

Stakes sensitivity and credit rating: a new challenge for regulators

Article (Published Version)

Booth, Anthony and de Bruin, Boudewijn (2019) Stakes sensitivity and credit rating: a new challenge for regulators. *Journal of Business Ethics*. 0-0. ISSN 1573-0697

This version is available from Sussex Research Online: <http://sro.sussex.ac.uk/id/eprint/86267/>

This document is made available in accordance with publisher policies and may differ from the published version or from the version of record. If you wish to cite this item you are advised to consult the publisher's version. Please see the URL above for details on accessing the published version.

Copyright and reuse:

Sussex Research Online is a digital repository of the research output of the University.

Copyright and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable, the material made available in SRO has been checked for eligibility before being made available.

Copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational, or not-for-profit purposes without prior permission or charge, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



Stakes Sensitivity and Credit Rating: A New Challenge for Regulators

Anthony Booth¹ · Boudewijn de Bruin²

Received: 10 April 2018 / Accepted: 31 July 2019
© The Author(s) 2019

Abstract

The ethical practices of credit rating agencies (CRAs), particularly following the 2008 financial crisis, have been subject to extensive analysis by economists, ethicists, and policymakers. We raise a novel issue facing CRAs that has to do with a problem concerning the transmission of epistemic status of ratings from CRAs to the beneficiaries of the ratings (investors, etc.), and use it to provide a new challenge for regulators. Building on recent work in philosophy, we argue that since CRAs have different stakes than the beneficiaries of the ratings in the ratings being accurate, what counts as knowledge (and as having ‘epistemic status’) concerning credit risk for a CRA may not count as knowledge (as having epistemic status) for the beneficiary. Further, as it stands, many institutional investors (pension funds, insurance companies, etc.) are bound by law to make some of their investment decisions dependent on the ratings of officially recognized CRAs. We argue that the observation that the epistemic status of ratings does not transmit from CRAs to beneficiaries makes salient a new challenge for those who think current regulation regarding the CRAs is prudentially justified, namely, to show that the harm caused by acting on a rating that does not have epistemic status for beneficiaries is compensated by the benefit from them acting on a CRA rating that does have epistemic status for the CRA. Unlike most other commentators, therefore, we offer a defeasible reason to drop references to CRAs in prudential regulation of the financial industry.

Keywords Credit rating agencies · Epistemology · Information and knowledge · Regulation

Introduction

In the aftermath of the global financial crisis, various governmental committees, journalists, economists, and ethicists have scrutinized the ethics of the credit-rating agencies (CRAs) (e.g., De Bruin 2017; Pagano and Volpin 2010; Scarlet and Kelly 2012; Strier 2008). *Guardian* journalist Joris Luyendijk has called the CRAs the ‘all purpose bogeyman’ of the global financial crisis, and Paul Krugman has even gone as far as saying that their ratings are ‘literally worse than useless’.¹ The Financial Crisis Inquiry Commission,

created by the US government in 2009, has rather boldly suggested that CRAs such as Standard & Poor’s, Fitch, and Moody’s were a necessary contributing cause of the crisis - claiming that the crisis ‘could not have happened without the rating agencies’ (p.xxv). And even if these comments may be slightly hyperbolic, the complaints that led to that damning assessment are well-known and numerous: CRAs were completely re-miss with their ratings of structured finance instruments (asset-backed securities, etc.); they were much too slow to adjust ratings in the light of new data (it is striking, for instance, that in the days directly preceding Lehman Brothers’ demise on 15 September 2008, the largest CRAs still gave it good ratings); and their techniques as *tying* (where a CRA provides an unsolicited rating for a fee) and *notching* (where a CRA issues low ratings on asset-backed securities unless further assets constituting those securities are also rated) are seen to be ‘anti-competitive’ (Langhor and Langhor 2009). Note, finally, that it is the issuers of the to-be-rated securities, not the beneficiaries of the ratings, that

✉ Boudewijn de Bruin
b.p.de.bruin@rug.nl

Anthony Booth
a.r.booth@sussex.ac.uk

¹ Department of Philosophy, School of History, Art History, and Philosophy, University of Sussex, Brighton BN1 9QN, UK

² Faculty of Philosophy, Faculty of Economics and Business, University of Groningen, Nettelbosje 2, 9747 AE Groningen, The Netherlands

¹ www.theguardian.com/commentisfree/joris-luyendijk-banking-blog/2012/dec/17/rating-agencies-bogeymen-william-j-harrington;krugman.blogs.nytimes.com/2013/02/23/little-statesmen-andphilosophers.

pay the CRAs; many authors have claimed that this underlying *issuer-pays* remuneration model leads to unavoidable conflicts of interest (Bolton et al. 2007).

In this paper, we raise a completely new kind of issue facing the CRAs that has to do with a problem about calculating risk and communicating risk judgements. We argue that CRA judgements of creditworthiness (the ratings) suffer from what we call *stakes sensitivity*, a phenomenon recently described by philosophers.² Since CRAs have different stakes than the beneficiaries of the ratings in the ratings being accurate, what counts as knowledge concerning credit risk for a CRA may not count as knowledge for the beneficiary. The so-called *epistemic status* of a rating is, then, not necessarily transmitted from the CRA to its beneficiary. We argue that while this is not necessarily an ethical failure on the part of the CRAs, it does present a challenge for the regulation of CRAs as we know it. Currently, many institutional investors (pension funds, insurance companies, etc.) are permitted to invest in securities above a certain threshold rating, set by law, where the rating is to be determined by a dozen officially recognized CRAs. A result of this is that good ratings are effectively ‘stamps of approval’ to invest (White 2010). We argue that if the law is to continue to make investment decisions dependent on the CRAs in contexts where stakes sensitivity arises then it is incumbent on the defenders of the law to demonstrate that the harm this may cause investors is compensated by the benefit from investors acting on CRA ratings that have epistemic status for the CRA. We are sceptical that this challenge can be met, but we leave the possibility that it can open in this paper.

Unlike such authors on this issue as Scalet and Kelly (2012), we therefore, argue for a potential bit of deregulation, and as such our paper may well buttress arguments put forward by Richardson and White (2009), and De Bruin (2017).

The structure of this article is as follows. Introducing some background facts about CRAs and their ratings, Sect. 2 argues that, alongside two other roles they play, ratings are to be seen as judgements concerning credit risk (default risk) of issuers of rated securities. In Sect. 3, we introduce some new ideas that have recently emerged in philosophy. We argue that CRAs operate in a context in which stakes sensitivity ensures that a knowledge claim about credit risk made by a CRA may not lead to the ratings’ beneficiaries’ possessing knowledge about credit risk. Section 4 concludes by examining the policy repercussions of our diagnosis.

