AI & SOCIETY
https://doi.org/10.1007/s00146-022-01551-7

ORIGINAL PAPER

Artificial intelligence and the secret ballot

Jakob Thrane Mainz[1] · **Jørn Sønderholm**[2] · **Rasmus Uhrenfeldt**[3]

Received: 7 March 2022 / Accepted: 17 August 2022

Abstract
In this paper, we argue that because of the advent of Artificial Intelligence, the secret ballot is now much less effective at protecting voters from voting related instances of social ostracism and social punishment. If one has access to vast amounts of data about specific electors, then it is possible, at least with respect to a significant subset of electors, to infer with high levels of accuracy how they voted in a past election. Since the accuracy levels of Artificial Intelligence are so high, the practical consequences of someone inferring one's vote are identical to the practical consequences of having one's vote revealed directly under an open voting regime. Therefore, if one thinks that the secret ballot is at least partly justified because it protects electors against voting related social ostracism and social punishment, one should be morally troubled by how Artificial Intelligence today can be used to infer individual electors' past voting behaviour.

Keywords Artificial Intelligence · Secret Ballot; Voting · Democracy

1 Introduction

In the past decade, many academic publications have been devoted to showing how Artificial Intelligence (AI) has had a disrupting influence in many important public policy domains.[1] These domains include law enforcement (predicting future crime spots), college admission (predicting which students are likely to drop out), incarceration (predicting which inmates are likely to reoffend when released) and epidemiology (predicting outbreaks of disease around the world).[2] Recently, several articles have shown how AI has entered the domain of politics, and sometimes with worrying effects (Christiano 2021; Erman and Furendal 2022; Sønderholm et al. 2021). One example of such an effect is that algorithms can be used to spread fake news which, in turn, can undermine some electors' ability to reach informed conclusions on political issues (Christiano 2021: 1). Another example is microtargeting. Microtargeting is often an effective tool in changing behaviour or opinion because the algorithms used in microtargeting are trained on data about what members of the targeted group are likely to be moved by. Microtargeting can, for example, be used in get-out-the-vote campaigns and in efforts to shape political opinion in a certain direction. Microtargeting can pose a threat to democracy because it can be used to manipulate electors. Messages sent to electors can be emotional in nature and thereby trigger anger or fear (Christiano 2021: 5). Microtargeting can be particularly problematic when the groups targeted are constituted by low-information electors who are not deliberatively sophisticated. When such groups are targeted, microtargeting can contribute to political inequality (Christiano 2021: 1).

In this article, we add to the emerging literature in political philosophy on how AI affects politics and how AI can be a threat to our democratic institutions. We have a narrow focus in virtue of zooming in on the relationship between the

Jakob Thrane Mainz, Jørn Sønderholm and Rasmus Uhrenfeldt contributed equally to the paper.

✉ Jakob Thrane Mainz
jtm@cas.au.dk

Jørn Sønderholm
jorns@ikl.aau.dk

Rasmus Uhrenfeldt
rauh@jm.dk

[1] Aarhus University, Aarhus, Denmark

[2] Aalborg University, Aalborg, Denmark

[3] Danish Ministry of Justice, Copenhagen, Denmark

[1] Throughout the paper, we use the term 'AI' rather loosely. It denotes several different data analytics techniques, including machine learning and deep neural networks. Nothing in the argument hinges on this, however.

[2] For a sample of such articles, see (Rona-Tas 2020, Brayne and Christin 2021, Skeem and Lowenkamp 2020, Khemasuwan and Colt 2021, Roberts 2019).

Published online: 01 September 2022

secret ballot and AI utilized to infer the past voting behaviour of individual electors.

The secret ballot is an entrenched political institution in liberal democracies. In referendums and elections in such democracies, electors cast their votes in secret. The secret ballot is often taken to be an integral part of the fairness and legitimacy of democratic elections. This view has famously been expressed by Robert Dahl who claims that countries without the secret ballot cannot be judged to have fair and free elections (Dahl 1998: 96). A similar verdict is reached by Dennis Thompson, who attributes the reduction of illicit voter influences to the success of the secret ballot (Thompson 2002: 66). The institution of the secret ballot is even encoded in the United Nations' Universal Declaration of Human Rights:

> The will of the people shall be the basis of the authority of government; this will shall be expressed in periodic and genuine elections which shall be by universal and equal suffrage and shall be held by secret vote or by equivalent free voting procedures (United Nations (General Assembly) 1948: article 21).

