

7 | *Communicating virtues: the raters*

Credit rating agencies publish assessment (ratings) of the creditworthiness of issuers of corporate and government bonds and structured debt securities. Many investors use their services. If you visit the web site of Moody's, one of the big three credit rating agencies in the world, you acknowledge that you agree with its terms of use, which include the condition that you will

make your own study and evaluation of each credit decision or security, and of each issuer and guarantor of, and each provider of credit support for, each security or credit that you may consider purchasing, holding, selling, or providing.

You also agree that

any tools or information made available on [Moody's web site] are not a substitute for the exercise of independent judgment and expertise. You should always seek the assistance of a professional for advice on investments, tax, the law, or other professional matters.¹

Could Moody's be any clearer in encouraging you to be inquisitive? Their statements would probably fail to persuade most commentators. In an interview with *Guardian* journalist Joris Luyendijk, a senior analyst who had worked for Moody's described raters as the 'all-purpose bogeymen' for the global financial crisis.² Paul Krugman, the Nobel-winning economist and *New York Times* columnist, called their judgments 'literally worse than useless'.³ Other commentators certainly do not mince their words either when they compare the agencies to alchemists or astrologers, as we shall shortly see. The two quotes from

¹ www.moodys.com/termsofuseinfo.aspx.

² <http://www.theguardian.com/commentisfree/joris-luyendijk-banking-blog/2012/dec/17/rating-agencies-bogeymen-william-j-harrington>.

³ <http://krugman.blogs.nytimes.com/2013/02/23/little-statesmen-and-philosophers/>.

Moody's terms of use may suggest, however, that if raters are astrologers they are considerably more explicit about the limits of their predictions than most of their star-gazing colleagues: for horoscopes generally do not come with disclaimers as detailed as those that Moody's provides.

More scholarly sides have fervidly criticized the rating agencies, too. The most prominent criticism concerns misratings of structured finance. While in 2007 the bulk of mortgage-backed securities received top ratings, most of them are now considered junk bonds. The agencies are also blamed for their inordinately slow revisions. A day before Lehman Brothers blew up in September 2008, the bank still had good ratings from the big three credit rating agencies: Standard and Poor's, Moody's, and Fitch. In addition, the agencies have been accused of exacerbating the European government bond crisis; the increased costs of borrowing that Greece incurred after Moody's downgraded the country in June 2010 caused significant additional problems to an economy that was already in serious trouble. Raters have also failed to predict disastrous defaults (WorldCom, Tyco, Enron are only a few examples); they have been unwilling to disclose the methodological assumptions that underlie their judgements (methods are considered trade secrets); they have been accused of dubious sales techniques such as *tying* (threatening to downgrade an issuer if no additional services are bought from the agency), *notching* (only offering a rating of a security if other assets are rated as well) and helping issuers to design securities with a particular intended rating by providing them with the software they themselves use in their rating process. In addition, some authors decry the alleged conflicts of interest that arise when issuers instead of investors pay the agencies for rating (the *issuer-pays* compensation scheme), or when issuers solicit ratings from many credit rating agencies and decide to publish only the best (a phenomenon called *ratings shopping*). A plain fact of enormous ethical relevance, moreover, is that the market for credit rating is highly concentrated. Around 95 per cent of the market is in the hands of the big three American credit rating agencies. The Herfindahl-Hirschman index, a standard measure of market concentration, edges over 3,000, which is higher than for almost any other sector.⁴

Implicit in most of these criticisms is the claim that rating agencies do not do their work well enough. Moody's calls its ratings mere 'opinions' about the credit quality of debt obligations, which must not be viewed

⁴ www.gao.gov/products/GAO-10-782.

as ‘statements of fact’ or ‘recommendations to purchase, sell or hold any securities’.⁵ But most commentators find this purely underhand and uncandid. It is as though one were to sell toys with the disclaimer that determining the risk to children is the buyer’s responsibility. More should be done to guarantee that the toys are not hazardous. In one of the few publications devoted to the ethics of credit rating agencies, Steven Scalet and Thomas Kelly even argue that

reasonably accessible investing information is not merely a public good . . . but an important component for creating conditions of justice in a capitalist society, akin to making voting reasonably accessible to all in a democratic society.⁶

If they are right, rating agencies do not even resemble toy manufacturers very much. What agencies do comes close to realizing human rights. Such an important function, it seems, requires a high level of epistemic virtue, not only in the *production* of the ratings, but also in their *communication*. What dishonest disclaimers about ‘opinions’ seem to be motivated by is little more than a desire to evade liability. Consequently, the rating agencies need other-regarding epistemic virtues such as honesty and generosity; they need an ethics of communication. Or so it would seem.

The distinction between self-regarding and other-regarding virtues is not new. Primary examples of self-regarding virtues are courage and patience, because they are directed at ensuring our personal wellbeing; other-regarding virtues, by contrast, further the good life of others, and include benevolence, justice and honesty.⁷ Virtue epistemologists do have a view of other-regarding virtues, although most authors have only discussed them fairly briefly. Jason Kawall, and Robert Roberts and Jay Wood are the authors of quite elaborate accounts of other-regarding virtue, to which I turn shortly. Linda Zagzebski lists the ‘teaching virtues’ among the intellectual virtues, defining them as ‘the social virtues of being communicative, including intellectual candor and knowing your audience and how they respond’.⁸ Jason Baehr spends some time on generosity.⁹ Heather Battaly examines ways in

⁵ <https://www.moodys.com/termsfuseinfo.aspx>.

⁶ Scalet and Kelly, ‘Ethics of credit rating’, 489.

⁷ See, e.g., Taylor and Wolfram, ‘Self-regarding and other-regarding virtues’, which critically examines the way self-regarding virtues have been appraised.

⁸ Zagzebski, *Virtues of the mind*, 114. ⁹ Baehr, *Inquiring mind*, 110–11.

which teachers may encourage students to show concern for epistemic virtue.¹⁰

In this chapter I discuss other-regarding epistemic virtues in more detail. One reason is that I have already referred to them before when, for instance, I showed why CEOs should be epistemically generous. Another reason is perhaps more surprising. Unlike most commentators I do not think that accusing the credit rating agencies of disingenuous communication is so straightforwardly plausible; rather, I believe, the problems surrounding them are to be seen in the light of a form of regulation that has led to unjustifiable outsourcing of epistemic responsibility. Governments have singled out the rating agencies as nearly official sources of information about credit risks, whose verdicts investors are legally bound to take seriously. As a result, investors have become less interested in forming their own judgements about these risks. Instead of encouraging epistemic virtues, regulation has dumbed investors down, and inexcusably so. That is what I argue here at any rate.

I start with a brief discussion of Jason Kawall's view of other-regarding epistemic virtues. I show that for all its ingenuity his view misses an essential difference between epistemic and non-epistemic other-regarding virtues: the need for the beneficiary to cooperate. I introduce the concept of *interlucency* to show what this requirement amounts to, and illustrate this by means of a case about stock market recommendations that is also interesting in its own right. I then turn to the credit rating agencies and regulation.

Other-regarding epistemic virtues

Kawall groups the other-regarding epistemic virtues in three categories.¹¹ Two categories are, I believe, best seen as 'meta-virtues'; echoing Zagzebski's suggestion, Kawall calls them the virtue of being a *good teacher* and the virtue of being a *good listener* or *critic*. The third category contains honesty, sincerity, integrity, creativity and other traits inspiring people to communicate in virtuous ways. Like non-epistemic other-regarding virtues, these virtues are constitutive of the good life. Kawall expresses himself slightly more conditionally here than in the

¹⁰ Battaly, 'Teaching intellectual virtues'. Also see *Journal of Philosophy of Education*, 47, 2, 2013 (special issue).