Judging Risk

CRAs rate debt. Using letters ranging from AAA (almost riskless) to the D of default, CRAs publish their judgements on the credit risk of debt issued by companies, countries, provinces, territories, water authorities, municipalities, and, more recently, of structured finance instruments such as mortgage-backed securities and collateralized debt obligations.³ They are often claimed to wield enormous power. The *New York Times* journalist Thomas Friedman (1995) presciently observed more than two decades ago that ‘the U.S. can destroy a country by levelling it with bombs, [but] Moody’s can destroy a country by downgrading its bonds’. To illustrate, it is unsurprising, for instance, that when Standard & Poor’s downgraded the US on 5 August 2011 (the first downgrade of the US debt ever), the Obama administration undertook a concerted media campaign to dissuade investors from taking the downgrade seriously. Dovetailing this point is the fact that, following the 2008 financial crisis, part of the UK government’s justification of austerity measures (spending cuts and tax rises) was to prevent a downgrading of the UK ratings. And the CRAs not only backed this up, but explicitly endorsed the government policy. In September 2010, for instance, Moody’s claimed that ‘[t]he stable outlook on the UK’s AAA rating is largely driven by the Government’s commitment to stabilise and eventually reverse the deterioration in its financial strength’.⁴

When we look back at the history of the CRAs, we see that they existed in an earlier guise as so-called *credit reporting firms* (Olegario 2006; Poon 2012; Sinclair 2008). After the US Civil War, as trade distances increased, merchants became increasingly unable to always evaluate for themselves their business partners’ credibility and creditworthiness. So they had to resort to third-party information. Credit reporting firms tapped into a demand for such information—a demand that was large because a great deal of trade at the time involved trade credit. Credit reporting firms therefore started collecting information from sheriffs, cashiers, and other reliable ‘correspondents’, and made this available to business people to assist them in estimating the trustworthiness of individual merchants, small- and mid-sized

² The terminology is ours. Philosophers use terms such as ‘subject sensitivity’, ‘interest relativity’, ‘context sensitivity’.

³ Credit risk is the risk that an issuer fails to pay interest and/or does not pay back the loan, and is generally distinguished from market risk, liquidity risk, country risk, etc. The methods used to rate these bonds differ in obvious respects. Our argument applies to the three main categories of assets that CRAs rate: government bonds, corporate bonds, and structure debt. Moreover, ratings of government bonds (ratings of countries, provinces, etc.) also capture an estimation of the risk that the willingness to pay deteriorates as the result of particular political developments. See White (2010) for an overview.

⁴ www.telegraph.co.uk/finance/economics/8013617/Austerity-plan-helps-UK-retain-top-AAA-rating.html.

businesses, and sometimes even large corporations, as witnessed by Henry Varnum Poor's well-known 1860 *History of Railroads and Canals of the United States* (Poor 1860).

For a long time, the main business model of credit reporting firms was the following: they collected information from correspondents, transmitted it to their clients, who then made their own judgements about creditworthiness ultimately based on that information. But in the 1910s, this model underwent a change the importance of which we can hardly overestimate. The credit-reporting firms now started publishing their own assessments of creditworthiness, introduced the letter system that we have today, and rapidly expanded into bond rating. That is to say, credit reporting firms no longer left it to the users of their information to make their own judgements concerning credit risk; instead, they provided these users with *their* judgements about credit risks.⁵

The importance of this change in business model has been highlighted by various authors (e.g., Olegario 2006; Sinclair 2008). It is also key to the argument we develop in this paper. It makes salient the fact that the CRAs are not so much informational intermediaries, transmitting information to beneficiaries who subsequently make up their own minds. On the contrary, CRAs are themselves responsible for evaluating and interpreting the information they collect, and for aggregating the information into one judgement about credit risk (captured in the lettered scores from AAA to D). Unlike the earlier credit *reporting* firms, what CRAs have published since the beginning of the twentieth century is best understood as *judgements* that are the result of the agencies' *own* evaluation, interpretation, and aggregation of information (Olegario 2006). A rating therefore represents a knowledge claim, or at the very least a claim carrying 'epistemic status' (we explain this later).

A second and equally important step was to follow, however. Two years after the crash of 1929, the US Office of the Comptroller of the Currency decided to regulate on the basis of a distinction between *investment* and *non-investment* grade bonds. It required banks to treat the latter as more risky, and to make a clear distinction between the two types of securities on their balance sheets. Half a decade later an outright proscription against investing in what is now called *speculative* (non-investment) bonds followed. In itself, this may seem a reasonable measure, aimed at protecting US citizens against the effects of a next crisis. What is important for our argument, however, is that what counted as non-investment or speculative was to be determined solely by a handful of officially recognized credit rating agencies. As

a leading expert on credit rating agencies describes it, the CRAs' judgements of creditworthiness 'attained the force of law' (White 2010, p. 213). It is important to underscore that these were measures primarily directed at financial stability, which concerns the financial system as a whole, and in particular its ability to absorb shocks. Protection of investors was not the main reason why ratings obtained the force of law.

And they never seem to have lost that force again. In 1975, a registry of Nationally Recognized Statistical Rating Organizations was created, selecting the few CRAs that may play a role in prudential regulation, and even in periods of *deregulation*, the judgements of CRAs have been legally binding to many investors. As Pagano and Volpin (2010) observe, when in 1989 it became possible for pension funds to invest in asset-backed securities, the conditions stipulated that they must only buy highly rated such securities. In line with White (2010), these authors therefore conclude that a result of the regulation was that the investment decisions of banks, pension funds, insurance companies, and other institutional investors are crucially restricted by what judgements CRAs arrive at.⁶

While this does not mean that institutional investors were legally mandated to fully 'outsource' their judgements of creditworthiness in all circumstances (De Bruin 2017), economists (see White 2010 for a survey) and sociologists (see Poon 2012 for a survey) observe that as a result of these regulatory developments ratings currently do a number of different things. Following the economics and sociology of finance literature we identify the following three functions.

- (1) Ratings express *judgements* about creditworthiness.
- (2) Ratings are used in *monitoring* issuers of rated securities.
- (3) Ratings can figure as *directives* to investors.

