

An Automatic Finite-Sample Robustness Metric: When Can Dropping a Little Data Make a Big Difference?

Tamara Broderick, Ryan Giordano*, and Rachael Meager^{†‡}

November 4, 2021

Abstract

We propose a method to assess the sensitivity of econometric analyses to the removal of a small fraction of the data. Manually checking the influence of all possible small subsets is computationally infeasible, so we provide an approximation to find the most influential subset. Our metric, the “Approximate Maximum Influence Perturbation,” is automatically computable for common methods including (but not limited to) OLS, IV, MLE, GMM, and variational Bayes. We provide finite-sample error bounds on approximation performance. At minimal extra cost, we provide an exact finite-sample lower bound on sensitivity. We find that sensitivity is driven by a signal-to-noise ratio in the inference problem, is not reflected in standard errors, does not disappear asymptotically, and is not due to misspecification. While some empirical applications are robust, results of several economics papers can be overturned by removing less than 1% of the sample.

1 Introduction

Ideally, policymakers will use economics research to inform decisions that affect people’s livelihoods, health, and well-being. Yet study samples may differ from the target populations of these decisions in non-random ways, perhaps because of practical challenges in obtaining truly random samples, or because populations generally differ across time and place. When these deviations from the ideal random

*Equal contribution primary author

†Equal contribution primary author

‡Tamara Broderick and Ryan Giordano were supported in part by an Office of Naval Research Early Career Grant, an NSF CAREER Award, and an Army Research Office YIP Award. We thank Avi Feller, Jesse Shapiro, Emily Oster, Michael Kremer, Peter Hull, Tetsuya Kaji, Heather Sarsons, Kirill Borusyak, Tin Danh Nguyen and the authors of all of our applications for their insightful comments and suggestions. All mistakes are our own. Corresponding Author: Rachael Meager, reachable at r.meager@lse.ac.uk.

sampling exercise are small, one might think that the empirical conclusions would still hold in the populations affected by policy. It therefore seems prudent to ask whether a small percentage of a study’s sample—or a handful of data points—has been instrumental in determining its findings. In this paper we provide a finite-sample, automatically-computable metric of how dropping a small amount of data can change empirical conclusions. We show that certain empirical results from high-profile studies in economics can be reversed by removing less than 1% of the sample even when standard errors are small, and we investigate why.

There are several reasons to care about whether empirical conclusions are substantially influenced by small percentages of the finite sample. In practice, even if we can sample from the population of direct interest, small percentages of the data are missing; either surveyors and implementers cannot find these individuals, or they refuse to answer our questions, or their answers get lost or garbled during data processing. As this missingness cannot safely be assumed random, researchers might care whether their substantive conclusions could conceivably be overturned by a missing handful of data points. Similarly, consumers of research who are concerned about potentially non-random errors in sample construction at any stage of the analysis might be interested in this metric as a measure of the exposure of a study’s conclusions to this concern. Conclusions that are highly influenced by a small handful of data points are more exposed to adverse events or errors during data analysis, including p-hacking, even if these errors are unintentional.

Yet even if researchers could construct a perfectly random sample from a given study population, the target population for our policy decisions is always different from the study population, if only because the world may change in the time between the research and the decision. For this reason, social scientists often aspire to uncover generalizable truths about the world and to make policy recommendations that would apply more broadly than to a single study population.

In this paper, we propose to directly measure the extent to which a small fraction of a data sample has influenced the central claims or conclusions of a study. For a particular fraction α (e.g., $\alpha = 0.001$), we propose to find the set of no more than $100\alpha\%$ of all the observations that effects the greatest change in an estimator when those observations are removed from the sample, and to report this change. For example, suppose we were to find a statistically-significant average increase in household consumption after implementing some economic policy intervention. Further suppose that, by dropping 0.1% of the sample (often fewer than 10 data points), we instead find a statistically-significant average *decrease* in consumption. Then it would be challenging to argue that there is strong evidence that this intervention would yield consumption increases in even slightly different populations.

To quantify this sensitivity, one could consider every possible $1 - \alpha$ fraction of the data, and re-run the original analysis on all of these data subsets. But this direct

implementation is computationally prohibitive.¹ We propose a fast approximation that works for common estimators—including Generalized Methods of Moments (GMM), Ordinary Least Squares (OLS), Instrumental Variables (IV), Maximum Likelihood Estimators (MLE), Variational Bayes (VB), and all minimizers of smooth empirical loss. Roughly, we give each data point a weight and apply a Taylor expansion in the weights (Section 2.1 and Section 2.2). Our approximation is fast, automatable, and easy to use, and we provide an R package on GitHub called “zaminfluence.”²

We show that our approximation performs well using theoretical analyses, simulation studies, and applied examples. We demonstrate that the approximation error is low when the percentage of the sample removed is small (Section 3.3). Moreover, for the cost of a single additional data analysis, we can provide an exact lower bound on the worst-case change in an analysis upon removing $100\alpha\%$ of the data (Section 2.2.1). We check that our metric detects combinations of data points that reverse empirical conclusions when removed from real-life datasets (Section 4). For example, in the Oregon Medicaid study (Finkelstein et al., 2012), we can identify a subset containing less than 1% of the original data that controls the sign of the effects of Medicaid on certain health outcomes. In the Mexico microcredit study (Angelucci et al., 2015), we find a single observation, out of 16,500, that controls the sign of the ATE on household profit.

We investigate the source of this sensitivity when it arises, and we show that it is not captured in conventional standard errors. We find that a result’s exposure to the influence of a small fraction of the sample need not reflect a model misspecification problem nor the presence of gross outliers. Sensitivity according to our metric can arise, even if the model is exactly correct and the data set large, if there is a low *signal-to-noise ratio*: that is, if the strength of the claim (signal) is small relative to a quantity that consistently estimates the standard deviation of the limiting distribution of root- N times the quantity of interest (Section 3). For example, in OLS this “noise” is large when we have a high ratio of residual variance to regressor variance (Section 3.1). This noise can be large even when standard errors are small, because it does not disappear as N grows.

We examine several applications from empirical economics papers and find that the sensitivity captured by our metric varies considerably across analyses in practice. In many cases, the sign and significance of certain estimated treatment effects can be reversed by dropping less than 1% of the sample, even when the t-statistics are very large and inference is very precise; see, e.g., the Oregon Medicaid RCT (Finkelstein et al., 2012) in Section 4.1. In Section 4.2, we examine the Progres

¹Indeed, Young (2019) finds it computationally prohibitive to re-run their analysis when leaving out every possible subset of two data points. To illustrate, consider an analysis that takes 1 second to run; checking removal of every 4 data points from a data set of size 400 would take over 33 years. See Section 2 for more detail.

²<https://github.com/rgiordan/zaminfluence>. The name stands for “Z-estimator approximate maximum influence.”

Cash Transfers RCT (Angelucci and De Giorgi, 2009) and show that trimming outliers in the outcome data does not necessarily reduce sensitivity. In Section 4.3 we examine a simple two-parameter linear regression on seven Microcredit RCTs (Meager, 2020) and, in Section 4.4, we examine a Bayesian hierarchical analysis of the same data; these final two analyses show that neither very simple nor relatively complex Bayesian models are immune to sensitivity to dropping small fractions of the data. However, not all analyses we examine are non-robust. Certain results across the applications we examine are robust up to 5% and even 10% removal.

We recommend that researchers use our metric to complement standard errors and other robustness checks. Our goal is not to supplant other analyses, but to provide an additional tool to be incorporated into a broader ecosystem of systematic stability analysis in data science (Yu, 2013). For example, since our approximation is fundamentally local due to the Taylor expansion, practitioners may also consider global sensitivity checks such as those proposed by Leamer (1984, 1985); Sobol (2001); Saltelli (2004), or the breakdown frontiers approach of He et al. (1990); Masten and Poirier (2020). Our method is also no substitute for tailored robustness checks designed by researchers to investigate specific concerns about sensitivity of results to certain structures or assumptions. And practitioners may benefit from robustifying their analysis (Mosteller and Tukey, 1977; Hansen and Sargent, 2008; Chen et al., 2011) even if they pass our check. Our metric is also complementary to classical robustness measures, although we are able to connect our metric to these measures via the influence function. We see that gross error sensitivity is set up for designing estimators and arbitrary adversarial perturbations to the population distribution, whereas our metric is set up for assessing sensitivity to dropping a small subset of the data at hand once an analysis has been performed. We do not recommend researchers discard results that are not robust to removal of a small, highly-influential subset of data, but rather adjust their interpretation of such results as being less generally applicable to somewhat differing populations. We do not yet recommend any specific alterations to common inferential procedures based on our metric, but we believe this direction is promising for future research.

2 A proposed metric

Suppose we observe N data points d_1, \dots, d_N . For instance, in a regression problem, the n -th data point might consist of covariates x_n and response(s) y_n , with $d_n = (x_n, y_n)$. Consider a parameter $\theta \in \mathbb{R}^P$ of interest. Typically we estimate θ via some function $\hat{\theta}$ of our data. The central claim of an empirical economics paper is typically focused on some attribute of θ , such as the sign or significance of a particular effect or quantity. A frequentist analyst might be worried if removing some small fraction α of the data were to

- Change the sign of an effect.

- Change the significance of an effect.
- Generate a significant result of the opposite sign.

To capture these concerns, we define the following quantities:

Definition 1. Let the *Maximum Influence Perturbation* be the largest possible change induced in the quantity of interest by dropping no more than $100\alpha\%$ of the data.

We will often be interested in the set that achieves the Maximum Influence Perturbation, so we call it the *Most Influential Set*.

And we will be interested in the minimum data proportion $\alpha \in [0, 1]$ required to achieve a change of some size Δ in the quantity of interest, so we call that α the *Perturbation-Inducing Proportion*. We report NA if no such α exists.

In general, to compute the Maximum Influence Perturbation for some α , we would need to enumerate every data subset that drops no more than $100\alpha\%$ of the original data. And, for each such subset, we would need to re-run our entire data analysis. If m is the greatest integer smaller than 100α , then the number of such subsets is larger than $\binom{N}{m}$. For $N = 400$ and $m = 4$, $\binom{N}{m} = 1.05 * 10^9$. So computing the Maximum Influence Perturbation in even this simple case requires re-running our data analysis over 1 billion times. If each data analysis took 1 second, computing the Maximum Influence Perturbation would take over 33 years to compute. Indeed, the Maximum Influence Perturbation, Most Influential Set, and Perturbation-Inducing Proportion may all be computationally prohibitive even for relatively small analyses.

2.1 Setup: Notation and Assumptions

To address this computational issue, we propose to use a (fast) approximation instead. We will see, for the cost of one additional data analysis, our approximation can provide a lower bound on the exact Maximum Influence Perturbation. More generally we provide theory and experiments to support the quality of our approximation. We provide open-source code³ and show that our approximation is fully automatable in practice (Section 2.2).

Our approximation is akin to a Taylor expansion, so it will require certain aspects of our estimator to be differentiable. We now summarize the assumptions under which our approximation exists, and note that many common analyses satisfy these assumptions—including, but not limited to, OLS, IV, GMM, MLE, and variational Bayes. Below, in Section 3.3, we will state stricter sufficient conditions that guarantee not only the existence but also the finite-sample accuracy of the approximation.

³<https://github.com/rgiordano/zaminfluence>

Assumption 1. $\hat{\theta}$ is a Z-estimator; that is, $\hat{\theta}$ is the solution to the following estimating equation,⁴ where $G(\cdot, d_n) : \mathbb{R}^P \rightarrow \mathbb{R}^P$ is a twice continuously differentiable function and 0_P is the column vector of P zeros.

$$\sum_{n=1}^N G(\hat{\theta}, d_n) = 0_P. \quad (1)$$

Assumption 2. $\phi : \mathbb{R}^P \rightarrow \mathbb{R}$, which we interpret as a function that takes the full parameter θ and returns the quantity of interest from θ , is continuously differentiable.⁵

For instance, the function that picks out the p -th effect from the vector θ , $\phi(\theta) = \theta_p$, satisfies this assumption.

To form our approximation, we introduce a vector of data weights, $\vec{w} = (w_1, \dots, w_N)$, where w_n is the weight for the n -th data point. We recover the original data set by giving every data point a weight of 1: $\vec{w} = \vec{1} = (1, \dots, 1)$. We can denote a subset of the original data as follows: start with $\vec{w} = \vec{1}$; then, if the data point indexed by n is left out, set $w_n = 0$. We can collect weightings corresponding to all data subsets that drop no more than $100\alpha\%$ of the original data as follows:

$$W_\alpha := \{\vec{w} : \text{No more than } \lfloor \alpha N \rfloor \text{ elements of } \vec{w} \text{ are 0 and the rest are 1}\}. \quad (2)$$

Our main idea will be to form a Taylor expansion of our quantity of interest ϕ as a function of the weights, rather than recalculate ϕ for each data subset (i.e., for each reweighting).

