# Systems and Markets: Instability and Irrationality

K. Giles*, C. Martin*, R. Martin**, X. Lin*

**Abstract-- New financial products are difficult to price. Often the products suffer through an initial period of price volatility as the market searchers for an equilibrium value. In this paper, we extend our previous results by showing that the rate of convergence is extremely slow, less than sqrt(N), where N is the number of draws. As a result with a sufficiently low discount rate, an investor gambling on draws from the urn faces possibility of time-relevant unbounded losses. Why do these losses fail to prevent the market from forming? Are investors then inherently irrational? Because the market price is an average of the beliefs of many investors rather than a reflection of the truth. Therefore, any investor who believes their priors better than the market's also believes they have unbounded earning potential. As a result, the market thrives, volatility persists, and some investors win at the expense of others.**

## I. INTRODUCTION

Consider the problem of a single investor betting against an urn holding an unknown distribution of black and white balls. The investor assumes a prior distribution over the distribution of balls in the urn. He updates his prior using Bayes rule with each draw. Before each draw, the investor places a bet on either black or white. The investor may bet on either black or white and is given odds that correspond to his current priors. The odds the investor faces lie somewhere between his beliefs and the true odds. Therefore, according to the investor's beliefs, the bet has a weakly positive expected gain.

* Texas Tech University, Lubbock, Texas: ** Federal Reserve Board of Governors, Washington DC (email: robert.f.martin@frb.gov): The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors or of any person associated with the Federal Reserve System.

Yet, the investor following this strategy faces an unbounded expected loss. Even with discounting larger than is normally accepted in macro models (discounting implying a risk free rate greater than 20 percent) the discounted expected value of playing the game is negative. The game yields an expected utility loss independent of the curvature of the agent's utility function. In other words, even risk loving agents would not play this game willingly. Still, new markets form.

New markets seem to be an integral part of the world. From the invention of money to the latest exotic derivative contract, new products and new markets are constantly created. Even when the driving process for these new products is understood (think of a contingent claim on a real underlying asset), the parameters of the process are not known and must be estimated. Given a long time series of returns data, this estimation is not easy and uncertainty over both mean and variance is inevitable. Investors are aware of the difficulty and know the expected loss when playing against the urn.

New markets form because investors do not play against the urn, but rather, play against other investors; investors who themselves do not know the distribution within the urn but are themselves using Bayesian updating to learn the true distribution. Competing against other players rather than the truth, improves the odds: an investor does not have to be right; he simply has to be better. Optimistic or risk loving investors will enter the market. Because losses are unbounded for the loser, gains are unbounded for the winner.

Agents must play an additional game to determine the betting odds. Although any price between the priors of the agents will work, we choose the market price that represents the median threshold price. In this

ISBN 978-963-311-370-7

sense, the market price reflects the average information in the market and is efficient.

Market price swings are characteristic in this setup. The price swings occur as investors' priors are overly influenced by early draws from the urn and as some investors leave the market as their funds run low or as they stop desiring trade at the market price. We find that limiting the size of investor positions reduces the size of the initial volatility.

Before an investor can trade, he must form beliefs over the potential return. For concreteness, focus on the case of a derivative security. Since this is a contingent claim with no intrinsic value, not only do investors need beliefs over the return structure but investors must sufficiently disagree over this structure to generate trade. Any two identical investors who agree over the return structure will both desire the same side of the transaction and the trade will not occur.

Most of the time, new markets behave normally. On occasion, however, these markets either boom or bust for a period of time and often these initial swings are corrected over time. What causes these swings? Many observers attribute the swings, or bubbles, to irrational behavior on the part of agents. In a strict sense, this assertion must be true: many trades occur far from the equilibrium price. Agents that know the equilibrium price will not trade at these prices, but what happens when that equilibrium price is not knowable?