To begin with the first, a rating stands for the CRA's *judgement* about the default risk attached to a particular security. The standard way in which CRAs describe ratings is as 'opinions'. Standard & Poor's webpage devoted to explaining what a rating is, for instance, states: 'Our ratings express our opinion about the ability and willingness of an issuer'.⁷ Authors such as Partnoy (1999) believe, however, that the use of the word 'opinion' is primarily meant to lay off liability for potential negative consequences of flawed ratings. Staging themselves as publishers of opinions on a par with news media such as newspapers, radio, and

⁵ CRAs also produce *outlooks* and *watchlists*. We ignore these instruments here, but come back to them later when we discuss the implicit contract between issuers and agencies.

⁶ Internal documents and contracts may also specify conditions referring to CRAs. We are here only concerned with restrictions that are set by regulators.

⁷ www.spratings.com/en_US/understanding-ratings.

television (or the press more generally) has so far successfully ensured that CRAs enjoy far-reaching First Amendment (freedom of expression) protection.

In fact, there is a great deal of evidence that in more technical contexts (documents, briefings, etc.) CRAs do describe themselves as ‘judging-to-be-true’ rather than ‘opining’. Describing the final stages of the rating procedure, Standard & Poor’s write, for instance, that they use ‘forward-looking analysis and analytic *judgment* to determine the ultimate rating conclusion with the goal of transparency and rating comparability’⁸ [our emphasis]. Secondly, specifically regarding the way in which they arrive at a rating for banks, Moody’s lists a number of factors that ‘inform [their] judgment about the level of support willingness for each major debt class’, and observe that the rating committee ‘may then employ its *judgment* of the specific circumstances in question to assign a given number of notches of support’ [our emphasis].⁹

For a statement to count as a judgement (or as the result of judging) it has to be based on careful consideration of relevant evidence. It involves assessing the merits of potential sources of information, appraising the strength of available evidence, settling for particular analytical tools to apply to the case at hand, and so forth. Judging has a decisional element to the effect that when one judges something to be the case, one has deemed the evidence backing up the judgement sufficient to warrant one to proceed on the basis of the assumption that it is the case. And, indeed, if CRAs did not issue such judgements (and only expressed genuine opinions), one might reasonably wonder whether the role they play in regulation is at all legitimate; for, in the language we explain in Sect. 3, ratings need to have *epistemic status* in order for them to be of any legitimate use, which roughly means that they need to be more than just a ‘wild guess’. The upshot of this is that we should understand CRAs to be providing investors with a judgement about creditworthiness.

The second way we can understand credit rating is to be *monitoring* the issuers of the rated securities. Economists such as Boot et al. (2006) have found evidence of what they call *implicit contracts* that are present between many CRAs and issuers. The CRAs can inform issuers (through published statements including so-called ‘watchlists’ and ‘outlooks’) about conditions that, if satisfied, will ensure that the security is not downgraded (or is upgraded). Issuers honouring this ‘contract’ will then be assured that their rating remains the same.¹⁰ Boot et al. (2006) give the following

instructive example. In the early years of the new millennium, telecom companies passed through a number of ratings downgrades that were accompanied by statements to the effect that a CRA—Moody’s in this case—left them a period of about one and a half years to reduce their debt in order to prevent another downgrade. And where Boot et al. (2006) focus only on corporate debt, the turmoil surrounding the rating of US debt illustrates how similar (yet often much less explicit and outspoken) contracts may operate between CRAs and other issuers of debt. For instance, after Standard & Poor’s downgraded the US in 2011, Moody’s in 2012 announced that it would follow if Congress failed to reach a deal to reduce the budget deficit—a fairly explicit formulation of a ‘contract’ in the sense of Boot et al. (2006).

Following White (2010), the third way in which we can understand credit ratings has to do with the role that ratings play with respect to institutional investors. We call this its *directive* function. As we saw, legally binding guidelines for institutional investors such as pension funds and insurance companies often only allow investments in securities with sufficiently high ratings. This has a striking effect on investment behaviour. There is robust evidence that investors generally respond to downgrades, but not to upgrades (White 2010). If ratings were *merely* passing on judgements about credit risk to investors, one would expect investors to respond to downgrades and upgrades alike: from a financial point of view, one should *ceteris paribus* adjust one’s investment strategies just as much when credit risk increases as when it decreases.¹¹ That downgrades count more heavily is therefore generally seen as supporting the view that ratings do not just reflect credit risk, but also give directions to investors—a role they would ultimately not have were it not for legislation (White 2010). Downgrades close to the legal threshold, that is, are directives to sell.¹²

Footnote 10 (continued)

outlooks to specify the conditions that might lead to a downgrade or upgrade.

¹¹ We note in passing that we are well aware of the fact that investment decisions are not only determined by considerations of risk and expected return, but also by such factors as a concern for diversification, or positive and negative screening in socially responsible investing. But these factors do not explain the asymmetrical responses to downgrades and upgrades.

¹² The extent to which a rating is used as a directive depends on whether the rating is close to the legal threshold. The pronounced response by investors to downgrades is, however, also visible for higher rated debt (see Gonzales et al. 2004 for a literature review). Clearly, an investor’s response to a downgrade also depends on other factors, such as the role the asset plays in the portfolio’s diversification, internal risk management procedures, etc. Close to the legal threshold, however, a downgrade will trump most of these considerations, though.

⁸ www.spratings.com/documents/20184/774196/Corporate+Ratings+Methodology.pdf/8fd4392a-4aae-4669-bd74-a9b86e18d781.

⁹ www.moody.com/microsites/brm/SP29913_AnalyticalApproach.pdf?WT.z_referringsource=BRM~Overview.

¹⁰ Boot et al. (2006) do not intend *contract* to refer to something explicitly written down. Rather a contract implicitly arises due to the fact that CRAs constantly monitor firms *and* use watchlists and

So far we have largely focused on ratings of corporate and government debt. But CRAs obtain a significant part of their revenues from rating structured finance instruments. Here it seems that there cannot be an implicit contract between the CRA and the issuer, because once the security has been designed, there is nothing the issuer can do to prevent the risks of a downgrade.

To set the stage: A typical mortgage-backed security starts with a mortgage lender (the *originator*) providing funds to house owners. An investment bank (called *arranger*) buys a collection of mortgages from the originator, and creates a Special Purpose Vehicle that buys these mortgages from the arranger, securitizing them (issuing mortgage-backed securities). The CRAs, then, provide ratings of (tranches of) these securities, which the investment bank can then sell to investors.¹³

Analysing the information CRAs use to arrive at these judgements, Pagano and Volpin (2010) observe that until 2007 Moody's only considered information from the prospectuses and annual reports provided by the Special Purpose Vehicle to arrive at their ratings. These documents do not provide information concerning individual mortgages and borrowers (such as the loan-to-value ratio or the borrower's credit score), despite the fact that, as the authors show, such information is essential to get a reliable rating. Hence they conclude that at least until 2007 (when the subprime mortgage meltdown started), Moody's used decidedly suboptimal rating strategies for important classes of structured debt. They offer no reasons why Fitch and Standard & Poor's should be expected to be different.

