



Market Inefficiency as an Absorbing State

Joel C Gibbons
Logistic Research & Trading Co.
St. Joseph, Michigan, 49085

*“Nothing is more irrational than a rational policy in an irrational world.”
John Maynard Keynes.*

In another essay, “The Perfectly Inefficient Market,” I argued that under plausible conditions a market can continue to function in the complete absence of information flow at all. In this essay I will pursue that topic further. This point that I will make is that inefficiency is an absorbing state. Market inefficiencies devalue information. If at a point in time information trades¹ contribute a sufficiently small portion of the total variance of price, information will become irrelevant and in time the market will iterate toward the perfectly inefficient equilibrium without regard to actual information.

The market I will examine is populated by the three kinds of participants that I introduced in the earlier essay: market makers, speculators, and non-economic agents. A market maker is a profit-seeking agent that takes long and short positions in the items being traded, and that makes its profits primarily from the bid-ask spread. A speculator is also a profit-seeking agent, but it pays the bid-ask spread. Its profits come from anticipating price changes. Non-economic agents are indifferent to trading profits and losses, and their trading activity is initiated by entirely different stimuli than the prospect of gain. The essential difference between market makers and speculators is that speculators initiate trading orders while market makers respond to orders.

Information Trading.

Information trades are ones that are motivated by some insight into the expected price change of the underlying instrument. I will assume for the sake of argument that speculators as a group have insight. Each individual speculator believes, when he enters a position, that he possesses valuable insight, but that may not in fact be correct. In any case, insight is only partial, and what each speculator believes he knows is a predictive distribution of the future price at a predetermined horizon. He acknowledges from the start that his information only narrows the range of plausible outcomes, and that it generally has little or nothing to specify how long he will have to wait for his ideas to be validated. The length of the horizon is part of the information set, and differs across speculators. Thus, one speculator might think that he has an insight into a report that will be issued the next morning, while another is convinced that he has insight into the capabilities of corporate management that will be revealed in consistently superior profits for a long time to come.

In an efficient market this sort of unique insight is rewarded. Efficiency actually requires that the party that possesses unique insight be rewarded immediately and in full. When the possessor of unique insight comes to the market to enter a position based on his ideas, the efficient market theory asserts that whatever knowledge he brings is immediately revealed to everyone. The strict demands of the efficient market theory preclude anyone else earning trading profits by trading on information that has already been revealed to the market. This is not to say that the market equilibrium shift instantly

¹ Information trades are transactions motivated by ideas and knowledge about matters external to the market itself. Technical trading – trades motivated by the trading record of the market – is not a kind of information trading. Similarly, noise trading – trades motivated by considerations specific to the trader – are not information trades.

to the price that the informed speculator anticipated. We can illustrate this point by using a single issue of stock as an example. Let the stock be priced at \$20, and suppose an informed speculator has become convinced that the shares are worth \$40. When he buys some shares, the price will move higher, but not necessarily all the way to \$40. The reason is that he is informed but he is not omniscient. It is very possible that there is a rival speculator in the market who has reason to think that while the shares are perhaps cheap at \$20, they would be wildly overvalued at \$40. The bull, having taken his long position when the shares were at \$20 will note with some chagrin that the rally seems to stall out well short of his \$40 target, but he would not be surprised. Experience would have taught him that very rarely do positions work out exactly as he had anticipated.

The foregoing example makes clear that an efficient market is a realistic possibility. At one time there was a simplistic characterization of efficiency which seemed to imply that everyone already knows everything and therefore has no reason at all to every buy or sell anything. The absurdity of this characterization – the absurdity of assuming that a market was perpetually in a state that would make the market itself irrelevant – was regularly turned back on the Efficient Markets Hypothesis and offered as proof that efficiency is a logical impossibility. Markets can however be efficient, and this well-known characterization is nothing more than a canard. Efficiency only requires that all speculative profits go to those who bring new information to the market, and that no one can profit by simply exploiting the market process itself.

