Prediction Market Policies

Saitulaa Naranong *

Abstract

Prediction markets are a mechanism for collating participants' beliefs into a probability that a future outcome will take place. Their historical purpose was to provide governments, management, and other leadership with accurate forecasts, in order to inform future decision-making. Participants of prediction markets are financially rewarded for accurate predictions, incentivizing the communication of well-researched, unbiased views. Using qualitative research methods and meta-synthesis of prediction market implementations, this study examines practical approaches to prediction markets, using them to inform future practices. The provision of market liquidity by Automated Market Makers (AMM) is described, and examples of AMMs are detailed.

This article provides a policy framework for prediction markets, with a view to preventing negative side effects that may arise from their practice. Areas of potential concern are examined, with the example of a bank run as a case study, and restrictions on prediction market structure and on permissible contracts are devised as necessary to avoid adverse outcomes.

Keywords: Prediction Market, Automated Market Maker, Policy, Bank Run

---

* Department of Mathematics and Statistics, Faculty of Science and Technology, Thammasat University

Email: saichu.classes@gmail.com

Received: June 17, 2025. Revised: June 27, 2025. Accepted: June 29, 2025

Introduction

Effective policymaking benefits significantly from accurate predictions as to the results of proposed or alternative courses of action. Prior knowledge about the effects on economic growth, unemployment, company failures, and other factors can help a government, regulators, or even the management of individual companies select preferable policies. As such, the ability to synthesize existing information into accurate predictions, which in turn may be used as inputs for further predictions, is of great theoretical value in governance. Prediction markets provide a mechanism for such, and have advantages over existing systems for at least some applications. At the same time, there has been considerable concern over the possible negative side effects of prediction markets (Wolfers & Zitzewitz, 2004; Shi, Conitzer, Guo, 2009; Bell, 2011). This paper proposes a cautious framework for prediction markets with a priority of preventing such side effects.

For some mechanisms of truth-seeking, there is no formal method of resolving a spectrum of views into a single probability that may be used for the purpose of policy (Van Bouwel, 2015): imagine a publication full of divergent predictions of economic growth based on different models, none of which contain any obvious errors. Prediction markets are generally designed to output probabilities, with the market clearing price (or some similarly determinable attribute) being the market's best guess of the likelihood of the outcome in question, thereby providing the service of incorporating and evaluating all models. Certain mechanisms for truth-seeking may also be vulnerable to conflicting incentives, such as conformity pressure (Goncalo & Duguid, 2012; Strickland & Crowne, 1962). Participants in economic experiments have been known to report beliefs that were more conservative than their actual beliefs, a result of pull-to-center effects from risk aversion (Schotter & Trevino, 2014). Prediction markets align the financial incentives of the predictor with the accuracy of their predictions, obviating these and many other issues: forecasters gain profit based on correctness, and no other considerations.

These theoretical advantages to prediction markets have borne fruit: studies have found prediction markets to be more accurate than polls for predicting election results (Berg, Forsythe, Nelson, & Rietz, 2008), more accurate than surveys for collecting and filtering out promising product ideas in an organization (Soukhoroukova, Spann, & Skiera, 2012), more accurate than expert predictions for business forecasting (Spann & Skiera, 2003), and more accurate than official company sales forecasts (Chen & Plott, 2002). Probability forecasts from the Intelligence Community Prediction Market were found to be more accurate than those of intelligence analysts, even adjusting for the possibility that more recent information was accessible to market participants (Stastny & Lehner, 2018). Atanasov et al., 2017, studying the prediction of geopolitical events, found prediction markets to be more accurate than prediction polls, though less accurate than team prediction polls after the latter were statistically transformed, including by weighting against the past performance of teams—particularly when the resolution date was still far in the future.

Although prediction markets have existed in some form for hundreds of years, usually in the form of betting on political outcomes (Rhode & Strumpf, 2004), recent digitization and automation have made it practical to pose a wide variety of research questions in these markets. Indeed, systems such as the Automated Market Maker (AMM) now allow participants to propose their own contracts (i.e. their own prediction questions) while ensuring at least some level of liquidity for their contract market. (The “contract market” refers to the sub-market that trades an individual contract.) This paper will not take a position on the ideal level of centralization of such markets, except to provide frameworks to minimize potential negative side effects in the event of participants posing their own questions. It may be decided that prediction markets should conservatively contain only carefully chosen policy questions of wide-ranging import, or that a wider variety of questions should be permitted; that decision is beyond our current scope. As long as the minimum is obtained, the goal of a digital government can be implemented with accurate forecasts leading to beneficial outcomes (Bell, 2011).

This paper employs qualitative research methods and meta-synthesis to examine prediction markets, their implementations and approaches, and potential benefits and drawbacks. The overarching theory of prediction markets is developed, along with the effects of implementation decisions such as prediction market adjudication. As prediction markets have historically been thinly traded, possible sources of liquidity are examined, including different types of AMM. With the meta-synthesis of these various aspects in hand, this paper considers thoroughly the issue of adverse side effects, providing a policy framework to avoid such. This paper examines the case study of a possible bank run as a large-scale potential example of such side effects.

Overview of Prediction Market Design

Prediction markets are financial markets – the items being traded are not physical assets, but rather contracts that pay out based on the outcome of a future event. The prototypical example is of a binary contract, referring to an event taking place or not, and can be worded as a question (e.g. “Will Candidate X win the 2028 U.S. Presidential Election?”) or otherwise (“Candidate X will win the 2028 U.S. Presidential Election”). In fact, binary events such as this generally have two associated contracts, YES and NO (or analogous labels), which pay out if the event in question takes place or does not take place, respectively. Contracts may also be created for events with three or more possible outcomes, the most basic of which still pays out only for the "winner". Contracts may also pay different amounts based on the result of some measure e.g. the number of U.S. Senate seats won by a political party or the percentage of a vote won by a political party (Wolfers & Zitzewitz, 2004). Conditional contracts, whose output is a conditional probability, are only triggered when a condition is met ("if policy A is adopted, economic growth will increase by B"; the contract only takes effect if policy A is actually adopted, otherwise failing to resolve and liquidating), and are useful for dealing with hypotheticals that may end up not taking place.