One might think that this shows that the ratings of structured securities are intended to play a different role than ratings concerning corporate and government bonds. As we remarked above, it is certainly true that ratings of structured debt are unlikely to have a monitoring function. The Special Purpose Vehicle is not an entity that can change policies in response to ratings, as firms and states can do, so it seems that there can be no implicit contract between it and the CRA.¹⁴ But, nonetheless, these ratings are *judgements* of creditworthiness, and they offer *directives* to investors. Ratings of structured debt are meant to capture credit risk in exactly the same way as they capture credit risk of

corporate and government bonds. For instance, Standard & Poor's writes that their 'ratings represent a uniform measure of credit quality globally and across all types of debt instruments... [i]n other words, an "AAA" rated corporate bond should exhibit the same degree of credit quality as an "AAA" rated securitized issue' (quoted by Pagano and Volpin 2010, p. 407). (Standard & Poor's was wide off the mark, but that is a different story.) Similarly, as Pagano and Volpin (2010) show, deregulation in the 1980s allowed pension funds to invest in asset-backed securities on the condition that they buy highly rated securities only. These authors persuasively argue that the ratings of structured debt contributed significantly to investor demand for such securities. They observe that even as late as 2008, for instance, there were more than sixty thousand structured finance instruments of the highest possible rating, but there were only twelve triple A rated corporations. It is unsurprising that with such 'stamps of approval', as White (2010) calls the ratings, many institutional investors bought structured finance.

Let us take stock. By using insights from the economics and sociology literatures on CRAs we have seen that because of historical and legal developments, CRAs as we know them today have three functions: they publish judgements about credit risk; they monitor issuers through implicit contracts; and they provide directives to institutional investors. For the time being we have not made any normative judgement regarding their practices, but having made our target clear in terms of what credit ratings do, we now move to raise a problem for the CRAs brought into relief by recent research in philosophy.

Stakes Sensitivity of Knowledge Transmission

There is widespread agreement among philosophers that we tend to be less inclined to attribute knowledge to people with regard to a proposition when there is a lot at stake as regards to its the truth.¹⁵ To give a slightly simplified example for

¹³ See Pagano and Volpin (2010) for an excellent more detailed exposition.

¹⁴ CRAs have been accused of having played a role in the *design* of certain structured finance instruments, with the goal of constructing a security that *minimally* qualifies for a particular rating. One might consider this to be an implicit contract in the sense of Boot et al. (2006), because the CRA somehow commits itself to giving the security a particular rating. But we believe that the temporal aspect that is essential to the model of Boot et al. (2006) is missing here (the issuer cannot change the security after it has been designed). That is why we are inclined not to consider this to be an instance of monitoring.

¹⁵ One of the first studies in contemporary epistemology pointing this out was DeRose (1992). See, e.g., DeRose (2011) for a more up to date case and survey of its reception. However, the issue of how practical matters affect what is known is a very old one in philosophy and may, arguably, go all the way back to Plato. Certainly, the issue is important to thinkers in the so-called 'continental tradition' as Nietzsche, Foucault and Habermas, as well as contemporary thinkers in philosophy of science working on inductive risk (for a good overview see Elliot 2017), and of course the well-known 'underdetermination (of evidence on theory choice) thesis' can be read as showing how practical values are ultimately guiding what is known (e.g. Kuhn 1974; Laudan 1984). We focus on the contemporary epistemology literature, which given its focus on *thresholds* for knowledge illustrates the problem with CRAs best.

ease of exposition, ask yourself: do you know whether you locked your car this morning in the car park? Suppose that you answer in the affirmative. Then suppose a colleague tells you, and you believe her, that she saw two youths vandalizing an apparently unlocked car in the car park just a few minutes ago. If you experience even a little more reluctance to say that you know whether you locked the car once your colleague tells you how high the stakes may be, then you have experienced the intuition in question. For the sake of exposition in this article, we call this phenomenon *stakes sensitivity*. (A more precise definition follows below.)

Most philosophers accept that any good analysis of knowledge ought to take stakes sensitivity into account.¹⁶ There are broadly three kinds of response called *contextualism*, *subject-sensitive invariantism*, and, as mentioned, *purist invariantism*.¹⁷ It should be noted that the latter (*purist invariantism*) is a way of maintaining the traditional view of knowledge, despite stakes sensitivity (where the ‘traditional view of knowledge’ is one according to which what, and how much, is at stake makes no difference at all to what one knows). This should hopefully clarify to people unfamiliar with this literature that ‘stakes sensitivity’ itself is not taken to be controversial—it is how one responds to it that is debated. We will now seek to briefly describe these views, aiming to steer a middle course between making the material intelligible while appealing to the general reader, and doing justice to the specialists’ concerns.¹⁸

According to *contextualism*, what explains stakes sensitivity of knowledge transmission is that the semantic value of the word ‘knowledge’ (and its correlates) is context-sensitive (e.g., DeRose 1992; Lewis 1996). Context-sensitivity is a very familiar idea among linguists, and the semantic value of words such as ‘tall’ are uncontroversially considered to be context-sensitive. The idea is that just as it is not right to call someone ‘tall’ in a context which ‘contains’ only taller people (even though in other contexts she may be adequately described as ‘tall’), it is accurate to say that you do not know whether you locked the car in a context where you have just heard that two youths may be vandalizing your car.

A different way of responding, which is going to be the focus of this article, is to appeal to what has become known

as *subject-sensitive invariantism* (Fantl and McGrath 2009; Hawthorne 2013; Stanley 2005, 2008). According to this view, the semantic value of the word ‘knows’ does not vary depending on the context it is used in (hence ‘invariantism’ instead of ‘contextualism’). Yet the *standards* for knowledge are sensitive to whether the stakes are high or low for a given knowing subject. Subjects for whom the stakes are high with respect to some proposition must satisfy higher evidential demands in order to possess knowledge of the proposition than subjects for whom the stakes are low. An illustrative analogy is to compare standards that professionals have to meet. While ‘coastal guard’ in Britain and Australia refers to the same profession, the standards that coastal guards have to satisfy differ, because Australian waters contain other dangers than British (sharks, for instance).

