

Article

Dynamic Scoring: Probabilistic Model Selection Based on Utility Maximization

Jan Vecer ^{1,2} 

¹ Department of Probability and Mathematical Statistics, Charles University, Sokolovska 83, 18675 Praha, Czech Republic; vecer@karlin.mff.cuni.cz; Tel.: +420-221-913-419

² Vysoka Skola Aplikovaneho Prava, Chomutovicka 1443, 14900 Praha, Czech Republic

Received: 30 November 2018; Accepted: 28 December 2018; Published: 8 January 2019



Abstract: We propose a novel approach of model selection for probability estimates that may be applied in time evolving setting. Specifically, we show that any discrepancy between different probability estimates opens a possibility to compare them by trading on a hypothetical betting market that trades probabilities. We describe the mechanism of such a market, where agents maximize some utility function which determines the optimal trading volume for given odds. This procedure produces supply and demand functions, that determine the size of the bet as a function of a trading probability. These functions are closed form for the choice of logarithmic and exponential utility functions. Having two probability estimates and the corresponding supply and demand functions, the trade matching these estimates happens at the intersection of the supply and demand functions. We show that an agent using correct probabilities will realize a profit in expectation when trading against any other set of probabilities. The expected profit realized by the correct view of the market probabilities can be used as a measure of information in terms of statistical divergence.

Keywords: model selection; utility maximization; explicit supply and demand functions; equilibrium; implied probability; statistical divergence; prediction markets

1. Introduction

Measuring the quality of the fit of a model to the data is a central question in statistics. Standard models for probabilistic estimates based on several explanatory variables include the logit and probit models, where the optimal fit is found using maximum likelihood methods. Adding more explanatory variables gives a better fit, but the tradeoff is that the system may become over parametrized and the added predictive power can be limited for variables with small statistical significance. One approach for model selection relies on the use of the Akaike information criterion [1], which penalizes the addition of every explanatory variable, thus finding an optimal balance by selecting only the variables that improve the quality of the model. An up to date overview of model selection based on information-theoretic approaches in the analysis of empirical data can be found Burnham and Anderson [2] or in Guo et al. [3]. However, analysis of large data sets may be best performed by more complex nonlinear models outside of the traditional framework of the regression analysis. For instance, techniques used in machine learning may not directly yield any measure representing the statistical quality of the model in the sense of the Akaike. In such cases, one can evaluate the quality of the model by using the proper scoring rules, such as the Brier score or the log loss.

The quality assessment of a probabilistic estimate of an outcome is traditionally done by a scoring rule. The proper scoring rule compares the probabilistic estimates with the actual outcomes. Different types of scoring have been extensively studied in the previous literature dating back to the 1950's. Brier [4] introduced a quadratic score function (now called the Brier score) and Good [5] introduced a logarithmic scoring rule. They both represent a proper scoring rule, which gives the highest expected

reward to the true probability distribution. This means that the use of a proper scoring rule encourages the forecaster to be honest and maximize the expected reward. There is extensive subsequent work that focused on various extensions, theoretical properties of the scoring rules and on the mathematical representation of such rules. The main results can be found in papers of Savage [6], Matheson and Winkler [7], Schervish [8], or Gneiting and Raftery [9].

Scoring rules compare one probability estimate with one-time observation. However, it is not clear how to measure the quality of the probabilistic estimate if such an estimate evolves in time. Many practical applications (weather predictions, elections, game probabilities, financial estimates) typically involve time varying estimates. The problem of dynamically evolving probability estimates is that any new information that updates such probabilities comes itself at random times, which prevents a natural time grouping of various estimates required for a quality assessment. In addition, a small initial difference in two models may lead to a long-lasting discrepancy that came from a single source if no new information arrives in the meantime, and thus adding (or equivalently averaging) scoring rules over time would disproportionately amplify such differences that should count only once. This means that a correct approach to measuring the quality of probability time series must adequately account for only newly arriving information and isolate its impact on the current estimate.

Our paper gives an entirely new approach to measure the quality of the fit obtained from several alternative probabilistic models that can be used in any setting. The idea is to compare the performance of two models against each other directly, so it gives a pairwise comparison which makes it distinct from the techniques of the proper scoring rules. This approach generalizes the concept of divergence in statistics (such as Kullback-Leibler [10] later studied by Eguchi and Copas [11]) that compares two probabilistic distributions to the time varying setting. Having two different probability estimates opens a possibility of setting a bet that would make the respective agents optimal in the sense of improving their utility functions. The problem of determination of the optimal size of the bet of an agent with true odds against an agent with biased odds was already studied in Kelly [12], but his study was limited to a static setup with fixed odds, one time observation and a logarithmic utility function. We substantially expand this idea and develop an approach that can be applied to a dynamic setting when the probabilistic estimates are evolving in time.

This technique is relatively simple and straightforward. If we have two alternative probability estimates, the discrepancy can be traded off on a hypothetical betting market. The trading mechanism is based on a utility maximization of the two agents. This mechanism finds both the trading price (probability where the two agents agree to trade) and the corresponding volume of such a bet. We show that an agent using the true probabilities has a non-negative expectation of the profit resulting from trading against any other probability estimate, meaning that the true model cannot be beaten in the long run. The only condition for this result to hold is that the agent with the true probabilities uses an increasing concave utility function. In particular, it does not matter what strategy is used by the second agent. When the two agents agree on a matching algorithm, such as the one based on the equilibrium, the resulting expected profit can be used as a definition of divergence of two probabilistic distributions in the sense of information theory in a dynamic setting. The resulting profit from this trading algorithm can be viewed as information gained from the comparison of two probabilistic models.

Our approach can also be applied in situations when there is no explicit model for the resulting probability estimates in terms of some explanatory variables. This happens for instance in cases when the probabilities are quoted by external subjects, such as in the case of weather predictions where we see only the final probabilities of some events such as rain or snow. In such situations, we have no way to assess the quality of the estimates by Akaike's criterion, but we still can use our approach in terms of the direct comparison of two estimates. Given that we have several models, one cannot statistically identify if the best model corresponds to the true model, which is similar to Akaike's criterion. Interestingly, our approach can statistically conclude that all studied models are incorrect. That follows from the fact that the expected utility of the trading exposures should be an

increasing function of time for the true model and it is possible that this condition is not satisfied for any studied model.

The idea of comparing two probabilities by market trading is novel and only limited related literature exists. Some relevant work has been done in the study of prediction markets that trade probabilities. Wolfers and Zitzewitz [13] give an overview of prediction market types, and their follow up work Wolfers and Zitzewitz [14] determines the size of the optimal position for the agent that maximizes logarithmic utility function for a number of option binary contracts. Arrow et al. [15] argue that prediction markets can help to produce forecasts of event outcomes with a lower prediction error than conventional forecasting methods. Other works focus on checking the quality of probability estimates in different markets using empirical evidence. Ait-Sahalia et al. [16] studied the question whether option markets correctly price the probabilities of movement of the underlying asset. Spann and Skiera [17] compared the forecast accuracy of different methods, namely prediction markets, tipsters and betting odds for soccer games in a static setup. Croxson and Reade [18] studied the efficiency in betting markets after arrival of major news, such as a goal. Forsythe et al. [19] analyzed the Iowa Political Stock Market to ascertain how well markets work as aggregators of information. Rothschild [20] compared the quality of election forecasts from prediction markets and from polls by analysis of the respective probit models. A recent paper of Taleb [21] that studies the evolution of election predictions notes that the resulting probabilistic time series must be a martingale and he argues that the historical time series should have exhibited a fit to a martingale model based on no-arbitrage arguments.

The structure of the paper is as follows. First, we describe the optimal trading behavior of an agent that uses specific probability estimates. This itself is interesting as the maximization of the utility of this agent leads to explicit supply and demand functions for the choice of logarithmic and exponential utility functions. Supply and demand functions are fundamental concepts in economics. However, to the best of our knowledge, there are no known examples of explicit supply and demand functions that are results of the maximization of a utility function. We show that the expected profit for an agent that uses the true probability measure is non-negative regardless of the trading strategy of the rest of the market that match the quotes of the first agent. This result is analogous to a property of a proper scoring rule.

If we have two agents with competing probability models, we also need to describe the matching algorithm that would make the size of the bets acceptable for both of them. The natural choice of the trading price is the point where the supply/demand functions of both agents intersect at an equilibrium point, giving us both the trading price and the trading volume of this match. This itself creates profit-loss statistics from the comparison of the two models, including the statistical significance that measures whether one model is better than a second model. The resulting expected profit can be used as a definition of a statistical divergence in a dynamic setting. We summarize the matching algorithm so that the reader can implement it. The trading on the probability market leads to another novel concept, namely there is a price corresponding to the probability at which the agent is unwilling to trade, meaning that his supply/demand function is equal to zero. Intuitively, the agent should not be willing to change his positions if the quoted price equals his own estimate of the probability, but the agent may have uneven exposures and thus the no-trading point may be shifted. We call this point an implied probability and discuss it in a separate section. We conclude the paper with an illustrative simulated example that compares risk neutral probabilities from the Black-Scholes model [22] with a true volatility 0.25 with models using a low volatility 0.2 and a high volatility 0.3. The true model statistically outperforms both low and high volatility models in less than 100 scenarios. Interestingly, the high volatility model performs slightly better than the low volatility model. The direct comparison of the high volatility model and the low volatility model provides an example of identification of two bad models as neither of the two models has an increasing expected utility function.

