Reputation Markets

Xiang Yan
Stanford University
xyan@stanford.edu

Benjamin Van Roy
Stanford University
bvr@stanford.edu

ABSTRACT

A reputation system should incentivize users to obtain and reveal estimates of content quality. It should also aggregate these estimates to establish content reputation in a way that counters strategic manipulation. Mechanisms have been proposed in recent literature that offer financial incentives to induce these desirable outcomes. In this paper, to systematically study what we believe to be fundamental characteristics of these mechanisms, we view them as information markets designed to assess content quality, and refer to them as *reputation markets*. Specifically, we develop a rational expectations equilibrium model to study how incentives created by reputation markets should influence community behavior and the accuracy of assessments. Our analysis suggests that reputation markets offer a number of desirable features:

- As the quality of information improves or the cost of information acquisition decreases, reputation assessments become increasingly robust to manipulation.
- If users can pay to acquire information, errors in reputation assessments do not depend on uncertainty in the manipulator's intent.
- Reputation distortion incurs cost to the manipulator, resulting in cash transfers to other users.
- Pseudonyms do not help a manipulator distort reputations.

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: Electronic Commerce

General Terms

Economics, Theory

Keywords

Online ratings, Strategic manipulation, Reputation market, Information market, Rational expectations equilibrium

1. INTRODUCTION

Overwhelmed by the enormous amount of online content, users increasingly rely on reputation systems to help them find quality content. Given a piece of content, a reputation system typically elicits and aggregates reviews from its past consumers to arrive at a numerical rating. The research community has extensively studied how reviews should be aggregated, under the assumption that truthful reviews are available [1].

In practice, however, users may assess content quality carelessly when reviewing or choose not to review at all. An additional challenge is that agents that maintain vested interest in content reputation may review dishonestly. For instance, a content provider may choose to submit many positive ratings for his low-quality content under different pseudonyms, and benefit from its inflated reputation. To counter such distortions on reputation, Bhattacharjee and Goel [3] proposed a mechanism that aims to rank pieces of content by their ability to generate revenue. Their mechanism shares revenue with users in such a way that inaccurate rankings give rise to arbitrage opportunities for users that detect these inaccuracies.

It is unlikely that any user will be absolutely certain when ratings are inaccurate. After all, ratings are meant to be aggregate views, and though an individual may be well-informed, other users may offer additional information of value. Understanding how effective an incentive mechanism will be when users are uncertain calls for equilibrium analysis. In equilibrium, reputations aggregate information across the community, the impact of manipulation is abated by user activity, and users are incentivized to acquire information when worthwhile.

In this paper, we formulate and analyze equilibrium models to better understand the equilibrium behavior of mechanisms that enable users to profit by correcting inaccurate ratings. We propose to view such mechanisms as information markets that are designed specifically to assess the quality of content. We refer to them as *reputation markets*. Rather than focusing on a specific mechanism, we study a simplified, abstract model of reputation market and identify what we believe to be fundamental characteristics shared by reputation markets.

We consider a setting similar to a financial economic model proposed by Grossman and Stiglitz [5]. Users trade, as pricetakers, contracts with payoff contingent upon content quality, which is measured by the future advertisement revenue per view that the content generates. The market price of contracts then drives content reputation, which estimates quality. To trade optimally, each user may choose to obtain information about content quality at a fixed cost. Users that do so acquire partial observation of content quality, while the rest infer quality from the contract price. A manipulator receives a payoff contingent upon content reputation, and manipulates the price and consequently, reputation, to his advantage. The user community's uncertainty over the manipulator's objective leads to uncertainty over his position. To facilitate analysis, we consider a special case where all random variables involved are Gaussian, users exhibit constant absolute risk aversion, and the manipulator is risk neutral. A closed-form price function is obtained by solving for a rational expectations equilibrium where cost of information acquisition offsets its advantage, and trading activity among users offsets the manipulator's position. We then analyze drivers for the effectiveness of reputation as an estimate of quality, measured by the amount of information about quality contained in price, relative to that contained in observations. Note that this model only involves one piece of content. Although it is possible to extend the model to capture multiple competing pieces of content, we focus on the simple case as a starting point.

