

The origin of bubbles

Yue Shen *
Queen's University

May 8, 2015

Abstract

In this paper we explore the fundamental question of why bubbles exist. We construct a simple model of asset bubble and show that the model is equivalent to a reverse common value auction and that bubbles exist. By decomposing equilibrium bidding strategies into two components (incentives), we show that bubbles arise because this is a discriminatory price (first-price) auctions, i.e. a trader's payoff is exactly her bidding price, while in a uniform price (second-price) auction there is no bubble. Based on this finding, experiments are devised to offset these incentives and reduce or eliminate bubbles.

1 Introduction

Asset bubbles are usually defined as large price deviations from their fundamental or intrinsic values that can last for extended periods. Historical examples of bubbles include the Dutch tulip mania of the 1630s, the South Sea bubble of 1720 in England, the Mississippi bubble in France, the Great Recession of 1929 in the United States, DotCom bubble in the late 1990s, and the most recent housing bubble crash since 2008. Such phenomena have long been intriguing to economists because they not only affect the financial sector, but also have huge impact on the real economy. In addition, bubbles have been difficult to explain and generate theoretically. One major hurdle is the “no trade theorems”. One version of the theorem states that, if, the initial allocation is efficient, there is common knowledge that all traders are rational, and agents have common priors about the distribution of asset fundamental values, then agents have no incentive to trade¹. In dynamic settings, in particular, the standard neoclassical theory precludes the existence of bubbles by backward induction².

*Department of Economics, Queen's University, Kingston, Ontario, Canada. sheny@econ.queensu.ca

¹See Grossman and Stiglitz (1980) and Milgrom and Stokey (1982).

²See Santos and Woodford (1997).

Facing these impossibilities, some economists choose to deny that bubbles exist³. Meanwhile, many explanations to bubbles have been suggested (see the brief literature review at the end of section 1), but, even among those who believe that bubbles exist, the cause of bubbles remains disputed.

In this paper, we provide some new insight on why bubbles arise. We first present a simple model of asset bubbles. The bubble in our model arises from rational traders' asymmetric information and the interaction between rational and behavioral traders. We transform the framework of Abreu and Brunnermeier (2003) from the dimension of time to the dimension of value/price. The asset has an unobservable fundamental value and all rational traders have a common prior belief about the fundamental. In addition, each trader receives a private signal. The price keeps growing exogenously. The bubble, if any, will burst when a certain fraction of traders have sold. A trader's problem is to decide when to sell the asset. She understands that, by selling early, she can make a small profit; by selling late she might be able capture a large price appreciation but also have higher risk of being caught in the crash. A key ingredient is that a trader does not know how many others have belief higher or lower than hers. Therefore, every trader has a different posterior belief about the asset fundamental value and the size of the bubble. This dispersion of belief induces dispersion of exit strategies, which is what allows the bubble to arise. It is worth mentioning that three frictions are at work allowing the bubble to arise and grow: 1. limited short selling, 2. lack of common knowledge of the population's belief, and 3. unobservability of individual's trade.

Then we show that this model is equivalent to a discriminatory price (first-price) common value auction. The auction is one where traders simultaneously submit bids, and the lowest bids win and each receives their individual bidding prices, while the rest receive a unknown common value outside option. The former corresponds to those who sell before the crash and the latter to those who can only sell at the revealed fundamental value. Recent literature in auctions shows that prices⁴ do not converge to the true value in discriminatory price common value auctions⁵. This equivalence means that bubbles arise due to the same reason as price deviating from true value in the discriminatory price common value auctions. As a trading mechanism, discriminatory price common value auctions fail to aggregate information dispersed in the bidders and bidders' strategies systematically deviate from the true value. Equivalently, market often overreacts to news because of the inability of privately informed

³See Garber (1990)

⁴The price here is the marginal bid, i.e., the highest winning bid or lowest losing bid.

⁵A detailed review of this literature is postponed to Section 3.2.

investors to reach consensus and they optimally riding the bubble. As a result, this equivalence might be interpreted as a “bad news” for efficient market hypotheses, because it highlights that the bubbles and crashes can arise from the information asymmetry and failure of information aggregation.

Next, we identify two bid shading incentives in traders’ strategies that generate bubbles. In this discriminatory price reverse common value auction, two (reverse) bid shading⁶ incentives exist: 1. traders’ efforts to avoid both (reverse) winner’s and loser’s curse and 2. traders’ price-setting incentives (because they get what they bid upon winning) in first-price auctions compared to price-taking behavior in second-price auctions. We show that the size of the bubble equals the price setting component in the strategies. Compared with previous auction literature, the two incentives in my model can be explicitly separated in the bidding strategies due to the specification of prior and posterior belief in our model. This separation may also help explain why price deviate from true value in the discriminatory price common value auctions. In addition, the two incentives have opposite responses to the change of winning percentage in the auction (the ratio between “winning slots” and number of bidders), which can explain some differences between several models’ results in the literature.

Based on above insight, we test several tax schemes on our model and we find that if the tax and subsidy are designed such that the payoff structures become a uniform price (second-price) auction, then there is no bubble at all. This suggests that taxes and subsidies can be used to change the payoff structures and reduce the bubble.

Our model is related to a vast literature on bubbles. For surveys on bubbles, refer to Brunnermeier (2009), Brunnermeier (2001), Brunnermeier and Oehmke (2012) and Scherbina (2013).

One strand of the literature allows heterogeneous priors and rational agents to “agree to disagree”. Harrison and Kreps (1978) show that when agents disagree about the probability distributions of dividend streams, a trader may buy an asset at a price that exceeds her own valuation since she believes that she can find a more optimistic buyer. Scheinkman and Xiong (2003) justify this behavior by overconfidence as a source of belief disagreement. Allen et al. (1993) and Conlon (2004) allow agents to have heterogeneous priors and hold worthless assets in hopes of selling it to greater fools. One advantage of these models is that all agents are rational, but, to some extent, they lack a clear price path with steadily growth and then a sudden collapse. My model adopts a more conventional and parsimonious setting with a common prior belief and the upturn and crash are simple and clear.

⁶Bid shading is the practice of a bidder placing a bid below what she believes a good is worth.

Another strand of literature, including our model, studies bubbles that arise from the interaction between rational traders and behavioral traders. In De Long et al. (1990), rational traders buy the asset and induce positive-feedback traders to purchase. After the price is pushed up, the rational traders can profitably unload. The bubble exists because time is discrete in their model. If we allow continuous time and assume that, when rational traders unload together not all of them will receive the full price⁷, then the bubble will disappear. This is because rational traders will then have incentive to preempt each other, which diminishes the size of the bubble until it becomes zero. My model is robust in this sense and the lack of common knowledge prevent rational traders from preempting each other.

In a different strand of literature, bubbles arise from the herding behavior. See Brunnermeier (2001) and Chamley (2004) for surveys on this topic. Avery and Zemsky (1998) introduce a sequential trade model, where the true value of the asset is either 0 or 1, and the next trader in line is uninformed with probability $1 - \mu$ and informed with probability μ . Informed traders also have two private types: poorly and perfectly informed. Three dimensions of uncertainty work together to generate an upward informational cascade and a crash: the signal about the asset value, the quality of the signals, and the uncertainty about the μ itself. The market maker adjusts the price after observing each round of transaction. A “bubble” arises when many of the poorly informed traders are herding so that the market maker keeps raising the price, and then a perfectly informed trader shows up and behave differently, and the bubble bursts. Lee (1998) studies an “information avalanche”, where a trader must pay a fee but can decide when to trade. The market maker can adjust the price after every round of trade. All traders are informed but differ in the precision of their information. A partial informational cascade can occur if traders “wait and see” and then an extreme signal triggers all previous inactive traders to sell. In both models, the entire upturn arise from the uncertainty about the value, while in my model the upturn is exogenous though the bubble component is endogenous. In their models the event of bubbles followed by crashes happen with very small probabilities and traders are only given one opportunity to trade, while in my model the event has to happen and traders can buy and sell for as many times as they want.