To that end, we first reformulate our setup, now with the weights \vec{w} ; note that we recover the original problem (for the full data) above by setting $\vec{w} = \vec{1}$ in what follows. Let $\hat{\theta}(\vec{w})$ be our parameter estimate at the weighted data set described by \vec{w} . Namely, $\hat{\theta}(\vec{w})$ is the solution to the weighted estimating equation

$$\sum_{n=1}^N w_n G(\hat{\theta}(\vec{w}), d_n) = 0_P. \quad (3)$$

We allow that the quantity of interest ϕ may depend on \vec{w} not only via the estimator θ , so we optionally write $\phi(\theta, \vec{w})$ with $\phi(\cdot, \cdot) : \mathbb{R}^P \times \mathbb{R}^N \rightarrow \mathbb{R}$. Whenever we write $\phi(\cdot)$ as a function of a single argument, we will implicitly mean $\phi(\cdot, \vec{1})$. We require that $\phi(\cdot, \cdot)$ be continuously differentiable in both its arguments. For instance, we can use $\phi(\theta, \vec{w}) = \theta_p$ to pick out the p -th component of θ . Or, to consider questions of statistical significance, we may choose $\phi(\theta, \vec{w}) = \theta_p + 1.96\sigma_p(\theta, \vec{w})$, where $\sigma_p(\theta, \vec{w})$ is an estimate of the standard error depending smoothly on θ and \vec{w} ; this example

⁴Sometimes Eq. 1 is associated with “M-estimators” that optimize a smooth objective function, since such M-estimators typically take the form of a Z-estimator. However, some Z-estimators, such as IV regression, do not optimize any particular empirical objective function, so the notion of Z-estimator is in fact more general.

⁵Below, we will allow for additional dependence in ϕ on data weights.

is our motivation for allowing the more general \vec{w} dependence in $\phi(\theta, \vec{w})$.

With this notation in hand, we can restate our original goal of computing the Most Influential Set as solving

$$\vec{w}^{**} := \arg \max_{\vec{w} \in W_\alpha} \left(\phi(\hat{\theta}(\vec{w}), \vec{w}) - \hat{\phi} \right). \quad (4)$$

Here we focus on positive changes in ϕ since negative changes can be found by reversing the sign of ϕ and using $-\phi$ instead. In particular, the zero indices of \vec{w}^{**} correspond to the Most Influential Set: $S_\alpha := \{n : \vec{w}_n^{**} = 0\}$. And $\Psi_\alpha = \phi(\vec{w}^{**}) - \hat{\phi}$ is the Maximum Influence Perturbation. The Perturbation Inducing Proportion is the smallest α that induces a change of at least size Δ : $\alpha_\Delta^* := \inf\{\alpha : \Psi_\alpha > \Delta\}$.

2.2 A Tractable Approximation

Our approximation, then, centers on a first-order Taylor expansion (and thus linear approximation) in $\vec{w} \mapsto \phi(\hat{\theta}(\vec{w}), \vec{w})$ around $\vec{w} = \vec{1}$. Let $\hat{\phi} := \phi(\hat{\theta}(\vec{1}), \vec{1})$, the quantity of interest at the original dataset. Then:

$$\phi(\hat{\theta}(\vec{w}), \vec{w}) \approx \phi^{\text{lin}}(\vec{w}) := \hat{\phi} + \sum_{n=1}^N (w_n - 1)\psi_n, \text{ with } \psi_n := \left. \frac{\partial \phi(\hat{\theta}(\vec{w}), \vec{w})}{\partial w_n} \right|_{\vec{w}=\vec{1}}. \quad (5)$$

We can in turn approximate the Most Influential Set as follows. Let $\psi_{(n)}$ denote the order statistics of ψ_n , i.e., the ψ_n sorted from most negative to most positive. Let $\mathbb{I}(\cdot)$ denote the indicator function taking value 0 when the argument is false and 1 when true. Then

$$\begin{aligned} \vec{w}^{**} \approx \vec{w}^* &:= \arg \max_{\vec{w} \in W_\alpha} \left(\phi^{\text{lin}}(\vec{w}) - \hat{\phi} \right) = \arg \max_{\vec{w} \in W_\alpha} \sum_{n: w_n=0} (-\psi_n) \Rightarrow \\ \phi^{\text{lin}}(\vec{w}^*) - \hat{\phi} &= - \sum_{n=1}^{\lfloor \alpha N \rfloor} \psi_{(n)} \mathbb{I}(\psi_{(n)} < 0). \end{aligned} \quad (6)$$

To compute \vec{w}^* (analogous to the \vec{w}^{**} that determines the exact Most Influential Set), we compute ψ_n for each n . Then we choose \vec{w}^* to have entries equal to zero at the $\lfloor \alpha N \rfloor$ indices n where ψ_n is most negative (and to have entries equal to one elsewhere). Analogous to the Perturbation Inducing Proportion, we can find the minimum data proportion α required to achieve a change of some size Δ : i.e., such that $\phi^{\text{lin}}(\vec{w}^*) - \hat{\phi} > \Delta$. In particular, we iteratively remove the most negative ψ_n (and the index n) until the Δ change is achieved; if the number of removed points is M , the proportion we report is $\alpha = M/N$. Recall that finding the exact Maximum Influence Perturbation, Most Influential Set, and Perturbation-Inducing Proportion required running a data analysis more than $\binom{M}{\lfloor \alpha N \rfloor}$ times. By contrast, our approximation requires running just the single original data analysis, N additional fast calculations to compute each ψ_n , and finally a sort on the ψ_n values.

We define our approximate quantities, as detailed immediately above, as follows.

Definition 2. The *Approximate Most Influential Set* is the set \hat{S}_α of at most $100\alpha\%$ data indices that, when left out, induce the biggest approximate change $\phi^{\text{lin}}(\vec{w}) - \hat{\phi}$; i.e., it is the set of data indices left out by \vec{w}^* : $\hat{S}_\alpha := \{n : \vec{w}_n^* = 0\}$.

The *Approximate Maximum Influence Perturbation (AMIP)* $\hat{\Psi}_\alpha$ is the approximate change observed at \vec{w}^* : $\hat{\Psi}_\alpha := \phi^{\text{lin}}(\vec{w}^*) - \hat{\phi}$.

The *Approximate Perturbation Inducing Proportion* $\hat{\alpha}_\Delta^*$ is the smallest α needed to cause the approximate change $\phi^{\text{lin}}(\vec{w}) - \hat{\phi}$ to be greater than Δ . That is, $\hat{\alpha}_\Delta^* := \inf\{\alpha : \hat{\Psi}_\alpha > \Delta\}$. We report NA if no $\alpha \in [0, 1]$ can effect this change.

Below, we will sometimes emphasize that the AMIP is a sensitivity and refer to it as the *AMIP sensitivity*. We will say that an analysis is *AMIP-non-robust* if, for a particular α of interest, the AMIP is large enough to change the substantive conclusions of the analysis. Conversely, if the AMIP is not large enough, we say an analysis is *AMIP-robust*. And we generically use the AMIP acronym to describe our methodology even when calculating the Approximate Most Influential Set or Approximate Perturbation Inducing Proportion.

2.2.1 An exact lower bound on the Maximum Influence Perturbation

For any problem where performing estimation a second time is not prohibitively costly, we can re-run our analysis without the data points in the Approximate Most Influential Set and thereby provide a lower bound on the exact Maximum Influence Perturbation.

Formally, let \vec{w}^{**} be the weight vector for the exact Most Influential Set, and let \vec{w}^* be the weight vector for the Approximate Most Influential Set \hat{S}_α . We run the estimation procedure an extra time to recover $\phi(\hat{\theta}(\vec{w}^*), \vec{w}^*)$. Then, by definition,

$$\Psi_\alpha = \phi(\hat{\theta}(\vec{w}^{**}), \vec{w}^{**}) - \hat{\phi} = \max_{\vec{w} \in W_\alpha} \left(\phi(\hat{\theta}(\vec{w}), \vec{w}) - \hat{\phi} \right) \geq \phi(\hat{\theta}(\vec{w}^*), \vec{w}^*) - \hat{\phi}.$$

Since $\phi(\hat{\theta}(\vec{w}^*), \vec{w}^*) - \hat{\phi}$ is a lower bound for Ψ_α , we can use the Approximate Most Influential Set to conclusively demonstrate non-robustness. Of course, this lower bound holds for *any* weight vector and will be most useful if the Approximate Maximum Influence Perturbation is close to the exact Maximum Influence Perturbation. In Section 3.3 below, we establish the accuracy of the approximation for small α under mild regularity conditions.

2.2.2 Computing the influence scores

To finish describing our approximation, it remains to detail how to compute $\psi_n = \frac{\partial \phi(\hat{\theta}(\vec{w}), \vec{w})}{\partial w_n} \Big|_{\vec{w}=\vec{1}}$ from Eq. 5. We will refer to the quantity $\frac{\partial \phi(\hat{\theta}(\vec{w}), \vec{w})}{\partial w_n} \Big|_{\vec{w}}$ as the *influence score* of data point n for ϕ at \vec{w} since, as we show in Section 3.2 below, it is

the *empirical influence function* evaluated at the datapoint d_n . To compute the influence score, we first apply the chain rule:

$$\frac{\partial \phi(\hat{\theta}(\vec{w}), \vec{w})}{\partial w_n} \Big|_{\hat{\theta}(\vec{w}), \vec{w}} = \frac{\partial \phi(\theta, \vec{w})}{\partial \theta^T} \Big|_{\hat{\theta}(\vec{w}), \vec{w}} \frac{\partial \hat{\theta}(\vec{w})}{\partial w_n} \Big|_{\vec{w}} + \frac{\partial \phi(\theta, \vec{w})}{\partial w_n} \Big|_{\hat{\theta}(\vec{w}), \vec{w}}. \quad (7)$$

The derivatives of $\phi(\cdot, \cdot)$ can be calculated using automatic differentiation software such as Python's `autograd` library (Maclaurin et al., 2015; Baydin et al., 2017). And once we have $\hat{\theta}(\vec{1})$ from running the original data analysis, we can evaluate these derivatives at $\vec{w} = \vec{1}$: e.g., $\frac{\partial \phi(\theta, \vec{w})}{\partial \theta^T} \Big|_{\hat{\theta}(\vec{1}), \vec{w}=\vec{1}}$.

The term $\frac{\partial \hat{\theta}(\vec{w})}{\partial w_n} \Big|_{\vec{w}=\vec{1}}$ requires slightly more work since $\hat{\theta}(\vec{w})$ is defined implicitly. We follow standard arguments from the statistics and mathematics literatures (Krantz and Parks, 2012; Hampel, 1974) to show how to calculate it below.

Start by considering the more general setting where $\hat{\theta}(\vec{w})$ is the solution to the equation $\gamma(\hat{\theta}(\vec{w}), \vec{w}) = 0_P$. We assume $\gamma(\cdot, \vec{w})$ is continuously differentiable with full-rank Jacobian matrix; then the derivative $\frac{\partial \hat{\theta}(\vec{w})}{\partial w_n} \Big|_{\vec{w}}$ exists by the implicit function theorem (Krantz and Parks, 2012, Theorem 3.3.1). We can thus use the chain rule and solve for $\frac{\partial \hat{\theta}(\vec{w})}{\partial w_n} \Big|_{\vec{w}}$; in what follows, $0_{P \times N}$ is the $P \times N$ matrix of zeros.

$$0_{P \times N} = \frac{d\gamma(\hat{\theta}(\vec{w}), \vec{w})}{d\vec{w}^T} \Big|_{\vec{w}} = \frac{\partial \gamma(\theta, \vec{w})}{\partial \theta^T} \Big|_{\hat{\theta}(\vec{w}), \vec{w}} \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\vec{w}} + \frac{\partial \gamma(\theta, \vec{w})}{\partial \vec{w}^T} \Big|_{\hat{\theta}(\vec{w}), \vec{w}} \quad (8)$$

$$\Rightarrow \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\vec{w}} = - \left(\frac{\partial \gamma(\theta, \vec{w})}{\partial \theta^T} \Big|_{\hat{\theta}(\vec{w}), \vec{w}} \right)^{-1} \frac{\partial \gamma(\theta, \vec{w})}{\partial \vec{w}^T} \Big|_{\hat{\theta}(\vec{w}), \vec{w}}, \quad (9)$$

where we can take the inverse by our full-rank assumption.

We apply the general setting above to our special case with $\gamma(\theta, \vec{w}) = \sum_{n=1}^N w_n G(\theta, d_n)$ to find

$$\frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\vec{w}} = - \left(\sum_{n=1}^N w_n \frac{\partial G(\theta, d_n)}{\partial \theta^T} \Big|_{\hat{\theta}(\vec{w})} \right)^{-1} \left(G(\hat{\theta}(\vec{w}), d_1), \dots, G(\hat{\theta}(\vec{w}), d_N) \right), \quad (10)$$

which can again be computed with automatic differentiation software.

2.3 Example functions of interest

We end this section with some concrete examples of quantities of interest. Recall from the start of Section 2 that we are often interested in whether we can change the sign or significance of an estimator, or generate a significant result of the opposite sign. Recall that $\phi(\cdot)$ with only one argument is a function of θ , and $\phi(\cdot, \cdot)$ with two arguments is a function of both θ and the weights \vec{w} .