Since the instrument itself cannot answer questions about real-world outcomes, the outcome of the contract must be adjudicated—the adjudicator may be some authority, whether a designated "prediction judge" (Bell, 2011) or the original contract creator, or a decentralized oracle. An example of the latter is the UMA Optimistic Oracle used by Polymarket, which allows anyone to submit a proposed resolution to the contract, accepting said resolution as correct unless challenged within a time period. Challenges to the proposal are subsequently resolved through a debate/vote process, and both the proposer and challenger must issue bonds, which may be taken away depending on the outcome of the challenge (Heidt, Sandner, & Anders, in press).

Liquidity, for general financial markets, leads to an increase in informed trading (Lee & Chung, 2022), though studies concerning the effect of liquidity on the accuracy of prediction markets have been mixed. Dana et al., 2019 found a larger spread between the accuracy of prediction markets alone, against the accuracy of a weighted average of prediction market and transformed self-reported probabilities (in the latter's favor) during times of low engagement in the market: little trading activity, large bid-ask spreads, and contracts whose resolution is expected to be in months rather than earlier. Tetlock (2008) finds that, on the TradeSports exchange, increased liquidity can reduce the accuracy of contract markets, offering an explanation for the phenomenon concerning non-rational limit order traders. What is clear is that large bid-ask spreads make the market-clearing price, and therefore the market probability, itself less precise. Another important consideration is that liquidity increases the difficulty of various forms of market manipulation (Gu et al., 2024; Horst & Naujokat, 2011) larger markets are more difficult to move. Finally, decentralized oracles have increased vulnerability to oracle manipulation attacks in thinly traded markets, as they rely on feedback from other market participants (Aspembitova & Bentley, 2022). Highly liquid markets with numerous participants provide an ecosystem where false conclusions fed to an oracle are robustly challenged.

Table 1. Discretionary parameters in the design of a prediction market

Discretionary parameters in the design of a prediction market	**Examples**
Contract payoffs	• Binary contracts: YES or NO, with all-or-nothing payoff • Linear payoffs in terms of a continuous variable (Berg & Rietz, 2019)
Types of contracts permitted	• Pre-specified contracts only • Contracts following certain templates • Arbitrary contracts
Choice of adjudicator	• Original contract writer • Designated judge • Decentralized oracle (e.g. UMA Optimistic Oracle)
Automated Market Maker	• None • Dynamic Parimutuel Market Maker • Logarithmic Market Scoring Rule Market Maker • Constant Sum Market Maker

Source: author's analysis

Automated Market Makers

AMMs are computer programs with which a participant can trade even in the absence of other market participants (Bartoletti, Chiang, & Lluch-Lafuente, 2022). As prediction markets have historically been illiquid with few participants, AMMs have helped to provide liquidity and ensure that trade remains possible on an ongoing basis (Slamka, Skiera, & Spann, 2012). Although AMMs first gained popularity in prediction markets (Mohan, 2022), they have since been used and further developed in the field of decentralized finance, which faces similar liquidity concerns (Mohan, 2022; Pourpouneh, Nielsen, & Ross, 2020). As such, advancements of AMMs in that field may continue to be adopted for Prediction Markets as they appear.

Examples of Market Makers include the Logarithmic Market Scoring Rule Market Maker (Hanson, 2007), Constant Sum Market Maker (Schlegel, Kwaśnicki, & Mamageishvili, 2022), Uniswap V3 Market Maker (Adams, Zinsmeister, Salem, Keefer, & Robinson, 2021), and Dynamic Parimutuel Market Maker (Pennock, 2004; Agrawal, Delage, Peters, Wang, & Ye, 2009; Lin & Chen, 2009). This paper will elaborate below on the Dynamic Parimutuel Market Maker and the Constant Sum Market Maker.

Dynamic Parimutuel Market Makers

Dynamic Parimutuel Market Makers (Pennock, 2004; Agrawal et al., 2009) were based on traditional parimutuel systems seen in racetrack betting. In such traditional parimutuel systems, separate pools of money for every possible outcome are built up from participants' wagers on said outcomes, and the "winning bets" receive the sum of the money entered into the "losing bets", less transaction costs (Thaler & Ziemba, 1988). For example, if the outcome space is partitioned into five mutually exclusive outcomes, A, B, C, D, E, if the amount of currency wagered in the winning outcome A is 1,000 THB, and the amount of currency wagered in the other outcomes (less transaction costs) is 8,000 THB, then each 1 THB wagered into outcome A entitles the participant to 8 THB from the losing pools. The implicit specification is that each THB invested into outcome A is converted into a contract, and the total amount invested into losing contracts is distributed evenly amongst winning contracts, hence the reward is calculated as $\frac{8000\,THB}{1000\,Contract} = 8\,\frac{THB}{Contract}$.

Classic Parimutuel systems are elegantly simple in terms of computing market probabilities: at any particular moment, the market probability of outcome X is interpreted to be the proportion of money wagered on option X (Thaler & Ziemba, 1988). In the example above, assuming zero transaction costs for simplicity, the probability of option A during the close of the market would have been $\frac{1000\,THB}{1000\,THB + 8000\,THB} \approx 11.1\,percent.$

The main disadvantage of traditional parimutuel systems is that the payoff of the implicit "contract" (1 THB invested in the correct answer) is unknown ahead of time, as it depends on the total amounts wagered both for the winning option, and for all other options. A related disadvantage is that the time at which the wager is placed does not matter, as all holders of the resulting "contract" receive equivalent rewards regardless for any THB invested, encouraging later bets as more information becomes available and, in the extreme case, encouraging all bets to take place "last minute". The latter is incompatible with the goal of prediction markets, which is to predict events ahead of time. This extreme case can be prevented by closing the parimutuel market long before the event in question, but a prediction market should also ideally make updates to the probabilities as more information becomes available.