We build on the approach of Abreu and Brunnermeier (2003). Abreu and Brunnermeier (2003) (henceforth AB2003) challenge the efficient markets perspective by showing that it is optimal for rational

⁷This is a realistic assumption in that it takes time to sell a chunk of shares, especially when the market is near the brink of a crash.

traders to ride the bubble. In their model, the price grows exogenously and, at some random moment, the growth rate of fundamental value falls behind that of the price, hence a bubble emerges. Rational traders become aware of this sequentially. The bubble bursts when a certain fraction of traders have sold and trader is to decide when to sell the asset. We simplify AB2003 by transforming the uncertainty on the dimension of time to value/price. This transformation abstracts time away and removes the assumption of sequential awareness. Now the price is the only concern for agents, which makes the model comparable to an auction.

Results from recent empirical studies of stock market are consistent with AB2003. Temin and Voth (2004) show that a major investor in the South Sea bubble knew that a bubble was in progress and nonetheless invested in the stock and hence was riding the bubble profitably. Brunnermeier and Nagel (2004) and Griffin et al. (2011) both study the Tech bubble in the late 1990s. They show that, instead of correcting the price bubble, hedge funds turned out to be the most aggressive investors. They profited in the upturn, and unloaded their positions before the downturn.

Doblas-Madrid (2012) (henceforth DM2012) constructs a discrete-time version of AB2003 and addresses certain issues in AB2003. In AB2003 the price increase is exogenous and only the arbitrageur's selling decisions are modeled. In addition, the price does not respond to rational traders' sales until an exogenous threshold is reached. DM2012 removes the behavioral agents from the model, partially endogenize the upward price path and makes the price responsive to sales. His model inherit AB2003 These features make his model more realistic.

AB2003 and DM2012 provide a convincing framework of bubbles, but why bubbles exist is not entirely clear. AB2003 explains that the lack of common knowledge prevents traders from perfectly coordinating, and hence the backward induction has no bite and a bubble can exist. However, this is only a necessary condition for bubbles to exist. Brunnermeier and Morgan (2010) show that a discrete-trader version of AB2003 can be recast as an auction, but they do not pursue this direction. In particular, they are silent on how the emergence of bubbles is related to auction theory. I modify and preserve the backbone of AB2003's framework and focus on incentives in traders' strategy that generate bubbles.

The remainder of the paper is organized as follows. Section 2 introduces the model and characterizes the equilibrium. Section 3.1 shows the equivalence of the bubble model to a common value auction. Section 3.2 reviews the auction literature and shows that prices generally do not converge to true value

in a common value auctions. Section 3.3 identifies the two reverse bid shading incentives which generate the bubble, Section 3.4 devises several simple tax/subsidy experiments to show how to reduce the size of the bubble, and Section 3.5 shows the two incentives respond differently to changes in model parameters. Section 4 concludes.

2 The model

Time t is continuous and there is only one asset. The asset's fundamental value is θ , which is unobservable. There is a unit mass of risk neutral rational traders (henceforth *traders*), each holding 1 unit of the asset at the beginning. Limited short selling is allowed and without loss of generality each trader's asset position is restricted and normalized to $[0, 1]$. A trader can buy and sell shares at any time, which is unobservable to others. The price can be publicly observed and is denoted $p(t)$. Without loss of generality, assume that $p(0) = 0^8$. After $t = 0$, the price increases continuously and deterministically. At any time, when the price rises above θ , we say there is a *bubble*. There is no discounting.

As in AB2003, the backdrop is that the asset price keeps increasing, which can be interpreted as there are behavioral agents buying the asset. These behavioral agents are overly optimistic. They believe that a technology shock has permanently raised the productivity. They keep buying the asset, which pushes up the price, as was the case during the tech bubble in the late 1990s. However, the price will not rise forever. Rational traders start to sell when they gradually believe that the price is too high. We assume that when an exogenous fraction κ ($0 < \kappa < 1$) of the rational traders has sold, the price stops increasing and jumps instantly to its fundamental value and stays there thereafter, i.e., the bubble bursts. In line with AB2003, we call this an endogenous crash. This threshold can be interpreted as a point at which the selling pressure cannot be concealed by the price noise, and that the price is too high becomes a common knowledge among all rational traders and behavioral agents. Behavioral agents are not explicitly modeled here, and a rational trader is only concerned about how many other rational traders have sold. An example of the price path is depicted in Figure 1. The size of the bubble at t is the gap between the current price $p(t)$ and θ when $p(t) > \theta$. To rule out the nuance equilibrium where all traders hold the asset forever and never sell, we assume there is an exogenous upper bound \bar{B} for the bubble size, and the bubble will burst when the bubble size is larger than \bar{B} , even if the fraction

⁸As long as the initial price equals the lowest possible θ , results will not change (up to a shift).




Figure 1: An example of the price path

of the rational traders who has sold is less than κ . This is the exogenous crash in AB2003. Imposing such an exogenous upper bound is reasonable because, after all, when an asset's book value exceeds the GDP of the whole economy, no one (not even behavioral agents) believes that the price is a fair reflection of its fundamental value. We are only interested in the endogenous crash, so we assume that \bar{B} is large enough. We will show that in equilibrium \bar{B} is never binding and the bubble always bursts due to the threshold κ .

θ has a probability density $\phi(\theta)$ over $[0, \infty)$. We restrict to two alternative distributions: an improper uniform distribution on $[0, \infty)$ and an exponential distribution with density $\phi(\theta) = \lambda e^{-\lambda\theta}$, since they give simple solution forms⁹. At $t = 0$, each trader receives a private signal v , which is uniformly distributed on $[\theta - \frac{\eta}{2}, \theta + \frac{\eta}{2}]$. This is the only signal that a trader receives in the game. v can be regarded as a trader's type. Let $\Phi(\theta|v)$ be the posterior CDF about θ of a trader with signal v , and $\phi(\theta|v)$ be the corresponding PDF. Given signal v , the support of the posterior is $[v - \frac{\eta}{2}, v + \frac{\eta}{2}]$. In the exponential prior case, $\Phi(\theta|v) = \frac{e^{\lambda\eta} - e^{\lambda(v + \frac{\eta}{2} - \theta)}}{e^{\lambda\eta} - 1}$, and in the uniform prior case, $\Phi(\theta|v) = \frac{\theta - (v - \frac{\eta}{2})}{\eta}$ ¹⁰. Figure 2 depicts the posterior belief about θ for trader v , v' and v'' in the case of uniform prior belief. These different posterior beliefs reflect different opinions about the asset fundamental value. A trader is not sure about her signal's position within $[\theta - \frac{\eta}{2}, \theta + \frac{\eta}{2}]$, i.e., a trader does not know how many others' signals are lower than hers, and how many are higher than hers. This is an important element

⁹The improper uniform distribution on $[0, \infty)$ has well defined posterior belief when we specify how signals are distributed. The uniform prior case is adapted from Li and Milne (2014) and the exponential prior case is adapted from AB2003

¹⁰When $\theta < \frac{\eta}{2}$, some traders will receive negative signals, which is perfectly compatible with the assumption that θ is non-negative. When $\theta < \eta$, those with $v < \frac{\eta}{2}$ will have a truncated belief support because θ cannot be below zero. This cause traders with very low signals behave differently, which will be clarified in Proposition 2.1 where we characterize the equilibrium. The rest of this section will ignore this special case.




Figure 2: Posterior beliefs

in the model because the lack of common knowledge of the distribution of the population prevents traders from perfectly coordinating each other. In contrast, in the standard literature with a common posterior belief, perfect coordination leads to backward induction, which rules out the existence of the bubble. The above specifications are such that all traders' posterior belief have exactly the same shape except for a horizontal shift. Hence, traders will behave the same relative to their respective signals (except those close to the lower boundary), and everyone sells at a different price because they each has a different signal. It is worth mentioning that three frictions in this market are at work together to generate the bubble: 1. lack of common knowledge, 2. limited short selling, and 3. unobservability of individual's trade.

We make the following two assumptions, which can dramatically simplify the analysis while retaining the main flavor of AB2003.

Assumption 2.1. *A trader uses a trigger strategy: she sells only once, whereby she sells all her shares and will never buy back.*

Given all traders' signal profile, let $P^*(v)$ denote the selling price of a trader v .