For *purist invariantism*, finally, to account for stakes sensitivity, we need an explanation in terms of pragmatics or conversational implicature, rather than one grounded in observations concerning semantics (thought to be different from pragmatics) or standards of knowledge (Brown 2005; Reed 2010). The view is that we are less likely to attribute knowledge in high-stakes situations because by uttering the word ‘know’ in such situations we would often *pragmatically* (as opposed to *semantically*) entail certain things we do not wish to be held accountable for. In high-stakes situations, we are less inclined to assert that the subject should *act* on their knowledge about some proposition, and so this makes us less likely to want to say that the subject knows a proposition; for if we said that they knew, we would be pragmatically implying that the subject should be allowed to act on their knowledge.¹⁹

While this looks like a very promising response, contextualists typically reply that this violates the plausible principle to the effect that we should assert propositions only if we know them. This is captured by the so-called *knowledge-based norm of assertion* (as opposed to, say, a belief-based norm of assertion). Williamson (1996, 2000), probably the most cited proponent of such a knowledge-based norm of assertion, thinks that this norm is constitutive of any speech act. That is, to the extent we take ourselves to be saying something, we must take our utterance to be governed by a knowledge-based norm of assertion. Every assertion we

¹⁶ Some ‘experimental’ philosophers cast doubt on the claim that members of the general public have the same intuitions as professional philosophers about these kinds of case (e.g., Buckwalter 2010; Buckwalter and Schaffer 2015). The methodology of experimental philosophy (as well as its putative results in this case) have been subject to criticism (Nagel 2012; Mortensen and Nagel 2016).

¹⁷ For current purposes, we set aside variants of these three strands. See, e.g., DeRose (2011) for details.

¹⁸ We set aside various issues, including the question of whether one needs to distinguish a first-person and a third-person version of stakes sensitivity. See, e.g., DeRose (2011).

¹⁹ One might even think that rational choice theory (game theory) accords with this move, because on such an analysis it might seem that what is rational depends in part on what a subject’s utility function looks like, which among other things reflects the subject’s attitude towards risk. Given two different utility functions, but the same beliefs (based on the same evidence), it may be rational to perform two different actions. It is currently an open question in the literature whether we should understand stakes in this way, though, and we do not want to take a stand on this issue here. Anderson and Hawthorne (forthcoming) evaluate various different, formal ways of modelling stakes are explored, and find all wanting.

make is, then, a knowledge claim in the sense that you must take yourself to know it. Contextualism and the knowledge-based norm of assertion are, then, ‘two sides of the same coin’ (DeRose 2002, p. 169). Subject-sensitive invariantists, in turn, reply to the pure invariantist that knowledge rather figures in a *norm of action* on the grounds that it is counterintuitive to hold on to the view that if you know a proposition this might give you insufficient warrant to act on the proposition (Fantl and McGrath 2009). As a result, the purist invariantist thus either has to deny both of these plausible principles (the norm of assertion and the norm of action), or else try to show how her theory is compatible with them after all.

The claims about credit ratings we defend in this paper go through on all three views—since all accept the linguistic data we are calling stakes sensitivity. Yet for ease of exposition, we focus on how subject-sensitive invariantism explains the stakes sensitivity of knowledge transmission. Our argument therefore proceeds from the following principle:

(Knowledge Ascriptions) The higher the stakes are with respect to a proposition *p*, the more evidence is required for an agent *S* to know that *p*.

An immediate consequence of this principle is:

(Epistemic Status Transmission) If *S*’s stakes with respect to a proposition *p* are higher than *T*’s stakes with respect to *p*, and *T* knows *p*, then *S* may not know that *p* even if *S* has the *exactly the same evidence* as *T*.

Two clarificatory points are in order before we move to discuss how stakes sensitivity affects CRAs. First, we assume that agent *S* can be a group or an institution just as well as an individual. Secondly, while the principles above are parsed in terms of what *S* ‘knows’, it is meant to apply not only to knowledge in a strict sense, but to any state that has *epistemic status*, as it is called in the literature.

Let us explain the term. Epistemic status is an (always positive) appraisal of a cognitive attitude (knowledge, understanding, belief, wisdom, etc.) on the grounds of its satisfying certain conditions that make it justified, warranted, rational, etc. (We also use epistemic status in a derived sense to apply to what expresses these cognitive attitudes. In that sense, ratings can have epistemic status.) Something may have epistemic status, but yet fail to be knowledge, since one may have a justified, true belief that is not knowledge. Yet for a belief to have epistemic status it must be more than just a ‘wild guess’ or an unsupported opinion, even if by sheer coincidence the guess turned out to be true. Only *sufficient* evidence endows a belief with epistemic status. If you wildly guess on the basis of no evidence that Managua is the capital of Nicaragua your belief counts as having no epistemic status, even though it is true. A vague memory

that it is will give you *some* evidence, but not enough to give your belief that Managua is the capital of Nicaragua epistemic status. This is why your belief here cannot have *some* epistemic status—it either has it, or it does not. A friend’s telling you that it is Managua may give your belief epistemic status, but if a lot depends on it (a colleague asked you to meet in the capital of Nicaragua and you are about to book a flight), you will need stronger evidence. This illustrates that the higher the stakes for *S* are with respect to *p*, the higher is the evidential threshold required for *S* for her belief that *p* to have epistemic status. Moreover, epistemic status is relative to the agent. My friend may have visited Nicaragua in the past, and so her belief has epistemic status. As long as I do not know this fact about my friend (and do not know that she can be trusted as a source of information about this bit of geography), my belief may not have epistemic status. It is easier for us from now on to refer to this more technical notion of epistemic status. The reader may wonder why we talked about ‘knowledge’ at all and not began the paper immediately talking about epistemic status. The reason is that stakes sensitivity describes a linguistic phenomenon, and very few people in everyday English outside philosophy departments talk about epistemic status. The ordinary person uses the word ‘knowledge’ to describe what we are talking about.

In the preceding section, we argued that when a CRA issues a rating we must interpret that as the CRA’s decision to form a judgement concerning credit risk. If the ratings were just wild guesses or merely unsupported opinions we would have to very seriously question their purpose, and it would be straightforwardly inappropriate for regulators to bind institutional investors to the ratings. In other words, for the CRA’s judgements of creditworthiness to be of any use to anyone, they need to have epistemic status in the above sense of the word. This means that they have to be based on relevant and sufficient evidence. To our knowledge, even when they claim that they are offering ‘opinions’, no CRA has hitherto denied that their ratings need to be based on relevant and sufficient evidence. This is underscored by such things as that CRAs make significant investments in research methodology. The question surely is whether the evidence that informs them is such that it confers epistemic status to what they say, that is, brings it closer to satisfying the requirements of genuine knowledge. There is, then, no reason to assume that the ratings are not subject to both the norm of assertion and the norm of action.

