

Fable of the Deans: The Use of Market Norms in Academia

Jeppe von Platz

University of Richmond

1. Taylor's argument

In part III of *Markets with Limits*, James Stacey Taylor argues that “academic research should be primarily governed by academic, and not market, norms.”¹ Taylor’s argument gets its sting from the claim that academic research currently is governed mostly by market norms, a situation that might explain some of the problems in current academic research diagnosed and discussed in the earlier parts of his book. But Taylor’s main claim in part III is not about how academic research is done, but about how it should be done, *viz.*, that it should be governed primarily by academic norms. Taylor does not define the primacy academic norms should enjoy, but I take him to mean that academic norms (and attendant motives) should be the primary driver of what researchers do and how they do it.

Here is an overview of how I understand Taylor’s argument for this conclusion:

- A. The aim of academic research is and should be to advance understanding.
- B. Academic research conducted in accordance with academic norms advances understanding better than academic research conducted in accordance with market norms.
- C. Therefore, academic research should be primarily governed by academic norms.

If we assume that market norms and academic norms are exclusive options and are the only options worth considering (two questionable assumptions that I revisit below), then C follows from premises A and B.

¹ James Stacey Taylor, *Markets with Limits: How the Commodification of Academia Derails Debate* (New York: Routledge, 2022), 5.

The first premise can be interpreted in light of Hart's notion of a "general justifying aim."² The general justifying aim of an institution or practice is the aim for which it exists and in light of which it should be designed. In this vein, we can read Taylor as arguing that the reason we have a system of academic research – and incur the significant costs of maintaining such a system – is to advance our understanding. Our practices of academic research should, accordingly, be designed so as to best achieve this aim. In this essay I accept premise A, thus interpreted.

The second premise involves the concept of norms and relies on the distinction between academic and market norms. Practices and institutions are governed by systems of norms (some codified, some uncoded), and we have some control over these systems. The question is what norms we should seek to create and maintain – premise B answers this question.

A norm "identifies what type of behavior [...] would be appropriate to perform in a particular situation."³ The type of behavior here in question is not simply act-types, but involves their motives and social context. Taylor adopts a generic understanding of a norm as a "guideline for appropriate behavior."⁴ However, the distinction between market norms and academic norms is drawn in terms of the reasons motivating the prescribed behavior. Taylor follows Anderson in understanding market norms as "impersonal, egoistic, exclusive, want-regarding, and oriented to "exit" rather than "voice".⁵ Academic norms, by contrast, are governed by the "good internal to academic work [...] understanding."⁶ The distinction isn't entirely clear, and I'm unsure that Anderson's definition of market norms is adequate. That said, the part of Taylor's argument that I will be discussing does not draw argumentative force from this definition, but from the distinction between academic activities that are motivated by desires for extrinsic rewards such as money, fame, or comfortable employment conditions (market-like motivations) and academic activities that are motivated by the desire to advance understanding. The former sort of activities are,

² H.L.A. Hart, *Punishment and Responsibility: Essays in the Philosophy of Law* (Oxford: Oxford University Press, 1968), 4.

³ Taylor, *Markets with Limits*, 149.

⁴ Taylor, *Markets with Limits*, 150.

⁵ Taylor *Markets with Limits* 149 and 150; quoting E. Anderson, *Value in Ethics and Economics* (Cambridge, MA: Harvard University Press, 1993), 145.

⁶ Taylor, *Markets with Limits*, 150.

for the purposes of this paper, governed by market norms; the latter are governed by academic norms.

As I understand it, Taylor's argument for premise B has a negative and a positive part that together imply the comparison. The negative part is an argument that shows how some shortcomings of current academic research can be explained (at least in part) by how academic research is governed (at least in part) by market norms. The positive part is an argument aiming to show that academic research conducted in accordance with academic norms suffers from fewer of those shortcomings, without incurring other shortcomings sufficient to outweigh the comparative advantage. Together, these two parts imply that academic research that is conducted in accordance with academic norms would do better at advancing understanding than academic research conducted in accordance with market norms. So, the argument for premise B is as follows:

1. Academic research conducted in accordance with market norms tends to suffer various shortcomings: these include inaccurate references, misattribution of claims, and propagation of errors.
2. Academic research conducted in accordance with academic norms would suffer from fewer shortcomings, while not doing worse at advancing understanding in other respects.
3. Therefore, academic research conducted in accordance with academic norms advances understanding better than academic research conducted in accordance with market norms (= B).

