

ETHICS AND ASYMMETRY: SKIN IN THE GAME AS A REQUIRED HEURISTIC FOR ACTING UNDER UNCERTAINTY

C. Sandis & N.N. Taleb

Abstract: We propose a global and mandatory heuristic that anyone involved in an action that can possibly generate harm for others, even probabilistically, should be required to be exposed to some damage, regardless of context. We link the rule to various philosophical approaches to ethics and moral luck.

Early draft for comments - please do not cite

1. Acting Under Uncertainty

The chances of informed action and prediction can be seriously increased if we better comprehend the multiple causes of ignorance. The study of ignorance, then, is of supreme importance in our individual and social lives, from health & safety measures to politics and gambling (Rescher 2009). But how are we to act in the face of all the uncertainty that remains after we have become aware of our ignorance? The idea of *skin in the game* is crucial for the well-functioning of a complex world. In an opaque system there is, alas, an incentive for operators to hide risk: to benefit from the upside when things go well without ever paying for the downside when one's luck runs out. There is no possible risk management method that can replace skin in the game —particularly when informational opacity is compounded by informational asymmetry viz. the principal-agent problem that arises when those who gain the upside resulting from actions performed under some degree of uncertainty are not the same as those who incur the downside of those same acts. For example, bankers and corporate managers get bonuses for positive 'performance', but do not have to pay out reverse bonuses for negative performance. This gives them an incentive to bury risks in the tails of the distribution, thereby delaying blowups.

The ancients were fully aware of this incentive to hide risks, and implemented very simple but potent heuristics (for the effectiveness and applicability of fast and frugal heuristics both in general and in the moral domain, see Gigerenzer, 2010). About 3,800 years ago, Hammurabi's code specified that if a builder builds a house and the house collapses and causes the death of the owner of the house, that builder shall be put to death. This is the best risk-management rule ever. The ancients understood that the builder will always know more about the risks than the client, and can hide sources of fragility and improve his profitability by cutting corners. The foundation is the best place to hide such things. The builder can also fool the inspector, for the person hiding risk has a large informational advantage over the one who has to find it. The same absence of personal risk is what

motivates people to only appear to be doing good, rather than to actually do it. The problems and remedies are as follows:

First, consider policy makers and politicians. In a decentralized system, say municipalities, these people are typically kept in check by feelings of shame upon harming others with their mistakes. In a large centralized system, the sources of error are not so visible. Spreadsheets do not make people feel shame. The penalty of shame is a factor that counts in favour of governments (and businesses) that are small, local, personal, and decentralized versus ones that are large, national or multi-national, anonymous, and centralised.¹ When the latter fail, everybody except the culprit ends up paying the cost, leading to national and international 'austerity'.²

Second, we misunderstand the incentive structure of corporate managers. Counter to public perception, corporate managers are not entrepreneurs. They are not what one could call agents of capitalism. Since 2000, in the United States, the stock market has lost (depending how one measures it) up two trillion dollars for investors, compared to leaving their funds in cash or treasury bills. It is tempting to think that since managers are paid on incentive, they would be incurring losses. Not at all: there is an irrational and unethical asymmetry. Because of the embedded option in their profession, managers received more than four hundred billion dollars in compensation. The manager who loses money does not return his bonus or incur a negative one. The built-in optionality in the compensation of corporate managers can only be removed by forcing them to eat some of the losses.

Third, there is a problem with academic economists, quantitative modellers, and policy wonks. The reason economic models do not fit reality is that economists have no disincentive and are never penalized by their errors. So long as they please the journal editors, their work is fine. We use models such as portfolio theory and similar methods without any remote empirical reason. The solution is to prevent economists from teaching practitioners. Again this brings us to decentralization by a system where policy is decided at a local level by smaller units and hence in no need for economists.

Fourth, the predictors. Predictions in socioeconomic domains don't work. Predictors are rarely harmed by their predictions. Yet we know that people take more risks after they see a numerical prediction. The solution is to ask —and only take into account— what the predictor has done (what he has in his portfolio), or is committed to doing in the future. It is unethical to drag people into exposures without incurring losses. Fifth, to deal with warmongers, Ralph Nader has rightly proposed that those who vote in favour of war should subject themselves (or their own kin) to the draft.

We believe *Skin in the game* is the heuristics of a safe, just, and society. Opposed to this is the unethical practice of taking all the praise and benefits of good fortune whilst disassociating oneself

¹ This argument against 'big government' is not an argument against the state providing welfare but an argument against doing so in a centralized (as opposed to municipal) fashion that enables people to hide behind bureaucratic anonymity.

