

Market Epistemology

Michael Thicke

Forthcoming in *Synthese* (2017)

Abstract

According to Margaret Gilbert's collective epistemology, we should take attributions of beliefs to groups seriously, rather than metaphorically or as reducible to individual belief. I argue that, similarly, attributions of belief to markets ought to be taken seriously and not merely as reports of the average beliefs of market participants. While many of Gilbert's purported examples of group belief are better thought of as instances of acceptance, some collectives, such as courts and markets, genuinely believe. Such collectives enact truth-aimed processes that are beyond the control of any single individual. These processes produce beliefs that are distinct from any individual belief and do not merely report the "average" or "majority" view of the group. In the case of markets, beliefs are indicated by prices, though it is often difficult to infer beliefs from prices and those inferences are almost always uncertain. Market beliefs are justified when traders collectively possess sufficient evidence, there are sufficient incentives for participants to trade based upon the evidence they possess, sophisticated traders have enough power to counter common cognitive biases, and there is no manipulation of prices. Thus, in some cases markets can know.

1 Introduction

On a cold January morning in 1986 the space shuttle Challenger launched from the Kennedy Space Center only to explode moments later. It took several months for the Presidential Commission on the Space Shuttle Challenger Accident to conclude that defective rubber "O-rings" in the shuttle's solid rocket boosters were responsible for the explosion, contrary to the press's initial speculation that the main fuel tank was to blame.¹ In this paper I will contend

1. I first learned of this example from Surowiecki (2005).

that the stock market knew the cause of the crash long before the commission announced its findings or the press identified the culprit. That is, I will argue that the stock market is a collective entity capable of forming beliefs, and in instances like those of the Challenger disaster those beliefs constitute knowledge.

We regularly encounter statements such as “The union believes that management is being unreasonable” or “Our family believes in ghosts” (Gilbert 2004, p. 95)—statements that, taken literally, ascribe beliefs to collectives. Margaret Gilbert argues that we should indeed take such statements literally, rather than as shorthand for statements like “The members of the family believe...” or “The members of the union have agreed to claim...”—statements that assign beliefs to members of the collective rather than the collective itself. In Gilbert’s account, a group of individuals commits to believe as a group, and in doing so creates a collective belief that is distinct from the beliefs of any individual. She brands the study of such group beliefs “collective epistemology.”

We also regularly encounter statements such as “the market believes Apple’s best years are behind it” (Commins and Meyer 2014) and “the implication is the market believes the bond is less risky...” (FN Arena 2014).² These statements also attribute beliefs to a collective, but to a particular collective: the market. The primary aim of this paper is to argue that epistemologists ought to take these ascriptions of collective belief seriously as well.³ It argues for “market epistemology” as a subcategory of Gilbert’s collective epistemology. The market, however, is a very different sort of collective than unions or families, and is not amenable to Gilbert’s commitment-based account of collective belief.

Considering market beliefs as collective beliefs requires a redefinition of collective belief that does not involve commitment. I side with epistemologists such as Wray (2001) and Meijers (2002) who claim that Gilbert’s examples of collective belief are better thought of as collective acceptance, but I argue that there is another class of collective epistemic products, including market beliefs and court beliefs, that ought to be considered proper cases of collective belief. Collective beliefs, I argue, are produced by truth-aimed collective processes that

2. Gilbert and those following her follow the convention of using imagined, though plausible, ascriptions of collective belief to inform their arguments. Because statements ascribing beliefs to markets are less ubiquitous than those ascribing beliefs to other collectives, throughout this paper I will be appealing to concrete examples. I believe the danger of appealing to examples that may eventually feel outdated is more than compensated for by the benefit of being able to examine language as it is actually used.

3. I am aware of one other attempt to interpret statements of market belief as collective belief: Orléan (2004). However, Orléan is an economist and his account of collective belief is unconnected to any discussions in collective epistemology. Remarkably, he seems to have independently arrived at a summative /non-summative distinction just as Gilbert did, but has defined it in an entirely different way. He defines a collective belief to be “what the majority of the group thinks the collective belief is.”

are beyond the control of any single individual. The Supreme Court of the United States, for example, produces a verdict that is based most directly on the votes of the court justices, but also on the evidence and arguments supplied by each party, court procedures, rules of evidence, and previous case law. All of these components of the court work together to produce a belief, represented by the verdict, that is distinct from the belief of any individual member of the court.

Further, I argue that markets are not only capable of believing, but are also capable of knowing. Following the familiar notion of knowledge as justified true belief, I propose a reliabilist account of market justification as a condition for market knowledge. While the work of financial economists casts doubt on the notion that *all* market beliefs are justified, considering narrower reference classes can grant justified status to some market beliefs while excluding others.

The immediate aim of this paper is to expand the domain of collective epistemology into the realm of markets. Markets have the potential to produce knowledge about the world in a fundamentally different way than individuals or other social institutions, and this paper offers a way for philosophers to discuss markets using the familiar notions of belief and justification. It offers a way to engage with market beliefs philosophically. However, its ultimate goal is pragmatic: identifying market beliefs and distinguishing justified from unjustified beliefs has the potential to inform society about a range of topics, from the likely results of elections (Rothschild 2009), to the consequences of terrorism (Arvas 2015), to the future of the climate (Hsu 2011). This work is therefore similar in spirit to that of Solomon (2001), Miller (2012), and others in distinguishing between justified and unjustified consensus in science. Analyzing scientific knowledge has long been a core project of social epistemology, and I believe it is time for social epistemologists to devote similar attention to markets.

2 Believers and Rejectionists

Gilbert's argument that we ought to take attributions of beliefs to groups at face value has been controversial even among epistemologists who are sympathetic to her overall project of rejecting a purely individualist epistemology. The two chief camps within this sympathetic group are the "believers" (Gilbert 2002; Tollefsen 2003; Rolin 2010; Brown 2015) who agree with Gilbert that group beliefs ought to be considered proper beliefs and the "rejectionists" who argue that they are not, but rather ought to be considered as a kind of acceptance (Tuomela 2000; Wray 2001; Meijers 2002). In this section I will argue that the rejectionists are correct

to claim that many instances of group belief described by the believers are best described either as acceptance or summative belief. However, some instances of group belief should be considered instances of belief proper, including those held by courts and some groups of scientists. What these instances of group belief have in common are constraining epistemic procedures that are beyond the immediate control of any of the group's members. In the following sections I will argue that market beliefs should also be included in this category.

For Gilbert, the essence of group belief is commitment. When a group "believes" something, all of its members are committed to act, when they are acting as members of the group, as if it were true. She claims, "A population, P, believes that p if and only if the members of P are jointly committed to believe as a body that p" (Gilbert 2004, p. 100). For Gilbert, a commitment is a moral obligation to act in a particular way, and commitments are created through deliberate action. Gilbert applies this interpretation both to formal bodies such as courts and unions, as well as to casual, ad hoc groups. For instance, she describes a situation where two friends, Roz and Mark, discuss the justifiability of war. Roz is wholeheartedly against war, but verbally assents to Mark's statement that defensive war is justified. Mark's declaration and Roz's verbal assent constitute commitment. At that point, Gilbert argues, either could correctly state, "We believe that a defensive war is justifiable" (Gilbert 2002, p. 42). This could be the case even if Mark believed otherwise; group belief is based on members' public commitment, not on their internal beliefs.

Gilbert acknowledges that sometimes such ascriptions of group belief are instead shorthand for something like, "Both he and I believe that...": statements that are directly translatable into statements about individuals' beliefs. These sorts of group beliefs she describes as "summative" beliefs (Gilbert 2004, p. 97). However, she disagrees with the "simple summative account" of group belief, in which ascriptions of group belief are *always* translated to mean that all or most members of a group share that belief (p. 97). Instead, she argues, such statements should by default be interpreted according to her "plural subject account": "A population, P, believes that p if and only if the members of P are jointly committed to believe as a body that p" (p. 100). Once members of a group have made such a commitment they offend against the rest of the group by making unqualified statements to the contrary. If in conversation Roz stated, "Defensive war is not justifiable," she would be, in Gilbert's account, breaking her commitment and Mark would be justifiably offended. However, Roz could state without offense, "We believe that defensive war is justifiable, although personally I don't believe it is" (Gilbert 2002, p. 44).