Our analysis sheds light on the following benefits of reputation markets:

- As the quality of information improves or the cost of information acquisition decreases, reputation assessments become increasingly robust to manipulation.
- If users can pay to acquire information, errors in reputation assessments do not depend on uncertainty in the manipulator's intent.
- Reputation distortion incurs cost to the manipulator, resulting in cash transfers to other users.
- Pseudonyms do not help a manipulator distort reputations.

This paper is organized as follows. In the next section, we discuss some related work. In Section 3, we present our reputation market formulation and its equilibrium conditions. In Section 4, we establish results about equilibrium behavior. In Section 5, we perform comparative statics analysis and discuss drivers for effectiveness of reputation markets. We comment on potential extensions of our model in a final section.

2. RELATED WORK

Mechanisms that incentivize users to provide honest ratings emerged in recent literature. Preceding [3], Avery, Resnick, and Zeckhauser [2] viewed content ratings as underprovided public good and proposed mechanisms that, in a dynamic setting, direct cash transfer among users to coordinate schedule of ratings and allocation of costs incurred by ratings. Their model assumes honest ratings, however. Miller, Resnick, and Zeckhauser [9] later proposed the peer-prediction method. This mechanism uses proper scoring rules to direct cash transfer among users in a way that makes truthful revelation a Nash Equilibrium. Jurca and Faltings [8] studied a similar setting as [9] and analyzed cases where truthful revelation is the unique Nash Equilibrium and collusion never benefits users.

Since a reputation market is a special case of information markets, it is worth mentioning some prior work on the effectiveness of information markets to aggregate information in the presence of manipulation. Hanson [6] studied an information market setting where higher uncertainty over the objective of a price manipulator incentivizes traders to acquire additional information and consequently improve price informativeness. While [6] assumes that the manipulator's payoff is quadratic in price, in this paper, we assume that it is linear in content reputation, which is inferred from price. We believe that our formulation is well motivated because it establishes an explicit relationship between the manipulator's payoff and content reputation, and captures the setting that users infer reputation from price. Further work by Hanson, Oprea, and Porter [7] provides empirical evidence that traders that are aware of the presence of manipulators adjust their inference of asset payoff and their positions accordingly. The informativeness of the resultant price is then robust to manipulation.

In another related paper, Dellarocas [4] studies a setting where all users obtain the same noisy estimate of content quality from some exogenous source (e.g., an online forum). As this estimate impacts user demands for the content, multiple competing firms with vested interest in the content reputation manipulate this estimate at a cost (e.g., by posting untruthfully on the forum) to maximize revenue. Although user and firm behaviors at equilibrium in [4] are similar to the behaviors of their counterparts in our paper, in our opinion, we study a more general version of the setting in [4] and hence produce different qualitative insights. In particular, while [4] studies how users adjust their interpretation of online ratings in anticipation of manipulators' actions, we study how market incentives should impact users' and manipulators' behaviors.

3. EQUILIBRIUM MODEL

Consider a piece of content of interest. Our reputation market aims to estimate its quality d by aggregating the user community's observations. To fix ideas, let d be the revenue per view that the content generates in a future time. We assume that $d = \tilde{d} + \zeta$, where random variable $\tilde{d}$ represents the observable characteristics of quality, and random variable ζ is noise. $\tilde{d}$ can be observed at a cost $c > 0$, while d and ζ are not observable.

There are two time periods. Our mechanism provides a market where a contract that pays out d in the second period is traded. We assume that there is no limit on borrowing and short-selling, and no interest on borrowing is paid or earned. In the first period, each user chooses whether to observe $\tilde{d}$ and upon observation (or no observation), chooses an amount to invest in the contract. We refer to users that observe $\tilde{d}$ as *informed* and the rest, *uninformed*. We model the user community as a continuum, and denote the fraction that is informed with λ. An informed user invests in θ_I units of the contract and aims to optimize expected utility $\mathrm{E}[u(\theta_I(d-p)-c)|p,\tilde{d}]$ conditioned on contract price p and his observation $\tilde{d}$. An uninformed user invests in θ_U units of the contract and aims to optimize $\mathrm{E}[u(\theta_U(d-p))|p]$ conditioned on p. Here, we assume that all users have the same utility function u.

