

Financial Instruments in Recommendation Mechanisms

Markus Jakobsson

RSA Laboratories,
Bedford, MA 01730
www.markus-jakobsson.com

Abstract. We demonstrate how to use financial instruments to produce recommendation mechanisms. We describe how futures and futures options, both relating to the *perception* of a company or service, can be used to derive accurate recommendations that are secure against abuse. We suggest the notion of *economic reductions* to attribute a cost to the introduction of bias in the recommendation system. We demonstrate the use of such an approach using a simplified set of assumptions on the behavior of the market.

Keywords: futures, recommendations.

1 Introduction

Mary wants to buy cashmere socks, and has found two companies on the web that carries them. However, she has not heard of either company, and would like to know which, if any, she should make her purchase from.

This example, along with the recent expression “*On the Internet, nobody knows that you’re a dog*” [6] capture – in a nutshell – the lack of trust associated with large decentralized networks. Common wisdom has it that this lack of trust is inherent, that is, that it cannot be overcome without imposing a strong structure on the network. Such a structure, however, is almost guaranteed to quench many smaller initiatives in favor of large brand-named on-line institutions, given the likely lack of time and resources to fairly assess all but the largest choices. (In our example, the mom-and-pop cashmere socks store would not be likely to be reviewed by a centralized authority, whereas a large competitor would be likely to.) Interestingly enough, it appears that the same result also is being achieved by the *lack of* such a structure, likely to be caused by the fact that consumers only trust organizations that they have already heard reassuring things about. In other words, the lack of a structure would give large organizations with a name-brand recognition among consumers a definitive edge over smaller organizations, whether these have better products or not. The question we attempt to answer in this paper is how to structure light-weight and accurate mechanisms allowing consumers to assess the value of to them unknown services of varying market penetration, and in a way that is not vulnerable to abuse. In particular, we want to avoid that one organization “cooks” the ratings in its favor (which is a security concern) at the same time as we want changes in service to be quickly

reflected in the ratings (this requirement prevents static recommendations – this can be seen to be a requirement that aggravates the design of secure solutions.) Our proposed solution is based on the rational behavior of investors by extracting a recommendation from trends in investment patterns. While markets are known not to be fully rational at all times, one can see that investors will be incentivised to detect incorrect recommendations and correct these by providing upward or downward pressure on the corresponding commodity. If we revert to our example for a moment, Mary would decide what company to deal with by inspecting stock prices measuring the quality of service of the two companies. (Note that this is different from the normal stocks, which measure the money-making abilities of the companies.) If Mary realizes that the service she obtains (after making a decision) is not consistent with the recommendation she saw, then there is an opportunity for her to make money on knowing the true value of the “goodness stock”. In particular, if the socks are much better than suggested by the recommendation, then it is likely that more people soon will find out, and the stock prices measuring the sock quality will soon go up.

There are several concerns to be addressed. First and foremost, the recommendation mechanism must be *abuse-free* in that it must eventually reflect the impressions of users and buyers, and not allow deceitful bias to be introduced by companies with a stake in the outcome of the recommendation. As such, the system should defend against the effects of companies attempting to downgrade a competitor’s image. A good system should, for the same reason, also guard against companies boosting their own images in ways that does not involve improving their services and products. However, neither of these requirements can be expected to be met in *full*, as suggested by the role advertisement plays in shifting the public perception in a favorable way, but without improving services and products *per se*. Therefore, our aim is to develop mechanisms that achieve a protection *strong enough* that attempts at manipulating the recommendation mechanism are less effective and more costly than other ways of getting improved ratings, such as advertisements and improved service. Analogous to how reductions provide hardness relationships in complexity theory, one can (given the right model of the market) perform *economic reductions* to demonstrate the robustness of a recommendation mechanism, by attributing a cost to the effort of manipulation. Instead of computational hardness assumptions, these would be based on economic assumptions. We draft some reasonable assumptions in order to exemplify the technique and provide a rough analysis of the scheme. However, in order for a commercially meaningful reduction to be performed, much more careful modelling is required.

A second requirement on a recommendation mechanism is that it should reflect consumer opinions in a *timely* manner, that is, the recommendation should be representative of recent performance and perception. It can easily be seen that there is a conflict between the degree to which a mechanism obtains abuse-freeness, and the timeliness of the mechanism. In particular, if a mechanism only takes the last opinion into consideration, it is easily manipulated; on the other hand, if it averages opinions over too long of a time, there is a risk that the

trends are not clearly distinguishable.

