On the Observational Implications of Knightian Uncertainty

Hassett, Kevin and Zhong, Weifeng

Council of Economic Advisers, American Enterprise Institute

3 October 2017

On the Observational Implications of Knightian Uncertainty*

Kevin A. Hassett† and Weifeng Zhong‡

October 3, 2017

Abstract

We develop a model of a prediction market with ambiguity and derive testable implications of the presence of Knightian uncertainty. Our model can explain two commonly observed empirical regularities in betting markets: the tendency for longshots to win less often than odds would indicate and the tendency for favorites to win more often. Using historical data from Intrade, we further present empirical evidence that is consistent with the predicted presence of Knightian uncertainty. Our evidence also suggests that, even with information acquisition, the Knightian uncertainty of the world may be not "learnable" to the traders in prediction markets.

Keywords: ambiguity, Knightian uncertainty, prediction market, maxmin preferences

*We appreciate helpful comments from Julian Chan, Jon Hartley and Joe Sullivan, and we thank Cody Kallen for excellent research assistance. All errors are our own.

†Chairman, Council of Economic Advisers. The research and writing of this paper was completed while Mr. Hassett was at the American Enterprise Institute. This paper does not represent the findings or views of the Council of Economic Advisers or the United States. E-mail: Kevin.A.Hassett@cea.eop.gov.

‡(corresponding author) American Enterprise Institute, 1789 Massachusetts Avenue, NW, Washington, DC 20036. E-mail: Weifeng.Zhong@AEI.org.

1 Introduction

At least since the work of Knight (1921), economists have understood that economic agents may behave differently in risky circumstances, where outcomes are random but governed by known probabilities, as opposed to uncertain circumstances, where risks are unknown. Ellsberg (1961) provides examples that highlight the tendency for some decision makers to be averse to the presence of Knightian uncertainty—or, ambiguity.

In recent years, there has been an explosion of theoretical work developing models that incorporate ambiguity aversion, building off of the seminal contribution of Gilboa and Schmeidler (1989). In the literature to date, Knightian uncertainty has been a factor inserted in a model that could possibly explain puzzling observations. It has served a role analogous to that of dark matter in cosmological models, lurking behind the scenes to explain observed phenomena, never being directly observed. At the same time, a rich literature has evolved exploring the efficiency of betting and prediction markets that price specific events. Following on the early work of Kahneman and Tversky (1979) and Asch, Malkiel and Quandt (1982), the ability of these markets to predict future events has been studied extensively, and a number of empirical anomalies have been identified.

In this paper, we extend the theoretical literature and connect it to the prediction-market application. In so doing, we develop more direct observable implications of the presence of Knightian uncertainty than has been achieved previously in the literature, and a method to test for its presence.

While we below will formally derive a model that suggests our test, the intuition of our approach is quite straightforward and can be illustrated using an example from Ellsberg (1961). Suppose that we have two urns. In one urn, we have 50 black balls and 50 red balls. In another urn—the "Knightian urn"—we have 100 balls, but we have no information regarding the proportions. A subject is offered a game. If she pulls a black ball out of the urn, she wins \$1. If she pulls a red ball out she wins nothing. The literature has documented a tendency for individuals to prefer the urn with the known probabilities, suggesting that they exhibit the aversion to ambiguity discussed above.

Suppose that an econometrician could observe games played with both of the urns in Ellsberg's game. With a number of repeated trials, the sample proportions from the first urn would fairly rapidly indicate an estimate that the binomial probability of victory is 50 percent. With enough data, one would say that with great confidence. On the other hand, if one observed repeated play with the second, Knightian urn which, after all, has some number of black balls in it, then the sample proportion would also converge to an estimated binomial probability, but that probability would not necessarily be 50 percent.

The observation that motivates this paper stems from this thought experiment. Given a market derived *ex ante* probability of a binary event, as one frequently observes in betting markets, there will naturally be circumstances where information is extremely solid, and odds are quite far from 50 percent. There will also be situations where information suggests there is an even match (as with a coin flip), and the contract suggests there is close to a 50 percent chance of either outcome. This often happens, for example, in presidential futures markets in the U.S. after the conventions are over. But it is also possible that there are contracts that suggest that the odds of either outcome are 50 percent because the event is shrouded in ambiguity. If we were to estimate the *ex post* sample proportions from just these contracts with *ex ante* 50 percent probabilities, then they could, as in the Ellsberg example above, be anything. If we were to estimate the *ex post* sample proportions of the high information contracts with probabilities far from 50 percent, the proportions and *ex ante* probabilities should, if markets are efficient, align. But close to 50 percent, they might not, and if they do not, it is indication of the presence of Knightian uncertainty. Thus, the pattern by which the relationship between *ex post* proportions and *ex ante* probabilities deviates from the 45 degree line becomes informative regarding the presence of Knightian uncertainty. We also discuss the extent to which learning can occur in markets over time. If Knightian uncertainty induces knowledge acquisition, then the relationship between proportions and probabilities will evolve as a market matures, a possibility we explore in the paper.

The next subsection briefly reviews the literature. In Sections 2 and 3, we draw on the work of Gilboa and Schmeidler (1989) and Dow and Werlang (1992) and develop a model that suggests that the pattern described by our intuitive example would emerge in a market influenced by the present of significant Knightian uncertainty. In Section 4, we provide some high-level evidence that the relationship between *ex post* proportions and *ex ante* probabilities is consistent with the predictions of our model. Section 5 concludes.

1.1 Literature Review

This paper draws from two different strands in the literature. First, theorists have made remarkable strides in recent years incorporating Knightian uncertainty and ambiguity aversion into models of financial markets.

These models have, according to an exhaustive recent review, "implications for portfolio choice and asset pricing that are very different from those of SEU (subjective expected utility theory) and that help to explain otherwise puzzling features of the data."[1] Ambiguity aversion could help explain the tendency of markets to stop operating during financial crises, for prices to not be completely informative, and even for there to be bank runs.[2]

This branch of the literature has focused on financial markets in general. At the same time, an equally impressive literature has emerged exploring the functioning of prediction markets, which, for the most part, price in the probability of specific binary events. As Thaler and Ziemba (1988) first noted, these prediction markets may be a better laboratory to test cutting edge theories, as they contain contracts with known durations, and observable discrete events that stop the trading. While an equity might live on virtually forever, a presidential election future has a specific end date, and its ability to forecast the outcome can be precisely evaluated.

This second literature has advanced both empirically and theoretically. On the theoretical side, Manski (2004) first illustrated that the beliefs of bettors may not necessarily yield a market-based probability. More recently, Wolfers and Zitzewitz (2006) identify the conditions under which prediction-market prices coincide with bettors' mean beliefs about probabilities. On the empirical side, prediction markets have been found to be informative regarding the odds of events occurring. Berg et al. (2008), for example, find that the Iowa Electronic Markets outperformed polls in predicting election outcomes. At the same time, markets have been found to exhibit a favorite-longshot bias, with favorites outperforming their odds, and longshots under-

[1]See Epstein and Schneider (2010), p. 315.

[2]See Caballero and Krishnamurthy (2008), Caballero and Simsek (2013), Guidolin and Rinaldi (2010), Routledge and Zin (2009), and Uhlig (2013).

performing (see, e.g., Cain, Law and Peel, 2000). A number of possible explanations for this pattern include insider trading (Shin, 1992), risk loving behavior (Weitzman, 1965), and imperfect ability to process information (Snowberg and Wolfers, 2010).

The connection of these two literatures seems quite promising, as betting markets often exist for events, such as Brexit or elections, for which Knightian uncertainty may well be present. Since they also have finite and determinate life spans, they also allow the econometrician the ability to evaluate their performance *ex post*. We now turn to illustrating the utility of this approach.

2 A Model of Prediction Market with Ambiguity

2.1 Setup

Events and Contracts. Consider a prediction market for the occurrence of a binary event. There are two all-or-nothing contracts corresponding to the two possible realizations. One contract pays \$1 if event A occurs and \$0 otherwise, while the other contract pays \$1 if the complementary event A^c occurs and \$0 otherwise. Let π denote the price of contract A. No-arbitrage condition dictates that, in equilibrium, the price of contract A^c be $1 - \pi$.

Traders. There is a continuum I of competitive traders, each endowed with homogeneous initial wealth w. The "net" position on contract A held by trader i is denoted by $x_i \in \mathbb{R}$.[3] Given price π, the final wealth w_i of trader i is

$$w_i = \begin{cases} w + (1 - \pi)\, x_i & \text{if event } A \text{ occurs,} \\ w - \pi x_i & \text{if event } A^c \text{ occurs.} \end{cases}$$

All traders have log utility of their final wealth: $u\,(w_i) = \ln w_i$.

[3] In practice, trader i can long and/or short contract A and/or contract A^c, but some strategies are mathematically equivalent. For example, holding $m_i > 0$ units of contract A and $n_i > m_i > 0$ units of contract A^c would be equivalent to holding m_i units of cash, 0 unit of contract A , and $n_i - m_i > 0$ units of contract A^c. Therefore, without loss of generality, we let a single decision variable $x_i = m_i - n_i$ (which could be positive, zero, or negative) represent the "net" position held by trader i.