So-called rejectionists agree with Gilbert that there is more to group belief than the simple summative account allows. They disagree with Gilbert that these non-summative group beliefs are proper instances of belief. Instead, they argue, non-summative group beliefs ought to be considered as instances of acceptance. Acceptance potentially differs from belief along several axes, and there may even be several varieties of acceptance.⁴ However, for the rejectionists the most important distinguishing features between belief and acceptance are that acceptance is goal-oriented or pragmatic while belief is neither, and that acceptance is voluntary while belief is not. Tuomela (2000) notes, “acceptance can be intentionally performed action while belief cannot” (p. 122). According to Wray (2001), “a collective belief is adopted by a group as a means to realizing the group’s goals” (p. 319), and therefore collective beliefs are not instances of belief proper. And according to Meijers (2002), “Acceptance... is a pragmatic notion. To accept a proposition is to adopt a policy and use the accepted proposition as a premise in one’s practical deliberations about future actions” (p. 77). In the case of Roz and Mark, it seems as though their group belief is voluntary: Mark chose to make a statement and Roz chose to assent to it. Further, this belief seems to have been arrived at for reasons that are not purely epistemic: perhaps Mark made the initial statement in a mistaken effort to impress Roz and she assented because she wanted to avoid an argument. These are good reasons for interpreting Roz and Mark’s purported group belief as a kind of acceptance rather than as belief proper.

Wray and Meijers further distance themselves from Gilbert regarding the collectives to which they are willing to attribute group belief. As the Roz and Mark example demonstrates, Gilbert takes a permissive view in attributing non-summative group beliefs. Although she allows that some attributions of group belief might be summative, she claims that attributions of group belief ought to be interpreted as non-summative by default. Wray and Meijers are more restrictive. According to Meijers (2002), group belief is non-summative only in groups that form a “structured whole” (p. 73). In Meijers’ account, a group forms a structured whole rather than a mere aggregate when the group shares a “collective intention,” and group beliefs are non-summative only when they are relevant to that intention (p. 72). A group of employees might have group beliefs about a particular business strategy, but not about gay marriage, even if all of those employees agree about gay marriage (assuming gay marriage isn’t relevant to their business). With respect to the latter belief, the employees would form

4. Tuomela (2000), for instance, distinguishes between pragmatic acceptance (accepting a proposition because it is useful to do so) and “acceptance as true,” which is accepting for purely epistemic reasons. Meijers (2003) distinguishes between acceptance “for the sake of argument” and committal acceptance that has the binding nature of Gilbert’s group belief.

a mere aggregate, and their belief would be of the weak “opinion poll” (ie. summative) type rather than the stronger agreement-based type of group belief (p. 72).

Wray’s view of collective belief is even more restrictive. Like Meijers, he argues that non-summative group belief is inherently goal-oriented. Further, he argues that groups capable of forming beliefs must have what Emile Durkheim calls “organic solidarity” (Wray 2007, p. 341). The individuals within a group with organic solidarity are functionally interdependent in the way that organs in the body are functionally interdependent. No organ in the body can perform its task independently, and it makes little sense to assign a goal to any individual organ. Rather, all organs in the body must work together to accomplish the goal of the individual they compose. Similarly, the particle physicists described by Knorr-Cetina (1999) are functionally dependent and have a collective aim that is irreducible to individuals’ aims. In contrast, according to Wray, larger groups of scientists, such as “climate scientists” do not have organic solidarity, even though they depend on each other’s work in some way. This loose dependence is not sufficient for organic solidarity because there is no irreducible goal shared by all climate scientists. Scientific specialties do not have organic solidarity, and therefore, Wray argues, do not have collective beliefs. Finally, Wray claims that to have collective beliefs, groups must have some mechanism for “adopting a view” that is irreducible to that of its members. For example, a scientific paper might be said to represent a research team’s view, while reflecting the views of no individual members of the group (Wray 2007, p. 345).

The arguments against considering commitment based belief as belief proper are, in my view, decisive. Commitment is typically a voluntary action, and belief formation, in the individual case at least, is involuntary. Gilbert (2002) argues that inferring from the voluntary actions of individuals that the group itself is voluntarily believing is making the fallacy of composition: while the individuals composing the group might voluntarily commit to a belief, that does not imply that the group itself is voluntarily believing the claim. This may be so, but Gilbert offers no account of how a group, especially a group composed of just two friends, might plausibly be said to believe involuntarily while all of its members act voluntarily. The burden of proof is on Gilbert to show how beliefs are involuntarily formed by collectives when all their members commit to that belief voluntarily, not on rejectionists to prove that they cannot be. Tollefsen (2003) argues that belief formation might not be as involuntary as rejectionists imply. As a parent, for example, Tollefsen “chose to read Dr. Sears over Dr. Spock,” and as a consequence came to acquire “a radically different set of beliefs” than those she would have acquired if she had chosen Dr. Spock (p. 400). It is certainly true that people

often acquire beliefs in virtue of achieving some goal or as the result of conscious choices, but this is not what happens when a group commits to a proposition. The commitment is not a byproduct or consequence of a conscious choice, but a choice itself. There is a significant difference between choosing to read a book and therefore forming a belief about childrearing and directly choosing to believe something about childrearing. In the first case a voluntary choice leads to an involuntary belief, while in the latter the belief itself is the subject of choice. That the former occurs is uncontroversial, while the latter goes against fundamental intuitions about the nature of belief. While we may choose to set ourselves on a path that we expect will lead us to form one belief or another, the final step of actually forming that belief is beyond the grasp of voluntary action. If it weren't, there would be no need to read the book at all!

Nevertheless, Wray, Meijers, and other rejectionists are too quick to dismiss *all* group beliefs as cases of acceptance or summative belief. Believers and rejectionists share a common conception of group belief as based upon commitment. While I agree with Wray and Meijers that commitment based beliefs should be considered acceptance, I do not agree that all collective beliefs are commitment based. Consider this example of collective belief: "In the opinion of the court, this law is unconstitutional" (Gilbert 2004, p. 95). In Gilbert's account, the court believes that the law is unconstitutional in virtue of a commitment by the members of the court to collectively believe it to be so. But that is not how courts operate in practice. In the recent Supreme Court of the United States ruling extending the right of marriage to same-sex couples, Chief Justice John G. Roberts Jr. dissented, writing, "Although the policy arguments for extending marriage to same-sex couples may be compelling, the legal arguments for requiring such an extension are not. The fundamental right to marry does not include a right to make a State change its definition of marriage" (Supreme Court of the United States 2015, p. 41). Although Gilbert allows parties to a group commitment to dissent, she requires them to do so in a qualified way: "While acknowledging that she and her friends believe a certain thing, Rose might add, without a sense of fault, that she, *personally*, does not believe it" (Gilbert 2002, p. 45). Roberts does no such thing. Rather, he is "bluntly expressing a view contrary" (Gilbert 2004, p. 100) to the court's opinion, something that in Gilbert's view is considered an offense against the group. However, Roberts' dissent is entirely standard practice, and it is hard to imagine that the other justices consider him to have offended against them. Therefore it seems that Gilbert must either claim that the court's opinion is not a genuine case of collective belief (because there is no commitment), or she must abandon commitment as the basis of collective belief. As court beliefs are a paradigm case of collective belief, it is hard

Table 1: Pettit and List’s judges: individual votes. To be liable, the defendant must have caused the harm and have a duty of care.

Judge	Valid Contract?	Breach?	Liable?
A.	Yes	No	No
B.	No	Yes	No
C.	Yes	Yes	Yes

to imagine choosing the first alternative.

3 Redefining Collective Belief

I propose a non commitment based basis for group belief based on List and Pettit’s example of a breach of contract case decided by a panel of three judges (2011) . They us to imagine that the court will find the defendant liable if there was both a valid contract and that contract was breached (Pettit 2011, p. 93). Each judge, in their example, votes as in the table below:

List and Pettit observe that we could adjudicate the case in two possible ways. We could follow a “conclusion-centered procedure” based on each judge’s opinion of whether the defendant is liable, or we could follow a “premise-centered procedure” based on the majority view of each premise. In the first case, the defendant would be found not liable because judges A and B found him so. In the second, the defendant would be found liable, because the majority of judges found there both to have been a valid contract and that the contract was breached. If we follow the premise-centered procedure, the view of the court is not reducible to the view of the justices, as the court found the defendant liable while the majority of the judges would have judged the defendant not liable. In this sense, they argue, such groups have “minds of their own.”⁵

In List and Pettit’s example, the opinion of the court is not based on joint commitment. Rather, the court’s opinion is a brute fact, determined by a conjunction of the justices’ votes and the rules of the court. Once the outcome of the case is decided, the judges may act or speak however they like, and it will not affect in any way what the court, as a whole, “believes.” Further, because the court’s belief about liability does not follow straightforwardly from each judge’s belief about liability, it cannot be considered a case of summative group belief.

