothe (egrothe@uci.edu)

Mark Steyvers (msteyver@uci.edu)

Department of Cognitive Sciences, 3151 Social Sciences Plaza A
University of California, Irvine, CA 92697-5100 USA

Abstract

Prediction markets provide a mechanism for using groups of people to determine the probability of events. We ask whether these probability estimates, for logically related events, exhibit the irrationalities often found in individual judgments. In particular, we explore combinations of markets that provide tests of classic conjunction and disjunction fallacies. Across a number of markets, asking about a wide variety of events, we find a few interesting violations of probability theory, but mostly rational adherence. We discuss our exploratory analyses in terms of the relationship between group and individual probability estimation, and the effectiveness of prediction markets.

Keywords: Conjunction fallacy, disjunction fallacy, prediction markets, group estimation, heuristics and biases

Introduction

Few lines of research in cognitive science are as well known and influential in other fields and endeavours as the heuristic and biases program (Tversky & Kahneman, 1974). The idea that human decision-makers proceed on shaky grounds—making choices that are often incorrect, sometimes worryingly irrational, and occasionally close to bizarre—is treated seriously in all areas of academia and business that need to understand and use people’s judgments.

One of the best-studied and most central parts of heuristics and biases research explores how people with probabilities for events that bear some logical relationship to each other. Most famously, the conjunction fallacy shows that people can judge a specific event to be more probable than a more general class of events to which it belongs (e.g., Tversky & Kahneman, 1983). After learning about a person called Linda, described in a way consistent with being a feminist, she is often judged to more likely be a feminist bank teller than simply a bank teller.

In a similar violation of probability theory, disjunction fallacies show people judging a general event to have a different probability from the sum of the probabilities of its separate components (e.g., Wright & Whalley, 1983). Asking people the individual probability that they will be doing different, and mutually exclusive, activities at a

specific time in the future, then summing those probabilities, often returns a value greater than one.

This paper asks whether these sorts of classic conjunction and disjunction fallacies are also seen in the behavior of groups of people, when they work together in the unique context provided by prediction markets.

Prediction Markets

In a prediction market, people trade contracts related to events. These events can be about anything—politics, sport, entertainment, climate, and so on—for which there will be a clear true-or-false answer known in a reasonably short period of time. Events in a prediction market could include “Brazil wins the 2010 FIFA World Cup”, “Steve Jobs to depart as CEO of Apple on or before 31 Dec 2009”, and “Bernard Madoff to plead guilty or be found guilty of Securities Fraud.”

In a standard framework, the value of a contact is bounded between 0 and 100, and can be bought or sold. Traders owning a contact can specify a price at which they are willing to sell. Potential buyers can specifically a price at which they are willing to buy. When the price of a buyer and a seller meet, a transaction takes place, and the *value* of the contract is updated to represent this last traded price. When the event itself is decided, and known to be either true or false, the value of the contracts are realized. Those contracts for true events pay 100, while those that are false pay 0.

Under this trading system, it is natural and reasonable to view the value of a contract as corresponding to an estimate of the probability of the event being true. The Intrade¹ prediction market we study in this paper explicitly encourages a probabilistic interpretation of market values. It says “[s]ince our contracts trade between 0 and 100, you can think of the price at any time to be the percentage probability of that event occurring,” and gives a tutorial on the logic of trading in terms of buying if you think a contract is being offered at a value below the true probability of the event, and selling if you can find a buyer willing to pay more than what you judge to be the true probability.