Beliefs and Ambiguity. Suppose trader i has a subjective belief that event A occurs with probability $\widetilde{q} \in [0,1]$. Then, the subjective expected utility of trader i from holding position x_i at price π is given by

$$U(\pi, x_i; \widetilde{q}) = \widetilde{q} \ln(w + (1-\pi)x_i) + (1-\widetilde{q}) \ln(w - \pi x_i).$$

However, ambiguity exists, for traders may be uncertain about how likely event A is to occur. We follow Gilboa and Schmeidler (1989) and model ambiguity using the "multiple-prior" framework. Specifically, suppose each trader i considers every probability $\widetilde{q} \in [q_i - \epsilon, q_i + \epsilon]$, where $\epsilon \geq 0$, an admissible probability that governs the realization of the binary event. Under this framework, q_i represents the "mean" belief of trader i, while ϵ is interpreted as a measure of ambiguity. Given price π, trader i chooses position x_i to maximize the minimum—that is, the worst-case scenario—of all her admissible, subjective expected utilities:

$$\max_{x_i \in \mathbb{R}} \left[\min_{\widetilde{q} \in [q_i - \epsilon, q_i + \epsilon]} U(\pi, x_i; \widetilde{q}) \right]. \tag{1}$$

Traders are heterogeneous in mean belief. Let the distribution of traders' mean beliefs be characterized by a cumulative distribution function F over interval $[\epsilon, 1-\epsilon]$. That is, for the most pessimistic trader, the worst-case belief that A occurs is probability 0 while, for the most optimistic trader, the best-case belief that A occurs is probability 1.

2.2 Optimal Demand and Portfolio Inertia

Solving the inner minimization reduces the optimization problem (1) to

$$\max_{x_i \in \mathbb{R}} U(\pi, x_i; q_i - sgn(x_i)\epsilon),$$

where $sgn(\cdot)$ is an indicator function that takes the sign of its argument.

The intuition behind the above expression is straightforward. If trader i has a positive position on contract A, then the worst-case scenario would be that event A occurs with probability $q_i - \epsilon$, the lower bound. Similarly, if the position of trader i is negative, then, in the worst-case scenario, event A occurs with the upper-bound probability, $q_i + \epsilon$.

Solving the maximization problem gives the optimal (net) demand for contract A by trader i,

$$x(\pi; q_i) = \begin{cases} \frac{q_i - \epsilon - \pi}{\pi(1-\pi)} w & \text{if } \pi \in [0, q_i - \epsilon), \\ 0 & \text{if } \pi \in [q_i - \epsilon, q_i + \epsilon], \\ \frac{q_i + \epsilon - \pi}{\pi(1-\pi)} w & \text{if } \pi \in (q_i + \epsilon, 1], \end{cases} \tag{2}$$

as a function of price and mean belief. Therefore, trader i longs contract A when the price is lower than her most pessimistic belief, and shorts contract A when the price is higher than her most optimistic belief. For any price in the intermediate range $[q_i - \epsilon, q_i + \epsilon]$, trader i does not participate in the prediction market—the phenomenon of portfolio inertia.

That portfolio inertia arises when investors have maxmin preferences is well known in the finance literature since the work by Dow and Werlang (1992). The setup of this model replicates this phenomenon in the context of prediction markets. In particular, for each trader, the size of price region at which portfolio inertia occurs is given by 2ϵ. In other words, the higher the degree of ambiguity, the more inertial the traders' portfolios.

2.3 Equilibrium

Given price π for contract A and distribution function F of traders' mean beliefs, the aggregate (net) demand for the contract is given by

$$X_F(\pi) = \int_{\epsilon}^{1-\epsilon} x(\pi; q)\, dF(q). \tag{3}$$

The prediction market is in equilibrium when the aggregate demand for contract A equals zero, that is, $X_F(\pi) = 0$. The following proposition establishes the equilibrium price.[4]

Proposition 1 *Given distribution function F, the equilibrium price π_F^* is such that*

$$\pi_F^* = \mathbb{E}_F(q) + \int_{\pi_F^* - \epsilon}^{\pi_F^* + \epsilon} F(q)\, dq - \epsilon.$$

[4] We relegate all proofs to Appendix A.

When ambiguity is absent (i.e., $\epsilon = 0$) the prediction market aggregates the "wisdom of crowds:"

$$\pi_F^* \mid_{\epsilon=0} = \mathbb{E}_F(q).$$

That is, the equilibrium price of contract A corresponds to the *average* of traders' mean beliefs about the occurrence of event A.

In the presence of ambiguity, however, the prediction market does not necessarily aggregate the wisdom of crowds. In particular, it aggregates the wisdom of crowds if and only if the distribution function F is such that $\int_{\pi_F^*-\epsilon}^{\pi_F^*+\epsilon} F(q)\, dq = \epsilon$. The next proposition shows that the situations in which such equality happens to hold are topologically rare.

Proposition 2 *The prediction market "rarely" aggregates the wisdom of crowds. Formally, let Δ be the space of probability distributions over $[\epsilon, 1-\epsilon]$, endowed with the weak topology. Then, the subset of probability distributions such that the equilibrium price equals the average of traders' mean beliefs is nowhere dense in Δ.*

Propositions 1 and 2 together suggests that the presence of ambiguity renders the prediction market ineffective in aggregating the beliefs held by heterogeneous traders.

Proposition 3 *The equilibrium quantity of trades is strictly decreasing in the degree of ambiguity.*

Proposition 3 is a direct consequence of portfolio inertia. As the degree of ambiguity increases, the "inaction range" of each trader i, $[q_i - \epsilon, q_i + \epsilon]$, becomes wider. Since each trader is more likely to stay put in a more ambiguous environment, the aggregate trades must be fewer as well. This result is reminiscent of well-known models of ambiguity in financial economics (e.g., Caballero and Krishnamurthy, 2008; Guidolin and Rinaldi, 2010; and Routledge and Zin,2009), which suggest that a significant increase in Knightian uncertainty may contribute to liquidity hoarding and market breakdown.

Moreover, the deterrence of trades in a particular way is what causes the failure of the prediction market to aggregate beliefs. Specifically, for any prevailing price π, the traders who stay put are those with moderate beliefs such that their inaction ranges cover π. Those who trade have beliefs that are more extreme—either more optimistic or more pessimistic—than the abstainers. Suppose the abstaining traders did trade, the chance of the hypothetical market price, after aggregating the abstaining traders' beliefs, happens to be exactly the same as π is zero.

3 Testable Implications

The previous section has derived the equilibrium results under ambiguity. However, since the degree of ambiguity is not observable, those results cannot be tested directly. In this section, we impose more information structures on the model and derive implications that are testable with prediction-market data.

3.1 Information Structures

Suppose the true probability that event A occurs is given by $p \in [\epsilon, 1-\epsilon]$. No trader knows about p for certain. However, there is a mass $m \in (0,1)$ of traders whose mean beliefs equal p, while all the other uninformed traders' mean beliefs are continuously distributed over $[\epsilon, 1-\epsilon]$. Assumption 1 embeds these additional structures into the distribution function F.

Assumption 1 *The distribution function F takes the following form:*

$$F(q) \equiv \begin{cases} (1-m)\overline{F}(q) & \text{if } q \in [\epsilon, p), \\ (1-m)\overline{F}(q) + m & \text{if } q \in [p, 1-\epsilon], \end{cases}$$

where $\overline{F}$ is some continuous *distribution function of q over* $[\epsilon, 1-\epsilon]$.

The functional form of $\overline{F}$ is left unspecified. We let $\overline{\Phi}$ denote the integral of $\overline{F}$, i.e., $\overline{\Phi}(q) \equiv \int_\epsilon^q \overline{F}(q')\,dq'$.

The interpretation of Assumption 1 is that there are some traders who are informed while the others are not. One could provide a micro-foundation for this setup by assuming that the mass m of traders have the correct mean beliefs because they have received private signals informative of the true probability, while all other traders have received no such signals. With such a micro-foundation, the traders' (multiple) beliefs should be interpreted as their (multiple) posteriors. In this section, we adopt a reduced-form approach and build these details directly into the distribution function F.

As will be clear in the next subsection, the assumption of a mass m is not essential to deriving testable implications. Whether m is extremely small, in which case few traders are informed, or very large, which means most traders are informed, the main results, shown in Proposition 4, are unchanged. The essence of this setup is that a fraction of traders hold beliefs that are informative of the true state of the world. Therefore, when these informed traders do not participate in the market, their beliefs will not be reflected in the equilibrium price.

3.2 Implications

Since the distribution function F is given and parameterized by the true probability p, applying Proposition 1 allows us to solve for the equilibrium price as a function of p, as shown below.

Proposition 4 *Under Assumption 1, the equilibrium price $\pi^*(p)$ is:*

1. *continuous, with $\pi^*(\epsilon) > \epsilon$ and $\pi^*(1-\epsilon) < 1-\epsilon$;*
2. *such that $\pi^*(p) = \widehat{\pi}$ for any $p \in [\widehat{\pi}-\epsilon, \widehat{\pi}+\epsilon]$;*
3. *strictly increasing for $p \notin [\widehat{\pi}-\epsilon, \widehat{\pi}+\epsilon]$;*

where $\widehat{\pi}$ is identified by $\widehat{\pi} - \overline{\Phi}(\widehat{\pi}+\epsilon) + \overline{\Phi}(\widehat{\pi}-\epsilon) = 1 - 2\epsilon - \overline{\Phi}(1-\epsilon)$.

Figure 1 plots the equilibrium price in a p–π diagram, where the true probability $p = (\pi^*)^{-1}(\pi)$ is a *correspondence* of the equilibrium price π. Specifically, it attains a non-singleton set value when $\pi = \widehat{\pi}$, with the size of that set equal to 2ϵ.

The most important feature of the equilibrium (part 2) is that there exists a range of true probabilities, $[\widehat{\pi}-\epsilon, \widehat{\pi}+\epsilon]$, within which the market price is *not at all* responsive to any change in the underlying state of the world. That is, $\pi^{*\prime}(p) = 0$ for any p in that range. Instead of prediction, the prediction market simply assigns an uninformative number $\widehat{\pi}$—the mid-point of the range $[\widehat{\pi}-\epsilon, \widehat{\pi}+\epsilon]$—as the price. The reason for this result is straightforward: Since the traders who hold private information about p are not trading, what exactly those traders know about the true

Figure 1: Prediction Market Equilibrium in the $p-\pi$ Diagram.