Assumption 2.2. $P^*(\cdot)$ is continuous, strictly increasing and differentiable on $(-\frac{\eta}{2} + \eta\kappa, \infty)$.

This assumption guarantees that traders with higher signals must sell at higher prices. It also implies that $P^{*-1}(\cdot)$ is well defined.

Let p_T denote the highest price before crash. Since θ is a random variable, $p_T = P^*(\theta - \frac{\eta}{2} + \eta\kappa)$ is also a random variable which is perfectly correlated with θ , given $P^*(\cdot)$. Suppose a trader v decides to sell at price p , then if $p < p_T$, she will be able to flee the market before the crash; otherwise, she will be caught in the crash. By inverting this relationship, we know that $\theta = P^{*-1}(p_T) + \frac{\eta}{2} - \eta\kappa$, and she will get caught if $\theta \in [v - \frac{\eta}{2}, P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa]$, and she will flee the market successfully before the crash if

$\theta \in [P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa, v + \frac{\eta}{2}]$. Denote $\omega(p)$ the expected payoff from selling at price p . Given that all others use strategy $P^*(\cdot)$, the expected payoff for trader v is

$$\omega(p) = \int_{\theta=v-\frac{\eta}{2}}^{P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa} \theta\phi(\theta|v)d\theta + \int_{P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa}^{v+\frac{\eta}{2}} p\phi(\theta|v)d\theta \quad (1)$$

Note that the belief in this game is static. Although the lower boundary of the posterior belief may shrink upwards¹¹ over time when the price increases¹¹, this change is expected. Pre-crash sale price p in (1) is conditional on selling before the crash, and any price in the process up to p is expected and already considered in (1). Hence there is no need to update the belief. This invariability is similar to the equivalence between a common value first-price sealed bid auction and a common value Dutch auction.

A trader's problem is

$$\max_p \omega(p)$$

We first solve for the equilibrium strategy $P^*(\cdot)$, then we characterize the unique equilibrium in Proposition 2.1.

Differentiating $\omega(p)$ w.r.t p , imposing $P^*(v) = p$, and set $\frac{dE[R|v]}{dp} = 0$, we have the first order condition

$$1 = \left[P^* - \left(v + \frac{\eta}{2} - \eta\kappa \right) \right] \frac{1}{P^{*'}} \frac{\phi\left(v + \frac{\eta}{2} - \eta\kappa|v\right)}{1 - \Phi\left(v + \frac{\eta}{2} - \eta\kappa|v\right)} \quad (2)$$

This FOC can be interpreted in terms of marginal benefit (MB) and marginal cost (MC). For a trader who evaluates selling at p vs. $p + \Delta$ (Δ is small), the relevant marginal event is that the bubble bursts between p and $p + \Delta$, which implies that $\theta = v + \frac{\eta}{2} - \eta\kappa$. The benefit of selling at $p + \Delta$ instead of p is Δ (the price appreciation). The cost is that she could get caught in the crash if the bubble bursts in between p and $p + \Delta$, which equals the loss $p - \theta$ (due to bubble bursting) multiplied by $\frac{\Phi(P^{*-1}(p+\Delta)+\frac{\eta}{2}-\eta\kappa|v)-\Phi(P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa|v)}{1-\Phi(P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa|v)}$ (the probability of bursting between p and $p + \Delta$). Dividing both sides by Δ and letting $\Delta \rightarrow 0$, we have MB = 1 and MC = $(p - \theta) \frac{1}{P^{*'}} \frac{\phi(v+\frac{\eta}{2}-\eta\kappa|v)}{1-\Phi(v+\frac{\eta}{2}-\eta\kappa|v)}$ ($\theta = v + \frac{\eta}{2} - \eta\kappa$ if the bubble bursts between p and $p + \Delta$).

¹¹The initial belief $\phi(\theta|v)$ is uniform over $[v - \frac{\eta}{2}, v + \frac{\eta}{2}]$. Given that all others use equilibrium strategy $P^*(\cdot)$, the bubble will burst at $p_T = P^*(\theta - \frac{\eta}{2} + \eta\kappa)$ and hence $\theta = P^{*-1}(p_T) + \frac{\eta}{2} - \eta\kappa$. If the bubble bursts at any price p_e , then it must be that $\theta = P^{*-1}(p_e) + \frac{\eta}{2} - \eta\kappa$. When the current price p_e is such that $P^{*-1}(p_e) + \frac{\eta}{2} - \eta\kappa < v - \frac{\eta}{2}$, where $v - \frac{\eta}{2}$ is the lower bound of θ , the bubble will certainly not burst for trader v . When price p_e has increased such that $P^{*-1}(p_e) + \frac{\eta}{2} - \eta\kappa > v - \frac{\eta}{2}$, from trader v 's point of view, the bubble can burst any moment. The fact that the bubble has not burst below p_e implies that θ cannot be below $P^{*-1}(p_e) + \frac{\eta}{2} - \eta\kappa$, which shrinks the support of θ from below and the new support of θ is now $[P^{*-1}(p_e) + \frac{\eta}{2} - \eta\kappa, v + \frac{\eta}{2}]$.

2.1 Equilibrium

Proposition 2.1. *There is a unique equilibrium, in which a trader sells at price*

$$\begin{cases} B, & \text{if her signal } v < -\frac{\eta}{2} + \eta\kappa \\ P_s^*(v) = v + \frac{\eta}{2} - \eta\kappa + B, & \text{if her signal } v \geq -\frac{\eta}{2} + \eta\kappa \end{cases}$$

This gives rise to a bubble of the size B , where $B = \begin{cases} \eta\kappa, & \text{if prior belief is uniform} \\ \frac{1-e^{-\lambda\eta\kappa}}{\lambda}, & \text{if prior belief is exponential} \end{cases}$

See Appendix A for proof¹². The equilibrium strategy is depicted in Figure 3.




Figure 3: Equilibrium strategy

3 The origin of bubbles

3.1 Relationship to an auction

The model in Section 2 can be recast as a pure common value auction. Specifically, this is a reverse discriminatory price (first-price) sealed-bid multi-unit auction with a single unit demand and common value outside option, and bidders are continuous instead of discrete. We will explain the terms step by step. It is a reverse auction (also called procurement auction) because traders/bidders are selling instead of buying. In this procurement auction, the auctioneer has a continuous mass of κ ($\kappa < 1$) of (identical) projects/contracts and would like to buy (labor or service) from a mass of one of bidders who want to provide/sell (their labor or service). A mass of κ of the bidders who bid the lowest (selling at lowest prices) win and each winner receives a contract value equal to their respective bidding price (selling price)¹³. In contrast, in a normal auctions, bidders are buyers and winners are those who bid

¹²The proof of uniqueness for those traders in the lower boundary $[-\frac{\eta}{2}, -\frac{\eta}{2} + \eta\kappa]$ requires an additional technical assumption: when any positive mass of traders sell at the same price and the bubble bursts right at that price, only some of them (random draw) can sell at the pre-crash price, while others have to sell at the post-crash price θ .

¹³Winners need not exert any effort to fulfill the contract.

the highest. It is a discriminatory price auction because each winner gets her own bidding price, which corresponds to the first-price in the single object auction. The losers receive a common value outside option θ . Figure 4 depicts the bidding strategies. Winners in the auction are those whose selling/bidding




Figure 4: Bidding strategy

prices are low and who sell before the crash, which is the thick solid line. The losers are those who are caught in the crash, and the thick dashed line represents their planned, though never realized, selling prices, since the bubble bursts before they have a chance to sell. All traders are involved in this single auction.

In order to win, a trader necessarily bids low enough to be in the lower fraction of κ of all bids. But the lower she bids, the lower the payoff upon winning. Since an agent receives either θ or p (her bidding price), she only compares the two alternatives and her bidding strategy only depends on her belief about the common value outside option θ : the lower her belief about θ , the lower she bids. This is why we say it is a common value auction. In this case the expected payoff with signal v is exactly equation (1).

Equivalently, we can recast the bubble model as the same auction but without outside option. Instead, winners need to exert a common effort θ to fulfill the contract, which is uncertain at the time of bidding¹⁴. Then the expected payoff of trader v changes to

$$\omega(p) = \int_{P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa}^{v + \frac{\eta}{2}} (p - \theta)\phi(\theta|v)d\theta$$

This expected payoff is different from the auction with outside option in equation (1), but the first order condition is the same, as in equation (2), so the equilibrium strategy is the same.