To infer content quality, an otherwise uninformed user would collect all publicly available information from the market, namely price p, to establish a reputation assessment $\mathrm{E}[d|p]$. Note that the best reputation assessment that one could hope for is $\mathrm{E}[d|\tilde{d}]$. Hence, a reputation market is effective if $\mathrm{E}[d|p]$ equals, or is close to $\mathrm{E}[d|\tilde{d}]$.

We now introduce a manipulator that maintains vested interest in content reputation, and aims to manipulate price and distort $\mathrm{E}[d|p]$. For instance, he may be the content provider or a competitor, and hence would benefit from a high or low content reputation, respectively. We assume that he derives payoff $\psi\,\mathrm{E}[d|p]$, where coefficient ψ is known to him, but not to other users. The manipulator always pays c to observe $\tilde{d}$. He chooses to invest in γ units of the contract and optimizes expected utility $\mathrm{E}[v(\gamma(d-p)+\psi\,\mathrm{E}[d|p]-c)|p,\tilde{d}]$, where v is his utility function.

A rational expectations equilibrium of the market is characterized by fraction of users that are informed $\lambda^* \in [0,1]$, users' investment functions $\theta_I^* : \Re^2 \mapsto \Re$ and $\theta_U^* : \Re \mapsto \Re$, the manipulator's investment function $\gamma^* : \Re^3 \mapsto \Re$, and price function $p^* : \Re^2 \mapsto \Re$ such that for all $\tilde{d}$ and ψ,

$$\begin{aligned}
&\theta_I^*(p^*(\tilde{d},\gamma^*),\tilde{d}) \\
&\quad \in \operatorname*{argmax}_{\theta_I} \mathrm{E}[u(\theta_I(d-p^*(\tilde{d},\gamma^*))-c)|p^*(\tilde{d},\gamma^*),\tilde{d}], \\
&\theta_U^*(p^*(\tilde{d},\gamma^*)) \\
&\quad \in \operatorname*{argmax}_{\theta_U} \mathrm{E}[u(\theta_U(d-p^*(\tilde{d},\gamma^*)))|p^*(\tilde{d},\gamma^*)], \\
&\gamma^*(p^*,\tilde{d},\psi) \\
&\quad \in \operatorname*{argmax}_{\gamma} \mathrm{E}[v(\gamma(d-p^*)+\psi\,\mathrm{E}[d|p^*]-c)|p^*,\tilde{d}], \\
&\mathrm{E}[u(\theta_I^*(d-p^*(\tilde{d},\gamma^*))-c)|p^*(\tilde{d},\gamma^*),\tilde{d}] \\
&\quad = \mathrm{E}[u(\theta_U^*(d-p^*(\tilde{d},\gamma^*)))|p^*(\tilde{d},\gamma^*)], \text{ and} \\
&\lambda^*\theta_I^*(p^*(\tilde{d},\gamma^*),\tilde{d}) + (1-\lambda^*)\theta_U^*(p^*(\tilde{d},\gamma^*)) + \gamma^*(p^*,\tilde{d},\psi) \\
&\quad = 0.
\end{aligned}$$

The first three conditions require that at equilibrium, users and the manipulator each optimize their expected utilities. The fourth condition maintains that uninformed users do not have incentive to acquire additional information and become informed. The fifth condition requires the market to clear.