To form our motivating examples, suppose for the remainder of this section we are interested in the p -th component of $\hat{\theta}$, where $\hat{\theta}_p$ is positive and statistically significant. That is, let $\hat{\sigma}_p$ be an estimator of the variance of the limiting distri-

bution of $\sqrt{N}\hat{\phi}$, and let $\hat{\theta}_p - \frac{1.96}{\sqrt{N}}\hat{\sigma}_p$ be the lower end of our confidence interval. So we assume $\hat{\theta}_p > 0$ and $\hat{\theta}_p - \frac{1.96}{\sqrt{N}}\hat{\sigma}_p > 0$. Moreover, we will write $\hat{\sigma}_p(\theta, \vec{w})$ to emphasize that standard errors are typically given as functions of θ and the weights \vec{w} . For example, standard errors based on the observed Fisher information matrix $\frac{1}{N} \sum_{n=1}^N \vec{w}_n \frac{\partial G(\theta, d_n)}{\partial \theta} \Big|_{\hat{\theta}(\vec{w})}$ will, in general, depend on the weights both explicitly and through $\hat{\theta}(\vec{w})$.

To make $\hat{\theta}_p$ change sign, we can take

$$\phi(\theta) = -\theta_p. \quad (\text{Change sign}) \quad (11)$$

We use $-\theta_p$ instead of θ_p since we have defined ϕ as a function that we are trying to increase (cf. Eq. 4 and the discussion after). Increasing $\phi(\hat{\theta})$, for ϕ in Eq. 11, by an amount $\Delta = \hat{\theta}_p$ is equivalent to $\hat{\theta}_p$ changing sign from positive to negative.

To make $\hat{\theta}_p$ statistically non-significant, we wish to take the lower bound of the confidence interval to 0. To that end, we can take

$$\phi(\theta, \vec{w}) = -\left(\theta_p - \frac{1.96}{\sqrt{N}}\hat{\sigma}_p(\theta, \vec{w})\right). \quad (\text{Change significance}) \quad (12)$$

As in the previous case, we choose Eq. 12 with a leading negative sign because we are trying to increase ϕ (cf. Eq. 4). Increasing $\phi(\hat{\theta}, \vec{w})$, for ϕ in Eq. 12, by an amount $\Delta = \hat{\theta}_p - \frac{1.96}{\sqrt{N}}\hat{\sigma}_p$ is equivalent to $\hat{\theta}_p$ becoming statistically insignificant.

Similarly, to change to a significant result of the opposite sign, we can take

$$\phi(\theta, \vec{w}) = -\left(\theta_p + \frac{1.96}{\sqrt{N}}\hat{\sigma}_p(\theta, \vec{w})\right) \quad (\text{Significant sign reversal})$$

and $\Delta = \hat{\theta}_p + \frac{1.96}{\sqrt{N}}\hat{\sigma}_p$, for if the upper end of the confidence interval is negative, then the estimator must be negative and statistically significant.

In each case above, the quantity Δ represents how far we must move ϕ in order to reverse our conclusions. In this sense, Δ is a measure of the amount of “signal” in the original dataset. As we will discuss in Section 3 below, the signal Δ is one of the three key quantities that determine AMIP robustness.

2.4 A real-world OLS regression example

Before continuing, we illustrate our method with an example. Economists often analyze causal relationships using linear regressions estimated via ordinary least squares (OLS), but a researcher rarely believes the conditional mean dependence is truly linear. Rather, researchers use linear regression since it allows transparent and straightforward estimation of an average treatment effect or local average treatment effect. Researchers often invoke the law of large numbers to justify the focus on the sample mean, and invoke the central limit theorem to justify the use of Gaussian confidence intervals when the sample is large. We now discuss an example from

recent economics literature showing how, in practice, the omission of a very small number of data points can have outsize influence on regression parameters in the finite sample even when the full sample is large. We will study AMIP sensitivity for OLS further using simulation and theory in Section 3.1 below.

Consider as an example the set of seven randomized controlled trials of expanding access to microcredit discussed by Meager (2019). For illustrative purposes we single out the study with the largest sample size: Angelucci et al. (2015). This study has approximately 16,500 households. A full treatment of all seven studies is in Sections 4.3 and 4.4 along with tables and figures of the results discussed below.

We consider the headline results on household business profit regressed on an intercept and a binary variable indicating whether a household was allocated to the treatment group or to the control group. Let Y_{ik} denote the profit measured for household i in site k , and let T_{ik} denote their treatment status. We estimate the following model via OLS with the regression formula $Y_{ik} \sim \beta_0 + \beta_1 T_{ik}$. In the notation of Section 2.1, we have $\theta = (\beta_0, \beta_1)^T$, $d_{ik} = (Y_{ik}, T_{ik})$ with $n = (i, k)$, and $G(\theta, d_{ik}) = (Y_{ik} - (\beta_0 + \beta_1 T_{ik}))(1, T_{ik})^T$.

We confirm the main findings of the study in estimating a non-significant average treatment effect (ATE) of -4.55 USD PPP per 2 weeks, with a standard error of 5.88. We are interested in whether we can change the sign of β_1 from negative to positive, so we take $\phi(\theta) = \beta_1$. We compute ψ_n for each data point in the sample, which takes only a fraction of a second in R using our `Zaminfluence` package.

Examining $\vec{\psi}$, we find that one household has $\psi_n = 4.95$; removing that single household should flip the sign if the approximation is accurate. We manually remove the data point and re-run the regression, and indeed find that the ATE is now 0.4 with a standard error of 3.19. Moreover, by removing 15 households we can generate an ATE of 7.03 with a standard error of 2.55: a significant result of the opposite sign.

How is it possible for the absence of a single household to flip the sign of an estimate that was ostensibly based on all the information from a sample of 16,500? It may be tempting to suspect the use of sample means, which are known to be non-robust to gross errors, or to speculate that such excess sensitivity is simply symptomatic of ordinary sampling noise which is captured adequately by standard errors. In Section 3 to follow, we show that such intuition is not correct. On the contrary, AMIP robustness is in fact fundamentally different than both standard errors and classical robustness to gross errors.

3 Underlying theory and interpretation

We now establish the determinants and accuracy of AMIP robustness. We begin by deriving the key quantities of AMIP robustness in the simple case of correctly specified univariate OLS regression (Section 3.1). For this simple case, we show

with theory and simulations that AMIP robustness is not necessarily driven by misspecification, that AMIP non-robustness does not vanish asymptotically, and that AMIP robustness is distinct from standard errors. Next, we formally extend these conclusions to general Z-estimators in Section 3.2. Finally, in Section 3.3, we establish conditions under which the approximation is provably uniformly accurate for small α , both in finite sample and asymptotically.

We will see that a central equation in our understanding of AMIP robustness is its decomposition into three key quantities: the signal, noise, and shape. First, the *signal* Δ is the size of change in our quantity of interest that would reverse our substantive conclusion (see Section 2.3 above). Large values of the signal Δ indicate that large changes are needed to make a different decision. Second, the *noise* $\hat{\sigma}_\psi$ is defined by

$$\hat{\sigma}_\psi^2 := \frac{1}{N} \sum_{n=1}^N (N\psi_n)^2 \quad (13)$$

We call $\hat{\sigma}_\psi$ the noise because $\hat{\sigma}_\psi^2$ is typically a consistent estimator of the variance of the limiting distribution of $\sqrt{N}\phi(\hat{\theta})$, a fact that will follow below from the relationship between AMIP robustness, robust standard error estimators, and the influence function (see Section 3.2.1, paragraph (b) or, more generally, Section 3.2.3, paragraph (d)). Third, the *shape* $\hat{\mathcal{T}}_\alpha$ is defined as

$$\hat{\mathcal{T}}_\alpha := -\frac{1}{N} \sum_{n=1}^{\lfloor \alpha N \rfloor} \frac{N\psi_{(n)}}{\hat{\sigma}_\psi} \mathbb{I}(\psi_{(n)} < 0), \quad (14)$$

where $\psi_{(n)}$ refers to the n -th order statistic of the influence scores, and $\mathbb{I}(\cdot)$ denotes the indicator function taking value 1 when its argument is true and 0 otherwise. The shape $\hat{\mathcal{T}}_\alpha$ depends in a complicated way on the shape of the tail of the distribution of the influence scores, but we show that $0 \leq \hat{\mathcal{T}}_\alpha \leq \sqrt{\alpha(1-\alpha)}$ with probability one, and that $\hat{\mathcal{T}}_\alpha$ converges in probability to a nonzero constant under standard assumptions (see Section 3.2.1, paragraph (c)). Given these three quantities, we will show in Section 3.2.1, paragraph (a) that

$$\text{An analysis is AMIP non-robust} \quad \Leftrightarrow \quad \frac{\Delta}{\hat{\sigma}_\psi} \leq \hat{\mathcal{T}}_\alpha. \quad (15)$$

We refer to the quantity $\Delta/\hat{\sigma}_\psi$ as the *signal-to-noise ratio*. For a given α , Eq. 15 suggests that it is the signal-to-noise ratio that primarily determines AMIP robustness. Additionally, this decomposition allows us to succinctly compare AMIP robustness to standard errors and gross-error robustness, as well as to analyze the large- N behavior of AMIP robustness.

This section will use the following notation. Let the symbol \xrightarrow{p} denote convergence in probability, and \rightsquigarrow denote convergence in distribution, both as $N \rightarrow \infty$. Let $\|\cdot\|_{op}$ denote the operator norm of a matrix.

3.1 Theory and interpretation for Ordinary Least Squares

We begin by focusing on the simple case of correctly-specified univariate linear regression, both to provide intuition and motivate the more general results that follow.

3.1.1 Problem setup for Ordinary Least Squares example

(a) Model. Let $X = (x_1, \dots, x_N)^T$ denote a vector of N continuous mean-zero regressors, drawn IID from a distribution with finite variance σ_x^2 . Let $\varepsilon = (\varepsilon_1, \dots, \varepsilon_N)$ be a vector of IID draws from a $\mathcal{N}(0, \sigma_\varepsilon^2)$ distribution, where we will assume σ_ε is known. For some unknown $\theta_0 \in \mathbb{R}$, let $y_n = \theta_0 x_n + \varepsilon_n$, so that the vector $Y = (y_1, \dots, y_N)$ given X is drawn from a correctly specified regression model with true coefficient θ_0 .

(b) Weighted estimating equation. The OLS estimator $\hat{\theta}$ is traditionally found by maximizing the (log) likelihood: $\log p(y_n | \theta, x_n) = -\frac{1}{2\sigma_\varepsilon^2} (y_n - \theta x_n)^2 + C$, where C does not depend on θ . In particular, setting the derivative of the log likelihood to zero yields the estimating equation $G(\theta, d_n) = -\frac{1}{\sigma_\varepsilon^2} (y_n - \theta x_n) x_n = 0$. That is, $\hat{\theta}$ is a Z-estimator with this choice of G (see Eq. 1). Typical Z-estimators do not have closed-form solutions. But in this case, the solution to the estimating equation returns the usual OLS estimate. A similar derivation returns the solution to the weighted estimating equation given in Eq. 3: $\hat{\theta}(\vec{w}) = \left(\frac{1}{N} \sum_{n=1}^N \vec{w}_n x_n^2 \right)^{-1} \frac{1}{N} \sum_{n=1}^N \vec{w}_n y_n x_n$.

(c) Quantity of interest. Suppose we are interested in the sign of θ_0 . Without loss of generality, we assume $\hat{\theta} < 0$. Then our quantity of interest is $\phi(\theta) = \theta$.

(d) Signal and noise. For our quantity of interest, the signal is $\Delta = |\hat{\theta}|$ since, if we can increase $\hat{\theta}$ by an amount $|\hat{\theta}|$, its sign will change. To compute the noise, we compute the influence scores. Directly differentiating the explicit formula for $\hat{\theta}$ gives, as it must, the same value for ψ_n as the implicit function theorem result of Eq. 9. Letting $\hat{\varepsilon}_n := y_n - \hat{\theta} x_n$ and $S_X := \frac{1}{N} \sum_{n=1}^N x_n^2$, we see, either by direct differentiation or by Eq. 9, that $\psi_n = N^{-1} S_X^{-1} x_n \hat{\varepsilon}_n$. For intuition about the noise $\hat{\sigma}_\psi$, we observe its asymptotic behavior. Standard results for OLS give:

$$\hat{\sigma}_\psi^2 = \frac{1}{N} \sum_{n=1}^N (N \psi_n)^2 = S_X^{-2} \frac{1}{N} \sum_{n=1}^N x_n^2 \hat{\varepsilon}_n^2 \xrightarrow{p} \frac{\sigma_\varepsilon^2}{\sigma_x^2}. \quad (16)$$

Note that the noise includes a contribution from both the residual and regressor variance—we describe $\hat{\sigma}_\psi$ as the “noise” because it estimates the variability of $\sqrt{N} \hat{\theta}$, not of the residuals (see Section 3.1.2, paragraph (e) below). Finally, we emphasize that, although we will be using asymptotics to provide intuition, by “noise” we will always mean the finite-sample quantity $\hat{\sigma}_\psi$, not its asymptotic limit.

3.1.2 What determines AMIP robustness for Ordinary Least Squares?

Now that we have translated OLS into our framework, we can analyze the AMIP for OLS. To that end, we use both theory and a simulation study. We outline the simulation study before describing our main conclusions. For $N = 5,000$ data points, and for a range of σ_x and σ_ε , we drew normal regressors $x_n \sim \mathcal{N}(0, \sigma_x^2)$ and residuals $\varepsilon_n \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. For $\theta_0 = 0.5$, we set $y_n = \theta_0 x_n + \varepsilon_n$. We computed the OLS estimator $\hat{\theta} = \sum_{n=1}^N y_n x_n / \sum_{n=1}^N x_n^2$.