¹¹ Kawall, 'Other-regarding epistemic virtues'.

non-epistemic case, writing that ‘the development of other-regarding virtues *may* constitute part of the epistemic flourishing and wellbeing of an epistemic agent’ and that ‘[a]n epistemic agent who focuses exclusively on self-regarding epistemic virtues (gaining knowledge and justified beliefs for herself alone) *could* be a deficient epistemic agent to the extent that she is a member of a community’.¹²

Kawall advances a number of arguments in defence of this claim. The first argument that other-regarding epistemic virtues are essential elements of *eudaimonia* looks at science. Scientists typically think of themselves as contributing to a ‘common body’ of knowledge rather than a mere ‘personal stock’ of knowledge’.¹³ Kawall seems to imply that their doing so is essential. It is, he thinks, part of a scientist’s good life to work for the sake of the scientific community. Secondly, communities value acquiring new knowledge more than acquiring old or irrelevant knowledge. Kawall illustrates this claim by comparing a person discovering a new species in the Amazon basin with a person memorizing an entire encyclopaedia. The latter’s cognitive accomplishments may, if anything, be admired; but the former’s epistemic contributions will be genuinely valued; and what we value in the former’s contributions is, according to Kawall, other-regarding virtues. Kawall’s third argument for other-regarding epistemic virtues uses a case due to Jonathan Kvanvig.¹⁴ Kvanvig asks us to imagine two agents S and T who are completely identical with respect to the knowledge they possess. What S knows T knows, and what T knows S knows. The only difference between the two is that S has acquired the knowledge all by herself, whereas T has learnt everything from S. Kvanvig claims that S is a ‘superior cognitive being’ than T.¹⁵ Kawall agrees, and he believes that S’s cognitive superiority can be adequately explained by other-regarding epistemic virtues: for S ‘has developed other-regarding epistemic virtues which [T] appears to entirely lack’.¹⁶ Furthermore, Kawall seems to suggest, without other-regarding epistemic virtues it would be impossible to explain Kvanvig’s judgement, and that is why we need them. Kawall’s fourth and final argument takes a case of a very good teacher inspiring students to become genuinely interested in and curious about the topics

¹² Kawall, ‘Other-regarding epistemic virtues’, 260; emphasis added.

¹³ Kawall, ‘Other-regarding epistemic virtues’, 268.

¹⁴ Kvanvig, *Value of knowledge*, 148. ¹⁵ Kvanvig, *Value of knowledge*, 148.

¹⁶ Kawall, ‘Other-regarding epistemic virtues’, 271.

she teaches. What Kawall values about the teacher is that she ‘contributes to a surplus of true beliefs over false beliefs . . . among her students and community’.¹⁷ Again, other-regarding epistemic virtues are needed to explain this.

One may find fault with the diagnosis Kawall gives of specific cases. I do not, for example, think that Kvanvig’s case necessarily suggests that T lacks other-regarding epistemic virtue. T may just as well have failed to carry out investigative actions because of a lack of self-regarding epistemic virtues or a lack of opportunity for research. (In the latter case, S is not rightfully called *cognitively superior* to T.) Most people know most of what they know about maths the way T knows things. They learn maths from others; but this does not mean that they lack other-regarding epistemic virtue. Furthermore, Kawall’s approach to epistemic virtue is, I think, rather highbrow, making it difficult to apply it to simpler forms of knowledge we attempt to acquire. Only a few of us are scientists; most of us, however, want to know how to prepare a meal or drive a car. We encountered these problems in Chapter 2, where I proposed an alternative to Jason Baehr’s view of personal intellectual worth because of its being too intellectualist to capture knowledge acquisition outside the domain of science. It is not so much a desire for wisdom, but a desire for profit that leads businesses to engage in research activities, and I do not see any reason to judge these activities as less virtuous. The idea of instrumental epistemic value I submitted as an alternative to Baehr’s view is not immediately applicable to other-regarding virtues, however, if people gain knowledge as instrumental to their own goals only. Up to now the instrumental value of epistemic virtues has been in their contribution to gaining knowledge. What knowledge others acquire was important only in so far as it influenced our own knowledge or our ability to gain knowledge. I have somewhat neglected this point at various stages of the book, perhaps rather carelessly speaking about the epistemic generosity of a CEO, for instance, without making it clear that, as a virtue, the generosity of CEOs does not directly contribute to realizing their own private goals but rather those of the company. Generosity may truly be a nuisance to a CEO whose mainspring is to get rich. Yet generosity can be consistently viewed as an instrument to reaching particular goals. Despite my disagreement with

¹⁷ Kawall, ‘Other-regarding epistemic virtues’, 271.

Kawall, his theory of other-regarding virtues does contain a suggestion as to how these goals could be developed further. We have to think of these goals as arising out of a *community*.¹⁸ Kawall's example is the scientific or academic community furthering science. But nothing in the concept of community prohibits us from applying it to business. Indeed, this is exactly what a flourishing Aristotelian tradition in ethics has begun to examine, viewing firms as communities contributing to the common good in ways that transcend individual *eudaimonia*.¹⁹ What Kawall refers to as *community*, in business consequently becomes the corporation, the partnership, the firm. A firm acquires epistemic virtue among other things by ensuring that the individual epistemic virtues of employees match the demands placed upon these employees by the specific way in which their job contributes to realizing corporate goals set by its directors. Some of the required epistemic virtues will be self-regarding. A person working in the research and development department cannot do without love of knowledge. Some of the virtues, however, will be other-regarding, such as CEO generosity.

One might object that generosity does not necessarily contribute to every corporate goal, strictly speaking. Hiding things from investors may sometimes be better advice if strict maximization of shareholder value is one's goal. When the aim is to derive other-regarding epistemic virtues from community goals, we should not therefore consider the corporation as a community in isolation. Corporations operate within larger environments. As we have seen, even Milton Friedman, who is generally viewed as one of the most uncompromising advocates of shareholder value maximization, assigned lexicographic priority to two other goals: namely, law and ethics. It is the responsibility of corporate executives to earn as much as possible for the owners of the firm provided they conform to 'the basic rules of the society, both those embodied in law and those embodied in ethical custom'.²⁰ What I should say therefore is that the epistemic virtue of generosity for a CEO originates in the corporate goals together with these 'basic rules of society'. The corporation is still the community from which other-regarding epistemic virtues arise; the rules of the society in which the corporation functions place conditions on the goals this community can develop.

¹⁸ Kawall, 'Other-regarding epistemic virtues', 272.

¹⁹ Sison, 'Common good theory'. ²⁰ Friedman, 'Social responsibility', 33.

Generosity

Let me now turn to generosity. Kawall seems to claim that other-regarding virtues do not require us to ensure that our audiences understand what we say:

Honesty is a virtue, and we have duties to testify clearly, etc. in a fashion which should help others to gain true beliefs. But we need not guarantee that our testimony will be accepted. Compare – there is a moral other-regarding virtue of benevolence, even if we cannot guarantee that, e.g., money we donate will be used for food and not bombs.²¹

First difference from non-epistemic virtue

This is plausible if it refers to our inability to *force* our testimony or beliefs upon another person. It is implausible if the position stems from a reticence to explore the further consequences of other-regarding epistemic virtues. One way to see this is to turn to Robert Roberts and Jay Wood, who define generosity as a disposition to give freely, for the purpose of benefiting the receiver. Their definition includes generosity (giving) and good stewardship, the two dimensions of Aristotle's liberality.²² Despite the definition's straightforwardness, what epistemic generosity motivates and enables one to do is far from obvious. When epistemically generous people give information to others, they do not lose what they give, unlike non-epistemically generous givers. This does not mean, however, that giving information comes at no cost. The costs of sharing knowledge about music or tennis with my neighbours will probably amount to nothing more than the time spent on it. Buyers who freely share with a dealer in second-hand cars the maximum price they want to pay, however, will certainly end up paying too much. Adopting the austere picture of epistemic virtue propounded by the personal intellectual worth view makes it rather difficult to develop a concept of generosity that is sensitive to this issue. This is one of the reasons why I explore a view of epistemic virtue based on instrumental epistemic value. What type of information sharing generosity amounts to in business typically depends on the particular non-epistemic ends that generosity is

²¹ Kawall, 'Other-regarding epistemic virtues', 274.