We construct a trading model where agents use all of the information available each period to update their beliefs over the return structure. In the model, agents with different priors bet on draws from a statistical urn. Agents with optimistic priors take long positions and agents with pessimistic priors take short positions. With each draw from the urn, agents update their beliefs using Bayes rule. The distribution of beliefs converges to the true distribution of the urn but may take a large number of observations to do so.

This paper explores the properties of the transition path. The unconditional model shows wide price swings in the first several periods as investor's initial diffuse beliefs are easily moved by the draws from the urn; not knowing the true distribution the first few draws contain substantial information, although this information is often misleading. When the first several draws are tilted toward either pessimistic or optimistic beliefs, the movement in price is extreme. For example, an initial run of 4 (either black or white), can more than triple the initial price. Yet, even when the two outcomes are equally likely, this event occurs over 12 percent of the time.

Importantly, with Bayesian learning, the ordering of initial beliefs is weakly preserved independent of the initial draws. In an environment with agents tilted towards optimistic beliefs, a necessary condition for overpricing, the most optimistic agents always lose on average over the course of the game. They lose because the payoff they are betting on is too low relative to the objective beliefs. If the traders also face budget constraints, the market can instantly implode as optimistic traders can no longer trade and hence no longer influence prices. The only traders remaining have pessimistic beliefs.

## II. LITERATURE REVIEW

Many papers have studied the importance of learning in macroeconomic systems. The majority of this literature has focused on finding stability under learning, typically within a rational expectations framework. [6] extends this work to Bayesian learning. They conclude that sufficiently sophisticated Bayesian learning schemes do not prevent stability; the economy under Bayesian learning converges to the rational expectations equilibrium and does so faster than under other learning rules such as recursive least squares. Although Bayesian economies converge to the same point, the transition path can be much different, even when agent's priors are centered at the rational expectations equilibrium.

Because our interest is in Bayesian updating alone, we abstract from a formal economic model. However, the implications of our results carry over to a broad class of macro models including the linearized models found in [3] This paper studies the problem through the lens of a traditional urn problem.

An agent must decide whether or not to bet, at posted prices, on draws from the urn. We show that investors must have sufficiently optimistic priors on the probability of success to enter the market. We then explore different methods by which the investor updates his priors over time. We find conditions on the initial draws from the urn which lead to immediate market collapse, stable markets, and boom markets. For boom markets, we then give further conditions for either subsequent stability or subsequent collapse.

## III. THE MODEL

The model is quite simple. A large number of agents must decide whether or not to bet on the outcome of a draw from an urn and conditional on betting they must decide black or white. The agents have no information over the distribution of black and white balls within the urn and will trade whenever they expect to profit from the transaction.

### *A. The Urn*

Let U be an urn filled with a mixture of white and black balls and let B be the number of black balls and W be the number of white. The true probability of drawing a black ball is

$$r^T = \frac{B}{B+W}.$$

Both B and W are large, integer random numbers and $r^T$ has support from 0 to 1. Each period, a single ball is drawn from the urn and the color of the ball is public information. The balls are drawn with replacement.\

*Bayesian Updating*

Let $f_{it}(r)$ be the subjective probability density of the of the percentage r of black balls in the urn at time t by agent i. The density at time t is a function of the history of draws between time 1 and time t as well as the agent's prior beliefs denoted, $f_{i0}(r)$. The agent updates the subjective density function following each draw from the urn using to Bayes rule.

Let n be number of black balls drawn and m the number of white balls drawn from the urn between time 0 and date t Define the object I as follows: $I_{it}(r|n,m) \equiv r^n(1-r)^m f_{it}(r)$. I adjusts the agents current density using the information of the number of black and white balls drawn as of time t. Notice, I can be computed recursively. The density is updated according to the following:

$$f_{it+1}(r|n,m) = \frac{I_{it}(r|n,m)}{\int_0^1 I_{it}(r|n,m)dr}$$

The denominator is a scaling constant that the new density is a probability density.