Dynamic parimutuel systems attempt to solve several of the latter problems. The market maker at any time prices the contract for outcome X as $P_X = \frac{Q_X}{\sqrt{\sum_{Y\in\Omega} Q_Y^2}}$, where each Q_Y is the amount of THB that has been invested into outcome Y (for all Y in the outcome space Ω). The market probability of outcome X is then $\pi_X = \frac{Q_X^2}{\sum_{Y\in\Omega} Q_Y^2}$ (Lin & Chen, 2009). In words, the market probability of outcome X is the proportion of the quantity invested in X, but with all quantities in the proportion being squared; the price of a contract for X is then the square root of said probability (i.e. $\pi_X = P_X^2$). Because the denominator starts at zero (with no wagers), the market must initially be subsidized, with funds added to at least one Q_Y. This can, in theory, allow for market participants to achieve a net gain at the expense of the subsidizer, but is a valid tradeoff if the subsidizer gains value from the resulting market probability information, as in many prediction markets.

For example, if participants have previously wagered 1000 THB, 2000 THB, 3000 THB, and 4000 THB into outcomes A, B, C, D, respectively, then the current price of a contract for outcome A is:

$$P_A = \frac{1000\,THB}{\sqrt{(1000\,THB)^2 + (2000\,THB)^2 + (3000\,THB)^2 + (4000\,THB)^2}} \approx 0.183$$

which is interpreted as a price of 0.183 THB per contract. The current market probability of outcome A is:

$$\pi_A = \frac{(1000\,THB)^2}{(1000\,THB)^2 + (2000\,THB)^2 + (3000\,THB)^2 + (4000\,THB)^2} \approx 0.033$$

i.e. a 3.3 percent market probability of outcome A occurring (note the invariant $\pi_X = P_X^2$; in this case $0.033 \approx 0.183^2$).

Although this adjustment provides for a variable price and an instantaneous market probability, it does not solve the issue that the payoff per contract is unknown ahead of time.

Constant-Product Market Maker

Polymarket and Manifold use variations of the constant-product market maker (Halawi, Zhang, Yueh-Han, & Steinhardt, 2025). In its base specification, the constant-product market maker manages the reserves of security X and Y by ensuring the following invariant:

$$R_X R_Y = k$$

Where R_X , R_Y are the reserves of security X and Y, respectively, and k is a constant value for the contract market based on the current level of liquidity (Zhang, Chen, & Park, 2018). In words, the market maker ensures that the product of the number of reserves of X and Y remains constant. Thus, a participant looking to purchase reserves of X for their portfolio is obligated to deposit some amount of Y into the reserves to maintain the above invariant; vice versa, a participant looking to sell reserves of X is paid in terms of reserves of Y. If Δ_X and Δ_Y are the changes in X and Y, respectively, after the transaction, then we must have:

$$(R_X + \Delta_X)(R_Y + \Delta_Y) = k$$

Hence

$$R_X R_Y + R_Y \Delta_X + R_X \Delta_Y + \Delta_X \Delta_Y = k$$

$$R_X R_Y + R_Y \Delta_X + (R_X + \Delta_X)\Delta_Y = k$$

$$\Delta_Y = \frac{k - R_X R_Y - R_Y \Delta_X}{R_X + \Delta_X} = \frac{k - R_Y(R_X + \Delta_X)}{R_X + \Delta_X} = \frac{k}{R_X + \Delta_X} - R_Y$$

By symmetry:

$$\Delta_X = \frac{k - R_X R_Y - R_X \Delta_Y}{R_Y + \Delta_Y} = \frac{k - R_X(R_Y + \Delta_Y)}{R_Y + \Delta_Y} = \frac{k}{R_Y + \Delta_Y} - R_X$$

In the context of prediction markets, X and Y are usually YES and NO contracts paying out if the prediction comes to pass or fails to do so, respectively. Thus, YES contracts are priced in terms of NO contracts, and vice versa. The market probability for the outcome taking place is the proportion of YES reserves: $\frac{R_{YES}}{R_{YES}+R_{NO}}$. Finally, there exist mechanisms, which depend on the exact market implementation, for adding liquidity by converting currency into additional contracts, or removing liquidity by converting contracts into currency; said mechanisms change the value of the constant k as well (Zhang et al., 2018; Angeris & Chitra, 2020).

Table 2. Possible approach to relaxing prediction market rules as the regulator obtains increased practical experience of market behavior

	Pilot Program	Large-scale testing	Post-testing
Provider of Liquidity	Automated Market Maker	(In addition) Financial firms testing market making platforms	(In addition) All financial firms
Prediction Contract Templates	Preset questions only, after period of suggestion	Only certain templates allowed	Cautious lifting of restrictions in selected categories, particularly “safer”, more macro-focused issues
Participants	Invited Participants Only	Participants who pass qualification examinations	Possibly all participants (with possible exceptions, such as “insiders”), or simplified qualification examinations

Source: author’s analysis

Policy Recommendations for Prediction Markets

Operation of Prediction Markets

Prediction markets having been historically illiquid, it is essential for one or more Automated Market Makers (as described earlier) to be used as fallback in case any contract market ends up thinly traded. Ideally, prediction markets should be developed into liquid markets. In addition to increasing the difficulty of manipulation (Gu et al., 2024; Horst & Naujokat, 2011), including decentralized oracle manipulation, liquidity also supports participants who carry higher fixed costs, such as research-focused organizations that dedicate time and resources to obtaining accurate answers. Even small purchases can cause significant price movements in illiquid markets, limiting the profit available; in order for research organizations to be adequately funded by potential profit, liquid prediction markets are a necessity.