¹⁴The common effort can be thought of as each winning bidder to build an identical house with the same design.

In what follows, we will show that the above auction, as well as some other common value auctions, fails to aggregate information and bidding prices deviate from the true value of the objects in auction.

3.2 A review of common value auction literature: price convergence

In Figure 4 the marginal bid, which is the highest winner's bid and also the lowest loser's bid, is the bid that equals the bursting price. In our bubble model (a discriminatory-price common value auction), the marginal bidding price is higher than the fundamental value. In the case of uniform prior belief, for instance, the marginal bid is $\theta + \eta\kappa$, which is higher than the true value θ . If somehow the marginal bidding price equals the true value θ , then there will be no bubble. Similarly, in the literature of uniform-price (second-price) common value auctions, if the bidding prices converge to the true value of the object, then it is said that the information aggregation holds. The price convergence, if holds, demonstrates that, while no one knows for sure the true value of the objects, the auction as a price formation process can aggregate information diffused in the economy and reveal the true value.

Early literature in common value auction shows that bidding prices converge to true value of object and that auction is a form of trading that achieves market efficiency¹⁵. However, the assumptions adopted in the early literature turn out to be quite strong. They not only require the monotone likelihood ratio property (MLRP), but also require that the likelihood ratio approaches to zero¹⁶.

As recent literature in common value auctions shows¹⁷, when the requirement that the likelihood ratio converges to zero does not hold, the price generally fails to converge to the true value. Kremer (2002) shows that both first and second-price single object common value auctions fail to aggregate information. Jackson and Kremer (2007) show that, in the discriminatory price common value auction with k identical objects for sale and each bidder desiring only one item, the information aggregation also fails and the price does not converge to asset value even when both $k \rightarrow \infty$ and $n \rightarrow \infty$. In particular, they show that the marginal/pivotal bid in the auction is lower than the true value of the

¹⁵Wilson (1977) and Milgrom (1979) showed that, in a first-price common value single unit auction, the winning bid converges in probability to the value of the object as the number of bidders n becomes large. Milgrom (1981) showed that, in a uniform-price (second-price) common value auction with k identical objects for sale and each bidder only desiring one item, where all winning bidders pay the (same) $k + 1$ th bid (which corresponds to the second price in single object auction), if we fix k and let the total number of bidders $n \rightarrow \infty$, then the price (the highest loser's bid, which is also the $k + 1$ th highest bid) also converges to the true value.

¹⁶Let $f(v|\theta)$ be the density distribution of a bidder's estimate when true value is θ . MLRP requires that, if $\theta_1 < \theta_2$, $\frac{f(v|\theta_1)}{f(v|\theta_2)}$ decreases in v , which most distributions satisfy. But the price convergence also requires that $\frac{f(v|\theta_1)}{f(v|\theta_2)} \rightarrow 0$ as $v \rightarrow \bar{v}$, where \bar{v} is the upper bound of the support of v .

¹⁷Bidders and objects are all discrete in these models.

asset. The only situation where the price does converge to true value is Pesendorfer and Swinkels (1997), who show that when both $k \rightarrow \infty$ and $n \rightarrow \infty$, the price converges to true value in the uniform price common value auctions.

To apply these results, we expand our model (a reverse discriminatory price common value auction) to a uniform-price setting and compare its result with our discriminatory case. Under uniform-price, all winners receive the marginal bidder's bidding price. We also compare our models (continuous bidder and object, reverse auction) with the standard auction models (discrete n bidders and k objects, normal auction) under the framework of common value auctions. The price convergence are summarized in Table 1.

	Uniform-price	Discriminatory price
Standard auctions (normal, $k \rightarrow \infty$, $n \rightarrow \infty$)	$P \rightarrow V$	marginal $P < V$
Our models (reverse, continuous)	$p_T = \theta$	$p_T > \theta$

Table 1: Price convergence (p_T is bursting price)

Under discriminatory-price, standard auctions (Jackson and Kremer (2007)) have marginal bidding price that is lower than true value. Since our model is a reverse auction, this means that the marginal price in our reverse auction should be higher than the true value, which is true. This conformity shows that bubbles arise are not due to any peculiar assumption in our model and they exist in more general settings. Why bubbles arise in our model has the same root as why the marginal bidding price does not equal the true value in the discriminatory price common value auctions. When private information is dispersed in market participants and no one is sure about the asset fundamental value, the market generally fails to aggregate these information and price fails to reflect average belief. Riding the bubble becomes optimal to each investor. Compared to previous literature in asset bubbles, this interpretation highlight and further clarify that the bubbles and crashes can arise from the information asymmetry, which is an intrinsic characteristic of the market. This is a bad news for market efficiency hypothesis in the sense that, every time a news arrives and belief differ, there could potentially be a coordination failure and hence a bubble arises.

Under uniform-price, the bidding price converges to the true value in the stand auctions (Pesendorfer and Swinkels (1997)), which is compatible to our model where there is no bubble. This result suggests that there are potential opportunities to reduce the bubble, and we will discuss them in Section 3.4.

3.3 Two bid shading incentives

Up to this point, it is still unclear why bubbles necessarily arise in our model, or equivalently why the marginal bidding price is higher (lower) than the true value in the reverse (normal) discriminatory price common value auctions. One can further ask that why it is necessarily an upward instead of a downward deviation. There does not seem to be any particular reason in AB2003, or DM2012, or our model that forces traders to sell or wait till the price is high enough such that a bubble emerges. We try to answer these questions by decomposing traders' strategies in our original model into two (reverse) price shading incentives: the *marginal bidder incentive* and the *price setting incentive*, then we show how they affect the marginal bid so that the bid deviates upwards from true value and generates a bubbles.

The equilibrium selling price $P^*(v)$ can be decomposed as follows.

$$\begin{aligned} P^*(v) &= v + \text{marginal bidder incentive} + \text{price setting incentive} \\ &= v + \frac{\eta}{2} - \eta\kappa + \frac{1 - \Phi(v + \frac{\eta}{2} - \eta\kappa|v)}{\phi(v + \frac{\eta}{2} - \eta\kappa|v)} \end{aligned}$$

3.3.1 Marginal bidder incentive

The marginal bidder incentive is that bidders try to avoid both winner's and loser's curse, so that the marginal bidder bids exactly at the true value. It is defined as bidding by conditioning on being the marginal bidder. This incentive exists in all common value auctions and in common value auctions only.

This incentive is best illustrated in uniform-price common value (normal) auction, where there is no interference of the other incentive. The winner's curse is that, if a bidder turns out to be the winner, it means that her signal is one of the highest among all bidders and hence biased upward. Hence, if she has bid naively, it is very likely that she overestimated and overbid. This bias is more striking in the classic auction models where there are n bidders and only one object for auction. Hence sophisticated bidders in uniform-price common value auctions shade their bids conditioning on winning. The *loser's curse* is much less well-known¹⁸ because in most auction models there is only one object, while loser's curse matters only when there are a significant number of objects compared with the number of bidders¹⁹. The loser's curse is that, upon losing, a bidder finds that her signal is among the lowest and hence

¹⁸Pesendorfer and Swinkels (1997) discuss this point in detail.

¹⁹Symmetrically, the winner's curse does not matter when there are n bidders and $n - 1$ objects.

she may have underestimated the true value and underbid. This is more striking when there are n bidders and $n - 1$ objects for auction. To avoid this curse, sophisticated bidders should shade their bids conditioning on losing.

In general cases where there are n bidders and $k < n$ objects, the equilibrium bidding strategy in uniform-price common value auction is to bid conditioning on being the marginal bidder²⁰. This is because, if the marginal bidder's bid is higher than true value, then all winners will be paying a price higher than the true value due to the uniform-price setting. Then all bidder will revise their bids downwards. If the marginal bid is lower than the true value, then a bidder has incentive to raise her bid to have higher probability of winning while (essentially) not affecting the price she pays upon winning. Ex ante, no one knows whether she will be the marginal bidder. But, in equilibrium, everyone behaves as if she is the marginal bidder²¹.