Not only is the idea that prediction market values represent probability estimates well-established, but the values themselves are widely trusted. There is empirical

¹See <http://www.intrade.com>

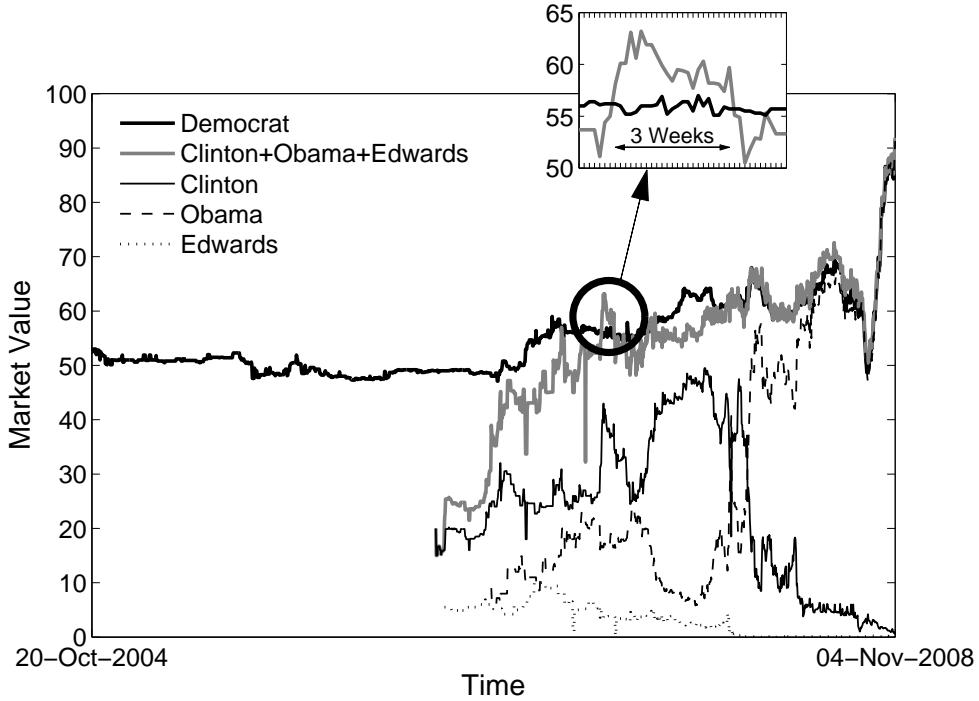


Figure 1: Prediction market values for three individual presidential candidates (Clinton, Obama and Edwards), their total (Clinton+Obama+Edwards) and for the Democratic nominee. A three week period in which the sum of the markets for the individual Democratic candidates is greater than the general Democratic nominee market is highlighted.

evidence (see, for example Christiansen, 2007) that the probabilities produced by prediction markets can be consistently excellent guides to the final outcome for a wide variety of events, and important decision-makers and organizations endorse the predictive ability of these markets. Nevertheless, since these probabilities are, ultimately, produced by people, we wondered whether they are subject to the same sorts of logical fallacies found in heuristics and biases research.

Do Prediction Markets Show Fallacies?

One of the most heavily traded Intrade prediction markets asked whether or not Hillary Clinton would win the 2008 US Presidential Election. Another heavily traded market asked the related, but more general question, of whether the Democratic party nominee would win the 2008 US Presidential Election. The close relationship between these markets invites the possibility of a conjunction fallacy. Just as Linda is more likely to be a bank teller than a feminist bank teller, the Democratic nominee is more likely to win than Clinton. If these two prediction markets are making rational probability judgments, the value for Clinton should never exceed the value of the general Democratic nominee.

Nevertheless, just as it might be easier to imagine Linda as a feminist bank teller than a generic bank teller, it might be easier to imagine the inauguration of Clinton than the more abstract idea of the inauguration of the

Democratic nominee. So, it seems worth checking how the specific Clinton and the general Democratic nominee markets are valued at the same points in time, over their trading histories. The same analysis makes sense for other individual Democrats—including Barack Obama and John Edwards—who had their own heavily traded markets for winning the 2008 US Presidential Election. Figure 1 shows the change in market value for the individual Clinton, Obama and Edwards markets, as well as the Democratic nominee market, over the period in which they existed. It is clear that none of the individuals ever exceed the value of their party, and so there is no conjunction fallacy.