Taylor's argument for premise 1 is delivered (with relish) by his discussion of the various shortcomings found in the current debate about the moral limits of markets. If, for the sake of the argument here, we grant premises A and 1, we are left with the validity of inferences in both arguments and the truth of premise 2 as the yet questionable elements of Taylor's argument.

I argue that premise 2 is ambiguous. On one reading, it is likely true, but the argument to 3 is invalid. On the other reading, premises 1 and 2 imply 3 (the argument is valid), but premise 2 is unsupported by Taylor's arguments and is likely false.

I first (section 2) introduce a distinction between the aims of a practice, the norms of the practice, and moves made in that practice. I then (section 3) argue that premise 2 is likely true when we compare *instances* of academic research conducted in accordance with market

norms and academic norms, but that such a reading renders the inference to 3 invalid, since this conclusion is about academia as a practice, not instances in that practice. Next (section 4), I argue that the needed reading of premise 2, which says is that the academic practice would better advance understanding if those engaging in it were governed by academic norms, is unsupported by Taylor's argument and is likely false.

2. Practices, norms, and moves

We can distinguish between the aims of practices, the rules or norms of a practice, and the moves made within a practice. Let's assume that moves made within a practice are justified when they comply with the norms for that practice. How should we then think about the relation between the aims of the practice and the justification and content of the rules of the practice? I am granting Taylor's position that the aim of the practice(s) of academic research is to advance understanding. But what does that entail with respect to the norms of this practice?

That a practice is justified by reference to some aim does not imply that the rules of that practice or the deliberations of the agents engaging in that practice must make explicit reference to or have that same aim. In particular, that a practice has some aim does not entail that the agents engaging in that practice must make their choices with that aim in mind, or that the rules that structure their activities should make them incorporate that aim into their preference order (even less, that they must give that aim high priority).

As Rawls wrote in the opening sentence of "Two Concepts of Rules", there is a "distinction between justifying a practice and justifying a particular action falling under it."⁷ To illustrate the point, Rawls sketches how we might justify our practices of punishment by reference to their utility, but justify particular instances of punishment by reference to the retribution appropriate to the crime committed.⁸ Thus understood, the rules of our practices of punishment are justified by, *but do not make explicit reference to*, the utilitarian aim; nor need the utilitarian aim figure into the preferences or decision procedures for any of the roles inhabited by the members of that practice (judges, lawyers, etc.). The same distinction appears in Hart's discussion of punishment

⁷ J. Rawls, "Two Concepts of Rules" in *John Rawls: Collected Papers*, S. Freeman ed. (Cambridge, MA: Harvard University Press, 1999), 20.

⁸ Rawls, "Two Concepts of Rules", 21-4.

referenced above with the concept of a ‘general justifying aim’. Hart begins that discussion with a distinction between the definition of punishment, the justification of practices of punishment, and the distribution of punishment.⁹ While our answers to the relevant questions – What is punishment? What justifies practices of punishment? Who should be punished and how? – are conceptually linked, they are distinct questions, and we can consistently defend a utilitarian justification of practices of punishment, while adopting a retributivist answer to questions of distribution. In this manner, Rawls and Hart share the idea that we can justify our practices of punishment by their consequences, while rejecting that consequentialist considerations should appear in the rules of those practices, or should govern the choices of those who inhabit the various roles of those practices. The rules and decision procedures of the practice can be retributivist, even if the aim of the practice is utilitarian.

The distinction is not particular to discussions of punishment, but also shows up at crucial junctures in debates about the choice of economic system. Thus, defenders of using market norms to structure economic behavior often justify their position by reference to the desirable consequences of those norms governing that behavior, rather than by the value of such behavior in itself. From Adam Smith through Friedrich Hayek, the most influential argument for governing economic behavior by market norms is that this unleashes the “creative powers of a free civilization” – the spontaneous order process by which individuals pursuing their own gain act in a manner (as if guided by an invisible hand) that is beneficial to others.¹⁰ The general justifying aim of our economic practices, on such an account, is to create social utility, but the rules of the practice (free market capitalism) make no reference to that aim, and the moves in that practice aim at *individual* utility, *not* at social utility.