2. Results

2.1. Market Model for Trading Probabilities

Consider a situation when we have two possible binary outcomes, (y_1, y_2) with $y_1 + y_2 = 1$. At time t , the probability of these outcomes is given by $(p_1(t), p_2(t))$. As a process, $p_1(t)$ is the conditional expected value of the final outcome:

$$p_1(t) = \mathbb{E}_t[y_1]$$

under the true probability measure \mathbb{P} . In particular, $p_1(t)$ as a process is a martingale under the true probability measure. The true values $(p_1(t), p_2(t))$ are typically unobserved and they are a subject of statistical estimation. Since there can be several competing models for these probabilities, the question is what model is the best among the available alternatives. Let us consider the situation when we have two models that give two probability measures \mathbb{P}^a and \mathbb{P}^b . Thus we have two probability estimates at time t : $p^a(t) = (p_1^a(t), p_2^a(t))$ and $p^b(t) = (p_1^b(t), p_2^b(t))$ of the outcomes (y_1, y_2) . Obviously, $p_2^a(t) = 1 - p_1^a(t)$ and $p_2^b(t) = 1 - p_1^b(t)$.

Let us limit our analysis to the situation of two outcomes, the generalization to an arbitrary number of the number of outcomes is possible, but the formulas involved become rather complicated which we will address it in subsequent work. Given two different estimates, a natural question is which probability series is better in the sense that it is closer to the true probability. We present an approach to compare these estimators based on the following. When there is any discrepancy between the two estimates $p^a(t)$ and $p^b(t)$ at time t , it introduces a possibility of placing a bet. This bet will have two exposures: $N_1(t)$ that corresponds to the win of the first player if the first selection happens and $N_2(t)$ that corresponds to the win of the first player if the second selection happens. These bet sizes will be added to the existing cumulative exposures from the past $(V_1(t), V_2(t))$, resulting in

$$V_1(t + dt) = V_1(t) + N_1(t), \quad V_2(t + dt) = V_2(t) + N_2(t).$$

The market is a zero sum game, and the second agent will collect the negative value of these: $-V_1(t)$ and $-V_2(t)$ respectively. We now discuss how to determine the changes in the exposures $(N_1(t), N_2(t))$ and the cumulative exposures $(V_1(t), V_2(t))$ that would be acceptable for both agents.

An agent can evaluate the current exposures $V(t) = (V_1(t), V_2(t))$ by using some utility function $U(x)$ that is increasing and concave. The standard choices are

Logarithmic Utility: $U(x) = \log\left(1 + \frac{x}{B}\right)$,
Exponential Utility: $U(x) = 1 - \exp\left(-\frac{x}{B}\right)$,
Power Utility: $U(x) = \frac{\left(1 + \frac{x}{B}\right)^{1-a} - 1}{1-a}, a > 0$.

The parameter B is a free parameter and it can be interpreted as a bankroll. In the case of logarithmic utility, the exposures cannot fall below $-B$ as the evaluation of such a position would be negative infinity, thus bounding the resulting exposures. This is not the case for the exponential utility, where exposures below $-B$ are evaluated with a finite value, but such positions are still evaluated highly negatively, meaning that exposures below $-B$ are possible, but unlikely to be realized. Note that logarithmic utility is a limiting case of a power utility when $a \rightarrow 1$. As it turns out, the choice of power utility does not lead to analytical formulas for the below described trading algorithm, and thus we will give formulas only for the logarithmic and exponential utilities.

The agent with probability estimates $p^a(t)$ that correspond to probability measure \mathbb{P}^a assigns an evaluation of his positions $V(t) = (V_1(t), V_2(t))$ as

$$u^a(V_1(t), V_2(t)) = \mathbb{E}^a U(V(t)) = p_1^a(t) \cdot U(V_1(t)) + p_2^a(t) \cdot U(V_2(t)).$$

Similar evaluation of $V(t)$ appears in Kelly [12] and more recently in Wolfers and Zitzewitz [14], but these works are limited to the logarithmic utility function and to static positions. Note that the evaluation of the trading positions $V(t), u^a(V_1(t), V_2(t))$, is based on a subjective probability measure \mathbb{P}^a . For any given odds of the first selection that corresponds to the inverse of the traded probability p , an agent can add or subtract some quantity from the existing exposures which will result in a more favorable utility value. This creates a supply/demand function for each agent, which represents a volume that the agent is willing to bet given the available market odds $\frac{1}{p}$.

More specifically, if the first agent accepts an additional bet of $N(t)$ units on the first selection at price p , his exposures will change to

$$V_1(t + dt) = V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N(t), \quad V_2(t + dt) = V_2(t) + N(t),$$

so the changes in the exposures are given by

$$N_1(t) = -\left(\frac{1}{p} - 1\right) \cdot N(t), \quad N_2(t) = N(t).$$

If the first selection materializes, the first agent is obliged to pay off $\left(\frac{1}{p} - 1\right) \cdot N$ units that correspond to the bet size N at the $\frac{1}{p}$ odds. If the first selection does not materialize, the agent that accepted the bet will collect the original bet size N . Thus for every probability p that represents the trading price, we can find the corresponding N that maximizes the utility function, thus creating a supply/demand function. The reader should note that the parameter p plays a dual role of representing both the trading price and the probability of the first selection. In the following text, parameter p is used exclusively for the trading price as quoted by the market, while parameters p^a and p^b represent subjective probabilities of the respective agents.

The optimal betting size N maximizes

$$p_1^a(t) \cdot U(V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N) + p_2^a(t) \cdot U(V_2(t) + N).$$

This defines the supply function:

Definition 1. A supply function $S(p_1^a, V_1, V_2, p)$ is equal to N that maximizes

$$p_1^a(t) \cdot U(V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N) + p_2^a(t) \cdot U(V_2(t) + N). \tag{1}$$

for a given utility function U .

Given a utility function U , the current exposures $V = (V_1, V_2)$ and the opinion about the probability p_1^a , the agent finds an optimal bet size for any quoted price p . This is how we interpret the supply function.

The following theorem assures that the supply function is well defined:

Theorem 1. The supply function $S(p_1^a, V_1, V_2, p)$ is uniquely defined.

Proof of Theorem 1. This is a simple consequence of concavity of the utility function U and the proof can end here. For more details, consider the case when U has a second derivative: $U''(x) < 0$. Define

$$f(N) = p_1^a(t) \cdot U(V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N) + p_2^a(t) \cdot U(V_2(t) + N).$$

The second derivative of $f(N)$ is

$$f''(N) = \left(\frac{1}{p} - 1\right)^2 p_1^a(t) \cdot U''(V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N) + p_2^a(t) \cdot U''(V_2(t) + N) < 0,$$

so f is also concave and thus attains a unique maximum. \square

The supply function $S(p_1^a, V_1, V_2, p)$ gives the size of the bet N on the second selection. This is directly the change of the size of the second selection N_2 :

$$N_2 = S(p_1^a, V_1, V_2, p).$$

The change in the first selection N_1 is given by

$$N_1 = -\left(\frac{1}{p} - 1\right) \cdot N_2 = -\left(\frac{1}{p} - 1\right) \cdot S(p_1^a, V_1, V_2, p).$$

However, the betting positions can be flipped from (V_1, V_2) with probabilities (p_a^1, p_a^2) to (V_2, V_1) with probabilities (p_a^2, p_a^1) . This also flips N_1 for N_2 , meaning they can be directly determined from the same supply function S , but with flipped positions:

$$N_1 = S(1 - p_1^a, V_2, V_1, 1 - p).$$

Finding the maximum in Equation (1) may not lead in general to a simple closed form formula. In particular, a supply function is not closed form for power utilities even for the simplest choices of the coefficient a . On the other hand, the choice of logarithmic or exponential utility functions leads to convenient formulas. We can take the derivative of Equation (1) and find where it equals zero, which is equivalent to

$$\left(\frac{1}{p} - 1\right)p_1^a(t) \cdot U'(V_1(t) - \left(\frac{1}{p} - 1\right) \cdot N) = p_2^a(t) \cdot U'(V_2(t) + N). \tag{2}$$

For the choice of the logarithmic utility function, we find

$$S^{\log}(p_1^a, V_1, V_2, B, p) = \frac{B(p_1^a - p) + p(p_1^a V_1 - p_1^a V_2 - V_1) + p_1^a V_2}{p - 1}. \tag{3}$$

Similarly, exponential utility gives:

$$S^{\exp}(p_1^a, V_1, V_2, B, p) = -Bp \log\left(\frac{(1-p)p_1^a}{(1-p_1^a)p}\right) + pV_1 - pV_2. \tag{4}$$

We are adding the bankroll B as a parameter since the utility functions depend on it. To the best of our knowledge, there are no previous results in the literature that produce analytical formula for a supply function which is a result of utility maximization.