Our method draws on the economic incentives of investors to combine a consciousness of trends with a memory of the past, by extracting recommendations from investment statistics. It rewards investors with good foresight and punishes mistaken investors, the latter of which translates manipulative behavior into economic losses. Of course, it is important to remember that profits as well as losses would be restricted to people willing to put their money at stake, while recommendations can be provided to *anybody*.

Finally, a third issue of importance is the cost of maintaining the service. It can be seen that the costs of our mechanism are close to negligible, and its operation light-weight in that it does not require any noticeable maintenance effort. (This is the case since its costs are defrayed by the trading fees.)

Outline. We begin (in section 2) by reviewing a host of traditional solutions used for the purpose of consumer feedback. This description is interleaved with discussions of weaknesses of and requirements on these solutions, were they to be employed in a setting such as the Internet.

We explain the intuition of our solution in section 3. Then, in section 4, we review the structure of some financial instruments underlying our solution. In particular, we will discuss the principles behind stocks, futures, and futures options. A reader familiar with these financial primitives can go directly to the next section (section 5), in which we present our protocols for generating, evaluating and presenting recommendations. Our method relies on the principle that in a free market, the price of each commodity corresponds to the common understanding of its value. By letting commodities track aspects of companies and services that we wish to rank in the recommendation mechanism, we can simply use the ranking of their respective market values for the recommendation. Both individuals and institutional brokers may invest money in such a “market of opinions” where their investments are then translated into recommendations.

Our mechanism is therefore related to the Iowa Electronic Markets [4], but with a recommendation engine placed on top. While the Iowa Electronic Markets failed to project the outcome of the year 2000 elections – one of the markets suggested a republican victory, while the other a democratic – the mechanism still bears promise. Possible problems causing the failed prognosis may be the closeness of the race, and that politics “infected” the game. Another problem may be the limited size of the markets. (Which suggests that the precision of our recommendation mechanism depends on the associated market size.) We refer to [2] for a more thorough discussion of these issues.

In section 6, we discuss possible user interfaces to be built on top of the recommendation mechanisms. We mostly consider the user interfaces for the “common user” (as opposed to the corporate investor.)

In section 7, we analyse the quality of our solution by providing bounds on its accuracy, and study – under a set of working assumptions we establish – the cost of maintaining artificially high (or low) recommendations by manipulating the system.

2 Existing Solutions

We will here discuss existing mechanisms for recommendations, employed in “the real world”, and explain the weaknesses of these mechanisms, were they to be employed in an Internet setting. We will, in particular, discuss the trust requirements and the risk for abuse in these schemes, along with the timeliness and the cost of collecting and maintaining the feedback.

No Recommendation Mechanism. A system not explicitly employing recommendation mechanisms has to rely on advertisements and methods for retaining customers. Due to the lack of trust, it benefits already name-branded players, whether in their “known business” or when entering a new niche.

Personal Recommendations. In a system relying on *personal* recommendations (without any centralized control of these), a user may either choose only to take into consideration recommendations from to him known users (potentially using a friend-of-a-friend chain consisting of a few links); or to consider recommendations from larger sets of people. The former option suffers from a likely shortage of recommendation material (at least relating to merchants that are not household names) and an associated lack of timeliness, while the latter allows abuse. In order to curtail abuse, one can imagine a tiered recommendation structure, where users can post feedback not only about services, but also about users submitting recommendations. However, this mechanism is still exposed to abuse, and may cause “mutual over-rating” (as seen in the structure for rating buyers/sellers in eBay’s system [3].) Finally, it requires mechanisms for compiling large amounts of feedback and extracting the essence of these.

Better Business Bureau. The Better Business Bureau (BBB, [1]) compiles complaints, evaluates (to some extent) the validity of these, and posts warnings when thresholds are reached. One could imagine a service of this type that not only handles negative feedback and warnings, but also positive feedback and suggestions. In either case, though, this type of structure is fraught with the problem of biased feedback. Moreover, the overhead for evaluating the feedback (and its veracity in particular) may be substantial. There is a relationship between the cost of producing a recommendation, and its accuracy and timeliness.