Multiple justifications have been offered for the secret ballot.[3] We focus on just two of the most common ones. The first justification is that the secret ballot protects electors against being ostracized because of having voted in a particular way.[4] This includes electors being ostracized by their neighbours, family members and/or co-workers because they have voted in a way that goes against the political views of the ostracizing party. Let us call this phenomenon 'social ostracism'. The second justification is that the secret ballot protects electors against being formally punished for having voted in a particular way.[5] This includes employers firing employees, or not hiring them in the first place, because they have voted in a way that goes against the employer's political views. It also includes landlords evicting tenants from their tenancies, or not letting the tenancy out to them in the first place, because they have voted in a way that goes against the landlord's political views. Let us call this phenomenon 'social punishment'.

Neither social ostracism nor social punishment involves influencing electors *ex ante* to vote in a particular way. It only involves punishing electors *ex post* for having voted in a particular way. Other common justifications for the secret ballot concern the *ex ante* influencing of electors. One such justification is that the secret ballot prevents vote buying because open voting makes it possible to pay someone *ex ante* to vote in a particular way, and then verify *ex post* whether the elector in question voted accordingly. While the prevention of vote buying and other forms of influencing voters *ex ante* to vote in particular ways are indeed an important justification for the secret ballot, we shall not address them at any length in this article. The reason is that even very accurate AI predictions cannot guarantee that a particular voter will indeed vote in accordance with the prediction. A voter may fit the profile of someone who tends to vote for a particular party, and still vote for a completely different party. A vote buyer can therefore not be guaranteed that the vote seller in fact votes as agreed upon. But even though vote buying and other types of *ex ante* influences of voters have not become possible because of the advent of AI, social ostracism and social punishment may indeed become more common due to the advent of AI. The reason is—as we shall argue—that these types of *ex post* punishment of voters do not require that the ballot is completely revealed or that the AI is completely accurate.[6]

[3] See (Mares 2015: 3–4, Elklit and Maley 2019: 65, Lever 2007, 2015). See (Brennan and Pettit 1990, Sturgis 2005, Engelen and Nys 2013) for arguments for open voting.

[4] Elklit and Maley, for instance, write: "The secret ballot remains important in protecting voters from forms of pressure, especially within families, that fall into something of a grey area between illegal coercion and legitimate persuasion"(Elklit and Maley 2019: 66). Similarly, Vandamme writes: "As I already suggested, there are prudential reasons for defending the secret ballot. They make us prefer secret ballots for their beneficial effects on vulnerable voters. If, for example, we believe that the voices of vulnerable people are of primary importance in the democratic quest for justice, we might want to make sure that they will be in a position to defend their legitimate interests or express a judgement free from domination" (Vandamme 2018: 392). A voter is certainly not free from domination if she knows that the price for voting in accordance with her political preference is that she will be looked upon with disdain by her co-workers and excluded from social activities at her workplace.

[5] Elklit and Maley write: "Second, the secret ballot can be seen as an instrument that protects voters from the possibility of violence or other coercive action intended to influence their voting decision or to punish them for having voted in a particular way (or indeed for having voted at all)" (Elklit and Maley 2019: 65). Manin claims that

Footnote 5 (continued)

open voting in general elections has three undesirable implications, one of which is that it increases the importance of private rewards and punishments in elections (Manin 2015: abstract). J. S. Mill also recognized this justification for the secret ballot. At one point he writes "Thirty years ago it was still true that in the election of members of Parliament the main evil to be guarded against was that which the [secret] ballot would exclude—coercion by landlords, employers, and customers" (Mill 1861: chapter 10). Coercion by, say, a landlord can take either of two forms. First, it can be an *ex ante* threat to the effect that the tenant will be evicted if she does not vote in accordance with the landlord's precepts. Second, it can be an *ex post* punishment to the effect that the tenant will be evicted if she has not voted in accordance with the landlord's precepts.

[6] Our argument also potentially affects other common justifications of the secret ballot. For example, Annabelle Lever has defended the secret ballot on the grounds that voting is a private matter and that open voting therefore violates electors' privacy rights (Lever 2007, 2015).