This paper leaves unspecified the entities that should have the power to open contract markets (i.e. to pose the questions whose outcomes are being predicted). It is possible for government entities to reserve this power for themselves, to delegate it to trusted organizations, to allow any person with a license to open a contract market, or to allow any person to open a contract market with relatively few restrictions. The advantage of a permissive policy is the greater variety of questions posed by independent researchers; an outcome likely to increase innovation in research areas covered by prediction market questions. The disadvantage of a permissive policy is the greater likelihood of side effects and the greater difficulty in regulating against potentially problematic contract markets before the fact. While it is out of the scope of this paper to determine the optimal policy mix, one possible argument is that initially-centralized contract market opening would allow for unforeseen issues to be noticed and corrected while prediction markets are still limited (and thus, hopefully, while the possible dangers are still limited). See Table 2 for one possible approach. As prediction markets become better understood and the practical problems are addressed, conservative expansions of market creation powers to other parties could then be enacted.

At minimum, contracts should be created to monitor the possible futures of important economic indicators. Contracts should exist concerning future Gross Domestic Product (GDP), Employment Rates, Import and Export Rates, Incomes, and similar measures. It is possible for specially-designed contracts to be created for the same measure in order to determine their probability distributions to greater detail. For example, a multi-outcome contract could be created, the various outcome contracts corresponding to possible ranges of GDP (e.g. a sub-contract might pay out if 20,000,000,000,000 THB ≤ GDP< 21,000,000,000,000 THB); the prediction market would then assign a probability to each of these ranges. Additionally, a contract can be created whose payout linearly corresponds to GDP; the market clearing price then corresponds (after a linear transformation) to the market's "best guess" for that future GDP, which amounts to *E*(GDP), the expected value of GDP. This can be further supplemented by a contract that pays as a function of GDP^2, from which we can derive $E(GDP^2)$, and thereby the variance $E(GDP^2) - E(GDP)^2$ and standard deviation $\sqrt{E(\mathrm{GDP}^2) - E(\mathrm{GDP})^2}$ (Wolfers & Zitzewitz, 2004). The number and types of contract markets to be created should be a question of the usefulness of these various statistics as well as the actual observed liquidity of the resulting markets; if current liquidity is too low to support some of these sub-measures, they should be shelved until after the industry has further developed.

Requirements for a Robust Contract

Since prediction market contracts are resolved by adjudication (whether by authority or decentralized oracle), an essential aspect of a successful contract market is the clarity of said contract (Wolfers & Zitzewitz, 2004). Participants and observers should agree, after the fact, whether the terms of the contract have been met, and therefore agree on the correct outcome, making the adjudicator's role a mere formality. Otherwise, the contract market would be mainly predicting the adjudicator's judgment, rather than the event in question. Contentious resolutions would also reduce confidence in the market as a whole, possibly leading to reduced participation and liquidity.

As such, several criteria should be adhered to in formulating a prediction market contract. First, a contract should be bounded in time; there should be an obvious "end date" after which the contract has been resolved, one way or the other. This can sometimes be less obvious than originally intended. An election is commonly considered completed after all votes have been tallied, but it may be contested in court even after the fact, leading to an unclear extension of the contract end date. A less ambiguous contract in terms of duration might be "Candidate X will win the 2028 U.S. Presidential Election as of November 20, 2028, 11:59:59 PM" (although, of course, this does not remove other sources of ambiguity, such as what is considered a "win"). Such durational clarification may not even be desired to begin with, if the intention of the market participants and observers is to determine the actual outcome, regardless of time.

Second, the terms of the contract should be specific and based on unambiguous criteria. A clear example of an ambiguous criterion is that "if policy X is implemented, it leads to a good result": this is wholly dependent on the adjudicator's judgment of the "good result". On the other hand, a criterion of "if policy X is implemented, Gross Domestic Product (GDP) increases by at least 10 percent within two years of implementation" is much clearer, although there still remain other ambiguities potentially subject to adjudication. Generally speaking, a measurable and measured criterion is less ambiguous than a subjective criterion.

It is not generally possible to eliminate *all* ambiguity within a contract, and it may not even be desirable, as in the "end date" example given above. The contract maker often further specifies the details of the contract (GDP data is taken from a selected source, other sources may be used if unavailable; the contract ends 9 quarterly GDP reports after the policy has come into action, rather than the vaguer “two years”). There is no widely-adopted standard for the level of clarification, making it a decision of the contract maker in practice.

A related problem is the sudden appearance of ambiguity where none had been originally anticipated. An election whose results are ordinarily accepted is an example of this: if the election is won by candidate X in terms of votes, but legally challenged, and the challenge ends with candidate Y determined to be the victor, it is ambiguous (unless this situation was written in the contract) whether X’s original victory or Y’s subsequent victory determines the outcome. Wolfers & Zitzewitz (2004) provide additional examples: a client changing a schedule adds ambiguity to a prediction contract concerning whether a project would be “on schedule”, and a senator changing parties the day after an election adds ambiguity to the number of seats actually won by said parties.

Prevention of Adverse Side Effects from Prediction Markets

By its nomenclature, prediction markets were originally intended to predict, but under some circumstances, it is also possible for them to influence the outcome being “predicted”. Say, for example, that a prediction question is “Company X will go bankrupt [in some time period].” It is theoretically possible for the management of Company X to buy a large number of YES shares and cause the event to take place, personally profiting in the process. In extreme cases, prediction markets can even incentivize the harm of people related to the outcome being predicted—an unacceptable side effect. In order to prevent even the possibility of this happening, this paper proposes that prediction markets must deal exclusively with large-scale outcomes, such as the economy as a whole, that are not dependent on the actions of any one person or any small group of people. This prevents both the issue where the integral group affects the outcome, and also prevents the targeting of said integral group by malicious actors.