The role of market makers has to be looked at in this light. They in fact do profit by exploiting the market process itself. That is to say, they earn a fee for simply effecting trading without necessarily having any unique insight of their own. Now, in practice market makers do speculate in the shares in which they make markets, so all actual market makers also seek to exploit information, but in theory it is not necessary for them to do so. If market makers earn excessive profits by maintaining excessively wide bid – ask spreads, or if their actions introduce transient bubbles in prices that take some time to be corrected, then the market is not very efficient. While it is logically possible that a market would be almost perfectly efficient – that transitory price bubbles be small and disappear rapidly – that is by no means certain. We want to investigate in this essay what happens when inefficiencies begin to account for a large share of the price variance.

Partially Efficient Markets.

It has always been understood that actual markets are more or less efficient. The standard of perfect efficiency was considered to be a benchmark against which to judge the performance of markets. Initial tests of market efficiency generally supported the view that the stock market is rather highly efficient, but the statistical model used for those tests were later shown to be rather weak. Over the last two decades or more a vast body of empirical research has overthrown the initial conclusion. At the present time it would be fair to conclude that the stock market appears to be rather efficient, but by no means perfectly efficient. Similar tests applied to other markets have come to rather different conclusions, in the sense that they consistently appear to be far much less

efficient. The market for foreign exchange in particular has also been studied extensively, with the conclusion being that foreign exchange markets are not very efficient at all. This body of empirical research has been focused on characterizing the degree of efficiency on average – of estimating the correlation between unobservable “fair value” and actual market prices – and the results are very valuable for that purpose. We will take it as given that actual markets exist in some intermediary state, between perfect efficiency and perfect inefficiency. Our attention focuses elsewhere. Specifically, we ask

Is there a dynamic in a partially efficient market that drives it either toward greater efficiency or toward greater inefficiency?

The answer we propose is that

Inefficiency rewards inefficiency and punishes efficiency, and that consequently markets tend, other things being equal, to become less efficient over time.

Speculative Behavior and Efficiency.

As mentioned above, we group market players conceptually into three groups: market makers, speculators, and non-economic agents. We need not dwell on the role of non-economic agents because their actions are not motivated by what is going on in the underlying fair values of the things that are traded in the market. The truly non-economic agents, like central banks, are motivated by policy concerns that are very often adopted precisely because they defy economic logic. After all, central banks and industry regulators would hardly be needed simply to ratify what the economic agents bring about by their trading decisions. The role of governmental authority in this context is precisely its capacity to defy the tide, and hopefully to reverse it.

In my essay on the perfectly inefficient market, market makers played a central role. That is because a perfectly inefficient market is in essence one in which information has become useless and observed price behavior is simply the byproduct of a very wide bid – ask spread. In that essay I assumed, for purposes of discussion, that the market in question is one in which there is in fact no information at all. That assumption reduced that analytical problem to one of showing that price variance would not disappear. The essential dynamic is that irrational price volatility is pure risk – risk for which there is no corresponding return – and that the defensive response of market makers would actually exacerbate inefficiency.

It is not necessary however to assume that there is in fact no information flow. What is essential is that trading decisions – the decisions made by market makers, speculators, and non-economic agents alike – is not responsive to available information. A market will be perfectly inefficient if the players simply ignore the kinds of information that would, in better circumstances, result in efficient pricing.

More generally, a market will be rather inefficient if the players choose rationally to ignore information because it has become too risky to use it.

Since it is speculators – a category that includes the more familiar category of “investors” – who initiate information trades, this analysis must be directed at their incentives and their behavior. If they worsen inefficiency, market makers will react accordingly by widening their bid – spread. What we need to understand is how inefficiencies – random, transitory price shocks – devalue information.

Information and Information Trades.