Historically, the secret ballot has protected electors from social ostracism as well as social punishment. By voting in secret, members of the electorate could hide their voting behaviour from their employer, landlords, and family—and by doing so avoid social ostracism and social punishment. In this paper, we argue that for a significant subset of electors, the secret ballot is now much less effective at protecting individual voters against social ostracism and social punishment. The reason is that AI and the mass-collection of personal information by companies/organizations that interact with their customers/members via online platforms, now make it possible to accurately infer how specific electors have voted. For some electors, it is now possible to infer, based on just a few data points and with extraordinary accuracy levels, how they have voted. Given that this is now possible, the act of casting your ballot in secrecy is no longer sufficient for others not to have a high credence in how you voted. Even when casting your vote is done in secret, it does not mean that your vote is de facto secret, because other people can accurately infer *ex post* how you voted. While the secret ballot in the past did protect electors against social ostracism and social punishment, it no longer does so because it is now possible to make some electors' ballot de facto non-secret using AI. Of course, the secret ballot never protected electors against *all* instances of social ostracism and social punishment. People can easily be socially ostracized or socially punished for reasons unrelated to their voting behaviour. Nevertheless, because of the advent of AI, electors are now less protected against social ostracism and social punishment than they used to be. Moreover, people can be socially ostracized and socially punished for reasons related to their voting behaviour, even without the use of AI. We often have a relatively accurate picture of people's political orientation based on their personal appearance, consumption patterns, accommodation choice, arrangement of family life and choice of hobbies. However, with the advent of AI, it is now possible to infer with much higher levels of accuracy the political orientations of individuals. And, it is now possible to infer the political orientations not only of people with whom we interact with in person or online, but also of millions or hundreds of millions of people whom we have never met and never will meet.

The article is structured as follows. In section II, we explain how AI can be used to infer how a specific elector has voted in an election. In section III, we explain how this possibility makes the secret ballot ineffective in protecting a significant subset of electors against social ostracism and social punishment. In section IV, we discuss and reject an objection to our argument. Finally, in section V, we make a few concluding remarks.

2 How AI can be used to infer how a significant subset of electors have voted

In this section, we explain how AI can be used to infer how some electors have voted in an election. But first, let us briefly describe a voting situation that renders it probable that AI can be used to make *accurate* inferences about how some electors have voted. Throughout the article, we shall assume that we find ourselves in the following situation: recently there was a referendum on a state constitutional amendment intended to ban same-sex marriage. Electors faced a binary voting choice: they could either vote for a ban on same-sex marriage, or they could vote against it. Moreover, the referendum took place in the United States (US), where there are publicly available voter registration lists in place.[7] These voter registration lists contain information about individual electors, including information about whether the elector in question participated in a specific election. In every state, the list contains the full voting history of all electors in that state (Sønderholm et al. 2021: 121). This means that after the referendum, a data analytics company could purchase access to the voter registration list to confirm that a specific elector voted in the referendum.

In a scenario like the one described above, where the voting choice is binary, and where it is possible to verify that a specific elector has voted, it will *ceteris paribus* be easier to accurately infer how a specific elector has voted than it would be in an election, in a country without voter registration lists, where electors faced a non-binary voting choice.[8] Our argument does not rest on whether we find ourselves in the scenario described above. We only assume that we do, to showcase in what type of real-life situation someone might decide to use an AI to predict past voting behaviour, and in what type of situation doing so might be particularly fruitful.

So, how can a data analytics company infer how a specific elector voted in the referendum? Suppose that a data analytics company scrapes publicly available data from millions of electors' social media profiles, and purchases access to all voter registration lists in the US. This gives the company access to data about electors' ethnicity, religion, gender, sexuality, age, job, income, address, space–time location, authored blog posts, Facebook likes, music and movie downloads, marital status, family size, language skills and education level. The company then trains a machine learning

[7] For an overview of the content of these lists, and where to purchase access to them, see http://voterlist.electproject.org/home (Accessed May 3, 2022).

[8] Elections where three or more parties/candidates are on the ballot constitute elections where electors face a non-binary voting choice. The same is true of elections where there are two parties/candidates and where electors have the option of returning a blank ballot.

model to find statistical correlations in the data set, which in turn might reveal significant statistical correlations between non-political features like, say, marital status, and voting behaviour in past elections. Once the model has been trained and tested, the company can then use the model to make inferences about specific electors' past voting behaviour.