By now it should be clear where this is going: while we can grant Taylor’s claim that the general justifying aim of our academic practices is to advance understanding, it might well be that the best way to advance this aim is to have the behavior of those engaging in these practices (academics) structured by market norms so that their activities aim at extrinsic rewards (money, fame, comfortable working

⁹ Hart, *Punishment and Responsibility*, Chapter 1.

¹⁰ F.A.v. Hayek, *The Constitution of Liberty* (Chicago IL: Chicago University Press, 1960), 22.

conditions), rather than at the aim of the practice as such (the intrinsic aim of advancing understanding).

A correlate of this possibility is that we need to distinguish between the value of individual acts or instances of academic research and the value of the practice of academic research as it is structured by any given set of norms. Even if instances of research by any given researcher have higher expected (social) utility if she aims to advance understanding than if she aims to secure extrinsic rewards, it might also be true that the practice of academic research advances understanding better when it is structured by market norms rather than academic norms.

3. First reading of premise 2: true premise, invalid inference

Taylor's argument for premise 2 is that academic research that is motivated by the desire to advance understanding will be less prone to various types of mistakes and, therefore, of higher quality in terms of advancing understanding:

This painstaking *academic* approach to academic work will likely result in fewer publications than the adoption of a more market-oriented approach. But it will also likely increase the quality of the work produced. Work produced by academics who direct their action in accord with academic norms is more likely to be exegetically accurate than the work of more market-oriented academics. This, in turn, will make it more likely that their criticisms of others' work will be sounder than those of their more market-oriented counterparts. They will thus be more likely than them to produce work that contributes to – rather than derails – the debates they engage in.¹¹

In short, academic research that is motivated by the desire to advance understanding will better advance understanding than academic research that is motivated by desire for extrinsic rewards such as money, comfortable working conditions, or prestige. Or, in terms of norms, that “academic research conducted in accord with market norms will be more prone to error than that conducted in accord with the norms of the academy. Since this is so, academic research should not be primarily driven by the norms of the market.”¹²

¹¹ Taylor, *Markets with Limits*, 151.

¹² Taylor, *Markets with Limits*, 160.

However, this inference is a fallacy. Even if we agree that instances of academic research conducted in accord with market norms are more prone to error and, therefore (and other things equal) better advance understanding, *it does not follow* that academic research should not be governed by market norms.

It may be that any piece of academic research produced with the aim to advance understanding will do so better than a piece by the same researcher aiming to secure money, pleasure, and prestige. It also may be that if we were to run pairwise comparisons of pieces of research produced by some member of the academic profession, where one stack contains pieces of research that aim to advance understanding and the other stack has pieces of research that aim for extrinsic rewards, we'd find that the pieces in the former pile are of higher quality – in the relevant sense of better advancing our understanding.

I don't mean that any piece in the stack with academically motivated research is better than any piece in the extrinsically motivated stack. Rather, the best piece in the academically motivated stack is better than the best piece in the extrinsically motivated stack, and the worst in the extrinsically motivated stack is worse than the worst in the intrinsically motivated stack, and there's an even distribution in quality of members of each stack between lowest and highest. Alternatively, we might compare counterparts: each piece is compared as intrinsically motivated and extrinsically motivated, in which case each piece that is intrinsically motivated would be better than its extrinsically motivated counterpart. However, no matter how we compare them, a generalization to the comparative value of the stacks (rather than instances in the stacks) has to consider that *there'll be more pieces in the extrinsically motivated stack*. So, neither method of comparison supports a general judgment about the aggregate value of the stacks.