Figure 1 shows another analysis, however, that does point towards an irrationality in these prediction markets. Since only one person can win the election, the probabilities of Clinton, Obama and Edwards combined should not exceed the probability for the Democratic nominee. But, as the gray line shows, by adding the market values of the three individuals, total probabilities are obtained that exceed the Democratic value at several points in time. The most severe of the violation is highlighted by the circle in Figure 1, and a zoomed-in look at the curves in this region is shown. For about a 3 week period, the probability that Clinton or Obama or Edwards would win the election was judged by the Intrade markets as higher—by more than 5% for much of the time—than

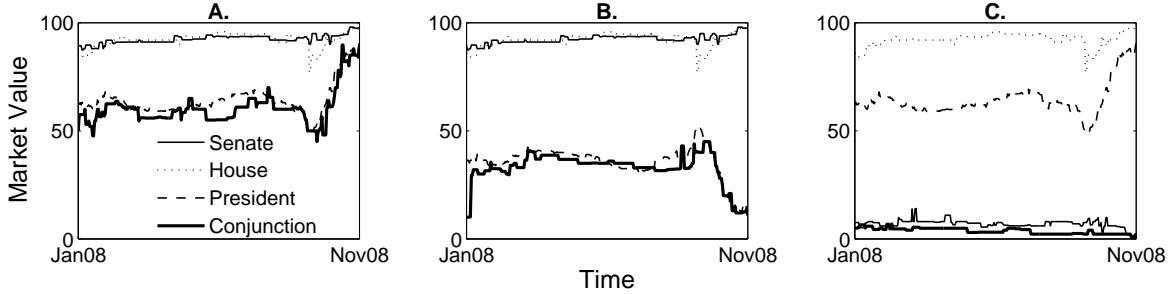


Figure 2: Prediction market values for individual propositions about whether Democrats or Republicans win the Presidential, House and Senate elections, and the value for the market whose event is the logical conjunction of these three individual markets. The three different analyses correspond to three different scenarios, as described in the text.

the probability of the Democratic nominee winning. This pattern is essentially a form of the disjunction fallacy, since the sum of the probabilities of mutually exclusive events add to more than the probability of their union.

To be fair, it is possible to imagine circumstances in which the violation we just observed is not so irrational. It would have been possible, for example, for Clinton to have won the US Presidential election as an independent. Under that improbable, but not impossible, scenario, the sum of the three individuals could exceed the Democratic nominee. We think that it is unlikely this sort of thinking is responsible for the data in Figure 1, and we will present a little evidence later to support our doubts. But it is not logically impossible. So, to provide sounder tests of whether prediction markets fall prey to conjunction and disjunction fallacies, we need to consider different combinations of markets that do make various combinations of probabilities logically impossible.

Overview

In this paper, we use some Intrade prediction markets to look for possible conjunction or disjunction fallacies. Our approach is exploratory, intended to provoke thinking about the rationality of markets, and the applicability of heuristics and biases to group decision-making. We examined a large range of logically related markets, most of which were heavily traded in the real-world high-stakes setting provided by Intrade, and present the highlights of those explorations here.

The most basic question we ask is whether we can find evidence of conjunction or disjunction fallacies in prediction markets. We are especially interested in substantial violations, where market values show large violations for sustained periods of time. We are interested in when and why such violations occur, and whether or how quickly they are corrected. An important theoretical question raised by our line of study involves the relationship between individual and group estimation of probabilities, especially as they are made for logically related events. An important applied question asks about the rationality of prediction markets. We will discuss these issues in the light of our exploratory findings.