5. However, Pettit’s view of group belief does not fall neatly into either the believer or rejectionist position. He classifies courts and similar groups as “intentional subjects,” but claims that such a group “is incapable of forming degrees of belief and desire in the ordinary fashion of animal subjects” (Pettit 2011, p. 256).

There are three important features of List and Pettit's example that can be generalized into an account of group belief that is not based on joint commitment. First, the court produces an *epistemic product* that contains and reports the court's belief. Whether the defendant is liable is something that an individual could have a belief about. Second, there is some constrained process that produces this product. Laws and court procedures are designed to produce a verdict that is epistemically justified, and both the court justices and the public recognize that trial procedure is aimed at producing a verdict. Each individual in the trial process has a well defined role governed by explicit rules; none can simply act as they please. Third, the process is neither dictatorial nor simply summative. If two of the justices were merely advisors and the third made the final decision, or if the case was adjudicated according to the conclusion-centered procedure, then we would say that "the court believes" is just shorthand for "the chief justice believes" or "the majority of the justices believe."

Regarding the first feature, a belief, in ordinary language, is a mental state that can be true or false, correct or incorrect. Philosophers have debated just what beliefs are, how they relate to other mental states such as acceptance and desire, and how they connect to behavior. For the purposes of this paper, I take a belief simply as something that could be expressed as a proposition and that aims at the truth. If I believe that humans will walk on Mars within the next 25 years, it means that I regard it as true that humans will do so. If I make that claim while regarding it as unlikely that humans will walk on Mars within that time or having no real opinion about the matter, then I might be speaking deceptively, or I might merely accept but not believe. An epistemic product is an identifiable object that contains and reports a belief. For the US Supreme Court, this product is a published opinion, which is published in several formats, both physical and electronic. The opinion is not *identical* with the court's belief, but rather contains and reports it. The court's belief does not reside in the heads of individual justices. For individuals, beliefs presumably reside in our minds, but for the court belief resides in published documents. Even if every court justice changed their minds about a case the minute the opinion was delivered, the court's belief would be unaffected until a subsequent contrary opinion containing and reporting a new belief was published. Similarly, for a body like the Intergovernmental Panel on Climate Change, group beliefs reside in their published reports, regardless of what any individual scientist who contributes to those reports believes. For a union or committee, belief might reside in official minutes.

The second generalizable feature of the court's beliefs is that they are the result of a constrained process. Legal verdicts are produced according to very strict rules. Gilbert claims that

groups can form beliefs involuntarily while all of their members act voluntarily, but offers no explanation of how this might occur. The constrained process criterion allows for this division between individual and group intent. It is the rules, implicit and explicit, of a collective such as the court that channel the efforts of individuals (each with their own objectives) toward the group's aim. The group's aim is encoded in the procedures of the collective; it is not a function of the aims of any of its members. If this aim is to produce a true proposition, and it operates as expected, then the result will be a collective belief.

The constrained process criterion disqualifies many of Gilbert's own examples as instances of proper group belief. For instance, she cites, "Our family believes that the best way to spend Christmas is at the beach" as an example of collective belief (Gilbert 2002, p. 25). It is *possible* that this family has instituted a procedure for getting to the truth of where best to spend Christmas, but it seems more likely that statements like this are either cases of acceptance or simple summative group belief. In the case of Roz and Mark's position on war discussed above, there is clearly no constrained procedure at work; they are just having a casual conversation. We should therefore take Mark's statement of their collective belief as either a mistaken statement of summative belief (because Roz's agreement that some wars are justified was not genuine) or as a statement of collective acceptance.

The third and final criterion for group belief suggested by the court example is that group beliefs must be non-summative and non-dictatorial. Non-summative group beliefs are those that cannot be expressed as a simple function of individuals' beliefs. Non-dictatorial beliefs do not depend solely on the belief or action of a single individual; the group belief is not just a stand-in for one person's belief. In Pettit's example, the verdict of the court could not be straightforwardly reduced to the conclusions of the judges. Two out of three judges would have found the defendant not liable if they ruled alone, while the court itself found the defendant liable. This shows that the court's verdict is not simply a shorthand consensus statement or based upon some threshold of agreement. Further, no judge in the case is able to simply dictate the verdict. We could imagine a situation where the "head" judge was the only judge whose vote mattered, and the votes of the other judges were considered merely advisory. If the court's view *was* summative or dictatorial, its statement could still be a case of belief, but it would not be a case of *group* belief. Rather, it would either be summative belief (shorthand for "a majority of judges believe that...") or just an individual belief ("Judge Roberts believes that...").

Table 2 (Table 2) summarizes the criteria for group belief. Summative beliefs are, as be-

Table 2: Criteria for group belief.

State	Epistemic Product	Constrained	Non-Summative
Summative Belief	yes	?	no
Acceptance	?	no	yes
Group Belief	yes	yes	yes

beliefs, epistemic products. They may be the result of constrained procedures (opinion polls, for example), but might also be arbitrary (as with Roz and Mark). They are, by definition, summative. Groups may accept epistemic propositions, but doing so is voluntary—unconstrained. It is this direct, free choice that separates acceptance from belief. We cannot choose to believe, but we do choose to accept. A proper group belief is an epistemic product that is constrained and non-summative. It should be uncontroversial that being an epistemic product and being non-summative are necessary conditions for group belief. Beliefs are by definition epistemic products, and summative beliefs are shorthand for individual beliefs. The constrained criterion could be more controversial. It seems clear that it is a necessary criterion—beliefs cannot be the product of direct, unconstrained choices—but is it sufficient to demarcate belief from acceptance? Perhaps, for instance, a criterion requiring group beliefs to be non-instrumental is also required. Beliefs, as we normally think of them, are not adopted just because it would be useful to hold them. A union, following strict rules of order, might conclude that the employer can afford a 5% salary increase not because of any salient evidence but just because it strengthens the union’s bargaining position. This appears to satisfy all three criteria, but should surely be deemed acceptance. Indeed, not all constraints will be sufficient to produce group belief rather than acceptance. The constraints must be truth-aimed. Such constraints are what *cause* beliefs to be non-instrumental. An organization with constraints that are not truth-aimed will not produce collective beliefs.

Alternative criteria to demarcate group belief from acceptance could be considered. One option is to demarcate by fiat: group beliefs are non-summative and not acceptance. This approach has the virtue of simplicity, but fails to offer any guidance of *how* to tell belief from acceptance. Alternatively, one could try to further specify what characteristics are necessary for belief. Non-instrumentality, as discussed, is one option. Others include requiring that the procedure be designed with the intention of generating true propositions, or reliably producing epistemically justified propositions. Such emendations run the risk of excluding too many possible cases. Markets, for example, were not designed for the purpose of revealing truth. If they do so, it is as a byproduct of some other intention. If applied in the individual

case, requiring conscious design would also disqualify human beliefs as proper beliefs. A reliability criterion would disqualify unjustified beliefs, which is also undesirable; one might wish to argue that courts produce unjustified beliefs, for example. The constraint criterion is designed to be as permissive as possible while identifying the most significant way in which acceptance differs from belief.

In my account, group beliefs are epistemic products of constrained processes that are non-dictatorial and non-summative. This account is distinct from Wray's and Meijer's accounts, but there are important points of similarity. As discussed above, both Wray and Meijers hold the nature of the group to be a crucial factor in whether it is capable of having group belief. In Wray's account, the group must have "organic solidarity" and in Meijer's they must form a "structured whole." Both of these conceptions suggest a group where each individual is not fully autonomous. Just as an organ in the body is bound to the whole body's purpose, an individual in a group with organic solidarity is bound to the group's purpose. My requirement that a group belief be an identifiable product of a non-arbitrary process is similar to Wray's requirement that a group have a recognized way of "adopting a view." I also agree with Wray and Meijers that such groups "cannot exist without some goal" (Wray 2003, p. 369). I disagree with their argument however, that these properties of (some groups) imply that all groups are incapable of forming beliefs as such, but can only accept. For groups such as the courts, which have explicit epistemic aims, their statements of belief ought to be taken at face value. Some apparent belief statements may be better described as acceptances, for instance when there is evidence that the statement was made for pragmatic rather than epistemic reasons, but if all evidence points to epistemic intentions, "belief" is the better descriptor.

4 Market Prices

In this section I will give an account of how markets, and specifically the stock market, generate prices. The primary aim of this account is to demonstrate that market prices satisfy the essential features of collective belief: they are constrained, non-dictatorial and non-summative.