Therefore, granting that a comparison of instances of research would favor research conducted to advance understanding does not show (or even suggest) that academic research in general would better further understanding, if it was generally motivated to advance understanding rather than extrinsic rewards. (This would be true *even if* every member of the intrinsically motivated stack was better than every member of the extrinsically motivated stack, which seems unlikely.) The question is about the aggregate – what norms of academic research would create the system of academic research that best furthers understanding – and the imagined pairwise comparisons do not answer that question. It may be that academia advances human understanding best if academics are

governed by extrinsic desires and not by the desire to advance understanding. Accordingly, the following syllogism is a fallacy, for the conclusion is about the practice while the premises are about the activities conducted within the practice:

- i. Instances of academic research conducted in accordance with market norms suffer from a number of shortcomings.
- ii. Instances of academic research by the same researcher conducted in accordance with academic norms would suffer fewer shortcomings, while not doing worse at advancing understanding in other respects.
- iii. Therefore, academic research conducted in accordance with academic norms advances understanding better than academic research conducted in accordance with market norms (= premise B).

The fallacy is also on display in the section titled “Defending the primacy of academic norms”. There, Taylor argues that the view that academic research should be governed by market norms is “self-defeating.”¹³ It is self-defeating because it asserts at the same time that “the aim of academic research is the production of academic publications” and that “The primary purpose of these publications is [...] to secure professional advancement [...] for their authors.”¹⁴ Taylor continues:

[T]he primary purpose of academic research *cannot* be to function as a sorting mechanism to allocate the extrinsic rewards of academic research. The primary reason for an activity performed cannot be to determine who should receive the rewards of performing it. Offering this as a reason for its performance would not answer the question of why the activity was judged valuable in the first place. [...] If the proper aim of academic research is to enhance understanding, then it should be directed by norms that would make this more likely to be achieved.¹⁵

The argument runs together the aim of the practice and the aims of those engaging in the practice. It is perfectly consistent to maintain that the aim of the practice is to advance understanding, while the aim

¹³ Taylor, *Markets with Limits*, 172.

¹⁴ Taylor, *Markets with Limits*, 172.

¹⁵ Taylor, *Markets with Limits*, 172.

of the individuals engaging in the practice are the extrinsic rewards of research activities, and, therefore, that the motives for engaging in the activities of the practice are those rewards, without the aim of the practice being to distribute rewards. The position is analogous to the Hayekian justification of free markets: the aim of the practice is social utility, which is best served by the norms of the free market, wherein each pursue the extrinsic rewards of economic activity, while the whole is designed to maximize social utility, *not* to distribute the rewards. This defense of market norms is not self-defeating. Analogously, we can reject Taylor's quick inference from the aim of academic research being to advance human understanding to the conclusion that academic norms, rather than market norms, would increase the degree to which academia achieves this aim.

In sum, if we understand premise 2 to say that individual pieces of research by a researcher are likely to better advance understanding if their production is governed by academic norms than if their production is governed by market norms, then this premise is likely true: but it does not imply the conclusion that our practices of academic research would advance understanding better, if conducted in accordance with academic rather than market norms. To warrant that inference, and hence to establish premise B, Taylor has to show that academic output in general or in the aggregate would better advance human understanding when conducted in accordance with academic norms.

4. Second reading of premise 2: valid inference, unsupported and likely false premise

Taylor needs to show that academic output in general or in the aggregate would advance understanding better when conducted in accordance with academic norms than when conducted in accordance with market norms. He doesn't, and there are reasons to believe that it wouldn't.

First, then, is the problem of a missing argument. As described in the previous section, Taylor argues that instances of academic research produced by any given researcher would be better if she aims to advance understanding than if she aims for external rewards. But Taylor offers no argument that allows him to generalize from this claim to the needed conclusion about the value of aggregate output, either for that individual researcher, or for the practices of academia in general.

Second, there are reasons to believe that aggregate output might better advance human understanding, if it is to some substantial degree conducted in accordance with market norms. Conceptually, there are several ways this might happen: The degree to which research advances understanding is a function of both the quantity and the quality of research output. These, in turn, are functions of the number of researchers, the time they allocate to research, and the time they spend on any given piece of research. Since, as Taylor recognizes, academically-oriented research takes more time than market-oriented research,¹⁶ the variables of quality and quantity are inversely correlated – and there is no *a priori* way to establish that research governed primarily by academic norms (maximizing for quality) does better on net than academic research governed equally by academic and market norms or academic research governed primarily by market norms (maximizing for quantity). In addition, holding the community of researchers constant, they might allocate less of their time to research if the practice is governed primarily by academic norms. Finally, the community of researchers might not be constant. There might be more researchers if research is conducted substantially in accordance with market norms, and the quality of researchers might be higher if the academic community is governed substantially by market norms.¹⁷

In short, there might be more researchers producing more and better research, if academic research is governed substantially by market rather than primarily by academic norms. These are, of course, mere possibilities, and whether and when they would be realized depends on multiple empirical factors. That said, there are reasons to think that some of these conceptual possibilities are realized.