Let us consider a concrete example. The data analytics company might decide to use the model to infer how a specific elector—call her Mary—voted in the referendum on same-sex marriage. By running the model, it turns out—strangely enough—that certain non-political data points about Mary are strong predictors of voting behaviour. For example, the car that Mary drives reveals a relatively accurate picture of Mary's political views. If Mary drives a sedan, she is much more likely to have voted for the Democrats than for the Republicans in a previous Presidential election. If Mary drives a pickup truck, she is much more likely to have voted for the Republicans than for the Democrats in a previous Presidential election.[9] And if Mary voted Republican in the most recent election, she is highly likely to have voted *against* same-sex marriage in the referendum. If she voted for the Democrats in the most recent election, she is highly likely to have voted *in favour* of same-sex marriage in the referendum.

If the data analytics company has access to tens of thousands of data points about Mary, they are likely to have a very accurate picture of who Mary is, and how she voted in the referendum. While an individual data point like the car she drives is a strong predictor, combinations of tens of thousands of data points can drive up the accuracy of the prediction significantly. Combinations of various data points about Mary can render it probable, with a high degree of accuracy, how Mary voted in the referendum.[10]

Already in 2014, David Nickerson and Todd Rogers published a paper in *Journal of Economic Perspectives*, in which they showed that it is possible to accurately predict voting behaviour at the individual level (Nickerson and Rogers 2014). That is, AI models not only make it possible to predict the voting behaviour of groups of voters, but also of particular voters. With techniques similar to those used by Nickerson and Rogers, various websites, like Google and Facebook, now routinely employ AI models specifically designed to predict individual people's voting behaviour (Hinds and Joinson 2019). Links we click on online, news we read, posts we 'like' on social media sites etc., generate an accurate digital picture of our political orientation, which, in a de facto two-party system like the one in the US, is highly correlated with voting for a particular party (Peters 2022). Even the pictures of ourselves that we (or someone else) post online may be used to predict our voting behaviour. An already infamous study published in *Scientific reports* by Michal Kosinski found that existing facial recognition software can predict the political orientation of individual voters, because the faces of conservatives and liberals apparently differ enough to predict with relatively high levels of accuracy if individual people belong to one or the other group (Kosinski 2021).

To illustrate just how easy it is to predict the voting behaviour of individual electors without using obviously political data, and without using facial recognition software, consider the voting prediction tool developed by The Economist and YouGov. Although this prediction tool does not qualify as an AI, it illustrates just how few (and seemingly trivial) data points it takes to predict the voting behaviour of individual electors with very high levels of accuracy. According to The Economist and YouGov prediction tool, there is a 97 percent chance that an elector votes for the Democrats, if she fits the following description: is a black female, who does not believe in reincarnation, is an atheist, is heterosexual, is between 45 and 64 years old, is not a regular church-goer, is married, has children, has a college degree, earns $30–64 thousands a year, lives in a city in the Midwest and speaks Spanish.[11] Assuming that this prediction tool is roughly correct, the probability of 97 percent is reached by having access to a relatively limited number of data points about an individual elector. Imagine how accurate the inferences can be if a data analytics company has access to tens of thousands of data points about Mary. If we assume, as we do throughout the paper, that we are dealing with a binary voting choice, in a setting where voter registration lists are publicly available, and where the data analytics company has access to massive amounts of data about Mary, we contend that the company can infer how Mary voted in the referendum with a very high level of accuracy.

It is important to note that the level of accuracy with which a data analytics company can infer how specific electors voted in an election varies from elector to elector. The more stereotypical an elector is, the easier it is to infer her voting behaviour. For some electors, the company might be able to infer with almost certainty how they voted, while for

[9] See (Gebru et al. 2017).

[10] A recent study shows that based on Facebook likes alone, it is possible to predict individual electors' voting intentions in a multi-party system with 60-70 percent accuracy (Kristensen et al. 2017). A recent article asks "Can we predict the voting behavior of Facebook users from their public Facebook profile?"(Idan and Feigenbaum 2019: 816). One model has an accuracy score of 82,5% when it comes to predicting whom an elector would vote for in the 2016 US Presidential election (Idan and Feigenbaum 2019: 823).

[11] The interested reader can try *The Economist* prediction tool herself at: https://www.economist.com/graphic-detail/2018/11/03/how-to-forecast-an-americans-vote (Accessed May 4, 2022). *The Economist* prediction tool is not an AI in any way, but it highlights how few data points it requires to reach high levels of predictive accuracy.

other electors, the company will end up with significantly lower credence in how these specific electors voted. However, note that this is compatible with our claim that *for a significant subset of electors*, the secret ballot no longer protects against social ostracism and social punishment.