Stock markets use a "continuous auction" system to match sellers and buyers. On exchanges such as the New York Stock Exchange (NYSE), traders can place two types of orders: limit orders and market orders.⁶ A limit order is an order to buy or sell at a particular price.

6. There is a third type of order, the "stop order," where a trader can place an order to sell or buy only when the market price reaches a certain point. It is unimportant for the present discussion.

It might take the form, “Buy 100 shares at 32 or better.” A market order is an order to buy or sell at the best available price. Specialists keep track of the markets for individual stocks in an order book (once actually a book, now electronically) that records filled and unfilled orders. Unfilled orders are often represented by market depth charts, as in [Figure 1](#).

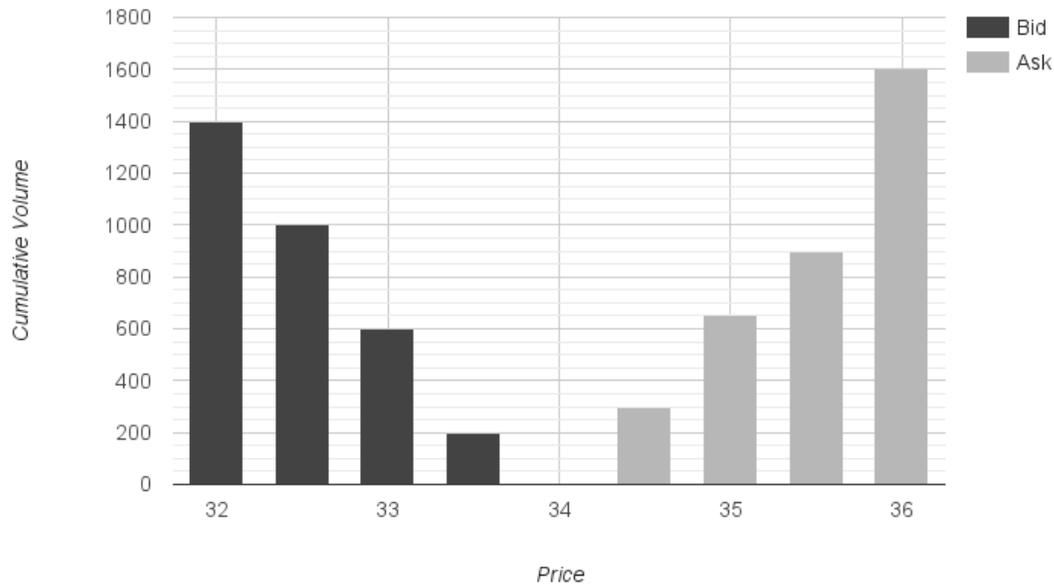


Figure 1: Market depth chart for a stock with a price of 34.

In [Figure 1](#), unfilled buy orders (bids) are represented in dark grey, while unfilled sell orders (asks) are represented in light grey. The height of each bar represents the total volume of unfilled orders at that price or better. For instance, there are 650 total shares being offered for sale at a price of 35 or greater, and there are orders to buy 1000 shares at a price of 32.5 or less. In this example there are no orders to buy at a price greater than 33.5 and no orders to sell at a price less than 34.5, so no orders will be filled until new orders come in.

The reported price of a share is the last price of a filled order. In this example, we can imagine that the last filled order was at 34. Now suppose a market buy order for 500 shares is placed. The best available asking price is 34.5, at a volume of 300 shares (the leftmost light grey bar). So 300 shares will be sold for 34.5, and an additional 200 shares will be sold for 35, resulting in a new market price of 35. Conversely, if a market sell order for 500 shares is placed, the new price will be 33. Alternatively, new limit orders could be placed, but these

will usually either cancel out (for instance, a bid at 34 and an ask the same), or narrow the bid-ask spread (for instance, a bid at 33.75 and an ask at 34.25). Therefore, limit orders will normally determine the magnitude of price movements and market orders will determine the direction of movements.

Price formation on the stock market is clearly very constrained: there are precise mechanisms for placing limit and market orders and for how the market mechanism is to resolve them. There are occasionally unpredictable human interventions in the process, such as when exchange operators decide to halt trading during precipitous declines, but these are rare exceptions. It is less clear that the constraints are truth-aimed. A plausible interpretation of stock market prices is that they simply reflect investor sentiment; prices rise and fall as stocks fall in and out of favor, with little or no connection to the value of their associated businesses. Savvy investors profit not based upon superior knowledge of the underlying business, but instead based upon their ability to predict the behavior of other investors. Accepting this view of the market, there would be little reason to infer any beliefs from market prices. However, this view cannot be *entirely* correct. Occasionally reality does assert itself on the market: if a company ceases to exist, so will its shares; if a company realizes profits it may issue per-share dividends proportional to that profit; if a company is acquired by another company, it will be by purchasing its shares. In each of these cases, investors are rewarded for making correct predictions about a company: what is the likelihood that it will go bankrupt, what are its expected dividends, and what price may it fetch in an acquisition? If the answers to such questions suggest that the company is over-priced, investors seeking to maximize their returns ought to sell its shares, which will in turn cause the price to decline, and vice-versa. Because of this corrective feedback effect, the constraints of the price formation process are truth aimed.⁷

It is easy to see that stock market prices are almost always non-dictatorial. Prices reflect actual transactions, and each transaction requires two parties. Any individual can raise or lower the price by placing market buy or sell orders, but how much the price rises changes depends on others' limit orders. An individual with an unlimited budget could in theory either set a floor or ceiling on price with large limit orders, or even set the price by taking both sides of a trade, but such manipulations are generally both extremely expensive and illegal.

7. Though as [section 6](#) discusses, this does not mean that market prices always, or even often, imply true beliefs. Perhaps, as suggested by Fredrich A von Hayek (1937), the best that can be said is that there is a "tendency" for market beliefs to become more correct.

Prices are also non-summative. Market price is not a simple average of what bidders are willing to pay and askers desire to receive. Nor is it a weighted average of willingness and desire. No matter the volume, a bid of 31 in the above example will have *no effect* on the market price. Neither is market price an average of the desire or willingness of market purchasers, as a market order does not indicate what price the trader was willing to pay, only that a buyer was willing to pay at least the lowest asking price or a seller was willing to accept the highest bid. Indeed market traders need not know themselves—many buyers of Apple stock probably have no precise opinion about how much they are willing to pay for an individual share.

5 Inferring Beliefs from Prices

It should be obvious that beliefs play *some role* in determining market prices. When traders buy or sell a financial security, it is hard to imagine that they do so having no beliefs whatsoever about that security. Determining the content of those beliefs would be an intractable problem, but that is not required for constructing a theory of market belief. Rather, the aim is to infer what *the market*, as a collective entity, believes about a security. Even this task, however, requires significant assumptions. This section will outline successively more ambitious approaches for inferring market belief. All of these approaches require making assumptions that are at best approximately true, and often not even that. Nevertheless, in some circumstances inferring market beliefs should be relatively non-controversial. One such case is the Challenger disaster. Another is prediction markets, which are specifically designed to facilitate belief inferences.

Inferring beliefs from prices is not new to epistemology. Subjective probabilities are often interpreted as dispositions to bet. Bruno de Finetti, for example, referred to the “obvious idea that the degree of probability attributed to a given event is revealed by the conditions under which he would be disposed to bet on that event” (De Finetti 1964, p. 101). Willingness to bet can be expressed as the price an individual is willing to pay for a given payoff should that event occur. This approach is problematic for real people: willingness to bet will depend upon the size of the bet, attitudes towards gambling, and many other factors other than beliefs about the probability of an event (Howson and Urbach 2006, p. 52). For example, if we find a person who is willing to pay \$50 for a wager that pays \$100 if France wins the World Cup, it might be the case that she believes that there is a 50% chance that France will win, but it could also mean that she is a fan of the French team, or that she really likes to gamble and would take nearly any bet offered. Making an inference from betting behavior to subjective

probability is necessarily problematic and uncertain.

Similar problems exist for inferring beliefs from prices in all cases. Economists view prices as arising from a combination of people's constraints (budgets), preferences, and beliefs (Hausman 1992, p. 13). For any given price, it is impossible to know, from the standpoint of economic theory, how each individual factor affects prices. If the price of orange juice rises, it might be because people expect poor growing conditions next season, that people all of a sudden really like orange juice, or that orange juice lovers have more money to spend. The same problem exists for the stock market: if the price of Apple rises, it might be because of investors' beliefs about Apple's profitability, or it might be just because investors *like* Apple—it is fashionable to own shares in Apple.