That being said, although this paper does not recommend it, there is an argument that existing financial markets implicitly tolerate a certain risk of adverse side effects, providing us with an already existing upper bound. In the absence of insider trading laws, management of Company X *would* be able to short the stock (holding a short position similarly allows them to profit from the company’s failure), then negatively influence the future of Company X. As such, it may be an alternative general guideline that prediction markets, correctly implemented, should be designed to do no worse harm than the existing financial markets. This may be done either by restricting the allowed contracts, and/or similar laws against “insider trading” of prediction contracts. The problem with this idea is that prediction markets (prior to restriction) allow a much wider variety of instruments than existing financial markets.

The number of possible loopholes is necessarily multiplied immensely; it becomes far less practical for any rule to catch all possible issues hidden away in contracts devised by arbitrary market participants. By comparison, laws protecting existing financial markets were developed over decades or centuries (depending on the exact market) and do not prevent all malicious action even now. As such, this paper ***does not*** recommend that arbitrary prediction contracts be permitted in the current policy; the author expects any number of unintended consequences from such a policy.

Contracts should, instead, be built from certain allowed templates only. Said templates should be carefully scrutinized for potential side effects, with restrictions posed as necessary to prevent said effects. Generally speaking, the broader the macroeconomic indicator, the less likely the contract is to cause side effects.

- Global or International indicators are probably largely immune from the actions of small groups of people, particularly when the indicator cannot be directly adjusted by policymakers. For example, a policymaker may adjust the amount of trade between two nations or the price of a commodity through international policy, but they are less likely to be able to directly adjust, say, the size of the global economy or the amount of global trade.
- National-level statistics may still be influenced by individual policymakers, although even then, they are more difficult to change than, say, that of a company. Macroeconomic policy is necessary, wide-scale and indirect, and there is a greater possibility of pushback from policy opponents.
- Industry-level contracts (e.g. "industry revenue will increase by…", "industry profits will increase by …") may be broad enough to avoid side effects for diversified, well-established industries. Potential problems appear, however, in the use of similar contracts for nascent or concentrated industries, whose trajectories may be highly dependent on the actions of only a few people.

Any sub-industry level contract becomes vulnerable to side effects. At the very least, they are vulnerable to the same side effects that might come about from, for example, the stock, bond, or options market—only to an even greater degree, as contracts generally have more specificity than those of the ordinary financial instruments. As an example, natural questions for prediction markets would be whether:

- Company X adopts policy Y
- Company X meets its sales target

These and numerous other examples can be easily influenced by company management (simply by actually adopting policy Y or by shutting down sales, for example), and thus are not aligned with this paper's cautious approach to prediction markets. One possible approach may be to extend insider trading rules to such contracts, but this fails to account for other possible side effects, such as the sabotage of management who advocate for policy Y, or the spread of malicious rumors to cut sales—both these examples might theoretically be enacted by an "outsider". Importantly, these examples are not exhaustive: the greater the variety of contracts, the greater the number of unforeseen side effects possible, including from otherwise unrelated persons.

Case Study: Bank Runs

Topics or events where the prediction itself is likely to drastically affect the outcome, a concept sometimes called reflexivity (Soros, 2013), are problematic for prediction markets as an analysis tool. A prediction of certain poor outcomes, when numerous watchers of the market realize and act upon it, can lead to even poorer outcomes, leading to even worse predictions in a negative feedback loop. While any mechanism that provides information must end up with some degree of this phenomenon, extreme or exceptional cases can and should be avoided and carefully guarded against.

One example of such an extreme case is a bank run (Oesterheld, Treutlein, Cooper, & Hudson, 2023), during which mere rumors of a bank's insolvency may cause depositors to rapidly withdraw cash, causing the bank to run out of liquid cash; in such a circumstance the bank may default, freeze withdrawals, or liquidate investments at disadvantageous values and potentially render itself insolvent in truth (Andolfatto, Nosal, & Sultanum, 2017, He & Manela, 2016). Prediction market probabilities are generally more credible than mere rumors; it is therefore a reasonable worry that a market predicting a bank failure with medium or high probability may cause that very same bank failure. This may be the case even if the market prediction was relatively unreliable—if the market was thinly traded, or manipulated, or temporarily irrational—considering issues of exaggeration, omission, and accuracy in the communication of news to the layman, both through official media sources (Moore & Singletary, 1985; Sumner et al., 2014; Chang, 2015) and unofficial social media sources (Collins, Hoang, Nguyen, & Hwang, 2021; Olan, Jayawickrama, Arakpogun, Suklan, & Liu, 2024).

A run on an individual bank can further lead to contagion, as depositors grow concerned about similar banks (Kaufman, 1988), interbank lending (Morrison & White, 2013), or companies holding similar assets (Webel, 2013), leading to side effects not only with the bank whose fate is being predicted, but with the industry itself. The speed and contagiousness of bank runs has potentially increased in recent years due to the increased speed of modern withdrawal methods (Ofir & Elmakiess, 2025), along with the proliferation of social media, as seen with the recent collapses of Silicon Valley Bank (SVB) and Signature Bank (Cookson, Fox, Gil-Bazo, Imbet, & Schiller, 2023). Social media rumors may amplify further loss even in unrelated markets (Dosumu, Sakariyahu, Oyekola, & Lawal, 2023).