The updating multiplier, $r^n(1\text{-}r)^m$, in front of last period's density has the effect of shrinking the portion of the density that is inconsistent with the previous draws. For example, after 5 whites and 5 blacks are drawn, the weight this multiplier attaches to an r value of 0.9 is on the order 10e-06. The majority of the agent's new density must then be to the left of 0.9. For the same reasons, initial runs have an undue influence on the distribution. This is true even from the very first draw.

### *B. Initial Beliefs*

The initial beliefs of the agent determine the expected dynamics of the game. Agents have no knowledge of the true distribution of white and black balls with in the urn. As is shown by, any unbiased estimate of r^{T} must be centered at 0.5. Therefore, we expect that, on average, reasonable priors must be centered at 0.5. Indeed, given the properties of the game and given risk neutral agents, we can derive the optimal initial beliefs. They have the following form.

In principle, given that there is limited initial information. Agents on average should share these beliefs. However, if every agent has the same priors, their subjective beliefs are identical and they do not trade. Every agent weakly prefers the same side of the gamble and the market cannot form. Of course, since the agents are also indifferent between black and white at the market price, they could randomly assign themselves to black and white and create a market. We assume there is either an ε cost of trading or that agents are slightly risk averse, precluding this solution.

Therefore, for a market to form, some agents must have priors sufficiently different from the average person to make trade worthwhile. That is, some group of agents must form beliefs relatively tilted towards black. We label these arbitrarily as optimistic beliefs. These agents will always end up betting on black.

*C. The Agent's Problem*

The agents are very simple economic actors. They seek only to maximize expected winnings at each date taken as given their beliefs over r. The agents are subject to a potentially binding budget constraint below which their cumulative winnings may not fall. Each agent solves the following problem at each date t:

$$\max_G \int_0^1 rP_{Bt}Gf_{it}(r)dr - \int_0^1 (1-r)P_{Wt}Gf_{it}(r)dr$$

$$st$$

$$A_{it} > -\underline{A}$$

$$A_{i(t+1)} = A_{it} + Winnings_t$$

$$G = \left\{ \begin{array}{c} 1 \text{ bet on Black} \\ -1 \text{ bet on White} \end{array} \right\}$$

$P^B$ and $P^W$ are the payoffs associated with black and white at each date. $A_{it}$ are the assets of agent i at date t.

The agent bets each period so long as their assets are above the lower bound and so long as the max function is strictly positive. Any agent that does not perceive a gain sits out in that period. The only purpose of this assumption is to ensure that markets do not form when all agents have the same beliefs. We could formalize this assumption by adding either an ε cost of trading or by making the agents slightly risk averse. Neither of these assumptions would change the game in any fundamental way.

*D. Equilibrium*

We find the equilibrium payoff period by period as the payoff that makes the average agent in the economy indifferent between betting Black and betting White. That is when the maximization in the agent's problem above is exactly zero when G=1 for the average agent.

Then, after normalizing $P^W$=-1, the equilibrium value of $P^B$ can be found by solving the following equation:

$$\int_{i\in I}\left\{\int_0^1 rP_{Bt}Gf_{it}(r)dr - \int_0^1 (1-r)P_{Wt}Gf_{it}(r)dr\right\}di = 0$$

In many of our setups other equilibrium payoff could exist, but we assume that a market maker enforces this particular equilibrium. For example, with two types of agents, any price between the priors of the optimistic and pessimistic agent would suffice. Only with a continuum of agents is the equilibrium payoff necessarily unique.