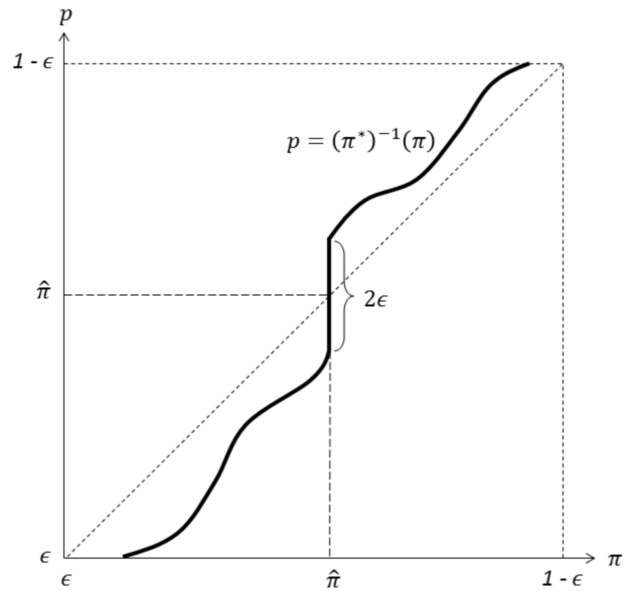

state of the world must not be reflected in the market price. Morever, note that the size of the segment $[\widehat{\pi} - \epsilon, \widehat{\pi} + \epsilon]$, measuring the degree of ambiguity, does *not* depend on m, the mass of traders who are informed of p. Even if m is infinitesimal, the market price stops reflecting the true probability once these traders abstain from betting.

Outside the range $[\widehat{\pi} - \epsilon, \widehat{\pi} + \epsilon]$, however, the prediction market works (part 3). Specifically, if all parameters of the model were known, one would be able to infer the true probability p from the equilibrium market price $\pi^*(p)$. The higher the true probability, the higher the price.

Part 1 of the proposition also shows that, for a true probability that is very high (near $1 - \epsilon$) or very low (near ϵ), the equilibrium price exhibits a favorite-longshot bias commonly observed in the literature (e.g., Cain, Law, and Peel, 2000): favorite events are under-priced while longshot events are over-priced. The intuition is as follows. For a longshot event where $p = \epsilon$, for example, if the market price was as low as ϵ, that would imply all traders' mean beliefs were greater than the prevailing price and, hence, all traders would long the contract, which cannot be an equilibrium. Therefore, the equilibrium price of a longshot must be significantly larger than the longshot's odds.

Note that, although the degree of ambiguity, ϵ, is not directly observable in reality, Proposition 4 yields implications of the presence of ambiguity that are testable with prediction-market data. Suppose an econometrician could conduct a large number of repeated trials for each value of the true probability. Then, with enough data, the *ex post* sample proportion, denoted by P, would converge to the corresponding true probability, p. It follows that the estimated relationship between P and the market price, π, would converge to the graph of the correspondence $p = (\pi^*)^{-1}(\pi)$. As in Figure 1, such ideal trials would show a big jump at price level $\pi = \widehat{\pi}$, with P being generally below the 45-degree line below but close to $\widehat{\pi}$, and above it just thereafter. Moreover, since the relationship between P and π fundamentally shifts between the two continuous segments, our result suggests a testable structural change near the jump in at $\widehat{\pi}$.

Before conducting the test, the econometrician may not know where the jump would appear, because $\widehat{\pi}$, given by

$$\widehat{\pi} - \overline{\Phi}(\widehat{\pi} + \epsilon) + \overline{\Phi}(\widehat{\pi} - \epsilon) = 1 - 2\epsilon - \overline{\Phi}(1 - \epsilon),$$

depends on the distribution of mean beliefs among all prediction-market traders. It follows from the above equation that $\widehat{\pi}$ would be smaller than 0.5 if $\overline{F}$ is skewed towards the lowest mean belief ϵ, and larger than 0.5 if $\overline{F}$ is skewed towards the highest mean belief $1 - \epsilon$. But when $\overline{F}$ is symmetrically distributed over $[\epsilon, 1 - \epsilon]$, $\widehat{\pi}$ would be equal to 0.5, which is the following corollary.

Corollary 5 *Under Assumption 1,*

$$\widehat{\pi} = 0.5$$

if $\overline{F}$ is a symmetric distribution function over $[\epsilon, 1 - \epsilon]$ (i.e., $\overline{F}(1 - x) = 1 - \overline{F}(x)$ for any $x \in [\epsilon, 1 - \epsilon]$).

In practice, the empirical chart would precisely follow Figure 1 with the jump at 0.5 in the case of symmetry, but not if asymmetries were present. But even if one might expect skewness to be present for some contracts but not others, the range for the crossover point could be scattered about the neighborhood of 0.5. The aggregation of a large number of contracts, therefore, could push the average $\widehat{\pi}$ to be in the neighborhood of 0.5. Since each contract would exhibit a similar (if slightly shifted) pattern, the overall pattern should loosely follow Figure 1 if Knightian uncertainty is important in these markets, even though some reflect symmetry whereas others do not.

Accordingly, the theory suggests that the empirical relationship between P and π would contain a testable structural break about the neighborhood of 0.5, where one would expect to see observations scattered below the 45-degree line to the left of the break-point, and above the 45-degree line to the right. The presence of a structural break adjacent to 0.5, therefore, would be an indication that Knightian uncertainty is a factor in the market, and would be consistent with the intuition provided in the introduction.

4 Empirical Evidence

In this section, we provide some high-level evidence that is consistent with the theoretical predictions.[5]

We use the historical data from Intrade, a popular online prediction platform which operated from 2003 to 2013. The platform hosted prediction contracts across wide-ranging categories of events, such as business (e.g., whether the CEO of a certain company would step down), current events (e.g., which city would host the Olympic), entertainment (e.g., which movie would win the Academy Award for the Best Picture), politics (e.g., which candidate would be elected the U.S. president), etc. We collect all those contracts that are on binary events, regardless of their categories, and record how each binary event had turned out.

The aim of the empirical analysis is to estimate the *ex post* sample proportion, P, of event A's occurrence as a function of the *ex ante* price, π, of contract A. We process the data in the following way. The observations are sorted by price and evenly partitioned into a number of percentile bins. For each percentile bin, we calculate the sample proportion of event A's occurrences whose corresponding prices fall into that bin. Finally, we plot the sample proportions against the mid-points of the corresponding price bins.

[5] See Appendix B for the details of the empirics.

If the theory developed in the previous section holds, the following is what one would expect in the empirics. Recall that the value of $\widehat{\pi}$ depends on the distribution of mean beliefs among traders. Since each observation in the dataset is from a certain market with a certain distribution of mean beliefs held by the participating traders, we can interpret each observation as a single draw from the data-generating process associated with a certain version of Figure 1. For a price bin closer to 0, therefore, it is more likely that the observations contained in the bin have been drawn from the left part of Figure 1, i.e., below the break-point. Similarly, for a price bin closer to 1, the observations are more likely to have been drawn from the right part of Figure 1, i.e., above the break-point. More important, when the price bin is near 0.5, the observations are more likely to be from just around the jump, suggesting a structural break.

We start with the empirical evidence from political events, one of the largest categories in the Intrade dataset. These events, like Brexit and U.S. presidential elections, often see a high volume of transactions between bettors. Figure 2, based on a partition into 50 bins (i.e., 2% of observations per bin), plots the sample proportion for all bins against the corresponding price. Since prices evolve over time in the prediction markets until the random events are realized, the two panels of the figure together capture the effect of timing by showing the estimation for two different dates: (a) the first day market opens to bettors, and (b) the last trading day before the event is realized, respectively.

For each panel, we conduct three analyses. First, a linear regression assuming *no* structural breaks is shown as the dashed line in the diagram. Next, we run two types of break-point tests—an F test[6] and a "moving sum of residuals" (MOSUM) test[7]—against the null hypothesis that there is no structural breaks for the entire sample. Lastly, we re-run the linear regression by estimating the location of one break-point (as suggested by our theory). The estimation returns (i) a linear segment on each side of the estimated structural break, plotted as the solid lines in the diagram, as well as (ii) the location of the structural break, identified by two red dots in the diagram corresponding to, respectively, the last observation of the first segment and the first

[6] The F test is an extension of the "Chow test" (1960), against the alternative hypothesis of an *unknown* break-point. See, e.g., Andrews (1993) and Andrews and Ploberger (1994) for details.

[7] The MOSUM test analyzes the moving sum of residuals and detects whether a strong shift of the fluctuation process exists. See, e.g., Chu, Hornik, and Kuan (1995a, b) for details.

Figure 2: Prediction Market Data in the P-π Diagram: Politics (50 bins).

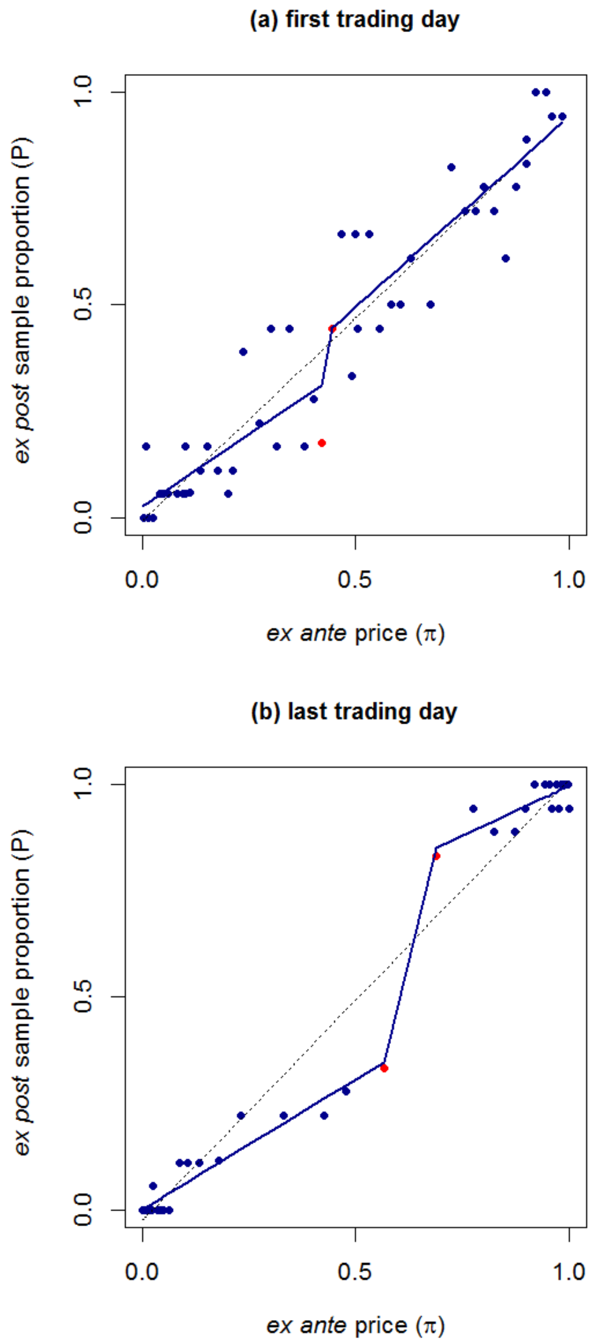

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

observation of the second segment. The details of the three analyses are shown in the column "50 bins" of Table 1.