First, let us look at the simple case where we hold the community of researchers and the time they allocate to research

¹⁶ E.g. Taylor, *Markets with Limits*, 182.

¹⁷ In comments on a draft of this essay, Shawn Klein suggested that an epistemic asymmetry might also count in favor of market norms: researchers might have limited knowledge of how to best advance understanding in general, but know well how to maximize their individual interests, given their abilities and local situation. Taylor might respond, first, that he is not recommending that researchers aim to advance understanding in general, but only in the specific domain that is their area of specialty, and, second, that the care with citations and so forth that he recommends generally prevents *misunderstanding*.

constant. Here, the issue is one of optimizing the allocation of a fixed amount of time. This question cannot be approached as an exclusive choice (either academic or market norms; either quality or quantity), but as one of finding the right balance of concerns for quality and quantity. That is, the question is one of balancing the quality of the pieces in the stack of research with the quantity of pieces (the height of the stack). Taylor argues that research should be governed primarily by academic norms, which means that he favors a point on the continuum towards the quality end of the scale (very short stack, very high quality). But there are reasons to expect that this would be a suboptimal choice – that is, that a higher stack of lower quality would do better at maximizing the relevant value (advancing understanding).

Taylor describes how he read all of Marcuse’s work in three languages in order to investigate the claim that Marcuse originated the humorous distinction between scholarship and scholarshit.¹⁸ The result is that it is unlikely that the distinction originates in Marcuse’s writings.¹⁹ Applying the same degree of diligence to all one’s academic attributions of claims, references, and so forth would mean that (holding time constant) researchers would produce much less output than if they permit themselves to commit some minor mistakes. The rigorously checked output would, of course, be of higher quality. However, it seems clear that, after an initial investment sufficient to make the argument and avoid gross misattributions of claims (straw-manning) and such, there’s a declining marginal return of increased understanding from the increase in time allocated to any one piece of research (the last fruits picked are very hard to reach!), and the opportunity costs of time allocated to any one piece of research remain constant. So, there’ll be a point lower than the point favored by Taylor beyond which the increase in quality will be insufficient to outweigh the decrease in quantity.

The optimum allocation of time for diligent checking of references is tricky to compute, but Taylor seems to adopt a maximizing requirement – we should aim for as few errors in references or mistaken attributions of claims as possible – which means that following his prescription is suboptimal, for it focuses entirely on just one of the two variables of the cost-benefit analysis that determines the optimum point.²⁰ That is, it requires us to maximize benefits without a view to the

¹⁸ Taylor, *Markets with Limits*, 144, 148, 153 n24, and 157-8 n64.

¹⁹ Taylor, *Markets with Limits*, 157-8 n64.

²⁰ At one point, Taylor recognizes the diminishing marginal returns: “Once an

(opportunity) costs, and so is likely to miss the optimum point. Thus, Taylor summarizes his argument:

To the degree that an academic researcher conducts her work in accord with academic norms, it is likely to have fewer exegetical errors than that produced in accord with market norms – including fewer misrepresentations of others’ views. Since such misrepresentations will impede (or even derail) productive debate, avoiding them is a consummation devoutly to be wish’d. Thus, given that the internal aim of the practice of academic research is to further understanding, academics should primarily (if not exclusively) direct their research in accord with academic, rather than market, norms.²¹

However, it wouldn’t be “a consummation devoutly to be wish’d,” if the costs outweigh the benefits. Taylor does not show that they wouldn’t, and it seems likely that they would, insofar as maximizing for one variable (quality) and neglecting another (quantity) is likely to lead to a suboptimal allocation.

