


Artificial intelligence and algorithmic exclusion

Catherine Tucker



ACKNOWLEDGMENTS

The author thanks The Hamilton Project team for shepherding the proposal and author's conference participants for generous feedback. Asha Patt provided valuable research assistance.

MISSION STATEMENT

The Hamilton Project seeks to advance America's promise of opportunity, prosperity, and growth.

We believe that today's increasingly competitive global economy demands public policy ideas commensurate with the challenges of the 21st century. The Project's economic strategy reflects a judgment that long-term prosperity is best achieved by fostering economic growth and broad participation in that growth, by enhancing individual economic security, and by embracing a role for effective government in making needed public investments.

Our strategy calls for combining public investment, a secure social safety net, and fiscal discipline. In that framework, the Project puts forward innovative proposals from leading economic thinkers—based on credible evidence and experience, not ideology or doctrine—to introduce new and effective policy options into the national debate.

The Project is named after Alexander Hamilton, the nation's first Treasury secretary, who laid the foundation for the modern American economy. Hamilton stood for sound fiscal policy, believed that broad-based opportunity for advancement would drive American economic growth, and recognized that "prudent aids and encouragements on the part of government" are necessary to enhance and guide market forces. The guiding principles of the Project remain consistent with these views.





Artificial intelligence and algorithmic exclusion

Catherine Tucker

Massachusetts Institute of Technology Sloan School of Management

December 2025

NOTE: This policy proposal is a proposal from the author(s). As emphasized in The Hamilton Project's original strategy paper, the Project was designed in part to provide a forum for leading thinkers across the nation to put forward innovative and potentially important economic policy ideas that share the Project's broad goals of promoting economic growth, broad-based participation in growth, and economic security. The author(s) are invited to express their own ideas in policy proposal, whether or not the Project's staff or advisory council agrees with the specific proposals. This policy proposal is offered in that spirit.

BROOKINGS

Abstract

Artificial intelligence (AI) can be a force for good, especially if it is less biased than human judgment. To automate decisionmaking, AI relies on data being available as an input for these decisions. However, if data inputs are not available, people can be excluded from the benefits of AI because this means that AI outputs about them are missing. This proposal argues that addressing these “data deserts” and consequent algorithmic exclusion should be a priority for policymakers concerned with algorithmic fairness.

Contents

Introduction.....	1
The challenge.....	2
Some examples of missing data and algorithms.....	2
The proposal.....	5
Incorporating algorithmic exclusion considerations into legislation.....	5
A concrete case study: Credit reporting.....	6
Questions and concerns.....	9
What should happen when algorithmic exclusion is identified?.....	9
Are real-time input algorithms really so widespread?.....	9
Can statistical imputation address the issues raised in this proposal?.....	9
Are missing data always a bad thing?.....	9
Could this proposal create an incentive to try to obtain data from people who would actively prefer their data not to be collected?.....	10
Would this proposal be costly for firms?.....	10
Would this proposal be difficult to administer?.....	10
Why isn't it sufficient or preferable to focus on the accuracy or bias of model outputs?.....	11
Could this proposal create an incentive to produce inaccurate output values in the absence of adequate data?.....	11
Conclusion.....	12
References	13

Introduction

Advances in artificial intelligence (AI), in terms of its ability to make predictions at scale, will transform the economy (Agrawal, Gans, and Goldfarb 2016). Rapid advances in AI's capabilities have also raised questions about algorithmic discrimination or bias, and about whether algorithms can be fair. There have been numerous examples of AI systems making decisions that appear unfair and that are likely to reinforce existing inequality.

Economists view AI generally as a set of algorithms applied to data that reduce the costs of making predictions from data. As AI continues to take on roles that have historically been held by human decision-makers, AI predictions increasingly define the boundaries of opportunity. Gándara et al. (2024) show that algorithms used to predict college student success are less accurate when used to predict outcomes for racial minorities. In addition, Lambrecht and Tucker (2019) show that algorithms are less likely to show science, technology, engineering, and mathematics (STEM)-related content to women. Conversely, some work in economics has shown that AI systems can make apparently fairer predictions than humans when it comes to screening resumes (Cowgill 2017; Cowgill and Tucker 2019) or making bail decisions (Kleinberg et al. 2018). Access to education, health services, financial credit, and matters of public safety are all increasingly mediated through AI systems that process large volumes of individual-level data.

So far, policymakers have focused on the fairness of AI by proposing policies that would prevent apparent discrimination. These efforts presume that the central equity issue for AI is the quality of the inferences algorithms make about people from training data. It is true that biased training data will lead to biased outputs: AI systems are only as good as the data that they learn from.

But what if the training data are not only biased, but also incomplete? AI systems can fail not only because they make biased predictions, but also because they make no meaningful predictions at all for certain individuals or populations. That is, unfairness may arise not only because an algorithm sees individuals incorrectly, but also because it does not see them at all when it tries to make real-time predictions.

This proposal is focused on instances when AI is being used to make real-time predictions about an individual and requires data about that individual to be available as an input. Many software algorithms follow a two-step process. In the first step, algorithms are trained on existing data. The potential for bias in this first step has received plenty of well-deserved attention (Obermeyer et al. 2019; O'Neil 2016). In the second step, they use what they have learned from the training data to make a prediction about an individual in real time. This proposal focuses on the second step because, in order for AI to return a prediction about an individual, data about that individual must be available to the system. I refer to the phenomenon of missing predictions due to missing real-time data as "algorithmic exclusion," defined as failure or harm arising from insufficient input data.