Reviews. A review (such as [5]) is based on surveys or feedback, and have a functional structure similar to the suggestions by a BBB. Not surprisingly, reviews suffer the same shortcomings as structures based on a BBB, but may additionally suffer from problems relating to trust. Namely, users would have to trust not only the veracity of feedback underlying the recommendation, but also the lack of bias introduced by the reviewer, particularly so of it is not clear from where the reviewing organization receives its funding. To some extent, the quality of recommendation mechanism based on reviews depends on the number of independent reviews, and on the quantity of feedback from users. The Zagat [7] restaurant guide is a noteworthy example of a recommendation mechanism that

has gained enough momentum – both among reviewers and users – to gain trust with these. Current review-based recommendation systems typically charge the user for access to recommendations. The cost of producing a recommendation relates closely to its accuracy and timeliness.

3 Intuition

In order to achieve our goals, we will take the novel approach of employing financial instruments to extract recommendations. Thus, recommendations will be based on the current market value of *opinions* about a company or service (as opposed to the company stock in general). The recommendations come with the implicit guarantee that any measureable error in the recommendation (and its timeliness) corresponds to a financial opportunity for anybody who discovers this fact. This will serve to quickly correct recommendations and to keep them as honest as can be.

Our mechanism is secure against “biased buying” by parties interested in thwarting the recommendation outcome. This follows from the fact that the economic power of any company is miniscule in comparison to the economic power of the market place. This is particularly the case for small and medium-sized companies, which are probably also more likely than larger companies to be tempted by such tactics. Therefore, if a large body of investors were to disagree with a rating or recommendation, this would soon be reflected in the market value it correspondings to. (Additionally, standard measures against insider trading would apply.)

A first approach could be to create a stock associated with the perception of each service, or each aspect of each service. Thus, one could imagine stocks tracking the perception of the *quality* of a company’s products, the perception of the *value* of these, and the perception of the *service* provided. However, stocks have the drawback of being less volatile than other financial instruments, which translates into a lower timeliness of a recommendation mechanism built of stocks. Also, the lower profits achievable by stocks may make the recommendation less accurate even in a (hypothetical) stable state of the system. Finally, stocks offer less flexibility than some other instruments – we will see examples of the usefulness of such flexibility onwards in our description.

In order to achieve increased volatility, we propose the use of *futures* and *futures options*. Unlike common futures and futures options, ours would relate not to the *expected price of a commodity* but to the *perceived quality or value of a service*. While either futures or futures options may be used for our mechanism, and both may be employed at the same time, the user interfaces will differ for the two. We will discuss this in detail after having presented the workings of the recommendation mechanism.

The market values of the futures (resp. futures options) indicate the perception of the associated services or companies. The relationship between the prices associated with two competing companies will, similarly, specify a ranking of the companies. We will show that financial arbitrage will automatically cause a

linear ordering of all of the companies being compared. Such an ordering may be performed with respect to each aspect (such as quality, value, service) that corresponds to a future or futures option.

4 Overview of Relevant Financial Instruments

Futures: Long and Short. Traditionally, a *futures contract* is a promise to buy or sell a certain quantity of goods at a given time. To be *long* means to have agreed to *obtain delivery* at the delivery month of the contract, while being *short* means to have agreed to *make delivery* according to the contract. Either way, it is the case that the price of the delivery is agreed upon at the time the position is taken. Futures were introduced as a type of insurance: If, in May, a farmer takes a *short* position for delivery of wheat in September (corresponding to the quantity of wheat he anticipates obtaining at harvest), then he is able to guarantee a profit for his wheat corresponding the contractual price at the time of taking the position. Similarly, a wheat consumer (such as a baker, perhaps) may take a *long* position to avoid that price fluctuations of wheat alter his calculated profits.

Speculation. It is not necessary to either *have* or *want* wheat in order to buy wheat futures. If a trader believes that the wheat price is about to fall, he will go *short*, at which time he promises to deliver at the price specified in the contract. Later, he could either buy the wheat and make delivery, or more commonly, go *long* to cancel out his previous position. If the price fell during this time, the delivery price he is offered at the early point in time is going to be higher than the price he has to pay to *avoid* making delivery at the latter point in time. The trader will thus make a profit. Of course, this goes both ways, and if the price were to go up, then our trader will accrue a corresponding loss. Similarly, an investor who believes that the price of some merchandise will go up would take a long position at first, and later cancel his position by going short with the same quantity.

Arbitrage. Assume that the exchange rate of pounds to dollars is 1.3; that the rate of dollars to marks is 1.2; and that the rate of marks to pounds is 1.1. Clearly, this is unsustainable, since an investor could start with a small amount of pounds; exchange those for dollars, the dollars for marks, and the marks for pounds; after which he would end up with more pounds than he started with. This process, called arbitrage, is what will immediately impose a linear ordering of the currencies by applying increased upward or downward pressure on the value of at least one of them.