When our model is extended with a uniform-price setting, the bidding/selling strategy is $v + \frac{\eta}{2} - \eta\kappa$. Conditional on being the marginal bidder, i.e. her signal is $\theta - \frac{\eta}{2} + \eta\kappa$, a trader's bid will be exactly the true value θ . Therefore, $\frac{\eta}{2} - \eta\kappa$ is bidders' effort (on top of v) to avoid both of the winner's and loser's curses and to bid exactly at θ if being marginal bidder (See Table 2 for a summary).

3.3.2 Price setting incentive

The price setting incentive is that traders try to set the price to seize extra surplus in discriminatory price auctions, compared to the price-taking behavior in uniform price auctions. This incentive exists in all discriminatory price auctions and in discriminatory price auctions only.

In uniform-price auctions, a bidder does not pay what she bids and her bid essentially has no impact on her payment. Hence she behaves like a price-taker. In contrast a bidder in normal (reverse) discriminatory price auctions pays (gets) exactly what she bids upon winning. Hence, in normal (reverse) auctions she has incentive to lower (raise) her bid.

This incentive is best illustrated in private value auctions. In normal first-price private value auctions, if a bidder bids exactly at her private value, then obviously she always has zero surplus. To extract positive surplus, she shades her bid and her equilibrium bidding strategy is higher than her private value²². To illustrate the price setting incentive, we extend our model along a second dimension

²⁰In fact, in both of the single object and $n - 1$ objects cases, we have already conditioned on being the marginal bidder.

²¹This is why in the normal discrete bidder uniform-price common value auctions, the equilibrium strategy is $E[v|X_1 = v, Y_k = v]$, where X_1 is my signal, and Y_k is the k th highest signal among all other bidders.

²²In normal first-price common value auctions, this incentive still exists. But when she lowers her bid, she also lowers

and apply the private value setting. In the private value setting, a winner needs to exert an effort equal to her signal v instead of an common value θ . When our model is set under the private value setting, the price setting incentive in the uniform prior case is (positive) $\eta\kappa$, and in the exponential prior case $\frac{1-e^{-\lambda\eta\kappa}}{\lambda}$ (See Table 2 for a summary). These terms are actually the inverse of the hazard rates in the FOCs (recall that the hazard rates are $\frac{1}{\eta\kappa}$ and $\frac{\lambda}{1-e^{-\lambda\eta\kappa}}$ in the uniform prior and exponential prior cases, respectively). If we re-write Equation (2), we have²³

$$P^*(v) = \left(v + \frac{\eta}{2} - \eta\kappa\right) + \frac{1 - \Phi\left(v + \frac{\eta}{2} - \eta\kappa|v\right)}{\phi\left(v + \frac{\eta}{2} - \eta\kappa|v\right)}$$

where the term $\frac{1-\Phi(v+\frac{\eta}{2}-\eta\kappa|v)}{\phi(v+\frac{\eta}{2}-\eta\kappa|v)}$ is the general form of the price setting incentive. In comparison, the FOCs in the uniform-price cases (uniform and exponential priors) do not have this price-setting term. In uniform-price (reverse) auctions, since a bidder does not pay what she bids, she would bid infinitely high if she is guaranteed to win. When she has to consider the possibility of losing, then the situation where she is the marginal winner matters. In this case, $\frac{1-\Phi(v+\frac{\eta}{2}-\eta\kappa|v)}{\phi(v+\frac{\eta}{2}-\eta\kappa|v)}$ is exactly what she can add to her bid to balance between seizing extra value and not forgoing too much opportunity of winning. Finally, notice that the price setting incentive under both prior beliefs, $\eta\kappa$ and $\frac{1-e^{-\lambda\eta\kappa}}{\lambda}$ are increasing in κ , which means that higher κ allows more winning slots and therefore relieves the competition and results in more aggressive bids (higher selling prices).

3.3.3 A comparison

These two incentives are well-known in the auction literature (may be under different names), but are entangled in the bidding strategies and have not been separated before. Our model provides a unique opportunity where we can separate them explicitly, which is due to the information structure and distribution specifications in our model²⁴. Table 2 summarizes the selling/bidding strategies from her winning probability. So the equilibrium strategy has to balance these two forces.

²³For simplicity, we assume that $P^{*'} = 1$.

²⁴The reason why our model has a simple solution form and always has bubbles is due to the lack of common knowledge adopted from AB2003. Specifically, except at the lower boundary, there is no common knowledge about the lower or upper bound for current population's signals. In contrast, Jackson and Kremer (2007) assume that signals are always distributed within $[0, 1]$, irrespective of the realization of object value. The object value only affects the shape of the signal distributions. This specification introduces an extra complication: for those whose signals are high enough, they know for sure that they are among the highest types and they will win, hence they have incentive to further shade their bids; symmetrically, for those whose signals are low enough, they have incentive to bid even higher. As a result, bids are somewhat concentrated in the middle. When the realized object value is high, the average transaction price is below the true value; when the object value is very low, the average transaction price is higher than the true value. However, on average, the expected price is lower than the expected value, and in particular, the marginal bid is always lower than the true value.

our model, which is extended along two dimensions and hence has four cases. One can observe that

	Uniform prior		Exponential prior	
	Uniform price	Discriminatory price	Uniform price	Discriminatory price
Private value	v	$v + \eta\kappa$	v	$v + \frac{1-e^{-\lambda\eta\kappa}}{\lambda}$
Common value	$v + \frac{\eta}{2} - \eta\kappa$	$v + \frac{\eta}{2} - \eta\kappa + \eta\kappa$	$v + \frac{\eta}{2} - \eta\kappa$	$v + \frac{\eta}{2} - \eta\kappa + \frac{1-e^{-\lambda\eta\kappa}}{\lambda}$

Table 2: Bidding strategies in continuous bidder/object reverse auctions

the marginal bidder incentive exists in all common value auctions only, and the price setting incentive exists in all discriminatory price auctions only.

In our original model (the discriminatory price common value case), both incentives exist in the selling strategy. So the bubble bursts at

$$\begin{aligned}
 p_T &= P^*(\theta - \frac{\eta}{2} + \eta\kappa) = (\theta - \frac{\eta}{2} + \eta\kappa) + \text{marginal bidding incentive} + \text{price setting incentive} \\
 &= \theta - \frac{\eta}{2} + \eta\kappa + \frac{\eta}{2} - \eta\kappa + \frac{1 - \Phi}{\phi} = \theta + \frac{1 - \Phi}{\phi}
 \end{aligned}$$

By definition, the marginal bidding incentive is to ensure the marginal bidder to bid true value θ , hence the marginal bidding incentive is canceled out by the gap between marginal bidder's signal and true value θ . Therefore, both the marginal bidder and the price setting incentive affect traders' selling strategies and are the origin of bubbles, but the size of the bubble is determined by price setting incentive alone. So the size of the bubble is

$$B_{\text{price setting incentive}} = \frac{1 - \Phi(v + \frac{\eta}{2} - \eta\kappa|v)}{\phi(v + \frac{\eta}{2} - \eta\kappa|v)}$$

Table 3 shows the bursting price in the case of uniform and discriminatory price under common value. Since the price setting incentive is increasing in κ (introduced in next subsection), under both

	Uniform prior		Exponential prior	
	Uniform price	Discriminatory price	Uniform price	Discriminatory price
Common value	θ	$\theta + \eta\kappa$	θ	$\theta + \frac{1-e^{-\lambda\eta\kappa}}{\lambda}$

Table 3: Marginal bids (“bursting prices”) in continuous bidder/object reverse auctions

uniform and exponential prior, bubble size is increasing in κ .

3.4 Reduce the size of the bubble: transaction tax and subsidy

Now we impose the proportional transaction tax (Tobin tax) and subsidy to our original model and evaluate their effects on the size of the bubble. The results demonstrate how taxes and subsidies could affect traders' incentives, and provide theoretical grounds and warnings to certain tax policies.