To facilitate solving for an equilibrium analytically, we study a special case of the general model. In particular, we assume Gaussian distribution of observation and noise. That is, $\tilde{d} \sim \mathcal{N}(0,\sigma_{\tilde{d}}^2)$ and $\zeta \sim \mathcal{N}(0,\sigma_\zeta^2)$. We assume that users exhibit constant absolute risk aversion (CARA) and have utility function $u(w) = -\exp(-aw)$, where $a > 0$. We assume that the manipulator is risk neutral and has utility function $v(w) = w$. Note that our analysis applies to the case where the manipulator has CARA utility as well, although the algebra would be more involved and results more difficult to interpret. We also assume that users share a Gaussian prior over manipulator's incentive. That is, $\psi \sim \mathcal{N}(0,\sigma_\psi^2)$. Note that if random variables $\tilde{d}$, ζ, and ψ have non-zero means, our analysis also carries through and achieves the same results.

4. ANALYSIS

In this section, we establish results about equilibrium behavior.

Proposition 1. *At equilibrium, the ratio between the expected utilities of informed and uninformed users satisfies, for any* p^**,*

$$\frac{\mathrm{E}[u(\theta_I^*(d-p^*)-c)|p^*]}{\mathrm{E}[u(\theta_U^*(d-p^*))|p^*]} = e^{ac}\,\frac{\mathrm{Stdv}[d|\tilde{d}]}{\mathrm{Stdv}[d|p^*]}.$$

When price reflects content quality poorly, or information cost is low, or users have low risk aversion, the right hand side of the equation is small. As utility is negative, this implies large incentive to acquire information.

To assess the effectivess of the reputation market, we use the ratio $\mathrm{Stdv}[d|p^*]/\,\mathrm{Stdv}[d|\tilde{d}]$ to measure distortion. Here, the numerator is the root-mean-squared error of reputation as an estimate of quality. The denominator is the root-mean-squared error of quality conditioned on observation. The ratio is one if p fully reveals $\tilde{d}$. It is larger otherwise. We also denote the correlation coefficient between p and $\tilde{d}$ as ρ, and use it to measure the informativeness of price in absolute terms. The following proposition helps us study these measures in the next section.

Proposition 2. *The ratio and correlation coefficient measures satisfy*

$$\frac{\mathrm{Stdv}[d|p^*]}{\mathrm{Stdv}[d|\tilde{d}]} = e^{ac}, \ \textit{and}$$

$$\rho = \sqrt{1 - \left(\frac{\sigma_\zeta^2}{\sigma_{\tilde{d}}^2}\right)(e^{2ac}-1)}.$$

The next two propositions concern price and the manipulator's payoff.

Proposition 3. *An equilibrium exists where market price is linear in observation* $\tilde{d}$ *and manipulator's incentive coefficient* ψ*. In particular,*

$$p^* = \frac{1}{2}\tilde{d} + \eta\psi, \ \textit{where}$$

$$\eta = -\frac{ae^{2ac}\sigma_\zeta^2(\sigma_{\tilde{d}}^2 - (e^{2ac}-1)\sigma_\zeta^2)}{(1-\lambda^*)(\sigma_{\tilde{d}}^2 - 2(e^{2ac}-1)\sigma_\zeta^2)}.$$

The manipulator's equilibrium position γ^* *is linear in* $\tilde{d}$ *and* ψ *as well:*

$$\gamma^*(p,\tilde{d},\psi) = \frac{\alpha}{2}\tilde{d} + \left(\frac{a\sigma_\zeta^2\alpha}{2\lambda^*}+1\right)\beta\psi,$$

where α *and* β *are constants.*

Proposition 4. *The manipulator's expected payoff in the market is*

$$\mathrm{E}[\gamma^*(d^*-p^*)] = \frac{\alpha}{4}\sigma_{\tilde{d}}^2 + \frac{1}{4\alpha}\left(\beta^2 - \frac{\sigma_{\tilde{d}}^2}{4\mathrm{Var}[p^*]^2}\right)\sigma_\psi^2,$$

where α *and* β *are constants. The expected payoff is negative when* $\rho > 1/\sqrt{2}$.

Note that the payoff in Proposition 4 does not include the manipulator's gain $\psi\,\mathrm{E}[d|p]$, which is external to the market.