² See McQuillan (2013) and Orr (2013).

from the results of bad luck or miscalculation. In what follows we situate our view within the framework of ethical debates relating to the moral significance of actions whose effects result from ignorance and luck. In what follows demonstrate how the idea of skin in the game can effectively resolve debates about (a) moral luck and (b) egoism vs. altruism, whilst successfully bypassing (c) debates between subjectivist and objectivist norms of action under uncertainty, by showing how their concerns are of no pragmatic concern.

2. Symmetrical Constraints in Moral Philosophy

The skin in the game heuristic is best viewed as a rule of thumb that places a pragmatic *constraint* on normative theories. Whatever the best moral theory (consequentialism, deontology, contractualism, virtue ethics, particularism etc.) or political ideology (socialism, capitalism, libertarianism) might be, the 'rule' tells us that we should be suspicious of people who appeal to it to justify actions that pass the cost of any risk-taking to another party whilst keeping the benefits for themselves.

At the heart of this heuristic lies a simple moral objection to negative asymmetry that lies at the heart of some of the oldest and most famous moral ideas, as illustrated in Figure 1 below.

Rules of Moral Symmetry				
<p>1. Lex Talionis: 'An eye for an eye, a tooth for a tooth' (Exodus 21.24)</p>	<p>2. 15th Law of Holiness and Justice: 'Love your neighbor as yourself' (Leviticus 19.18)</p>	<p>3. Silver Rule: 'Do not do unto others what you would not have them do unto you' (Hillel the Elder)</p>	<p>4. Golden Rule: Do unto others as you would have them do unto you' (Matthew 7:12)</p>	<p>5. Formula of the Universal Law: 'act only in accordance with that maxim through which you can at the same time will that it become a universal' law.' (Kant 1785: 4:421)</p>

Fig. 1. Moral Symmetry

Of course the clearest examples of any rule are likely to stem from a deontological approach, but the skin on the game constraint is not committed to deontology. Indeed, moral symmetry is one of the key ideas behind many forms of social contract theory (e.g. 'I scratch your back, you scratch mine'), and different emphases on symmetry may also be found in consequentialism (which places the overall good above that of the agent) and virtue ethics (which looks for an ethical mean between excess and deficiency).

As worded, all of the principles in the table above are problematic. Take, for example, fourth principle of reciprocity in fig.1 above. This 'golden rule' seems to suggest that if I would like you to come up and kiss me then I should go up to you and kiss you (regardless of whether *you* would like this). But while the precise principles may be faulted, the spirit of symmetry behind them (and arguably every moral tradition)³ contains much insight. Indeed, the very plausibility of Derek Parfit's recent attempt to demonstrate that the best versions of the most popular normative theories converge (Parfit 2012), must ultimately hang upon a common spirit of this kind. As we shall we shall see, however, there can be positive asymmetries in our behaviour, as well as negative ones.

3. Altruism vs. Egoism

Psychological Egoists claim that we always do what we most desire (Mandeville: 1714). Those who believe in the possibility of altruism tend to either deny this (Nagel 1970) or to distinguish between self-centred desires and the desire to benefit others (Butler 1726). So while it is not false to think that whether or not we ever act altruistically is an empirical question (Slote 1964), we must not forget that its answer will partly depend upon a priori distinctions between notions such as those of 'desire', 'motivation', 'reason', and so on. It is such distinctions, rather than experimental research, that allow us to recognise that while anyone who is not a sociopath will feel contentment in helping others, only a pervert would help others *in order* to acquire this feeling (Sandis 2012:75; cf. Broad 1930).

The most pragmatic way of distinguishing between egoists and altruists is to ask whether someone has ever voluntarily (a) paid a cost for someone else's benefit or (b) been willing to reap the rewards of risk whilst passing the cost to another. The first, altruistic, action is one where the agent has skin in another person's game (Taleb 2013:375) The second, egoistic act, is one where the person has no skin in the game. People we call 'saints' are frequently disposed to act in the former way. Those who tend to act in the latter way we typically call 'assholes'. In reality, most of us are neither saints nor assholes: most of the time have skin in our own game and nobody else's. On occasion, however, even the most average of people is liable to either slip up or rise to the occasion. Such moments are marked by asymmetry, be it of a negative or positive kind (See fig. 2)

³ See Blackburn (2001:101).

