

Beyond Manipulation? Revisiting Technopaternalism in the Context of News Recommenders

By Marijn Sax (University of Amsterdam)

Beyond Manipulation? Revisiting Technopaternalism in the Context of News Recommenders

By Marijn Sax (University of Amsterdam)

Introduction

I am going to play with fire here: I am going to discuss paternalism, as a European academic, for an American audience.

Although this may sound like a terrible idea, please bear with me. I will not be writing about Big Ideological or Political Divides. Rather, I want to focus on two underlying conceptual issues: 1) how autonomy tends to be underdefined in political paternalism discussions, and 2) the underappreciated importance of epistemic questions. I will argue that in paternalism discussions an individualistic, rationalistic account of autonomy is usually *assumed*, but not explicitly argued for. This is a problem, because individualistic accounts of autonomy – as opposed to relational accounts of autonomy – preconfigure paternalism discussions in a very particular, contingent manner. I will also argue that especially in the context of data-driven persuasive choice environments, it is increasingly difficult to see empirical questions as distinct from the underlying principled paternalism discussions.

Why did I start to look into paternalism (again)? Not because I actively wanted to, to be honest. I am very aware of the overly politicized and often counterproductive discussions it incites. However, while working on a research project on news diversity in news recommender systems, paternalism-like questions keep popping up. In the project, I (a philosopher) work together with communication scientists and computational linguists. We try to do three things: 1) use democratic theory to better understand the role and value of news diversity; 2) use computational methods to try to operationalize new understandings of news diversity (informed by 1)); and 3) use communication science theory and experiments to find out how users can be nudged towards actually consuming more diverse news. An obvious question we must ask ourselves is: “How proactive can we be in promoting diversity with the use of clever user interface design choices?”

Having worked on questions of manipulation and autonomy in digital choice architectures for quite some time,¹ my first impulse was to suggest that this is *obviously* a manipulation question. There is of course a manipulation angle here. But given the special *democratic* role that is often assigned to the media, nudging (or not) for news diversity doesn't strike me as just a manipulation question. The media can – or should – be considered as a public interest institution. This is the moment the question of paternalism presents itself: if we believe diverse news consumption to be in the public interest, how 'aggressive' can news recommenders be in pushing diverse content?

Paternalism and News Recommenders

In the context of public interest media, paternalism has always been identified as an important question. I don't want to claim that the paternalism angle is original. What is interesting though, is the fact that the paternalism question resurfaces in a recent wave of ethnographic/interview studies on the subject of news recommender design and, more specifically, their on the ground implementation at media organizations. Bodó (2019) and Hilden (2021), to name two examples, investigated how different actors *within* media organizations (editors, developers, marketers, data scientists) discuss the implementation of news recommenders and negotiate how different objectives and values should be balanced. Interestingly, most actors involved in news recommender design and implementation seemed to realize that their specific way of implementing a news recommender functions as a de facto 'paternalism dial'. They *have* to make concrete choices when deciding how the recommender is going to recommend a small subset of the available news items to news readers. The following quotes from Bodó's and Hilden's studies illustrate how paternalisms (implicitly) presents itself as a relevant consideration in news recommender implementation:

"Many editors see algorithmic recommenders as a tool to execute editorial decisions. If editors think something is important and should reach as many people as possible, they can use algorithms to "push articles to the public" (IP2 ['IP2' is an anonymized interviewee ID, MS]) to make sure that the news item gets through people's natural filters and reaches them even if they had it algorithmically filtered out. This approach considers recommender algorithms as a tool whose usefulness goes beyond the short-term satisfaction of user interests, and uses them to address societal level considerations." (Bodó 2019: 1065-1066)

"The discussions on what content recommender systems should put forth partly resembles old debates on the purposes of public service broadcasting—to what extent is it in the

¹ Please allow me shamelessly to plug my book *Between Empowerment and Manipulation* here (Sax 2021).

public interest to show light entertainment? Is the paternalistic approach to public service broadcasting being revived in the digital age?" (Hilden 2021: 5)

[Interviewee 04:] "There is a danger as your end user might not be interested in looking at the content you recommend, and we need to respect that. It is difficult to draw the line between respecting a viewer's behavior and widening his view." (Hilden 2021: 13)

This restatement of the age-old paternalism question got me thinking how we can navigate these paternalism questions in contemporary digital media discussions. Our first intuition tends to be to see this as a purely ideological/political question – How much 'interference with autonomy for good' is one comfortable with, to put it a bit simplistically.