5. RESULTS

Equipped with results in the previous section, we analyze drivers for the effectiveness of the reputation market. From Proposition 2, we observe that

- Decrease in information cost results in lower reputation distortion and higher price informativeness. Indeed, distortion vanishes as information cost approaches zero.

- Improvement in the quality of information (i.e., increase in $\sigma_{\tilde{d}}/\sigma_\zeta$) increases price informativeness. It does not impact reputation distortion, however. To see why, examine the right hand side of the equality in Proposition 1. As $\sigma_{\tilde{d}}/\sigma_\zeta$ increases, both Stdv$[d|\tilde{d}]$ and Stdv$[d|p^*]$ decrease. Stdv$[d|p^*]$ cannot drop "faster" than Stdv$[d|\tilde{d}]$, however, because it would result in the left hand side ratio greater than one, which would imply that informed users derive less utility than uninformed users at equilibrium. And this cannot happen.

- If users can pay to acquire information, price informativeness and reputation distortion do not depend on uncertainty over the manipulator's incentive ψ. To appreciate the significance of this result, note that intuitively, high uncertainty over ψ leads to high uncertainty over the manipulator's position γ. Users should then be less sure whether, say, a high market price is due to observations that suggest high content quality, or due to a large position by the manipulator. Consequently, one might think that high σ_ψ leads to high reputation distortion. Our analysis shows, however, that uncertainty in ψ and γ introduces incentive to acquire information, and the equilibrium fraction of informed users λ is adjusted due to this incentive. The net result is constant price informativeness and constant reputation distortion.

- Decrease in risk aversion leads users to take larger positions, improves price informativeness, and lowers reputation distortion.

From proposition 4, we observe that when price is sufficiently informative, strategic trade on average incurs cost to the manipulator, resulting in cash transfer to the user community. This property is aligned with our objective to discourage reputation manipulations.

Our reputation market also addresses the "cheap pseudonym problem" experienced by other systems, where manipulators behave strategically under multiple pseudonyms to circumvent per-account restrictions (e.g., maximum number of reviews per user). In our setting, trades executed by a single agent operating multiple accounts can be equivalently executed through a single account. The number of accounts does not change market price, and hence does not benefit a manipulator.

It is also worth noting that if there is no manipulator (i.e., $\gamma = 0$), or the manipulator's incentive is perfectly observed by users (i.e., $\sigma_\psi = 0$), or content quality observation is perfect (i.e., $\sigma_\zeta = 0$), no equilibrium as we defined exists. To see why, consider an equilibrium where a positive fraction of users is informed. The price function fully reveals observations and gives uninformed users, for free, the same information obtained by the informed ones at a cost c. The informed users would feel that they have over-paid. The outcome where nobody is informed is not an equilibrium, either, because taking price as given, a user could then benefit from acquiring information and investing accordingly.

6. EXTENSIONS

There are multiple ways to generalize our model in order to study the effectiveness of reputation markets in more realistic settings. Natural extensions include cases with multiple competing pieces of content and multiple manipulators. Other realistic restrictions include limits on borrowing and short-selling in the market. It would also be interesting to build and study a real Internet service that operates a reputation market to rate content.

Acknowledgments

The authors thank Ashish Goel and Rajat Bhattacharjee for helpful discussions which inspired this paper, and thank Christina Aperjis and Beomsoo Park for their reviews. This research was supported in part by the National Science Foundation through grant IIS-0428868.

7. REFERENCES

[1] G. Adomavicius and A. Tuzhilin. Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6): 734-749, 2005.

[2] C. Avery, P. Resnick, and R. Zeckhauser. The Market for Evaluations. *American Economic Review*, 89(3):564-584, 1999.

[3] R. Bhattacharjee and A. Goel. Algorithms and incentives for robust ranking. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2007.

[4] C. Dellarocas. Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms. *Management Science*, 52(10): 1577-1593, 2006.

[5] S. Grossman and J. Stiglitz. On the Impossibility of Informationally Efficient Markets. *The American Economic Review*, 70(3): 393-408, 1980.