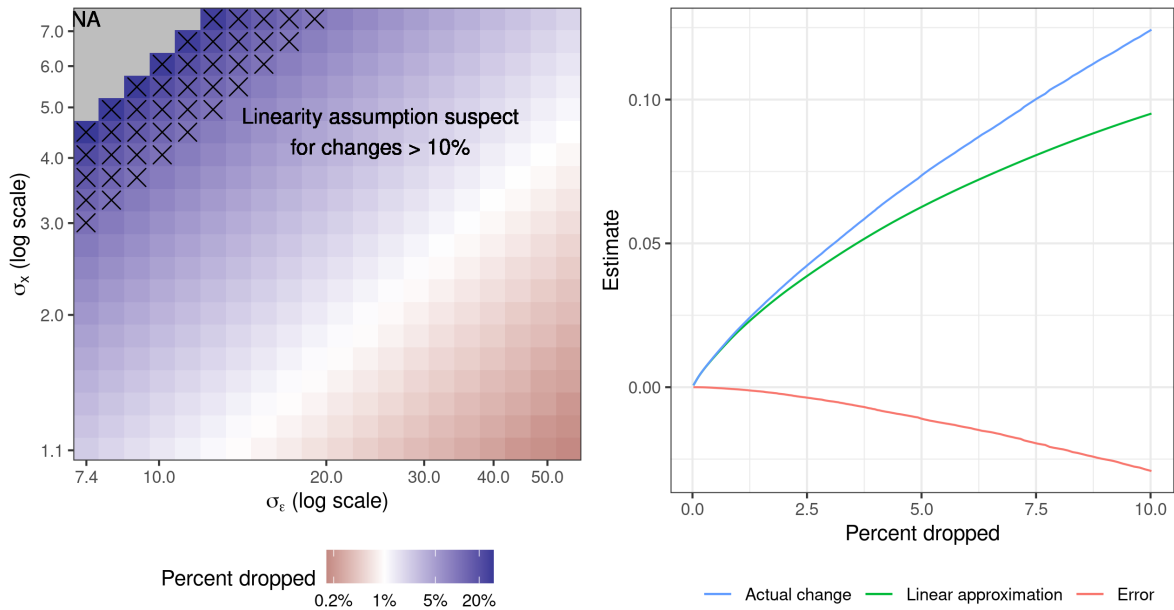


Figure 1: Simulation results for univariate linear regression with $N = 5,000$ observations. **Left panel:** The approximate perturbation inducing proportion at differing values of σ_x and σ_ε . Red colors indicate datasets whose sign can be predicted to change when dropping less than 1% of datapoints. The grey areas indicate $\hat{\Psi}_\alpha = \text{NA}$, a failure of the linear approximation to locate any way to change the sign. **Right panel:** The actual change, linear approximation to the change, and approximation error for $\sigma_x = 2$ and $\sigma_\varepsilon = 1$.

(a) Signal-to-noise ratio drives AMIP robustness. From our discussion at the start of Section 3, we expect that the signal-to-noise ratio drives whether an analysis is AMIP-robust or not. In our simulation, N is large and we keep θ_0 fixed, so we expect that the signal does not change substantially over the simulation. Therefore, signal-to-noise is controlled by the noise. Following the asymptotic argument above, we approximate the noise as $\sigma_\varepsilon/\sigma_x$. In the left panel of Figure 1, we vary σ_ε and σ_x and plot the resulting Approximate Perturbation Inducing Proportion α^* to change the sign of $\hat{\theta}$. As expected, we see that the simulations with the largest approximate noise $\sigma_\varepsilon/\sigma_x$ are the least robust, in the sense that one can reverse the sign of $\hat{\theta}$ by dropping a very small proportion of points.

(b) Influential data points have both a large residual and large regressor. Let $(\hat{\varepsilon}x)_{(n)}$ denote the products $\hat{\varepsilon}_n x_n$, sorted from most negative to most

positive, so that the sorted influence scores are $\psi_{(n)} = N^{-1}S_X^{-1}(\hat{\varepsilon}x)_{(n)}$. From this formula, we observe that influential datapoints have both a large residual and a large regressor (relative to the regressor variance).⁶ A typical influence score goes to zero at rate N^{-1} , though extreme values such as $\max_n |\psi_n|$ may obey a different rate. However, since $\frac{1}{N} \sum_{n=1}^N x_n^2$ and $\frac{1}{N} \sum_{n=1}^N \varepsilon_n^2$ are finite with high probability, even $\max_n |\psi_n|$ does not diverge in this case.⁷

(c) AMIP sensitivity does not vanish as $N \rightarrow \infty$. Standard results for OLS give that $S_X \xrightarrow{p} \sigma_x^2$ and $\hat{\varepsilon}_n - \varepsilon_n \xrightarrow{p} 0$. So $N\psi_n - \sigma_x^{-2}x_n\varepsilon_n \xrightarrow{p} 0$. Consequently, the empirical distribution of $N\psi_n$ converges to a non-degenerate distribution with finite variance. Let q_α denote the α -th quantile of the distribution of the random variable $\sigma_x^{-2}x_1\varepsilon_1$. Since x_n and ε_n are independent, and about half of the ε_n will be negative, we expect about half of the influence scores to be negative. So for $\alpha \ll 1/2$, with high probability at least αN influence scores are negative. Then, by Eq. 6 and Slutsky’s theorem, we have

$$\phi^{\text{lin}}(\vec{w}^*) - \hat{\phi} = - \sum_{n=1}^{\alpha N} \psi_{(n)} = - \frac{1}{N} \sum_{n=1}^{\alpha N} S_X^{-1}(\hat{\varepsilon}x)_{(n)} \xrightarrow{p} \mathbb{E} \left[- \frac{x_1\varepsilon_1}{\sigma_x^2} \mathbb{I} \left(\frac{x_1\varepsilon_1}{\sigma_x^2} \leq q_\alpha \right) \right].$$

The right hand side of the preceding display is strictly positive for finite α . So, for fixed α , we expect that AMIP sensitivity does not vanish as $N \rightarrow \infty$.⁸

(d) AMIP non-robustness is not due only to misspecification. Our simulations are well specified. Yet we see from Figure 1 that different cases can still be robust or non-robust under various robustness cut-offs—according to their differing signal-to-noise ratios.

Asymptotically as $N \rightarrow \infty$, even in a well-specified model, we in fact expect AMIP non-robustness at any α for a sufficiently small $|\theta_0|$. The limiting value of the AMIP sensitivity does not depend on θ_0 . Thus, as $N \rightarrow \infty$, our quantity of interest (for changing the sign of the estimator) will be AMIP non-robust with high probability if and only if $|\theta_0| < \mathbb{E} \left[- \frac{x_1\varepsilon_1}{\sigma_x^2} \mathbb{I} \left(\frac{x_1\varepsilon_1}{\sigma_x^2} \leq q_\alpha \right) \right]$. If we are interested in the sign of θ_0 , and $|\theta_0|$ is small relative to the tail means of $\sigma_X^{-2}x_1\varepsilon_1$, then the problem will be AMIP non-robust with probability approaching one, no matter how large N is—despite the fact that the model is correctly specified and there are no abnormalities in the data.

(e) Though both are scaled by the noise, standard errors are different from—and typically smaller than—AMIP sensitivity. In what may seem at first like a remarkable coincidence, the variance of the limiting distribution of $N\psi_n$ (which determines AMIP sensitivity—see Eq. 5) is the same as the variance of

⁶Indeed, if we had taken $\phi(\theta) = \hat{\theta}x_n = \hat{y}_n$, then the n -th influence score would have been $S_X^{-1}x_n^2\hat{\varepsilon}_n$, which is precisely the leverage score times the residual. This expression formalizes the conceptual link made by Chatterjee and Hadi (1986) between influence, leverage, and large values of $\hat{\varepsilon}_n$.

⁷The finiteness follows from the inequality $\frac{1}{N} \max_n x_n^2 \leq \frac{1}{N} \sum_{n=1}^N x_n^2 \xrightarrow{p} \sigma_x^2$, with an analogous inequality for ε_n . However, since we know ε_n is Gaussian, we actually have a stronger result in this case: $\max_{n \in \{1, \dots, N\}} |\varepsilon_n|$ grows at rate $\sqrt{\log(2N)}$ (Rigollet, 2015, Theorem 1.14).

⁸As desired, though, the expectation does go to zero as $\alpha \rightarrow 0$ since $\mathbb{E} [|x_1\varepsilon_1|] < \infty$.

the limiting distribution of our quantity of interest $\sqrt{N}(\hat{\theta} - \theta_0)$ (which determines classical standard errors). The two distributions are not the same—the limiting distribution of $N\psi_n$ is not, in general, normal—but they have the same scale. In particular, compare the noise limit in Eq. 16 with the following limit, which follows by standard results for OLS.

$$\sqrt{N}(\hat{\theta} - \theta_0) \rightsquigarrow \mathcal{N}\left(0, \frac{\sigma_\varepsilon^2}{\sigma_x^2}\right).$$

As we discuss below in Section 3.2.1, paragraph (b) and Section 3.2, paragraph (d), this equality is no coincidence, but a general (and well-known) relationship between influence scores and the limiting distributions of quantities of interest.

For large N , use of standard errors will admit the hypothesis that θ_0 might be 0 whenever $|\theta_0| < \frac{1.96}{\sqrt{N}} \frac{\sigma_\varepsilon}{\sigma_x}$. Thus, for every $\theta_0 \neq 0$, using standard errors always rejects $\theta_0 = 0$ for sufficiently large N . By contrast, as we saw above, using the AMIP will admit a change large enough to move $\hat{\theta}$ to 0 whenever

$$|\theta_0| \leq \left(\mathbb{E} \left[-\frac{x_1}{\sigma_x} \frac{\varepsilon_1}{\sigma_\varepsilon} \mathbb{I} \left(\frac{x_1}{\sigma_x} \frac{\varepsilon_1}{\sigma_\varepsilon} \leq \frac{\sigma_x}{\sigma_\varepsilon} q_\alpha \right) \right] \right) \frac{\sigma_\varepsilon}{\sigma_x} \neq \frac{1.96}{\sqrt{N}} \frac{\sigma_\varepsilon}{\sigma_x}.$$

Thus, we see that both the AMIP sensitivity and standard errors admit larger possible values for $\hat{\theta}$ when the limiting value $|\theta_0|/(\sigma_\varepsilon/\sigma_x)$ of the signal-to-noise ratio is large. But AMIP sensitivity is determined by the tail mean of the standardized influence scores, and standard errors are determined by a quantity that goes to zero as $N \rightarrow \infty$. Thus AMIP sensitivity is distinct from, and typically larger than, standard errors. The tail behavior of the unit-variance random variable $\frac{x_1}{\sigma_x} \frac{\varepsilon_1}{\sigma_\varepsilon}$ is exactly the shape we introduced at the start of Section 3. The shape captures the scale-independent shape of the tails of the distribution of the influence scores; see Section 3.2.1, paragraph (c) below for a detailed and general analysis.

(f) Our approximation is accurate for small α . The expression for $\hat{\theta}(\vec{w})$ depends on two terms, $\left(\frac{1}{N} \sum_{n=1}^N \vec{w}_n x_n^2\right)^{-1}$ and $\frac{1}{N} \sum_{n=1}^N \vec{w}_n y_n x_n$, both of which are uniformly smooth functions of \vec{w}/N with high probability for sufficiently small $\|\vec{w} - \vec{1}\|_2/N$. As a consequence of smoothness, we expect a linear approximation formed at $\vec{w} = \vec{1}$ to be accurate when $\|\vec{w} - \vec{1}\|_2/N$ is small. And when \vec{w} contains no more than $\lfloor \alpha N \rfloor$ zeros and the rest ones, we have that $\|\vec{w} - \vec{1}\|_2/N \leq \alpha$, so we expect a linear approximation to be accurate when α is small. We make this intuition precise and general in Section 3.3 below.

We check the accuracy of the approximation empirically in Figure 1. For the right hand plot in Figure 1, we fixed $\sigma_\varepsilon = 1$ and $\sigma_x = 2$. We computed the Approximate Most Influential Set for a range of left-out proportions α from 0 to 10%. For each α , we computed the linear approximation, re-ran the regression to compute the actual change, and computed the error of the linear approximation as the difference of the two. The right panel of Figure 1 shows how the relative error of

the approximation vanishes for small α , and that, qualitatively, the approximation is very good for removal proportions less than 2.5%.

3.2 Theory and interpretation for general Z-estimators

We next show that the conclusions of Section 3.1 hold not just for OLS but in considerable generality for Z-estimators. In the present section, we will establish more generally that AMIP sensitivity is not a product of misspecification, does not vanish as N goes to infinity, and is distinct from standard errors. To that end, in Section 3.2.1 we first formally decompose the AMIP into the shape and noise terms defined at the beginning of Section 3, and we establish that the shape is roughly constant across distributions. Then, in Section 3.2.2, we use this decomposition to revisit our OLS conclusions about AMIP sensitivity but now more broadly. Finally, in Section 3.2.3, we connect the AMIP to the influence function, showing how AMIP robustness is different from gross error robustness.