Economists modelling the behavior of stock market prices get around this problem by assuming that traders are motivated by only two factors: risk and return.⁸ This two factor account of stock prices is most often attributed to Markowitz's classic security analysis textbook, *Portfolio Selection* (Markowitz 1959). It is formalized by the Capital Asset Pricing Model (CAPM), which is the simplest and most well known model of stock market prices. The CAPM assumes that all investors are risk-averse (they require extra incentive to risk their money) and maximize their expected returns for a desired level of risk (Shih et al. 2013, p. 416), where expected returns measure the income a particular stock will generate. Risk is quantified by beta (β), the covariance between the returns of a particular security and the overall market. A high-beta security will have a high variance with respect to the market. If the overall market increases its returns a little, the security is likely to increase its returns a lot, and if the market declines, a high-beta security will likely decline more. The CAPM assumes that β can be estimated to a sufficient degree of accuracy from past observed returns (Fama and French 2004), and is therefore relatively stable. Therefore movements in a stock's price can be interpreted as changes in its expected returns.⁹

Who or what is doing the "expecting" in such cases? In the original formulations of the CAPM, all investors were assumed to have homogeneous expectations and access to all relevant information (Shih et al. 2013, p. 416). Other versions of the CAPM with

8. Purchasing a security on the stock market entitles the purchaser to an income stream proportionate to the underlying company's profits. Expected returns are the expected magnitude of that income stream. Risk measures the variability of that income stream. The CAPM assumes investors are "risk averse": they require extra compensation for assuming increased risk.

9. Economists generally agree that the CAPM is not an accurate model of stock market prices; factors beyond risk and expected returns, such as firm size, price-earnings ratios, and levels of debt all appear to have a significant influence on market prices (Fama and French 2004, 50). Further, past volatility might not be a good indicator of present risk. Nevertheless, the CAPM offers a good first approximation for connecting market prices to underlying beliefs.

heterogeneous beliefs also exist (J. T. Williams 1977; Levy, Levy, and Benita 2006). In both cases, it makes sense to think of the market as a collective entity whose beliefs and preferences do not reduce to any single individual, and to think of this collective entity as “expecting” the returns of securities to change. If investors really did have homogeneous expectations and access to the same information, then market belief would be a simple summation of those homogeneous beliefs (and since every individual’s beliefs are identical, the market’s belief would be exactly the same as the beliefs of each individual). Investors do not really have homogeneous expectations, however. For the more realistic heterogeneous models, if the market price was merely an average of each individual’s assessment of the fair price, then market expectation would represent the average expectations of investors. As the previous section argued, though, prices are not merely an average of individual assessments. Therefore in either interpretation of the CAPM there is no way to reduce market expectation to the expectations of individuals, and so the best option is to see the market itself as expecting.

The CAPM is indisputably an inaccurate model of investor belief. Few, if any, real investors are likely to consider *only* expected returns and historical variance in making their stock market trades. For the purposes of inferring market belief, however, the CAPM need not be an accurate model of individual investors. Rather, it needs to offer an account of how market beliefs (as collective beliefs) are connected to market prices, and that account ought to, in many cases, plausibly connect the response of market prices to economically relevant events. The CAPM accomplishes this goal. Other pricing models, with additional or different parameters could also be used to infer market beliefs. The chief advantage of the CAPM over other potential models is that it best captures economists’ intuitive understanding of market prices; when an economist attributes a belief to the market, it is based on a CAPM-like view of market prices as determined by expected returns. Additionally, because this view of price formation is so pervasive, many investors are likely to think in these terms as well and therefore to base their decisions solely on expected returns—market pricing theories are performative (MacKenzie 2006).

It is important to distinguish between what the market, as a collective entity, *believes* and what we can *infer* about its beliefs. If we accept that beliefs influence prices in some way, and that prices are non-summative and non-dictatorial, then it is a small step to accepting that prices reflect collective beliefs. This step does not rely on any particular economic theory or model. A model such as the CAPM is only required if we want to infer something about the *content* of those collective beliefs. The CAPM constrains the range of possible market beliefs to

beliefs about risk and return, and it makes it possible to identify price changes with changes in belief about expected returns. The CAPM and similar models only license inferences about risks and returns, not about their causes. When Apple's price decreases, the CAPM tells us that the market believes Apple's future returns will be lower than previously expected, but not *why* it believes they will be lower. Is it because Apple's new product isn't as great as expected, because one of Apple's competitors has a better than expected new product, or because the market believes Apple's cultural cachet has declined? The CAPM offers no guidance for deciding between these or any other possible explanations. Therefore a modest account of market belief as collective belief will stop at beliefs about expected risks and returns.

However, few attributions of market belief are so modest. Some are close, such as: "The market believes Amazon will beat its guidance and report \$25.78 billion in operating earnings and \$24.90 billion in revenue" (H. Williams 2015), but many are far more ambitious: "The market believes that stocks would be worth more if [Trump] were to lose the election" (Wolfers 2016). In the first case, beliefs are attributed to the market about earnings and revenue, which are closely related to returns, but in the second the link is much more tenuous.

Gilbert advises us to take such attributions of collective belief seriously, absent compelling reasons not to do so. A significant difference between market belief and Gilbert's collective belief, though, is that Gilbert's account is based upon joint commitment. For Gilbert, collective beliefs are formed when members of the collective commit to believing. This makes many statements of collective belief performative: by stating, "We believe," the members of the collective are simultaneously creating the group belief through joint commitment and asserting that the belief exists. Even when performativity is not at work, Gilbert's version of group belief requires some sort of explicit statement for the collective to collectively believe. This makes inferring the content of the group belief straightforward. No such resources are available for market belief, making attributions of market belief much more tentative.

Attributing a belief to a market can be thought of as an instance of inference to the best explanation (IBE). According to IBE, we infer the existence of a particular cause if that cause would be the best explanation of our evidence. Based on footprints through the snow in our backyard, for example, we infer that somebody walked through our yard last night. As Lipton (2004) argues, we can make such inferences even if we cannot even in principle check whether someone did in fact walk through the yard: if the walker is fundamentally unobservable. We can do this, according to Lipton, because in cases where we *can* verify the results of our inferences, IBE tends to be reliable. Market beliefs are exactly this sort of cause: they

are fundamentally unobservable, but if they did exist they would often serve as the best explanations for market price movements. Along similar lines, Dennett argues that we are justified in attributing beliefs (and other cognitive states) when doing so leads to successful predictions (Dennett 1998, 15).¹⁰ Taking Dennett's "intentional stance," it makes sense to attribute beliefs to markets when doing so is useful for predicting their behavior. For instance, attributing beliefs to markets could be used to predict how prices will change should new evidence arise that should change those beliefs.

A perennial challenge for IBE is to specify exactly what counts as a "best explanation." Here I have no significant contribution to offer. Good explanations are commonly supposed to possess virtues such as simplicity, generality, and coherence with other theories (Douven 2011). Lipton argues that causal explanations are best suited to IBE, and thus a good explanation will be a plausible causal story explaining the evidence: an unobserved walker *caused* the observed footprints and cadaverous matter *caused* the childbed fever deaths Semmelweis observed. Further, Lipton argues that often good explanations will be contrastive explanations: they explain why we observe footprints in the snow rather than a pristine whiteness.

A prime case of inferring market belief comes from the stock market's reaction to the Challenger disaster. Although NASA is a federal institution, it contracts out much of its work to business. In the case of the space shuttle, multiple corporations were responsible for its construction and operation. Rockwell International constructed the shuttle and main engines, Martin Marietta supplied the main fuel tank, Morton Thiokol made the solid rocket boosters, and Lockheed was responsible for ground support (Maloney and Mulherin 2003, p. 455). Each of these four was a publicly traded company. The stock price of all four companies dropped in the immediate aftermath of the disaster, but one stood out:

Morton Thiokol's 1-day return was -11.86%, more than 6 standard deviations greater than the firm's average daily stock return in the 3 months prior to the crash. By contrast, the stock returns of Lockheed, Martin Marietta, and Rockwell, while all negative, were less than 2 standard deviations different than the average return for the firms in the 3 months preceding the crash. (p. 456)

Following the modest approach to belief inference, all we can conclude from this is that the market believed that the disaster would affect Morton Thiokol's expected returns more than for the other companies. However, a more ambitious conclusion can be made: the market believed that Morton Thiokol, and consequently the solid rocket boosters that it manufactured, were responsible. Other interpretations are possible; for instance, a higher percentage of Mor-

10. Thank you to an anonymous reviewer for pointing me to Dennett's discussion of belief attribution.

ton Thiokol's business depended on NASA than did that of the other companies (p. 459). Indeed, the press at the time didn't place any particular weight on the stock market reaction and Thiokol's responsibility wasn't publicly established until several months later (p. 455). But these other factors aren't sufficient, according to subsequent analysis, to account for such a large decline (p. 459). Therefore, that the market believed Thiokol responsible is the best available explanation for Thiokol's price decline. Only that explanation enables us to account both for Thiokol's decline relative to the other companies involved in shuttle operation and the large magnitude of the decline. If the market did believe Thiokol was responsible, then we would expect to see the price behavior that was actually observed.