²² Roberts and Wood, *Intellectual virtues*, 286–304.

supposed to contribute to, and obviously no business enterprise has as an end the hastening of its own demise by helping its competitors.

This is one difference between epistemic and non-epistemic generosity: one does not necessarily lose what one gives if the gift is knowledge, but that is not to say that giving epistemic gifts can never harm the giver. Another difference is that one only succeeds in giving an epistemic gift if the recipient of the gift cooperates in certain ways. Money given to a charitable organization is a gift, even if it the organization misspends it; it is a gift once the charity's bank account has been credited. Sending an item of information to a person does not, however, entail that the sender has made an epistemic gift.

Second difference from non-epistemic virtue

To see this, I move to investment recommendations, which are provided by stock market analysts. Analysts give recommendations about company equity. The format is quite rigid, allowing them to choose exactly one of the following five possible recommendations: *strong sell*, *sell*, *hold*, *buy* and *strong buy*. When an analyst has a hold recommendation on Royal Dutch Shell it suggests, one would think, that one should not sell shares in Royal Dutch Shell if one owns them, but should not buy them either. What else can *holding* shares mean? All the same, one interprets the hold recommendation as a recommendation to *sell*. Analyst recommendations show a shift of scale (called *stock recommendation bias*) not unlike that of a tennis coach consistently characterizing terrible shots as 'not bad'. Less than 5 per cent of all recommendations are recommendations to sell. Of all recommendations, 95 per cent are as a consequence either neutral or positive. This cannot be what the analysts mean. In reality hold recommendations are recommendations to get rid of the shares, and only 'very bad' shares get sell or strong sell recommendations.²³

Institutional investors (insurance companies, pension funds, large endowments, etc.) are fully aware of the bias.²⁴ They sell after a hold recommendation, buy after strong buy recommendations and do nothing after buy recommendations. Small, non-professional investors trading

²³ Malmendier and Shanthikumar, 'Are small investors naive?'

²⁴ S. Iskoz, 'Essays in financial economics', MIT Sloan School of Management (2003), dspace.mit.edu/bitstream/handle/1721.1/16969/53484012.pdf.

on their personal accounts do not discount the bias, however. They take analyst recommendations literally, to their potential disadvantage. (The issue of stock recommendation bias is very different from the issue of whether analysts can outperform the market.) There seems to be no evidence to support the claim that small investors' lack of knowledge of the bias is the result of analysts intentionally deceiving them.²⁵ I offer the diagnoses that their lack of knowledge is caused by miscommunication and failures of epistemic generosity and what I shall call *interlucency*.

I first sketch a game-theoretical model of investment recommendations. Using modelling techniques from linguistics, the interaction between analysts and investors may be viewed as one between speakers (or senders) and hearers (or recipients).²⁶ Analysts have three 'strategies' to choose from, which they can use to communicate their advice. They may use an upwardly biased strategy *U*, a literal strategy *L*, and a downwardly biased strategy *D*. Investors, in turn, may interpret analysts at face value and use a strategy *l*, or they may interpret them as being upwardly or downwardly biased, with corresponding strategies *u* and *d*. The most natural outcome arises when both analysts and investors 'play' their literal strategies; but adopting biased strategies *U* and *u*, or *D* and *d*, in no way disrupts communication. Converging on *U* and *u* is exactly what analysts and large investors do.

I now turn to a defence of the second claim about the difference between epistemic and non-epistemic virtues. Generous people share knowledge with others. Sharing knowledge is more than merely sending a particular message in a linguistic game; it is sending a message that hearers are in the position to use to increase their knowledge. This is no different from non-epistemic virtue. I am not really generous if in response to a demand for transportation I offer my car for use by a person who is unable to drive. Genuine generosity in such a case would lead me to offer the person a lift. Similarly, epistemically generous people adjust the way they communicate to their audience and try to ensure they use the speaker strategy that the hearers are likely to match.

This leads to an interesting difference from non-epistemic generosity because it also requires active cooperation from the recipient. For

²⁵ Malmendier and Shanthikumar, 'Are small investors naive?'

²⁶ Traces of such models can be found in Lewis, *Convention* and Schiffer, *Meaning*. Also see Stalnaker, 'Common ground'.

beneficiaries of non-epistemic generosity to benefit from generous gifts, they only need to accept them. If you accept the lift someone offers you, or if you do not pay back the money you receive in your bank account, generosity has done its work in an unmediated or immediate way, whatever use you may make of the money. Epistemic generosity, by contrast, uses language as a medium only and the gift is not the mere utterance of words. If someone gives you advice, you have to interpret the linguistic utterances in which the advisor has cast the advice. This may go wrong because you may interpret the message incorrectly.

The reader may object that this is also true of non-epistemic generosity. Non-epistemically generous financial aid to, say, famine victims causes similar problems if it fails to reach the victims. This problem is more accurately described, however, as one in which the gift was not used in the way the giver intended. The problem with epistemic generosity is not that the gift is misused, but rather that no gift has been given as long as the recipient fails to interpret the linguistic utterance correctly. The investor first has to interpret a hold recommendation as a recommendation to *sell*. It is subsequently up to the investor to decide whether to use this 'gift' as it was intended, to misuse it or not to use it at all.

A consequence of reasoning along these lines is that to be epistemically generous, people must express themselves in ways that the beneficiaries of their generosity understand. This in turn requires that the recipients provide the senders with relevant feedback, especially when, as in the case of stock recommendations, common words acquire uncommon meanings. (Uncommon words acquiring common or uncommon meanings is much less of a problem because recipients can easily spot uncommon words and ask for clarification.)

Let us return to the example of analyst recommendations, and let us suppose that a particular analyst believes that investors should rid their portfolios of Royal Dutch Shell equity. To communicate this advice the analyst has to choose a communication strategy such as *U*, *L* or *D*. Epistemically virtuous analysts choose a strategy they believe the recipients interpret as a recommendation to sell the shares. It is important to note that this does not exclude any of the three strategies. As we have seen, when analysts and institutional investors communicate and interpret via upwardly biased ways and coordinate on choosing *U* and *u*, analysts get the recommendation across. A true mark of epistemic generosity is that the sender has reasons to think that the recipient uses the correct strategy; and to examine whether such reasons

are available requires that the sender actively track the recipient's understanding. The sender cannot do this, however, unless the recipient is sufficiently open about her interpretation. The recipient has to acknowledge receipt of the message and must try to make clear how she understands the message. Both sender and recipient have to contribute to sufficient openness concerning the communication and interpretation strategies they use in order that epistemic generosity gets off the ground. Contributing to such openness by tracking understanding, acknowledging receipt, providing feedback and so on is what I call *interlucent* senders and recipients do.

Personal, one-to-one communication between finance practitioners and customers, and to a lesser extent telephone conversations and email correspondence, are ways of communication that allow interlucency. Advisors talking to clients have ample opportunity to track understanding. Carefully listening to clients is often sufficient to spot errors in understanding. It is evident that a client's stated intention to sit still after having received a hold recommendation betrays a clear misunderstanding, and a virtuous advisor seizes the opportunity to set this right. By contrast, unilateral communication using web sites, information leaflets and other forms of written documentation offer less space for interlucency. Senders never know whether the intended recipients read the web sites and brochures. They have little room for tracking the recipients' understanding. Recipients who fail to understand have no way to gain clarification, except by face-to-face communication. Interlucency may be conceived of as an epistemic virtue. To avoid communicative misunderstanding, interlucient people try to place themselves in the position of others and adopt their perspective. They pay due attention to what others say, but they also actively signal their own interpretation in order to allow their communication partners to provide feedback on these interpretations or to adapt their communication strategies.