Second, incentives don’t just matter for how persons use their time in the activities of a given practice, but also for the time they allocate to that practice. So, the expected rewards of doing research matter to how much time a researcher will allocate to research. Advancing understanding is one sort of reward: we all care about that, but not exclusively. We can have the reward of advancing understanding and also seek to be rewarded by extrinsic goods. The combined motivation is stronger than either by itself. So, we spend more time on research if we are motivated *also* by external rewards – meaning that

academic who has produced work in accordance with these norms believes she can no longer improve on it – *or that any attempt to do so will merely result in diminishing marginal returns* – then she should submit it for publication.” (*Markets with Limits*, 173, my emphasis) The emphasized portion seems to recognize that a maximizing-quality-stance is mistaken, but it’s hard to make sense of, insofar as there’ll be diminishing marginal returns at any given point of time invested. Moreover, from the context it appears that the returns here are measured entirely in terms of *quality*, and so the quoted passage does nothing to counter the problem that Taylor’s argument overlooks the costs in terms of alternative research projects foregone for the sake of continual improvement of this one publication.

²¹ Taylor, *Markets with Limits*, 174.

any given researcher would allocate more time to research in a system that ties extrinsic rewards to academic output.

Third, incentives also attract people to a profession. Again, academics generally care about advancing understanding, but we also have other desires – including for material gain, recognition, and comfortable working conditions. If academia did not tie external rewards to research output, then it is likely that some good researchers would seek employment elsewhere. Thus, if there are extrinsic rewards tied to academic research, then the position of academic researcher is more competitive to alternative careers, so, we can expect more competition for academic jobs which, in turn, will lead to more qualified researchers.

In sum, it is likely that academic research in the aggregate will advance human understanding better if it is governed to some significant extent by market norms and not, as Taylor argues, when it is governed primarily by academic norms.

5. An objection

Taylor might object that I am overstating his position. That is, he might object that he is not arguing that research should be governed entirely by academic norms, but merely that it *should not* be governed primarily by market norms. Here is a passage that states his conclusion in those terms:

It is possible that understanding would be best furthered were academic research to be conducted in accord with some mix of market and academic norms. [...] However, it is likely that if market norms were to dominate the production of academic research, its overall quality would suffer [...] Academic research should thus be guided primarily (but not exclusively) by the norms of the academy. [...] The norms of the market should not be allowed to dominate in the academic realm.²²

This passage offers two variants of the conclusion. One that says that academic research should be guided primarily by academic norms, the other that academic research should not be dominated by the norms of the market. These are consistent, but differ in strength and interest. The former is a stronger (it implies the latter and not *vice versa*) and more interesting conclusion – and is also the one that Taylor states

²² Taylor, *Markets with Limits*, 161.

repeatedly in other parts of the book (as in the passages quoted above). The critique of Taylor's argument that I raised in the previous section addresses the former and does not refute the latter (weaker and less interesting) conclusion – indeed, it supports the latter, insofar as my argument suggests that neither academic nor market norms should dominate, since understanding is best furthered by a combination of both sorts of norms.

Since Taylor repeatedly states his thesis / conclusion as the stronger claim – that academic research should be governed primarily by academic norms – and since this claim is more interesting, I do not believe that he could retreat to the weaker claim without changing his position and losing much of what makes his argument interesting.

6. Conclusion

Taylor argues that academic research should be governed primarily by academic norms (and not by market norms). He defends this conclusion by identifying the aim of academic research as advancing human understanding, and then arguing that academic research governed by academic norms better advances human understanding than research governed by market norms. He argues that instances of research conducted in accord with academic norms will tend to have more accurate references, fewer mistaken attributions of claims and arguments, and be less prone to reproduce errors than research conducted in accord with market norms. Alas, this argument does not suffice to establish Taylor's conclusion, since the conclusion is not about instances of research, but that academic research in general or in the aggregate would do better at advancing human understanding when conducted in accord with academic norms. Taylor offers no argument for the general claim, and it seems likely that it is false.²³

²³ I am grateful to James Stacey Taylor for being continuously thought-provoking, to Shawn Klein for organizing this symposium, and to Shawn and Lauren McGillicuddy for helpful comments on drafts of this essay.