3.2.1 The decomposition of the AMIP

(a) The AMIP is the noise times the shape. Let $\psi_{(1)}, \dots, \psi_{(N)}$ denote the order statistics of the influence scores. Recall that the Approximate Maximum Influence Perturbation is given by the negative of the sum of the $[\alpha N]$ largest influence scores. So we can write

$$\hat{\Psi}_\alpha = \phi^{\text{lin}}(\bar{w}^*) - \hat{\phi} = - \sum_{n=1}^{[\alpha N]} \psi_{(n)} \mathbb{I}(\psi_{(n)} < 0) = \hat{\sigma}_\psi \hat{\mathcal{J}}_\alpha. \quad (17)$$

The first equality follows from the definition of the AMIP $\hat{\Psi}_\alpha$ (Definition 2). The second equality follows from Eq. 6. The third equality follows from the definitions of noise $\hat{\sigma}_\psi$ and shape $\hat{\mathcal{J}}_\alpha$ at the start of Section 3.

(b) The noise is an estimator of the standard deviation of the limiting distribution of the quantity of interest (Z-estimator version). In the case of Z-estimators, we can show by direct computation that $\hat{\sigma}_\psi^2$ is the estimator of the variance of the limiting distribution of $\sqrt{N}\phi(\hat{\theta})$ given by the delta method and the “sandwich” or “robust” covariance estimator (Huber, 1967; Stefanski and Boos, 2002). To see this, observe first that $\frac{1}{N} \sum_{n=1}^N \left. \frac{d\hat{\theta}(\bar{w})}{d\bar{w}_n} \right|_{\bar{1}} \left(\left. \frac{d\hat{\theta}(\bar{w})}{d\bar{w}_n} \right|_{\bar{1}} \right)^T$, as given by Eq. 10, is precisely the sandwich covariance estimator for the covariance of the limiting distribution of $\sqrt{N}\hat{\theta}$. In turn, the sample variance of the linear approximation given in Eq. 7, given by $\hat{\sigma}_\psi^2$, is then the delta method variance estimator for $\sqrt{N}\hat{\phi}$. Note that we came to the same conclusion in the special case of OLS in Section 3.1.2, paragraph (e) above.

It follows that we can use $\hat{\sigma}_\psi$ to form consistent credible intervals for ϕ , a fact that will be useful below when comparing AMIP robustness to standard errors.

Specifically, if $\hat{\sigma}_\psi \xrightarrow{p} \sigma_\psi$ and $\hat{\theta} \xrightarrow{p} \theta_\infty$, then

$$\sqrt{N}(\phi(\hat{\theta}) - \phi(\theta_\infty)) \rightsquigarrow \mathcal{N}(0, \sigma_\psi^2). \quad (18)$$

As we discuss in Section 3.2.3, paragraph (d) below, this relationship between asymptotic variance and the influence scores is in fact a consequence of a general relationship between influence functions and distributional limits.

(c) The shape depends primarily on α , not on the model specification.

More precisely, we next show that the shape $\hat{\mathcal{T}}_\alpha$ satisfies the following properties. (1) With probability one, $0 \leq \hat{\mathcal{T}}_\alpha \leq \sqrt{\alpha(1-\alpha)}$. (2) Typically, $\hat{\mathcal{T}}_\alpha$ converges in probability to a nonzero constant as $N \rightarrow \infty$. (3) $\hat{\mathcal{T}}_\alpha$ is largest when the influence scores of the left-out points are all equal. Conversely, heavy tails in the distribution of ψ_n result in smaller values of $\hat{\mathcal{T}}_\alpha$. (4) Empirically, $\hat{\mathcal{T}}_\alpha$ varies relatively little among common sampling distributions.

To prove the lower bound in (1), we observe that the indicator $\mathbb{I}(\psi_{(n)} < 0)$ accounts for the fact that the adversarial weight would leave out fewer points rather than drop a point with positive $\psi_{(n)}$. Because of this, $\hat{\mathcal{T}}_\alpha \geq 0$. We show the upper bound of (1) as part of the extremization argument for (3) below.

To prove (2), notice that $\hat{\mathcal{T}}_\alpha$ is a sum of $\lfloor \alpha N \rfloor$ positive terms, divided by N . In general, then, we expect $\hat{\mathcal{T}}_\alpha$ to converge to a nonzero constant for fixed α as long as the distribution of $N\psi_n$ converges marginally in distribution to a non-degenerate random variable. And indeed, by Eqs. 7 and 10, we expect such convergence from Slutsky's theorem as long as $\hat{\theta}$ and $\frac{1}{N} \sum_{n=1}^N \left. \frac{\partial G(\hat{\theta}, d_n)}{\partial \theta} \right|_{\hat{\theta}}$ converge in probability to constants, since $N\psi_n$ is proportional to $G(\hat{\theta}, d_n)$, which itself has a non-degenerate limiting distribution.

We next show (3), that $\hat{\mathcal{T}}_\alpha$ takes its largest possible value when all the influence scores $\psi_{(1)}, \dots, \psi_{(\alpha N)}$ take the same negative value. To that end, take αN to be an integer for simplicity. By the definition of $\hat{\sigma}_\psi$ (Eq. 13), $\frac{1}{N} \sum_{n=1}^N \left(\frac{N\psi_{(n)}}{\hat{\sigma}_\psi} \right)^2 = 1$, and by properties of the influence function detailed below, $\sum_{n=1}^N \psi_n = 0$ (Section 3.2.3, paragraph (c)). So $\hat{\mathcal{T}}_\alpha$ is a tail average of scalars with zero sample mean and unit sample variance. Therefore, it is equivalent to consider scalars z_1, \dots, z_N with $\frac{1}{N} \sum_{n=1}^N z_n = 0$ and $\frac{1}{N} \sum_{n=1}^N z_n^2 = 1$ and to ask how to maximize the average $-\frac{1}{\alpha N} \sum_{n=1}^{\alpha N} z_{(n)}$.

To perform this maximization we divide datapoints into a set D of dropped indices, and set K of kept indices. To be precise, $D := \{n : z_{(n)} \leq z_{(\alpha N)}\}$ and $K := \{1, \dots, N\} \setminus D$. We write the sample means and variances within the sets respectively as $\mu_D := \frac{1}{\alpha N} \sum_{n \in D} z_n$ and $v_D := \frac{1}{\alpha N} \sum_{n \in D} (z_n - \mu_D)^2$, with analogous expressions for μ_K and v_K . In this notation, our goal is to extremize μ_D , the mean in the dropped set. The constraints on the distribution can then be written as $\frac{1}{N} \sum_{n=1}^N z_n = 0 \Rightarrow \alpha \mu_D + (1-\alpha)\mu_K = 0$, and $\frac{1}{N} \sum_{n=1}^N z_n^2 = 1 \Rightarrow \alpha(v_D + \mu_D^2) + (1-\alpha)(v_K + \mu_K^2) = 1$. Given these constraints, we extremize μ_D by setting $v_K = v_D = 0$,

in which case we achieve $\mu_D = -\sqrt{(1-\alpha)/\alpha}$. Identifying $N\psi_n/\hat{\sigma}_\psi$ with z_n , and $\hat{\mathcal{T}}_\alpha$ with $\alpha\mu_D$, we see that the worst-case value of $\hat{\mathcal{T}}_\alpha$ occurs when all the influence scores $\psi_{(1)}, \dots, \psi_{(\alpha N)}$ take the same negative value. This observation completes our argument for (3). It also follows from this argument that $\hat{\mathcal{T}}_\alpha \leq \sqrt{\alpha(1-\alpha)}$ with probability one, a bound that is achieved in the worst-case. This observation supplies the upper bound in (1).

To establish point (4), we fix a representative α , simulate a large number of IID draws \tilde{z}_n from some common distributions, standardize to get $z_n := \frac{\tilde{z}_n - \bar{\tilde{z}}}{\sqrt{\frac{1}{N} \sum_{n=1}^N (\tilde{z}_n - \bar{\tilde{z}})^2}}$, and compute the shape $\hat{\mathcal{T}}_\alpha = -\frac{1}{N} \sum_{n=1}^{\lfloor \alpha N \rfloor} z_{(n)}$. We find that, across common distributions, $\hat{\mathcal{T}}_\alpha$ varies relatively little. For example, for $\alpha = 0.01$, a Normal distribution gives $\hat{\mathcal{T}}_\alpha = 0.0266$, a Cauchy distribution gives $\hat{\mathcal{T}}_\alpha = 0.0022$. As expected based on the reasoning of the previous paragraph, the heavy-tailed Cauchy distribution has a smaller shape than the Normal distribution. The worst-case distribution, for which all left-out z_n are equal, gives $\hat{\mathcal{T}}_\alpha = 0.0995 \approx \sqrt{\alpha(1-\alpha)}$ as expected.

3.2.2 What determines AMIP robustness?

We now use the decomposition of the AMIP into noise and shape, and the relative stability of the shape, to derive a number of general properties of AMIP robustness.

(a) Signal-to-noise ratio drives AMIP robustness. We argued above that we do not expect $\hat{\mathcal{T}}_\alpha$ to vary radically from one problem to another. By contrast, the noise $\hat{\sigma}_\psi$ can, in principle, be any positive number. We conclude then, that the signal-to-noise ratio, rather than the shape, principally determines AMIP robustness.

This relationship also suggests what might be done if the analysis is deemed AMIP non-robust. Since, as we showed in Section 3.2.1, paragraph (b), $\hat{\sigma}_\psi$ is thus the same quantity that enters standard error computations, analysts are typically attentive to choosing estimators with $\hat{\sigma}_\psi$ as small as possible while still guaranteeing desirable properties like consistency. Meanwhile, the signal Δ is determined by the question being asked and the true state of nature as estimated by $\hat{\theta}$. In light of these observations, consider a case where $\Delta/\hat{\sigma}_\psi$ is too small to ensure AMIP robustness. Then it seems necessary for the investigator to ask a different question, or investigate different data, to find an AMIP robust analysis.

(b) AMIP sensitivity does not vanish as $N \rightarrow \infty$. Both $\hat{\sigma}_\psi$ and $\hat{\mathcal{T}}_\alpha$ converge to nonzero constants. So $\hat{\sigma}_\psi \hat{\mathcal{T}}_\alpha$, the estimated amount by which you can change an estimator, does not go to zero, either. If the signal Δ is less than the probability limit of $\hat{\sigma}_\psi \hat{\mathcal{T}}_\alpha$, then the problem will be AMIP non-robust no matter how large N grows. As we discuss below, this behavior contrasts sharply with the behavior of standard errors.

(c) AMIP non-robustness is not due only to misspecification. Consider

a correctly-specified problem with no aberrant data points. As we discussed above in Section 3.2.1, paragraph (b), the noise will still have some non-zero probability limit. We showed in Section 3.2.1, paragraph (c) that the shape will have a non-zero probability limit. And the quantity of interest $\phi(\hat{\theta})$ can generally be expected to have a non-zero probability limit. So by the decomposition of Eq. 15, if the user is interested in a question whose signal is small enough, their problem will be AMIP non-robust, despite correct specification.

(d) Though both are scaled by noise, standard errors are different from—and typically smaller than—AMIP sensitivity. Recall that classical standard errors based on limiting normal approximations also depend on $\hat{\sigma}_\psi$, in that we typically report a confidence interval for ϕ of the form $\phi \in \left(\phi(\hat{\theta}, \vec{1}) \pm q_N \frac{\hat{\sigma}_\psi}{\sqrt{N}}\right)$, where q_N is some quantile of the normal distribution, e.g. the 0.975-th quantile $q_N \approx 1.96$. In this sense, using standard errors allow that ϕ may be as large as $\phi + \Delta$ whenever $\Delta/\hat{\sigma}_\psi \leq \frac{1.96}{\sqrt{N}}$. By contrast, AMIP robustness allows that ϕ may be as large as $\phi + \Delta$ when $\Delta/\hat{\sigma}_\psi \leq \hat{\mathcal{J}}_\alpha$. Since $\hat{\mathcal{J}}_\alpha \neq \frac{1.96}{\sqrt{N}}$ in general, these two approaches will yield different conclusions. Indeed, typically $\hat{\mathcal{J}}_\alpha$ converges to a non-zero constant as $N \rightarrow 0$, while $\frac{1.96}{\sqrt{N}}$ converges to zero.

(e) Statistical non-significance is always AMIP-non-robust as $N \rightarrow \infty$. This observation follows as a corollary of the discussion above. In particular, we might conclude statistical non-significance if $|\phi(\hat{\theta}, \vec{1})| \leq \frac{1.96\hat{\sigma}_\psi}{\sqrt{N}}$. To produce a statistically significant result, and so undermine the conclusion, it suffices to move $\phi(\hat{\theta}, \vec{1})$ by more than $\frac{1.96\hat{\sigma}_\psi}{\sqrt{N}}$. Take any α . As we have seen above, we can produce a change of $\hat{\sigma}_\psi \hat{\mathcal{J}}_\alpha$, which is greater than $\frac{1.96\hat{\sigma}_\psi}{\sqrt{N}}$ whenever $\hat{\mathcal{J}}_\alpha > 1.96/\sqrt{N}$. Thus, for any fixed α , there always exists a sufficiently large N such that statistical non-significance can be undermined by dropping at most α proportion of the data. By contrast, statistical significance can be robust if $\phi(\hat{\theta}, \vec{1})$ converges to a value sufficiently far from 0.