I propose a concrete, policy-relevant addition to regulations and proposals on AI fairness: incorporate algorithmic exclusion as a class of algorithmic harm equal in importance to bias and discrimination (Tucker 2023). When a regulation aims to prevent "bias, discrimination, and other harms" (Eliminating Bias in Algorithmic Systems Act of 2023, S.3478), it should also specifically address the ways in which certain people are excluded from, or unable to derive benefits from, having information relating to them included in an algorithm's training data or assumptions.

The challenge

For AI, the absence of data—whether the data are fragmented, outdated, low quality, or missing—can be just as harmful as the inclusion of overtly biased data or the use of models trained on biased training data. These data deserts lead to systematic under-recognition of the very populations that an equity-focused policy would be designed to protect (Neumann et al. 2024). Some attention has been paid in fairness debates to issues of representativeness in training data, but less attention has been paid to the implications of missing data at the application stage of algorithmic decisionmaking—that is, when a model tries to evaluate an individual based on real-time or situational inputs. Algorithmic exclusion formally describes failure when an AI-driven system lacks enough data on an individual to return an output about them.

This phenomenon is not a side effect of poor design; rather, it is a structural feature of digital inequality. The same economic and social forces that marginalize individuals offline—including lack of internet access, fewer interactions with institutions that collect data, and lower participation in activities that generate traceable data—also reduce their visibility online (Neumann et al. 2024). These gaps create what are referred to as data deserts: zones where AI cannot function effectively.

The digital era has allowed firms to collect, store, and parse data at far lower cost than ever before (Goldfarb and Tucker 2019). Though the costs of collecting and storing data and feeding it to algorithms have fallen, that does not mean that data are collected evenly across individuals. Often, having a broad digital footprint is a matter of privilege, which reflects having reliable access to digital technologies. Therefore, as discussed further below, lower-income households often have less data amassed about them.

The way that many instances of AI or algorithms work is that they are pre-trained on a certain dataset to make predictions. Much work on algorithmic bias is focused on the training of the algorithm and how bias in data can lead the parameterization of such a model (or weights) to reflect existing bias that was in the training data (Obermeyer et al. 2019; O’Neil 2016). Then, after the algorithm is trained, in order to make a prediction about a particular individual, that person’s data need to be inputted into the algorithm. If there

are insufficient data, then the algorithm will work less well and may fail to return a prediction.

This means that algorithms are more able to make predictions about those individuals who have deep digital footprints. Those who do not have sufficient data for a prediction to be made may face either an algorithm unable to make a prediction about them, with potentially detrimental consequences, or a less accurate prediction, if the algorithm imputes the missing data.

Consider the hypothetical example of an AI system used to screen loan applications. A well-designed fairness audit might assess whether the system produces disparate loan approval rates by race or gender. But if applicants from certain groups have missing data—i.e., they are less likely to have credit histories, formal employment records, or stable addresses—they may never even reach the stage where the algorithm decides. The applicants are not misclassified, though—a prediction about them is simply missing.

Some examples of missing data and algorithms


Focusing on algorithmic auditing of outputs alone will identify some kinds of biases that affect those whose data are fed into AI systems.

But what of the needs of those whose data are not fed into AI systems? I recognize that the average American interacts with digital infrastructure dozens—if not hundreds—of times per day, from mobile apps to payment systems to social media and location services. In this context, the idea that there are people about whom insufficient digital data exist may seem outdated. However, such a perspective overlooks the uneven distribution of digital participation.

Prior research in digital economics shows that the volume of data generated is not just a function of the number of interactions, but also of the type of infrastructure available, the incentives for data collection, and the institutional channels through which data enter algorithmic systems (Lambrech and Tucker 2019). Many people—such as the elderly and those with lower incomes—have less access to the internet and smartphones and thus generate data at lower rates (figure 1). This absence of data is not due to oversight or bias in the traditional sense, but rather to an incompatibility

FIGURE 1

Digital connectedness, by age and household income, 2023



Source: U.S. Census Bureau 2023; author's calculations.

THE
HAMILTON
PROJECT

BROOKINGS

between how people live and how data are captured (Chiou and Tucker 2020; Sen and Tucker 2022).

Here I provide conceptual and illustrative examples of algorithmic exclusion that stems from this data disparity and is distinct from algorithmic bias.

Conceptual illustration: A GPS system

Imagine your car or phone's GPS system. The algorithmic exclusion issue is akin to what happens when you enter a tunnel. At that point, the GPS system no longer has access to data about your real-time location. As a result, it can no longer make accurate predictions

about what lane you should be in or how far you are from your exit. This effectively excludes you from the benefits of algorithmic prediction. It is precisely this scenario, writ large across the economy for millions of people, that this proposal is attempting to prevent.

This is different from the accuracy of the map that the GPS system uses or the sophistication of the system. Algorithmic exclusion refers to when a system entirely fails to deliver because it is missing the crucial real-time data it needs (in this case, your location).

COVID-19 vaccine allocation algorithm

The U.S. COVID-19 vaccine allocation algorithm, named Tiberius, used data from the American Community Survey of 2018 to allocate vaccines in the early days of vaccine availability, when demand for COVID-19 vaccines outstripped supply (Simunaci 2020). The issue is that the American Community Survey has been shown to undercount certain groups; analysis of the 2020 Census found that:

- Renters are undercounted by 1.5 percent;
- Black Americans are undercounted by 3.3 percent; and
- Hispanic Americans or Latinos are undercounted by 4.5 percent (U.S. Census Bureau 2022).