Selected Analyses

Looking through the Intrade markets available at the end of 2008, we found a rich set of possibilities to look for logical fallacies in the values of related markets. We restricted our attention to those markets that were heavily enough traded on a day-by-day basis for the market value to be a meaningful running probability estimate. These markets were typically found in the realms of politics, especially in the context of the 2008 US elections, and in markets requiring predictions of the values of various financial and economic measures. Some also involved more general world events, including military- and climate-related markets. In this section, we report on three lines of analysis in some detail, giving a set of results that are indicative of what we found.

Political Conjunctions

Not surprisingly, Intrade had prediction markets for events like ‘Democrats win Presidency’, ‘Republicans win House’, and ‘Democrats win Senate’. Besides these separate markets, Intrade had a number of markets dealing with combinations of results, like ‘Democrats win the Presidency and House, but Republicans win the Senate’. Comparing the market values for the individual events with the market for the intersection of those events provides an obvious and direct test for conjunction fallacies.

We looked at all of the possible combinations, and show three representative analyses in Figure 2. In Figure 2a, the individual markets for the Democrats winning the presidency, the Democrats winning the house, and the Democrats winning the senate are shown, together with the single market for the event “The Democrats win the presidency, house and senate”. This final market, representing the conjunction of the first three, cannot logically be more likely than the individual markets. This requirement seems to be met by the prediction markets values. There are certainly no large and significant violations.

Similarly, Figure 2b corresponds to the markets for a Democratic president, but Republican control of the house and senate, and their logical conjunction. Figure 2c corresponds to Democratic presidency and con-

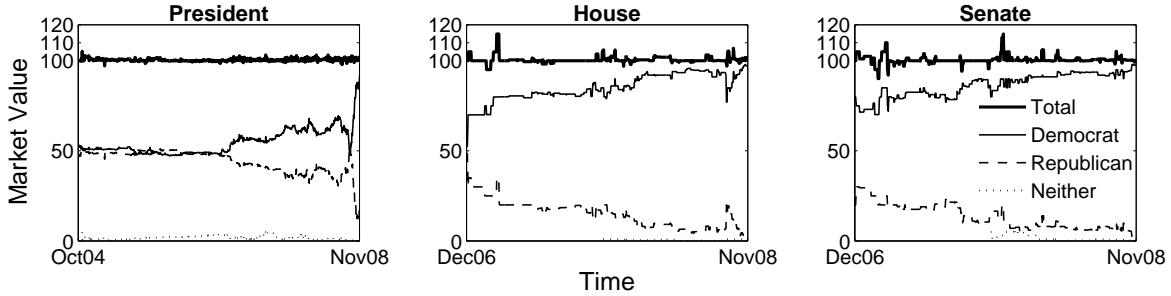


Figure 3: Prediction market values for the ‘Democrat’, ‘Republican’ and ‘Neither’ choices, and the sum total of values for these three markets, for the presidential, house and senate races.

trol of the house, but Republican control of the senate, and their conjunction market. Looking at all three analysis together gives the basic conclusion we drew from looking at all possibilities. This is that we could find no evidence of a conjunction fallacy, even though the actual pattern and level of the probabilities vary significantly under different scenarios.

Political Disjunctions

Figure 3 shows three analyses dealing with political markets that directly test for disjunction fallacies. All three relate to the 2008 US Elections, and the probabilities of the Democratic party, the Republican party, or neither party, achieved electoral success. The three analyses relate to the presidential, house and senate races. In each case, of course, the ‘Democrat’, ‘Republican’ and ‘Neither’ choices partition the total possibilities, and so the sum of their individual markets should be 100.

The left panel of Figure 3 suggests that this additive relationship is approximately satisfied in the presidential markets. The line showing the sum of the three market values hovers consistently close to 100 throughout, even though in the last year or two there are large and sometimes sudden shifts in how the Democratic and Republican components contribute to the sum. The house and senate races tell a different story, with quite large violations evident in short bursts in both totals. These violations seem to be more often over- than under-estimations of the total probability (i.e., they tend to show the supra-additivity usually found in individuals), although some instances of both are evident.