In the later text, we will specify the market behavior of the second agent, leading to an automated procedure of market matching. However, the following result is true regardless of the matching behavior of the second agent. This theorem is analogous to the concept of a proper scoring rule, which means that the true probability estimate cannot be beaten in expectation when compared with any other alternative model. In addition, it does not matter what utility function is used, the true model is expected to generate a profit regardless of the utility used. Thus the exact choice of the utility function is secondary.

Theorem 2 (True probability has a non-negative expectation). *The exposures $V(t)$ of an agent trading with the true probability measure \mathbb{P} , $p(t) = (p_1(t), p_2(t))$, satisfy the following inequalities:*

$$U^{-1}(\mathbb{E}[U(V(s))]) \leq U^{-1}(\mathbb{E}[U(V(t))]) \leq \mathbb{E}[V(t)]. \tag{5}$$

In particular, $\mathbb{E}[V(t)] \geq 0$ for all $t \geq 0$.

Proof of Theorem 2. This is consequence of a submartingale property of the expected utility and concavity of the utility function. First, note that under the true measure, the process $\mathbb{E}[U(V(t))]$ is a submartingale, meaning

$$\mathbb{E}[U(V(s))] \leq \mathbb{E}_s[\mathbb{E}[U(V(t))]]$$

for $s \leq t$. This is by design. The expected utility moves either by trading actions of the trader or by the evolution of the probabilities. The trading action of the agent always improves the value of the utility with respect to its probability measure, which is true even for agents with other probability estimates. However, in contrast to other probability estimates, the evolution of true probabilities has a neutral impact on the expected utility as the probability p is a martingale under the true probability measure \mathbb{P} . Note that other probability estimates $p^b(t)$ are not martingales under the true measure \mathbb{P} and the evolution of the probabilities does not have a neutral impact on the conditional expected value, meaning other probabilities are not guaranteed to lead to submartingale evolution.

Second, due to the Jensen’s inequality, we have

$$\mathbb{E}[U(V(t))] \leq U(\mathbb{E}[V(t)]).$$

Combining the two properties, we get

$$\mathbb{E}[U(V(s))] \leq \mathbb{E}[U(V(t))] \leq U(\mathbb{E}[V(t)]),$$

and the inequality stated in the theorem is a result of taking the inverse of the utility function U^{-1} , which is increasing. In particular, when $s = 0$, we have $V(0) = (0, 0)$ and we find

$$\mathbb{E}[V(t)] \geq U^{-1}(\mathbb{E}[U(V(0))]) = U^{-1}(U(0)) = 0.$$

This finishes the proof. \square

Note that while the expected utility process, $\mathbb{E}[U(V(t))]$, is a submartingale with an increasing expected value, the same is not true for the expected profit, $\mathbb{E}[V(t)]$. Theorem 2 gives lower bounds for the expected profit, but this does not guarantee that the expected profit is non-decreasing in time. An agent maximizing the utility function may prefer a position with a smaller expectation and a smaller variance over a position with a higher expectation and a higher variance. As an illustrative example, consider an agent with a utility function $U(x) = \log(1 + x)$ that corresponds to the logarithmic utility function with a choice of parameter $B = 1$. Consider the situation when the true probability is $p = \frac{1}{2}$ and the agent uses this probability for trading. The supply function is positive for probabilities above 0.5, and negative for probabilities below 0.5 with the intercept at 0.5. Consider a second agent asking for a trade at $p^t(1) = 0.55$. As we discuss in the later text, this trading price would be chosen if the second agent uses the same utility function, but his probabilistic view is $p^b(1) = 0.6$. The supply function gives $N = 0.1111\dots$, leading to exposures

$$V_1(1) = N_1(0) = -0.0909\dots, \quad V_2(1) = N_2(0) = 0.1111\dots$$

This leads to a favorable position for the first agent with the expected utility

$$\mathbb{E}[U(V(1))] = 0.5 \times \log(1 - 0.0909\dots) + 0.5 \times \log(1 + 0.1111\dots) = 0.005025\dots,$$

and the expected profit

$$\mathbb{E}[V(1)] = 0.5 \times (-0.0909\dots) + 0.5 \times 0.1111 = 0.0101\dots$$

After this trade, the supply function of the first agent shifts down, the supply function is positive for probabilities above 0.55 and negative below 0.55. The first agent accepts additional bets on the

first selection only for prices above 0.55. This trade itself is a good illustration of the entire trading procedure. The true probability $p = \frac{1}{2}$ corresponds to the situation of a coin toss with odds $\frac{1}{p} = 2$. At these odds, the agent is unwilling to trade. However, at odds equal to $\frac{1}{0.55} = 1.8181$, he is willing to accept the bet, creating the exposures $(V_1, V_2) = (-0.0909, 0.1111)$. After this trade, the agent would accept additional bets only above the price 0.55, or equivalently, for odds below 1.8181. Although accepting bet prices in the interval $[0.5, 0.55]$ would further increase his expected profit, it would reduce the expected utility. The first agent would rather reduce his exposure on the first selection at the expense of the expected value. He has already sold the first selection at the higher price of 0.55, so any trade at a buying price below 0.55 would reduce variability of the payoff, which would be reflected in a higher expected utility. Figure 1 shows the supply/demand functions corresponding to the changes of both exposures (N_1, N_2) before and after the trade, illustrating the shift of the offered volume.

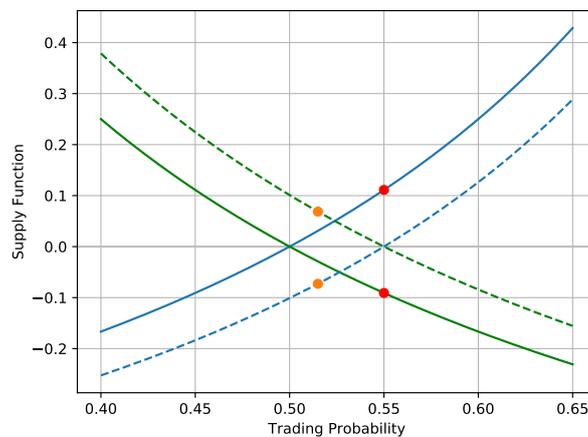


Figure 1. The solid lines represent the supply functions corresponding to the exposure N_2 for both positions (N_1, N_2) in green and in blue respectively as a function of the trading price p . Note that for $p = \frac{1}{2}$, the agent’s optimal bet size is simply $(0, 0)$, the agent is not willing to trade if the odds correspond to the true probability. The first trade happens at odds 0.55 for $(N_1(0), N_2(0)) = (-0.091, 0.111)$ (red dots). This shifts the supply function to the dashed lines. Note that the no trading point $(0, 0)$ shifts to probability 0.55 as a result of a trade on that price, which will be discussed in the section on implied probability. This means that the sign of the trade changes at the interval $[0.5, 0.55]$, resulting in a loss of the expected value. We show the second trade at price 0.515 (orange dots at $(N_1(1), N_2(1)) = (0.069, -0.073)$).

Consider a follow up trade on a price $p^t(2) = 0.515$. The existing exposures update to

$$V_1(2) = V_1(1) + N_1(1) = V_1(1) + 0.06864 \dots = -0.02226 \dots,$$

$$V_2(2) = V_2(1) + N_2(1) = V_2(1) - 0.07289 \dots = 0.03821 \dots,$$

leading to

$$\mathbb{E}[U(V(2))] = 0.5 \times \log(1 - 0.02226 \dots) + 0.5 \times \log(1 + 0.03821 \dots) = 0.00749 \dots,$$

and

$$E[V(2)] = 0.5 \times (-0.02226 \dots) + 0.5 \times 0.03821 \dots = 0.00797 \dots$$

Thus the expected utility has increased, but the expected value has decreased. Theorem 2 gives a lower bound for the expected value. The inverse of the utility function $U(x) = \log(1 + \frac{x}{B})$ is given by

$$U^{-1}(x) = B \cdot (e^x - 1) \geq B \cdot x,$$

which guarantees that the expected value will be in any case at least B times the expected utility value (set $x = U(\mathbb{E}[V(t)])$). When $B = 1$, the expected profit is higher than the expected utility. While the expected utility is monotonically increasing, the expected profit may decrease as illustrated by the above example.

Theorem 2 is a key result for model comparison. The expected profit for the model with true probabilities is positive against any other alternative model. This result does not depend on the choice of the utility function, all that is needed here is that the utility function is increasing and concave. It also does not depend on the potential trading strategy of the second agent. In particular, the above described betting strategy with true probabilities is expected to generate a positive profit against the rest of the market even with multiple actors in the market.