However, starting to read and think about this paternalism question – again – I never arrived at this overarching, ideological 'how much paternalism do we want' question. Why? Because the ideological paternalism question doesn't quite make sense to me. Why? Because paternalism discussions tend to be based on a range of poorly specified assumptions. And so in what follows I want to address two underlying questions that we first need to address before we can even discuss paternalism in a sensible manner.

Two Questions: Autonomy and the Epistemic Dimension of Paternalism

Let me first quickly get something out of the way. An often seen and helpful first move in discussions that take place in the conceptual minefield that is paternalism, is to insist that one must differentiate between different *types* of paternalism (e.g., hard versus soft, broad versus narrow, weak versus strong, pure versus impure, moral versus welfare). Dworkin (see, e.g., 1972, 2013, 2020) has done a lot of work to identify different types of paternalism. These distinctions are helpful to break down big, explosive political debates into smaller puzzles that are easier to digest. In order not to redo this helpful types-of-paternalism exercise, I want to focus on two different issues: underdefining autonomy and lack of attention for paternalism's epistemic dimension.

Autonomy

The tension between autonomy and the promotion of some – to be defined – good the realization of which might require a (partial) restriction of autonomous decision-making, lies at the heart of paternalism discussions. And so autonomy is, of course, not an ignored concept in paternalism discussions. But autonomy tends to be defined in a narrow, individualistic, rationalistic manner, where autonomy is associated with negative freedom and the absence of interferences with

behavior. My use of the word ‘defined’ is generous here, since autonomy is more often than not simply *assumed* to be ‘that one concept that is basically aligned with personal freedom’.²

I think we can and need to do better. Because some level of infringement of autonomy is a *definitional* feature of paternalism, the *precise* understanding of underlying notions of autonomy is essential to properly analyzing paternalism. By implicitly relying on a rather specific and contingent individualistic, rationalistic understanding of autonomy, the following *political* discussions on paternalism are already pre-structured by the conceptual assumptions. It’s a shame that this happens not through explicit argumentation in favor of a rather specific understanding of autonomy, but rather by lazily assuming that we all know what autonomy is, roughly speaking, about.

To illustrate my point, consider how a theory of relational autonomy fits into paternalism arguments. The literature on relational theory – developed by theorists such as Nedelsky (1989), Code (1991), Mackenzie & Stoljar (2000) – argues against overly individualistic, rationalistic accounts of autonomy. They argue that the dominant procedural (individualistic, rationalistic) accounts of autonomy presuppose an “[a]utonomous man [who] is – and should be – self-sufficient, independent, and self-reliant, a self-realizing individual who directs his efforts towards maximizing his personal gains” (Code 1991: 78). If one departs from such a presupposition, autonomy is about securing maximum independence in reasoning and decision-making capabilities.

Relational autonomy theorists point out that social reality doesn’t work this way. People are never and will never be fully independent, and striving for maximum independence is not only unattainable, but also simply unattractive. Our ability to critically reflect on our own lives and think about ways to provide some structure and direction to our own life is always already dependent on our relationships with others. This makes us dependent, which is both a fact of life and an attractive feature of life.

Consider the fact that no one is born autonomous; everyone needs their parents, siblings, friends and so on to *develop* one’s capacity for practicing autonomy. The same holds for grown-ups. We constantly rely on – are dependent on – others for advice, for providing a safe space to vent or to break down. As Nedelsky (1989: 92) puts it succinctly: “If we ask ourselves what actually enables

² Thaler and Sunstein’s entire *nudge* oeuvre is founded on an inability or unwillingness to provide clear, stable definitions of key terms like ‘autonomy’, ‘freedom’, and ‘liberty’. See Engelen & Nys 2020.

people to be autonomous, the answer is not isolation, but relationships – with parents, teachers, friends, loved ones – that provide support and guidance necessary for the development and experience of autonomy”.