A few remarks on the results follow. First, in both panels, the regression lines without structural breaks fall very close to the 45-degree line, suggesting the overall efficiency of markets in pricing the probabilities of random events. The evidence of market efficiency *on the first trading day* is remarkable because, for politics, a lot of markets opened a long time—sometimes years—ahead of the resolution of the events. Yet, as the regression table shows, the slopes are statistically significant and very close to 1.

Second, although panel (a) is relatively noisier, panel (b) shows a clear pattern as predicted by our theory: The null hypothesis of no structural breaks is cleared rejected by the tests, and the break-point estimation shows a significant jump near price level 0.5. The diagram, hence, resembles our prediction shown in Figure 1. For observations in the intermediate price range, one might think that the price is close to 0.5 because traders have solid information suggesting an "even match" between outcome A and A^c. It is also possible, however, that the market is shrouded in ambiguity as some traders, albeit partially informed, are reluctant to trade. Just like in the example of a Knightian urn, an intermediate price in this case could mean a wide range of true probabilities. In panel (b), for a price in the break region between 0.57 and 0.69, the sample proportion could be as low as 33%, or as high as 83%. In other words, the degree of ambiguity, ϵ, in this particular example is about 0.25 (i.e., half of 83% − 33%). Such a magnitude is significant not only statistically—as the rejection of null hypothesis "no structural breaks" implies the rejection of "ϵ equals to zero"—but also economically. According to the multiple-prior framework, it would mean that a typical trader would consider all the probabilities within an interval of length 0.5 equally admissible in governing the realization of the binary event. A significant jump near price level 0.5 like the one in panel (b), therefore, is an indication of the presence of Knightian uncertainty. As the linear regression without breaks shows, the specific pattern of observations also causes the regression line to have a slope larger—albeit only slightly—than 1.

Another important difference between the two panels is that, in panel (b), more observations are clustered near price levels 0 and 1. This means, by the last day, more traders hold (posterior) beliefs that some outcome—either A or A^c—is very

Table 1: Estimation and Test of Structural Breaks: Politics

Dependent variable: *ex post* sample proportion
Independent variable: *ex ante* price

	50 bins	30 bins
(a) first trading day		
slope (no breaks)	0.953***	0.950***
	(22.7)	(21.5)
tests of "no breaks"		
F test	not rejected	not rejected
MOSUM test	not rejected	rejected*
structural break estimation		
slope (segment 1)	0.727***	0.690***
	(5.75)	(5.13)
break region	[0.45, 0.47]	[0.43, 0.47]
slope (segment 2)	0.863***	0.867***
	(7.70)	(7.56)
(b) last trading day		
slope (no breaks)	1.03***	1.04***
	(44.8)	(44.8)
tests of "no breaks"		
F test	rejected***	rejected***
MOSUM test	rejected*	rejected*
structural break estimation		
slope (segment 1)	0.617***	0.767***
	(15.6)	(21.4)
break region	[0.57, 0.69]	[0.60, 0.75]
slope (segment 2)	0.464***	0.229*
	(5.10)	(2.13)

t statistics in parentheses
$^{\wedge}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

likely to be realized, suggesting a decrease in *risks* over time. Such a decrease in risks can be a result of information acquisition by the traders, who, until the random events resolve, may have the incentives to learn about the events and update their bets accordingly. Since risks have decreased while ambiguity remains, our empirical evidence also suggests an *observational* distinction between the concepts of risk and Knightian uncertainty.

Furthermore, the above empirical patterns are to some extent robust against the choice of the number of bins. Figure 3 reproduces the diagram by partitioning the data into 30 bins instead. The column "30 bins" of Table 1 shows the details of the regression and break-point analysis. Overall, the observations drawn earlier still hold.[8]

The observation that ambiguity remains until the last trading day suggests, unlike risks, Knightian uncertainty may be *not* "learnable" in practice to the traders. The intuition can be illustrated using the Knightian urn where the composition of black and red balls is unknown. Imagine two different scenarios. In the first scenario, a subject observed repeated draws from the same Knightian urn. In this case, the sample proportion over time would reveal the true composition of the two colors because, after all, the composition is fixed over time. In the second scenario, there was an experimenter who replaced the Knightian urn with a new one every time a ball was drawn by the subject. In this case, the sample proportion may not inform the subject of what to expect in the next Knightian urn, simply because the composition of black and red balls in the new urn could be anything of the experimenter's choosing. If the underlying data-generating process—that is, the way the experimenter changed every other Knightian urn—was not learnable to the subject, then the degree of ambiguity would not decrease over time.[9] One might think that, in politics, it is intuitively easy for traders to acquire knowledge—from polls, news reports, political analyses, etc. Yet, our empirical evidence, which is based on a large number of prediction markets about various political events, seems to fit the second scenario, suggesting that the Knightian uncertainty of politics may indeed be *not* "learnable" through information acquisition.

[8] We have checked other variations between 30 and 50, which yield similar results (omitted to limit space). Obviously, the number of bins should be neither too small (which would leave too few points in the diagram), nor too large (which would leave too few observations per bin).

[9] See Epstein and Schneider (2007) for a theoretical treatment of learning under ambiguity.

Figure 3: Prediction Market Data in the P-π Diagram: Politics (30 bins).

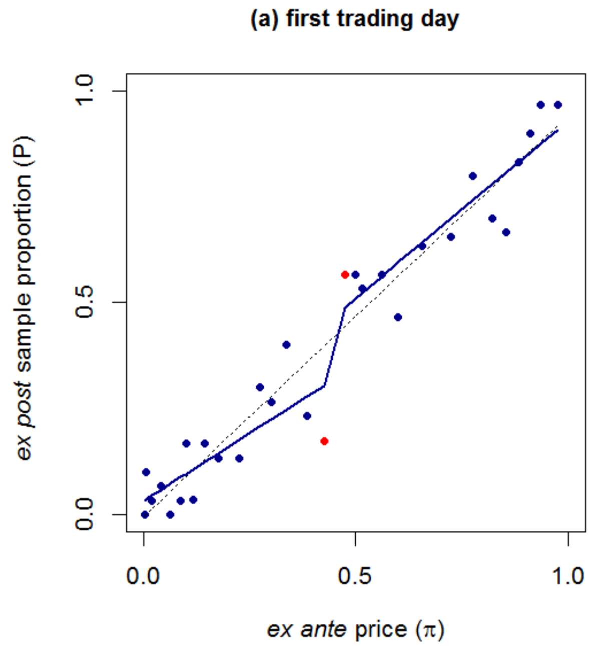

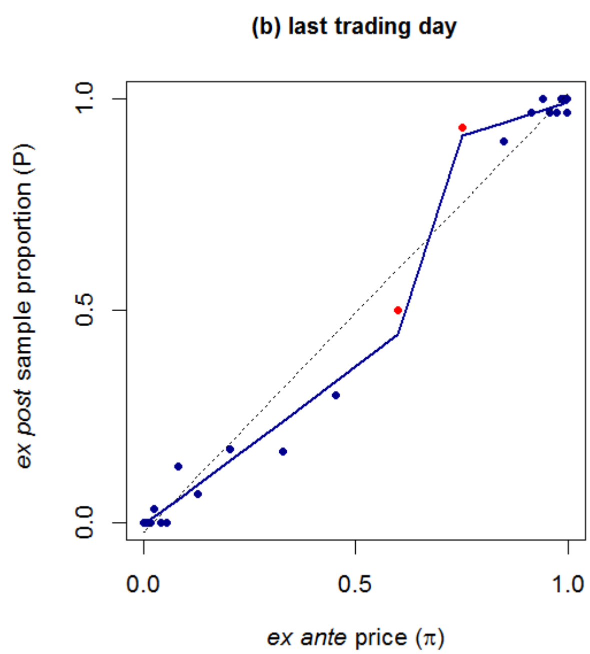

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

We now turn to another major category: entertainment events, such as the winners of cinematic awards or the box offices of movies. Figures 4 and 5 reproduce the P-π diagram for 30 bins and 50 bins, respectively, and Table 2 reports the details of the regressions and break-point tests. Although qualitatively similar, the patterns are less pronounced compared to politics. The jump near 0.5 is less clear and, interestingly, the clustering near 0 and 1 is less marked. This evidence suggests less learning in entertainment than in politics, which is understandable since it is more difficult for bettors to acquire information about the general public's personal tastes of movies and music.

Politics and entertainment together accounts for over 80% of the Intrade dataset. However, for completeness, we reproduce the empirical evidence with the full sample, as shown in Figures 6 and 7, as well as Table 3. The patterns, essentially by construction, are similar to what we establish above.