Should we conclude that stock market analysts show insufficient concern for interlucency? That would be going fast. In describing my shots as 'not bad', my tennis coach by no means fails to help, as long as I understand what he means. Epistemic generosity is entirely compatible with understatement, hyperbole or other figures of speech, where they do not obscure communication. Given that institutional investors are perfectly capable of understanding analysts' recommendations, the case against the analysts is fairly weak. Secondly, epistemically temperate

private investors realize that they do not fully grasp much of what they read. They know that recommendations from such consumer organizations as Which? or Consumer Reports should be taken with a grain of salt. Temperate people do not make their purchasing choices entirely dependent on what others say. Only a mild degree of curiosity suffices for private investors to consult web sites and articles explaining the stock recommendation bias (and also, by the way, the sheer lack of evidence backing the added value of analyst recommendations). Private investors following analyst recommendations without any further thought are in any case somewhat naive.²⁷

This conclusion may be disappointing; for why do we need a theory of other-regarding epistemic virtues in business if stock market analysts can get off so easily? Let me clarify. In earlier chapters I have already shown that other-regarding epistemic virtues are crucial to business, but I did so without turning to recent work in virtue epistemology. As the discussion of knowledge sharing in Chapter 5 made clear, no business enterprise can do without epistemically generous employees. In some way, the present chapter is more concerned with the *limits* of epistemic virtue. It is tempting to use the theory of epistemic virtue to make grand claims about the informational duties of professionals in the financial services industry towards clients and prospective customers. It is tempting to blame accountancy firms, banks, credit rating agencies, insurance companies, mortgage lenders, pension funds and governments for having provided us with so little and such obscure information, and it is equally tempting to find fault with analysts who fail to ensure that their audience understands their recommendations. I shall defend the view that though the temptation is understandable, it is misplaced. This, I hope, is not only interesting in and of itself, but also provides insights that are relevant to regulation. I argue that outsourcing epistemic responsibility is something that regulators should be reluctant to do.

Credit ratings

While stock market analysts are an important source of information for financial markets, credit rating agencies and accountants play more pronounced roles. It is chartered accountants who write the official auditors' reports that corporate annual reports are legally bound to

²⁷ Malmendier and Shanthikumar, 'Are small investors naive?'

include to make the documents valuable to banks, shareholders and tax officials, among others; and it is credit rating agencies that are designated by many governments as the sole authoritative source of credit risk. If an argument for other-regarding epistemic virtue among stock market analysts fails, one may still hope to make a case for such virtues in credit rating and accountancy. This chapter considers the raters, and shows that the case for other-regarding virtues is weak because governments have rather clumsily outsourced epistemic responsibility. The next chapter turns to the accountant, showing that the case for outsourcing epistemic responsibility is stronger once one considers that management and accountant form a joint epistemic agent.

Credit risk: asserting creditworthiness

It is useful to distinguish three functions of credit rating agencies: namely, estimating credit risk, monitoring issuers and, thanks to regulation, exerting influence on the management of regulated institutional funds. Firstly, their role is to furnish investors with estimations of the credit risk. *Credit risk* captures the risk that the issuer of a security (e.g., the corporation borrowing money) will fail to pay interest and/or repay the loan. It excludes such things as the risk that markets will turn unfavourable (market risk) or that no one will want to buy or sell the securities (liquidity risk). Credit rating agencies express their judgements of credit risk in letter combinations, ranging from the top-ranking AAA (for Standard and Poor's, and Fitch) and Aaa (for Moody's) to the D of default or bankruptcy. In the case of government debt, credit rating agencies also incorporate an estimation of the willingness to pay because, unlike companies, countries may decide not to pay back their loans when they think this will prevent social or political unrest.

Martha Poon describes the rating procedure in four steps.²⁸ The process starts with a primary analyst developing a preliminary rating on the basis of the financial statements provided by the issuer of the security. The credit rating agency then meets the issuer's representatives for discussion. During the third step of the process the credit rating agency develops a short report detailing and motivating the decision. The final step is that a committee is set up, including the primary analyst and the managing director, as well as other analysts, managers and staff

²⁸ Poon, 'Rating agencies', 283.

members with relevant knowledge. The committee votes on the final rating. The agency sends the final rating together with the report to the issuers. In principle issuers can appeal to the ratings decision, but they hardly ever do so. A press release finally publishes the rating.

Monitoring: directing management

A second role is that of *monitoring* the issuers. Credit rating agencies attempt to influence corporate or political decision making and they do this, not by participating in the issuer's decision making process, but by verbal means only: their ratings. The agencies review ratings every twelve to eighteen months. In the meantime, however, the primary analyst can put issuers on *watch lists* and provide *outlooks* about them, showing the concerns the agency has about the short- and medium-term development of the issuer's creditworthiness. Warning investors of potential ratings changes, these instruments may be perceived by the issuers as signals of problems that must be resolved to prevent a real downgrade.²⁹ Perhaps this sounds rather far-fetched as a method of active monitoring. Theoretical and empirical work in economics, however, shows that agencies use watch lists and outlooks as part of an implicit contract between agencies and issuers, the terms of which stipulate that issuers shall do their best to avoid future downgrades.³⁰ Particularly for issuers with low perceived creditworthiness, watch lists fulfil this coercive function rather well.³¹

Stamps of approval: directing investors

But how much value do ratings have to investors? Standard and Poor's emphasizes that its ratings have to be interpreted as providing information on the relative ranking of issuers, and so does Fitch. Moody's states that '[t]here is an expectation that ratings will, on average, relate to subsequent default frequency, although they typically are not defined as precise default rate estimates'; perhaps slightly inconsistently it describes its ratings also as 'relative'.³² Empirical work on credit rating

²⁹ Bannier and Hirsch, 'The watchlist'.

³⁰ Boot et al., 'Coordination mechanisms'.

³¹ Bannier and Hirsch, 'The watchlist'.

³² www.moodys.com/ratings-process/Understanding-Moody-s-Corporate-Bond-Ratings-And-Rating-Process/002005001.

agencies demonstrates that the ratings the big three agencies give to corporate and government bonds correspond rather accurately with default probability, suggesting that they offer more than a mere relative ordering of credit risk. Triple A amounts to a 0.5 per cent probability of default, whereas the highly speculative B- (Standard and Poor's) and B3 (Moody's) amount to 49.2 and 48.3 per cent.³³ But in contrast to what many investors thought before the subprime mortgage meltdown started, ratings do not have the same meaning across different classes of securities. Baa corporate bond ratings from Moody's were associated with a default probability of 2 per cent over the period 1983 to 2005; collateralized-debt obligations with the same rating had a twelve times higher likelihood of 24 per cent that they would default.³⁴

It is important to realize that the fact that ratings accurately reflect default probabilities offers no proof of their added value. Research on the determinants of bankruptcy shows that numerous measures may be used to approximate credit ratings rather accurately on the basis of publicly available information. This is a severe blow to the accomplishments of the agencies, given that they claim to have superior information, because it has been obtained privately in off-the-record conversations with the issuers themselves. These publicly available determinants include standard financial ratios of a firm's profitability, liquidity, solvency and size, but also measures of corporate governance (ownership structure and the way the firm is managed) and board independence, and a number of macroeconomic factors such as the growth of gross domestic product.³⁵ Lawrence White observes more technically that the correlation between credit ratings and default rate referred to above can also be obtained by looking at publicly available information about bond spreads, which is roughly the difference between what one gets from the bond and what one gets from a 'risk-free' benchmark such as US treasury bonds or Libor. As White concludes, '[t]he question of what true value the major credit rating agencies bring to the financial markets remains open and difficult to resolve'.³⁶ For all we know, they may, as Paul Krugman suggests above, be useless.