The above-mentioned claim that a Trump presidency would be bad for the stock market involves more complex reasoning. Wolfers (2016) offers four reasons for attributing this belief to the market. First, aggregate stock prices as tracked by the S&P 500 index rose and fell in correspondence with news about Trump's chances. For example, during one of the presidential debates "When Mrs. Clinton pummeled [Trump] over his tax returns, stocks rose." Next, there were no competing explanations for the price changes: "The [price increases] occurred between 9 and 11pm on a Monday [during the debate]... a stretch when there was no other important economic or financial news." Third, price changes were large compared to typical changes at the time of the debate. Fourth, other indicators, such as the price of the Mexican peso displayed similar patterns to those observed for the S&P 500. These are all good reasons to attribute the price changes on the S&P 500 to changing beliefs about the probability of Trump winning the presidential election. If the market believed that Trump would be bad for business, it explains all of Wolfers' observations—it is the simplest explanation for all of these abnormal patterns in stock market prices. Conversely, the lack of a significant decline in the S&P 500 after Trump's surprise victory suggests that Wolfers' inference may have been incorrect; Wolfers now must explain why the market apparently changed its mind about Trump so suddenly, and so his explanation no longer has the appeal of simplicity that it once did. Therefore we have less reason to take it as the best explanation for those price changes.

Prediction markets make belief inferences far more straightforward. Most prediction markets are double auction markets similar to the stock market. Market participants submit *bids* to buy at a certain price or *asks* to sell at a price. The market mechanism then completes transactions by matching compatible bids and asks. Whereas in stock markets it is shares of corporations being traded, in prediction markets the traded commodities are contracts promising payments based on the fulfillment of predictions.

The best known and most studied prediction markets are political markets, and especially the Iowa Electronic Markets (IEM) run by researchers at the University of Iowa College of Business. A typical instance of such markets is a presidential election market where contracts pay \$1 if the specified candidate is elected and \$0 otherwise. If contracts priced at \$0.70 pay out \$1 70% of the time, then their expected rewards are \$0.70—their price. This means that the price of a contract can be interpreted as the probability that its associated candidate will be elected president. All prediction markets have precisely defined procedures for evaluating whether a contract will pay out, and so making inferences from prices to probabilities that those conditions will be fulfilled is straightforward.

Interpreting stock market prices, as discussed above, is rarely a straightforward process. In some cases, such as the Challenger disaster, there is an obvious interpretation for price changes. In other cases, such as with Apple's share price, there are multiple possible interpretations. Prediction markets avoid all of these difficulties. If the price of a contract rises, it is because the market believes its associated outcome is more likely. *Why* it believes that outcome might be unclear, but *that* it does is straightforward.

6 From Belief to Knowledge

As I discussed above, it is reasonable to infer that on January 28, 1986 the market believed that Morton Thiokol, and by extension the solid rocket boosters it produced, were responsible for the Challenger disaster. It is also true that the solid rocket boosters, and therefore Morton Thiokol, were responsible. The market got it right. But did the market *know* that Morton Thiokol was to blame? Philosophers since Plato have observed that mere true belief does not amount to knowledge. The gap between the two is usually bridged by justification: knowledge is *justified* true belief.

The two leading accounts of justification are evidentialism and reliabilism. According to evidentialism, a person's belief is justified to the extent that their evidence fits their belief (Feldman and Conee 1985). According to reliabilism, a person's belief is justified to the extent that the cognitive process that produced the belief frequently produces true beliefs (Goldman 1979). Both these accounts of justification have been proposed for group beliefs. Tuomela, for example, argues for an evidentialist view of justification: a group belief is justified only if it is supported by the evidence available to group members (Tuomela 2012, p. 119). Goldman, in contrast, argues that group beliefs are justified only if produced by a reliable "belief-aggregation process" (Goldman 2014, p. 22).

Regardless of the merits of each account for other sorts of belief, there is only one viable account of the justification of market beliefs: reliabilism. For an individual's belief, it is conceivable that one might analyze the evidence possessed by that individual and thereby judge whether it is justified. This might also be possible for a panel of judges or a scientific research group. It is not reasonable, though, to expect to perform such an analysis for market beliefs that could be generated by hundreds or thousands of geographically dispersed individuals whose only interaction is through buying and selling securities and whose identities might not even be known. As Friedrich Hayek famously claimed, the correctness of market prices "can never be empirically verified for those cases in which it is of interest" (Friedrich A von Hayek 2002, p. 10). If it could, Hayek argued, there would be no need for markets in the first place; whatever method of empirical verification was used to verify market prices could equally be used to set prices by fiat, obviating the need for a decentralized price system. If market beliefs are to be justified, reliabilism appears to be the only option.

According to Goldman's "social process reliabilism," group beliefs are justified only if the individual beliefs responsible for them are themselves justified (Goldman 2014). Justification, according to Goldman, is *transmitted* from individuals to groups. Given that group beliefs are aggregations of individual beliefs, it makes some sense to require that the individual beliefs must themselves be justified for their aggregation to be justified. Where else could justification come from? However, this intuition is misleading: there are clear cases where there is no reason to think that any individual's belief is justified while the aggregated belief is reliable. For example, James Surowiecki tells the story of fairgoers attempting to guess the number of jellybeans in a giant jar, where the closest guess will win some prize (Surowiecki 2005). While the accuracy of any individual guess is quite low, given enough fairgoers the average estimate is extremely accurate. There are many such examples, from guessing the weight of a pig to the location of a lost submarine on the ocean floor, and the idea that collectives can produce guesses more accurate than their individual members can be seen as far back as Aristotle (Lyon and Pacuit 2013). In each case, the aggregate estimate can be almost eerily accurate. If there are enough fairgoers, some of them are bound to guess the number of beans correctly, but we would never call their true beliefs knowledge. There is little reason to expect that the next time those individuals guessed they would be any more accurate than their fellows. However, there is reason to believe that fairgoers will collectively continue to make accurate estimates; this is a reliable process for forming collective beliefs. We should be able to say that, although no individual knew how many beans were in the jar, collectively

they did know, at least approximately. What Goldman misses is that aggregation processes need not be epistemically idle: aggregation is capable of more than *transmission*—it can filter, weight, and de-bias individual beliefs. Not all aggregation is epistemically beneficial, though; group deliberation can lead to consensus views that are more polarized and less reliable than the beliefs of most individual members before the deliberation began, for example (Sunstein 2006).

If markets work similarly to fairgoers guessing about jellybeans, it is plausible that, following a reliabilist account of justification, all market beliefs are justified. Eugene Fama's Efficient Market Hypothesis (EMH) essentially makes this claim. Fama defines an efficient market as a market where, "given the available information, actual prices at every point in time represent very good estimates of intrinsic values" (Fama 1965, p. 90).¹¹ Describing a market as "efficient," by Fama's definition, is equivalent to describing its beliefs as justified. Prices represent estimates and if a market is efficient, its estimates will be "very good" at "every point in time." Just as the estimates of competing fairgoers cause their collective estimate to be more reliable than any individual's, Fama argues that analysts competing for profits on the stock market drive its prices toward accurate estimations of companies' values. If Fama is correct that the U.S. stock market is efficient, then its beliefs are always justified. However, evidence for market efficiency is mixed. Initial tests of the EMH, based on its implication that stock prices ought to follow a random walk, were very successful (Fama 1970). Further, stock markets seemed to correctly anticipate events such as stock splits and earnings announcements (Fama et al. 1969). As early as the 1970s, though, evidence of significant anomalies began to mount (Jensen 1978), and today there are a wide range of well documented violations of market efficiency. De Bondt and Thaler (1985), for instance, found that the market systematically overreacts to new information, and Shiller, Fischer, and Friedman (1984) argues that the stock market is subject to recurring fluctuations (bubbles) that do not correspond to future changes in earnings. Therefore, the case for accepting all market beliefs as justified based on the EMH appears insufficient.