For simplicity, in the rest of this section we focus on the uniform prior case only. Write the expected payoff for trader v in a general form:

$$\omega(p) = \int_{\theta=v-\frac{\eta}{2}}^{P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa} B\phi(\theta|v)d\theta + \int_{P^{*-1}(p)+\frac{\eta}{2}-\eta\kappa}^{v+\frac{\eta}{2}} A\phi(\theta|v)d\theta$$

where A is the pre-crash payoff, and B is the post-crash payoff. Without tax or subsidy, $A = p$ and $B = \theta$. Let τ denote the tax rate. We consider the following five scenarios: 1. taxing on all sales, 2. taxing only on pre-crash sales, 3. taxing only on post-crash sales, 4. subsidizing only pre-crash sales and 5. taxing and subsidizing pre-crash sales.

3.4.1 Experiment 1: taxing on all sales

If we levy a tax of rate τ on both pre and post-crash sale revenues, then $A = (1 - \tau)p$ and $B = (1 - \tau)\theta$. Solving a trader's problem, we find that the tax does not change traders' strategies relative to the original model, and the bubble size is unchanged. This is because the FOC is now $1 - \tau = [(1 - \tau)P^* - (1 - \tau)\theta] \frac{1}{P^{*'}(1 - \Phi)}$. Normalizing the FOC by dividing $(1 - \tau)$ on both sides, we see that the FOC is the same as the one without tax, hence the tax does not distort a trader's incentive. This result is consistent with empirical evidences and some models which show that transaction taxes failed to reduce price volatilities (see literature review).

3.4.2 Experiment 2: taxing only on pre-crash sales

If we levy the tax only on pre-crash sales, then $A = (1 - \tau)p$ and $B = \theta$, and we have a surprising result that the bubble actually becomes larger. This scenario is relevant when there is a significant tax cut after the crash to fight the recession and stimulate the economy, and when traders *anticipate* this tax cut before the crash. This gives the equilibrium strategy $P^*(v) = \frac{v+\frac{\eta}{2}}{1-\tau}$ and the bubble size becomes $\frac{\tau\theta+\eta\kappa}{1-\tau}$, which is larger than $\eta\kappa$ in the original model, i.e. the tax works against its intention to deflate the bubble. The intuition for this result is that, the FOC is now $1 - \tau = [(1 - \tau)P^* - \theta] \frac{1}{P^{*'}(1 - \Phi)}$. After normalization by dividing $(1 - \tau)$ on both sides the FOC becomes $1 = [P^* - \frac{\theta}{1-\tau}] \frac{1}{P^{*'}(1 - \Phi)}$, where the

difference between pre and post-crash sales (in the brackets) is smaller than that without tax. This difference is the the loss due to burst. A smaller loss makes being caught in the crash less scary, and this encourages traders to sell more aggressively at higher prices²⁵, which further inflates the bubble. It also shows that the bubble size actually grows with θ : the higher the realization of θ , the larger the bubble is. Therefore policy makers should be wary of this counterproductive effect of transaction taxes when dealing with bubbles, especially when it is foreseeable that the tax will be abolished after the crash due to political or economic considerations.

3.4.3 Experiment 3: taxing only on post-crash sales

If we levy the tax only on post-crash sales, the bubble will be smaller. In this case $A = p$ and $B = (1 - \tau)\theta$, and the equilibrium strategy becomes $P^*(v) = (1 - \tau)(v + \frac{\eta}{2})$ and the bubble size becomes $(1 - \tau)\eta\kappa - \tau\theta$, which is smaller than the original bubble size $\eta\kappa$ in the original model. The reason is that, with tax only imposed on θ , the loss due to burst becomes larger in the normalized FOC, which advises that traders sell more cautiously and conservatively, i.e. they forego some upcoming appreciation, and sell earlier at lower prices to secure the realized appreciation. Although the bubble becomes smaller, this tax may be politically unacceptable.

3.4.4 Experiment 4: subsidizing only pre-crash sales

Now we consider a scenario with only subsidy on pre-crash sales, which entirely eliminate the price setting incentive and hence can completely remove the bubble. Instead of proportional to sales, the subsidy is a “complement” in the form of a lump sum: government guarantees that all pre-crash sellers’ (but not post-crash sellers) final payoff equals the highest selling price p_T . In terms of implementability, the tax and subsidy have to wait until the crash has happened and the highest selling price is known, which requires that the government be able to pull out transaction records and know who have sold before the crash and at what prices. In this case, $A = P^*(\theta - \frac{\eta}{2} + \eta\kappa)$ and $B = \theta$, and the equilibrium bidding strategy is $P^*(v) = v + \frac{\eta}{2} - \eta\kappa$. The lowest type will sell at $P^*(\theta - \frac{\eta}{2}) = \theta - \eta\kappa$, and the

²⁵More formally, when the gap between pre and post-crash sales becomes smaller due to the tax, the marginal cost of holding the asset is smaller than its marginal benefit at this price, and a trader’s best response is to hold a bit longer and sell at a higher price. Since each trader is infinitesimal, one trader’s deviate does not change the loss (the bubble size), but will change her individual hazard rate $\frac{1}{P^*} \frac{\phi}{1-\phi}$. This reasoning correctly prescribes the direction of how selling strategies should adjust in response to the tax. But since this game is not strategic complementary, new equilibrium may not be found by having all traders make the same deviation and then find the new best response and iterating the above steps in hope that it converges.

pivotal type $v = \theta - \frac{\eta}{2} + \eta\kappa$ will sell at θ . It turns out that there is no bubble at all. See Figure 5. The intuition for this result is that the subsidy encourages selling at lower prices, and a trader does




Figure 5: No bubble scheme with unbalanced budget

not have to worry about selling too early. In equilibrium all traders will adjust their selling strategies such that the bubble simply disappears. Notice that we guarantee that every pre-crash seller gets the pivotal bidding price, not θ^{26} . Such a plan, however, transfers huge public funds to private investors, which may not be economically or politically acceptable.

3.4.5 Experiment 5: taxing and subsidizing pre-crash sales

Now we look at a scenario with both tax and subsidy, which can partially eliminate the price setting incentive to reduce the size of the bubble while maintaining a balanced budget. Suppose that the government announces that all pre-crash sellers will get the selling price of the median pre-crash sellers, i.e. we subsidize the lower half of pre-crash sellers and tax the higher half of the pre-crash sellers (both in lump sum). See Figure 6. Again the tax and subsidy cannot be applied until the crash has happened. Since pre-crash sellers are uniformly distributed in $[\theta - \frac{\eta}{2}, \theta - \frac{\eta}{2} + \eta\kappa]$, the median pre-crash seller is




Figure 6: Half bubble scheme with balanced budget

$\theta - \frac{\eta}{2} + \frac{\eta\kappa}{2}$, and every pre-crash seller gets $P^*(\theta - \frac{\eta}{2} + \frac{\eta\kappa}{2})$ after the tax/subsidy. Then $A = P^*(\theta - \frac{\eta}{2} + \frac{\eta\kappa}{2})$ and $B = \theta$, which gives equilibrium strategy $P^*(v) = v + \frac{\eta}{2} - \frac{\eta\kappa}{2}$. As a result, the bubble size reduces

²⁶This scenario and the next one are closely related to turning a discriminatory price common value auction to a uniform price one, which is discussed in the early version of this paper.

to $\frac{\eta\kappa}{2}$, which is half of the original size, and we can maintain a balanced budget. There is still a bubble because some pre-crash trades are taxed (as in Scenario 2), and the bubble becomes smaller because some pre-crash trades are subsidized (as in Scenario 4), i.e. we have combined Scenario 2 and 4. Sale prices before the crash are uniformly distributed in $[\theta - \frac{\eta}{2}, \theta + \frac{\eta}{2}]$. Since everyone effectively gets θ , taxing a trader $v \in [\theta - \frac{\eta}{2}, \theta - \frac{\eta}{2} + \eta\kappa]$ a lump sum amount $v + \frac{\eta}{2} - \frac{\eta\kappa}{2} - \theta$ will balance the budget (a negative amount means subsidizing instead of taxing).

3.5 Opposite responses of the two incentives to κ

Now we show that the two incentives have opposite responses to a change in κ , which can explain the difference between several models' results in the literature. In our original model (a reverse auction),

- When κ increases, the marginal bidder incentive decreases bidding prices.