2.2. Market Matching for Two Models

If we have two models to compare, the question is how to set the market matching algorithm for both agents representing the two alternative probabilistic views. While the results of the previous sections are valid regardless of the actions of the remaining market actors, this section focuses on a trading algorithm for two agents, where the trading price and the volume is acceptable for both sides. Let us assume that the second agent is also maximizing a utility function, but since he uses a different set of probabilities, the valuation of his exposures is different. The exposures of the second agent are $(-V_1, -V_2)$ as he is on a negative side of the trade. Thus both agents have a supply/demand function based on how much they are willing to bet, and it is natural to set the bet size N where the two functions intersect at the equilibrium. As discussed in the previous text, the trading behavior of the second agent is irrelevant for the profitability of an agent that uses true probabilities. For sake of symmetry, it makes sense that both agents use the same utility function. The supply function for the second agent must be taken with a negative sign as he is filling a negative position $-V$.

A general approach to a matching algorithm requires one to determine the trading price and the trading volume where the two supply and demand functions for the two agents intersect. The trading price p^t solves

$$S(p_1^a, V_1, V_2, p^t) = -S(p_1^b, -V_1, -V_2, p^t), \tag{6}$$

which is a point on the x -axis where the supply of the first agent with probabilistic view p_1^a and exposures V_1 and V_2 matches the negative supply of the second agent with probabilistic view p_1^b and exposures $-V_1$ and $-V_2$. Similarly, the exact volume N to trade is given by the supply function itself

$$N = S(p_1^a, V_1, V_2, p^t). \tag{7}$$

Since we do not have a supply function for an arbitrary utility function, we will limit ourselves to the situations of the logarithmic and exponential utility functions. It does not matter for the statistical comparison of different probability models as the true model will always outperform any other model regardless of the choice of the utility function. The trading price for the logarithmic utility function is given by

$$p^{t,log} = \frac{Bp_1^a + Bp_1^b + p_1^a V_2 - p_1^b V_2}{2B - p_1^a V_1 + p_1^b V_1 + p_1^a V_2 - p_1^b V_2}, \tag{8}$$

with the corresponding volume

$$N^{log}(p_1^a, p_1^b, V_1, V_2, B) = \frac{B^2(p_1^a - p_1^b) + B(p_1^a(2p_1^b - 1) - p_1^b)(V_1 - V_2) + V_1 V_2(p_1^b - p_1^a)}{B(p_1^a + p_1^b - 2) + V_1(p_1^a - p_1^b)}. \tag{9}$$

Note that the initial values when $V_1 = V_2 = 0$ simplify to

$$p^{t,log} = \frac{p_1^a + p_1^b}{2},$$

meaning that the trading price is at the arithmetic average of the two probability estimates and the corresponding exposures are

$$N_1 = B \cdot \frac{p_1^a - p_1^b}{p_1^a + p_1^b} = B \cdot \frac{\frac{p_1^a}{p_1^b} - 1}{\frac{p_1^a}{p_1^b} + 1}, \quad N_2 = B \cdot \frac{p_2^a - p_2^b}{p_2^a + p_2^b} = B \cdot \frac{\frac{p_2^a}{p_2^b} - 1}{\frac{p_2^a}{p_2^b} + 1}.$$

This corresponds to the model comparison of two static models where we have only one reading, and where one can use the traditional scoring function.

The trading price for the exponential utility function is given by

$$p^{t,exp} = \frac{\sqrt{p_1^a} \sqrt{p_1^b}}{\sqrt{1 - p_1^a} \sqrt{1 - p_1^b} + \sqrt{p_1^a} \sqrt{p_1^b}}. \tag{10}$$

and the volume is

$$N^{exp}(p_1^a, p_1^b, V_1, V_2, B) = \frac{\sqrt{p_1^a} \sqrt{p_1^b} \left(-B \log \left(\frac{\sqrt{p_1^a} \sqrt{1 - p_1^b}}{\sqrt{1 - p_1^a} \sqrt{p_1^b}} \right) + V_1 - V_2 \right)}{\sqrt{1 - p_1^a} \sqrt{1 - p_1^b} + \sqrt{p_1^a} \sqrt{p_1^b}}. \tag{11}$$

In contrast to the logarithmic utility, the trading price for the exponential utility model does not depend on the existing exposures (V_1, V_2) . The trading price is a scaled geometric average of the two probability estimates.

Remark 1 (Relationship to statistical divergence). *The resulting expected profit $\mathbb{E}[V(T)]$ from the above described matching algorithm based on utility maximization can be used as a definition of statistical divergence D between two probability measures \mathbb{P} and \mathbb{Q} as*

$$D(\mathbb{P}||\mathbb{Q}) = \mathbb{E}^{\mathbb{P}}[V(T)] \geq 0,$$

where the \mathbb{P} is the true measure based on Theorem 2. Moreover, the only way that the expected profit $\mathbb{E}[V(T)]$ is exactly zero is when there is no trade, meaning that the two probabilistic measures must be identical $\mathbb{P} = \mathbb{Q}$, which is a condition required for statistical divergence. The important generalization is that this divergence is defined for time evolving probabilities.

For instance, consider two probabilistic measures $\mathbb{P} = \mathbb{P}^a, \mathbb{Q} = \mathbb{P}^b$ and a logarithmic matching algorithm with $B = 1$ in a static case. We have seen that the resulting exposures $(p = p^a, q = p^b)$ are

$$V_1 = N_1 = \frac{p_1 - 1}{\frac{p_1}{q_1} + 1}, \quad V_2 = N_2 = \frac{p_2 - 1}{\frac{p_2}{q_2} + 1},$$

so the expected profit under the true measure \mathbb{P} is given by

$$\begin{aligned} \mathbb{E}^{\mathbb{P}}[V] &= p_1 \cdot \frac{p_1 - 1}{\frac{p_1}{q_1} + 1} + p_2 \cdot \frac{p_2 - 1}{\frac{p_2}{q_2} + 1} \\ &= \frac{p_1}{q_1} \cdot \frac{p_1 - 1}{\frac{p_1}{q_1} + 1} \cdot q_1 + \frac{p_2}{q_2} \cdot \frac{p_2 - 1}{\frac{p_2}{q_2} + 1} \cdot q_2 \\ &= \int_{\Omega} f \left(\frac{d\mathbb{P}}{d\mathbb{Q}} \right) d\mathbb{Q} \end{aligned}$$

for function

$$f(x) = \frac{x^2 - x}{x + 1}.$$

This is a type of *f*-divergence.

The divergence which is a result of maximization of the exponential utility function is not an *f*-divergence. However, the exponential utility plays a role in the Kullback-Leibler divergence. The Kullback-Leibler divergence can be viewed as a bet with resulting exposures

$$V_{1,KL} = N_{1,KL} = \log(p_1) - \log(q_1), \quad V_{2,KL} = N_{2,KL} = \log(p_2) - \log(q_2).$$

After such a trade, both agents would keep zero exponential utility with $B = 1$ as

$$\begin{aligned} \mathbb{E}^P[U(V)] &= p_1 \cdot (1 - \exp(-V_{1,KL})) + p_2 \cdot (1 - \exp(-V_{2,KL})) \\ &= p_1 \cdot (1 - \frac{q_1}{p_1}) + p_2 \cdot (1 - \frac{q_2}{p_2}) = 0, \end{aligned}$$

and

$$\begin{aligned} \mathbb{E}^Q[U(-V)] &= q_1 \cdot (1 - \exp(V_{1,KL})) + q_2 \cdot (1 - \exp(V_{2,KL})) \\ &= q_1 \cdot (1 - \frac{p_1}{q_1}) + q_2 \cdot (1 - \frac{p_2}{q_2}) = 0. \end{aligned}$$

This means that in a static setting, the exposures from the Kullback-Leibler leave both agents ambivalent with a situation of no trade. Since the expected exponential utility function has an approximately parabolic shape, the maximum is achieved close to the midpoints of the two zeroes, namely at $0.5 \times N_{1,KL}$ and $0.5 \times N_{2,KL}$. In particular, the initial exposure N_2 satisfies

$$\frac{\sqrt{p_1} \sqrt{q_1} \left(-\log \left(\frac{\sqrt{p_1} \sqrt{1-q_1}}{\sqrt{1-p_1} \sqrt{q_1}} \right) \right)}{\sqrt{1-p_1} \sqrt{1-q_1} + \sqrt{p_1} \sqrt{q_1}} \approx \frac{1}{2} \times \log \left(\frac{1-p_1}{1-q_1} \right),$$

which is true with a remarkable precision.

One can define a statistical divergence in a dynamic setting by adding the trading positions from the divergence in a static setting. For the case of Kullback-Leibler divergence, we can define

$$\begin{aligned} V_1(t + dt) &= V_1(t) + N_1(t) = V_1(t) + (\log(p_1(t)) - \log(q_1(t))), \\ V_2(t + dt) &= V_2(t) + N_2(t) = V_2(t) + (\log(p_2(t)) - \log(q_2(t))). \end{aligned}$$

This is a natural construction, which moreover guarantees that the process $E[V(t)]$ is non-decreasing under the true measure. This is a simple consequence of a divergence property as $\mathbb{E}^P[N(t)] \geq 0$ at each time t . However, this construction ignores the information flow in time as any long lasting discrepancy is added repeatedly. As a consequence, methods based on utility maximization that properly reflect the information flow in time turn out to be superior as we illustrate in the final section.