5 Concluding Remarks

Knightian uncertainty—an important theoretical concept in the literature that is often used to explain observed phenomena—has never been directly evidenced in an empirical setting. In this paper, we have developed a model of a prediction market with ambiguity, where traders have maxmin preferences. We have derived more direct, observational implications of the presence of Knightian uncertainty. Using the historical betting data from Intrade, we have further presented some high-level evidence that is consistent with the prediction of our model. In particular, for price levels close to 0.5, the market-implied, *ex ante* probability of a random event is *not* indicative of the *ex post* sample proportion, suggesting the presence of Knightian uncertainty.

Moreover, our empirical evidence has shown that, although traders seem to have acquired information which leads to a decrease in risks, ambiguity remains until the last trading day, suggesting that the Knightian uncertainty of the world may be not "learnable" to traders. By comparing political events and entertainment events, we have also shown that the empirical patterns we identified are more pronounced in politics than in entertainment.

The evidence we have provided is only preliminary, since the empirics of this

Figure 4: Prediction Market Data in the P-π Diagram: Entertainment (50 bins).

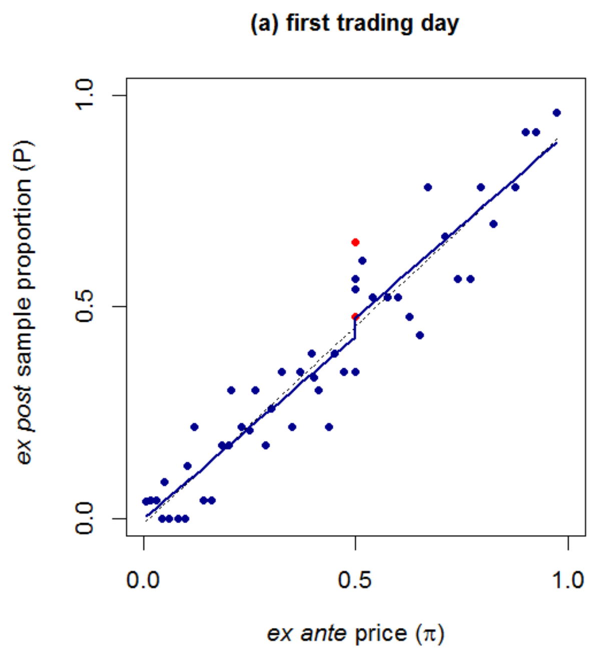

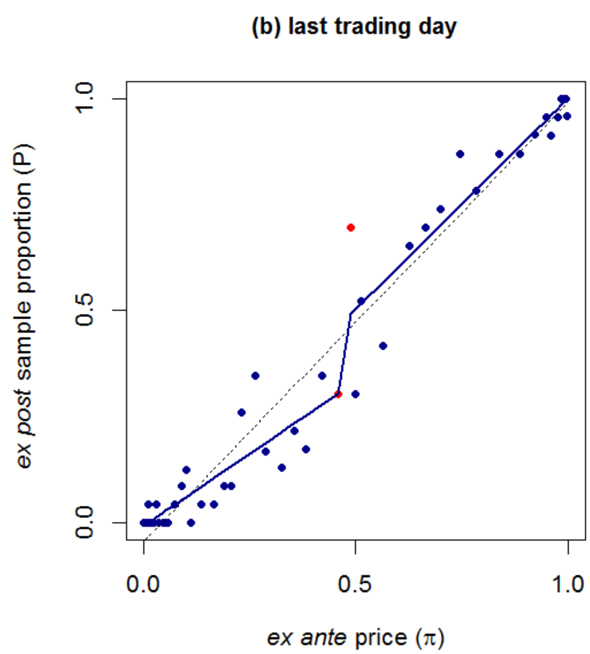

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

Figure 5: Prediction Market Data in the P-π Diagram: Entertainment (30 bins).

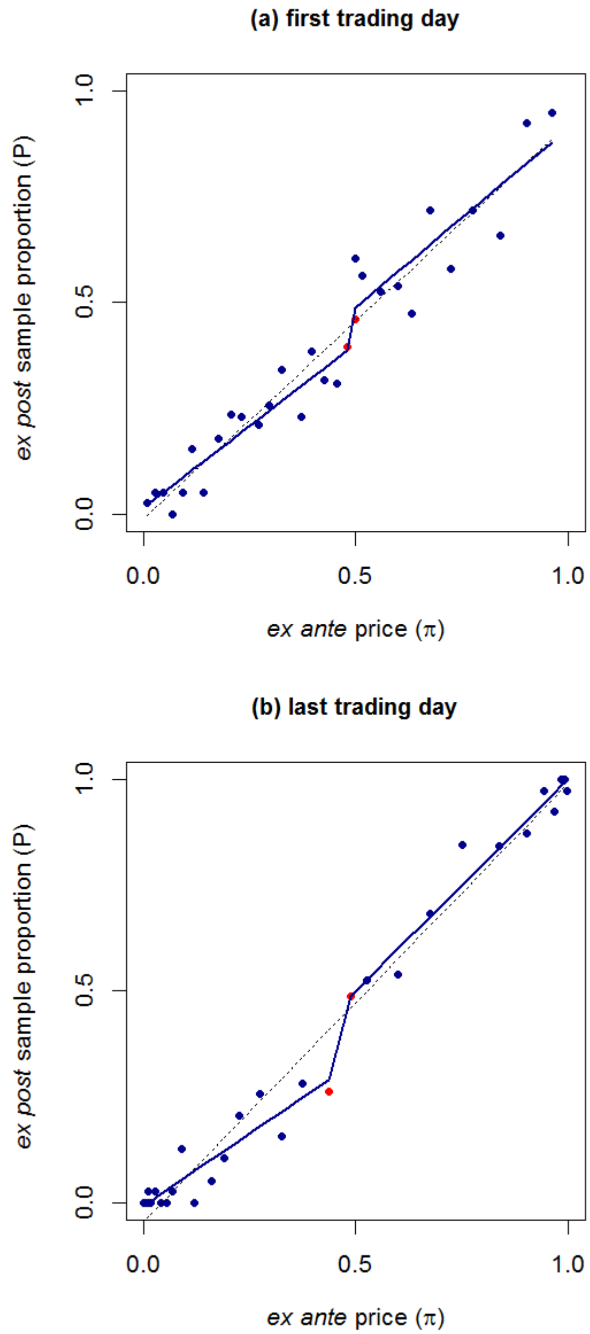

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

Table 2: Estimation and Test of Structural Breaks: Entertainment

Dependent variable: *ex post* sample proportion
Independent variable: *ex ante* price

	50 bins	30 bins
(a) first trading day		
slope (no breaks)	0.933^{***}	0.936^{***}
	(20.9)	(23.3)
tests of "no breaks"		
F test	not rejected	not rejected
MOSUM test	not rejected	not rejected
structural break estimation		
slope (segment 1)	0.859^{***}	0.929^{***}
	(8.66)	(16.2)
break region	[0.50, 0.50]	[0.68, 0.73]
slope (segment 2)	0.881^{***}	1.49^{***}
	(6.80)	(4.89)
(b) last trading day		
slope (no breaks)	1.04^{***}	1.04^{***}
	(33.7)	(37.8)
tests of "no breaks"		
F test	rejected***	rejected***
MOSUM test	rejected$^{\wedge}$	not rejected
structural break estimation		
slope (segment 1)	0.744^{***}	0.697^{***}
	(10.3)	(10.3)
break region	[0.50, 0.51]	[0.44, 0.49]
slope (segment 2)	0.866^{***}	1.00^{*}
	(9.02)	(16.2)

t statistics in parentheses
$^{\wedge}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 6: Prediction Market Data in the P-π Diagram: Full Sample (50 bins).

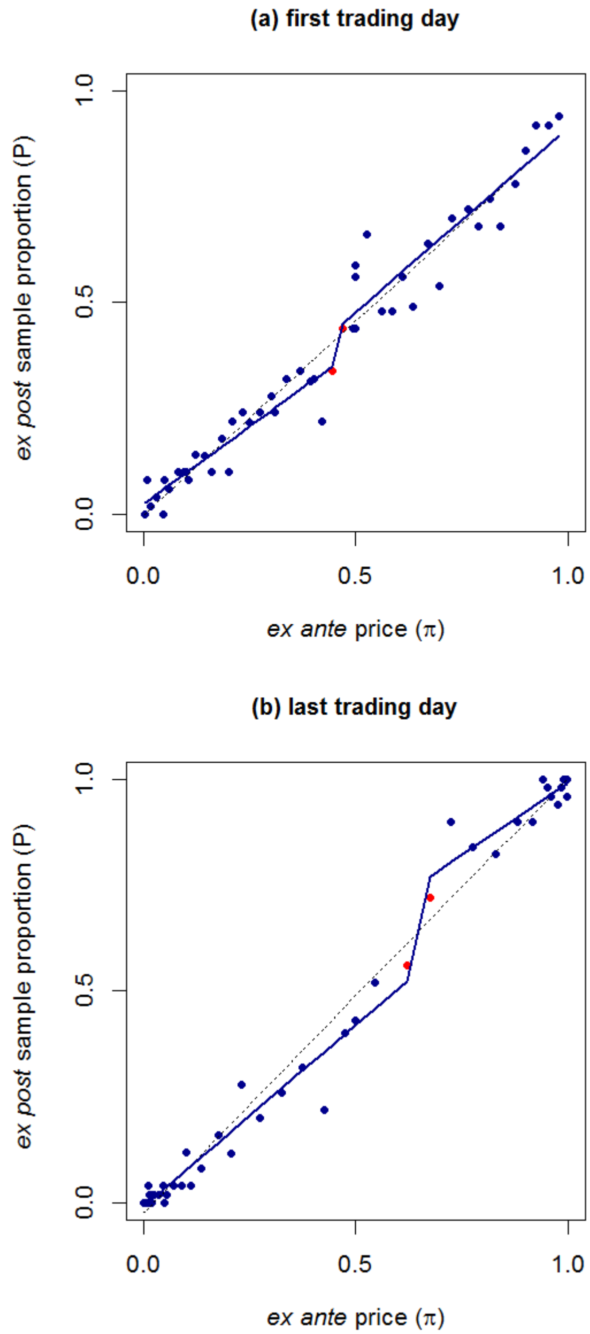

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

Figure 7: Prediction Market Data in the P-π Diagram: Full Sample (30 bins).