It is also important to note that we are not dealing with predictions about how people are going to vote in *future* elections. As the—by now—infamous scandal involving Facebook and Cambridge Analytica showed, political parties and other political organizations are interested in using AI to predict how individual electors are going to vote in an upcoming election and then use microtargeting to persuade them to vote in a particular way. While it may indeed be possible to predict future voting behaviour relatively accurately, it is *ceteris paribus* easier to infer how individual electors voted in the past, than it is to predict how they are going to vote in the future. This is one factor that explains why it is possible to reach levels of accuracy approaching 100% when it comes to inferring past political behaviour.

Even though we shall not discuss the issue of predicting voting behaviour in future elections, we should stress that a solid picture of how individual voters have voted in previous elections may give a relatively solid picture of how people will vote in future elections. Indications of how individual voters will vote in the future is of course useful information that can be used to manipulate or convince voters to vote in particular ways. Moreover, punishing voters *ex post* for having voted in a particular way may well effect how, and if, they decide to vote in the next election. So, *ex post* punishment (and even the prospect of such punishment) of voters may well serve, in itself, as a way to manipulate voters *ex ante* in the next election.[12] However, we restrict the scope of our claim to only include past voting behaviour, because AI-inferred indications of how individual voters will vote in future elections cannot be used to *force* people to vote in particular ways, or to verify *ex post* if they voted for the candidate agreed upon. As things are now, AI-driven methods of inferring individual voter behaviour are primarily, although not exclusively, an issue related to the secret ballot when it comes to *past* elections.

Our argument is still interesting and important, even though it is restricted in the ways described above. This is so because of the normatively entrenched nature of the secret ballot. The secret ballot is a jewel in the crown of liberal democracies (up there with the one-person-one-vote principle), and any undermining of the secret ballot for a significant subset of electors is likely to be met with alarm.

[12] Thanks to an anonymous reviewer for bringing this point to our attention.

3 Why the secret ballot no longer protects all electors against social ostracism and social punishment

Let us now revisit Mary by going through two hypotheticals involving her. Hypothetical I shows why the secret ballot no longer protects electors like Mary against social *ostracism*, while Hypothetical II shows why the secret ballot no longer protects electors like Mary against social *punishment*.

> *Hypothetical I*: Mary is a tenant in a privately-owned apartment complex named 'St. John's Apartments'. The landlord is devoted to Catholicism and wishes that his tenants share his faith. Promotional material for St. John's emphasizes that it is a place where people of faith can feel at home and that a commitment to Catholicism is something that unites the tenants. In interviews with prospective tenants, the landlord emphasizes the Catholic faith of the tenants, shows St. John's chapel, and mentions that a local priest comes by several times a week and that tenants help each other with transportation to Mass on Sundays. At some point, Mary's landlord becomes suspicious of his tenants' religious commitments. Having read about the recent developments in AI, he decides to contract with a data analytics company and asks it to find out if the tenants voted in the referendum on same-sex marriage, and if so, what they voted. The company purchases access to the publicly available voter registration lists, and it scrapes vast amounts of publicly available data about Mary and the other tenants from their social media profiles. It also purchases data about the tenants from other companies. By agreeing to the terms and conditions for use of the services of these other companies, the tenants have typically allowed these companies to sell the data they collect about their customers. For a long time, Mary has been trying to hide her real political beliefs in the offline world, but she has expressed her progressive views on online political fora. Having gained access to all the tenants' data, the company then runs a machine learning model on all the data. It turns out that Mary voted, and that there is a 92% chance that she voted in favour of same-sex marriage. For all the other tenants who voted, there was, for each of them, a high likelihood that they voted *against* same-sex marriage. Receiving the result about Mary, the landlord shares the information about her with the other tenants of St. John's, and he lays out all the information about Mary that the company utilized in arriving at its prediction about how Mary voted. After the report about Mary's voting behaviour is shared among the tenants, things are different for Mary. She is socially ostracized by the other tenants.

They neither talk to her, invite her to social gatherings, or offer her a lift to mass on Sunday. They do not even look at her. This quickly becomes unbearable for Mary who moves out after three months. Mary did indeed support same-sex marriage through her voting, but the tenants of St. John's do not *know* this. They only have access to the information presented to them, and for them a 92% probability of "culpability" is enough to sway them to think that Mary supports social- and legal change that is antithetical to their most deeply held convictions. For them, the difference between, say, a 92%, 95% and 99.1% probability of a certain event having occurred is insignificant in their decision-making about how to treat Mary.[13] So, having heard that there is a 92% probability that Mary did vote in favour of same-sex marriage is enough for them to treat Mary as if they knew that Mary did vote in this manner.