1. No skin in the game	2. Skin in the game	3. Skin in someone else's game
Selfish/egoistic Negative asymmetry Individualistic Morality	neither egoistic nor altruistic Symmetry (neutral) Conventional morality	Selfless/altruistic Positive asymmetry Other-based morality

Fig. 2 Egoism vs. Altruism

The middle column in the table is the largest because most of the actions of the average person tend to fall within it. It is not wonder, then, that the 'eye for an eye' reciprocity it epitomises is - for better or worse - a conventional morality. To its left lies the sort of individualistic morality frequently associated with Nietzsche but most clearly ascribable to 'rational' egoism of Ayn Rand (1964) and others who maintain that 'greed is good'. To its right lies the morality of self-sacrifice. This comes in all sorts of stripes: Christian, socialist, utilitarian, and so on. Needless to say, these divisions are never as sharp in practice as they are in theory. Rand's egoistic heroes, for example, subscribe to the symmetrical thought that one should never demand that others take a risk one wouldn't take oneself. Conversely, most welfare states are run by bureaucrats with no skin in the game. Both sides are fooling themselves.

The symmetrical constraint entails that we act wrongly when we open ourselves to great harm that could have reasonably been foreseen and avoided, but the wrongness isn't a moral one. We act immorally when we open *others* to great risk but are only willing to be considered as responsible for our actions if the risk turns out not to harm anyone. Such actions involve the malignant transfer of fragility and anti-fragility from one party to another with the aim of getting any possible benefits of our actions without being liable for any possible harms (Taleb 2012:375). This agency problem is an asymmetry.

Those who responsible for such transfers (most predictive analysts, bankers, bureaucrats, consultants, editors, politicians, risk vendors, and sophists) attempt to justify their hypocrisy by

appealing to bad luck and uncertainty. They offer excuses of the 'we acted on information we believed was correct at the time' or 'obviously it fell way short of expectations' but refuse to accept any liability for their actions and protest wildly at the mere thought that they should pay the cost. These may be contrasted with who have skin in the game viz. those who take risks for themselves and keep their downside. Typical examples are activists, artisans, citizens (as opposed to '*idiotes*'), entrepreneurs, traders, and writers. The greatest contrast, however, is with the who put their own skin in the game for the sake of others. We call such people heroes and saints but they include not only knights and warriors but also some maverick artists, journalists, scientists, and writers who put their livelihood reputations on the line for the sake of others (Taleb 2012: 377). This all brings us to the so-called 'problem of moral luck' In the next section we show why the correct approach to this evaluative problem is bad news for those who live with no skin in the game.

4. Moral Luck

Consider the case of two equally reckless drivers, only one of which kills a pedestrian. According to Bernard Williams the unlucky driver is morally guilty of something worse than the other driving (namely manslaughter). According to Kant, both drivers are at most only liable for reckless driving. Both views are confused.⁴ What we should say is that from the moral point of view, a certain kind of reckless driving is as bad as manslaughter. When a person drives recklessly they *take upon them* the risk of manslaughter and so are responsible for it if it happens, and for opening themselves up to it (which is morally just as bad) if it doesn't (see Sandis 2010). Hegel got it right, then when he wrote not only that 'The laurels of mere willing are dry leaves that never were green' but also:

It happens of course that circumstances may make an action miscarry to a greater or lesser degree. In the case of arson, for instance, the fire may not catch or alternatively it may take hold further than the incendiary intended. In spite of this, however, we must not make this a distinction between good and bad luck, since in acting a man must lay his account with externality. The old proverb is correct: 'A flung stone is the devil's'. To act is to expose oneself to bad luck. Thus bad luck has a right over me and is an embodiment of my own willing. (Hegel, *Philosophy of Right*, § 119A).

We are not only responsible for *known* characterizations of our actions and their effects but also for those which we *ought* to be aware of (even if we are not). Our ignorance does not always relieve us of responsibility for things we have done, because others can claim that, as rational beings we *should*

⁴ Kantians claim that the realm of what we are responsible for starts and ends with the will, everything else (including all outward behaviour) in principle lying out of our control. They are opposed by Humeans such as Bernard Williams who argue that there is so much luck in the world (in particular luck relating to our characters, the situations we find ourselves in, and the effects of our actions) that we have to accept the counter-intuitive view that whether or not one is morally liable for an action is largely a matter of luck.

have known what we were doing even if we did not. Such is the knowledge involved in putting other peoples' lives at risk with no skin (of our own) in the game.⁵

In sum, 'bad luck' is no excuse when it could have been reasonably foreseen. Foresight should not be restricted here to a particular event. If I know that 1/1000 actions of type A will have a tragic result it is not acceptable to perform thousands of these actions on the grounds that for each one there is only a probability of 1/1000 that something will go wrong. The greater the potential disaster the smaller the probability has to be for an act that could bring it about to be immoral. There is an inverse symmetry between the acceptable probability of risk and the weight of the potential damage being assessed.