## IV EXPECTED DIVERGENCE

*A. The Mean*

We are interested in the mean value of the posterior estimates of the percentage of white balls in an urn when the prior was given by a normal-like distribution,

$$e^{\frac{(x-\mu)}{\sigma}^2}$$

The mean is difficult to compute and so in this paper we will approximate this by a binomial distribution

$$kx^\alpha(1-x)^\beta$$

and will calculate the expected value of the quantity

$$E(w,b) = \frac{\int_0^1 x^{w+1}(1-x)^b x^\alpha(1-x)^\beta dx}{\int_0^1 x^w(1-x)^b x^\alpha(1-x)^\beta dx}$$

These integrals can be evaluated in closed from using the fact that

$$\int_0^1 x^n(1-x)^m dx = \frac{n!m!}{(n+m+1)!}$$

Thus we have

$$E(w,b) = \frac{\frac{(w+\alpha+1)!(b+\beta)!}{(w+b+\alpha+\beta+2)!}}{\frac{(w+\alpha)!(b+\beta)!}{(w+b+\alpha+\beta+1)!}} = \frac{w+\alpha+1}{w+b+\alpha+\beta+2}$$

This is the value after n draws where w white balls have been drawn. If the true percentage of white balls is given by p then the average value of E(w,b) is given by the sum

$$E_p(E(w,b)) = \sum_{w=0}^n p^w(1-p)^b \binom{n}{w} E(w,b).$$

We begin by letting b=n-w in the sum and in the expression for E(w,b). After doing this and simplifying we have

$$E_p(E(w,b)) = \sum_{w=0}^{n} p^w(1-p)^b \binom{n}{w} E(w,b)$$
$$= \frac{(1-p)^n}{n+\beta+\alpha+2} \sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} (w+\alpha+1)$$

We evaluate the two sums

$$\sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} w$$

And

$$\sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} (\alpha+1)$$

separately. We begin by just noting that

$$\sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} (\alpha+1) = \left(\left(\frac{p}{1-p}\right)+1\right)^n (\alpha+1)$$
$$= \frac{\alpha+1}{(1-p)^n}.$$

to evaluate the first sum we let

$$F(t) = \sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} t^w$$

and note that the first sum is given by

$$f'(t)|_{t=1}$$

Now we have

$$F(t) = \sum_{w=0}^{n} \left(\frac{pt}{1-p}\right)^w \binom{n}{w}$$
$$= \left(\frac{pt}{1-p}+1\right)^n$$
$$= \left(\frac{p(t-1)+1}{p-1}\right)^n$$

Now evaluating the derivative at 1 we have

$$\sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} w = \frac{np}{(p-1)^n}$$

$$E_p(E(w,b)) = \frac{(1-p)^n}{n+\beta+\alpha+2} \sum_{w=0}^{n} \left(\frac{p}{1-p}\right)^w \binom{n}{w} (w+\alpha+1)$$
$$= \frac{(1-p)^n}{n+\beta+\alpha+2} \left(\frac{np}{(p-1)^n} + \frac{\alpha+1}{(1-p)^n}\right)$$
$$= \frac{np+\alpha+1}{n+\beta+\alpha+2}$$

We are now prepared to play the game. We will bet E(w,b) cents that the next ball drawn will be white. If we win we will receive 1 dollar and nothing if we lose. Thus our expected gain is

$$p(1-E(w,b)) - (1-p)E(w,b) = p - E(w,b$$

Now averaging this over all possible draws we have the sum

$$E_n = \sum_{w=0}^{n} p^w(1-p)^b \binom{n}{w} (p - E(w,b))$$
$$= p - \frac{np+\alpha+1}{n+\beta+\alpha+2}$$
$$= \frac{p(\beta+\alpha+1)-\alpha-1}{n+\beta+\alpha+1}$$

The total losses are given by the sum

$$\sum_{n=0}^{\infty} E_n$$

Thus we see that our total loses do not converge but diverge very slowly at the rate of (1/n). Thus our losses are unbounded.

## V. AN EXAMPLE

For ease of exposition, assume half of the agents are endowed with priors centered at 0.5 and half of the agents have beliefs centered at 0.75. Given random fills of the urn, the agents with beliefs centered at 0.5 will, against an average urn, drive the optimistic agents assets to the lower bound. This occurs because the market clearing price in the initial period will imply a true probability of 0.625. Notice, that although this is an equilibrium price, it is not at all informative of the true distribution of the urn. This feature will be a consistent outcome of the model.