[6] R. Hanson and R. Oprea. Manipulators increase information market accuracy. Working paper. http://hanson.gmu.edu/biashelp.pdf.

[7] R. Hanson, R. Oprea, and D. Porter. Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior and Organization*, 60(4): 449-459.

[8] R. Jurca and B. Faltings. Collusion-resistant, incentive-compatible feedback payments. In *Proceedings of the Seventh ACM Conference on Electronic Commerce*. ACM, 2007.

[9] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51(9): 1359-1373, 2005.

APPENDIX

A. PROOFS OF PROPOSITIONS

Proof of Proposition 1. Given CARA utility function u, we verify that given $\tilde{d}$ and γ,

$$\mathrm{E}[u(\theta_U^*(d-p))|p] = -\exp\left(-\frac{(\mathrm{E}[d|p]-p)^2}{2\,\mathrm{Var}[d|p]}\right), \text{ and}$$

$$\begin{aligned} & \mathrm{E}[u(\theta_I^*(d-p)-c)|\tilde{d},\gamma] \\ = & -\exp\left(ac - \frac{(\mathrm{E}[d|\tilde{d},\gamma]-p)^2}{2\,\mathrm{Var}[d|\tilde{d},\gamma]}\right) \\ = & -\exp\left(ac - \frac{(\tilde{d}-p)^2}{2\sigma_\zeta^2}\right). \end{aligned}$$

Noting that p is determined by d and γ, we have, for any p,

$$\begin{aligned} & \mathrm{E}[u(\theta_I^*(d-p)-c)|p] \\ = & \mathrm{E}[\mathrm{E}[u(\theta_I^*(d-p)-c)|\tilde{d},\gamma]|p] \\ = & \mathrm{E}[-\exp\left(ac - \frac{(\tilde{d}-p)^2}{2\sigma_\zeta^2}\right)|p] \\ = & -\frac{\exp(ac)}{\sqrt{1+\frac{\mathrm{Var}[\tilde{d}|p]}{\sigma_\zeta^2}}}\exp\left(-\left(\frac{\frac{\mathrm{Var}[\tilde{d}|p]}{2\sigma_\zeta^2}}{1+\frac{\mathrm{Var}[\tilde{d}|p]}{\sigma_\zeta^2}}\right)\frac{(\mathrm{E}[\tilde{d}|p]-p)^2}{\mathrm{Var}[\tilde{d}|p]}\right) \\ = & \exp(ac)\sqrt{\frac{\sigma_\zeta^2}{\sigma_\zeta^2+\mathrm{Var}[\tilde{d}|p]}}\left(-\exp\left(\frac{-(\mathrm{E}[\tilde{d}|p]-p)^2}{2\,\mathrm{Var}[d|p]}\right)\right) \\ = & \exp(ac)\sqrt{\frac{\mathrm{Var}[d|\tilde{d}]}{\mathrm{Var}[d|p]}}\,\mathrm{E}[u(\theta_U^*(d-p))|p]. \end{aligned}$$

Hence, for all p^*,

$$\frac{\mathrm{E}[u(\theta_I^*(d-p^*)-c)|p^*]}{\mathrm{E}[u(\theta_U^*(d-p^*))|p^*]} = e^{ac}\frac{\mathrm{Stdv}[d|\tilde{d}]}{\mathrm{Stdv}[d|p^*]}.$$

For the third equality, we used the fact that

$$\mathrm{E}[\exp(-bx^2)] = \frac{1}{\sqrt{1+2b}}\exp\left(-\frac{b}{1+2b}\,\mathrm{E}[x]^2\right)$$

for random variable $x \sim \mathcal{N}(0,1)$ and constant b. For the fifth equality, we used the fact that $\mathrm{Var}[d|\tilde{d}] = \sigma_\zeta^2$ and $\mathrm{Var}[d|p] = \sigma_\zeta^2 + \mathrm{Var}[\tilde{d}|p]$. □