³³ Zhou, 'Credit rating and corporate default'.

³⁴ Strier, 'Rating the raters', 539.

³⁵ A classic paper is Altman, 'Corporate bankruptcy'. See Bhojraj and Sengupta, 'Bond ratings and yields' and Löffler, 'Rating through the cycle'.

³⁶ White, 'Credit rating agencies', 219.

That some investors do respond to changes in a security's rating despite the fact that ratings can be approximated on the basis of publicly available information seems hard to square with the hypothesis of efficient markets. (One version of the efficient market hypothesis is roughly that prices reflect all publicly available information.) A possible explanation of why investors respond (and also why, as they do, they respond more intensely to downgrades than to upgrades) leads us to a third function that credit rating agencies fulfil, besides informing investors about credit risk and monitoring the issuers of securities.³⁷ The letter judgements (AAA, AA+, etc.) play this third role as a consequence of a peculiar bit of financial regulation. In the 1930s US state governments started referring to credit ratings in their prudential regulation of pension funds. They also developed regulations prohibiting banks from investing in speculative investment securities, the sort of things popularly called *junk bonds*. This development has never stopped. Today the investment decisions of pension funds, health insurance companies, banks and many other financial services firms are severely curtailed, throughout the world, by rules that refer directly to the ratings published by a relatively small group of officially registered and accredited rating agencies.³⁸ When security's rating changes, managers of such institutional funds may consequently have to change their positions, even in cases where they have formed a different estimate of credit risk from the rating agency's.

The three roles that credit rating agencies play can be neatly summarized in philosophical terminology deriving from speech act theory developed by John Austin and John Searle.³⁹ We use words and sentences to carry out many disparate sorts of things such as asserting, ordering, promising, expressing emotions, pronouncing a couple 'man and wife', or directing people.⁴⁰ Most straightforwardly, ratings are *assertions of creditworthiness*. When Standard and Poor's gives a B+ rating to Austin Martin this is nothing other than the statement that the default probability of this company is around 32 per cent. Secondly, rating agencies provide *directives of management*. An example is Standard and Poor's informing Sainsbury about the measures that management should take to avert a potential downgrade:

³⁷ White, 'Credit rating agencies', 219.

³⁸ White, 'Credit rating agencies', 213.

³⁹ Austin, *How to do things*. Searle, 'Illocutionary acts'.

⁴⁰ Searle, 'Illocutionary acts'.

A weakening of [Sainsbury's] financial profile due to poor trading or capital investments and capital returns not fully mitigated by improvements in earnings could lead us to lower the ratings. Conversely, we could consider a positive rating action if Sainsbury achieved and maintained [funds from operations to debt ratios] of more than 25%.⁴¹

Thirdly, the agencies issue *directives of investment*. If Standard and Poor's gives Hilton a rating of BB-, as it once did, investors bound by regulation must be particularly careful if a one notch downgrade to B+ leads them sell the bonds, because then the hotel chain will verge close to *junk*.

Compromising epistemic virtue

Why did regulators endow the rating agencies with the authority to issue directives of investment? A little history may help us here. The predecessors of credit rating agencies were *credit reporting firms* such as the famous Mercantile Agency, founded in the United States in 1841. They expanded their activity particularly after the US Civil War, when demand increased for reliable information about the 'credit behaviour' of companies and individual businesspeople. Trade and mercantile exchanges started flourishing during that period, and much of the trade took place on the basis of trade credit. A buyer receives *trade credit* when a seller sells something but does not require the buyer to pay upon delivery but gives her, say, ninety days to pay. Trade credit is essential when, due to seasonal fluctuations in the buyer's cash flow, no payment can be made right away but only after the buyer has sold products to her own customers. Sellers only extend trade credit to buyers they have reason to trust. Credit reporting agencies therefore started gathering information that merchants could use to determine the trustworthiness of companies, using sheriffs, businesspeople, bank cashiers and other 'correspondents' as sources of information.⁴²

In the first two decades of the twentieth century, credit reporting firms changed in important respects. Until then they had specialized in providing information, leaving the ultimate judgements about creditworthiness to their clients. Around 1910, however, they began to publish

⁴¹ www.standardandpoors.com/ratings/articles/en/us/?assetID=1245193708812#ID2603.

⁴² Olegario, *Culture of credit*.

their own verdicts of creditworthiness and to adopt the letter system still in use today. The mid-1920s ratings from Moody's, covering almost all of America's corporate bonds, are an example. Governments found these verdicts reliable enough to include them in prudential regulation aimed at mitigating the effects of the crash of 1929 and the subsequent depression. In 1931, for example, the US Office of the Comptroller of the Currency, an important American regulator, introduced the distinction between *investment* and *non-investment grade* securities and determined that non-investment grade securities must be treated differently (as bearing higher risk) on a bank's balance sheet. Only five years later, an outright prohibition of banks investing in speculative securities followed, where the meaning of *speculative* had to be determined by officially recognized credit rating agencies.⁴³

It may be suspected that this development was inspired by the 1933 Glass-Steagall Act, or else by the general regulation-friendly sentiment that gave rise to the Act. But we see increased reliance on rating agencies in times of deregulation too. Since 1989, for example, American pension funds have been allowed to invest in asset-backed and mortgage-backed securities with high ratings, and in 2001 the Federal Deposit Insurance Corporation significantly weakened the capital requirements that banks in the United States have to satisfy concerning mortgage-backed securities receiving ratings of AA and above, from 8 per cent to only 1.6 per cent.⁴⁴ Even though outside the United States the role of credit rating agencies is of a more recent date, rather similar pictures come out of Europe and elsewhere. All in all 'the creditworthiness judgements of [credit rating agencies have] attained the force of law', as Lawrence White once said.⁴⁵

One may suspect that when agencies are granted such power it places enormous epistemic responsibility on them. This may change, however, once we recall that the added information value of credit ratings is dubious because they can be replicated on the basis of publicly available information. If ratings are just like horoscopes in that they do not add new information to what we already know, requiring epistemic virtue of raters is wide of the mark. One may blame astrologers for a lack of almost any epistemic virtue, but this is appropriate only if they seriously

⁴³ White, 'Credit rating agencies'.

⁴⁴ Pagano and Volpin, 'Credit ratings failures', 416.

⁴⁵ White, 'Credit rating agencies', 213.

conceive of what they deliver as genuine predictions. Most writers of astrology columns, however, seem to understand quite well that the game they play is a different one.

The comparison is perhaps a bit tendentious, but it does suggest that rather than blaming raters for a lack of epistemic virtue, the pressing issue is whether we should endow them with such epistemic powers. Lloyd Blankfein, then CEO of Goldman Sachs, once stated that

too many financial institutions and investors simply outsourced their risk management. Rather than undertake their own analysis, they relied on the rating agencies to do the essential work of risk analysis for them ... This overdependence on credit ratings coincided with the dilution of the coveted triple A rating. In January 2008, there were 12 triple A-rated companies in the world. At the same time, there were 64,000 structured finance instruments, such as collateralized debt obligations, rated triple A.⁴⁶

This indictment sounds largely true. Most astrologers only acknowledge their limited aims quite implicitly, but most people do not take horoscopes seriously. Credit rating agencies are, as we saw above, rather clear about their stated ambitions, but most of their clients use the ratings in ways that go beyond these ambitions. It is worth stressing that, like the readers of astrology columns and the users of stock recommendations, investors could have known more about the limitations of the ratings. Treating triple A rated structured debt securities as though they had a yield curve commonly associated with triple A rated corporate bonds was, as Philippe Jorion states, an 'act of blind faith in the credit rating', which is an expression of a lack of epistemic temperance.⁴⁷ Jorion made this comment in the context of a discussion of the Swiss financial services firm UBS. UBS employees rashly believed that the agencies were capable of deriving ratings of the quality they were used to obtaining for corporate and government bonds. The employees knew, however, that the agencies were much less experienced at rating structured finance than rating corporate debt.