Now this is a simple example of an algorithm that was set up to try to allocate scarce resources fairly—unlike many more-dynamic AI systems that absorb data in real time. But this straightforward case makes it apparent how missing data can affect algorithmic outcomes for some groups that were hardest hit by the COVID-19 pandemic.

The proposal

In the U.S., as yet, most policy action addressing algorithmic fairness has taken place through executive orders and enforcement actions by federal agencies. For example, the Federal Trade Commission (FTC) has begun using its authority to examine algorithmic bias, starting with a 2023 case against Rite Aid Corporation for its use of facial recognition technology without appropriate safeguards, which was used to identify shoplifters (Bedoya 2023; Federal Trade Commission 2023; O'Neill and Hutchinson 2024).

By contrast, the European Parliament has passed broad and sweeping laws that apply across the European Union (EU). The EU Artificial Intelligence Act (EU AI Act) was enacted in 2024 and will gradually come into force in a variety of AI contexts through 2026–27 (Ernst & Young 2024; European Commission 2025). This law introduced a class of prohibited uses of AI and a categorization of high-risk AI systems, where AI is used in areas such as health, justice, and education that have large potential consequences (European Commission 2025). These high-risk systems are subject to particular scrutiny and regulation, such as mandatory risk assessments, quality management systems, and conformity assessments (Ernst & Young 2024; European Commission 2025).

Some proposed legislation in the U.S. has some similarities to the EU AI Act. Proposed legislation introduced in the Senate, the Eliminating Bias in Algorithmic Systems Act of 2023 (S.3478), would require that any federal agency that uses or oversees AI have an office of civil rights focused on combating AI bias and discrimination. This bill calls for the offices of civil rights of relevant agencies to report on the steps these agencies have taken to mitigate “bias, discrimination, and other harms from covered algorithms,” including outreach to stakeholders. It further calls on these offices to provide recommendations on legal or administrative actions that could mitigate “bias, discrimination, and other harms from covered algorithms” (Eliminating Bias in Algorithmic Systems Act of 2023, S.3478).

The Algorithmic Accountability Act of 2025 (S.2164) was introduced in the Senate in June 2025; it seeks to extend the role of the FTC in regulating AI, focusing on several key sectors. The key sectors (e.g., health, justice, and education) are those that are more likely to make critical decisions; the act defines a critical decision as

a decision or judgment that has any legal, material, or similarly significant effect on a consumer’s life relating to access to or the cost, terms, or availability of (A) education and vocational training, including assessment, accreditation, or certification; (B) employment, workers management, or self-employment; (C) essential utilities, such as electricity, heat, water, internet or telecommunications access, or transportation; (D) family planning, including adoption services or reproductive services; (E) financial services, including any financial service provided by a mortgage company, mortgage broker, or creditor; (F) healthcare, including mental healthcare, dental, or vision; (G) housing or lodging, including any rental or short-term housing or lodging; (H) legal services, including private arbitration or mediation; or (I) any other service, program, or opportunity decisions about which have a comparably legal, material, or similarly significant effect on a consumer’s life as determined by the Commission through rulemaking. (Algorithmic Accountability Act of 2025, S.2164)

Incorporating algorithmic exclusion considerations into legislation

When policymakers draft legislation or regulations to address potential algorithmic bias, they should not just focus on bias in outputs, but also on the potential for missing or inaccurate predictions that occur because the algorithm did not have data available. Therefore, when legislation or regulation refers to algorithmic bias or discrimination, I propose that it should also refer to instances when the algorithm did not return any prediction because data were missing (i.e., algorithmic exclusion). Neither the Eliminating Bias in Algorithmic Systems Act of 2023 nor the Algorithmic Accountability Act of 2025 address algorithmic exclusion.

I propose that policymakers include algorithmic exclusion when seeking to address these issues, for example, “mitigat[ion of] bias, discrimination, exclusion, and other harms from covered algorithms” (Eliminating

TABLE 1

Addressing algorithmic exclusion in US and EU AI laws

Aspect	US Eliminating Bias in Algorithmic Systems Act (S.3478; H.R.10092)	Tucker proposal for US	EU AI Act	Tucker proposal for EU
Definition of bias	Refers to “bias, discrimination, and other harms” from covered algorithms.	Add “algorithmic exclusion,” meaning harms to those for whom there are no input data.	Focuses on unjustified detrimental treatment or exploitation of vulnerabilities.	Expand to include the idea of unjustified exclusion from algorithmic decisionmaking due to missing data and missing output.
Risk assessment tools	Requires risk assessments on bias and discrimination.	Add representation audits to identify who may not have the requisite data for the algorithm to make predictions.	Risk assessment mandatory for high-risk AI.	Include evaluation of the risks of missing data and missing output.
Audit standards	Focuses on documenting mitigation of known bias.	Add obligation to identify risks of missing data and missing output.	Mandates accuracy, robustness, and human oversight.	Integrate missing data and missing output checks to human oversight.



BROOKINGS

Bias in Algorithmic Systems Act of 2023, S.3478). Without attending to who is missing from the system altogether, well-meaning regulatory efforts risk reinforcing a two-tiered structure, where some individuals are profiled and others are not. Table 1 provides comparisons between extant EU and U.S. legislation and laws as well as my proposals.