Information is knowledge about the course of future events. We can usefully adopt the statistical model in which information, as it is modeled in statistics, is knowledge about the distribution of some future outcome. The quantity of information is, in effect, the degree to which the informed distribution is more precise, narrower, and the uninformed one. A familiar example, pitching pennies, illustrates this idea. Whenever you guess correctly the next throw, you keep the penny. When you guess wrong, you lose one of your own pennies. A priori, heads and tails are equally likely. Suppose however that the pennies are actually weighted to land heads 75% of the time. Information consists of knowing this actual probability of a head. It is limited information, in the sense that the next actual outcome is still unpredictable, but over time one can be fairly sure to make money by always betting on heads.²

Another example of information brings the concept closer to a financial perspective. The stock market encompasses thousands of individual corporations. We naturally think of the shares of a single company in terms of two factors: the expected price change of the shares over some period of time and the expected variance of actual outcomes around the expectation. The actual results for any given company are uncertain, and the variance quantifies just how uncertain they are. The company is modeled statistically as the mean and variance of its specific return distribution.³ When we look at all the companies together, another source of uncertainty or variation arises: variation between different companies. When we have gained specific insight about a particular firm, it is the variance of the firm-specific distribution that applies. When we merely survey the whole field of firms without having any information that is specific to one of them, then all of them appear to be the same. The variance of outcomes for any particular firm in that case is the sum of the specific variance of a typical firm and the variance across firms. To gain information about a single firm is, then, to differentiate it from the rest and to specialize the predictive distribution of its outcomes.

² It is not absolutely certain that one would make money, because depending on how much capital one has initially to gamble with it is possible to lose the entire stake before the law of averages comes to the rescue.

³ There is a third set of parameters, which consists of the correlations between the results for the particular firm and external, conditioning events. Among the correlations is the set of correlations between the particular firm and all other firms. These correlations play a central role in how one gains information about the firm, but it is not central to the point we are making here.

Fundamental research about a business or an industry fits neatly into this pattern. When we start to analyze the firms in a given industry, all of them appear to be the same. That is to say, while we know that firms differ, we have no specific information about these firms that would identify the more successful from the less successful. There is no alternative but to treat them as equals until we get a clearer picture of their differences. If we were to try to project the outcome of a given firm we would have to allow that it might be the most productive firm in the industry and that moreover it might be especially favored. On the other hand, it might be consistently the worst laggard, and suffer bad luck over the horizon to which we are making projections. There is an unpredictable element that no amount of research will anticipate, but we can narrow the range of projected outcomes by discerning where the firm ranks on the spectrum from most productive to least productive in the industry.

Information is by its nature incomplete. It never happens that the even the most informed speculator would know that future with perfect certainty. Some variance of outcomes remains, and it has two implications. First, of course, is that the speculative position will actually turn out to be more or less profitable than expected. Second, and for our purposes more importantly, the speculator can't know in advance how reliable his insight is. If the results are disappointing, there is no way to know whether the disappointment was the result of bad luck or bad information. He has his sources of fundamental research, as I have outlined it, but he also confronts a constant flow of feedback from the market itself. He knows that his fundamental research is incomplete, and may be highly incomplete. He knows furthermore that he has no objective basis to judge how complete or incomplete it is. He must ultimately guess how valuable is his initial body of information. He must also guard himself against the worst cases. His fortunes are suspended between two competing forces: the value of the insight that he brings to the market on the one side, and the inexorable logic of gambler's ruin on the other.

The law of gambler's ruin is a mathematical law which states that almost surely anyone who continues to wager his limited capital will eventually lose it all. However much capital any one speculator has, it is very small compared with the combined capital of the rest of the world, and so it is appropriate to assume that he is playing against an opponent that has unlimited capital. If the speculator has no special insight at all – if he is merely guessing – then it is certain that he will eventually lose all his capital. If he has some real insight, he will at least survive longer, and if his edge is greater than some threshold he can expect to survive indefinitely. His edge has to be quite large, as measured in units of expected profit per play of the game, for that to happen. The typical speculator is understandably, and perhaps justifiably, convinced that he has some valuable knowledge, and that he has therefore a positive expected profit per trade, but if he is realistic, he accepts that he is unlikely to defy the law of gambler's ruin indefinitely. It is advisable for him to try to supplement his trading offensive – which utilizes his knowledge – with a solid risk-management defensive. The defensive – various forms of stop-loss tactics – must necessarily be independent of his offense. He must be willing to take defensive measures even against the counsel of his trading knowledge because the very purpose of the defense is to prevent an over-commitment to his beliefs and insight.