Hypothetical II: up until the point where the landlord receives the report about Mary, everything is the same as in Hypothetical I. The landlord does not share the report with the tenants of St John's. Rather, he wants to evict Mary but worries that evicting her for (likely) having voted in a particular manner is legally impermissible and makes him vulnerable to a lawsuit from Mary. Therefore, he officially evicts her for some trivial but legally sound reason that has nothing to do with Mary's alleged voting behaviour.[14] Mary concludes that she cannot do anything about the landlord's decision and moves out after three months.

What happens to Mary in Hypothetical II is an example of an, unfortunately typical, kind of discrimination where the real reason an employer fires a female employer is that she is pregnant, and wants to take maternity leave, but where the official reason for firing her is a different one that neither refers to pregnancy nor maternity leave. The two hypotheticals in conjunction show that for electors like Mary, the secret ballot no longer protects her from social ostracism and social punishment. Had Mary lived in an age where AI did not exist, it is less likely that Mary would have suffered the fate that she suffers in the two hypotheticals. The two hypotheticals are empirically unrealistic if they were supposed to depict events from before the advent of AI. So, if one thinks that the secret ballot is justified because it protects electors against social ostracism and social punishment, then one should be troubled by how AI can be used to infer individual electors' voting behaviour today. AI can be used to undo, for a significant subset of electors, what the secret ballot originally was (successfully) put in place to secure for all electors.[15]

One might worry how plausible it is that a landlord would contract with a data analytics company to screen tenants (including potential ones) for particular types of voting behaviour. Judging from recent trends in the US renting market, it seems to be very likely. Data analytics companies routinely provide landlords with services that help them screen potential tenants using AI. These companies screen potential tenants based on all sorts of criteria decided by the landlords, such as credit history, criminal records, and eviction checks. Although the data analytics companies do not explicitly say so, it is easy to imagine that one of these criteria they could screen for could be political orientation or outright voting behaviour.[16] Mary's landlord could easily contract with one of these data analytics companies, and ask them to find out how Mary voted in previous elections. And, if the landlord would not like Mary's voting behaviour to be the official reason for evicting Mary, he can ask the company to find more innocent features of Mary that he can use as a reason for evicting her. According to NBC News, it is not yet legally settled what types of data may be included in a tenant screening process, and what types of data may be predicted based on the input data.[17] But even if it turns out that it is illegal for landlords to predict the voting behaviour of tenants, it is probably naive to think that landlords will not do so anyway. After all, as noted above, it is also illegal in most jurisdictions to fire women because they are pregnant, but it is still a commonly known problem that employers do so

[13] For a great overview of empirical studies showing how bad people generally are at probabilistic reasoning, see (Reani et al. 2019).

[14] For example, Mary was trivially late with her latest payment of rent or Mary has failed to furnish her balcony according to the exact regulations of St. John's. According to the lease, both offences are grounds for eviction.

[15] AI constitutes a relatively new technology, and we conjecture that such technology will become more accurate, powerful and accessible in the future when it, with increased speed and reduced costs, can analyse an increased number of data points. It is likely that there will be some diminishing marginal accuracy at some point because each extra data point will begin to correlate so highly with already existing data that each new point does little to increase accuracy. However, this does not undermine our main point: namely that a future Mary's landlord can contract with a data analytics company that has at its disposal an improved AI technology as compared to the technology of the company Mary's landlord contracts with. See (Parikh, Obermeyer, and Navathe 2019: 810) for an example of the impressive development path of AI-based predictive algorithms within the field of medicine.

[16] See for example this: https://newmillsproperties.com/artificial-intelligence-rental-market/#:~:text=Using%20Artificial%20Intelligence%20Rentberry%2C%20in,to%20recommend%20appropriate%20rental%20prices (Accessed May 3, 2022).

[17] See https://www.nbcnews.com/tech/tech-news/tenant-screening-software-faces-national-reckoning-n1260975 (Accessed May 3, 2022).

anyway, and simply find legally innocent ways to officially justify their decision.