Open Interest. For each long position created, one short position is also created. If a person holding one type takes a position of the opposite type, we say that the two positions *cancel*. The *open interest* is a count of the number of non-cancelled positions held for each type of future. As such, the open interest indicates the activity of the market; the trading volume is another such measure.

Earnest money. We mentioned that the payment for the commodity is performed at delivery. However, at the time the position is taken, both sides of the contract deposit *earnest money*, which is typically a fraction of the contractual price. If an investor takes a long position and the prices go up, then he will be able to *withdraw* against his earnest money (since less is needed with the new rate). The same holds for a short position and a falling price. On the other hand, a long investor would under falling prices have to *deposit more* earnest money to keep a security margin. If this margin is ever reached, the clearinghouse would have to limit the number of positions held by the investor, i.e., sell some of the positions at market prices.

Spreads. A *spread* is one long position and one short position – for two different but related types of commodities. The spread is a useful tool for the investor who believes that he knows how the prices of the two types of merchandise will develop *relative to each other* – but without wanting to make bets on how their individual values develop over time. As an example, an investor who takes a position *long Deutsche Mark / short Swiss Franc* believes that the former currency will gain in relation to the latter. As long as this happens, the investor will make a profit, independently of whether they both should go up or both should fall. A spread will typically require less earnest money than a single futures position, as the losses of one side will be balanced to a large extent by the profits of the other.

Futures Options. A *futures option* is a contract that gives the *possibility* of purchasing (resp. selling) a quantity of a commodity at a price specified in the contract. However, it is – unlike normal futures – not *forcing* the buyer of the option to do so. The delivery date of a futures option represents the last point in time when a buyer may exercise the option. The price of the futures option is related to the anticipated price developments of the underlying commodity. Thus, in a bullish market, the price of the *long* futures option is going to be high, and the price of the corresponding *short* position low. (This is similar to how the odds of the favorite race horse will be better than those of a relative newcomer, and the potential payoff the opposite.) It is possible to require the seller of a long (resp. short) futures option to be in possession of the corresponding long (resp. short) future to limit the amount of earnest money demanded by him.

5 Building a Recommendation Mechanism

Delivery. Delivery rarely takes place in a market where a large portion of the investors are speculators. In our setting, this will be even more pronounced, since for “perception futures”, there will be no commodity to be delivered. Therefore, the investors *always* have to cancel their positions at or before the contractual delivery date. It is possible to imagine a futures contract with an *infinite delivery date*. In a futures system with infinite delivery dates, investors would never be forced to cancel out positions, and money would be made by withdrawing against the earnest money, or by voluntarily cancelling out positions. Similarly,

an investor holding a losing position has the choice of depositing more earnest money, or to close out the position and get some portion of the deposited earnest money back. If he does not deposit more earnest money to an investment with continuously falling value, the clearinghouse will cancel out the position before the earnest money is depleted.

Choice of delivery dates. We propose the use of infinite delivery dates to track behavior of a non-seasonal type, such as the hit ratio of browsers, or the services offered by film developing companies. On the other hand, it may be beneficial to retain normal delivery dates for services of seasonal or periodic nature, such as the value of a tourist resort with different seasonal activities offered. For simplicity, we will focus on futures with infinite delivery dates, as these seem to be more useful for smooth tracking and recommendations for the applications we have in mind. On the other hand, we only consider futures options with normal delivery dates, since this simplifies the risk analysis for the seller of the option, and therefore increases trading volume. We note that it is possible to combine the use of futures having infinite delivery dates with futures options having normal delivery dates (although sellers of the options will find it harder to hedge properly.)

Buying individual futures. If a user (or a corporate investor) believes that a certain service is improving, then he will take a long position for the service. Should he be right, then other investors will follow, and the price will increase, giving the investor a profit. (Unless it is common knowledge that the service is improving, in which case it is to some extent already factored into the price.) Similarly, if the investor believes that a service is becoming worse, he will sell it short.

Using multiple perspectives. It is possible to have two or more sets of futures describing one set of company or services, but from different demographic perspectives. One of these perspectives, for example, can be “as perceived by black and latino men between the ages of 20 and 25” while another may be “as perceived by white teenage women.” We note the direct application of the corresponding rankings not only for recommendation systems, but also for purposes of directed advertisements, and research on demographics and consumer behavior.