It is noteworthy that the potential contagion of SVB and Signature Bank to other banks was halted with the Federal Deposit Insurance Corporation (FDIC) guarantee that all deposits in those two banks would be insured, not only the $250,000 per account usually covered by the FDIC (Spitler, 2024), as it removed the incentive for depositors to make further withdrawals. One possible argument, then, is that in the event that full deposit insurance is enacted, the possibility of a bank run prediction market inducing an actual bank run would be reduced—though not completely eliminated in the general case, as it would depend on the depositors' confidence in said insurance. However, such a system would lead to one of two possible cases:

1. The prediction market concerning the bank run would be continually priced at the NO option, rendering it of limited utility (this is the less likely option), or
2. The prediction market would reflect the probability that both the bank and the insurance program fail, which would lead to similar side effects, only amplified, as the failure of the insurance program is likely to be systemic.

As such, the bank run prediction market in the case of full deposit insurance would be either of limited usefulness, or incur similar (and perhaps worse) side effects. In all cases, this paper recommends *against* the allowance of company-level bankruptcy-related prediction questions.

It may be observed that the prices of stocks, bonds, credit default swaps, etc., in existing financial markets serve a similar function of predicting bankruptcy (Hull & White, 2003), limiting the incremental danger of adding a new financial instrument. On the other hand, many of these involve an indirect calculation of risk, reducing their accessibility compared to prediction markets, most of which are intended to provide a large ecosystem of observers with the actual probabilities. While the arguments concerning danger levels are noteworthy, their evaluation is beyond this paper's scope, which is to try to avoid side effects altogether for our proposed framework.

Neither are industry-level bankruptcy-related predictions free from issue. The tendency of banks to undertake correlated investments (Acharya & Yorulmazer, 2008), partially due to common regulation (Morrison & White, 2013), is known to market participants, allowing them to infer that what is true for some banks may be true for others as well. Indeed, Morrison and White (2013) argue that potential failures diminish the reputation of the regulatory system, causing other similarly regulated banks to be called into question even in the absence of any other failure mechanism. Psychologically, the group contagion effect is known to cause aversion if an otherwise-anonymous object in a group is known to have a negative quality (Mishra, Mishra, & Nayakankuppam, 2009); it is plausible that predicting "some otherwise-anonymous bank in the industry is insolvent" would similarly cause issues for the industry as a whole.

On the other hand, questions may be allowed that gauge the general health of the industry, rather than being tied directly, or too closely, to bankruptcy. Return on Equity, loan size, and size of the deposit base for the industry as a whole are likely to be admissible. Future expansions beyond the scope of our current suggestions may well extend similar measures to single companies, if, after initial trials, the possible side effects become better understood.

Conclusion

Prediction markets provide participants with financial incentives for researching the correct answer to posed questions, leading to improved accuracy over other methods of prediction and a resulting benefit for the study of policy and governance. This paper provides a proposed policy framework for practical prediction markets, with an emphasis on preventing adverse side effects. An overarching theory of prediction markets is posed, and an initial meta-synthesis of previous examples examines variations of AMMs and Decentralized Oracles as solutions to prediction market issues. Discretionary parameters in the design of a prediction market are specified (Table 1), as are possible policy variations of the prediction market industry as regulators obtain increased practical experience (Table 2). This study's cautious approach necessarily rules out prediction market questions whose answers are easily affected by a small number of people, avoiding many of the common side effects. With the prediction market framework restricted to questions about the industry level or larger, the paper examines the informative case study of banks and the banking industry, which is affected by crises of confidence. Prediction markets must be cautiously designed not to cause bank runs or worse: contagion of a single bank's failure to the entire industry. Through this case study, we examine less obvious restrictions that must be applied to prediction market questions in order to prevent this and similar side effects.

References

Acharya, V. V., & Yorulmazer, T. (2008). Information contagion and bank herding. *Journal of Money Credit and Banking*, *40*(1), 215–231. https://doi.org/10.1111/j.1538-4616.2008.00110.x

Adams, H., Zinsmeister, N., Salem, M., Keefer, R., & Robinson, D. (2021). *Uniswap v3 core*. uniswap https://app.uniswap.org/whitepaper-v3.pdf

Agrawal, S., Delage, E., Peters, M., Wang, Z., & Ye, Y. (2009). A unified framework for dynamic pari-mutuel information market design. In *Proceedings of the 10th ACM conference on Electronic commerce* (pp. 255-264). Association for Computing Machinery.

Andolfatto, D., Nosal, E., & Sultanum, B. (2017). Preventing bank runs. Theoretical Economics, 12(3), 1003–1028. https://doi.org/10.3982/te1970

Angeris, G., & Chitra, T. (2020). Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies* (pp. 80-91). Association for Computing Machinery.

Aspembitova, A. T., & Bentley, M. A. (2022). Oracles in decentralized Finance: attack costs, profits and mitigation measures. *Entropy*, *25*(1), 60. https://doi.org/10.3390/e25010060

Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., et al. (2017). Distilling the wisdom of crowds: Prediction markets vs. prediction polls. *Management Science*, *63*(3), 691-706.

Atanasov, P., Rescober, P., Stone, E., Swift, S. A., Servan-Schreiber, E., Tetlock, P., Ungar, L., & Mellers, B. (2016). Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls. *Management Science*, *63*(3), 691–706. https://doi.org/10.1287/mnsc.2015.2374

Bartoletti, M., Chiang, J. H., & Lluch-Lafuente, A. (2022). A theory of Automated Market Makers in DeFi. *Logical Methods in Computer Science*, *Volume 18, Issue 4*. https://doi.org/10.46298/lmcs-18(4:12)2022

Bell, T. W. (2011). Government prediction markets: why, who, and how. *Penn State Law Review*, *116*, 403.