Figure 2 illustrates the situation when the two agents trade two probability estimates, one with $p_1^a = 0.5$ and one with $p_1^b = 0.6$ with zero exposures $V_1 = V_2 = 0$ for both situations of the logarithmic and exponential utility functions. Both agents produce the supply /demand functions. In the situation of the logarithmic utility, the supply /demand functions intersect at the trading price $p^{t,log} = 0.55$ with the corresponding volume $N^{log} = 0.1111 \dots$. The exponential utility gives slightly different results, the trading price is at $p^{t,exp} = 0.55051 \dots$ and the resulting volume is $N^{exp} = 0.1116 \dots$

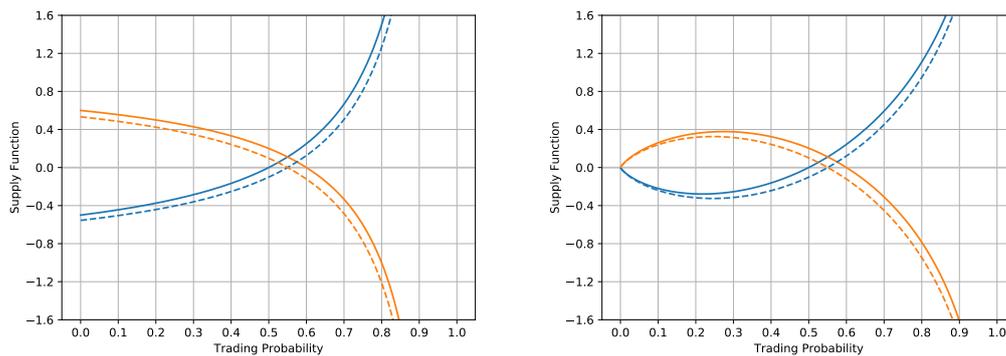


Figure 2. The supply functions corresponding to the exposure N_2 for an agent with $p_1^a = 0.5$ (solid blue lines) and the demand functions for an agent with $p_1^b = 0.6$ (solid orange lines) for log utility (**left**) and exponential utility (**right**) for $B = 1$ as a function of the trading price p . The functions intersect at $p = 0.55$ for the log utility at the level corresponding to $N = 0.1111 \dots$ and at $p = 0.55051 \dots$ for the exponential utility at the level corresponding to $N = 0.1116 \dots$. When the two agents trade off the discrepancy in the probability estimates, their respective supply and demand functions shift to the dashed lines. In particular, the shifted supply and demand functions intersect at the same probability levels as before, but the intersection is now at the level 0, so the discrepancy is fully traded.

For the comparison statistics, all we need to check is whether the profit statistic,

$$V^p(t) = y_1 \cdot V_1(t) + y_2 \cdot V_2(t),$$

is statistically positive. This is the exposure corresponding to the winning side, which is observed only after observing the ultimate outcome in one comparison. If we have n scenarios that are independent and identically distributed, the ratio

$$\frac{\bar{V}^p(t) - \mathbb{E}[V(t)]}{\frac{\sigma(V(t))}{\sqrt{n}}}$$

will converge to a standard normal distribution $N(0, 1)$ according to the Central Limit Theorem. Since the standard deviation of $V(t)$ is typically not observed, we can introduce

$$S^2(t) = \frac{1}{n-1} \sum_{i=1}^n (V_i^p(t) - \bar{V}^p(t))^2,$$

and the test statistic

$$T = \frac{\bar{V}^p(t) - \mathbb{E}[V(t)]}{\frac{S(t)}{\sqrt{n}}}$$

has t -distribution with $n - 1$ degrees of freedom. In particular, we can test the hypothesis that $\mathbb{E}[V(t)] = 0$, giving the statistical significance to the comparison of two alternative models. The absolute magnitude of $\bar{V}^p(t)$ is not important as different utilities may scale this value up or down, but rather the statistical significance depends on its ratio with the corresponding standard deviation.

2.3. Implementation of the Matching Algorithm

Let us summarize the algorithm that matches two probability estimates p^a and p^b needed for the implementation. The procedure is rather simple. The algorithm needs to give the change of the exposures for both $N_1(t)$ and $N_2(t)$, which is given by the function $N(p_1^a, p_1^b, V_1, V_2, B)$. Here, one can choose either the function N obtained from the logarithmic utility or the function obtained from the exponential utility. Since the expected profit of the true probability is non-negative regardless of the

choice of the utility function, it does not matter which one is chosen. We list both of them, but the actual implementation should choose one.

The functions of the matching volume are given by Equation (9) for the logarithmic utility and by Equation (11) for the exponential utility. Next, we need to determine the evolution of the exposures $(V_1(t), V_2(t))$ as time series. This is done by setting the initial values:

$$V_1(0) = 0, \quad V_2(0) = 0$$

and the updated exposures are given by

$$V_1(t + 1) = V_1(t) + N(1 - p^a(t), 1 - p_1^b(t), V_2(t), V_1(t), B), \tag{12}$$

$$V_2(t + 1) = V_2(t) + N(p_1^a(t), p_1^b(t), V_1(t), V_2(t), B). \tag{13}$$

3. Implied Probability

This section applies to an arbitrary number of possible outcomes and thus we present it in the full generality. For a given selection i , current exposures $V = (V_1, \dots, V_n)$, we have a supply function $S(p_i^a, V, p)$ which gives an optimal trading volume for any given price p . There is a price p_i^{imp} for which the supply function gives exactly zero, meaning that the agent is already at the maximal expected utility and thus unwilling to trade any volume N for that price quote. It is natural to call this price an implied probability.

Definition 2. The implied probability p_i^{imp} is a solution of the equation

$$S(p_i^a, V, p^{imp}) = 0. \tag{14}$$

Theorem 3. The implied probability is given by

$$p_i^{imp} = \frac{p_i^a \cdot U'(V_i)}{\sum_{k=1}^n p_k^a \cdot U'(V_k)}. \tag{15}$$

Proof of Theorem 3. We want to find the trading price p_i^{imp} for which the volume $N = 0$ maximizes the expected utility after rebalancing the exposures $V = (V_1, \dots, V_n)$. The i -th exposure obliges to pay out $(\frac{1}{p} - 1) \cdot N$, so the exposure V_i updates to $V_i - (\frac{1}{p} - 1) \cdot N$, all the other exposures collect N , so they update to $V_k + N$. The expected utility of the updated exposures is given by

$$\sum_{k=1, k \neq i}^n p_k^a \cdot U(V_k + N) + p_i^a \cdot U(V_i - (\frac{1}{p} - 1) \cdot N).$$

Solving for the maximum when $N = 0$ gives after differentiation

$$\sum_{k=1}^n p_k^a \cdot U'(V_k) = \frac{p_i^a \cdot U'(V_i)}{p},$$

which gives

$$p_i^{imp} = \frac{p_i^a \cdot U'(V_i)}{\sum_{k=1}^n p_k^a \cdot U'(V_k)}.$$

□

Example 1. The case of the logarithmic utility function gives $U'(x) = \frac{1}{1+\frac{x}{B}} \cdot \frac{1}{B}$, leading to

$$p_i^{imp,log} = \frac{\frac{p_i^a}{1+\frac{V_i}{B}}}{\sum_{k=1}^n \frac{p_k^a}{1+\frac{V_k}{B}}}.$$

When all exposures are zero, we simply retrieve $p_i^{imp} = p_i^a$, meaning that the agent would not engage in any trade when the market quote would agree with his own beliefs. The same is true when all exposures are equal. However, unbalanced exposures would shift the implied probability in an intuitive way. When the market previously bet on exposure i more frequently than on other selections, the agent would end up with a negative exposure $V_i << 0$. When all the other exposures are insignificant or even positive, the fraction $\frac{1}{1+\frac{V_i}{B}}$ would be

greater than 1, thus leading to inequality $p_i^{imp} > p_i^a$. Higher implied probability means that the agent would accept the bet only for prices above p_i^{imp} , meaning that he would quote lower odds. This makes perfect sense and this procedure describes an algorithmic approach how to adapt the odds to the market conditions. Analogously, when the exposure V_i becomes unproportionally positive, the market does not believe in materialization of that selection, leading to smaller implied volatility $p_i^{imp} < p_i^a$.

Exponential utility gives $U'(x) = \frac{1}{B}e^{-\frac{x}{B}}$, leading to implied probability

$$p_i^{imp,exp} = \frac{p_i^a e^{-\frac{V_i}{B}}}{\sum_{k=1}^n p_k^a e^{-\frac{V_k}{B}}}.$$

Again, the unbalanced exposures will shift the implied probability in a similar fashion as in the case of logarithmic utility.

For the sake of completeness, the power utility function has $U'(x) = \frac{1}{B}(1 + \frac{x}{B})^{-a}$, giving

$$p_i^{imp,pow} = \frac{\frac{p_i^a}{(1+\frac{V_i}{B})^a}}{\sum_{k=1}^n \frac{p_k^a}{(1+\frac{V_k}{B})^a}}.$$

The choice of $a = 1$ corresponds to the logarithmic utility, so the formulas indeed agree.