Lower-income Americans are especially vulnerable to exclusion-by-absence. First, as discussed, they are less likely to interact with the types of digital systems that generate persistent, high-quality data trails: credit systems, formal employment networks, subscription platforms, or high-frequency online consumption behaviors. Second, when they do interact, the data that are generated are often incomplete, irregular, or siloed in ways that prevent meaningful aggregation across contexts. As I explain in this proposal, the aggregation of datasets requires the ability to match data across individuals using a consistent “key” such as a phone number or name, and the consistency of this “key” is also a matter of privilege.

The policy implication is that lower-income individuals might not only suffer from algorithmic misclassification due to biased data, but might also be invisible to these systems entirely and have outputs completely missing. For example, a credit algorithm that requires five years of formal income history will simply fail to evaluate a worker who operates in informal or gig-based economies. A resume-screening algorithm that draws on structured employment records will skip over applicants with nonlinear or undocumented career paths. These are not edge cases,

but rather examples that are symptomatic of a broader pattern: Algorithmic systems work best for those already embedded in the data-rich infrastructures of formal economic participation.

By explicitly incorporating the problem of data nonavailability and missing predictions into algorithmic auditing and regulatory frameworks, this proposal ensures that the absence of data is treated as a substantive fairness issue, rather than as a mere technical inconvenience. Addressing algorithmic exclusion benefits not only lower-income Americans, but also anyone operating outside dominant data-generating regimes: new immigrants, survivors of domestic violence using alias names, the digitally disconnected, and others with complex data footprints.


A concrete case study: Credit reporting

Let us take the concrete example of the Protecting Your Credit Score Act of 2020 (H.R.5332) that was passed in the House of Representatives but not taken up by the Senate. Despite bipartisan support in the House, the major credit bureaus and industry opposed the bill because the bill established a free, one-stop portal for access to all three agency credit reports (Consumer Data Industry Association 2020).

The bill as written focused on improving the accuracy of outputs, and on the accuracy of inputs into the credit scoring decision. It had a focus on ensuring that incorrect information not be wrongfully included

FIGURE 2

Share of homes that are owner-occupied or renter-occupied, by race and ethnicity, 2023



Source: U.S. Census Bureau 2023; author’s calculations.

Note: White, Black, and Asian categories do not include people identifying as Hispanic or Latino. “Other” includes American Indian, some other race, and two or more races.



in an individual’s credit history. However, there was no language to ensure that equal weight is given to what information or data is excluded from an individual’s credit history.

A simple example of absent information can be seen in the form of rent payments. Classically, in the U.S., rent payments are not included in an individual’s credit score. Instead, only mortgage payments are included. To see the potential unequal impact of this omission, consider the data in figure 2 on who rents and who owns houses in the U.S. by race and ethnicity. It is clear there are far more white Americans who own houses, proportionally relative to Black or Hispanic Americans. In other words, the fact that mortgage data are missing from a credit bureau report is systematically correlated with race.

This pattern continues when looking at income, as shown by figure 3. That figure shows that, as you might expect, mortgage payment data are concentrated in those who are well off and are likely to be missing among low-income households.

In figures 2 and 3, renters are missing a crucial input into their credit score, input data that would potentially establish a record and history of payments. The missing inputs of data will necessarily lead to uneven outputs for these scores. Future legislation could ask that, when an individual receives their credit report and asks for an explanation for the credit agency’s decision, the agency would tell the person what data the algorithm was missing when it made that decision, which could have helped it come to a different decision. The credit applicant’s housing history

or borrowing history, for example, might be missing. It would not be expensive to provide this extra piece of information explaining a credit decision, but it would be helpful for people trying to understand their credit score.

This example also illustrates some of the tradeoffs of addressing missing data. It is currently possible for renters to have their rent reported to credit bureaus. However, rent reporting faces significant cost barriers that limit its widespread adoption. For landlords, especially landlords managing small or informal properties, the administrative and technological burden of setting up reporting systems can be prohibitive. For example, RealPage charges a renter \$5 a month to report rent (RealPage.com n.d.). Those landlords that do report rent they receive rely on third-party platforms that charge monthly fees for transmitting rent payment data; those fees could either be passed on to tenants or absorbed as operating costs. Moreover, because credit bureaus are not required to include rental data, and lenders are not obligated to consider it, the impact of rent reporting may be uneven and limited. Meanwhile, mandating that landlords report to credit bureaus could increase rents, and the value for improving the accuracy of credit scoring is unknown, as is the impact on racial disparities in credit scores.

I emphasize that rental data are not the only kinds of missing data in Americans’ credit reports. For example, many Americans use informal credit, such as payday loans, in a way that is unreported for the purposes of credit scoring (Consumer Financial Protection Bureau 2024).

FIGURE 3

Share of homes that are owner-occupied or renter-occupied, by household income, 2023



Source: U.S. Census Bureau 2023; author's calculations.

Note: Household income is in 2023 dollars.

THE
HAMILTON
PROJECT

BROOKINGS

Questions and concerns

What should happen when algorithmic exclusion is identified?

This proposal does not make recommendations for how to respond to algorithmic exclusion. Instead, my focus is on ensuring that algorithmic exclusion as a concept is not omitted when policymakers are seeking to address matters of bias and representativeness in AI.

Are real-time input algorithms really so widespread?

Another area of skepticism may center on the scope of real-time algorithmic decisionmaking. It is true that not all algorithmic systems rely on user-specific inputs at the moment of decision. Some operate on pre-trained models with limited dynamic input. However, when looking at the AI systems that generate much of the debate around algorithmic fairness—recidivism prediction tools, automated credit scoring, targeted digital advertising, and resume-screening platforms—these systems all depend fundamentally on receiving individual-level data in real time or near real time.