Berg, J. E., & Rietz, T. A. (2018). Longshots, overconfidence and efficiency on the Iowa Electronic Market. *International Journal of Forecasting*, *35*(1), 271–287. https://doi.org/10.1016/j.ijforecast.2018.03.004

Berg, J., Forsythe, R., Nelson, F., & Rietz, T. (2008). Chapter 80 Results from a Dozen Years of Election Futures Markets Research. In *Handbook of experimental economics results* (pp. 742–751). https://doi.org/10.1016/s1574-0722(07)00080-7

Chang, C. (2014). Inaccuracy in Health Research News: A typology and predictions of scientists' perceptions of the accuracy of research news. *Journal of Health Communication*, *20*(2), 177–186. https://doi.org/10.1080/10810730.2014.917746

Chen, K. Y., & Plott, C. R. (2002). *Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem.* California Institute of Technology. https://authors.library.caltech.edu/records/n9y0a-a5y79

Collins, B., Hoang, D. T., Nguyen, N. T., & Hwang, D. (2021). Trends in combating fake news on social media–a survey. *Journal of Information and Telecommunication*, *5*(2), 247-266. https://doi.org/10.1080/24751839.2020.1847379

Cookson, J. A., Fox, C., Gil-Bazo, J., Imbet, J. F., & Schiller, C. (2023). *Social media as a bank run catalyst*. Federal Deposit Insurance Corporation (FDIC). https://www.fdic.gov/analysis/cfr/bank-research-conference/annual-22nd/papers/cookson-paper.pdf

Dana, J., Atanasov, P., Tetlock, P., & Mellers, B. (2019). Are markets more accurate than polls? The surprising informational value of "just asking." *Judgment and Decision Making*, *14*(2), 135–147. https://doi.org/10.1017/S1930297500003375

Dosumu, O. E., Sakariyahu, R., Oyekola, O., & Lawal, R. (2023). Panic bank runs, global market contagion and the financial consequences of social media. *Economics Letters*, *228*, 111170. https://doi.org/10.1016/j.econlet.2023.111170

Goncalo, J. A., & Duguid, M. M. (2012). Follow the crowd in a new direction: When conformity pressure facilitates group creativity (and when it does not). *Organizational Behavior and Human Decision Processes*, *118*(1), 14–23. https://doi.org/10.1016/j.obhdp.2011.12.004

Gu, A., Wang, Y., Mascioli, C., Chakraborty, M., Savani, R., Turocy, T. L., & Wellman, M. P. (2024, November). The Effect of Liquidity on the Spoofability of Financial Markets. In *Proceedings of the 5th ACM International Conference on AI in Finance* (pp. 239-247). Association for Computing Machinery.

Halawi, D., Zhang, F., Yueh-Han, C., & Steinhardt, J. (2025). Approaching human-level forecasting with language models. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet et al. (Eds), *Proceedings of the 38th International Conference on Neural Information Processing Systems* (pp. 50426-50468). Curran Associates.

Hanson, R. (2012). LOGARITHMIC MARKETS CORING RULES FOR MODULAR COMBINATORIAL INFORMATION AGGREGATION. *The Journal of Prediction Markets*, *1*(1), 3–15. https://doi.org/10.5750/jpm.v1i1.417

He, Z., & Manela, A. (2014). Information acquisition in Rumor-Based bank runs. *The Journal of Finance*, *71*(3), 1113–1158. https://doi.org/10.1111/jofi.12202

Heidt, C., Sandner, P., & Anders, M. (2024). *Missing Links: Current trends and future potential in the application of blockchain oracles.* https://www.researchgate.net/publication/380419503_Missing_Links_Current_Trends_and_Future_Potential_in_the_Application_of_Blockchain_Oracles

Ho, A. T. Y., Polgreen, P. M., & Prendergast, T. (2016). Prediction Market for disease surveillance: A case study of influenza activity. *The Journal of Prediction Markets*, *10*(1), 68–82. https://doi.org/10.5750/jpm.v10i1.1162

Horst, U., & Naujokat, F. (2011). On derivatives with illiquid underlying and market manipulation. *Quantitative Finance*, *11*(7), 1051–1066. https://doi.org/10.1080/14697688.2011.552517

Hull, J. C., & White, A. D. (2003b). The valuation of credit default swap options. *The Journal of Derivatives*, *10*(3), 40–50. https://doi.org/10.3905/jod.2003.319200

Kaufman, G. G. (1988). The truth about bank runs. In C. England & T. F. Huertas (Eds), *The financial services revolution: Policy directions for the future* (pp. 9-40). Springer Netherlands.

Lee, A. J., & Chung, K. H. (2022). Hidden liquidity, market quality, and order submission strategies. *Journal of Financial Markets*, *61*, 100739. https://doi.org/10.1016/j.finmar.2022.100739

Lin, Q., & Chen, Y. (2009). Gaming dynamic parimutuel markets. In S. Leonardi (Ed.), *Internet and Network Economies: Vol 5929, Lecture Notes in Computer Science*. Springer. https://doi.org/10.1007/978-3-642-10841-9_64

Mishra, A., Mishra, H., & Nayakankuppam, D. (2009). The group-contagion effect: The influence of spatial groupings on perceived contagion and preferences. *Psychological Science, 20*(7), 867–870. https://doi.org/10.1111/j.1467-9280.2009.02371.x

Mohan, V. (2022). Automated market makers and decentralized exchanges: a DeFi primer. *Financial Innovation*, *8*(1). 20. https://doi.org/10.1186/s40854-021-00314-5

Moore, B., & Singletary, M. (1985). Scientific sources' perceptions of network news accuracy. *Deleted Journal*, *62*(4), 816–823. https://doi.org/10.1177/107769908506200415

Morrison, A. D., & White, L. (2013). Reputational contagion and optimal regulatory forbearance. *Journal of Financial Economics*, *110*(3), 642–658. https://doi.org/10.1016/j.jfineco.2013.08.011

Oesterheld, C., Treutlein, J., Cooper, E., & Hudson, R. (2023). Incentivizing honest performative predictions with proper scoring rules. In *Proceedings of the 39th Conference on Uncertainty in Artificial Intelligence: PMLR* (pp. 1564-1574). Association for Uncertainty in Artificial Intelligence.