All action must be understood as *exposition* to luck and to be judged accordingly. When we take a risk we cannot wash our hands of the consequences on others and hide behind masks of expectation, intention, ignorance, luck, uncertainty, and so on. The asymmetry in taking risks without having skin in the game is an unethical one. Any system deemed 'too big to fail' not only encourages but demands that we live according to such skinless asymmetry. The real black swan event of the 21st century is not that any financial crisis occurred (which was predictable) but that there was no full-blown revolution against the governments which continue to encourage 'idiotes' to gamble with other peoples' lives and money. In the next and final section we demonstrate how philosophical debates about how to act under uncertainty seem to ignore the simple lessons of the symmetrical approach.

5. Objectivism vs. Subjectivism

The ethics of risk is frequently thought of as a branch of moral philosophy concerned with abstract principles that tell us how we ought to act when we lack (or do not know whether or not we lack) information that is relevant to our choice (e.g. Altham 1984 who makes a technical distinction between mere risk and general uncertainty). Far from being infrequent, such scenarios are the norm and can only be excluded in controlled thought experiments. In an important sense, then, all acts are performed under uncertainty, which is not to say that we know what the consequences of our actions will be (see Dancy 2002: 233). This raises the deontic problem of how we ought to act in the face of known ignorance. As we shall see, the ethics of having skin in the game bypasses this problem, revealing it to be pragmatically irrelevant.

⁵ Hegel's solution famously offers two aspects of any given act: *Tat* (deed) corresponding to the objective (which I am causally responsible for), and *Handlung* (action) corresponding to the subjective (which can be morally imputed to me); rights relating to the latter in turn dividing into one relating to various elements of the self such as knowledge, intention, and purpose (PR 115, 117, & 120; see also 118A).

The worry, as raised in H.A. Prichard's influential paper 'Duty and the Ignorance of Fact', is that of whether a person's obligation to perform (or omit from performing) some action depends 'on certain characteristics of the situation in which he is, or on certain characteristics of his thought about the situation' (Prichard 1932: 84). Objectivists (such as Sidgwick and Parfit; cf. Ross 1930) claim that we ought to do whatever is *in fact* best, even when we cannot be reasonably expected to know what this is. By contrast Subjectivists (including Ross and Prichard 1932; cf. Ross 1939) claim that we ought to do whatever we *believe* will be best.

The difficulty of choosing between these positions is supposed to be that on the one hand, we want to leave room for the thought that we can be *wrong* about what we ought to do. The fact that what we *believe* we ought to do and what we *actually* ought to do can come apart in this way would seem to lend credence to objectivism. On the other hand, there is the procedural obstacle of the impossibility of stepping out of one's own mind in order to compare reality with one's impressions of it. Thus the objective view would seem to imply the absurd view that 'although we may have duties, we cannot know but only believe that we have; and therefore we are rendered uncertain whether we, or anyone else, has ever had, or will ever have a duty' (1932:89).⁶ A parallel absurdity is implied in this rhetorical question posed by Jonathan Dancy: 'Suppose that, unknown ...to me, someone has been buried alive in my garden during the night. Could this make it wrong of me to go away for a fortnight's holiday?' (Dancy 2000:57). As we saw in the previous section, moral obligation is constrained not by descriptive epistemology but by normative epistemology: it is not a mere question of what we know or could come to know but of what we *ought* to know, a truth mirrored in the popular legal thought that ignorance is no excuse. Prospectivists, most prominently Michael Zimmerman, attempt to avoid this dilemma by arguing that we ought to perform whichever action it is most reasonable to *expect* will be the best.

The above debates are merely academic and have no pragmatic weight. In fact, all three views share the common mistaken assumption that they are each motivated by the same notions of 'what one ought to do' when there are actually three different concepts at play:

- (i) **Objectivists** equate *what* we ought to do with whichever action turns out to be best. This is what we should *aim* at when we act.
- (ii) **Subjectivists** equate what we ought to do with whatever we judge to be best. This the only way *through which* we can aim at what is best.
- (iii) **Prospectivists** equate what we ought to do with what we can rationally expect to be best. This view attempts to reconcile objectivist and subjectivist intuitions that are only in tension because of the aforementioned assumption.

⁶ Ross (1939) rightly (but for the wrong reason) suggests that objectivists and subjectivists are talking at cross purposes. Cf. Zimmerman (2008: 1-2).

Whereas Objectivists are concerned with the rightness of the things we do (typically thought to be universals), Prospectivists and Subjectivists are concerned with the rightness of our acts of *doing* these things (typically thought to be particulars). Yet it is possible that one rightly acts in doing something that turns out not to be best and, by the same token, that one acts wrongly in doing something that turns out to be best.