What distinguishes these systems is not just their use of algorithms, but also their operational dependence on contemporaneous data from users, applicants, or subjects. These algorithms are not static in nature and do not make a single prediction at a single time; instead, they attempt to tailor their outputs to the specifics of a situation or individual at that moment. That very tailoring is what increases their vulnerability to failures when input data are missing, absent, or incomplete.

Moreover, the trend in AI deployment is toward more, not less, real-time personalization. As a result, the prevalence of algorithms that cannot function without individual input in real time is likely to grow.

Can statistical imputation address the issues raised in this proposal?

To a data scientist, the description of individuals being excluded due to missing data may sound like a familiar technical problem with well-documented fixes. The standard toolset includes various forms of imputation,

or statistical methods that aim to fill in missing values in a dataset. These range from simple approaches, like replacing missing entries with a mean or median, to more sophisticated techniques, such as forward or backward filling or k-nearest neighbors (abbreviated as KNN) imputation, where data from similar individuals are used as a proxy. Of course, these techniques all have downsides and can themselves lead to the reinforcing of inequality.

While these techniques are robust in many settings, they are not designed for algorithmic systems operating in real time with no data from the subject in question. For imputation to function, there must be some existing correlation structure in the available data. But many exclusion cases are not just about a single variable being missing but are also about a complete absence of relevant inputs for the algorithm to operate. In effect, there is no row in the dataset for the individual.

Even though some might argue that imputation is at least one way of dealing with this missing problem, my research suggests that, in many real-life systems, including predictions made in companies that are owned or adjacent to credit bureaus, the presence of missing data does not trigger sophisticated imputation, but rather simply defaults to not returning a prediction (Neumann et al. 2024).

Are missing data always a bad thing?

So far in this policy proposal, I have assumed that having an AI prediction made about an individual is better than having no AI prediction for that individual due to missing data. However, there are two contexts where this assumption might be false.

First, I have been assuming that an AI system would be put in place with the intention of benefiting those for whom it generates predictions, for example, by increasing access to government benefits. However, systems may not be designed with this intent in mind. For example, Palantir's Investigative Case Management system (EPIC v. ICE 2018) fuels detection by ICE of undocumented people and gathers data exhaustively on the segment of the population that they predict to be plausibly undocumented or connected to undocumented people for the purpose of excluding those people from the benefits associated with

residency in the U.S. The Department of Government Efficiency is working with U.S. Department of Housing and Urban Development data to try to identify U.S. citizen residents of public housing who have undocumented people as members of their household for the purpose of excluding those U.S. citizens along with their household members from public housing (Siegel, Natanson, and Meckler 2025). A potential nuance, therefore, for well-intentioned policy is that this policy proposal should be caveated with the emphasis that it applies only when algorithmic predictions' existence increases an individual's welfare rather than harming that individual's welfare.

Second, the missing data might reflect intentionality on the part of the individual. This policy proposal has been centered on cases where data on individuals are missing for systemic reasons, and not because of conscious choices made by those individuals to not create that data. However, I recognize that there are some individuals who consciously abstain from using credit cards or smartphones to purposely not leave a digital footprint. I should be clear that this policy proposal is focused on instances where individuals' missing data are not a result of their conscious choice to not create that data. My proposal is not suggesting, for example, that the National Security Agency should be allowed to collect data on every U.S. citizen. In the government context, there are often constraints on whether the state gathers data on people, and there are contexts, such as criminal justice, where the objective of accurate prediction or decisionmaking should not be paramount in the sense of requiring uniform and universal gathering of as many fields as possible on all people, because that in itself might be undesirable along other axes.

Could this proposal create an incentive to try to obtain data from people who would actively prefer their data not to be collected?

Relatedly, there is a more general worry that such a proposal would give incentives for mass surveillance. However, this concern can be resolved by making it clear that contexts where data are missing because a person has consciously chosen for their data to be missing are exempt from any repercussions.

Would this proposal be costly for firms?

A central concern in any policy proposal that modifies algorithmic regulation is the potential for added compliance costs. Here, the proposed intervention—explicitly accounting for missing data in algorithmic inputs—may appear to add complexity to firms' existing fairness audits. However, there are three reasons to

expect that the cost burden will be modest, and possibly even efficiency-enhancing, over the longer term.

First, the proposal does not require firms to collect new types of data, nor to change the underlying objectives of their algorithmic systems: It merely requires transparency and documentation around a class of failures that are already occurring silently.

Second, many firms already conduct internal audits to detect and mitigate data sparsity in model features—particularly for performance reasons. Embedding a lens into what data may be missing from these algorithms is simply a reframing of a typical data science activity, rather than something that represents incremental costs.

Third, the failure to handle missing data adequately can itself create costly legal and reputational risks. Users who are misclassified or excluded due to missing data could challenge decisions, stop interacting with the firm, or litigate. By proactively identifying and remediating exclusion risks due to missing data, firms not only might meet regulatory expectations, but also might strengthen the robustness of their systems (White House 2025).

Would this proposal be difficult to administer?

One might also worry that identifying and regulating missing input data could introduce administrative complexity. However, this concern overlooks a key asymmetry: Missing data are, by nature, observable. Unlike structural bias—where the challenge often lies in identifying subtle patterns in model outputs—missingness can be programmatically detected. An input variable required by the model but not present at the time of prediction is a measurable and auditable event.