The precise extent of credit risk seems to have left them cold at any rate. Despite being large enough to assign a team of economists to the task of comparing structured finance and corporate debt ratings, UBS apparently did not have the corporate curiosity to do so, nor to

⁴⁶ Quoted by Pagano and Volpin, 'Credit ratings failures', 404.

⁴⁷ Jorion, 'Lessons from the credit crisis', 929.

investigate the creditworthiness of issuers itself. This research would have been costly because obtaining information about all the underlying mortgages of a mortgage-backed security requires data that were only available from commercial data providers. But UBS could have done it.

Moreover, as Jorion also observes, UBS employees failed to ask even the simplest questions. How, for example, can a mortgage-backed security be assigned the triple A status of a riskless security and at the same time deliver a yield *much* higher than the Libor, a shining example of risklessness? The correlation between risk and return is the most fundamental principle of finance. A lack of epistemic courage may have led financial economists at UBS (and many other financial services firms) not to ask the obvious question: how can structured debt securities increase expected return and simultaneously stay almost risk free?

The discussion of the UBS case has drawn us into the topic of outsourcing epistemic responsibility. I defend the view that when regulation forces business to outsource epistemic responsibility to other organizations, epistemic virtues are in danger. Credit rating is used here as an example. If governments prohibit investors from investing in bonds characterized in terms of the credit risk as estimated by officially designated credit rating agencies, evaluating credit risk is no longer something that investors have reason to do themselves. This affects epistemic virtues. One might object that this is not very relevant as long as it does not influence investment behaviour among investors. Economists provide evidence, however, that the inflated ratings of structured debt securities contributed to a greater appetite for these products among investors.⁴⁸ Structured finance products are hard to disentangle, and without the ratings many investors would have found them too intricate to trade. Without the ratings, there would probably have been much less demand for them.

Love

Let me now turn to the virtues, love of knowledge to begin with. Outsourcing epistemic activities to credit rating agencies leads to a situation where regulated investors have little incentive themselves to probe the credit risk of securities they trade. The aggregate result of this is a decrease in epistemic activity, because absent such regulation more

⁴⁸ Pagano and Volpin, 'Credit ratings failures'.

parties would research credit risk themselves. Moreover, the methods a rating agency employs are largely unknown outside the agency. This decreases the quality of the research. Unlike the academic ideal of peer review and openness fostering informed and rational discussion, rating agencies keep their methods to themselves. This makes it more difficult to put their hypotheses to the test and is also likely to lead an unnecessary doubling of work.

The sheer complexity of structured finance products exacerbates this. A typical mortgage-backed security comprises hundreds or thousands of mortgages with different sorts of real estate as collateral. To assess the risks of such securities, raters have to assess, among other things, the magnitude of the correlation between the risks of the underlying assets (the collateral). It matters whether, say, all real estate is from Florida or from places scattered throughout the United States. The documentation that comes with mortgage-backed securities (the *prospectuses*) generally only contains statistical information about the average underlying mortgage, not about all individual mortgages. This is not enough to determine the correlation of risk, far from it, and data have to be purchased from data providers. Credit rating agencies were not very keen on doing research here, and investors themselves had no incentive either. Not until 2007, for example, did Moody's start requesting the simplest detailed data about the borrowers of mortgages such as the loan to value ratio, the borrower's credit score, and the borrower's debt to income level. These, however, are the most important indicators of a mortgage's credit risk.⁴⁹ An additional complexity is that the credit risk of mortgage-backed securities is determined not just by the risk that borrowers will default on their mortgage (the risk that they cannot repay), but also by the risk that they will pay back too early (and that the lender earns less interest than expected). Estimating prepayment risk is, however, mathematically complex.⁵⁰

Moreover, overwhelming evidence shows that a large majority of triple A rated structured debt had underlying loans (the things out of which the structured bonds were constructed) that barely made it to investment grade, which Efraim Benmelech and Jennifer Dlugosz aptly

⁴⁹ J. Mason and J. Rosner, 'Where did the risk go? How misapplied bond ratings cause mortgage backed securities and collateralized debt obligation market disruptions' (2007), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1027475.

⁵⁰ Agarwal et al., 'Optimal mortgage refinancing'.

describe as sheer *alchemy*, deriving as it does gilt-edged ratings out of junk bonds.⁵¹ More empirical research has to be carried out to examine this suggestion; as it stands, however, we have some initial indications that generosity and interlucency were not omnipresent.

Justice

Secondly, open-mindedness and epistemic justice are hardly fostered by a regime in which regulators bestow epistemic authority on particular companies. Regulated investors are forced by law to consider the rating agencies as the official source of information concerning credit risk. This largely obviates the need to consider what other sources say. A rating above the junk bond status is the only mark of approval an investor needs. This is aggravated by the issuer-pays compensation model. Issuers pay to get their securities rated; they are effectively the sponsors of the research that credit rating agencies carry out. To see why this is unlikely to contribute to epistemic virtue, consider pharmaceutical research. Drug studies funded by pharmaceutical companies show a systematic bias towards outcomes that favour the sponsor. Sponsored research is more likely to report positively on tested drugs.⁵² This phenomenon has not been thoroughly investigated in other industries, but a recent study by Andreas Milidonis suggests that bond ratings suffer from similar biases.⁵³ Milidonis investigated bond ratings for the American insurance industry, where both issuer-pays and investor-pays ratings are available. He did not directly examine whether issuer-paid agencies should be described as merely interested in currying the favours of the issuers, but he did find something that is epistemologically relevant all the same. Changes in ratings from issuer-paid agencies follow upon changes in ratings from investor-paid agencies; in other terms, issuer-paid agencies are not in the epistemic vanguard.

Independent evidence bolstering this claim may be obtained from observing the ways in which credit rating agencies developed the mathematical modelling techniques that play a fundamental role in rating structure securities. (I should point out that we do not know very much

⁵¹ Benmelech and Dlugosz, 'The alchemy of CDO ratings'.

⁵² Lexchin et al., 'Pharmaceutical industry sponsorship' is a frequently cited meta-analysis.

⁵³ Milidonis, 'Compensation incentives'.

about these models because they are trade secrets, but we know enough to develop a reasonable hypothesis.) In 2004, Moody's decided to introduce a new model for particular structured debt securities. Interviews between a Bloomberg journalist and former employees reveal that rather than stemming from a desire to increase accuracy, the reason for the shift was a desire to ease ratings standards. More structured securities would receive gilt-edged ratings to please their issuers. After Moody's had split from Dun and Bradstreet it became listed on the New York Stock Exchange in 2000. From then on, concerns about profitability and shareholder interests took centre stage, and for the first time in the history of the firm senior management received compensations partly in terms of stock options.⁵⁴

I do not wish to suggest that it is beyond dispute that gaining market share was the prime motivation underlying the revision of the rating models; the urge for reform may well have come from a realistic assessment that the traditional techniques of binomial expansion used for many structured products had become less suited to novel products having less diversified and more correlated collateral. That the new models lent themselves very nicely to doling out higher ratings attracting a new clientele, enlarging a hitherto rather small market share, does not make this suggestion very plausible, though. An unpublished study by Simi Kedia, Shivaram Rajgopal and Xing Zhou indeed suggests a strong link.⁵⁵ They discovered that after Moody's flotation on the stock market its ratings became decidedly more favourable than Standard and Poor's ratings.