The tenant screening industry is growing, and nearly all landlords now use some sort of tenant screening software.[18] So, while the two hypotheticals involving Mary may seem far-fetched at first sight, they are in fact rather realistic, and we will probably see more similar real cases in the future.

4 An objection

Here is one objection to our argument: it was never the secret ballot that protected electors against voting related instances of social ostracism and social punishment in the first place. Rather, it was the employment laws and rent acts that have been in place for a long time. Even in the absence of the secret ballot, current employment laws prohibit employers from firing employees because of their voting behaviour, and current rent acts prohibit landlords from evicting tenants because of their voting behaviour. So, even if the secret ballot no longer protects specific electors from social ostracism and social punishment because of the existence of AI, we should not be too worried because current employment laws and rent acts *still* protect electors from social ostracism and social punishment.

At least three replies to this objection are available to us. The first reply is that current employment laws and rent acts may protect electors against social *punishment,* but they do not protect electors from social *ostracism.* It may protect them against being formally fired or evicted, but it does not protect them against being ostracized by co-workers, family members, or friends. As Hypothetical I shows, even if Mary is not formally evicted from her apartment, she still suffers from the social consequences of other people knowing that it is likely that she voted in favour of same-sex marriage. In the absence of the secret ballot, current rent acts thus do nothing to protect Mary from social ostracism; she is socially ostracized anyway.

The second reply is that current employment laws and rent acts do in fact not even protect electors against social *punishment.* As Hypothetical II shows, current employment laws and rent acts can be easily bypassed. The landlord simply lies about the real reason for why Mary is evicted. So, in the absence of the secret ballot, current rent acts do nothing to protect Mary from social punishment; she is socially punished by being evicted from her apartment anyway.

The third reply is that not all electors are protected by employment laws and rent acts in the relevant way. To be protected by employment laws in the relevant way, you must be an employee. And to be protected by rent acts in the relevant way, you must be a tenant. But not all electors are either an employee or a tenant. For instance, some electors are *employers,* and some electors are *landlords.* So, at least for those electors who are *neither* an employee, nor a tenant, it is a problem if the secret ballot no longer protects against social ostracism and social punishment. This is so because it leaves them even more unprotected against social ostracism and social punishment.

What these replies show is that—just like the secret ballot—current employment laws and rent acts do not protect electors from social ostracism and social punishment. The advent of AI leaves some electors more unprotected against social ostracism and social punishment. Of course, before the advent of AI, it was also possible to infer, with varying levels of accuracy, how particular electors had voted using simple statistical calculations. But with AI, the speed and accuracy at which it is now possible to infer how electors have voted leaves many electors much more unprotected against social ostracism and social punishment.

5 Concluding remarks

We have argued that because of the advent of AI, the secret ballot does not effectively protect all voters from voting related instances of social ostracism and social punishment. Because AI makes it possible to infer with high levels of accuracy how (some) individual voters have voted, friends, family members, co-workers, etc. can now use this knowledge to socially ostracize people. And, landlords, employers, etc. can now use the knowledge to socially punish people. As mentioned, social ostracism and social punishment that is related to voter behaviour was possible even before the advent of AI. But with AI, it is now possible to reach levels of predictive accuracy that are much higher, and it is now possible to infer the voting behaviour of millions of people relatively easy and cheaply. Therefore, the secret ballot now leaves voters much less protected from social ostracism and social punishment than they used to be. This is important in itself, but it is also important because risks of social ostracism and social punishment are some of the most important justifications for introducing the secret ballot into our democratic institutions in the first place.[19]

Acknowledgements The authors would like to thank audiences at Aalborg University, Amsterdam University and Aarhus University for very useful discussions. Also thanks to David Paaske, Beate Rössler, Axel

[18] See https://www.wsj.com/articles/SB9967024419266674 10 (Accessed May 3, 2022).

[19] We wish to thank two anonymous reviewers from *AI & Society: Knowledge, Culture and Communication* for insightful and constructive comments on an earlier version of this paper.

Gosseries, and James Stacey Taylor for useful written comments on an earlier version of this paper.

Funding JTM's work is written as part of a research project funded by the Carlsberg Foundation (CF20-0257).

Availability of data and material N/A.

Declarations

Conflict of interest N/A.

Ethical approval and consent to participate N/A

Consent for publication All authors consent.