With these priors, we compute the initial equilibrium payoff for black using the equilibrium equations. Since the number of agent types is discrete, we may compute the price explicitly:

$$\left\{\begin{array}{l} 0.5\left\{\int_0^1 rP_BGf_o(r)dr - \int_0^1(1-r)P_WGf_o(r)dr\right\} \\ +0.5\left\{\int_0^1 rP_BGf_p(r)dr - \int_0^1(1-r)P_WGf_p(r)dr\right\} \end{array}\right\} = 0$$

where we have substituted o for the optimistic agent and p for the pessimistic agent and dropped the subscript t to ease exposition.† The 0.5 in front of each bracket indicates the relative population weights of the

† For the purposes of this example but not in the computations of equilibria described below, we ignore the truncation of the distributions at 0 and 1. We take the center of the distribution as the mean.

two groups. Integrating over each and substituting for $P^W$, we have:

$$0.5\{.75P_B-.25\}+0.5\{.5P_B-.5\}=0$$
$$.375P_B+.25P_B-.125-.25=0$$
$$.625P_B=.375$$
$$P_B=0.6$$

So, the equilibrium payoff under the initial priors is 0.6. An observer with knowledge of the true distribution of the urn, or perhaps an econometrician viewing the game with the benefit of hindsight and many draws from the urn, would puzzle over the seeming discount of betting Black. Yet at time zero, conditional on the priors, this is an equilibrium price which is simply reflecting the priors of the agents.

In the two agent case, neither agent has expected value of zero: both have positive expected winnings. The pessimistic agents will bet on white with expected return:

$$\{-.5P_B+.5*1\}=.2$$

and the expected winnings of the optimistic agents is:

$$\{.75P_B-.25*1=.2$$

Clearly, this trade should occur in equilibrium as both agents believe they have a positive expected payoff under their subjective beliefs.

This example aptly illustrates the multiple payoffs consistent with trade occurring. Pessimistic agents are willing to bet White for any payoff, P_{B}, less than 1. Optimistic agents are willing to bet black for any payoff, P_{B}, greater than. 0.33.

Let's assume, entering the $13^{th}$ period as our prior that

$$f(r)=K\exp\left(-\frac{(r-.75)}{.1}\right)^2$$

The denominator of the updating formula is just a constant so we don't need to worry about it and the constant k likewise divides out. Let’s assume now that after 13 periods we have drawn 10 black balls 3 white balls. Our new distribution is

$$f(r|10,3)=C_{12}r^{10}(1-r)^3\exp\left(-\frac{(r-.75)}{.1}\right)^2$$

We now differentiate this and set it equal to 0 and solve for r. Let’s say we find that the solution to the problem is r=0.78. In this case, the draws from the urn have caused us to update our prior upwards. In the agent's problem above, if the market interest rates does not change much, the investor would increase the size of his bet between period 13 and period 14.

## VI. SIMULATIONS

In this section, we simulate the results of the game played over a large number of draws. The two players are endowed with different priors over the contents of the urn: one player believes the urn to be filled primarily with black balls, the other believes either that the urn is filled with an even number of balls (that is, his beliefs are centered around the truth) or that there is a large number of white. In the first two games, the results are biased by a run of initial draws of black. In the last game, the initial draws are unrestricted.

### *A. Mean Value versus Optimists when the Initial Runs are Restricted*

Figure 1a shows the evolution of beliefs in a game where the first five draws from the urn are restricted to be black. The top panel shows the initial distributions. True-value agents have beliefs centered around 0.5, the actual odds of drawing black. Optimistic agents believe most of the balls are black and have beliefs centered around 0.85. Both agents are endowed with diffuse beliefs.

The second panel shows the beliefs after the first five draws. Because these draws are all black, both agents shift their beliefs towards 1. The initial draws move the agents beliefs substantially; however, the ordering of beliefs is preserved as it is under any sequence of draws.