Goldman suggests that to evaluate the justification of a belief, we ought to look to the reliability of the process that led to that belief. However, no two processes are exactly the same: they occur in different contexts and encounter different sets of evidence. When evaluating the reliability of a murder verdict, do we consider the court's past performance judging all crimes,

11. Fama (1970) defines an efficient market differently: "a market in which prices always "fully reflect" available information" (p. 383). This subsequent definition is the one that Fama and most others have adopted. However, it is easier to see the connection to belief in the 1965 version and they share the same essential view of markets.

all murders, all murders in the last number of years, or all murders allegedly committed by a certain ethnic group? Goldman calls this the “generality problem”: “A critical problem concerning our analysis is the degree of generality of the process-types in question. Input-output relations can be specified very broadly or very narrowly, and the degree of generality will partly determine the degree of reliability” (Goldman 1979, p. 12). For many beliefs, Goldman’s casting of the problem understates its difficulty; the problem is not just to decide how broadly or narrowly to define a process, but how to identify relevant instances of a process at all. As Conee and Feldman (1998) argue, when presented with a belief such as, “Smith knows that there is a maple tree nearby” (p. 23), how could one plausibly identify a process that led to this belief? What other of Smith’s beliefs might have been formed by the same process? We couldn’t simply point to *all* of Smith’s beliefs, as this would imply that humanity is divided into two groups: those whose beliefs are always justified and those whose beliefs never are. That would be an unacceptable result. Due to this and other considerations, () argue that reliabilism is an untenable account of justification.

On the other hand, Conee and Feldman grant that, for some beliefs, a process can be plausibly inferred. For example: “The process by which I just started my car is reliable”(22). In such cases they seem to concede that reliabilism may indeed be a plausible account of justification. It is because beliefs do not typically identify processes that reliabilism cannot serve as a full account of justification. Market beliefs are of this latter sort: all market beliefs (along with other proper group beliefs) are produced by an identifiable process and therefore the generality problem is tractable. Nevertheless, choosing an appropriate reference class will always be value-laden. At the least, there will be a tradeoff between the size of a reference class and its specificity, and how to balance the competing demand for a significant sample size with the requirement of sufficient similarity between process instances requires considering what sorts of errors are most important for us to avoid. For example, while prediction markets have a long track record of accurately predicting election results (Wolfers and Zitzewitz 2004), they have recently failed to predict the 2016 U.S. Republican nomination, the result of the “Brexit” referendum, and the 2016 U.S. Presidential election (*Who said Brexit was a surprise?* 2016). When considering prediction market beliefs for the next major election, should they be evaluated in reference to *all* prediction market beliefs, or just those in the recent past? There good reasons for both options. Prediction markets work the same way today as they did in 2010 and they are considering the same sort of question. This suggests that we ought to consider their full track record, which is still very good. Further, by considering their

full track record we have a larger sample size and therefore are less likely to be deceived by outliers. However, there is good reason to believe that the context of political elections today is fundamentally different than in the past: different demographics of people are voting and there is evidence that people are misreporting their intentions to pollsters. Therefore perhaps the results of prediction markets prior to 2015 are no longer relevant to predicting their performance in the future. Are we more concerned to avoid being deceived by outliers or to avoid missing relevant contextual differences? This is a value decision.

That there is no definitive way to demarcate justified from unjustified market beliefs according to reliabilism should not disqualify a reliabilist account of market justification. Unlike in the individual case, there is an identifiable process associated with each market belief and therefore the generality problem is less severe than for the individual case. Further, accepting that standards of justification are value laden does not mean that “anything goes.” Any reference class demarcation is ultimately answerable to future results. A reliabilist account implies that a set of beliefs classified as a single process ought to have a consistent track record. Any reason to believe that a subset of beliefs will not have the same reliability as other beliefs in their purported reference class is a reason to claim that they are not actually part of that class. If there is reason to believe that political prediction markets post-2015 will not be as reliable as political prediction markets pre-2015, there is equally reason to assign pre- and post-2015 markets to different reference classes. Similarly, if there is reason to believe that court verdicts for white defendants prove more reliable than for black defendants, there is equal reason to assign each to a different class. Conversely, if there is no reason to believe that stock prices on the New York Stock exchange are more or less reliable than those on the Tokyo Stock Exchange, there is no reason to consider price formation on each exchange as a separate process. Such reasons could be either causal or empirical, though causal reasons ought to manifest empirically and genuine empirical regularities ought to have some causal explanation.

Studies of financial markets provide both causal and empirical reasons for dividing markets into narrower reference classes. According to Paul Samuelson, “individual stock price variations are dominated by actual new information about subsequent dividends, but aggregate stock market variations are dominated by bubbles” (Private correspondence cited in Shiller 2014, p. 476). If this is the case, then evidence that markets reliably adjust beliefs about individual companies is not relevant to justifying market beliefs derived from aggregate market prices. In other words, if the price of an individual stock adjusts in response to some event, there is good reason to infer that the associated market belief is justified, but if

the overall market rises or falls rapidly, such as for technology stocks during the mid-1990s, the associated market belief (that internet companies will be wildly successful) should be treated with suspicion. For a variety of reasons, including the opportunity cost of investing, prices for short term predictions can be expected to be more reliable than those for long term predictions, and empirical tests have borne out this expectation (Rothschild 2009). Therefore short and long term market predictions ought to be considered different reference classes. In general, any behavioral finance study or test of market efficiency potentially offers a reason for demarcating markets into narrower reference classes.

These considerations suggest the market's belief in Morton Thiokol's responsibility was justified. While the stock market as a whole may not produce beliefs sufficiently reliable for justification, the market's belief about Thiokol falls within a narrower reference class of market beliefs that have been found to be more reliable than stock market prices overall. First, as Samuelson observed, prices for individual companies are more reliable indicators of value than market prices in general. Second, short term market predictions are more reliable than long-term, and investors would have known that an investigation of the disaster would produce an official verdict in a relatively short time, which in turn would have had immediate consequences for Morton Thiokol's profitability. Whether these considerations increase the market's belief about Thiokol sufficiently to count as justification is of course a matter of what threshold of reliability one requires for justification, but it seems fair to claim that if this instance is not justified, few if any market beliefs are; the market's reaction to Thiokol is nearly an ideal case.

In this section I have argued for a reliabilist account of justification for market beliefs. While for individuals or small groups it might be feasible to directly examine the evidence responsible for belief formation, this is not practical for large, decentralized groups such as markets. Further, market beliefs may be justified even when no individuals whose actions are responsible for those beliefs themselves hold relevant justified beliefs. When operating well, the epistemic power of markets is to generate beliefs that are more reliable than those of which any individual is capable. However, the track record of markets overall at forming justified beliefs is not sufficiently reliable to grant market beliefs universal justification, contrary to Fama's Efficient Market Hypothesis. Therefore when evaluating a particular market belief it is best to compare it to a narrower reference class than the market as a whole, though how exactly that reference class is chosen is inevitably ambiguous. Nevertheless, research on stock market anomalies and the performance of prediction markets suggest principled demarcation

criteria, such as preferring short term predictions and those for individual stocks rather than the entire market.

7 Conclusion

In this paper I argued that market beliefs should be considered a species of collective belief, though not one based upon commitment. Instead, market beliefs, and other collective beliefs, are epistemic products produced by constrained, non-summativ and non-dictatorial processes. Market beliefs are indicated by market prices. An inference from market price to market belief is an instance of inference to the best explanation. Reliabilism is the only tenable account of market justification; not only is evidentialism impractical for a market of any size, but market reliabilism is less vulnerable to critique than other forms of reliabilism because market beliefs are always associable with a concrete process. Though not all market beliefs should be considered justified, those concerning individual companies and short-term predictions are the best candidates for justified status. As the market's belief about Morton Thiokol satisfies both of these criteria, and it was true, we can say the market knew that Morton Thiokol, and its solid rocket boosters, were responsible for the Challenger disaster.

References

- Arvas, Taha Meli. 2015. Financial markets predict return to stability despite terror attack. *Daily Sabah* (). (Cit. on p. 3).
- Brown, Matthew J. 2015. A Critical Appreciation of Ronald N. Giere's "Distributed Cognition without Distributed Knowing". *Social Epistemology Review and Reply Collective* 4 (6): 1–7. (Cit. on p. 3).
- Commins, Patrick, and Jens Meyer. 2014. Markets Live: Gold dives, TPG shines. *The Sydney Morning Herald* (): 1–17. (Cit. on p. 2).
- Conee, E, and R Feldman. 1998. The generality problem for reliabilism. *Philosophical Studies* 89, no. 1 (): 1–29. (Cit. on p. 25).
- De Bondt, Werner FM, and Richard Thaler. 1985. Does the Stock Market Overreact? *The Journal of Finance* XL, no. 3 (): 793–805. (Cit. on p. 24).