If he places stops at points that seem to be a good place to get out based on his knowledge, then he is not reducing his risk, he is doubling it.

These observations apply even in a perfectly efficient market. Even in such a market, any given speculator will win if he possesses a great deal of information or he will at least survive for some period of time before his capital is gone. If inefficiencies are present, then the odds shift radically against even the most insightful speculator. Two forces are at work to diminish the value of information, and or even to make it a liability. The first is that the speculator can not be certain of the accuracy or reliability of his information and the conclusions he has drawn from it. When price movement is influenced by a transitory price run, he has no way of knowing if this is merely a transient event, or if it is evidence that his assumptions are wrong. Suppose, for definiteness, that the speculator is long, and that the price is rising much faster than he had anticipated. He knows that either he underestimated how good the news would be – that is to say, actual developments have exceeded the prior expected value – or he is witnessing a transient price bubble. Experience has taught him however that it is wise to take a profit because profits often prove to be fleeting. On the other side, if the price has fallen he again cannot know if he is witnessing a permanent or transitory event. In both cases, the price path frequently, or even generally, reflects outcomes that were not contemplated by his initial information. He is forced rationally to question his prior ideas.

At the same time that the speculator is depreciating the accuracy of his information, the variance of outcomes is forcing him to tighten his risk management, to reinforce stop – losses and to tighten profit-taking. This is prudent and necessary because of the logic of gambler's ruin, and it operates even if the speculator refuses to let his commitment to his initial research weaken. It happens because the ratio of expected price variance to total variance is lower. To summarize, transitory – but often prolonged – price bubbles cause the speculator to lose confidence in his information while at the same time causing him to apply risk management methods more aggressively. His expected return from entering a position is no longer the expected value given that what he knows is true and accurate, because of the cost of risk management and because of the cost of giving up on positions that are in fact justified by all available information. But though his expected value is lower, the variance of outcomes is greater. Depending on the magnitude of the transitory surges in relation to the variance of the permanent component of price, the speculator reaches a point where even though his ideas are sound, the expected value of his trades is negative. He simply cannot afford to utilize his knowledge.

This logic operates at the level of the individual, but when we aggregate across the whole market of speculators, a third ingredient become apparent, which further reduces the value of information. Since speculators are as a whole informed traders, the position taken by the average speculator is correlated with the direction of incoming information. In other words, most speculators will be positioned on the side of the market that gives the positive expected profit in light of what is known. That implies however that most speculators will at any given time and in any given market be on the same side of the trade, either long or short. Events that defy the prior knowledge and that force defensive actions therefore strike most traders at the same time! There is an

understandable exodus from positions that have lost support. This greatly enlarges the transient component of price. Many speculators, clinging to some commitment to their prior convictions, will refuse to panic. They will hold out in the belief that they are witnessing a temporary panic, and for the most part they will be right. They can expect to be right often, but they also know that they cannot afford to be wrong very often at all.

Depending on the contribution of transitory bubbles to the total variance of price, information can become a liability. Speculators attempting to survive in such a market environment will be well advised – assumed that they do not wish to simply find another career – to find technical trading methods that are intended to exploit the inefficiencies of the market. At that point, the market ceases to respond to information at all, because no one is willing to take a position based on it. In that situation the market is effectively perfectly inefficient, because while underlying market equilibrium is not constant, and there is revealing information available, it does not pay anyone to try to exploit that information. This analysis suggests the following hypothesis, which I think merits definitive modeling work:

There is a critical threshold for the variance of the transitory component of price variance, above which the market becomes effectively perfectly inefficient. The threshold is a function of the variance of the informational component of value and of the expected drift rate of price. The lower the drift rate – the lower the maximum expected return from information trading – the lower is the threshold.