Ofir, M., & Elmakiess, T. (2025). *Bank runs in the digital era: Technology, psychology, and regulation.* SSRN. https://papers.ssrn.com/sol3/Delivery.cfm?abstractid=5190205

Olan, F., Jayawickrama, U., Arakpogun, E. O., Suklan, J., & Liu, S. (2022). Fake news on Social Media: the Impact on Society. *Information Systems Frontiers*, *26*(2), 443–458. https://doi.org/10.1007/s10796-022-10242-z

Pennock, D. M. (2004). A dynamic pari-mutuel market for hedging, wagering, and information aggregation. *Proceedings of the 5th ACM Conference on Electronic Commerce*. https://courses.cs.duke.edu/spring07/cps296.3/pennock-ec-2004-dynamic-parimutuel.pdf

Pourpouneh, M., Nielsen, K., & Ross, O. (2020). *Automated market makers*. IFRO Working Paper. https://moodle.epfl.ch/pluginfile.php/2865848/mod_resource/content/1/IFRO_WP_2020_08.pdf

Rhode, P. W., & Strumpf, K. S. (2004). Historical presidential betting markets. *The Journal of Economic Perspectives*, *18*(2), 127–142. https://doi.org/10.1257/0895330041371277

Schlegel, J. C., Kwaśnicki, M., & Mamageishvili, A. (2023). Axioms for constant function market makers. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)* (p. 1079). Association for Computing Machinery. https://doi.org/10.1145/3580507.3597720

Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, *6*(1), 103–128. https://doi.org/10.1146/annurev-economics-080213-040927

Shi, P., Conitzer, V., & Guo, M. (2009). Prediction mechanisms that do not incentivize undesirable actions. In S. Leonardi (Ed.), *Internet and Network Economies: Vol 5929, Lecture Notes in Computer Science*. Springer. https://doi.org/10.1007/978-3-642-10841-9_10

Slamka, C., Skiera, B., & Spann, M. (2012). Prediction Market performance and Market Liquidity: A comparison of Automated Market makers. *IEEE Transactions on Engineering Management*, *60*(1), 169–185. https://doi.org/10.1109/tem.2012.2191618

Snowberg, E., Wolfers, J., & Zitzewitz, E. (2013). Prediction markets for economic forecasting. In *Handbook of economic forecasting* (pp. 657–687). https://doi.org/10.1016/b978-0-444-53683-9.00011-6

Soros, G. (2013). Fallibility, reflexivity, and the human uncertainty principle. *Journal of Economic Methodology*, *20*(4), 309–329. https://doi.org/10.1080/1350178x.2013.859415

Soukhoroukova, A., Spann, M., & Skiera, B. (2011). Sourcing, filtering, and Evaluating new product ideas: An Empirical exploration of the performance of idea markets. *Journal of Product Innovation Management*, *29*(1), 100–112. https://doi.org/10.1111/j.1540-5885.2011.00881.x

Spann, M., & Skiera, B. (2003). Internet-Based virtual stock markets for business forecasting. *Management Science*, *49*(10), 1310–1326. https://doi.org/10.1287/mnsc.49.10.1310.17314

Spitler, E. J. (2024). *Yelling "Fire" in the financial theater: bank runs in the social media age and the threat to financial stability*. Carolina Law Scholarship Repository. https://scholarship.law.unc.edu/ncbi/vol28/iss1/5/

Stastny, B. J., & Lehner, P. E. (2018). Comparative evaluation of the forecast accuracy of analysis reports and a prediction market. *Judgment and Decision Making*, *13*(2), 202–211. https://doi.org/10.1017/s1930297500007105

Strickland, B. R., & Crowne, D. P. (1962). Conformity under Conditions of Simulated Group Pressure as a Function of the Need for Social Approval. *The Journal of Social Psychology*, *58*(1), 171–181. https://doi.org/10.1080/00224545.1962.9712366

Sumner, P., Vivian-Griffiths, S., Boivin, J., Williams, A., Venetis, C. A., Davies, A., Ogden, J., Whelan, L., Hughes, B., Dalton, B., Boy, F., & Chambers, C. D. (2014). The association between exaggeration in health related science news and academic press releases: retrospective observational study. *BMJ, 349*, 7015–7015. https://doi.org/10.1136/bmj.g7015

Tetlock, P. C. (2008). Liquidity and prediction market efficiency. *SSRN Electronic Journal*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=929916

Thaler, R. H., & Ziemba, W. T. (1988). Anomalies: Parimutuel betting markets: racetracks and lotteries. *The Journal of Economic Perspectives*, *2*(2), 161–174. https://doi.org/10.1257/jep.2.2.161

Van Bouwel, J. (2014). Towards Democratic Models of Science: Exploring the case of scientific Pluralism. *Perspectives on Science*, *23*(2), 149–172. https://doi.org/10.1162/posc_a_00165

Webel, B. (2013). *Troubled asset relief program (TARP): Implementation and status.* Congressional Research Service. https://elischolar.library.yale.edu/ypfs-documents/8062/

Wolfers, J., & Zitzewitz, E. (2004). Prediction markets. *The Journal of Economic Perspectives*, *18*(2), 107–126. https://doi.org/10.1257/0895330041371321

Zhang, Y., Chen, X., & Park, D. (2018). *Formal specification of constant product (x*y= k) market maker model and implementation*. Champaign, IL: Runtime Verification. https://safefiles.defiyield.info/safe/files/audit/pdf/Uniswap_V1.pdf