The final shows the beliefs after 100 draws. After this many draws, the beliefs are no longer diffuse and about 99 percent of the mass lies between 0.51 and 0.62. Despite observing 95 draws from the true distribution and despite the evenness of the true odds, the agents place essentially zero probability on the true value of the distribution. This will be reflected in the market prices.

Figure 1a

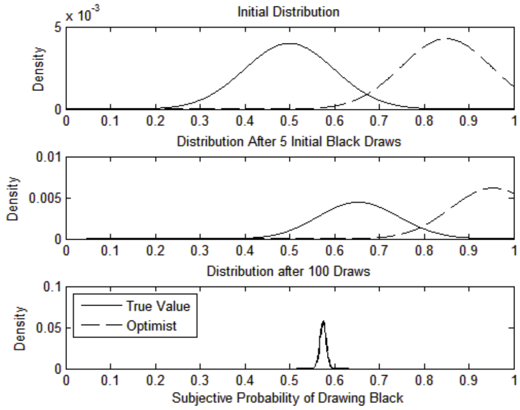


Figure 1b shows the evolution of subjective payoffs and cumulative earnings. The top panel shows the subjective expected payoff of each agent when betting on black at the market price. The beliefs converge over time. But, for any finite number of draws, optimists place a higher likelihood on black than do true value agents. The bottom panel shows the cumulative winnings of each agent. The optimistic agents make a lot of money in the first 5 periods as they bet on black. Thereafter, however, the optimist's winnings quickly turn negative as the market price, which places higher odds on black than white, works against them.

Figure 1b

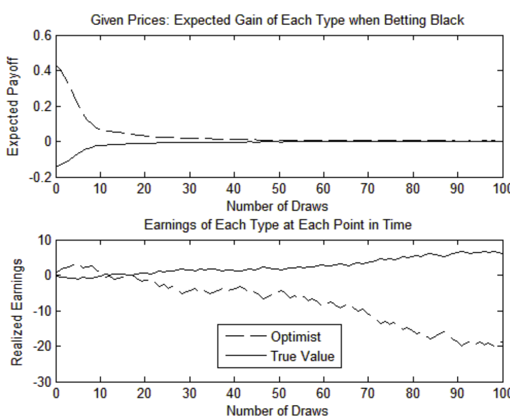


Finally, figure 1c shows the evolution of the market price, which is constructed as the average valuation between the two agents. The initial price reflects the pull of the optimists and the rapid increase in the price of black reflects the pull of the initial draws. The price then only converges very slowly to the objective value of zero. After 100 periods, the price is still more than ten percent above the true value.

Figure 1c

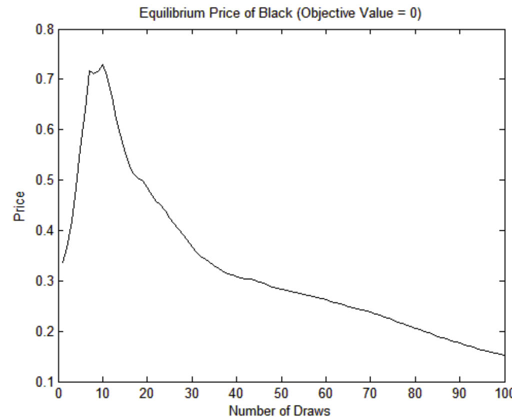


*B. Pessimists versus Optimists when the Initial Runs are Restricted*

Moving the initial beliefs of the true-value agents towards zero does not substantively change the game. Figures 2a through 2c show the same information as in figure 1. The agent's beliefs after 100 periods still remain centered nearly 0.56 and the agents still place essentially zero weight on the true value of 0.5. The main difference between the two scenarios lies in the market price. In this case, initial increase in the market price is slightly muted as the pessimistic agents weigh on the value.