Temperance

A third virtue to suffer is epistemic temperance. Philippe Jorion has pointed out that many risk management approaches have difficulties incorporating *unknown unknowns*. Examples are regulatory interventions in the form of trading restrictions or other market developments inspired by regulation, but also socio-political events or environmental

⁵⁴ K. Selig, 'Greed, negligence, or system failure? Credit agencies and the financial crisis', Case Studies in Ethics, The Kenan Institute for Ethics at Duke University, <https://kenan.ethics.duke.edu/wp-content/uploads/2012/07/Case-Study-Greed-and-Negligence.pdf>.

⁵⁵ S. Kedia, S. Rajgopal and X. Zhou, 'Did going public impair Moody's credit ratings?' (2013), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2343783.

catastrophes. Epistemically temperate risk assessment always leaves open the possibility that no decent quantifiable estimation of risk can be delivered for lack of information. Credit rating agencies did not, however, decide to withhold a judgement of credit risk on the grounds that the security was too complex; they did not characterize securities as *not rateable*.⁵⁶ They always rated. But where saying ‘We don’t know’ is not a possible outcome of inquiry, organizational support for epistemic temperance is severely decreased.

Courage

Excluding the possibility of ending up with no rating at all decreases the scope for practising epistemic temperance. In a similar way, excluding certain ratings changes compromises epistemic courage. Moody’s, for example, stated that it will never engage in ‘unnannounced multi-notch ratings changes’.⁵⁷ The firm will never radically change its mind about an issuer’s creditworthiness, allegedly to avoid disturbing financial markets or risking their relationship with issuers or investors. This is an intriguing, if flawed, argument. Epistemologists discuss whether one might adopt certain beliefs or hold on to certain beliefs for practical rather than epistemic reasons. Is it acceptable that I adopt a belief that, say, someone was killed accidentally rather than murdered if this avoids the riots that may result from bringing the murderer to justice? And if so, is it morally justifiable to do so on such grounds? This case may be difficult, and when ratings changes may lead to riots or even to wars, the agencies are certainly in an unenviable position. The answer to the questions, however, is easy to answer if their motivation stems from concerns about the risk of losing their clients. One needs epistemic courage to downgrade an issuer when one’s business depends on the issuer’s willingness to pay the business.

Generosity

Finally, I turn to other-regarding virtues. Regulation has led to a situation where the need for genuine communication between the senders

⁵⁶ Diomande et al., ‘Public credit rating agency’. Jorion, ‘Lessons from the credit crisis’.

⁵⁷ Quoted by Dooley, ‘Overhaul ratings process’.

and recipients of information has almost entirely disappeared. Agencies endowed with official epistemic authority that are paid by the issuers of the securities rather than the investors have little in the way of motivation to obtain feedback from the end users of their ratings, very much unlike the predecessors of the credit rating agencies, the credit reporting agencies. Moreover, the credit rating agencies may find it difficult to imagine what it means for investors not to understand the rating. The rating being only the letter combination it is (and regulation being quite clear about what that requires), what topics are there for them to discuss? A lack of generosity and interlucency, however, may be discerned in the fact that investors do not respond to ratings in the way they would rationally be expected to do if ratings had a completely unequivocal meaning. A fair amount of evidence indicates that investors respond asymmetrically to ratings changes. Most studies find that updates have no effects, but downgrades do.⁵⁸ Several theories are in the frame for an explanation of this phenomenon. Downgrades are more informative than upgrades if raters search more intensely for ‘bad news’ or if issuers provide ‘good news’ more readily themselves.⁵⁹ Others have suggested that investors respond to downgrades more than is rationally warranted.⁶⁰ The correct explanation need not detain us here. The fact is that raters do not interpret ratings literally as expressions of default probability.

To reiterate a point made earlier, it is true that around 2005 ample documentation was available showing that ratings were not comparable across asset classes, but time and again the rating agencies insisted that their models provided uniform rating measures. Standard and Poor’s stated in 2007 that

[o]ur ratings represent a uniform measure of credit quality globally and across all types of debt instruments. In other words, an ‘AAA’ rated corporate bond should exhibit the same degree of credit quality as an ‘AAA’ rated securitized issue.⁶¹

Moody’s and Fitch made similar claims. But as we saw, historical data reveal a very different story, making a rating of Baa from Moody’s for

⁵⁸ Gonzales et al., ‘Market dynamics’.

⁵⁹ Jorion and Zhang, ‘Information effects’.

⁶⁰ Dichev and Piotroski, ‘Bond ratings changes’.

⁶¹ Quoted by Pagano and Volpin, ‘Credit ratings failures’, 207.

structured bonds more than ten times as likely to default as a corporate bond.⁶²

Testimony

This conclusion is strengthened once we turn to *testimonial* knowledge. Within the theory of knowledge, two sources of knowledge and justification are distinguished. The most obvious source is perception. I know that I am sitting in front of my computer because I see that I am. Much of our knowledge, however, does not come to us through our senses. That sharks are dangerous and that Beethoven and Hegel were born in the same year I learned from other people. This kind of knowledge is called *testimonial knowledge* or *knowledge by testimony*. It is the sort of knowledge gained by reading books, asking experts, hiring consultants, listening to parents and teachers.

For testimony to be an acceptable ground for belief it has to be trustworthy, and it has to be perceived as trustworthy. But there are various kinds of obstacles to perceived trustworthiness. The most obvious obstacle arises when sources of testimony employ substandard belief formation policies. A source that does not possess genuine knowledge about a matter (owing to its not having carried out investigations in epistemically virtuous ways, for example) cannot help anyone to gain knowledge about it. Credit rating agencies using substandard research methodologies are therefore not trustworthy.

Another obstacle arises when the recipient fails to *perceive* the trustworthiness of the source. Sometimes this is caused by the recipient's being overly sceptical. Suppose that an unsubstantiated prejudice leads me to refuse to believe whatever analysts or raters tell me. Then I never trust their evidence and judgements, reliable though they may be. But in the case of the rating agencies, a more likely cause of a failure to establish a perception of trustworthiness is that the agencies are not particularly generous with information about the methodologies that underlie their ratings. To perceive an individual or organization as a trustworthy source of knowledge about a particular topic, one needs evidence of expertise. One needs indications that the source is knowledgeable in the relevant domain. That is quite difficult in the case of credit rating agencies. It is hard to find out exactly how agencies arrive at their ratings.

⁶² Strier, 'Rating the raters', 539.

Little information is available about the people responsible for a rating, their expertise and their rating success record. A great deal of the mathematical and computational methodology is hidden from our eyes. Rating agencies compete, among other things, on their methods, and consequently they consider their methods to be trade secrets.

But if we cannot determine an organization's trustworthiness, we should not trust it. Regulators can see that they cannot see how raters arrive at their judgements. Unlike methods in medicine, what credit ratings do to determine credit risk is something at which we can only guess. In such a case trust should be suspended. We simply lack the information we need to place our trust rationally. Outsourcing epistemic responsibility to parties that keep their methods secret flies in the face of common sense.