- De Finetti, Bruno. 1964. Foresight: Its Logical Laws in Subjective Sources. In *Studies in subjective probability*, ed. H E Kyburg and H E Smokler, 93–158. Wiley. (Cit. on p. 16).
- Dennett, Daniel C. 1998. *The Intentional Stance*. MIT Press. (Cit. on p. 19).
- Douven, Igor. 2011. Abduction. *Stanford Encyclopedia of Philosophy*. (Cit. on p. 20).
- Fama, Eugene F. 1965. The Behavior of Stock-Market Prices. *The Journal of Business* 38, no. 1 (): 34–105. (Cit. on p. 24).
- . 1970. Efficient capital markets: A review of theory and empirical work. *Journal of finance*:383–417. (Cit. on p. 24).
- Fama, Eugene F, Lawrence Fisher, Michael C Jensen, and Richard Roll. 1969. The adjustment of stock prices to new information. *International Economic Review*:1–21. (Cit. on p. 24).
- Fama, Eugene F, and Kenneth R French. 2004. The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*:49–51. (Cit. on p. 17).
- Feldman, Richard, and Earl Conee. 1985. Evidentialism. *Philosophical Studies* 48 (1): 15–34. (Cit. on p. 22).
- FN Arena. 2014. Australian Corporate Bond Price Tables. *9 News* (). (Cit. on p. 2).
- Gilbert, Margaret. 2002. Belief and Acceptance as Features of Groups. *Proto Sociology* 16:35–69. (Cit. on pp. 3, 4, 6, 7, 10).
- . 2004. Collective Epistemology. *Episteme* (). (Cit. on pp. 1, 4, 7).
- Goldman, Alvin I. 1979. What is Justified Belief? In *Justification and knowledge*, ed. George S Pappas, 1–24. (Cit. on pp. 22, 25).
- . 2014. Social Process Reliabilism. In *Essays in collective epistemology*, ed. Jennifer Lackey, 13–41. Oxford University Press. (Cit. on p. 23).
- Hausman, Daniel M. 1992. *The Inexact and Separate Science of Economics*. New York: Cambridge University Press. (Cit. on p. 16).

- Hayek, Friedrich A von. 1937. Economics and Knowledge. *Economica* 4, no. 13 (): 33. (Cit. on p. 15).
- Hayek, Friedrich A von. 2002. Competition as a discovery procedure. *The Quarterly Journal of Austrian Economics* 5 (3): 9–23. (Cit. on p. 23).
- Howson, Colin, and Peter Urbach. 2006. *Scientific Reasoning: The Bayesian Approach*. 3rd Edition. Chicago: Open Court. (Cit. on p. 16).
- Hsu, Shi-Ling. 2011. A Prediction Market for Climate Outcomes. *University of Colorado Law Review*. (Cit. on p. 3).
- Jensen, Michael C. 1978. Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, no. 6:95–101. (Cit. on p. 24).
- Knorr-Cetina, K. 1999. *Epistemic cultures : how the sciences make knowledge*. Cambridge, Mass.: Harvard University Press. (Cit. on p. 6).
- Levy, Haim, Moshe Levy, and Golan Benita. 2006. Capital Asset Prices with Heterogeneous Beliefs. *The Journal of Business* 79, no. 3 (): 1317–1353. (Cit. on p. 17).
- Lipton, Peter. 2004. *Inference to the Best Explanation*. Second Edition. Routledge. (Cit. on p. 19).
- Lyon, Aidan, and Eric Pacuit. 2013. The Wisdom of Crowds: Methods of Human Judgement Aggregation. In *Handbook of human computation*, 599–614. New York, NY: Springer New York. (Cit. on p. 23).
- MacKenzie, Donald A. 2006. *An engine, not a camera : how financial models shape markets*. Cambridge, Mass.: MIT Press. (Cit. on p. 18).
- Maloney, MT, and JH Mulherin. 2003. The complexity of price discovery in an efficient market: the stock market reaction to the Challenger crash* 1. *Journal of Corporate Finance* 9 (4): 453–479. (Cit. on p. 20).
- Markowitz, Harry M. 1959. *Portfolio Selection*. Efficient diversification of investments. John Wiley & Sons. (Cit. on p. 17).

- Meijers, Anthonie. 2002. Collective Agents and Cognitive Attitudes. *Proto Sociology* 16:70–85. (Cit. on pp. [2](#), [3](#), [5](#)).
- . 2003. Why Accept Collective Beliefs? Reply to Gilbert. *Proto Sociology* 18-19:377–388. (Cit. on p. [4](#)).
- Miller, Boaz. 2012. When is consensus knowledge based? Distinguishing shared knowledge from mere agreement. *Synthese* 190, no. 7 (): 1293–1316. (Cit. on p. [3](#)).
- Orléan, André. 2004. What is a Collective Belief. In *Cognitive economics*. (Cit. on p. [2](#)).
- Pettit, Philip. 2011. Groups with Minds of Their Own. In *Social epistemology: essential readings*, ed. Alvin I Goldman and Dennis Whitcomb. Oxford University Press. (Cit. on p. [8](#)).
- Rolin, K. 2010. Group Justification in Science. *Episteme* (). (Cit. on p. [3](#)).
- Rothschild, David. 2009. Forecasting Elections: Comparing Prediction Markets, Polls, and Their Biases. *Public Opinion Quarterly* 73 (5): 895–916. (Cit. on pp. [3](#), [27](#)).
- Shih, Yi-Cheng, Sheng-Syan Chen, Cheng-Few Lee, and Po-Jung Chen. 2013. The evolution of capital asset pricing models. *Review of Quantitative Finance and Accounting* 42, no. 3 (): 415–448. (Cit. on p. [17](#)).
- Shiller, Robert J. 2014. *Speculative Asset Prices* (). (Cit. on p. [27](#)).
- Shiller, Robert J, Stanley Fischer, and Benjamin M Friedman. 1984. Stock prices and social dynamics. *Brookings Papers on Economic Activity* 1984 (2): 457–510. (Cit. on p. [24](#)).
- Solomon, Miriam. 2001. *Social Empiricism*. Cambridge, Mass.: MIT Press. (Cit. on p. [3](#)).
- Sunstein, Cass R. 2006. *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press. (Cit. on p. [24](#)).
- Supreme Court of the United States. 2015. *Obergefell et al. v. Hodges, Director, Ohio Department of Health, et al.* (Cit. on p. [7](#)).

- Surowiecki, James. 2005. *The Wisdom of Crowds*. Random House LLC. (Cit. on pp. 1, 23).
- Tollefsen, Deborah. 2003. Rejecting rejectionism. *Proto Sociology*. (Cit. on pp. 3, 6).
- Tuomela, Raimo. 2000. Belief versus acceptance. *Philosophical Explorations* 3, no. 2 (): 122–137. (Cit. on pp. 3–5).
- . 2012. Group Knowledge Analyzed. *Episteme* 1, no. 02 (): 108–127. (Cit. on p. 23).
2016. *Who said Brexit was a surprise?* (). <http://www.economist.com/blogs/graphicdetail/2016/06/polls-versus-prediction-markets>. (Cit. on p. 26).
- Williams, Henry. 2015. Amazon.com, Inc. (AMZN): Will AWS Help In Topping Forecasts Again? *Business Finance News* (). (Cit. on p. 19).
- Williams, J T. 1977. Capital Asset Prices with Heterogeneous Beliefs. *Journal of Financial Economics* 5 (2): 219–239. (Cit. on p. 17).
- Wolfers, Justin. 2016. Debate Night Message: The Markets Are Afraid of Donald Trump. *New York Times* (). (Cit. on pp. 19, 21).
- Wolfers, Justin, and Eric Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18 (2): 107–126. (Cit. on p. 26).
- Wray, K Brad. 2001. Collective Belief And Acceptance. *Synthese* 129 (3): 319–333. (Cit. on pp. 2, 3, 5).
- Wray, K Brad. 2003. What Really Divides Gilbert and the Rejectionists? *Proto Sociology*:363–376. (Cit. on p. 12).
- . 2007. Who has Scientific Knowledge? *Social Epistemology* 21, no. 3 (): 337–347. (Cit. on p. 6).