Moreover, many of the tools and techniques developed for bias auditing—such as disparate impact analysis, error decomposition, or counterfactual simulation—can be adapted to examine patterns of missingness across demographic groups. In fact, from an administrative perspective, missing data might instead be easier to track and assess than outcome-based discrimination, which often requires sensitive or legally protected group labels that firms are reluctant or unable to collect.

The marginal administrative cost of incorporating missing data awareness into algorithmic oversight is likely to be low, relative to the potential fairness gains. By shifting the regulatory lens to include not only what data are used, but also who is excluded because their data are missing, we move toward a more comprehensive and more inclusive vision of algorithmic accountability.

Why isn't it sufficient or preferable to focus on the accuracy or bias of model outputs?

Most of the policy debate and regulation has focused on outputs. This makes sense a priori because ultimately, if we are worried about economic outcomes, then it makes sense to focus on the outputs of AI systems. As an economist, I completely agree with this perspective. The reason, however, that this proposal focuses on inputs rather than outputs is because, currently, policy attention has almost exclusively been about outputs that happen. They do not deal explicitly with predictions that are not made. In the end, audits focused on the biases of predictions will not catch instances where many disadvantaged people do not have predictions made about them at all.

Could this proposal create an incentive to produce inaccurate output values in the absence of adequate data?

Correspondingly, a worry could be that such a proposal that focuses on auditing and inspecting how many outputs are missing could lead to a rise in poor-quality outputs. The question then becomes whether it is worse to have missing outputs or to have no outputs at all. One potential argument, though, is that having outputs allows for a complete analysis of what is going wrong in the system, meaning that the root causes of the low quality of those outputs can be better assessed. And this may particularly be the case if organizations were potentially able to sidestep existing regulations or auditing that focus on output quality by simply having missing output.

Conclusion

This proposal seeks to augment the dominant frame of algorithmic fairness with a complementary, and often overlooked, lens: the role of data absence in shaping who benefits from algorithmic systems. While much of the discussion surrounding AI fairness has focused on the risks of biased representations within datasets, there has been relatively little attention paid to those who do not have data that can be an input for AI systems. These are individuals for whom no data are available at the moment of algorithmic interaction, and who therefore risk being excluded from the benefits or protections afforded by these AI systems.

This proposal explains how policies focused on algorithmic justice need to explicitly incorporate the potential for missing data when they seek solutions. This allows them to address a failure mode where individuals lack sufficient data to be effectively represented within AI-driven systems. Algorithmic exclusion is not simply the absence of data—it is the absence of recognition of this absence of data in a world where inclusion in digital systems is increasingly a prerequisite for access to resources, opportunities, and protections.