When κ is larger, winner's curse becomes smaller (or equivalently, loser's curse becomes larger), because winners now occupy a larger portion of the population and hence their signals are less biased compared to the population. Sophisticated bidders respond by easing their effort to offset the winner's curse, and hence bidding prices become lower.

- When κ increases, the price setting incentive increases bidding prices.

When winning slots are scarce, the competition is more intense and bidders undercut each other and this incentive is thwarted. When κ is large, the competition becomes less intense and bidders can ask higher prices losing much winning opportunities.

The responses in a reverse auction are summarized in Table 4. The responses in a normal auction are

	Reverse auction		WC: Winner's curse CP: Competition
	$\kappa \nearrow$	$\kappa \searrow$	
Marginal bidder incentive	WC $\searrow \Rightarrow$ bid lower	WC $\nearrow \Rightarrow$ bid higher	
Price setting incentive	CP $\searrow \Rightarrow$ bid higher	CP $\nearrow \Rightarrow$ bid lower	

Table 4: Opposite response to the change in κ

symmetric.

Depending on the specification of a model, when κ changes, the response of one incentive may dominate the response of the other, which can explain why several models in the literature have different

results. The strategy $P^*(\cdot)$ in our original model under exponential prior decreases in κ . That is, when more traders are allowed to sell before the crash, they become more cautious and sell earlier. This seemingly counterintuitive result emerges because the marginal bidder incentive dominates the price setting incentive under exponential prior decreases, hence the a trader’s overall response to κ is negative. In the uniform prior case, the strategy does not depend on κ , which suggests that the two incentives perfectly offset each other. In Brunnermeier and Morgan (2010), which is a discrete agent version of AB2003, strategies are increasing in κ . This is because a trader’s payoff in their model is an exponential function of the bidding/selling price, while a trader’s effort needed to fulfill the contract is not. So when κ changes, the coefficient (multiplying the inverse hazard rate) in the price setting incentive is larger, hence it dominate the marginal bidder incentive.

4 Conclusion

In this paper we construct a simple model of asset bubbles, where rational traders optimally ride the bubble. We further show this model is equivalent to a reverse discriminatory price common value auction, which implies that bubbles arise for the same reason of why bidding prices fail to reveal the true value in that type of auction. By decomposing bidding strategies into two components, we clarify traders’ incentives to “shade their bids”. We find that the two incentives, working together, are the fundamental causes of bubbles in our model. In particular, the size of the bubble equals the price setting component. Through the failure of common value auctions as a trading mechanism to reveal the true value of objects, our results demonstrate that bubbles and crashes are the consequences of inability of the market to aggregate private information dispersed among participants.

In addition, we devise simple tax/subsidy schemes that could (partially) offset traders’ incentives and reduce the size of bubbles. These experiments suggest directions for future policy studies.

References

- Abreu, Dilip, and Markus K. Brunnermeier (2003) ‘Bubbles and crashes.’ *Econometrica* 71(1), 173–204
- Allen, Franklin, Morris S., and Postlewaite A. (1993) ‘Finite bubbles with short sale constraints and asymmetric information.’ *Journal of Economic Theory* 61(2), 206–229
- Avery, Christopher, and Peter Zemsky (1998) ‘Multidimensional uncertainty and herd behavior in financial markets.’ *American Economic Review* 88(4), 724–48

- Brunnermeier, Markus K. (2001) *Asset Pricing under Asymmetric Information: Bubbles, Crashes, Technical Analysis, and Herding* number 9780198296980. In ‘OUP Catalogue.’ (Oxford University Press)
- Brunnermeier, Markus K. (2009) ‘Bubbles: Entry in new palgrave dictionary of economics’
- Brunnermeier, Markus K., and John Morgan (2010) ‘Clock games: Theory and experiments.’ *Games and Economic Behavior* 68(2), 532–550
- Brunnermeier, Markus K., and Martin Oehmke (2012) ‘Bubbles, financial crises, and systemic risk.’ NBER Working Papers 18398, National Bureau of Economic Research, Inc, September
- Brunnermeier, Markus K., and Stefan Nagel (2004) ‘Hedge funds and the technology bubble.’ *Journal of Finance* 59(5), 2013–2040
- Chamley, Christophe (2004) *Rational herds: Economic models of social learning* (Cambridge University Press)
- Conlon, John R. (2004) ‘Simple finite horizon bubbles robust to higher order knowledge.’ *Econometrica* 72(3), 927–936
- De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann (1990) ‘Positive feedback investment strategies and destabilizing rational speculation.’ *Journal of Finance* 45(2), 379–95
- Doblas-Madrid, Antonio (2012) ‘A robust model of bubbles with multidimensional uncertainty.’ *Econometrica* 80(5), 1845–1893
- Garber, Peter M (1990) ‘Famous first bubbles.’ *The Journal of Economic Perspectives* pp. 35–54
- Griffin, John M., Jeffrey H. Harris, Tao Shu, and Selim Topaloglu (2011) ‘Who drove and burst the tech bubble?’ *Journal of Finance* 66(4), 1251–1290
- Grossman, Sanford J, and Joseph E Stiglitz (1980) ‘On the impossibility of informationally efficient markets.’ *American Economic Review* 70(3), 393–408
- Harrison, J Michael, and David M Kreps (1978) ‘Speculative investor behavior in a stock market with heterogeneous expectations.’ *The Quarterly Journal of Economics* 92(2), 323–36
- Jackson, Matthew O., and Ilan Kremer (2007) ‘On the informational inefficiency of discriminatory price auctions.’ *Journal of Economic Theory* 132(1), 507–517
- Kremer, Ilan (2002) ‘Information aggregation in common value auctions.’ *Econometrica* 70(4), 1675–1682
- Lee, In Ho (1998) ‘Market crashes and informational avalanches.’ *Review of Economic Studies* 65(4), 741–59
- Li, Mei, and Frank Milne (2014) ‘The role of a large trader in a dynamic currency attack model.’ *Journal of Financial Intermediation* 23(4), 590–620
- Milgrom, Paul, and Nancy Stokey (1982) ‘Information, trade and common knowledge.’ *Journal of Economic Theory* 26(1), 17–27

- Milgrom, Paul R (1979) ‘A convergence theorem for competitive bidding with differential information.’ *Econometrica* 47(3), 679–88
- (1981) ‘Rational expectations, information acquisition, and competitive bidding.’ *Econometrica* 49(4), 921–43
- Pesendorfer, Wolfgang, and Jeroen M. Swinkels (1997) ‘The loser’s curse and information aggregation in common value auctions.’ *Econometrica* 65(6), 1247–1282
- Santos, Manuel S., and Michael Woodford (1997) ‘Rational asset pricing bubbles.’ *Econometrica* 65(1), pp. 19–57
- Scheinkman, Jose A., and Wei Xiong (2003) ‘Overconfidence and speculative bubbles.’ *Journal of Political Economy* 111(6), 1183–1219
- Scherbina, Anna (2013) ‘Asset price bubbles : a selective survey.’ *IMF Working Paper No. 13/45*.
- Temin, Peter, and Hans-Joachim Voth (2004) ‘Riding the south sea bubble.’ *American Economic Review* 94(5), 1654–1668
- Wilson, Robert (1977) ‘A bidding model of perfect competition.’ *Review of Economic Studies* 44(3), 511–18
The origin of bubbles

Appendix A Proof of Proposition 2.1

Part 1: Equilibrium verification for $v > \frac{\eta}{2}$

It is straightforward to substitute the equilibrium $P^*(\cdot)$ into (2) and show that it satisfies the equation, and verify that the $\text{SOC} < 0$ in equilibrium.

Part 2: Equilibrium uniqueness for $v > \frac{\eta}{2}$

We prove the uniqueness by four lemmas. Let $P^*(v)$ be the equilibrium strategy for an equilibrium, and $p^*(v) \equiv P^*(v) - v$.