Taken together, the house and senate totals show that disjunctive fallacies are observed in prediction markets. It is interesting to note that, while the violations are often significant in magnitude—exceeding 10% in several cases—they are not sustained. We cannot know whether this quick correction comes from traders observing the mis-calibration of the relevant markets, or is a natural rapid adjustment of the individual markets themselves to match expectations. Indeed, we do not know whether the same traders are actively engaged, or even aware, of all of the relevant markets displayed in Figure 3. What we do know is that, to the extent the market value represents

the belief a group of people have about the probability of a series of related events, those probabilities are often very well calibrated, but do sometimes briefly take irrational values.

There are two other observations to make about the analyses in Figure 3. The first is that the partitioning into Democratic, Republican, and a catch-all ‘Neither’ choice provides a logically sound test of disjunction. Unlike our introductory example with Clinton, Obama and Edwards, the current analyses do not have alternative scenarios under which the deviation of the total market value from 100 can be justified. The second point is the presidential analysis in the left panel of Figure 3 argues against the possible explanation we considered earlier, in which Clinton could win the presidency as an independent. If this were the thinking, we would expect the ‘Neither’ market would sharply increase at the time the sum of the Clinton, Obama and Edwards curves in Figure 1 exceeded the Democrat market, but it does not.

Conjunctions in Snow, War, Jobs and Tax

A general class of markets that allows for additional testing for the conjunction fallacy involves the estimation of whether a quantity is greater or less than a specified value. Perhaps the most intuitive example in Intrade rests on an event like “Central Park to get MORE than 10 inches of snowfall in Dec 2008”. Intrade has identically-worded markets for this snowfall, ranging from 10 inches, to 15 inches, 20 inches, and so on, all the way up to 40 inches. This means that, for example, if Central Park recorded 23 inches of snowfall in Dec 2008, the 10, 15 and 20 inches markets would all pay out.

Logically, if more than 15 inches of snow is recorded, more than 10 inches is also recorded. And so the market value, estimating the probability of each specific level of snowfall, should decrease as the numbers themselves increase. This is a form of logical conjunction, in which more than 15 inches is both itself, and more than 10 inches. The more general events, which are the lower levels of snowfall for this “more than” question, should be the most probable.

The markets dealing with snowfall in Central Park were only lightly traded, and so we did not analyze them.