References

- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2016. "The Simple Economics of Machine Intelligence." *Harvard Business Review* 17: 2–5.
- Bedoya, Alvaro M. 2023. "Statement of Commissioner Alvaro M. Bedoya on FTC v. Rite Aid Corporation & Rite Aid Headquarters Corporation, Commission File No. 2023190." Federal Trade Commission. December 19, 2023. https://www.ftc.gov/system/files/ftc_gov/pdf/2023190_commissioner_bedoya_riteaid_statement.pdf
- Chiou, Lesley, and Catherine Tucker. 2020. "In Times of Crisis, Digital Businesses Need to Think About 'Technology Laggards.'" *MIT Sloan Management Review*. <https://sloanreview.mit.edu/article/crisis-digital-business-technology-laggards/>
- Consumer Data Industry Association. 2020. "Protecting Your Credit Score Act." <https://www.cdiaonline.org/views/2020/07/21/protecting-your-credit-score-act/>
- Consumer Financial Protection Bureau. 2024. "Can Taking Out a Payday Loan Help Rebuild My Credit or Improve My Credit Score?" <https://www.consumerfinance.gov/ask-cfpb/can-taking-out-a-payday-loan-help-rebuild-my-credit-or-improve-my-credit-score-en-1611/>
- Cowgill, Bo. 2017. "Bias and Productivity in Humans and Algorithms: Theory and Evidence from Resume Screening." Working paper. https://conference.iza.org/conference_files/MacroEcon_2017/cowgill_b8981.pdf
- Cowgill, Bo, and Catherine Tucker. 2019. "Economics, Fairness and Algorithmic Bias." Working Paper in preparation for *The Journal of Economic Perspectives*, National Bureau of Economic Research, Cambridge, MA. <https://conference.nber.org/confer/2019/YSAIf19/SSRN-id3361280.pdf>
- EPIC v. ICE (Palantir Databases). 2018. FOIA Cases. US District Court for the District of Columbia. <https://epic.org/documents/epic-v-ice-palantir-databases/>
- Ernst & Young. 2024. "EU AI Act." Ernst & Young Global Limited. July 2024. <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-gl/insights/public-policy/documents/ey-gl-eu-ai-act-07-2024.pdf>
- European Commission. 2025. "AI Act." Shaping Europe's Digital Future. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- Federal Trade Commission. 2023. Complaint for Permanent Injunction and Other Relief. Federal Trade Commission v. Rite Aid Corporation, No. 2:23-cv-05023 (E.D. Pa. Dec. 19, 2023). https://www.ftc.gov/system/files/ftc_gov/pdf/2023190_riteaid_complaint_filed.pdf
- Gándara, Denisa, Hadis Anahideh, Matthew P. Ison, and Lorenzo Picchiarini. 2024. "Inside the Black Box: Detecting and Mitigating Algorithmic Bias Across Racialized Groups in College Student-Success Prediction." *AERA Open* 10.
- Goldfarb, Avi, and Catherine Tucker. 2019. "Digital Economics." *Journal of Economic Literature* 57 (1): 3–43.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. "Human Decisions and Machine Predictions." *Quarterly Journal of Economics* 133 (1): 237–93.
- Lambrecht, Anja, and Catherine Tucker. 2019. "Algorithmic Bias? An Empirical Study of Apparent Gender-Based Discrimination in the Display of STEM Career Ads." *Management Science* 65 (7): 2966–81.
- Neumann, Nico, Catherine E. Tucker, Levi Kaplan, Alan Mislove, and Piotr Sapiezynski. 2024. "Data Deserts and Black Boxes: The Impact of Socio-Economic Status on Consumer Profiling." *Management Science* 70 (11): 8003–29.
- Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." *Science* 366 (6464): 447–53.
- O'Neil, Cathy. 2016. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York, NY: Crown Publishing Group.
- O'Neill, Julie, and Elisabeth Hutchinson. 2024. "The FTC Brings Algorithmic Bias into Sharp Focus." Client Alert: Morrison Foerster. <https://www.mofo.com/resources/insights/240105-the-ftc-brings-algorithmic-bias-into-sharp-focus>
- RealPage.com. n.d. "RealPage Rent Reporting: Get Credit for Your On-Time Rent Payments." <https://www.realpage.com/resident-resource-center/rent-reporting/>
- Sen, Ananya, and Catherine Tucker. 2022. "Digital Disruption in Schooling and the Pandemic: Documenting the Digital Infrastructure Divide Among School Children." National Bureau of Economic Research, Cambridge, MA. https://www.nber.org/sites/default/files/2022-02/Tucker_Digital%20Disruption%20in%20Schooling%20and%20the%20Pandemic.pdf
- Siegel, R., H. Natanson, and L. Meckler. 2025. "DOGE Is Collecting Federal Data to Remove Immigrants from Housing, Jobs." *Washington Post*, April 15, 2025. <https://www.washingtonpost.com/business/2025/04/15/doge-ssa-immigration-trump-housing/>
- Simunaci, Lisa. 2020. "Tiberius Platform Aids COVID-19 Logistics, Delivery." U.S. Department of Defense.
- Tucker, Catherine. 2023. "Algorithmic Exclusion: The Fragility of Algorithms to Sparse and Missing Data." Working Paper, Center on Regulation and Markets, Brookings Institution, Washington, DC.
- U.S. Census Bureau. 2023. "American Community Survey." U.S. Census Bureau, Suitland, MD.
- U.S. Congress. House. 2020. Protecting Your Credit Score Act of 2020. HR 5332. 116th Cong., 1st sess. Introduced in House December 6, 2019. <https://www.congress.gov/bill/116th-congress/house-bill/5332>
- U.S. Congress. Senate. 2023. Eliminating Bias in Algorithmic Systems Act of 2023. S 3478. 118th Cong., 1st sess. Introduced in Senate December 12, 2023. <https://www.congress.gov/bill/118th-congress/senate-bill/3478>
- U.S. Congress. Senate. 2025. Algorithmic Accountability Act of 2025. S 2164. 119th Cong., 1st sess. Introduced in Senate June 25, 2025. <https://www.congress.gov/bill/119th-congress/senate-bill/2164>
- White House. 2025. "Restoring Equality of Opportunity and Meritocracy." Executive Order, April 23, 2025. White House, Washington, DC. <https://www.whitehouse.gov/presidential-actions/2025/04/restoring-equality-of-opportunity-and-meritocracy/>

GEORGE A. AKERLOF

University Professor,
Georgetown University

ROGER C. ALTMAN

Founder and Senior Chairman,
Evercore

KAREN L. ANDERSON

Managing Director,
Social Finance Institute

JARED BERNSTEIN

Senior Policy Fellow, Stanford Institute for
Economic Policy Research and the
Center for American Progress

ALAN S. BLINDER

Gordon S. Rentschler Memorial Professor of
Economics and Public Affairs,
Princeton University

KIMBERLY CLAUSING

Eric M. Zolt Chair in Tax Law and Policy,
UCLA School of Law

RAY DALIO

Founder, CIO Mentor, and
Member of the Bridgewater Board,
Bridgewater Associates

BRIAN DEESE

Institute Innovation Fellow,
Massachusetts Institute of Technology

JOHN M. DEUTCH

Emeritus Institute Professor,
Massachusetts Institute of Technology

WILLIAM C. DUDLEY

Senior Advisor, Griswold Center,
Princeton University;
Chair, Bretton Woods Committee;
Former President, Federal Reserve Bank of New York

WENDY EDELBERG

Senior Fellow, Economic Studies,
The Brookings Institution

BLAIR W. EFFRON

Co-Founder and Partner,
Centerview Partners

DOUGLAS W. ELMENDORF

Harvard University Distinguished Service Professor
and Lucius N. Littauer Professor of Public Policy,
Harvard Kennedy School

JUDY FEDER

Professor and Founding Dean,
McCourt School of Public Policy,
Georgetown University

MICHAEL FROMAN

President,
Council on Foreign Relations

JASON FURMAN

Aetna Professor of the Practice of
Economic Policy, Harvard University;
Senior Fellow, Peterson Institute for
International Economics;
Senior Counselor, The Hamilton Project