4. Discussion: Application to the Black-Scholes Model with Correct and Misspecified Volatilities

Let us illustrate how we can apply these techniques for model selection in a time evolving setup. For explanatory purposes, we use simulated data where we generate several competing models where one of them has exact probabilities and the others are misspecified. We choose the Black-Scholes model of the stock price driven by a Brownian motion $W(t)$:

$$S(t) = S(0) \exp(\sigma W(t) - \frac{1}{2}\sigma^2 t)$$

and a digital option contract that pays off \$1 if the terminal value of the stock $S(T)$ exceeds a specified strike price K . The price of this digital option is given by

$$\Phi\left(\frac{1}{\sigma\sqrt{T-t}} \cdot \log\left(\frac{S(t)}{K}\right) - \frac{1}{2}\sigma\sqrt{T-t}\right),$$

where t is the current time, T is the terminal time, $S(t)$ is the price of the stock at time t and K is the strike. The price of this digital option also depends on a parameter called the volatility σ .

The function $\Phi(\cdot)$ is the cumulative distribution function of the normal distribution.

Consider the situation when $S(0) = 100$, $K = 100$, $T = 1$ and the true volatility parameter is a constant $\sigma = 0.25$. Let us study two other competing models, one with a higher volatility 0.3 and one with a smaller volatility 0.2. The frequency of trading is set to $dt = \frac{1}{250}$, which in practice corresponds to one day. Figure 3 shows one scenario of a stock price evolution and Figure 4 shows the prices of the digital option for this scenario corresponding to a model with the true volatility 0.25 (in blue) and to a model with high volatility 0.3 (green). Note that both models lead to very similar probability estimates and the difference is barely visible. Moreover, the two models—one with the correct volatility and one with the high volatility—agree on the price of the digital option on a curve plotted in black, so each time the stock price reaches the black curve, the probabilities from both models are the same. Thus the difference between the two models is marginal when the probabilities are near the black curve, which is visible from the graph. The difference gets more pronounced when the probabilities depart from the indicated black curve. The two models trade off any difference on the price level corresponding to the implied probability indicated by the red curve. We use log utility in this example, the results from the exponential utility are nearly identical.

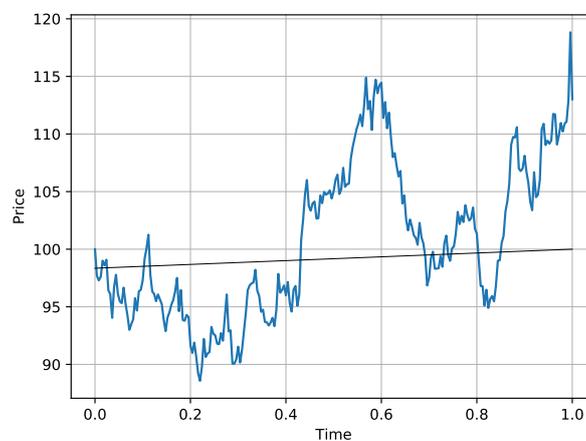


Figure 3. A simulated path of a stock price evolution in the Black-Scholes model with $S(0) = 100$ and the volatility parameter 0.25.

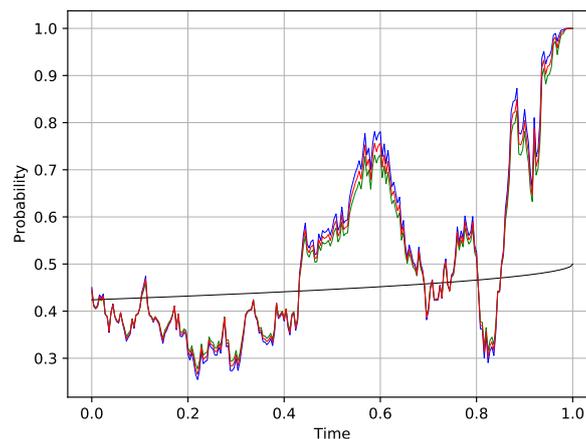


Figure 4. The corresponding prices of a digital option from a model with true volatility parameter (blue), with the high volatility parameter 0.3 (green) and the implied probability path where the two models trade in the equilibrium between the two competing models (red) for the stock evolution scenario from Figure 3. The black curves in both graphs indicate the point where the two models agree on the price of the corresponding digital option, so a close proximity to that curve makes the two models virtually indistinguishable.

For each simulated stock price scenario $S(t)$, we generate the corresponding evolution of the exposures $(V_1(t), V_2(t))$ by using the algorithm described in the implementation section. Theorem 2 assures that the true model will outperform any other model regardless of the choice of the utility function and for illustration, we choose the logarithmic utility function $U(x) = \log(1 + x)$, where we pick the parameter $B = 1$. The results based on the two other alternative approaches studied earlier, namely the matching algorithm based on the exponential utility function and the dynamic version of the Kullback-Leibler divergence, are studied at the end of this section for comparison. In order to have a deeper insight about the evolution of the exposures $V_1(t)$ and $V_2(t)$, let us make the following observation. When the two models agree on the probability, any possible trade must happen on the same price and as a result, the implied probability will agree with the estimate from the two models:

$$p_1^a(t) = p_1^b(t) = p_1^{imp}(t).$$

The implied probability agrees with the original probability only if we have

$$V_1(t) = V_2(t),$$

which is easy to see from the formulas for the implied probability. In our situation, it means that each time the stock price crosses the black line as in Figure 4, the two exposures V_1 and V_2 cross at that time. Note that each such cross is at the trading expense of the true model. The reason is that the true model uses a smaller volatility 0.25 in comparison to the high 0.3 volatility model and thus it sees the probability of the additional cross as smaller than that of a model with a higher volatility. However, in the region between the two crosses, the true model sells to the high volatility model the event that the probability will reverse, i.e., it is closer to 1 when above the black curve and closer to 0 when below the black curve in comparison to the high volatility model. Thus when the additional cross happens, the true model realizes a loss. On the other hand, when there is no additional cross, the true model realizes a profit as being closer the ultimate outcome in comparison to the model with high volatility.

Figure 5 shows the evolution of the exposures V_1 and V_2 for the scenario from Figure 4. Note that the exposures cross exactly at the times when the two models agree on the price of the digital option, or equivalently when the stock price crosses the black curve. Each subsequent crossing happens at a lower level since the true model buys the side which is not yet materialized and results in a loss if the probabilities return to the black curve. The most notable time between two visits of the black curve happens in the time interval $[0.432, 0.692]$. The model with the true volatility is on the buy side of the event that the digital option will end up in the money. The true model increases the exposure V_1 which will materialize if the option ends up in the money at the expense of the exposure V_2 which will materialize otherwise. This view is obviously correct in terms of the probabilities, but it still leaves a possibility of a random return back to the black curve. This reversal indeed happens at time $t = 0.692$, resulting in a trading loss for the true model at that time. The price process makes a few follow up crosses of the black curve with the last cross at time $t = 0.852$. From that point on, the true model is mostly on the buying side of V_1 that will ultimately win, improving the previously lost positions. When the model is sure with near certainty about the winning side, it effectively sells out the losing side, creating a relatively large negative exposure which does not materialize. The ultimate win is $V_1(1) = 0.012019 \dots$ and thus the true model realizes a profit over the model with high volatility in this particular scenario.

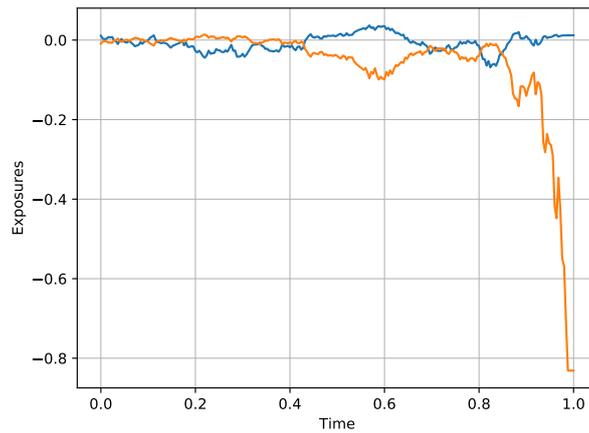


Figure 5. Evolution of V_1 (blue) and V_2 (orange) corresponding to the scenario from Figure 4. Note that the two exposures cross each time when the stock price crosses the black curve.