This conjecture raises in a very natural way a further question, which is whether any inefficiency at all triggers a descent to this threshold. It would obviously be regrettable in the extreme if that is true, but I do not think that it is. More likely it is that there is something like a second threshold with the property that any market which initially is that efficient or more is stable. While my analysis thus far has seemed to diminish the value of information, that is not true. Information is very valuable as long as it is recognized by the market as a whole in a reasonably prompt manner. Uncertainty about how rapidly a market will adjust to even the most insightful information is as important as the risk if intervening transitory runs in devaluing information. These risks are, as it were, partners in crime. There is a further consideration that tends to enhance the value of information. In any market in which it is possible to diversify specific risk, a speculator can remain fully, or almost fully, invested while taking quite small positions in individual trading instruments. The American stock market has historically been a model of the potential for diversification.

Diversification is real risk reduction, but it requires the speculator to have valuable information about many different specific risks, because he is exposed to many specific risks. A large institution with very deep capital – a mutual fund complex for instance – could probably afford to obtain and process such a broad agenda of information, and to build a highly diversified portfolio as a result. In recent times however even this recourse has been less effective because financial assets of all kinds have become much more highly correlated than they were historically. Not only are individual share issues more correlated with each other, they are more correlated with

foreign share, with domestic interest rates, and with foreign exchange. The commoditization of equities, which is itself a byproduct of the greatly increased appetite for speculative trading, has diminished the potential defensive value of diversification, I will not attempt to pursue this point further at this time. I conclude with the observation that a speculator who has valuable information about a many alternative instruments that are not too highly correlated with each other can very probably succeed even in the presence of significant inefficiencies. It goes without saying that no speculator can succeed in a perfectly inefficient market, regardless of what trading tactics he uses.

In the forgoing analysis I have equated markets that are in some respects very different. That is to say, I treated a market that is unresponsive to information as being the same as one in which there is simply no information at all. That needs some explanation. If the prospects of the instruments that are traded in the market changes over time, and if accordingly equilibrium price evolves over time, the market cannot perpetually defy the implications of information. If, to take a concrete example, the instrument is shares of a corporation that is rather successful and earns a high rate of return on its investments, then in the very long run the shares must rise. It does not follow however that the market is even slightly efficient in practice. If the shares react to good news in a discontinuous way – if they fail to respond for a long time and then adjust suddenly and explosively – the value of information about the operations of the firm will be almost without value. This example illustrates the edge that an informed and highly diversified speculator can have. He can simply position a small block of the shares and wait for the inevitable correction. As long as he is both right about the true underlying value of the shares and refuses to surrender to the market's rejection of his beliefs, he will eventually be paid for what he knows. If his position is small enough, he can presumably wait almost indefinitely for the truth to come out. The fact remains however that in general explosive or discontinuous corrections, even if they are corrections toward equilibrium values, diminish the value of information.

The argument that I have advanced here hinges on the prospect for speculators, but the analysis is not complete until we consider the situation of the other kinds of market participants: market makers, non-economic agents, and a third category that I will designate “buy and hold” players.

Market Makers.