MARK T. GALLOGLY

Co-Founder and Managing Principal,
Three Cairns Group

TED GAYER

President,
Niskanen Center

TIMOTHY F. GEITHNER

Chairman, Warburg Pincus;
Senior Counselor, The Hamilton Project

ROBERT GREENSTEIN

Visiting Fellow, The Hamilton Project,
Economic Studies, The Brookings Institution;
Founder and President Emeritus,
Center on Budget and Policy Priorities

MICHAEL GREENSTONE

Milton Friedman Distinguished Service
Professor in Economics,
Founding Director, Institute for
Climate and Sustainable Growth,
Director, Energy Policy Institute at UChicago (EPIC),
Kenneth C. Griffin Department of Economics,
University of Chicago

BEN HARRIS

Vice President and Director,
Economic Studies, The Brookings Institution

GLENN H. HUTCHINS

Co-Chair, The Brookings Institution;
Chairman, North Island

LAWRENCE F. KATZ

Elisabeth Allison Professor of Economics,
Harvard University

MELISSA S. KEARNEY

Director, Aspen Economic Strategy Group;
Gilbert F. Schaefer Professor of Economics,
University of Notre Dame;
Nonresident Senior Fellow, Economic Studies,
The Brookings Institution

EUGENE A. LUDWIG

27th Comptroller of the Currency;
Founder and CEO, Ludwig Advisors

KRISTON MCINTOSH

Chief Communications Officer,
American Civil Liberties Union

ERIC MINDICH

Founder and CEO,
Everblue Management

SUZANNE NORA JOHNSON

Co-Chair, The Brookings Institution; Former Vice
Chairman, Goldman Sachs Groups, Inc.

MICHAEL O'MARY

Managing Director, D. E. Shaw & Co., L.P.

PETER ORSZAG

CEO and Chairman, Lazard;
Senior Counselor, The Hamilton Project

PENNY PRITZKER

Chairman and Founder, PSP Partners

MEEGHAN PRUNTY

Senior Philanthropic Advisor, Schusterman Family
Philanthropies & Doris Duke Foundation;
Senior Advisor, Robert E. Rubin

ROBERT D. REISCHAUER

Distinguished Institute Fellow and
President Emeritus,
Urban Institute

NANCY L. ROSE

Charles P. Kindleberger Professor of
Applied Economics,
Massachusetts Institute of Technology

DAVID M. RUBENSTEIN

Co-Founder and Co-Chairman,
The Carlyle Group

ROBERT E. RUBIN

Former U.S. Treasury Secretary;
Co-Chair Emeritus,
Council on Foreign Relations

LEWIS A. (LEE) SACHS

Co-Founder and Managing Partner,
Gallatin Point Capital

LESLIE B. SAMUELS

Senior Counsel,
Cleary Gottlieb Steen and Hamilton LLP

NATASHA SARIN

Professor of Law, Yale Law School;
President and Co-Founder, The Budget Lab at Yale;
Senior Counselor, The Hamilton Project

DIANE WHITMORE SCHANZENBACH

McCourt Chair, McCourt School of Public Policy,
Georgetown University; Nonresident Senior Fellow,
Economic Studies, The Brookings Institution

RALPH L. SCHLOSSTEIN

Chairman Emeritus,
Evercore

ERIC SCHMIDT

Co-Founder, Schmidt Futures;
Former CEO and Chairman, Google

ERIC SCHWARTZ

Chairman and CEO,
76 West Holdings;
Partner, SSW Partners

JAY SHAMBAUGH

Professor of Economics and International Affairs,
The George Washington University;
Nonresident Senior Fellow, Economic Studies,
The Brookings Institution

JOSHUA L. STEINER

Partner, SSW Partners

THOMAS F. STEYER

Co-Executive Chair,
Galvanize Climate Solutions

MICHAEL R. STRAIN

Director of Economic Policy Studies,
American Enterprise Institute

ALLEN THORPE

Partner,
Hellman & Friedman

LAURA D'ANDREA TYSON

Distinguished Professor of the Graduate School,
University of California, Berkeley

OWEN ZIDAR

Professor of Economics and Public Affairs,
Princeton University

AVIVA ARON-DINE

Director

Artificial intelligence (AI) can be a force for good, especially if it is less biased than human judgment. To automate decisionmaking, AI relies on data being available as an input for these decisions. However, if data inputs are not available, people can be excluded from the benefits of AI because this means that AI outputs about them are missing. This proposal argues that addressing these “data deserts” and consequent algorithmic exclusion should be a priority for policymakers concerned with algorithmic fairness.

Addressing algorithmic exclusion in US and EU AI laws

Aspect	US Eliminating Bias in Algorithmic Systems Act (S.3478; H.R.10092)	Tucker proposal for US	EU AI Act	Tucker proposal for EU
Definition of bias	Refers to “bias, discrimination, and other harms” from covered algorithms.	Add “algorithmic exclusion,” meaning harms to those for whom there are no input data.	Focuses on unjustified detrimental treatment or exploitation of vulnerabilities.	Expand to include the idea of unjustified exclusion from algorithmic decisionmaking due to missing data and missing output.
Risk assessment tools	Requires risk assessments on bias and discrimination.	Add representation audits to identify who may not have the requisite data for the algorithm to make predictions.	Risk assessment mandatory for high-risk AI.	Include evaluation of the risks of missing data and missing output.
Audit standards	Focuses on documenting mitigation of known bias.	Add obligation to identify risks of missing data and missing output.	Mandates accuracy, robustness, and human oversight.	Integrate missing data and missing output checks to human oversight.