3.2.3 The influence function

We next review the influence function and give its particular form for Z -estimators. We establish a relationship between the influence scores and the empirical influence function. We use these connections to further justify the relationship between the noise and the limiting distribution of $\sqrt{N}\hat{\phi}$. Finally, we use the influence function to contrast AMIP robustness with gross error robustness and establish that outliers primarily affect AMIP robustness via the noise, rather than via the shape.

(a) Writing a statistic as a functional of the empirical distribution. Before defining the influence function, we set up some useful notation. Suppose we observe IID data, d_1, \dots, d_N . Each point is drawn from a data distribution $F_\infty(\cdot) = p(d_1 \leq \cdot)$, where the inequality may be multi-dimensional. For a generic distribution F , let T represent a functional of the distribution: $T(F)$. One example is the sample mean; for a generic distribution F , let $T_{mean}(F) = \int \tilde{d}dF(\tilde{d})$. Then

$T_{mean}(F_\infty) = \mathbb{E}[d_1]$ is the population mean. If we let \hat{F}_N denote the empirical distribution function $\hat{F}_N(\cdot) = \frac{1}{N} \sum_{n=1}^N \mathbb{I}(\cdot \leq d_n)$, then $T_{mean}(\hat{F}_N) = \frac{1}{N} \sum_{n=1}^N d_n$ is the sample mean.

Now consider Z-estimators. Define $T_Z(F)$ to be a quantity satisfying

$$\int G(T_Z(F), \tilde{d}) dF(\tilde{d}) = 0. \quad (19)$$

See, e.g., Hampel (1986, Section 4.2c, Def. 5). If we plug in \hat{F}_N for F in Eq. 19 (and multiply both sides by N), we recover the Z-estimator estimating equation from Eq. 1, with solution $\hat{\theta} = T_Z(\hat{F}_N)$. Similarly, let \hat{F}_w to be the distribution function putting weight $N^{-1}w_n$ at data point d_n . Plugging in \hat{F}_w for F in Eq. 19 yields the estimating equation in Eq. 3, for weighted Z-estimators, with solution $\hat{\theta}(\vec{w}) = T_Z(\hat{F}_w)$. Finally, we can define a new functional $T_\phi(F)$ by applying the smooth function ϕ , which picks out our quantity of interest, to $T_Z(F)$: $T_\phi(F) = \phi(T_Z(F), \vec{1})$.⁹

(b) The influence function. The influence function $\text{IF}(d; T, F)$ measures the effect on a statistic T of adding an infinitesimal amount of mass at point d to some base or reference data distribution F (Reeds, 1976; Hampel, 1986). Let δ_d be the probability measure with an atom of size 1 at d . Then

$$\text{IF}(d; T, F) := \lim_{\epsilon \searrow 0} \frac{T(\epsilon\delta_d + (1-\epsilon)F) - T(F)}{\epsilon}. \quad (20)$$

The influence function is defined in terms of an ordinary univariate derivative, and can be computed (as a function of d and F) using standard univariate calculus. In particular, our quantity of interest has the following influence function:

$$\text{IF}(d; T_\phi, F) = - \left. \frac{\partial \phi(\theta, \vec{1})}{\partial \theta^T} \right|_{\hat{\theta}(F)} \left(\int \left. \frac{\partial G(\theta, \tilde{d})}{\partial \theta^T} \right|_{\hat{\theta}(F)} dF(\tilde{d}) \right)^{-1} G(\hat{\theta}(F), d). \quad (21)$$

By comparing Eq. 21 with the definition of ψ_n in Eqs. 7 and 10, we can see that, formally,¹⁰,

$$N\psi_n = \text{IF}(d_n; T_\phi, \hat{F}_N). \quad (22)$$

Eq. 22 is not a coincidence. To see this, note that the set of distributions that can be expressed as weighted empirical distributions (\hat{F}_w above) is precisely the subspace of possible distribution functions concentrated on the observed data. So the derivative $N\psi_n = N\partial\phi(\hat{\theta}(\vec{w}), \vec{1})/\partial\vec{w}_n$ (Eq. 5) is simply a path derivative representation of the functional derivative $\text{IF}(d_n; T_\phi, \hat{F}_N)$.

⁹As in ordinary calculus in Euclidean space, we can also allow for explicit F dependence in ϕ by writing $\phi(\theta, F)$. Allowing this level of generality, though, is notationally burdensome and not typical in the analysis of the influence functions for Z-estimators. So we omit this dependence for simplicity.

¹⁰The factor of N arises to re-write the expectation as a sum over unit-valued weights.

We refer to the influence function applied with $F = \hat{F}_N$ as the *empirical influence function* (Hampel, 1986). We conclude that the ψ_n that we use to form our approximation are the values of the empirical influence function at the datapoints d_1, \dots, d_N . For this reason, we refer to the ψ_n as influence scores.

(c) The sum of the influence scores is zero. We can now use standard properties of the influence function to reason about $\vec{\psi}$. For instance, the fact that $\sum_{n=1}^N \psi_n = 0$ follows from Eq. 21 and the fact that $\hat{\theta}$ solves Eq. 1.

(d) The noise is an estimator of the standard deviation of the limiting distribution of the quantity of interest (influence function version). Observe that, by our influence function development above, we can write the squared noise as follows.

$$\hat{\sigma}_\psi^2 := N \|\vec{\psi}\|_2^2 = \frac{1}{N} \sum_{n=1}^N (N\psi_n)^2 = \frac{1}{N} \sum_{n=1}^N \text{IF}(d_n; T_\phi, \hat{F}_N)^2, \quad (23)$$

Recall that we saw above that $\hat{\sigma}_\psi^2$ consistently estimates the variance of the limiting distribution of $\sqrt{N}\hat{\phi}$, first in the special case of OLS (Section 3.1.2, paragraph (e)) and then for Z-estimators in general (Section 3.2.1, paragraph (b)). We can now see that those results are themselves special cases of the following well-known relationship between the influence function and the limiting variance of its corresponding functional:

$$\sqrt{N} \left(T(\hat{F}_N) - T(F_\infty) \right) \rightsquigarrow \mathcal{N} \left(0, \mathbb{E} \left[\text{IF}(d_1; T, F_\infty)^2 \right] \right), \quad (24)$$

where the expectation in the preceding display is taken with respect to $d_1 \sim F_\infty$ (see, e.g., Hampel (1986, Eq. 2.1.8)).¹¹ Specifically, if we can show that σ_ψ , the probability limit of $\hat{\sigma}_\psi$, is equal to $\mathbb{E} [\text{IF}(d_1; T, F_\infty)^2]$, then Eq. 24 would imply $\sqrt{N}(T_\phi(\hat{F}_N) - T_\phi(F_\infty)) \rightsquigarrow \mathcal{N}(0, \sigma_\psi^2)$, just as we showed in Eq. 18 using the sandwich covariance estimator. In our case, under standard assumptions, one can show directly from Eqs. 7 and 10 that $\text{IF}(d_n; T_\phi, \hat{F}_N) \xrightarrow{P} \text{IF}(d_n; T_\phi, F_\infty)$, almost surely in d_n . A law of large numbers can then be applied to Eq. 23 giving the desired result.

(e) AMIP robustness is different from gross error robustness. Roughly speaking, an estimator is considered non-robust to gross errors if its influence function is unbounded (Huber, 1981). For instance, the influence function arising from the OLS Z-estimator (Section 3.1) is classically known to be non-robust to gross errors. When an influence function is unbounded, one can produce arbitrarily large changes in the quantity of interest by making arbitrarily large changes to a single datapoint. Gross-error robustness is motivated by the possibility that some small number of datapoints come from a distribution arbitrarily different from the

¹¹Though Eq. 24 can provide useful intuition, as it does in our case, it is often easier in any particular problem to prove asymptotic results directly rather than through the functional analysis perspective of this section, since stating precise and general conditions under which Eq. 24 holds can be challenging. See, for example, the discussion in Serfling (2009, Chapter 6) or Van der Vaart (2000, Chapter 20).

model’s posited distribution. By contrast, to assess AMIP robustness, we do not make arbitrarily large changes to datapoints. We simply remove datapoints. And the analysis is AMIP-non-robust if a change of a particular size (Δ) can be induced, rather than an arbitrarily large change. Consequently, problems with unbounded influence functions (such as OLS in Section 3.1) can be AMIP-robust if $\Delta/\hat{\sigma}_\psi$ is sufficiently large. And perfectly specified problems with no outliers can be AMIP non-robust if $\Delta/\hat{\sigma}_\psi$ is sufficiently small.

(f) Outliers affect AMIP robustness through the noise. Consideration of gross-error robustness encourages users to examine their data for unusual “outliers” in the data; once outliers are removed or their influence diminished, the problem is considered gross-error robust. Since outliers are heuristically associated with heavy-tailed data distributions, one might expect the effect of outliers to affect AMIP robustness through the shape variable $\hat{\mathcal{T}}_\alpha$. However, our analysis of Section 3.2.1, paragraph (c) shows that gross errors actually *reduce* $\hat{\mathcal{T}}_\alpha$ and so render an estimator more robust for a fixed $\hat{\sigma}_\psi$. This observation does not imply that gross errors decrease AMIP sensitivity. Rather, gross errors increase AMIP sensitivity through the noise $\hat{\sigma}_\psi$. And, as we have seen, effects on $\hat{\sigma}_\psi$ also affect the computation of standard errors.

3.3 Accuracy of the approximation

In Section 3.1.2, paragraph (f) we argued that our approximation was accurate in OLS for small α . Now we extend that argument to the general case. In particular, we state sufficient conditions under which $\phi^{\text{lin}}(\vec{w})$ provides a good approximation to $\phi(\hat{\theta}(\vec{w}), \vec{w})$ for small α uniformly for $\vec{w} \in W_\alpha$. Our key result, Theorem 1, holds exactly in finite samples with bounds that are, in principal, computable. Additionally, the corresponding bounds can also be expected to hold with probability approaching one as $N \rightarrow \infty$ under standard assumptions.

3.3.1 Controlling the residual of a Taylor series

The linear approximation we use in Eq. 5 is a Taylor series, so its accuracy can be controlled by controlling the Taylor series residual. Giordano et al. (2019b) states conditions under which the first-order Taylor series approximation to $\hat{\theta}(\vec{w})$ is accurate—precisely when using the derivative as given in Eq. 9. Under additional smoothness assumptions on ϕ , we can extend those results to our present Eq. 5. Since the Taylor series expansion is expressed in terms of observable non-asymptotic quantities, the resulting error bounds hold exactly in finite sample and are, in principle, computable.

We first state assumptions under which the linear approximation is accurate for the vector $\hat{\theta}(\vec{w})$.

Assumption 3 ((Giordano et al. (2019b), Assumptions 1-4)). Let W_α be the set of weight vectors with no more than $\lfloor \alpha N \rfloor$ zeros as given by Eq. 2. Assume there exists a compact domain $\Omega_\theta \subseteq \mathbb{R}^D$ containing $\hat{\theta}(\vec{w})$ for all $\vec{w} \in W_\alpha$, such that

1. For all $\theta \in \Omega_\theta$ and all n , $\theta \mapsto G(\theta, d_n)$ is continuously differentiable with derivative

$$\left. \frac{\partial G(\theta, d_n)}{\partial \theta^T} \right|_\theta =: H(\theta, d_n).$$

2. For all $\theta \in \Omega_\theta$, there exists $C_{op} < \infty$ such that $\sup_{\theta \in \Omega_\theta} \left\| \frac{1}{N} \sum_{n=1}^N H(\theta, d_n) \right\|_{op} \leq C_{op}$.
3. There exists a constant $C_{gh} < \infty$ such that

$$\sup_{\theta \in \Omega_\theta} \max \left\{ \frac{1}{N} \sum_{n=1}^N \|G(\theta, d_n)\|_2^2, \frac{1}{N} \sum_{n=1}^N \|H(\theta, d_n)\|_2^2 \right\} \leq C_{gh}^2.$$

4. There exists a Δ_θ and an $L_h < \infty$ such that

$$\sup_{\theta: \|\theta - \hat{\theta}\|_2 \leq \Delta_\theta} \frac{1}{N} \sum_{n=1}^N \left\| H(\theta, d_n) - H(\hat{\theta}, d_n) \right\|_2^2 / \|\theta - \hat{\theta}\|_2^2 \leq L_h^2.$$

Roughly speaking, Assumption 3 states that the estimating equation is smooth and non-singular, that the sample averages are uniformly bounded, and that the estimating equation's derivatives are Lipschitz. Other than the size of the domain Ω_θ , Assumption 3 does not depend on W_α , nor on any asymptotic quantities; it states only (reasonable) assumptions on the actual problem at hand.