The above example describes the sources of the profit for both models considered. The true model realizes the profit only after the last visit of the point where both models agree on the price, which is represented by the black curve in Figure 4. The model with high volatility profits from returns to that curve. However, the model with high volatility expects more such visits than what occurs in reality, thus realizing a loss in expectation against the true model. Figure 6 shows a histogram of the realized profit of trading of the true model against the model with high volatility. Since the actual implementation of the model comparison is rather fast, we could have used $n = 1,000,000$ simulations of the price evolutions. The average profit was $\bar{V} = 0.00824 \dots$ with a standard deviation of $\sigma_V = 0.0457 \dots$. Thus the 95% confidence interval for the mean winning is

$$\bar{V} \pm 1.96 \times \frac{\sigma_V}{\sqrt{n}} \approx 0.00824 \pm 0.00009 = [0.00815, 0.00833].$$

These values also indicate how many scenarios must be analyzed before being reasonably sure that one model is better than the other. The one sided 95% confidence interval for the average win \bar{V}_n from n observations is given by

$$[\mu - 1.645 \frac{\sigma}{\sqrt{n}}, +\infty).$$

In order for this interval to be positive valued, the left end point needs to be above zero and the minimal such n solves

$$\mu - 1.645 \frac{\sigma}{\sqrt{n}} = 0.$$

The resulting n is given by

$$n = \frac{1.645^2}{(\frac{\mu}{\sigma})^2},$$

where $\frac{\mu}{\sigma}$ is an information ratio. Our estimated values give $n = 90$, so in this particular situation of comparing the models with true volatility and high volatility, it would take 90 scenarios to have positive \bar{V}_n statistics with a probability of 95%. Given that the probabilities representing the value of the digital option are close to each other to start with for this choice of volatilities, the number of such scenarios seems very reasonable. The closer the models, the more scenarios one would expect to be required and vice versa.

Let us also take a closer look at the exact shape of the distribution in Figure 6. As the losses of the true model come only from the repeated crossing of the black curve, the size of the loss should be proportional to the number of such crosses. Since the price process is Markov, it has no memory and thus the number of such crosses should be approximately geometric if we consider discrete monitoring, meaning that each additional cross happens only with probability $q < 1$. Thus the left tail of the

profit distribution that corresponds to the losses should be approximately geometric (or exponential), which is clearly visible in the graph. On the other hand, the profits are realized when the stock price stays for the most of the remaining time fully above or fully below the black curve, in which case the realized profit reflects the price fluctuations on the way, resulting in a normal distribution. These scenarios represent only a minority of possible situations, but enough to be visible in the right side of the histogram.

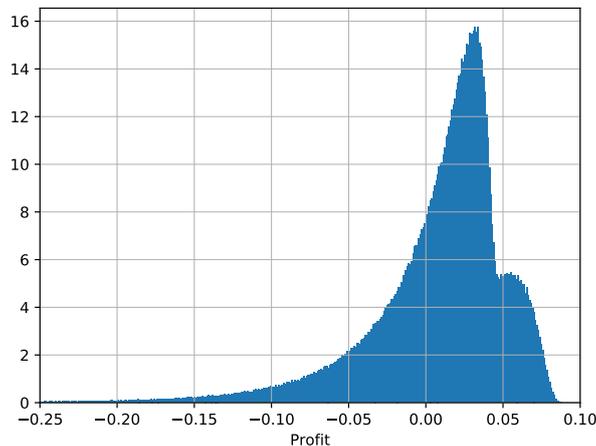


Figure 6. Distribution of the realized profit from comparison of the Black-Scholes model with the true volatility $\sigma = 0.25$ and the high volatility $\sigma = 0.3$, number of simulations $n = 1,000,000$.

After we studied one particular simulation, we can focus on the evolution of the expected profit at time t , $\mathbb{E}[V(t)]$, and the corresponding evolution of the utility function $\mathbb{E}[U(V(t))]$. We consider $n = 1,000,000$ different simulated scenarios, which results in very precise estimates. We have previously shown in Theorem 2 that the process $\mathbb{E}[U(V(t))]$ is a submartingale under the true measure, and thus the expected value should be a non-decreasing function of time. The monotonicity of the process $E[V(t)]$ is not guaranteed, but the specific choice of the utility function guarantees the inequality $E[V(t)] \geq \mathbb{E}[U(V(t))]$.

We study 3 model comparisons, true model versus the high volatility model, true model versus low volatility model and high volatility model versus low volatility model. The left pane of Figure 7 shows the time evolution of $\mathbb{E}[U(V(t))]$ (green) and $\mathbb{E}[V(t)]$ (blue) from the comparison of the true model with the high volatility model. The expected utility is clearly increasing and the expected profit dominates the expected utility, confirming the theoretical result that the true model should realize a profit against any other model in expectation. The right pane of the same figure presents the evaluation of the position $-V(t)$ from the perspective of the model with incorrect probabilities. The red graph is the expected utility

$$\mathbb{E}^H[U(-V(t))] = p_1^H(t) \cdot U(-V_1(t)) + p_2^H(t) \cdot U(-V_2(t))$$

from the perspective of probability weights $p^H = (p_1^H(t), p_2^H(t))$ that correspond to the value of the digital option with high volatility 0.3. The orange graph corresponds to the expected value $\mathbb{E}^H[-V(t)]$. Initially, the agent with the set of probabilities p^H subjectively experiences an increase in both his utility and expectation functions, which is the reason why the agent is willing to trade in the first place, but with increasing time, his position eventually becomes negative even when using his own shifted probabilities. The real profit as measured by the correct probabilities $\mathbb{E}[-V(t)]$ is graphed in blue, and it is the negative of the blue line from the left pane. Both models agree on the final payoff $V(T)$ and thus the blue and the orange curves meet at the end. The evolution of the expected utility under the true measure, $\mathbb{E}[U(-V(t))]$, is graphed in green.

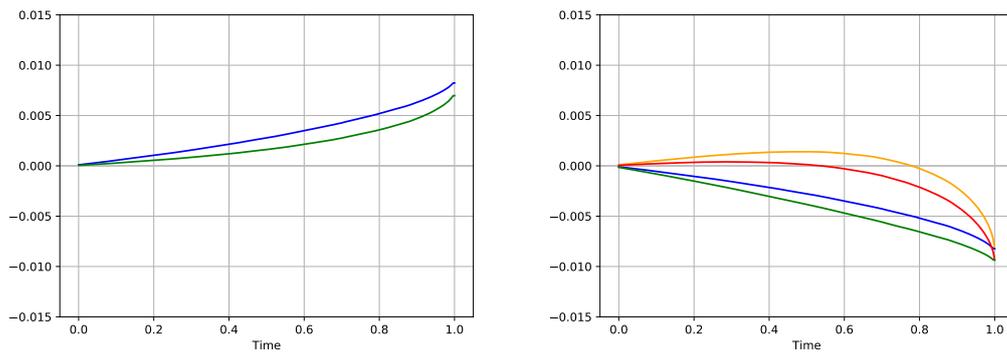


Figure 7. Comparison of the true model and the model with the high volatility. Left: expected profit (blue) and expected utility (green) under true probabilities. Both functions are increasing due to the submartingale property. Right: expected profit (orange) and expected utility of the position $-V(t)$ from the perspective of the probabilities from a model with high volatility. Since this is an incorrect model, we also plot the expected profit under the real measure (blue) and the expected utility under the real measure (green).

Figure 8 shows the same analysis for the comparison of the true model and the model with low volatility. The resulting graphs illustrate the same situation when the true model realizes a profit against the low volatility model. In fact, the true model realizes a slightly higher profit when trading against the low volatility model than the high volatility model. This opens a question whether the high volatility model would perform better in direct comparison with the low volatility model. Figure 9 illustrates this situation. Using the subjective probabilities, both models start to increase their respective expected utility functions. But since both models do not represent the correct probabilities, it is not expected that they would keep increasing for the entire time period. In fact, both models start to experience a decrease of their respective expected utilities after the first half of the trading period, which statistically indicates that both models in question are incorrect. This means that a direct comparison of two incorrect models may lead to a conclusion that neither of them is good.

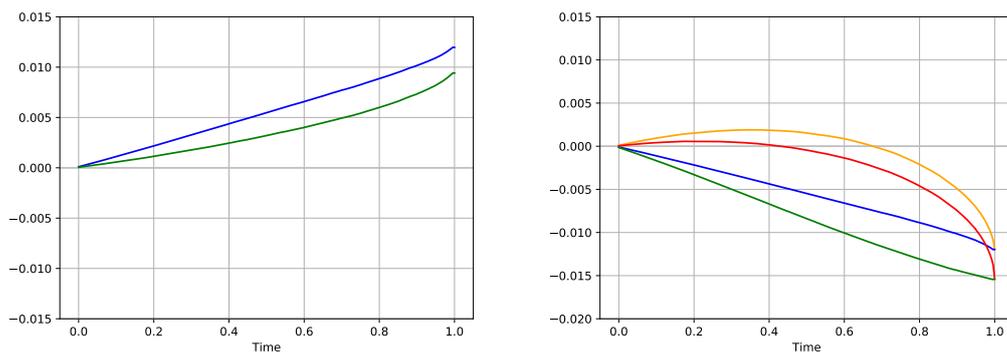


Figure 8. Comparison of the true model and the model with the low volatility analogous to the situation in Figure 7. The left pane shows the expected utility (green) and the expected profit from the perspective of the true model, the right pane shows the evaluations of the position $-V(t)$ both from the perspective of the low volatility model (expected utility in red, expected profit in orange) and from the perspective of the true probabilities (expected utility in green, expected profit in blue).