Figure 2a

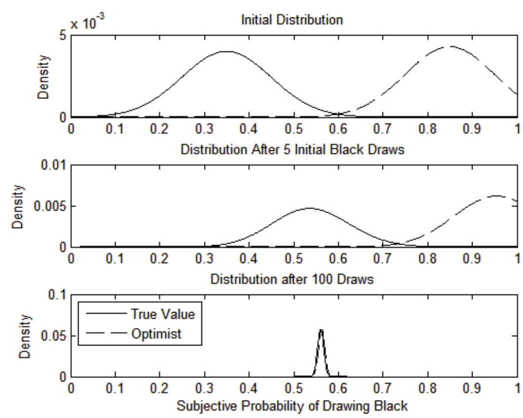

Figure 2b

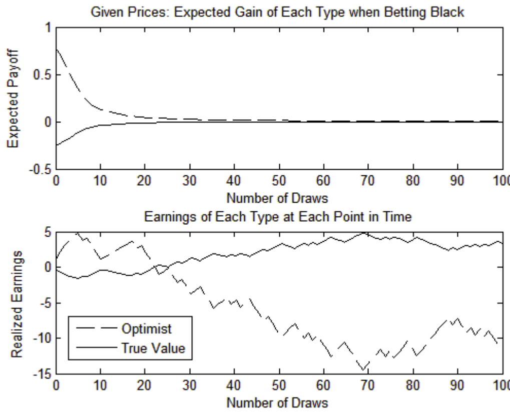


Figure 2c

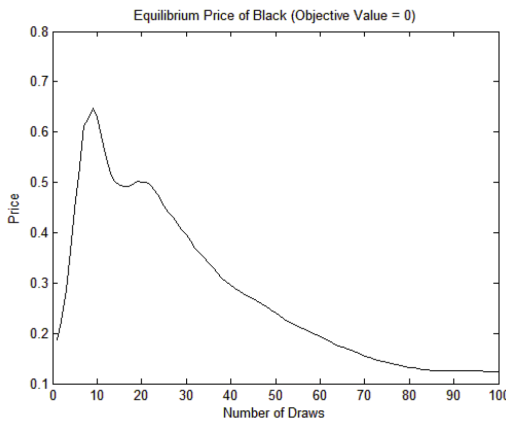


## Conclusion

Agents who believe their beliefs are willing to play an unfair game. Agents all know that if the game is played against agents with superior beliefs, they will lose. However, every agent believes their beliefs are closer to the truth and hence are willing to play. Their willingness to play keeps markets alive for a very long time, and with Bayesian updating, the market price never converges to the truth.

## Bibliography

[1] Bray, M., and N. Savin. (1986) "Rational Expectations Equilibria, learning, and Model Specification," Econometrica 54, pp 1129-1160.

[2]Bjork, Toma, and Agatha Valentian Mrugoci. “A General Theory of Markovian Time Inconsistent Stochastic Control Problems.” Working paper

[3] Bullard, James B., George W. Evean, and Seppo Honkapohja. (2007) "Near-Rational Exuberance," FRB of St. Louis Working Paper No. 2004-025B.

[4] Bullard, James B., George W. Evean, and Seppo Honkapohja. (2007) "Monetary Policy, Judgement, and Near-Rational Exuberance," FRB of St. Louis Working Paper No. 2007-008A.

[5] Bullard, James B., and Aarti Singh. (2007) "Learning and the Great Moderation," FRB of St. Louis Working Paper No. 2007-027A.

[6] Bullard, James B., and Jacek Suda. (2008) "The Stability of Macroeconomic Systems with Bayesian Learners," FRB of St. Louis Working Paper No. 2008-043A.

[7] Cogley, Tim, and Thomas Sargent. (2008) "Anticipated Utility and Rational Expectations as Approximations fo Bayesian Decision Making," International Economic Reiew, 49(1), pp 185-221.

[8] Dai Pra, Paolo, and Alessandro Andreoli. “Scaling and multiscaling in financial indexes: a simple model.” Working Paper.

[9]Gombani, Andrea, and Wolfgang Runggaldier. “Arbitrage-free multifactor term structure models: explicit computations for two factors.” Working paper.