Adapting To a Market Shock: Optimal Sequential Market-Making

Sanmay Das and Malik Magdon-Ismail

Rensselaer Polytechnic Institute

Abstract

We study the profit-maximization problem of a monopolistic market-maker. The sequential decision problem is hard because the state space is a function. We demonstrate that the belief state is well approximated by a Gaussian distribution. We prove a key monotonicity property of the Gaussian state update which makes the problem tractable. The algorithm leads to a surprising insight: an optimal monopolist can provide *more* liquidity than perfectly competitive market-makers, because a monopolist is willing to absorb initial losses in order to learn a new valuation rapidly so she can extract higher profits later.

Motivation

- Designing markets to achieve certain goals is gaining renewed importance with the prevalence of many **novel markets**, ranging from prediction markets [10] to markets for e-services [8].
- **Liquidity** can be a problem. New markets tend to be thin, and regular markets can also be illiquid immediately following informational shocks [6].
- Chicken-and-egg problem: people do not want to trade in thin markets, and yet, having many people trading is what creates liquidity. These markets therefore need to be bootstrapped into a phase where they are sufficiently liquid to attract trading. This bootstrapping is often achieved through **market-makers** [9].
- Should market-makers be monopolistic (NYSE model) or competitive (NASDAQ model)?
- It seems that competition would be better (wiping out MM profits), but NYSE's promotional literature used to tout the benefits of a monopolist for "maintaining a fair and orderly market" [4]
- Challenge to formal analysis: complexity of the monopolistic market maker's sequential decision problem. Path dependent exploitation-exploration dilema: set prices to extract maximum profit from the next trade or to get information about the new value of the instrument so as to generate larger profits from future trades?

Overview and Contribution

- We consider **optimal sequential price-setting** in the Glosten-Milgrom [5] model.
- MM's **State**: her belief about the instrument value. **Action**: set bid and ask prices. Must trade off profit taking (exploitation) with price discovery (exploration).
- **Functional state space**: Belief is a probability density function. Even with Gaussian assumptions, beliefs quickly become a complex product of error functions, and exact DP becomes intractable.
- **Continuous action space**: bid and ask prices are continuous control variables.
- **Our solution**: Maintain a Gaussian approximation to the state space
  - Allows for efficient solution of the dynamic program.
  - Extract mean and variance of the true state update, and use as mean and variance of the maintained state.
  - Simulations indicate that the approximation is good
- **Phenomenological implications**:
  - Surprise: in periods of extreme shock, when the market maker has large uncertainty relative to the traders, the monopolist provides greater liquidity than competitive zero-profit market-makers.
  - The monopolist increases liquidity, possibly taking short term losses, in order to learn more quickly, and in doing so offers the better social outcome.
  - Once the monopolist has adapted to the shock, she equilibrates at a higher bid ask spread than the the corresponding zero-profit market maker with the same beliefs.

The Model

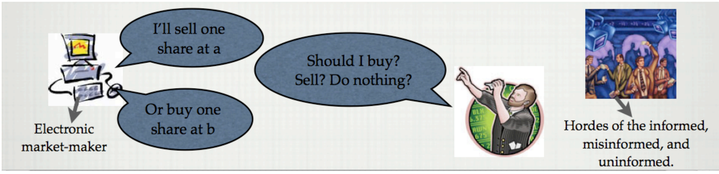

Price-Taking Traders: True value V is drawn from some underlying distribution. An arriving trader receives a signal w drawn from another distribution with mean V (assume it's normal). Trader buys if $w > a$ and sells if $w < b$

Market-Maker's State Space: MM's state is her belief distribution on V. The state space is functional!

Update is Bayesian:

$$p_{t+1}(v) = p_t(v)\frac{q_t(v; b_t, a_t)}{\mathcal{A}_t} \qquad q_t(V; b_t, a_t) = \begin{cases} 1 - F_\epsilon(a_t - V) & \text{Trader bought,} \\ F_\epsilon(a_t - V) - F_\epsilon(b_t - V) & \text{No trade,} \\ F_\epsilon(b_t - V) & \text{Trader sold,} \end{cases}$$

Zero-Profit and Myopic Market-Making: Bid and ask prices are the solutions to fixed point equations that depend only on the current state. So the problem is a one-step decision-making problem combined with figuring out the state update [2].

Optimal Market-Making

$$V(p_t; \pi) = E[r_0 | p_t, b_t^\pi(p_t), a_t^\pi(p_t)] + \gamma E[V(p_{t+1}; \pi) | p_t, b_t^\pi(p_t), a_t^\pi(p_t)]$$

Gaussian approximation to the state: maintain mean and variance. Then the value function is independent of the mean, and only needs to consider the variance.

$$V(\sigma_t) = \max_\delta \{r_t(\sigma_t, \delta) + \gamma E[V(\sigma_{t+1}) | \delta]\}$$

Optimal strategy: select the spread, and place the bid and ask around the current mean belief.