Lemma A.1. *Any equilibrium strategy $P^*(\cdot)$ must be such that $p^*(v)$ is bounded.*

Proof. The continuity of $P^*(\cdot)$ implies that $P^*(\cdot)$ is finite for finite v . Recall that $P^*(\cdot)$ is an increasing function. If $\lim_{v \rightarrow \infty} p^*(v) = \infty$, then $\forall \bar{B} > 0$, there exists θ such that $P^*(\theta - \frac{\eta}{2} + \eta\kappa) - \theta > \bar{B}$. Hence $P^*(v)$ cannot be an equilibrium strategy. \square

Lemma A.2. *In any equilibrium, $\text{FOC} = 0$ holds for all traders in $(\frac{\eta}{2}, \infty)$.*

Proof. Assume in an equilibrium strategy $P^*(\cdot)$, there exists a trader v such that her $\frac{d\omega}{dp} \neq 0$. If her $\frac{d\omega}{dp} < 0$, that means she can increase her expected payoff by decreasing her selling price p . The only lower boundary for price is zero. But selling at zero is never optimal because $\theta > 0$ for sure. Hence for all trader $v > \frac{\eta}{2}$, selling at the corner solution zero cannot be optimal. Therefore in any equilibrium a trader v ’s $\frac{d\omega}{dp}$ cannot be strictly negative. If her $\frac{d\omega}{dp} > 0$, that means she can increase her expected payoff by increasing her selling price p . But there is no upper boundary for price, so $\frac{d\omega}{dp} > 0$ itself implies this strategy is not optimal. Therefore in any equilibrium a trader v ’s FOC cannot be strictly positive. In addition, those p at which $\omega(p)$ is non-differentiable cannot be in the equilibrium for the following reason.
$$\frac{d\omega}{dp} = \frac{1}{P^*(P^{*-1}(p))} \phi(P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa|v) [P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa - p] + 1 - \Phi(P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa|v).$$

We know that $\frac{d\omega}{dp} = 1$ when p is such that $P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa < v - \frac{\eta}{2}$ and $\frac{d\omega}{dp} = 0$ when p is such that $P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa > v + \frac{\eta}{2}$. When $v - \frac{\eta}{2} < P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa < v + \frac{\eta}{2}$, $\omega(\cdot)$ is differentiable ($\frac{d\omega}{dp}$ is continuous), because $P^{*'} and ϕ are both and continuous in that range. $\omega(\cdot)$ is not differentiable ($\frac{d\omega}{dp}$ not continuous) at two points: $P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa = v \pm \frac{\eta}{2}$. But p at neither point can be equilibrium strategy, because $P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa = v - \frac{\eta}{2} \implies p = P^*(v - \eta + \eta\kappa)$ and $P^{*-1}(p) + \frac{\eta}{2} - \eta\kappa = v + \frac{\eta}{2} \implies p = P^*(v + \eta\kappa)$, while equilibrium requires that $p = P^*(v)$. Hence FOC=0 holds for all traders $v > \frac{\eta}{2}$ in equilibrium. $\square$$

Lemma A.3. *If $p^*(\cdot)$ is not a constant, then there exist \underline{v} and \bar{v} such that $p^*(\bar{v}) > p^*(\underline{v})$ and $p^{*'}(\bar{v}) \leq p^{*'}(\underline{v})$.*

Proof. Suppose otherwise. Then $\forall \underline{v}$ and \bar{v} , $p^*(\bar{v}) > p^*(\underline{v}) \implies p^{*'}(\bar{v}) > p^{*'}(\underline{v})$. But the continuity and differentiability of $p^*(\cdot)$ implies that $p^*(v)$ will diverge when $v \rightarrow \infty$, which contradicts Lemma A.1. \square

Lemma A.4. *$p^*(\cdot)$ is a constant.*

Proof. Suppose otherwise, and let \underline{v} and \bar{v} be such that $p^*(\bar{v}) > p^*(\underline{v})$ and $p^{*'}(\bar{v}) \leq p^{*'}(\underline{v})$. Consider the FOC in Equation (2). The hazard rate of bursting at $P^*(\underline{v})$ is $\frac{1}{P^{*'}(\underline{v})} \frac{\phi(\underline{v} + \frac{\eta}{2} - \eta\kappa | \underline{v})}{1 - \Phi(\underline{v} + \frac{\eta}{2} - \eta\kappa | \underline{v})}$ and the hazard rate of bursting at $P^*(\bar{v})$ is $\frac{1}{P^{*'}(\bar{v})} \frac{\phi(\bar{v} + \frac{\eta}{2} - \eta\kappa | \bar{v})}{1 - \Phi(\bar{v} + \frac{\eta}{2} - \eta\kappa | \bar{v})}$. For both uniform and exponential prior cases, $\frac{\phi(\underline{v} + \frac{\eta}{2} - \eta\kappa | \underline{v})}{1 - \Phi(\underline{v} + \frac{\eta}{2} - \eta\kappa | \underline{v})} = \frac{\phi(\bar{v} + \frac{\eta}{2} - \eta\kappa | \bar{v})}{1 - \Phi(\bar{v} + \frac{\eta}{2} - \eta\kappa | \bar{v})}$. However, $P^*(\underline{v}) - (\underline{v} + \frac{\eta}{2} - \eta\kappa) = p^*(\underline{v}) - \frac{\eta}{2} + \eta\kappa$ and $P^*(\bar{v}) - (\bar{v} + \frac{\eta}{2} - \eta\kappa) = p^*(\bar{v}) - \frac{\eta}{2} + \eta\kappa$, with $p^*(\bar{v}) - \frac{\eta}{2} + \eta\kappa > p^*(\underline{v}) - \frac{\eta}{2} + \eta\kappa$. Hence the FOC in Equation (2) cannot be satisfied for both trader \underline{v} and \bar{v} , a contradiction. \square

Part 3: Equilibrium verification and uniqueness for $v \leq \frac{\eta}{2}$

For a trader with signal $v \in [-\frac{\eta}{2}, \frac{\eta}{2}]$, the support of her posterior belief about θ is $[0, v + \frac{\eta}{2}]$, which is different than $[v - \frac{\eta}{2}, v + \frac{\eta}{2}]$ (the support of a trader with $v > \frac{\eta}{2}$). We show that, for a trader in $[-\frac{\eta}{2} + \eta\kappa, \frac{\eta}{2}]$, her strategy $P_s^*(v) = v + \frac{\eta}{2} - \eta\kappa + B$ can still be derived from FOC of the expected payoff. Denote $\Phi_s(\theta|v)$ the CDF of her posterior belief and $\phi_s(\theta|v)$ the corresponding p.d.f. In particular, $\Phi_s(\theta|v) = \frac{\Phi(\theta) - \Phi(0)}{\Phi(v + \frac{\eta}{2}) - \Phi(0)} = \frac{\Phi(\theta)}{\Phi(v + \frac{\eta}{2})}$ and $\phi_s(\theta|v) = \frac{\phi(\theta)}{\Phi(v + \frac{\eta}{2})}$, then $\frac{\phi_s(\theta|v)}{1 - \Phi_s(\theta|v)} = \frac{\phi(\theta)}{\Phi(v + \frac{\eta}{2}) - \Phi(\theta)} = \frac{\phi(\theta|v)}{1 - \Phi(\theta|v)}$. Hence the first order condition is exactly the same, which gives the same strategy. The proof of uniqueness in Step 2 goes through in this case as well so the uniqueness follows. For a trader in $[-\frac{\eta}{2}, -\frac{\eta}{2} + \eta\kappa)$, the bubble cannot burst before trader $v = -\frac{\eta}{2} + \eta\kappa$ sells (recall that $P_s^*(\cdot)$ is continuous and strictly increasing), hence she will sell together with trader $-\frac{\eta}{2} + \eta\kappa$. Actually all traders within $[-\frac{\eta}{2}, -\frac{\eta}{2} + \eta\kappa)$ will sell together with trader $-\frac{\eta}{2} + \eta\kappa$. To see why traders within $[-\frac{\eta}{2}, -\frac{\eta}{2} + \eta\kappa + \varepsilon)$ will not sell together at a price higher than $P_s^*(-\frac{\eta}{2} + \eta\kappa)$, notice that trader $v = -\frac{\eta}{2}$ knows for sure that $\theta = 0$ and she will not join the rest, since more than $\eta\kappa$ mass of traders selling together will make her worse off, so she has incentive to preempt.