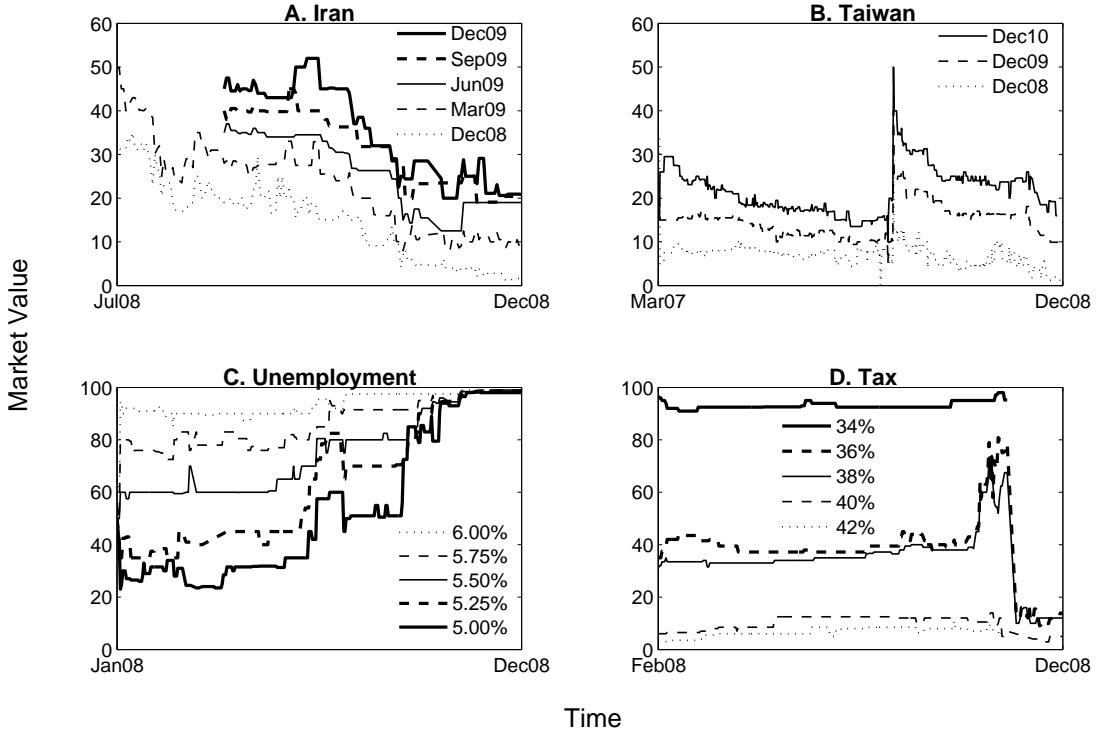


Figure 4: Prediction market values for (A) “USA and/or Israel to execute an overt Air Strike against Iran by [month]”, (B) “China to conduct overt military action against Taiwan on/before [month]”, (C) “US Unemployment Rate in Dec 2008 to be greater than [rate]”, (D) “Highest Marginal Single-Filer Fed Income Tax Rate to be Equal or Greater than [rate] in 2009 Tax Year”.

Instead, Figure 4 shows the results for four sets of heavily traded markets testing the conjunction fallacy in the same way.

Figure 4a deals with a possible US or Israeli air strike against Iran, with different markets specifying different months before which the strike must take place. The value of all of these markets generally declines over the time period shown, but the conjunctive orderings are almost always preserved. Other than a brief period around about October 2008 where the Sep09 value exceeded the Dec09 value—and this could possibly be the result of too little trading to keep the Sep09 estimate current—the market values maintain their logically required ordering.

Figure 4b involves a similar event, asking whether China will conduct overt military action again Taiwan before a specified month. The three markets show a clear step-change, in which the probability of action increases suddenly. Despite this rapid readjustment, however, the logically required ordering is preserved throughout.

Figure 4c considers markets for the US unemployment rate in December 2008, specifying different percentage rates that need to be exceeded. All of the specified values—ranging from 5.00 to 6.00%—define propositions that become progressively more likely in the economic situation leading up to December 2008, and so the markets increase in value towards 100. They do this al-

most always adhering to the logical ordering. There is only a brief and small violation for 5.25% and 5.50%.

Finally, Figure 4d looks at the US tax rate for filing individuals, with specific markets corresponding to rates between 34% and 42%. There is an interesting clumped structure to these markets, with 34% always being assessed as probably low enough to be true, 36% and 38% assessed very similarly, originally at about a value of 40, but then rising and falling dramatically and in unison, and 40% and 42% being similarly coupled, but with a steady and low probability of being true. But again, there are no significant conjunction fallacies.

Taken together, these analyses are representative of our broader explorations. We did not find any large and sustained conjunction fallacies by looking for violations of required orderings in multiple markets with different levels of the same quantity.

Discussion

The existence of conjunction and disjunction fallacies is one of the most robust findings in experimental psychology. Looking at the probabilities produced by prediction markets, we found some evidence for disjunction fallacies, but only the merest hint for conjunction fallacies. Of course, our explorations were limited. But, we did

not find the large and sustained effects that might be expected from experimental psychology. We now discuss some ideas that might help understand this discrepancy.