Using spreads. If an investor believes that company or service A is better than company or service B, he can create a spread position by buying A long and B short. (We note that this fuels the market price of A and cools the market price of B.) Of course, the decision has to be made after studying market values: if A trades at a much higher value than B, then it is commonly known that their service is better. However, an investor who can identify a situation where they are similarly priced, or even, where B is priced higher, would be likely to want to take the above position. If two different futures describes the same company or service, but from different points of view, then it is possible to create a spread

position between these two futures, corresponding to making a bet on what consumer group the service will advance the most onwards.

A note on price movements. As is normal in a free market, downward pressure on the price of a commodity will cause its price to drop. Similarly, upward pressure will result in prices going up. Therefore, the combined effect of investor purchases and sales will move the price of the commodity – in our case the future or futures option – to the level where, according to the market, it belongs.

Translating prices into recommendations. For each set of futures in the same type of market, it will be possible to rank the corresponding companies or services according to the price of the futures, giving the highest ranking to the company or service with the highest market price, etc. Here, the *same type of market* is used to mean when the corresponding services or companies can be *compared*. For example, one can compare the value of products of companies in the same business in a meaningful way, but one cannot compare the value of products between companies in entirely different businesses. Similarly, one cannot compare the value of products offered by one company to the delivery speed of a second company – whether they are in the same business or not. The clearinghouse, or any observer of the market, can create rankings of companies and services that can be compared. Each company or service can be ranked with respect to one or more aspects.

Average Rankings. If there are multiple futures describing one and the same service, but from different perspectives, one can clearly use this for recommendations geared towards the various consumer groups. One may also create *average rankings* by generating a weighted average (where the weights may be selected in proportion to the open interest of the individual futures) of the prices, which then would be translated into a ranking.

Determining the precision of a recommendation. The open interest is an indicator of the number of long and short positions held at the time. A large open interest, combined with a high volume of transactions, is an indicator of a high public interest in the corresponding future. This, in turn, translates to a high degree of precision of the ranking derived from the market prices. On the other hand, a large open interest without any noticeable trading suggests that the price may be about to move, but that the losing side of the trend is not yet convinced of the direction of the movements. If the open interest is very low, and transaction volume limited, then the precision one can obtain is low, as the opinion is based on only a few investors. Finally, if the open interest is low, but the transaction volume is high, then the trend can be seen as an indication of reasonable precision. (Thus, trends can be seen as a tie-breaking aspect used in the ranking, which primarily is based on the market prices.) In all of the above, large vs. small open interest must be seen as a fraction of the market of options related to the option in question; similarly, the trading volume must be seen in the perspective of total trading in the related market.

Interpreting and using futures options. If futures options are used in combination with futures, one can base the recommendation mechanism solely on the market prices of the futures, and allow the futures options merely to be another tool for trading and putting upwards and downwards pressure on the market values. Furthermore, one can use the discrepancy in the short and long prices for futures options to determine the trend, i.e., whether there is upwards or downwards pressure on the price of a future. The probably biggest benefit of futures options in our setting is that they do not require constant monitoring of the earnest money, but rather, once a position is taken, the investor may detach himself from further involvement until he decides to exercise the option (i.e., collect the profit, if any.) This makes them particularly practical for “casual investors”.

The effects of support purchases. If a company attempts to improve its image by means of performing support purchases of its own futures, then this will cause the value of competing futures to rise by means of arbitrage and new spread positions taken by investors who notice the discrepancy in futures prices and quality of service. In order to sustain the improved rating, the company therefore has to keep making support purchases to counter the market forces. We will study the cost of this in section 7. We note that a company can also improve its image by attempting to damage the image of its competitors. However, this results in a larger cost (to get the same relative improvement in the rating) as long as the cumulative market of the competitor’s futures has larger volume than the company’s own futures alone. Therefore, we will focus on the former threat.

6 User Interface

Displaying recommendations. According to the methods described above, one can create rankings of the various aspects of a service or company of interest to a user. These can be used for decision guides by allowing the user to prioritize the various aspects, and weighing the rankings obtained by the user-defined weights. One can also display the available services on a line, their positions set as a function of their traded value, with arrows indicating trends. Additionally, one may add a measure of the precision, using, e.g., a pie chart under each company on the axis. (Here, the precision, as mentioned, can be estimated from the open interest and the trading volume, as well as the current prices of its futures options.)