It is not necessary to add anything here to the analysis of market makers, above and beyond the analysis found in my previous essay, “The Perfectly Inefficient Market.” Market inefficiencies are a byproduct of their defensive tendency to widen the bid – ask spread. At any given time, market makers will quote what seem to be narrow spreads, but quoted bid – ask spreads are irrelevant for two reasons. First, they are good only for small orders, but transitory price bubbles are the result of persistent order imbalance representing large total volume of transactions. Quoted spreads are irrelevant also because they are biased. At any given point in time, the majority of orders that a market maker receives are on one side or the other. Suppose he is experiencing a run of sell

orders as speculators dump positions. If asked, the market maker will quote a bid and an ask, but they don't mean the same thing. The bid he quotes is real; he is probably going to have to buy at the price. But the ask is hypothetical. There is little interest on the part of the public in buying and so he is unlikely to be taken out no matter what ask he quotes. A few contrarians can exploit this imbalance to accumulate small positions at aggressive ask prices, and if the market stabilizes or recovers they will profit, but the profits that they earn – extracted from the market maker – are small compared with what the market maker earns from lowering his bid.

Non-economic Agents.

The situation of the non-economic agents is unchanged in many ways. They are able to continue to pursue their goals without regard to what is happening in the market. That is why they are a source of transitory price shocks. Some of the largest non-economic agents however are susceptible to a feedback of the following sort.

The very largest of the non-economic agents are the world's central banks. They pursue an agenda based on the needs and interests of the nation, or perhaps those of the governing party, that they serve. Those interests are themselves impacted by many markets, and for this reason the bankers will assume some responsibility for the success of the markets. The natural and almost inevitable result is that they will try to use their vast capital to steer the wayward market back to proper behavior. In Japan, for instance, the Treasury simply buys shares of stock in a futile attempt to increase their value. The problem with these actions is that they do not make the market More efficient – they do not make it more responsive to important information – they must make it Less efficient. Who indeed would waste time to learn anything about any of the instruments if the prices are simply going to be fixed by the central bank anyway? The presence of non-economic agents is a very difficult burden for a market to cope with, but it is not ordinarily impossible to do it. If those agents take it upon themselves to great increase their impact on market performance there is really little hope of restoring any degree of efficiency.

Buy and Hold Investors.

Do buy and hold investors make a market more efficient? That can hardly be the case, because they are actually absent from the market. The premise of buy and hold is that the world would be a better place if there were no market at all. Specifically, buy and hold has the effect of converting publicly traded instruments into private placements. We need not dwell here on a lengthy analysis of whether this approach works or not, though the rolls of corporate bankruptcy filings cannot be ignored on that point. Since buy and hold investors do not trade they actually have no impact on a market. For practical purposes, buy-and-hold is a policy for equity investors. Their importance in the context of this essay arises from how many of them there are and how much of the whole pool of outstanding shares they command. Individual investors are very patient investors, to the extent that the shares they own need to be subtracted from the tradable float of the stock

market. What remains is a much smaller market, populated by agents with much shorter decision horizons than is generally appreciated.

Conclusion.

The argument which I have advanced in this essay does not depend on any novel hypotheses or on any obscure mathematics. It rests two familiar consequences of randomness. One consequence is that speculators, like everyone else, learn from experience. When what we know is only a probability distribution of outcomes, it is impossible to be certain in advance of what we know. Since we are not certain about what is going to happen, we must constantly grope along, seeking confirming and disconfirming evidence. When that evidence arises, we adapt our expectations to it. Since transitory bubbles and shocks are indistinguishable from disconfirming evidence, they necessarily, and indeed Rationally, shake our confidence in what we initially thought we knew. These are the terms under which we all operate in life and speculators are by no exempt from their implications.

The other consequence is summed in the rule called Gambler's Ruin. Anyone who continues to gamble has limited capital against a market that has effectively unlimited capital is very likely to lose everything. A rather large expected profit on each play of the game is needed to make repeated gambling a viable activity. Faced with this fact, a rational speculator will attempt to supplement his stock of information, his trading asset, with risk management tools intended to conserve capital.

Under these circumstances, in the face of market inefficiencies even genuine information can easily become a liability rather than an asset. Successful speculators will not rely on information, but will rely on technical trading rules that are intended to exploit the inefficiencies in a purely mechanical way. At that point, market clearing prices cease to respond to information because, quite simply, actual trading decisions do not reflect available information.