One important consideration is that prediction markets produce probability estimates from groups of people, not the individuals tested in experimental psychology. But there is nothing inherent about using groups of people, rather than individuals, that guarantees the prevention of irrational decision-making. Rioting crowds and stock market bubbles are common events. Indeed, a stock bubble in a prediction market—so that, for example, Hillary Clinton becomes significantly over-valued—is exactly the sort of phenomenon that could potentially produce fallacies. And some empirical psychology studies have demonstrated fallacies in group settings (e.g., Tindale, 1993).

Surowiecki (2004) distinguishes between three types of group decision-making. The first are cognitive problems, such as deciding who will win a major sporting event. The second are coordination decisions, and include buyers and sellers working out how to trade at a fair price. The third are cooperation decisions, and include selfless feats like paying taxes. Traditional demonstrations of fallacies in heuristic and biases research have largely focused on cognitive problems. In contrast, prediction markets are more like coordination problems.

Thus, it is possible one of the key differences is not that the probability estimates we examined came from groups, but that they were made through a trading process. On the other hand, experimental psychology has used betting paradigms to replicate logical fallacies in individual decision-making (e.g., Bar-Hillel & Neter, 1993). This raises the question as whether coordinated group trading is somehow different from or better than individual betting. It would be interesting to know, for example, whether short selling (i.e., betting on a decrease in a stock, and so realizing a loss), which is notoriously psychologically difficult in financial markets (e.g., Odean, 1998), is psychologically easier in prediction markets.

A more general question is whether trading markets are the best way to solve group coordination problems. A number of different mechanisms exist, including stock pricing and futures contracts in economics, voting systems in politics, and points spreads in sports betting. It would be interesting to compare these approaches, and to relate them back to the constituent cognitive decision and estimates on which they must ultimately rest. This could tell us something about how individual and group decision-making are related, and how to extract good decisions from groups in different circumstances.

A key claim made Surowiecki (2004) is that good group decision-making requires four elements. The first two are a set of rules, and a mechanism for aggregating individual opinions. These are clearly provided by prediction markets like Intrade. The third element is that there is diversity, independence or decentralization in the individual decision-makers. While we do not have details on the traders in our markets, it seems likely they

had a variety of political dispositions, levels of expertise in finance and world affairs, and brought different information and opinions to bear on their individual decisions.

The final element advocated by Surowiecki (2004) is that individuals need to be able to communicate to some extent, but not too much. This is an interesting issue. Prediction markets in general, and possibly Intrade in particular, seem to strike this difficult balance in a natural way. The decisions made by individual traders to offer a buying or selling price are made in the context of current lists of offers. This gives individuals a calibrating perspective on the opinions of others. But to make successful trades, individuals must identify deviations from reality in existing offers. Contacts should be bought when undervalued, and sold when overvalued. Thus, individual traders must use their own expertise, opinion and data to provide fresh information to make trades, and so are encouraged to bring new information to the market. Perhaps the dynamics of trading allow group opinions to calibrate (rather than anchor) individual opinions, while encouraging improvement (rather than conformity) in decision-making.

Acknowledgments

This research was supported by a UROP grant from the University of California to EG and MDL. We thank members of the Memory and Decision Lab, and Danny Openheimer and Ben Newell for helpful discussions.

References

Bar-Hillel, M., & Neter, E. (1993). How alike it is versus how likely it is: A disjunction fallacy in probability judgments. *Journal of Personality and Social Psychology*, 65, 1119–1131.

Christiansen, J. D. (2007). Prediction markets: Practical experiments in small markets and behaviours observed. *The Journal of Prediction Markets*, 1, 17–41.

Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance*, LIII(5), 1775–1798.

Surowiecki, J. (2004). *The wisdom of crowds*. New York: Random House.

Tindale, S. (1993). *Decision errors made by individuals and groups*. Hillsdale, NJ: Erlbaum.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.

Tversky, A., & Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315.

Wright, G., & Whalley, P. (1983). The supra-additivity of subjective probability. In B. P. Stigum & F. Wenstop (Eds.), *Foundations of utility and risk theory with applications* (pp. 233–244). Dordrecht: Reidel.