References

Brayne S, Christin A (2021) Technologies of crime prediction: the reception of algorithms in policing and criminal courts. Soc Probl 68(3):608–624. https://doi.org/10.1093/socpro/spaa004

Brennan G, Pettit P (1990) Unveiling the vote. Brit J Polit Sci 20(03):311–333

Christiano T (2021) Algorithms, manipulation, and democracy. Can J Philos 1–16

Dahl RA (1998) On democracy. Yale University Press, New Haven

Elklit J, Maley M (2019) Why ballot secrecy still matters. J Democr 30(3):61–75

Engelen B, Thomas V, Nys R (2013) Against the secret ballot: Toward a new proposal for open voting. Acta Politica 48(4):490–507

Erman E, Furendal M (2022) The global governance of artificial intelligence: some normative concerns. Moral Philos Polit. https://doi.org/10.1515/mopp-2020-0046

Gebru T, Krause J, Wang Y, Chen D, Deng J, Aiden EL, Fei-Fei Li (2017) Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States. Proc Natl Acad Sci 114(50):13108. https://doi.org/10.1073/pnas.1700035114

Hinds J, Joinson A (2019) Human and computer personality prediction from digital footprints. Curr Dir Psychol Sci 28(2):204–211. https://doi.org/10.1177/0963721419827849

Idan L, Feigenbaum J (2019) Show me your friends, and I will tell you whom you vote for: predicting voting behavior in social networks. In: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining

Khemasuwan D, Colt HG (2021) Applications and challenges of AI-based algorithms in the COVID-19 pandemic. BMJ Innov 7(2):387. https://doi.org/10.1136/bmjinnov-2020-000648

Kosinski M (2021) Facial recognition technology can expose political orientation from naturalistic facial images. Sci Rep 11(1):1–7

Kristensen JB, Albrechtsen T, Dahl-Nielsen E, Jensen M, Skovrind M, Bornakke T (2017) Parsimonious data: how a single Facebook like predicts voting behavior in multiparty systems. PLoS ONE 12(9):e0184562. https://doi.org/10.1371/journal.pone.0184562

Lever A (2007) Mill and the secret ballot: beyond coercion and corruption. Utilitas 19(3):354–378

Lever A (2015) Privacy and democracy: what the secret ballot reveals. Law, Cult Hum 11(2):164–183. https://doi.org/10.1177/1743872112458745

Manin B (2015) Why open voting in general elections is undesirable. In: Elster J (ed) Secrecy and publicity in votes and debates. Cambridge University Press, New York

Mares I (2015) From open secrets to secret voting: democratic electoral reforms and voter autonomy. Cambridge University Press, New York

Mill JS (1861) Considerations on representative government. Parker, Son, and Bourn, London

Nickerson DW, Rogers T (2014) Political campaigns and big data. J Econ Perspect 28(2):51–73

Parikh RB, Obermeyer Z, Navathe AS (2019) Regulation of predictive analytics in medicine. Science 363(6429):810. https://doi.org/10.1126/science.aaw0029

Peters U (2022) Algorithmic political bias in artificial intelligence systems. Philos Technol 35(2):1–23

Reani M, Davies A, Peek N, Jay C (2019) Evidencing how experience and problem format affect probabilistic reasoning through interaction analysis. Front Psychol 10:1548–1548. https://doi.org/10.3389/fpsyg.2019.01548

Roberts SL (2019) Big data, algorithmic governmentality and the regulation of pandemic risk. Eur J Risk Regul 10(1):94–115. https://doi.org/10.1017/err.2019.6

Rona-Tas A (2020) Predicting the future: art and algorithms. Soc Econ Rev 18(3):893–911. https://doi.org/10.1093/ser/mwaa040

Skeem J, Lowenkamp C (2020) Using algorithms to address trade-offs inherent in predicting recidivism. Behav Sci Law 38(3):259–278. https://doi.org/10.1002/bsl.2465

Sturgis D (2005) Is voting a private matter? J Soc Philos 36(1):18–30. https://doi.org/10.1111/j.1467-9833.2005.00255.x

Sønderholm J, Mainz J, Uhrenfeldt R (2021) Big data analytics and how to buy an election. Public Aff Q 35(2):119–139. https://doi.org/10.2307/48616070

Thompson DF (2002) Just elections : creating a fair electoral process in the United States. University of Chicago Press, Chicago, London

United Nations (General Assembly) (1948) Universal declaration of human rights. United Nations, New York

Vandamme P-E (2018) Voting secrecy and the right to justification. Constellations 25(3):388