In line with the debate, two alternatives could at first sight be open. Recalling Krugman’s claim, quoted in the introduction, CRAs might be ‘literally worse than useless’; if that is true, ratings would be more akin to a wild guess. We set this alternative aside, however, for present purposes. If Krugman’s verdict is true, then this offers a clear and immediate argument against CRAs, making our argument here

superfluous. Moreover, as Richardson and White (2009) has observed, whether CRAs actually add value to the market can be considered to be a genuinely open research question. So we assume for the sake of our argument that ratings do have epistemic status, and hence are really evidence-based judgements.

Now to begin with, and in line with what CRAs themselves acknowledge, the following scenario may be true:

(Case 1) A CRA issues a rating based on (given what is at stake for the CRA) relevant and sufficient evidence to a user who acts on the assumption that the rating is accurate, but who ends up making a bad financial decision because, consistently with the CRAs judgement as well as the used evidence, a low probability event materializes.

A second alternative may also obtain, though. It is this scenario that we want to highlight in this paper, and which is not acknowledged by CRAs (or regulators for that matter).

(Case 2) A CRA issues a rating based on (given what is at stake for the CRA) relevant and sufficient evidence to a user for whom (given what is at stake for the user) the same evidence is not relevant and sufficient.

Currently, CRAs and regulators seem to be sensitive only to Case 1 type of problems ratings may run into. CRAs explicitly warn investors, for instance, that their ratings are relative probability judgements, and that things with low likelihood may still happen. So even as relative probability judgements, users must take their ratings with caution, and not be taken as a direct recommendations. As Fitch put it in their Terms of Use:

[it] is *not* making any *recommendation* or suggestion, directly or indirectly to you, or any other person, to buy, sell, make or hold any investment, loan or security or to undertake any investment strategy with respect to any investment, loan or security of any issuer' [our emphasis].²⁰

It is important to note, however, that there is a difference between saying that their ratings are not *recommendations*, and saying that their ratings have no *epistemic status*. It is clear that no CRA wants investors to believe that their ratings are *not* based on relevant and sufficient evidence—if they had no epistemic status, we would have no more reason to consult the CRAs than a Magic 8-Ball, so to speak.

We assume, then, that in Case 1 and 2, the ratings have epistemic status. Yet a difference remains; for unlike in Case 1, in Case 2 types of situation, the epistemic status of the CRA rating does *not* transmit to give the rating epistemic

status for the beneficiary of the rating. To see this, consider an institutional investor such as a pension fund. Clearly, the stakes in getting a correct risk assessment on a financial product they are thinking of buying or selling are much higher for the fund than for a CRA carrying out the rating, as their investment decisions affect their financial position and those of their clients (members). We are not denying of course that the CRA, in this particular case, comes to its judgement (the rating) based on relevant and sufficient evidence to meet the threshold required for it to have positive epistemic status *from the CRA's vantage point*. What we do observe, though, is that were the pension fund to marshal *exactly the same* evidence for themselves, that evidence would *not* meet the higher threshold entailed by the higher stakes. As such, the very same bits of evidence will yield, from the pension fund's point of view, a belief about credit risk without epistemic status.

Before proceeding to illustrate how this may lead to the pension fund's taking suboptimal financial decisions, it is important to note that we are not claiming that transmitting epistemic status is *always* bound to fail. The amount and quality of evidence that CRAs have to back their ratings differ. Some corporations or countries may provide more reliable information than others, which translates into the ratings being more reliable. In such cases, even though the stakes of the CRA and its beneficiaries are different, the evidence the CRA possesses is sufficient to give the rating epistemic status to the beneficiaries as well. If CRAs have evidence that is only *minimally sufficient* given their own stakes, knowledge will, however, not transmit to its beneficiaries.

When beneficiaries act on ratings that have little or no epistemic status for them, this may lead to suboptimal financial decisions. Consider the pension fund example again. Schematically, there are two scenarios here, depending on whether the fund would arrive at a higher or lower rating if it were to judge credit risk based on the same bits of evidence. If the fund's own rating were higher, acting on the basis of the CRA's rating leads to overly cautious investment behaviour. This type of situation can be witnessed when CRAs give ratings above but close to the legal threshold (close to 'speculative' or 'non-investment' grade). As we saw above, in that type of situation and typically after a downgrade, institutional investors bound by law display selling behaviour that, if they acted on their own judgement only, they would not have engaged in (White 2010).

In the other type of situation, if the fund's own ratings were lower than those of the CRAs, the fund is likely to assume risks that are inconsistent with its own risk management procedures. An example of this situation was given by Jorion (2009), who analysed purchases of structured securities at UBS, the Swiss bank. As we indicated above, the CRAs suggested that ratings are 'uniform' representations 'across all types of debt instruments', and that as a

²⁰ www.fitchratings.com/site/about/terms.

result, ‘an “AAA” rated corporate bond should exhibit the same degree of credit quality as an “AAA” rated securitized issue’ (quoted by Pagano and Volpin 2010, p. 407). UBS believed as much, for as Jorion (2009) shows, UBS assumed that highly rated tranches of certain asset-backed securities would behave as corporate debt with similar ratings. From the point of view of the CRA, there was sufficient evidence to back that claim (in the years before the 2008 crisis). Newly developed mathematical models were generally seen as highly reliable instruments to estimate these risks. The CRAs, however, would not stand to lose much if the ratings turned out to be less accurate. UBS, however, would lose a lot. (And indeed, it did suffer losses of \$19 billion in one year only as a result of taking the ratings as fully accurate.) As UBS’s stakes were higher, they would have needed additional evidence to back the rating. For instance, they would have needed an answer to the question of why, in 2008, there were around 64,000 structured debt instruments with a triple A rating as compared to only 12 corporate bonds globally (Pagano and Volpin 2010).

The situation is further compounded by the fact that the precise manner in which CRAs arrive at their judgements is not transparent to their external users: the algorithms they use are proprietary, and the resulting rating is informed by private information obtained through conversations with the issuer of the rated security. A result of this ‘privacy’ is that users of ratings cannot eliminate the possibility that they are *not* in a situation described in Case 2.