The interesting situation is to observe the evolution of the expected utility and the expected profit under the objective measure, which is available in our simulated experiment. If these probabilities are not available, we can still estimate the expected profit $\mathbb{E}[V(t)]$ from the observed average profit \bar{V}^p . As seen from the graph, the expected profit of the high volatility model is always above zero

when trading against the low volatility model. However, a substantial part of these gains is given up close to the expiration of the digital contract. Still, the final profit is slightly above zero, indicating that overestimating volatility would lead to better results than underestimating volatility.

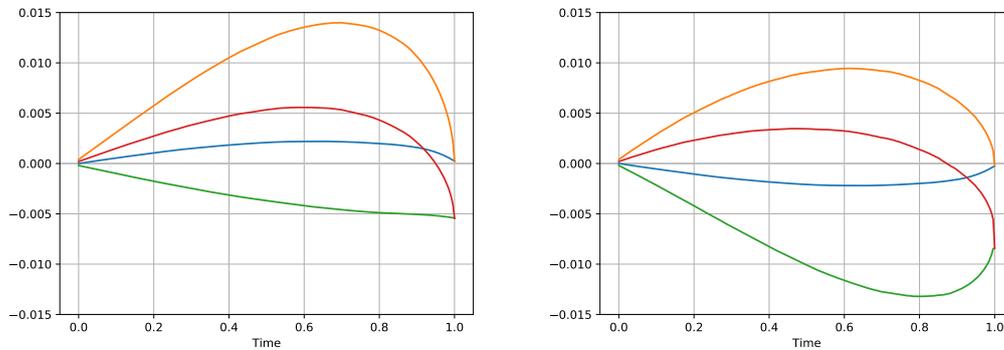


Figure 9. Comparison of two incorrect models, one with high volatility and one with low volatility. Left pane shows the evaluation of the position $V(t)$ from the perspective of the high volatility model (expected utility in red, expected profit in orange) together with the objective valuation of the same position from the perspective of the true model (expected utility in green, expected profit in blue). The right pane shows the evaluation of the position $-V(t)$ both from the perspective of the low volatility model with the analogous graphs.

Table 1 summarizes the profit-loss V statistics together with statistical significance of the results. The true model clearly outperforms both the models based on high or low volatility estimates. The statistical significance is substantial due to the fact of a huge number of tested scenarios. As mentioned earlier, the true model is expected to be significantly better in this situation for about 100 tested scenarios. The high volatility model outperforms the low volatility model, but the statistical significance of this result is small and it took a large number of simulations to show a small advantage of the high volatility model. However, for practical purposes, the two models do not substantially differ.

Table 1. Model Comparison, Different Volatilities.

Models	\bar{V}	s.d.	t Statistics	p -Value
True-High	0.00823793	0.00004757	173.16679	0.00000000
True-Low	0.01195865	0.00007445	160.61887	0.00000000
High-Low	0.00020800	0.00011238	1.850889	0.06418552

The final Table 2 illustrates the performance of the three different approaches presented in this paper—the matching based on the logarithmic utility function covered extensively earlier in this section, the matching based on the exponential utility function, and the Kullback-Leibler divergence—when applied to the comparison of the true versus high volatility model. We list the resulting estimate of $\mathbb{E}[V(T)]$, the estimate $\sigma(V(T))$, the information ratio $\frac{\mu}{\sigma}$ and the smallest number of observations n such that the comparison statistic \bar{V}_n is positive with at least 95% probability. The matching algorithms based on the logarithmic utility function and the exponential utility function are nearly identical, but the logarithmic utility matching turned out to be slightly better both in terms of a higher expectation and a smaller standard deviation. This means that the true model can be identified faster for the logarithmic utility in comparison to the exponential utility. The Kullback-Leibler divergence gives values that correspond to adding the exposures without trading off the past discrepancies, which means that the longer lasting discrepancies are added repeatedly. As a result, both the expectation and the standard deviation are inflated in comparison to the utility matching algorithms. The absolute values are not important, what matters for statistical estimation is the resulting information ratio, which is about one

third of the logarithmic utility matching information ratio. This means that it requires substantially more observations to correctly identify a better model. It would require 608 observations for the Kullback-Leibler divergence in comparison to the 90 observations for the logarithmic utility matching to assure that the event $\bar{V}_n > 0$ has more than 95% probability.

Table 2. True-High Model Comparison, Different Trading Approaches.

Divergence Model	\bar{V}	s.d.	Information Ratio	n: $\mathbb{P}(\bar{V}_n > 0) = 0.95$
Logarithmic	0.00823793	0.047572265	0.17314444	90
Exponential	0.00758190	0.048691108	0.15571435	112
Kullback-Leibler	0.81085788	12.15414317	0.06671452	608

5. Conclusions

This paper presents a novel method of measuring the quality of the probabilistic estimates that can be also applied in a time varying setting. This method is based on trading any discrepancy of the estimates on a hypothetical market, where both agents are maximizing their utility functions. The central result in this paper, Theorem 2, assures that the agent that uses true probabilities will realize a profit in expectation when trading against any other estimates. This result holds regardless of the choice of the utility function. The expected profit resulting from this procedure can be viewed as information gained in terms of statistical divergence.

Funding: This research was funded in part by Grant Agency of the Czech Republic grant number 16-21216S.

Conflicts of Interest: The author declares no conflict of interest.

References

1. Akaike, H. New look at statistical-model identification. *IEEE Trans. Autom. Control* **1974**, *AC19*, 716–723. [[CrossRef](#)]
2. Burnham, K.P.; Anderson, D.R. *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*; Springer Science & Business Media: Berlin/Heidelberg, Germany, 2003.
3. Guo, D.; Shamaï, S.; Verdú, S. The interplay between information and estimation measures. *Found. Trends Signal Process.* **2013**, *6*, 243–429. [[CrossRef](#)]
4. Brier, G.W. Verification of forecasts expressed in terms of probability. *Mon. Weather Rev.* **1950**, *78*, 1–3. [[CrossRef](#)]
5. Good, I.J. Rational decisions. *J. R. Stat. Soc. Ser. B (Methodol.)* **1952**, *14*, 107–114. [[CrossRef](#)]
6. Savage, L. Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* **1971**, *66*, 783–801. [[CrossRef](#)]
7. Matheson, J.E.; Winkler, R.L. Scoring Rules for Continuous Probability Distributions. *Manag. Sci.* **1976**, *22*, 1087–1096. [[CrossRef](#)]
8. Schervish, M. A general method for comparing probability assessors. *Ann. Stat.* **1989**, *17*, 1856–1879. [[CrossRef](#)]
9. Gneiting, T.; Raftery, A.E. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [[CrossRef](#)]
10. Kullback, S.; Leibler, R. On information and sufficiency. *Ann. Math. Stat.* **1951**, *22*, 79–86. [[CrossRef](#)]
11. Eguchi, S.; Copas, J. Interpreting Kullback-Leibler divergence with the Neyman-Pearson lemma. *J. Multivar. Anal.* **2006**, *97*, 2034–2040. [[CrossRef](#)]
12. Kelly, J. A New Interpretation of Information Rate. *Bell Syst. Tech. J.* **1956**, *35*, 917–926. [[CrossRef](#)]
13. Wolfers, J.; Zitzewitz, E. Prediction markets. *J. Econ. Perspect.* **2004**, *18*, 107–126. [[CrossRef](#)]
14. Wolfers, J.; Zitzewitz, E. *Interpreting Prediction Market Prices as Probabilities*; Technical report; National Bureau of Economic Research: Cambridge, MA, USA, 2006.
15. Arrow, K.J.; Forsythe, R.; Gorham, M.; Hahn, R.; Hanson, R.; Ledyard, J.O.; Levmore, S.; Litan, R.; Milgrom, P.; Nelson, F.D.; et al. Economics—The promise of prediction markets. *Science* **2008**, *320*, 877–878. [[CrossRef](#)] [[PubMed](#)]

16. Ait-Sahalia, Y.; Wang, Y.; Yared, F. Do option markets correctly price the probabilities of movement of the underlying asset? *J. Econom.* **2001**, *102*, 67–110. [[CrossRef](#)]
17. Spann, M.; Skiera, B. Sports Forecasting: A Comparison of the Forecast Accuracy of Prediction Markets, Betting Odds and Tipsters. *J. Forecast.* **2009**, *28*, 55–72. [[CrossRef](#)]
18. Croxson, K.; Reade, J.J. Information and efficiency: Goal arrival in soccer betting. *Econ. J.* **2014**, *124*, 62–91. [[CrossRef](#)]
19. Forsythe, R.; Nelson, F.; Neumann, G.; Wright, J. Anatomy of an experimental political stock-market. *Am. Econ. Rev.* **1992**, *82*, 1142–1161.
20. Rothschild, D. Forecasting elections comparing prediction markets, polls, and their biases. *Public Opin. Q.* **2009**, *73*, 895–916. [[CrossRef](#)]
21. Taleb, N.N. Election predictions as martingales: An arbitrage approach. *Quant. Financ.* **2018**, *18*, 1–5. [[CrossRef](#)]
22. Black, F.; Scholes, M. The Pricing of Options and Corporate Liabilities. *J. Polit. Econ.* **1973**, *81*, 637–654. [[CrossRef](#)]



© 2019 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).