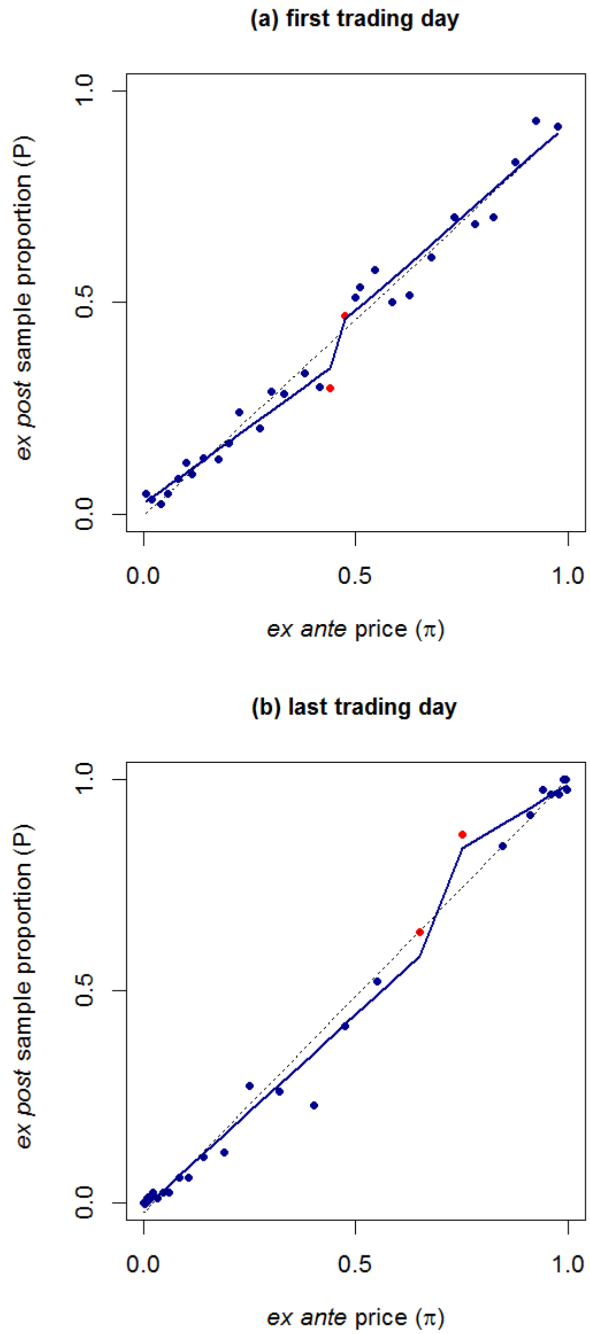

(Note: The dashed lines are regression lines without breaks. The solid lines are regression lines with one estimated break, with two red dots identifying the location of the break.)

Table 3: Estimation and Test of Structural Breaks: Full Sample

Dependent variable: *ex post* sample proportion
Independent variable: *ex ante* price

	50 bins	30 bins
(a) first trading day		
slope (no breaks)	0.917***	0.929***
	(63.6)	(27.7)
tests of "no breaks"		
F test	not rejected	not rejected
MOSUM test	not rejected	not rejected
structural break estimation		
slope (segment 1)	0.936***	0.739***
	(18.0)	(8.50)
break region	[0.59, 0.61]	[0.44, 0.48]
slope (segment 2)	1.26***	0.876***
	(9.56)	(10.5)
(b) last trading day		
slope (no breaks)	1.03***	1.03***
	(56.6)	(46.5)
tests of "no breaks"		
F test	rejected***	rejected***
MOSUM test	rejected*	rejected^
structural break estimation		
slope (segment 1)	0.853***	0.812***
	(25.2)	(20.9)
break region	[0.62, 0.68]	[0.55, 0.65]
slope (segment 2)	0.693***	0.765*
	(7.84)	(9.27)

t statistics in parentheses
^ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

paper are based on a single prediction platform that is skewed towards political and entertainment events. In a future, empirical study, we will collect more prediction-market data across different platforms and different event types, and we will examine more closely the relationship between the *ex post* sample proportion and the *ex ante* price by taking into account the type of events, the time ahead of the resolution of randomness, and other aspects of the betting markets.

References

[1] Andrews, Donald W. K. (1993). "Tests for Parameter Instability and Structural Change With Unknown Change Point," *Econometrica*, 61(4): 821–856.

[2] Andrews, Donald W. K. and Werner Ploberger (1994). "Optimal Tests when a Nuisance Parameter is Present Only Under the Alternative," *Econometrica*, 62(6): 1383–1414.

[3] Asch, Peter, Burton G. Malkiel, and Richard E. Quandt (1982). "Racetrack Betting and Informed Behavior," *Journal of Financial Economics*, 10(2): 187–194.

[4] Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz (2008). "Results from a Dozen Years of Election Futures Markets Research," in *Handbook of Experimental Economics Results*. Charles R. Plott and Vernon L. Smith, eds. Amsterdam; Oxford: North Holland.

[5] Caballero, Ricardo J. and Arvind Krishnamurthy (2008). "Collective Risk Management in a Flight to Quality Episode," *Journal of Finance*, 63(5): 2195–2230.

[6] Caballero, Ricardo J. and Alp Simsek (2013). "Fire Sales in a Model of Complexity," *Journal of Finance*, 68(6): 2549–2587.

[7] Cain, Michael, David Law, and David Peel (2000). "The Favourite-Longshot Bias and Market Efficiency in UK Football Betting," *Scottish Journal of Political Economy*, 47(1): 25–36.

[8] Chow, Gregory C. (1960). "Tests of Equality Between Sets of Coefficients in Two Linear Regressions," *Econometrica*, 28(3): 591–605.

[9] Chu, Chia-Shang J., Kurt Hornik, and Chung-Ming Kuan (1995a). "MOSUM Tests for Parameter Constancy," *Biometrika*, 82(3): 603–617.

[10] Chu, Chia-Shang J., Kurt Hornik, and Chung-Ming Kuan (1995b). "The Moving-Estimates Test for Parameter Stability," *Econometric Theory*, 11(4): 699–720.

[11] Dow, James and Sérgio Ribeiro da Costa Werlang (1992). "Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio," *Econometrica*, 60(1): 197–204.

[12] Ellsberg, Daniel (1961) "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, 75(4): 643–669.

[13] Epstein, Larry G. and Martin Schneider (2007). "Learning under Ambiguity," *Review of Economic Studies*, 74(4): 1275–1303.

[14] Epstein, Larry G. and Martin Schneider (2010). "Ambiguity and Asset Markets," *Annual Review of Financial Economics*, 2: 315–346.

[15] Gilboa, Itzhak and David Schmeidler (1989). "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18(2): 141–153.

[16] Guidolin, Massimo and Francesca Rinaldi (2010). "A Simple Model of Trading and Pricing Risky Assets under Ambiguity: Any Lessons for Policy-Makers?" *Applied Financial Economics*, 20(1-2): 105–135.

[17] Huber, Peter J. and Elvezio M. Ronchetti (2009). *Robust Statistics*. 2nd ed. Hoboken, N.J.: Wiley.

[18] Ipeirotis, Panos (2013). *Intrade-Archive*. GitHub repository, https://github.com/ipeirotis/Intrade-Archive.

[19] Kahneman, Daniel and Amos Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47(2): 263–292.

[20] Knight, Frank H. (1921). *Risk, Uncertainty and Profit*. Boston; New York: Houghton Mifflin Company.

[21] Manski, Charles F. (2006). "Interpreting the Predictions of Prediction Markets," *Economics Letters*, 91(3): 425–429.

[22] Routledge, Bryan R. and Stanley E. Zin (2009). "Model Uncertainty and Liquidity," *Review of Economic Dynamics*, 12(4): 543–566.

[23] Shin, Hyun Song (1992). "Prices of State Contingent Claims with Insider Traders, and the Favourite-Longshot Bias," *Economic Journal*, 102(411): 426–435.

[24] Snowberg, Erik and Justin Wolfers (2010). "Explaining the Favorite-Long Shot Bias: Is it Risk-Love or Misperceptions?" *Journal of Political Economy*, 118(4): 723–746.

[25] Sutherland, Wilson A. (1975). *Introduction to Metric and Topological Spaces.* Oxford: Clarendon Press.

[26] Thaler, Richard H. and William T. Ziemba (1988). "Anomalies: Parimutuel Betting Markets: Racetracks and Lotteries," *Journal of Economic Perspectives*, 2(2): 161–174.

[27] Uhlig, Maxim (2013). "Inflation Ambiguity and the Term Structure of U.S. Government Bonds," *Journal of Monetary Economics*, 60(2): 295–309.

[28] Wolfers, Justin and Eric Zitzewitz (2006). "Interpreting Prediction Market Prices as Probabilities," NBER Working Paper, No. 12200.

[29] Weitzman, Martin (1965). "Utility Analysis and Group Behavior: An Empirical Study," *Journal of Political Economy*, 73(1): 18–26.

Appendix A Proofs

Proof of Proposition 1. Note that any equilibrium price π has to satisfy (i) $\pi > 2\epsilon$ and (ii) $\pi < 1-2\epsilon$. If (i) does not hold, then $\pi^* \leq q_i + \epsilon$ for all i, which means any trader will have either a long position or a zero position—not an equilibrium. Similarly, if (ii) does not hold, no trader will have a long position, which cannot be an equilibrium either.