Under Assumption 3, we are able to apply Theorem 1 of Giordano et al. (2019b) for W_α and thereby prove the uniform accuracy of a linear approximation to $\hat{\theta}(\vec{w})$ for all $\vec{w} \in W_\alpha$. To extend the accuracy of an approximation of $\hat{\theta}(\vec{w})$ to our quantity of interest ϕ naturally requires smoothness assumptions on ϕ , which we now state.

Assumption 4. Define the re-scaled weights $\delta_n := \vec{w}_n / \sqrt{N}$, and assume that $\theta, \delta \mapsto \phi(\theta, \sqrt{N}\delta)$ has continuous partial derivatives, that the partial derivatives' $\|\cdot\|_2$ -norm evaluated at $\theta = \hat{\theta}(\vec{1})$ and $\vec{w} = \vec{1}$ is bounded by a finite constant C_ϕ , and that the partial derivatives are Lipschitz in $\|\cdot\|_2$ with finite constant L_ϕ .

We can now state our main accuracy theorem.

Theorem 1. Let Assumptions 3 and 4 hold. For sufficiently small α , there exist constants C_1 and C_2 , defined in terms of quantities given in Assumptions 3 and 4,

such that¹²

$$\sup_{\vec{w} \in W_\alpha} \left\| \phi^{\text{lin}}(\vec{w}) - \phi(\hat{\theta}(\vec{w}), \vec{w}) \right\| \leq C_1 \alpha \quad \text{and} \quad \sup_{\vec{w} \in W_\alpha} \left\| \phi(\hat{\theta}(\vec{w}), \vec{w}) - \hat{\phi} \right\| \leq C_2 \sqrt{\alpha}. \quad (25)$$

When α is small, we expect $\alpha \ll \sqrt{\alpha}$ (for example, when $\alpha = 0.01$, $\sqrt{\alpha} = 0.1 \gg 0.01$), so Theorem 1 states that the bound in the error of our linear approximation shrinks faster than the bound in the function itself as $\alpha \rightarrow 0$.

Theorem 1 is a finite-sample result, applying exactly to the problem at hand. All else equal, finite-sample results are preferable to asymptotic ones. Nevertheless, due to the many loose bounds employed in the proof, we do not expect the constants to be useful in practice. Additionally, Theorem 1 of Giordano et al. (2019b) may in theory require α to be smaller than $1/N$, resulting in a vacuous statement. Improving these shortcomings is an important avenue for future work (e.g. Giordano et al. (2019a); Wilson et al. (2020)). But it is therefore useful to observe that, when uniform laws of large numbers apply to $\theta \mapsto \|G(\theta, \cdot)\|_2$ and $\theta \mapsto \|H(\theta, \cdot)\|_2$, and the limiting functions are also non-singular, bounded, and Lipschitz, then one can expect Assumption 3 to hold with high probability and finite constants as $N \rightarrow \infty$. A precise statement of the necessary conditions for such asymptotics to apply is given in Lemma 1 of Giordano et al. (2019b).

3.3.2 Limitations of linear approximations

In every case we examine in our applications in Section 4, we manually re-run the analysis without the data points in the removal set \hat{S}_α ; in doing so, we find that the change suggested by the approximation is nearly always achieved in practice (a notable exception is given and discussed at the end of Section 4.4). However, linear approximations are only approximations, and intuition about the potential weaknesses of linear approximations in general apply to our approximation. The crux of Theorem 1 is that small α implies that $\vec{w} - \vec{1}$ is small, thus we can control the error of a linear approximation in \vec{w} evaluated at $\vec{1}$. Conversely, one would not expect the approximation to work well in general for large α and the correspondingly larger $\vec{w} - \vec{1}$.

As an extreme example, consider when the linear approximation reports that there is no feasible way to effect a particular change; i.e., when $\hat{\alpha}_\Delta^* = \text{NA}$ (see Definition 2). Such a result may seem to imply that, no matter how many datapoints one removes, the estimator will not change by an amount Δ , which is often absurd. However, such a result should be taken to mean that one would have to remove such a large proportion α of datapoints that the linear approximation on which we are basing the $\hat{\alpha}_\Delta^*$ is invalid. A more accurate interpretation of $\hat{\alpha}_\Delta^* = \text{NA}$ is that no

¹²We note that the rate $\sqrt{\alpha}$ is determined by a simple but coarse Cauchy-Schwartz bound (see Lemma 2). Tighter bounds may be achievable when the random variables $\|G(\theta, d_n)\|_2$ and $\|H(\theta, d_n)\|_2$ are uniformly integrable (see, e.g., Section 2.5 of Van der Vaart (2000)).

small proportion of points can be removed to produce a change Δ , for if there were such a small proportion, the linear approximation would have discovered it.

Similarly, linear approximations cannot be expected to work well near the boundary of parameter spaces. For example, if the quantity of interest is a variance, then the true parameter is constrained to be positive, but our linear approximation is not. It can help to linearize the problem using unconstrained reparameterizations (e.g., linearly approximating the log variance rather than variance). However, as we show in Section 4.4, simply transforming to an unconstrained space is still not guaranteed to produce accurate approximations near the boundary in the original, constrained space.

4 Applied experiments

4.1 The Oregon Medicaid experiment

In our first experiment, we show that even empirical analyses that display little classical uncertainty can be sensitive to the removal of less than 1% of the sample. We consider the Oregon Medicaid study (Finkelstein et al., 2012) and focus on health outcomes. The standard errors of the treatment effects are small relative to effect size; against a null hypothesis of no effect, most p values are well below 0.01. Yet we find that for most of the results, removing less than 1% of the sample can produce a significant result of the opposite sign to the full-sample analysis. In one case, removing less than 0.05% of the sample can change the significance of the result.

4.1.1 Background and replication

First we provide some context for the analysis and results of Finkelstein et al. (2012). In early 2008, the state of Oregon opened a waiting list for new enrollments in its Medicaid program for low-income adults. Oregon officials then drew names by lottery from the 90,000 people who signed up, and those who won the lottery could sign up for Medicaid along with any of their household members. This setup created a randomization into treatment and control groups at the household level. The Finkelstein et al. (2012) study measures outcomes one year after the treatment group received Medicaid. About 25% of the treatment group did indeed have Medicaid coverage by the end of the trial. The main analysis investigates treatment assignment as treatment itself (“intent to treat” or ITT analysis) and uses treatment assignment as an instrumental variable for take-up of insurance coverage (“local average treatment effect” or LATE analysis).

We focus on the health outcomes of winning the Medicaid lottery, which appear in Panel B from Table 9 of Finkelstein et al. (2012). Each of these J outcomes is denoted by y_{ihj} for individual i in household h for outcome type j . The data

sample to which we have access consists of survey responders ($N = 23,741$); some responders are from the same household. The variable LOTTERY_h equals one if household h won the Medicaid lottery, and zero otherwise. All regressions use a set of covariates X_{ih} comprised of household size fixed effects, survey wave fixed effects, and the interaction between the two. All regressions also use a set of demographic and economic covariates V_{ih} . To infer the ITT effects of winning the Medicaid lottery, the authors estimate the following model via OLS:

$$y_{ihj} = \beta_0 + \beta_1 \text{LOTTERY}_h + \beta_2 X_{ih} + \beta_3 V_{ih} + \epsilon_{ihj}.$$

To infer the LATE of taking up Medicaid on compliers, the authors employ an Instrumental Variables (IV) strategy using the lottery as an instrument for having Medicaid insurance. All standard errors are clustered on the household, and all regressions are weighted using survey weights defined by the variable `weight_12m`. We have access to the following seven outcome variables, presented in Panel B of Table 9 of the original paper (as well as our tables below) in the following order: a binary indicator of a self-reported measure of health being good or very good or excellent (not fair or poor), a binary indicator of self-reported health not being poor, a binary indicator of health being about the same or improving over the last six months, the number of days of good physical health in the past 30 days, the number of days on which poor physical or mental health did not impair usual activities, the number of days mental health was good in the past 30 days, and an indicator of not being depressed in last two weeks. We replicate Panel B of Table 9 of Finkelstein et al. (2012) exactly, both for the ITT effect ($\hat{\beta}_1$) for the entire population and for the LATE on compliers ($\hat{\pi}_1$). Both analyses show strong evidence for positive effects on all health measures, with most p values well below 0.01.

4.1.2 AMIP Sensitivity Results

For each health outcome in Panel B from Table 9 of Finkelstein et al. (2012), we compute the AMIP to assess how many data points one needs to remove to change the sign of the treatment effect, the significance of the treatment effect, or produce a significant result of the opposite sign. The sensitivity of the LATE analysis is shown in Table 1 and the sensitivity of the ITT analysis is shown in Table 2. In both cases we use exactly the models from the original paper, with all fixed effects and controls included and with clustering at the household level. For most outcomes, for both the LATE and ITT analysis, the sign of the treatment effect can be changed by removing around 0.5% of the data, or approximately 100 data points in a sample of approximately 22,000. The most robust outcome, “Health being better than fair” (“Health genflip 12m”), requires the removal of a little over 1% of the sample to change the sign. Across the various outcomes, we can drop even less of the sample to change the results from significant to non-significant. In some cases, we need

Study case	Original estimate	Target change	Refit estimate	Observations dropped
Health genflip 12m	0.133 (0.026)*	Sign change	-0.006 (0.025)	275 = 1.18%
		Significance change	0.044 (0.026)	162 = 0.69%
		Significant sign change	-0.043 (0.024)	381 = 1.63%
Health notpoor 12m	0.099 (0.018)*	Sign change	-0.003 (0.015)	155 = 0.66%
		Significance change	0.027 (0.016)	100 = 0.43%
		Significant sign change	-0.030 (0.015)*	219 = 0.94%
Health change flip 12m	0.113 (0.023)*	Sign change	-0.006 (0.022)	197 = 0.84%
		Significance change	0.039 (0.022)	106 = 0.45%
		Significant sign change	-0.049 (0.022)*	291 = 1.24%
Not bad days total 12m	1.317 (0.563)*	Sign change	-0.023 (0.535)	73 = 0.33%
		Significance change	1.078 (0.558)	10 = 0.05%
		Significant sign change	-1.009 (0.521)	144 = 0.66%
Not bad days physical 12m	1.585 (0.606)*	Sign change	-0.040 (0.577)	87 = 0.41%
		Significance change	1.131 (0.597)	20 = 0.09%
		Significant sign change	-1.141 (0.566)*	164 = 0.77%
Not bad days mental 12m	2.082 (0.640)*	Sign change	-0.062 (0.607)	123 = 0.57%
		Significance change	1.171 (0.625)	42 = 0.19%
		Significant sign change	-1.201 (0.594)*	212 = 0.98%
Nodep Screen 12m	0.078 (0.025)*	Sign change	-0.005 (0.024)	123 = 0.53%
		Significance change	0.046 (0.024)	42 = 0.18%
		Significant sign change	-0.050 (0.023)*	220 = 0.95%

Table 1: Medicaid profit results with IV for a range of outcome variables. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

Study case	Original estimate	Target change	Refit estimate	Observations dropped
Health genflip 12m	0.039 (0.008)*	Sign change	-0.004 (0.008)	286 = 1.22%
		Significance change	0.013 (0.008)	163 = 0.70%
		Significant sign change	-0.021 (0.008)*	422 = 1.81%
Health notpoor 12m	0.029 (0.005)*	Sign change	-0.001 (0.005)	156 = 0.67%
		Significance change	0.008 (0.005)	101 = 0.43%
		Significant sign change	-0.009 (0.004)*	224 = 0.96%
Health change flip 12m	0.033 (0.007)*	Sign change	-0.002 (0.006)	198 = 0.85%
		Significance change	0.011 (0.007)	106 = 0.45%
		Significant sign change	-0.015 (0.006)*	292 = 1.25%
Not bad days total 12m	0.381 (0.162)*	Sign change	-0.013 (0.157)	74 = 0.34%
		Significance change	0.306 (0.161)	11 = 0.05%
		Significant sign change	-0.309 (0.153)*	147 = 0.67%
Not bad days physical 12m	0.459 (0.175)*	Sign change	-0.017 (0.169)	88 = 0.41%
		Significance change	0.328 (0.172)	20 = 0.09%
		Significant sign change	-0.344 (0.165)*	166 = 0.78%
Not bad days mental 12m	0.603 (0.184)*	Sign change	-0.027 (0.178)	124 = 0.57%
		Significance change	0.340 (0.181)	42 = 0.19%
		Significant sign change	-0.381 (0.175)*	216 = 1.00%
Nodep Screen 12m	0.023 (0.007)*	Sign change	-0.001 (0.007)	123 = 0.53%
		Significance change	0.013 (0.007)	43 = 0.19%
		Significant sign change	-0.015 (0.007)*	225 = 0.97%

Table 2: Medicaid profit results with OLS for a range of outcome variables. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

remove only 10 or 20 data points to effect a change in significance. Finally, for most outcomes, we can remove less than 1% of the data to produce a significant result of the opposite sign. The only two exceptions, “Health genflip 12m” and “Health change flip 12m”, require the removal of slightly more than 1% to generate a significant result with the opposite sign.

We check the performance of the approximation for each analysis by re-running the model after manually removing the data points in the Approximate Most Influential Set. The result of this procedure is shown in the “Refit Estimate” column of Tables 1 and 2. For almost every result in each table, our approximate metric reliably uncovers combinations of data points that do deliver the claimed changes. As we discuss in Section 2.2.1, the changes recorded in the “Refit Estimate” column of Tables 1 and 2 form a lower bound on the true worst-case finite-sample sensitivity.