This underscores our claim that users are in a kind of double-jeopardy with respect to using CRA scores, exposed to two risks—namely, a risk from Case 1 to the effect that (i) a state of affairs, which the rating suggests has low probability, might still obtain; and a risk from Case 2 to the effect that (ii) the rating has no epistemic status for the users. As we argued above, the caveats currently issued by CRAs address the risk inherent in (i). They do not address the risk inherent in (ii), however, and *a fortiori* they do not address the issue of how that risk might synergize with the risk inherent in (i).

We envisage that one could respond by reverting to the claim that the CRAs simply provide ‘mere opinions’, supporting an interpretation along the following lines. Surely ratings are nothing more than opinions, as they have no epistemic status for the user, but all the same, because they do have epistemic status for the CRA, they still provide the users (investors) with some evidence at least. And while that evidence may not be sufficient to grant the user epistemic status for adopting a belief in the rating alone, the rating can certainly be a corroborating force in the midst of other bits of evidence available to the user.

We do think that this response has some force. The trouble is, however, that current legislation is such that it makes the response practically unavailable to a significant class of investors, namely, those investors that are bound by law to

heed the ratings. As we saw above, since the 1930s banks, insurance companies, pension funds, and other institutional investors have been obliged to divest from securities (or to take specific measures when investing in securities) with ratings below a certain threshold. The threshold is determined by the regulator, but the ratings are to be obtained from a handful officially recognized CRAs such as Standard & Poor’s, Fitch, and Moody’s. These investors do not have the option of interpreting a rating as just a bit of evidence concerning credit risk amidst other available evidence. They are forced to act on these ratings, specifically when securities are downgraded and obtain ratings close to the threshold. While the objection has some force, then, our response is that it is only rational (and legitimate) to endow ratings with such legal status if the epistemic status of ratings entirely transmits to investors. This of course then raises the question of whether it is appropriate to legislate in a way that gives CRA ratings what White (2010) described as ‘force of law’. We now turn to this issue in our next, and last, section.

A Challenge for Regulators

So far we have argued that ratings are, among other things, judgements about credit risk that direct law-bound institutional investors. Moreover, we have argued for the following three claims:

- (1) There is a likely misalignment between the stakes that CRAs have in their ratings being adequate and those of the investors using the ratings.
- (2) As a result, particularly since investors have no way of knowing that the extent to which stakes sensitivity impacts the adequacy of the ratings, this leads to a failure to transmit epistemic status from CRAs to their beneficiaries.
- (3) These transmission failures are material in relation to investment decisions—that is, they are of substantial import to the investors.

This section concludes our argument by reiterating our main policy conclusion, and defending it against two potential objections. Our policy conclusion is that if current regulation is to continue, regulators must first meet the following challenge: demonstrate that the harm caused by acting on a rating that does not have epistemic status for investors is compensated by the benefit from investors acting on CRA ratings that have epistemic status for the CRA (but not for the investor).

To bring this into relief, let us address two important ways one may be minded to respond to our case. The first is to wonder whether the users of CRA ratings see the practical value of the CRA ratings in terms of how accurate the

ratings model creditworthiness. An interesting example goes as follows: an investment bank selling a mortgage-backed security is legally obliged to have it rated by an accredited CRA. So it is not only potential investors interested in buying the security that will use the rating, but in an important sense the issuer of the security also counts as a user. One could now point out that the issuer is very likely to have already made a risk assessment using their own methods, and that getting a CRA to rate the security is practically useful in so far as it allows the issuer to issue the security (and because of the correlated attractiveness to potential investors). So one could say that the rating is of practical value to the issuer only in an indirect way.

This is certainly true, but we should like to respond by noting that, first, the rating's having this indirect practical value must be parasitic on the fact that regulation has made it a legal requirement for the bank to obtain the rating. If that requirement did not exist, the issuer would be completely free to choose what structured securities it would offer for sale. Further, surely the ratings having force of law gives them a legitimacy in the eyes of users that they would not otherwise have. As Pagano and Volpin (2010) show, the large appetite among investors for complex structured finance products witnessed before the subprime mortgage meltdown broke out in 2007 is to an important extent explained by their having obtained ratings from established CRAs. Consequently, the right question to ask when we are seeking to assess the practical value of CRA ratings is whether the existing legal requirements succeed as pieces of regulation.

The main reason put forward in this paper to answer this question negatively is that the epistemic status of ratings is not necessarily transmitted to investors using the ratings. Of course this reason may be defeasible, such that *all things considered* the regulation should be maintained. While we are open to this possibility, we are sceptical that this is the case if we bring to bear other problems with CRAs—so the challenge cannot be immediately ignored. Regulators are (or should be) well aware of the various complaints that ethicists and other commentators have voiced against CRAs, ranging from conflicts of interest in the issuer-pays model to such unacceptable tactics as tying, notching, and rate shopping (where an issuer of a security obtains ratings from various CRAs but only publishes the highest rating). These and other complaints, voiced by such authors as De Bruin (2017), Pagano and Volpin (2010), Scalet and Kelly (2012), and Strier (2008), are all to the effect that the ratings are unreliable and/or that investors are not in a good position to ascertain the reliability. While most of these commentators favour stricter regulation of CRAs, the additional problem we have highlighted in this paper makes it less probable that further regulation will be able to address all these issues. The stakes of CRAs and beneficiaries are often extremely misaligned,

and stricter regulation will not alter that. In such an environment, we believe, questions need to be asked whether it is right for the regulator to force investors to rely on ratings in some of their investment decisions. More transparency about how CRAs come to their ratings will help, but it will not fix things given the divergence of stakes. Moreover, we should recall the history of CRAs. As we saw, they developed from credit reporting agencies, who provided information about the creditworthiness of American businesses in the nineteenth century. Business people, investors, and other interested parties used the information as a basis for their own *judgements* concerning credit risk. Only around 1910 the letter system as we know it today was introduced. As Sinclair (2008, p. 24) writes: 'The transition between issuing compendiums of information and actually making judgments about the creditworthiness of debtors occurred between the 1907 financial crisis and the Pujo hearings of 1912'. Without this development, the issue of stakes sensitivity would not have arisen. An equally important step was taken in the 1930s when, as we saw, the US government gave credit ratings the 'force of law' (White 2010, p. 213). Without this second step, stakes sensitivity would not necessarily have become a concern for regulators. We hope to have shown that given current legislation it does.