Creating an account. Everybody who wants to start investing in *perception futures* or *perception futures options* needs to transfer money to a clearinghouse or a broker. A variety of well understood methods can be used for this, as well as for protecting the account against unauthorized access. Once an account has been created and a minimum balance established, the user can trade in available options and futures.




Fig. 1. Example interface.

Displaying investment information. Casual investors may use an interface as that depicted in figure 1. The investor may request more details, and can assess the approximate and relative market values of futures by looking at the related recommendation display. By selecting one or more companies, he can then obtain the prices for futures and futures options for each one of these, along with the prices associated with creating spreads for the various combinations. Delivery dates – when applicable – would be displayed as well. The user would also be given access to information on his existing positions, and be given the possibility of closing out or strengthening these. The more thorough investor would also be given information that allows him to assess the open interest, trading volume, and prices as functions of time.

7 Analysis

The following analysis based on a set of simplified assumptions meant to exemplify how to perform economical reductions. We do not attempt to derive the correct models (and we are aware of the likely complexity of these), and therefore advise the reader to take the conclusions of the analysis with a big grain of salt and merely think of these as indicative of the economic reductions possible, once market data is available. Having said that, let us now consider how an economic reduction can be made, given a model of the market behavior:

Let us fix a company or service to be analysed. We let x correspond to the number of clients it has per time unit, and y the total number of clients per time unit in its niche. We will let δ be the average perceived over-valuation of its future by consumers and investors comparing it to other companies. We will finally let ϵ be the actual amount of over-valuation of its future caused by investment of the company itself. Both δ and ϵ are assumed to be fractions larger than one. However, a similar argument will hold for depreciation of competitors' futures price.

Assumption 1. We assume that the number of *eyeballs* a future receives per time unit is $E = c_1x + c_2y$, where c_1 and c_2 are constants.

Assumption 2. We assume that the trading volume of a future on average is proportional both to its perceived over-valuation δ and to the number of eyeballs E . In particular, we assume that the volume is $V = \delta E$.

Assumption 3. For ease of analysis, we assume that $\delta = \nu\epsilon$, meaning that the average investor perceives the investment opportunity as a fraction ν of the real cost discrepancy.

Assumption 4. We assume that the market is rational in that the common investor tries to maximize his financial benefit.

Loss per time unit. The cost of improving the company's ranking by support-buying of its futures is the cost of the volume of futures traded by the cheater. Thus, the cheater's financial loss per time unit is $L = \epsilon V = \epsilon\delta E$ (wherein we rely on assumption 2.) Using assumption 3, we can simplify this to $L = \nu\epsilon^2 E$, which, according to assumption 1 is $L = \nu\epsilon^2(c_1x + c_2y)$. Consider now the ratio of the cost L taken per customer x . This is the *average support spending per customer* (in a steady state.) We plot the ratio L/x as a function of the discrepancy ϵ in figure 2, using $(c_1, c_2, \nu) = (1/100, 1/500, 0.9)$ as possible values and market shares (i.e., ratios x/y) between 1/100 and 1/2.




Fig. 2. The cost as a function of the degree of manipulation.

We see that – under our assumptions, and using our system – the losses increase rapidly with the distortion of the futures value. We also see that the smaller the market share, the larger the overhead of cheating. This is very convenient,

as larger companies are more likely to be carefully audited, and therefore less likely to even attempt purchasing their image on the futures market.

We note that the costs must be put in relation to the anticipated profits per time unit stemming from the improved rating. However, one must also consider alternative ways of obtaining this better rating, whether to advertise or to improve the products or services. The rating system is abuse-free if these latter costs are lower than those of support-buying futures.

8 Conclusions

We have presented a novel mechanism for producing recommendations. Our mechanism is light-weight, has resistance against manipulation, and is timely. It will not need any new software to be distributed to the average user (who may use a standard browser), and it will be easy to use.

Acknowledgements

Many thanks to Matt Franklin for inspiration and fruitful discussions, and to the anonymous referees for improving the presentation.

References

1. Better Business Bureau, www.bbb.org
2. D. Conley, "A free market election failure,"
www.salon.com/tech/feature/2000/11/16/election_prediction/index.html
3. eBay, www.eBay.com
4. Iowa Electronic Markets, www.biz.uiowa.edu/iem/index.html
5. Lonely Planet, www.lonelyplanet.com
6. Cartoon in the New Yorker, p. 61, July 5 '93.
7. Zagat reviews, www.zagat.com