Substitute (2) into (3) and rewrite the aggregate demand as

$$X_F(\pi) = \int_{\epsilon}^{\pi-\epsilon} \frac{q+\epsilon-\pi}{\pi(1-\pi)} w dF(q) + \int_{\pi+\epsilon}^{1-\epsilon} \frac{q-\epsilon-\pi}{\pi(1-\pi)} w dF(q).$$

Hence, $X_F(\pi) = 0$ if and only if

$$\begin{aligned}
& \int_{\epsilon}^{\pi-\epsilon} (q+\epsilon-\pi)\, dF(q) + \int_{\pi+\epsilon}^{1-\epsilon} (q-\epsilon-\pi)\, dF(q) = 0 \\
\Leftrightarrow\ & \int_{\epsilon}^{\pi-\epsilon} (q-\pi)\, dF(q) + \int_{\pi+\epsilon}^{1-\epsilon} (q-\pi)\, dF(q) + \int_{\epsilon}^{\pi-\epsilon} \epsilon dF(q) - \int_{\pi+\epsilon}^{1-\epsilon} \epsilon dF(q) = 0 \\
\Leftrightarrow\ & \mathbb{E}_F(q) - \pi - \int_{\pi-\epsilon}^{\pi+\epsilon} (q-\pi)\, dF(q) + \epsilon\left[F(\pi-\epsilon) + F(\pi+\epsilon) - 1\right] = 0 \\
\Leftrightarrow\ & \mathbb{E}_F(q) - \pi + \int_{\pi-\epsilon}^{\pi+\epsilon} F(q)\, dq - \left[(q-\pi) F(q)\right]_{\pi-\epsilon}^{\pi+\epsilon} + \epsilon\left[F(\pi-\epsilon) + F(\pi+\epsilon) - 1\right] = 0,
\end{aligned}$$

where the last step follows from integration by parts. Simplifying and rearranging terms yields the stated expression in the proposition. ■

Proof of Proposition 2. Let $\mathcal{G}$ be the space of distribution functions over $[\epsilon, 1-\epsilon]$, endowed with the Lévy metric ℓ, where

$$\begin{aligned}
& \ell(G_1, G_2) \\
\equiv\ & \inf\left\{\varepsilon > 0 \mid G_1(q-\varepsilon) - \varepsilon \leq G_2 \leq G_1(q+\varepsilon) + \varepsilon \text{ for all } q \in [\epsilon, 1-\epsilon]\right\}
\end{aligned}$$

for any $G_1, G_2 \in \mathcal{G}$. Let $\mathcal{F}$ be the subset of $\mathcal{G}$ that satisfies $\pi_F^* = \mathbb{E}_F(q)$ for any $F \in \mathcal{F}$. Since the Lévy metric metrizes the weak topology,[10] the proposition is equivalent to the claim that $\mathcal{F}$ is nowhere dense in $(\mathcal{G}, \ell)$.

[10] See, e.g., Huber and Ronchetti (2009), p. 28.

Note that $\mathcal{F}$ is closed. Since a set is nowhere dense if and only if the complement of its closure is dense,[11], it remains to be shown $\mathcal{G} \setminus \mathcal{F}$ is dense, that is, for any point in $\mathcal{G}$, there is a sequence from $\mathcal{G} \setminus \mathcal{F}$ converging to that point. It is thus enough to show, for any $F \in \mathcal{F}$ and any $\delta > 0$, there exists some $G \in \mathcal{G} \setminus \mathcal{F}$ such that $\ell(F, G) < \delta$.

F is non-decreasing since it is a distribution function. It follows that

$$\lim_{q \to [\mathbb{E}_F(q)+\epsilon]^-} F(q) \geq F(\mathbb{E}_F(q) - \epsilon).$$

We show prove the results by examining two cases.

Case 1: $\lim_{q \to [\mathbb{E}_F(q)+\epsilon]^-} F(q) > F(\mathbb{E}_F(q) - \epsilon)$.

Given $\delta > 0$, we construct a distribution function G from F as

$$G(q) \equiv \begin{cases} F(q) & \text{if } q \in [\epsilon, \mathbb{E}_F(q) - \epsilon - \delta_1), \\ F(\mathbb{E}_F(q) - \epsilon) & \text{if } q \in [\mathbb{E}_F(q) - \epsilon - \delta_1, \mathbb{E}_F(q) + \epsilon + \delta_2), \\ F(q) & \text{if } q \in [\mathbb{E}_F(q) + \epsilon + \delta_2, 1 - \epsilon], \end{cases}$$

where $\delta_1, \delta_2 > 0$ are such that function $g \equiv G - F$ satisfies conditions

$$\int_{\mathbb{E}_F(q)-\epsilon-\delta_1}^{\mathbb{E}_F(q)+\epsilon+\delta_2} g(q)\, dq = 0$$

and

$$\max\{g(\mathbb{E}_F(q) - \epsilon - \delta_1), -g(\mathbb{E}_F(q) + \epsilon + \delta_2)\} = \frac{\delta}{2}.$$

It is easily verified that G is a mean-preserving spread of F, with two new atoms created at points $\mathbb{E}_F(q) - \epsilon - \delta_1$ and $\mathbb{E}_F(q) + \epsilon + \delta_2$. By construction, this implies that

$$\begin{aligned} \int_{\mathbb{E}_G(q)-\epsilon}^{\mathbb{E}_G(q)+\epsilon} G(q)\, dq &= \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} G(q)\, dq \\ &= \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} [F(q) + g(q)]\, dq \\ &= \epsilon + \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} g(q)\, dq < \epsilon, \end{aligned}$$

[11] See, e.g., Sutherland (1975), p. 64.

where the last equality holds because $F \in \mathcal{F}$, and the inequality is due to $g(\mathbb{E}_F(q) + \epsilon) < 0$ which implies $\int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} g(q)\,dq < 0$. Since $\int_{\mathbb{E}_G(q)-\epsilon}^{\mathbb{E}_G(q)+\epsilon} G(q)\,dq < \epsilon$, $G \in \mathcal{G} \setminus \mathcal{F}$. Finally, let ρ be the uniform metric, that is,

$$\rho(G_1, G_2) \equiv \sup\{|G_1(q) - G_2(q)| \mid q \in [\epsilon, 1-\epsilon]\}$$

for any $G_1, G_2 \in \mathcal{G}$. By construction, $\rho(F, G) = \frac{\delta}{2}$. Since the Lévy metric is bounded by the uniform metric from above, that is, $\ell(G_1, G_2) \leq \rho(G_1, G_2)$ for any $G_1, G_2 \in \mathcal{G}$, we have $\ell(F, G) \leq \frac{\delta}{2} < \delta$.

Case 2: $\lim_{q \to [\mathbb{E}_F(q)+\epsilon]^-} F(q) = F(\mathbb{E}_F(q) - \epsilon)$.

Given $\delta > 0$, we construct a distribution function H from F as

$$H(q) \equiv \begin{cases} F(q) & \text{if } q \in [\epsilon, \mathbb{E}_F(q) - \epsilon), \\ F(\mathbb{E}_F(q) - \epsilon) + \delta_3 & \text{if } q \in [\mathbb{E}_F(q) - \epsilon, \mathbb{E}_F(q) + \epsilon + \delta_4), \\ F(q) & \text{if } q \in [\mathbb{E}_F(q) + \epsilon + \delta_4, 1 - \epsilon], \end{cases}$$

where $\delta_3, \delta_4 > 0$ are such that function $h \equiv H - F$ satisfies conditions

$$\int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon+\delta_4} h(q)\,dq = 0$$

and

$$\max\{\delta_3, -h(\mathbb{E}_F(q) + \epsilon + \delta_4)\} = \frac{\delta}{2}.$$

It is easily verified that H is a mean-preserving spread of F, with two new atoms created at points $\mathbb{E}_F(q) - \epsilon$ and $\mathbb{E}_F(q) + \epsilon + \delta_4$. By construction, this implies that

$$\begin{aligned} \int_{\mathbb{E}_G(q)-\epsilon}^{\mathbb{E}_G(q)+\epsilon} H(q)\,dq &= \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} H(q)\,dq \\ &= \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} [F(q) + h(q)]\,dq \\ &= \epsilon + \int_{\mathbb{E}_F(q)-\epsilon}^{\mathbb{E}_F(q)+\epsilon} h(q)\,dq \\ &= \epsilon + 2\epsilon\delta_3 > \epsilon, \end{aligned}$$

where the last but second equality holds because $F \in \mathcal{F}$, and the last equality follows from the construction of H. Since $\int_{\mathbb{E}_G(q)-\epsilon}^{\mathbb{E}_G(q)+\epsilon} H(q)\,dq > \epsilon$, $H \in \mathcal{G} \setminus \mathcal{F}$. Finally, similar to Case 1, we have $\rho(F,H) = \frac{\delta}{2}$ and, hence, $\ell(F,H) < \delta$. ■

Proof of Proposition 3. Decompose $X_F(\pi)$ into the aggregate supply (shorts) $S_F(\pi)$ and the aggregate demand (longs) $D_F(\pi)$, where

$$S_F(\pi) = \int_{\epsilon}^{\pi-\epsilon} -\frac{q+\epsilon-\pi}{\pi(1-\pi)} w dF(q)\,,\ D_F(\pi) = \int_{\pi+\epsilon}^{1-\epsilon} \frac{q-\epsilon-\pi}{\pi(1-\pi)} w dF(q)\,,$$

and $S_F(\pi_F^*) = D_F(\pi_F^*)$ in equilibrium. We show that an increase in ϵ shifts the supply curve inwards. That is,

$$\frac{dS_F(\pi)}{d\epsilon} = 0 + \frac{\epsilon+\epsilon-\pi}{\pi(1-\pi)} w dF(\epsilon) - \int_{\epsilon}^{\pi-\epsilon} \frac{\partial}{\partial\epsilon}\frac{q+\epsilon-\pi}{\pi(1-\pi)} w dF(q) < 0.$$

Similarly, an increase in ϵ shifts the demand curve inwards (i.e., $\frac{dD_F(\pi)}{d\epsilon} < 0$). It follows that the equilibrium quantity of trade—$S_F(\pi_F^*)$, or $D_F(\pi_F^*)$—has to be smaller as the degree of ambiguity increases. ■

Proof of Proposition 4. Let Φ denote the integral of F, i.e., $\Phi(q) \equiv \int_{\epsilon}^{q} F(q')\,dq'$. It follows from the definition of F that

$$\Phi(q) = \int_{\epsilon}^{q} F(q')\,dq' = \begin{cases} (1-m)\,\overline{\Phi}(q) & \text{if } q \in [\epsilon, p)\,, \\ (1-m)\,\overline{\Phi}(q) + m(q-p) & \text{if } q \in [p, 1-\epsilon]\,, \end{cases}$$

where $\overline{\Phi}$ is the integral of $\overline{F}$. The equilibrium condition becomes

$$\begin{aligned} \pi &= \mathbb{E}_F(q) + \Phi(\pi+\epsilon) - \Phi(\pi-\epsilon) - \epsilon \\ &= 1 - 2\epsilon - \Phi(1-\epsilon) + \Phi(\pi+\epsilon) - \Phi(\pi-\epsilon)\,, \end{aligned}$$

where the second equality follows from integration by parts. Since $\Phi(q)$ has a kink at point p, the equilibrium price depends on the position of p relative to $\pi+\epsilon$ and $\pi-\epsilon$.