This is not the only reason why outsourcing epistemic responsibility is misplaced. Determining trustworthiness is also made difficult by a second phenomenon: *ratings shopping*. The idea is simple. An issuer of a security applies for a rating to each of the three main agencies, compares the ratings and decides to publish the most favourable rating only. Ratings shopping, it seems, occurred quite widely. In an interview with *Wall Street Journal* reporters in 2008, Brian Clarkson, Moody's Investors Service President at the time, said: 'There is a lot of rating shopping that goes on . . . What the market doesn't know is who's seen certain transactions but wasn't hired to rate those deals.'⁶³ Even if rating agencies were entirely epistemically virtuous, ratings shopping would likely lead to inflated and untrustworthy ratings. Given sufficiently complex securities, even experts exercising epistemic virtue will disagree about credit risk. We find this in health care too, when medical specialists disagree in 'hard cases'. The public can accommodate differences whenever all views are made public and are easily accessible. In the case of ratings, however, issuers only publish the most favourable rating.⁶⁴ This makes it impossible to compare ratings and as a result this leads to a systematic upward bias among published ratings.⁶⁵

That regulators allow issuers of securities to shop for the best rating is, besides the methodologies being trade secrets, a serious obstacle to perceiving the trustworthiness of the agencies. A third argument against

⁶³ Quoted by Lucchetti, 'Bond-rating shifts'.

⁶⁴ Skreta and Veldkamp, 'Ratings shopping'.

⁶⁵ Griffin et al., 'Rating shopping or catering?'.

outsourcing epistemic responsibility to rating agencies is their issuer-pays compensation model. Several authors view this as a source of conflicts of interest. This may be too harsh. A conflict of interest arises, in John Boatright's useful definition, whenever 'a personal or institutional interest interferes with the ability of an individual or institution to act in the interest of another party, when the individual or institution has an ethical or legal obligation to act in that other party's interest'.⁶⁶ For there to be a conflict of interest in the present situation, credit rating agencies must have an ethical or legal obligation to act in the interest of potential buyers and sellers of rated securities. It is not clear, however, that such obligations exist. Legal obligations they probably do not have. Courts grant them First Amendment protection of free speech. Ethical obligations may follow from the fact that particular investors are by regulation forced to rely on the ratings, but the analogy with astrology, together with the fact that these investors have the resources to research credit risks themselves, does not make this immediately evident. It is hardly plausible to maintain that when a government decides to enforce laws obligating pilots to rely on their horoscopes instead of meteorologists when it comes to weather forecasts, this places ethical obligations on the astrologers writing the columns.

Even though the case for conflicts of interest is weak, the issuer-pays model still endangers trustworthiness because it leads to a situation where the interests of the testimonial sources of information and the recipients of information are not aligned. A recent article by John Griffin, Jordan Nickerson and Dragon Yongjun Tang addresses this issue under the heading of *ratings catering*.⁶⁷ Ratings catering is related to ratings shopping in the sense that it happens when issuers request ratings from more than one rater, but it differs in that the assumption of rater honesty is lifted. In the model of Griffin and his colleagues, rating agencies adjust initial ratings upwardly (and dishonestly) when the issuer shows that competing agencies have rated the security more favourably. The sample includes 716 collateralized triple A debt obligation tranches that were rated by Standard and Poor's and Moody's in the period 1997–2007, so the usual caveats apply. The conclusion is that a lenient Standard and Poor's is likely to be followed by Moody's,

⁶⁶ Boatright, 'Conflicts of interest', 219.

⁶⁷ Griffin et al., 'Rating shopping or catering?'

and vice versa. This is a consequence of the issuer-pays compensation scheme, which decreases the ratings' trustworthiness.

Outsourcing epistemic responsibility

We outsource epistemic responsibility more often than not. We rely on the judgements of accountants, legal advisers, doctors, consumer organizations and so on because we do not have the time, the skills and the money to do all the research ourselves. This is not wrong; testimony is an acceptable source of knowledge. But we should choose our sources of testimony with care, and when governments designate particular sources as the sole or ultimate source of information this is only justified if their trustworthiness is beyond rational doubt. It may be that some of the effects of outsourcing epistemic responsibility surveyed in this chapter are not as easy to detect as I suggest. Without knowledge of empirical research, for instance, it is not immediately evident that there is a mismatch between what ratings are claimed to express and how investors interpret them. No theoretical sophistication is needed, however, to see that one should not place trust in organizations whose methods one cannot check and compare with others.

The argument I develop here may still appear convoluted. It may be objected that I have only shown that outsourcing epistemic responsibility does not foster virtue without making the claim that this is wrong. It may be said that although legislators have made investors increasingly dependent on the published 'opinions' of credit rating agencies, one may object to blaming the agencies for a lack of generosity and interlucency. I think one can always defend the *prima facie* case in favour of epistemic virtue. Unless one is playing a game, and nothing else, one's claims should be backed by evidence obtained in epistemically virtuous ways. This applies to astrologers too. But I have a different aim. In Chapter 5 we saw that companies can help their employees practise epistemic virtue along three lines: virtue-to-function matching, organizational support for virtue and organizational remedies against vice. This chapter shows in a sense that, like companies, governments too influence epistemic virtues among citizens and companies. I did not develop a theory of how regulators can encourage epistemic virtue because the strategies they can use are very similar to the strategies that companies have at their disposal. Rather I looked at credit rating agencies. They have received ample criticism from commentators, and I do not wish to

downplay the relevance of the critiques. But if we are to blame a party in the first place, our blame should be directed at those governments that forced investors to outsource credit risk assessment to companies of which the trustworthiness is hard to determine.

Summary

Chapters 3 and 5 looked into a number of conceptual and empirical issues to do with individual and corporate epistemic virtues. I defended a view of epistemic virtue as instrumentally contributing to *eudaimonia*, and I analysed corporate epistemic virtue in terms of virtue-to-function matching, organizational support for virtue and organizational remedies against vice. The present chapter continued this investigation by looking at other-regarding epistemic virtues. But it also did something else. Not until the present chapter had I asked the question of whether we can normatively expect individuals or corporations to care for virtue. It may be quite nice to possess epistemic virtues as a character trait, but what could be the justification of requiring others to practise them, or to criticize others if they do not? It is true that from a job description epistemic virtues often readily follow; the minimal normative assumptions about the purpose of a firm, however, barred the derivation of corporate virtue from corporate purpose. If a corporation is merely a nexus of voluntary contracts of equal and freely consenting people, what reason could we have to blame them for running their business foolishly? They will soon be pushed out of the market by more virtuous competitors.

It is important to see that the applicability of the theory of epistemic virtues – and corporate epistemic virtues in particular – is independent of the minimal assumptions I prefer to make. Many commentators hold on to the view that banks are there to safeguard the private property of citizens and to foster their freedom in line with recent ethical ideals of corporate citizenship. If that is your view, then the case for epistemic virtue in finance is made more quickly.

But not too quickly. In this chapter I defended the claim that even though the credit rating agencies were far removed from being exemplars of epistemic virtue, government regulators deserve even harsher epistemic criticism. The analogy between horoscopes and credit ratings was perhaps a bit over the top. Yet it did serve the purpose of showing the recklessness of outsourcing epistemic responsibility to corporations

of which the testimonial trustworthiness regulators had not cared to examine sufficiently thoroughly. The more general lesson was that for it to be safe to outsource epistemic responsibility to a corporation with regards to a particular subject matter (assessment of credit risk in the case of the rating agencies), we have to ascertain two things. Firstly, of course, that the corporation is a trustworthy source of information concerning the subject matter. Part of the task here is also to establish that the corporation adds any informational value in the first place. This is all very plain, but already at this stage the regulators failed to pass the test. But secondly, we must be confident that outsourcing responsibility to the corporation will not have undesirable side effects. In the case of the rating agencies, outsourcing responsibility did have such effects, one of which was a lower than desirable degree of epistemic competition: when three American agencies are in the position to give official stamp of approval assessments of credit risk, what incentive would you have to assess these risks for yourself?

Next to the credit rating agencies and the stock market analysts discussed earlier in this chapter, accountants are viewed as an important source of information to financial markets. Recent and not so recent accounting scandals may suggest that outsourcing epistemic responsibility to accountants should meet similar scepticism. The next chapter examines this question. By doing so, it also places other-regarding epistemic virtues – the main theme of this chapter – in a more constructive light.