Given that one can do the right thing for the wrong reason, the deontic question of *what* the right thing to do is should therefore be distinguished from the evaluative question of when one is acting *rightly*. The latter is best answered via an account of how and when people and institutions are *liable* for choices they make under uncertainty. We have sought to answer this question (e.g. in the case of moral luck) via the skin in the game principle. This heuristic is not about action but about dispositions. Indeed, it relates directly to the virtue of being such that the system will not only *survive* uncertainty, randomness, and *volatility* but will actually benefit from it.⁷ Skin in the game heuristics follow directly from this principle of antifragility.

References

- Altham, J. E. J. (1984), 'Ethics of Risk', *Proceedings of the Aristotelian Society*, New Series, Vol. 84 (1983 - 1984), 15-29.
- Blackburn, S. (2001), *Ethics: A Very Short Introduction* (Oxford: Oxford University Press).
- Broad, C. D. (1930), *Five Types of Ethical Theory*. (London: Kegan Paul).
- Butler, J. (1726), 'Sermons I and XI' in his *Fifteen Sermons Preached at the Rolls Chapel*, in *The Works of Bishop Butler*, ed. J. H. Bernard (London: Macmillan).
- Dancy, J. (2000), *Practical Reality* (Oxford: Oxford University Press).
- _____ (2002), 'Prichard on Duty and Ignorance of Fact' in P. Stratton-Lake (ed.), *Ethical Intuitionism: Re-evaluations*. Oxford: Oxford University Press.
- Gigerenzer, G. (2010), "Moral Satisficing: Rethinking Moral Behavior as Bounded Rationality", *Topics in Cognitive Science* 2 (2010) 528–554
- Hegel, G. W. F. (1991), *Elements of the Philosophy of Right*, trans. H. B. Nisbet. Cambridge: Cambridge University Press.
- Kant, I. (1785), *Groundwork for the Metaphysics of Morals*, tr. A. Zweig, ed. T. E. Hill, Jr. & A. Zweig. (Oxford: Oxford University Press), 2002.
- Mandeville, B. (1714), *The Fable of the Bees* (London: J Roberts).
- McQuillan, M. (2013), 'The Severity of Austerity', *Times Higher Education*, 3-9 Jan, 30-35.
- Nagel, T. (1970), *The Possibility of Altruism* (New Jersey: Princeton University Press).

⁷ One can, of course, render this into a principle about action ('act in whatever way renders you anti-fragile') but such a principle treats anti-fragility as the ultimate end-in-itself whereas it is best to treat it as a property whose value is derived from its effects.

- Nussbaum, M. (2001), *The Fragility of Goodness: Luck and Ethics in Greek Tragedy and Philosophy*, rev. edition. Cambridge: Cambridge University Press.
- Orr, D. (2013), 'It wasn't labour who spent too much, it was the banks. How did we forget this?', *The Guardian*, Sat 5 Jan, p.39.
- Parfit, D. (2012), *On What Matters*, Vols. I & II (Oxford: Oxford University Press).
- Prichard, H. A. (2002), 'Duty and Ignorance of Fact' in *Moral Writings*, ed. J. MacAdam. Oxford: Oxford University Press.
- Rand, A. (1964), *The Virtue of Selfishness: A New Concept of Egoism* (New York: Signet).
- Rescher, N. (2009), *Ignorance: On the Wider Implications of Deficient Knowledge* (Pittsburgh: University of Pittsburgh Press).
- Ross, W. D. (1930), *The Right and the Good*, rev. ed. 2002, ed. P. Stratton-Lake.
- _____ (1939), *The Foundations of Ethics*. Oxford: Clarendon Press.
- Sandis, C. (2010), 'The Man Who Mistook his Handlung for a Tat: Hegel on Oedipus and Other Tragic Thebans', *Bulletin of the Hegel Society of Great Britain*, No. 62, 35-60.
- _____ (2012), *The Things We Do and Why We Do Them* (Palgrave Macmillan, 2012).
- Slote, M. A. (1964), 'An Empirical Basis for Psychological Egoism', *Journal of Philosophy* 61, 530-537..
- Taleb, N.N. (2013), *Antifragile: How to Live in a World We Don't Understand* (London: Allen Lane).
- Williams, B. (1993), *Shame and Necessity*. Cambridge: Cambridge University Press.
- Zimmerman, M.J. (2008), *Living with Uncertainty: The Moral Significance of Ignorance* (Cambridge: Cambridge University Press).