I cannot develop a full theory of relational autonomy here. (I encourage the reader to start with Mackenzie & Stoljar’s edited volume *Relational Autonomy*). But the first contours I sketched hopefully provide a clear contrast with standard procedural accounts of autonomy that value a person’s ability to make independent decisions in a calculated manner to optimize outcomes. On such accounts of autonomy, it is very easy to qualify influences (notice my neutral phrasing here) as potential *interferences* with autonomous decision-making. It also follows that if such an account of autonomy serves as a background assumption in discussions on paternalism, it is very easy to see/generate normative *tensions* that require our immediate paternalism-attention. If, however, one starts from a more relational understanding of autonomy, one works with a concept of autonomy that already has the importance and value of interdependencies built in. Heck, these interdependencies can even – under the right circumstances – be considered *constitutive of autonomy*. Here, I would argue, it also follows that a relational theory of autonomy renders different – and maybe even fewer – influences as normative tensions that require our paternalism-attention.

One can of course argue that theories of relational autonomy are not convincing. That is fine by me really (although I disagree). But, and this is the important point, every discussion of paternalism should start with an *explicit* and *substantive* discussion of the underlying concept of autonomy. Without a theory of autonomy, a theory of paternalism is blind (or dead).³

The Epistemic Dimension of Paternalism in the Digital Society

Let me now turn to what I have called ‘the epistemic dimension of paternalism’. In ‘classic’ paternalism debates there is a strong focus on the principled question of whether a person’s autonomy can be restricted for the promotion of some good. Practical, empirical considerations are often put to the side or discussed in the footnotes. Such empirical considerations include (but are not limited to): can we know or estimate how effective a proposed paternalistic measure is going to be? Or: how much can we know about the (true, informed) desires of people affected by a measure; and how certain can we be that our measure truly contributes to the idea of the good we are trying to promote here?

³ My phrasings in this section might suggest that I think theories of autonomy a pre-/non-political and can simply be adopted *before* we start digging into the political dimensions of paternalism. This is not the case. Theories of autonomy are themselves normative and can support (or be inspired by) political ideals. All the more reason to make the autonomy dimension of paternalism discussions more explicit and transparent.

As a philosopher, I can appreciate the fact that in a causal world filled with uncertainties, one must sometimes put empirical questions between parentheses to focus on disagreements regarding underlying values or principles. However, in the context of digital choice environments (such as news recommenders), I find it increasingly difficult to place the empirical questions pertaining to alleged paternalistic influences between parentheses. Why? For one reason, because the designers of digital choice environments pride themselves (correctly or not) for the fact that through data collection and analysis, they can acquire deep knowledge of the users of their digital choice environment as well as its causal effects on the world. If this promise of datafication holds true only partially, it might also imply a stronger duty to actually take into account the – allegedly – measurable effects of interventions on the world.

In such a context of datafication and digital choice environments, the ‘mere empirical’ considerations that normally figure as a footnote to paternalism discussions should, I would like to argue, be seen as relevant normative considerations. If designers/operators make strong knowledge claims about the efficaciousness of their choice environments and the psychological make-up of the users of these environments when proposing an intervention that might be seen as paternalistic, we can 1) require them, as a kind of proportionality requirement, to actually explain which effects they expect, and why; and 2) we should scrutinize the underlying knowledge claims made.

The second point, scrutinizing claims made, may seem trivial (and it is). But it is especially important in the very hype-sensitive context of (commercial) digital choice environments promising to be able to optimize all kinds of metrics. If we look at algorithmic news recommenders, we immediately see the importance of questioning the – positive, promising – claims of e.g. optimizing news recommenders ‘for democracy’. Optimizing for more *diverse* news consumption is often treated as a good proxy for optimizing for democratic values. (For good reason by the way, there is ample literature detailing the importance of news diversity for democracy. See, e.g., Helberger 2019).

To be able to steer users (paternalistically?) towards the consumption of a more diverse media diet, ‘diversity’ must be operationalized in a sensible manner. Natural language processing (NLP) tools are important here because types of diversity (e.g., standpoint, topic, tone) must be computationally ‘captured’ by algorithms. NLP scholars are working hard on this challenge and are making progress, but the current state of the art NLP solutions cannot really be said to ‘capture’

diversity in a conceptually meaningful and reliable manner (see, e.g., Vrijenhoek *et al.* (2021) for a first attempt to *approximate* operationalizations of different types of diversity).