The other important response to our policy conclusion that we should like to address is to ask the following: cannot CRAs assume a kind of mean (or median) stakes situation when making their assessments, and make it explicit that they do so? This would give a CRA rating epistemic status to any user, who was then able to assess the mean stakes situation vis-à-vis their own. We think this would be impractical, however, given the difficulty in assigning blanket accurate mathematical values to stakes. And even if this were possible, we would likely end up in an infinite regress: an answer to the question of whether the stakes in a given situation for a given user are 'average' one would itself be subject to stakes sensitivity. As such, we would have to predict variation regarding what counts as assessments about 'average' stakes situations with epistemic status within subjects for whom more, or less, was at stake with respect to getting the assessment right. This leads to a vicious regress, which we certainly should avoid.

To recapitulate, finally, we have shown that ratings are judgements concerning credit risk, and insofar as judgements are knowledge claims, they will display the phenomenon of stakes sensitivity: one and the same rating may have epistemic status for the CRA, but not for the investor, because their stakes are different. This is not to say that stakes sensitivity cannot be mitigated. Small 'boutique' CRAs can offer rating services that specifically take investor stakes into account, but such CRAs are organized in very different ways than the Nationally Recognized Statistical Rating Organizations whose ratings many institutional

investors are by law forced to follow. In other words, we do believe that ‘bespoke’ bond ratings can incorporate as much as possible the stakes of a specific investor. There is no principled impossibility to transmit epistemic status. But epistemic status transmission will generally fail with the ‘off-the-peg’ ratings offered by the CRAs that form the topic of this paper. As such, we conclude that it is imperative that regulators begin to seriously re-think the parameters under which which CRAs are given force of law.

Acknowledgements We would like to thank audiences at the Society of Applied Philosophy annual conferences in St Anne’s College, Oxford, 2014, and Utrecht, 2018. Special thanks are owed to Marco Meyer, Alex Oliver, Tom Simpson, Magdalena Smith, Chris Thompson, and Jens van ’t Klooster.

Funding This study was funded by Nederlandse Organisatie voor Wetenschappelijke Onderzoek 360-20-310.

Compliance with Ethical Standards

Conflict of interest Booth and De Bruin declare that they have no conflict of interest.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Anderson, C., & Hawthorne, J. (2019, forthcoming). Knowledge, practical adequacy, and stakes. *Oxford Studies in Epistemology* 6.
- Bolton, P., Freixas, X., & Shapiro, J. (2007). Conflicts of interest, information provision, and competition in the financial services industry. *Journal of Financial Economics*, 85(2), 297–330.
- Boot, A. W. A., Milbourn, T. T., & Schmeits, A. (2006). Credit ratings as coordination mechanisms. *Review of Financial Studies*, 19, 81–118.
- Brown, J. (2005). Adapt or die: The death of invariantism. *Philosophical Quarterly*, 55, 263–285.
- Buckwalter, W. (2010). Knowledge isn’t closed on Saturday: A study in ordinary language. *Review of Philosophy & Psychology*, 1(3), 395–406.
- Buckwalter, W., & Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Noûs*, 49(2), 201–234.
- de Bruin, B. (2017). Information as a condition of justice in financial markets: The regulation of credit-rating agencies. In L. Herzog (Ed.), *Just financial markets? Finance in a just society* (pp. 250–270). Oxford: Oxford University Press.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52, 913–929.
- DeRose, K. (2002). Assertion, knowledge, and context. *Philosophical Review*, 111(2), 167–203.
- DeRose, K. (2011). *The case for contextualism: Knowledge, skepticism and context*. Oxford: Oxford University Press.
- Elliot, K. C. (2017). *A tapestry of values: An introduction to values in science*. Oxford: Oxford University Press.

- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*. Oxford: Oxford University Press.
- Friedman, T. L. (1995). Foreign affairs: Don’t mess with Moody’s. *New York Times*, A19, 1.
- Gonzales, F., et al. (2004). Market dynamics associated with credit ratings: A literature review. *Banque de France Financial Stability Review*, 4, 53–76.
- Hawthorne, J. (2013). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Jorion, P. (2009). Risk management lessons from the credit crisis. *European Financial Management*, 15(5), 923–933.
- Kuhn, T. (1974). *Objectivity, value judgement, and theory choice in the essential tension*. Chicago: University of Chicago Press.
- Langhor, H., & Langhor, P. (2009). *The rating agencies and their credit ratings: What they are, how they work, and why they are relevant*. London: Wiley Finance.
- Laudan, L. (1984). *Science and values: The aims of science and their role in scientific debate*. Berkeley: University of California Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Mortensen, K., & Nagel, J. (2016). Armchair-friendly experimental philosophy. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy*. Oxford: Wiley-Blackwell.
- Nagel, J. (2012). Intuitions and experiments: A defence of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3), 495–527.
- Olegario, R. (2006). *A culture of credit: Embedding trust and transparency in American business*. Cambridge, MA: Harvard University Press.
- Pagano, M., & Volpin, P. (2010). Credit ratings failures and policy options. *Economic Policy*, 25(62), 401–431.
- Partnoy, F. (1999). The Siskel and Ebert of financial markets? Two thumbs down for the credit. *Washington University Law Quarterly*, 77(3), 619–714.
- Poon, M. (2012). Rating agencies. In K. Knorr Cetina, A. Preda (Eds.) *The Oxford Handbook of the Sociology of Finance* (pp. 1–16). Oxford: Oxford University Press.
- Poor, H. (1860). *History of railroads and canals of the United States*. New York: J.H. Schultz.
- Reed, B. (2010). A defence of stable invariantism. *Nous*, 44(2), 224–244.
- Richardson, M., & White, L. (2009). The rating agencies: Is regulation the answer?. *Financial Markets, Institutions & Instruments*, 18(2), 146–148.
- Scalet, S., & Kelly, T. F. (2012). The ethics of credit rating agencies: What happened and the way forward. *Journal of Business Ethics*, 111(4), 477–490.
- Sinclair, T. J. (2008). *The new masters of capital: American bond rating agencies and the politics of creditworthiness*. Ithaca: Cornell University Press.
- Stanley, J. (2005). *Knowledge and practical interests*. Oxford: Oxford University Press.
- Stanley, J. (2008). Knowledge and certainty. *Philosophical Issues*, 18, 35–57.
- Strier, F. (2008). Rating the raters: Conflicts of interest in the credit rating firms. *Business and Society Review*, 113(4), 533–553.
- White, L. J. (2010). Markets: The credit rating agencies. *Journal of Economic Perspectives*, 24(2), 211–226.
- Williamson, T. (1996). Knowing and asserting. *Philosophical Review*, 105, 489–523.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.