Proof of Proposition 2. Equilibrium condition requires the ratio of utilities between informed and uninformed users to be one. By Proposition 1, then, $e^{ac}\,\mathrm{Stdv}[d|\tilde{d}]/\,\mathrm{Stdv}[d|p^*] = 1$, which yields the first equation in Proposition 2. To prove the second equation, we verify that

$$\begin{aligned} & \left(\frac{\sigma_\zeta^2}{\sigma_{\tilde{d}}^2}\right)(e^{2ac}-1) \\ = & \left(\frac{\sigma_\zeta^2}{\sigma_{\tilde{d}}^2}\right)\left(\frac{\mathrm{Var}[d|p^*]}{\mathrm{Var}[d|\tilde{d}]}-1\right) \\ = & \left(\frac{\sigma_\zeta^2}{\sigma_{\tilde{d}}^2}\right)\left(\frac{\sigma_\zeta^2+\sigma_{\tilde{d}}^2-\frac{\sigma_{\tilde{d}}^4}{4\,\mathrm{Var}[p^*]}}{\sigma_\zeta^2}-1\right) \\ = & 1-\frac{\sigma_{\tilde{d}}^2}{4\,\mathrm{Var}[p^*]} \\ = & 1-\rho^2. \end{aligned}$$

□

Proof of Proposition 3. We start by noting that Proposition 1 implies that $\mathrm{Var}[d|p]$ and $\mathrm{Var}[p]$ are constants. In particular,

$$\mathrm{Var}[d|p] = e^{2ac}\sigma_\zeta^2, \text{ and}$$

$$\mathrm{Var}[p] = \frac{\mathrm{Cov}(d,p)^2}{\mathrm{Var}[d]-\mathrm{Var}[d|p]} = \frac{\sigma_{\tilde{d}}^4}{4(\sigma_{\tilde{d}}^2-(e^{2ac}-1)\sigma_\zeta^2)}.$$

Given a fixed $\lambda \in [0,1]$, we define constants

$$\alpha = -\left(\frac{\lambda}{a\sigma_\zeta^2}+\frac{(1-\lambda)(\sigma_{\tilde{d}}^2-2\,\mathrm{Var}[p])}{2a\,\mathrm{Var}[d|p]\,\mathrm{Var}[p]}\right), \text{ and}$$

$$\beta = -\frac{\lambda\sigma_{\tilde{d}}^2\,\mathrm{Var}[d|p]}{(1-\lambda)\sigma_\zeta^2(\sigma_{\tilde{d}}^2-2\,\mathrm{Var}[p])}.$$

It can be verified that

$$p(\tilde{d},\gamma) = \frac{\lambda}{2\lambda+a\sigma_\zeta^2\alpha}\tilde{d}+\frac{a\sigma_\zeta^2}{2\lambda+a\sigma_\zeta^2\alpha}\gamma,$$

$$\theta_I(p,\tilde{d}) = \frac{\tilde{d}-p}{a\sigma_\zeta^2},$$

$$\theta_U(p) = \frac{\mathrm{E}[d|p]-p}{a\,\mathrm{Var}[d|p]}, \text{ and}$$

$$\gamma(p,\tilde{d},\psi) = \frac{\alpha}{2}\tilde{d}+\left(\frac{a\sigma_\zeta^2\alpha}{2\lambda}+1\right)\beta\psi$$

satisfy the first four equilibrium conditions. Substituting these expressions into the market clearing condition, we could then solve for λ. Substituting γ into p yields the posited expression of p in Proposition 3. □

Proof of Proposition 4. The equality can be verified by substituting expressions for p^* and γ^* in the proof of Proposition 3 into the left hand side. By Proposition 2, $\rho > 1/\sqrt{2}$ is equivalent to $\sigma_{\tilde{d}} > \sqrt{2(e^{2ac}-1)}\sigma_\zeta$, which in turn implies $\alpha < 0$. It can then be shown that the expression in Proposition 4 is negative. We omit the algebra involved here, as it is complex and does not contribute to discussion. □