Coming back to my argument on the normative importance of the epistemic dimension of paternalism, one could argue that the current inability of news recommender technology to ‘capture’ news diversity in a truly meaningful manner should figure as a normative consideration in paternalism arguments. Given the centrality of diversity metrics in ‘nudging for diversity and democracy’-approaches, the currently sub-optimal diversity metrics should make it more difficult – but not impossible – to justify attempts to (paternalistically?) steer news consumption behavior.

With hype terms such as ‘AI’ and ‘algorithms’ being used very often to sell us techno-solutionism for age-old problems, epistemic modesty becomes an increasingly important virtue. When paternalism discussions enter the domain of hype-filled tech solutions, we should be especially cautious not to overestimate what the proposed tech solutions can *actually* do. Not falling into the trap of letting overly optimistic tech-related promises silently seep into our normative arguments on how the technology in question can and cannot be used requires constant philosophical vigilance.

Conclusion

I started this piece by announcing I would discuss paternalism not in an overly politicized way, but in a more boring conceptual way. Hopefully I succeeded in being more boring than political. My aim was not to provide a comprehensive discussion of paternalism, which would require an entire book (or two). Rather, my aim was merely to ask two conceptual questions about paternalism and develop nothing more than the first contours of possible answers. If discussions on the alleged paternalism of technology-supported interventions become heated and overly polarized in the near future, I hope we can take a step back, bracket the political questions for a minute, and first start to consider the underlying conceptual assumptions that tend to do a lot of the heavy lifting in these heated discussions. My quick suggestions in this piece won’t solve or help overcome polarized paternalism discussions, but maybe they help add a little more rigor as well as chill to the mix.

References

- Bodó, B. (2019). Selling News to Audiences – A Qualitative Inquiry into the Emerging Logics of Algorithmic News Personalization in European Quality News Media. *Digital Journalism*, 7(8), 1054-1075.
- Code, L. (1991). *What Can She Know? Feminist Theory and the Construction of Knowledge*. Cornell University Press.
- Dworkin, G. (1972). Paternalism. *The Monist*, 56(1), 64-84.
- Dworkin, G. (2013). Defining Paternalism. In C. Coons & M. Weber (Eds.), *Paternalism: Theory & Practice* (pp. 25-38). Cambridge: Cambridge University Press.
- Dworkin, G. (2020). Paternalism. In E. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. <<https://plato.stanford.edu/entries/paternalism/>>
- Engelen, B., & Nys, T. R. V. (2020). Nudging and Autonomy: Analyzing and Alleviating the Worries. *Review of Philosophy and Psychology*, 11, 137-156.
- Helberger, N. (2019). On the Democratic Role of News Recommenders. *Digital Journalism*, 7(8), 993-1012.
- Hilden, J. (2021). The Public Service Approach to Recommender Systems: Filtering to Cultivate. *Television & New Media*.
- Mackenzie, C., & Stoljar, N. (Eds.), *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. Oxford University Press.
- Nedelsky, J. (1989). Reconceiving Autonomy: Sources, Thoughts and Possibilities. *Yale Journal of Law and Feminism*, 1, 7-36.
- Sax, M. (2021). *Between Empowerment and Manipulation: The Ethics and Regulation of For-Profit Health Apps*. Wolters Kluwer. Open access: https://www.ivir.nl/sax_info_47-2/
- Vrijenhoek, S., Kaya, M., Metoui, N., Möller, J., Odijk, D., Helberger, N. (2021). Recommenders with a Mission: Assessing Diversity in News Recommendations. *CHIIR'21: Proceedings of the 2021 Conference on Human Interaction and Retrieval*, 173-183. <https://dl.acm.org/doi/abs/10.1145/3406522.3446019>



Marijn Sax
DLI Associate
Postdoc @ Institute for Information Law (IViR)
University of Amsterdam
M.Sax@uva.nl
[More Info >](#)