<u>Case 1</u>: $\pi - \epsilon \leq p \leq \pi + \epsilon$.

The equilibrium condition is rewritten as

$$\begin{aligned}\pi &= 1-2\epsilon-(1-m)\,\overline{\Phi}\,(1-\epsilon)-m\,(1-\epsilon-p)\\&\quad+(1-m)\,\overline{\Phi}\,(\pi+\epsilon)+m\,(\pi+\epsilon-p)-(1-m)\,\overline{\Phi}\,(\pi-\epsilon)\,.\end{aligned}$$

Rearranging terms and dividing both sides by $1-m$ yields

$$\pi-\overline{\Phi}\,(\pi+\epsilon)+\overline{\Phi}\,(\pi-\epsilon)=1-2\epsilon-\overline{\Phi}\,(1-\epsilon)\,.$$

<u>Case 2</u>: $p>\pi+\epsilon$.

The equilibrium condition is rewritten as

$$\begin{aligned}\pi &= 1-2\epsilon-(1-m)\,\overline{\Phi}\,(1-\epsilon)-m\,(1-\epsilon-p)\\&\quad+(1-m)\,\overline{\Phi}\,(\pi+\epsilon)-(1-m)\,\overline{\Phi}\,(\pi-\epsilon)\,.\end{aligned}$$

Rearranging terms yields

$$\frac{\pi}{1-m}-\overline{\Phi}\,(\pi+\epsilon)+\overline{\Phi}\,(\pi-\epsilon)=1-2\epsilon-\overline{\Phi}\,(1-\epsilon)+\frac{(p-\epsilon)\,m}{1-m}. \tag{4}$$

Note that the left-hand side of equation (4) is strictly increasing in π. Thus, the solution π^* to the equation is a continuous and strictly increasing function of p. Furthermore, as $p\longrightarrow\widehat{\pi}+\epsilon$, where $\widehat{\pi}$ is the equilibrium price in Case 1, the right-hand side of equation (4) converges to $1-2\epsilon-\overline{\Phi}\,(1-\epsilon)+\frac{\widehat{\pi}m}{1-m}$, and the solution to the equation converges to $\widehat{\pi}$. In other words, the equilibrium price is continuous at point $p=\widehat{\pi}+\epsilon$.

Next, we show $\pi^*\,(1-\epsilon)<1-2\epsilon$, which implies $\pi^*\,(1-\epsilon)<1-\epsilon$ in part 1 of the proposition. Let $LHS\,(\pi)$ and $RHS\,(p)$ denote the left- and right-hand sides of equation (4), as functions of π and p, respectively. Note that

$$\begin{aligned}LHS\,(1-2\epsilon)-RHS\,(1-\epsilon) &= \left[\frac{1-2\epsilon}{1-m}-\overline{\Phi}\,(1-\epsilon)+\overline{\Phi}\,(1-3\epsilon)\right]\\&\quad-\left[1-2\epsilon-\overline{\Phi}\,(1-\epsilon)+\frac{(1-2\epsilon)\,m}{1-m}\right]\\&= \overline{\Phi}\,(1-3\epsilon)>0.\end{aligned}$$

Since LHS is strictly increasing in π, the solution to the equation when $p = 1 - \epsilon$ must be smaller than $1 - 2\epsilon$.

Case 3: $p < \pi - \epsilon$.

The equilibrium condition is rewritten as

$$\begin{aligned}\pi &= 1 - 2\epsilon - (1-m)\,\overline{\Phi}\,(1-\epsilon) - m\,(1-\epsilon-p) \\ &\quad + (1-m)\,\overline{\Phi}\,(\pi+\epsilon) + m\,(\pi+\epsilon-p) \\ &\quad - (1-m)\,\overline{\Phi}\,(\pi-\epsilon) - m\,(\pi-\epsilon-p)\,.\end{aligned}$$

Rearranging terms yields

$$\frac{\pi}{1-m} - \overline{\Phi}\,(\pi+\epsilon) + \overline{\Phi}\,(\pi-\epsilon) = 1 - 2\epsilon - \overline{\Phi}\,(1-\epsilon) + \frac{(p+\epsilon)\,m}{1-m}. \tag{5}$$

Similar to Case 2, the solution π^* to equation (5) is continuous and strictly increasing in p, and it converges to $\widehat{\pi}$ as $p \to \widehat{\pi} - \epsilon$. Hence, the equilibrium price is continuous at point $p = \widehat{\pi} - \epsilon$ as well.

Next, we show $\pi^*(\epsilon) > 2\epsilon$, which implies $\pi^*(\epsilon) > \epsilon$ in part 1 of the proposition. Again, let $LHS\,(\pi)$ and $RHS\,(p)$ denote the left- and right-hand sides of equation (5). Note that

$$\begin{aligned}LHS\,(2\epsilon) - RHS\,(\epsilon) &= \left[\frac{2\epsilon}{1-m} - \overline{\Phi}\,(3\epsilon) + \overline{\Phi}\,(\epsilon)\right] \\ &\quad - \left[1 - 2\epsilon - \overline{\Phi}\,(1-\epsilon) + \frac{2\epsilon m}{1-m}\right] \\ &= \left[\overline{\Phi}\,(1-\epsilon) - \overline{\Phi}\,(3\epsilon)\right] - \left[(1-\epsilon) - 3\epsilon\right] < 0,\end{aligned}$$

where the last inequality holds because $\overline{\Phi}$ is the integral of distribution function $\overline{F}$ over $[\epsilon, 1-\epsilon]$. Since LHS is strictly increasing in π, the solution to the equation when $p = \epsilon$ must be larger than 2ϵ. ■

Proof of Corollary 5. Recall that $\widehat{\pi}$ is identified by equation

$$\widehat{\pi} - \overline{\Phi}\,(\widehat{\pi}+\epsilon) + \overline{\Phi}\,(\widehat{\pi}-\epsilon) = 1 - 2\epsilon - \overline{\Phi}\,(1-\epsilon)\,.$$

The symmetry of $\overline{F}$ implies $-\overline{\Phi}\,(1-x) = (x-\epsilon) - \overline{\Phi}\,(x)$ for any $x \in [\epsilon, 1-\epsilon]$. Thus, the equilibrium condition becomes

$$\begin{aligned} & \widehat{\pi} - \left[\widehat{\pi} + \overline{\Phi}\,(1-\widehat{\pi}-\epsilon)\right] + \overline{\Phi}\,(\widehat{\pi}-\epsilon) = 1-2\epsilon - \left[1-2\epsilon + \overline{\Phi}\,(\epsilon)\right] \\ \Leftrightarrow \quad & \overline{\Phi}\,(1-\widehat{\pi}-\epsilon) - \overline{\Phi}\,(\widehat{\pi}-\epsilon) = \overline{\Phi}\,(\epsilon) = 0, \end{aligned}$$

to which $\widehat{\pi} = 0.5$ is the only solution. ■

Appendix B Empirics

The historical data of Intrade was archived by Ipeirotis (2013) and is available on GitHub. Table 4 lists all the categories of events and the number of markets within each category. We complete the dataset by creating an outcome variable and recording how each random event had turned out. The outcome equals 1 if an event occurs, and it equals 0 if its complement event occurs.

Some markets have correlated outcomes, because they are about the same, uncertain circumstances. For example, concerning the 2012 U.S. Republican Party presidential nominee, there are 53 separate markets corresponding to 53 possible winners, including Mitt Romney, Rick Santorum, Ron Paul, Newt Gingrich, and "any other individual" not specified by the prediction platform. To avoid such correlation in the observations, for each group of these correlated markets, we randomly select one market into the aggregate sample and disregard the rest.

The total number of selected markets included in the final analysis also shown in Table 5. The table lists the number of observations—the total as well as the number of observations per percentile bin—for political events, entertainment events, and the full sample. The dataset is skewed towards political and entertainment events, as the two categories together accounts for 82% of the full sample.

Table 4: Intrade Data: Event Categories and Number of Markets.

Event category	Number of markets
Art	60
Business	43
Chess	52
Climate & Weather	861
Construction & Engineering	9
Current Events	1540
Education	1
Entertainment	8715
Fine Wine	5
Foreign Affairs	87
Legal	310
Media	10
Politics	5460
Real Estate	2
Science	20
Social & Civil	30
Technologies	65
Transportation	11

Table 5: Intrade Data: Number of Observations in Final Analysis.

Event category	Total observations	Observations per bin	
		(50 bins)	(30 bins)
Politics	897	18	30
Entertainment	1157	23	39
Full sample	2509	50	84