New mean and variance: extract mean and the variance of true (non-Gaussian) one-step-updated distribution.

$$\mu_{t+1} = \mu_t + \sigma_t \cdot \frac{B}{A} \qquad \sigma_{t+1}^2 = \sigma_t^2 \left(1 - \frac{AC + B^2}{A^2}\right)$$

where A, B, and C are various Gaussian integrals that can be easily computed.

The really nice fact: state update is monotonic! $\sigma_{t+1}^2 \le \sigma_t^2$

Now we can solve the Bellman equation efficiently by using linear interpolation on a fine grid of the real line.

The optimal action is then to choose the spread that maximizes the value function.

Results

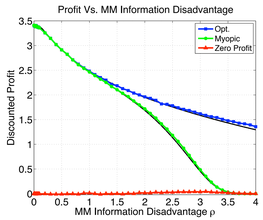

The realized and theoretical value functions match closely, validating the Gaussian framework. Information disadvantage is the ratio of MM's belief variance to the variance of trader beliefs

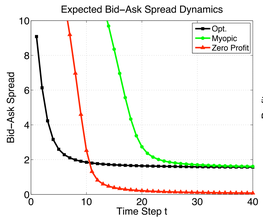

Realized spread over time (σ = 3). The monopolist begins with lowest spread and converges quickest to equilibrium

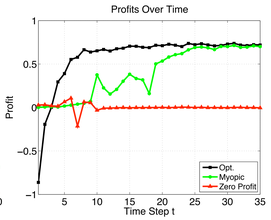

The monopolist is willing to take short-term losses to improve future profits through quicker price discovery

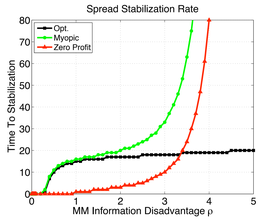

As the information disadvantage increases the relative advantage of the monopolist for price discovery goes up significantly

Conclusion

The MM is willing to take early losses by offering lower spreads in order to make significantly higher profits later. This is quantitative evidence that the optimal MM offers more liquidity than a zero-profit MM after a market shock, especially when the MM is at a large information disadvantage. In this regime, exploration is more important than exploitation. Competition may actually impede the price discovery process, since the market makers would have no incentive to take early losses for better price discovery – competitive pricing is not necessarily informationally efficient (there are quicker ways for the market to "learn" a new valuation).

Our solution is based on reducing a functional state space to a finite-dimensional one in which the Bellman equation can be solved efficiently. When the state is a probability distribution, updated according to independent events, we expect the Gaussian approximation to closely match the real state evolution. Hence, our methods may be generally applicable to problems of this form.

While this paper presents a stylized model, simple trading models have been shown to produce rich market behavior in many cases (Farmer *et al* 2005). The results presented here are an example of the kinds of insights that can be be gained from studying market properties in these models while approaching agent decision problems from the perspective of machine learning. At the same time, this paper is not purely theoretical. The eventual algorithm we present is easy to implement, and we are in the process of evaluating this algorithm in test prediction markets.

References

1. W.G. Christie and P.H. Schulz. Why do NASDAQ market-makers avoid odd-eighth quotes? *J. Finance 1994*.
2. S.Das. A learning Market-Maker in the Glosten Milgrom Model. *Quantitative Finance* 2005.
3. E. Even-Dar, S.M. Kakade, M. Kearns, and Y. Mansour. (In)stability properties of limit order dynamics. *Proc. ACM Conf. on Elect. Comm.* 2006.
4. L.R. Glosten. Insider trading, liquidity, and the role of the monopolist specialist. *J. Business* 1989.
5. L.R. Glosten and P.R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financial Economics*, 1985.
6. S.J. Grossman and M.H. Miller. Liquidity and market structure. *J. Finance* 1988.
7. S.M. Kakade, M. Kearns, Y. Mansour, and L. Ortiz. Competitive algorithms for VWAP and limit-order trading. *Proc. ACM Conf. on Elect. Comm.* 2004.
8. J.S. Lee and B. Szymanski. Auctions as a dynamic pricing mechanism for e-services. In *Service Enterprise Integration*, Kluwer, New York, 2006.
9. D. Pennock and R. Sami. Computational aspects of prediction markets. In *Algorithmic Game Theory*. Cambridge University Press,2007.
10. J.Wolfers and E. Zitzewitz. Prediction markets. *J.Economic Perspectives*, 2004.