By comparing Table 1 with Table 2, we see that the ITT results, estimated via OLS, are not notably more AMIP-robust than the LATE results, which are estimated via IV. This may seem at first counterintuitive based on a heuristic belief that IV is in some sense a less “robust” analysis than OLS in finite sample: for example, recent authors, including Young (2019), have suggested that the uncertainty intervals for IV may be more poorly calibrated in finite samples than the intervals for OLS. However, as we discuss in Section 3, the quality of being “robust” in the sense of a finite-sample estimator providing a good approximation to an asymptotic quantity is simply unrelated to AMIP robustness. Neither the size of the AMIP itself nor the accuracy of the AMIP approximation depends on asymptotic arguments (see, e.g., Section 3.2.2, paragraph (a) and the discussion of Theorem 1). The AMIP measures the sensitivity to data ablation of a particular procedure on a particular dataset and is indifferent to the fidelity of the chosen quantity of interest to some asymptotic limit. For this reason, a procedure such as IV may be “non-robust” in the sense of having poor coverage in finite sample (as reported by Young (2019)) and yet be AMIP-robust, or vice versa. The two notions of “robustness” are simply different.

4.2 Cash transfers

We next show that an empirical analysis can still be AMIP-non-robust even after outliers are removed. To that end, we apply our techniques to examine the robustness of the main analysis from Angelucci and De Giorgi (2009), one of the flagship studies showing the impact of cash transfers on ineligible (“non-poor”) households in the same villages, also known as “spillover effects.” The authors trimmed the consumption outcome for the non-poor households due to concerns about the influence of the largest values. Yet while the analysis on the poor households is quite robust, the analysis on the non-poor households—whom the trimming protocol actually affects—is much more sensitive.

4.2.1 Background and replication

Angelucci and De Giorgi (2009) employ a randomized controlled trial to study the impact of Progresa, a social program giving cash gifts to eligible poor households in Mexico. The randomization occurs at the village level. So one can estimate both a main effect on the poor households selected to receive Progresa and also the impact on the non-eligible “non-poor” households located in the same villages as Progresa-receiving poor households.

The main results of the paper show that there are strong positive impacts of Progresa on total household consumption measured as an index both for eligible poor households and for the non-eligible households; see Table 1 of Angelucci and De Giorgi (2009). The variable C_ind_{it} denotes total household consumption for household i in time period t . Values of C_ind_{it} above 10,000 are removed; such households are, by definition, non-poor. The authors study three different time periods separately to detect any change in the impact between the short and long term. They condition on a large set of variables (a household poverty index, land size, head of household gender, age, whether the household speaks an indigenous language, and literacy; at the locality level, a poverty index, and the number of households) to help ensure a fair comparison between households in the treatment and control villages. In this case these controls are important; the effects on the “non-poor” households are significant at the 5% level when the controls are included, but they are only significant at the 10% level in a simple regression on a dummy for treatment status.

The full data for the paper is available on the website of the *American Economic Review* thanks to the open-data policies of the journal and the authors. We can successfully replicate the results of this analysis with the controls and without, and we proceed with the controls in our present analysis in accordance with the original authors’ preferred specification. We consider the time periods indexed as $t = 8, 9, 10$ in the dataset provided, though we note that the authors do not rely on the results at $t = 8$ as the roll-out was still ongoing. We employ K control variables, where X_{itk} is the k -th variable for household i in period t . Then we run the following regression:

$$C_ind_{it} = \beta_0 + \beta_1 \mathbf{treat}_{poor,i} + \beta_2 \mathbf{treat}_{nonpoor,i} + \sum_{k=1}^K \beta_{2+k} X_{itk} + \epsilon_{it}.$$

Here, $\mathbf{treat}_{poor,i}$ refers to an interaction between the treatment indicator and an indicator for being a poor household; correspondingly, $\mathbf{treat}_{nonpoor,i}$ is an interaction between the treatment indicator and an indicator for being a non-poor household. We are able to exactly replicate the results of Table 1 of Angelucci and De Giorgi (2009), which exhibits positive effects of cash transfers.

Study case	Original estimate	Target change	Refit estimate	Observations dropped
Poor, period 8	17.312 (4.576)*	Sign change	-0.656 (3.745)	252 = 2.30%
		Significance change	7.284 (4.087)	83 = 0.76%
		Significant sign change	-7.212 (3.443)*	464 = 4.24%
Poor, period 9	27.924 (5.770)*	Sign change	-1.377 (4.406)	345 = 3.58%
		Significance change	7.077 (4.555)	146 = 1.52%
		Significant sign change	-8.951 (4.251)*	588 = 6.11%
Poor, period 10	33.861 (4.468)*	Sign change	-2.559 (3.541)	697 = 6.63%
		Significance change	4.806 (3.684)	435 = 4.14%
		Significant sign change	-9.416 (3.296)*	986 = 9.37%
Non-poor, period 8	-5.444 (7.133)	Sign change	0.260 (6.410)	5 = 0.11%
		Significance change	-12.845 (6.635)	16 = 0.35%
		Significant sign change	9.670 (5.573)	24 = 0.52%
Non-poor, period 9	22.852 (10.000)*	Sign change	-0.365 (7.542)	21 = 0.55%
		Significance change	16.506 (9.114)	3 = 0.08%
		Significant sign change	-11.733 (7.113)	53 = 1.38%
Non-poor, period 10	21.493 (9.405)*	Sign change	-0.573 (6.750)	30 = 0.70%
		Significance change	16.262 (8.927)	3 = 0.07%
		Significant sign change	-10.845 (6.467)	92 = 2.16%

Table 3: Cash transfers results for various periods and treatment groups. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

4.2.2 AMIP Sensitivity Results

We apply our methodology to assess how many data points one need remove to change the sign, the significance, or to generate a significant result of the opposite sign to that found in the full sample. We focus on the latter two time periods, as households had received only partial transfers in the first time period, but we show all three in order to replicate Table 1 from the original paper. Table 3 shows our results. Focusing on periods 9 and 10, we find that the inferences on the direct effects on the poor households are quite robust, but the inferences on the indirect effects are less so. For the analysis of the poor, one typically needs to remove much more than 1% of the sample to change conclusions. For the analysis of the non-poor, we can remove less than 0.5% of the data to change conclusions. In fact, we can remove only 3 data points in a sample of approximately 10,000 households to change the significance status for both $t = 9$ and $t = 10$.

We again check the quality of our approximation. The “Refit Estimate” column in Table 3 shows the results of manually re-running each analysis after removing the implicated data points. In most cases the AMIP correctly identifies a combination of data points that can make the claimed changes to the conclusions of the study. Although there are a few cases where re-running the analysis fails to produce the predicted statistically significant sign change, the observed changes are still large enough to be of practical interest. Furthermore, it is likely that the removal of a few additional points would in fact produce the desired statistically significant sign reversals.

Finally, we note that these results constitute an illustration of how gross error robustness is distinct from AMIP robustness (see Section 3.2.3, paragraph (e)). Recall that Angelucci and De Giorgi (2009) removed (non-poor) datapoints for which consumption was greater than 10,000. By removing outliers of the consumption variable in this way, the authors of this study made what is typically considered a conservative choice in view of classical robustness concerns about gross error sensitivity. Yet, as we have shown in Table 3, qualitative conclusions concerning the non-poor households remain non-robust to the removal of a small number of datapoints, which demonstrates empirically that one cannot necessarily make an analysis AMIP-robust by simply trimming outliers. Indeed, as we showed above in Section 3.1, even perfectly specified OLS regressions with no aberrant data points can be AMIP-non-robust if the signal to noise ratio is too low.

4.3 Seven RCTs of microcredit: Linear regression analysis

We now show that even a simple 2-parameter linear model that performs a comparison of means between the treatment and control group of a randomized trial can be highly sensitive. To that end, we consider the analysis of seven randomized controlled trials of expanding access to microcredit, first aggregated in Meager (2019). In Section 4.4 below, we will consider a more complicated Bayesian hierarchical model on the same data.

4.3.1 Background

Each of the seven microcredit studies was conducted in a different country, and each study selected certain communities to randomly receive greater access to microcredit. Researchers either built a branch, or combined building a branch with some active outreach, or randomly selected borrowers among those who applied. The selected studies are: Angelucci et al. (2015), Attanasio et al. (2015), Augsburg et al. (2015), Banerjee et al. (2015), Crépon et al. (2015), Karlan and Zinman (2011), and Tarozzi et al. (2015). Six of these studies were published in a special issue of the *American Economics Journal: Applied Economics* on microcredit. All seven studies together are commonly considered to represent the most solid evidence base for understanding the impact of microcredit.

We follow the original studies and Meager (2019) in analyzing the impact of access to microcredit as the treatment of interest. The studies range in their sample sizes from around 1,000 households in Mongolia (Attanasio et al., 2015) to around 16,500 households in Mexico (Angelucci et al., 2015). We first focus on the headline results on household business profit regressed on an intercept and a binary variable indicating whether a household was allocated to the treatment group or to the control group. For household i in site k , let Y_{ik} denote the profit measured, and let

Study case	Original estimate	Target change	Refit estimate	Observations dropped
Bosnia	37.534 (19.780)	Sign change	-2.226 (15.628)	14 = 1.17%
		Significance change	43.732 (18.889)*	1 = 0.08%
		Significant sign change	-34.929 (14.323)*	40 = 3.35%
Ethiopia	7.289 (7.893)	Sign change	-0.053 (2.513)	1 = 0.03%
		Significance change	15.356 (7.763)*	45 = 1.45%
		Significant sign change	-8.755 (1.852)*	66 = 2.12%
India	16.722 (11.830)	Sign change	-0.501 (8.221)	6 = 0.09%
		Significance change	22.895 (10.267)*	1 = 0.01%
		Significant sign change	-16.638 (7.537)*	32 = 0.47%
Mexico	-4.549 (5.879)	Sign change	0.398 (3.194)	1 = 0.01%
		Significance change	-10.962 (5.565)*	14 = 0.08%
		Significant sign change	7.030 (2.549)*	15 = 0.09%
Mongolia	-0.341 (0.223)	Sign change	0.021 (0.184)	16 = 1.66%
		Significance change	-0.436 (0.220)*	2 = 0.21%
		Significant sign change	0.361 (0.147)*	38 = 3.95%
Morocco	17.544 (11.401)	Sign change	-0.569 (9.920)	11 = 0.20%
		Significance change	21.720 (11.003)*	2 = 0.04%
		Significant sign change	-18.847 (9.007)*	30 = 0.55%
Philippines	66.564 (78.127)	Sign change	-4.014 (57.204)	9 = 0.81%
		Significance change	138.929 (66.880)*	4 = 0.36%
		Significant sign change	-122.494 (49.409)*	58 = 5.21%

Table 4: Microcredit regressions for the profit outcome. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

T_{ik} denote the treatment status. We estimate the following model via OLS:

$$Y_{ik} = \beta_0 + \beta T_{ik} + \epsilon_{ik}. \quad (26)$$

This regression model compares the means in the treatment and control groups and estimates the difference as $\hat{\beta}$. We follow Meager (2019) in omitting the control variables or fixed effects from the regressions in order to examine the robustness of this fundamental procedure. But in principle this omission should make no difference to the estimate $\hat{\beta}$, and indeed it does not (Meager, 2019).¹³

4.3.2 AMIP sensitivity results

The sensitivity results for the linear regression of profit on microcredit access appear in Table 4. In all cases, by removing less than 1% of the data points can change either the sign or the significance. In three of the studies, one can drop less than 1% of the data points to generate a result of the opposite sign that would be deemed significant at the 5% level. Mexico, the largest study, is the most sensitive: a

¹³The omission may in principle make a difference to the inference on β by affecting the standard errors. However, it turns out that in these studies the additional covariates make very little difference to the standard errors. We also do not cluster the standard errors at the community level for the same reason; the results are not substantially changed. Running the regression above in each of the seven studies delivers almost identical results to the preferred specification, as it should if intra-cluster correlations are weak and covariates are not strongly predictive of household profit.

Study case	Original estimate	Target change	Refit estimate	Observations dropped
Bosnia	-5.803 (2.819)*	Sign change	0.395 (2.135)	10 = 1.00%
		Significance change	-4.870 (2.693)	1 = 0.10%
		Significant sign change	5.130 (1.978)*	33 = 3.31%
India	-1.643 (0.576)*	Sign change	0.035 (0.506)	41 = 0.60%
		Significance change	-1.051 (0.536)*	8 = 0.12%
		Significant sign change	1.059 (0.487)*	85 = 1.25%
Mexico	-0.082 (0.094)	Sign change	0.000 (0.091)	12 = 0.07%
		Significance change	-0.180 (0.091)*	14 = 0.09%
		Significant sign change	0.176 (0.087)*	55 = 0.33%
Mongolia	1.523 (2.103)	Sign change	-0.033 (0.973)	3 = 0.31%
		Significance change	2.717 (1.027)*	10 = 1.04%
		Significant sign change	-2.623 (0.689)*	45 = 4.68%
Morocco	-0.420 (0.723)	Sign change	0.047 (0.669)	3 = 0.05%
		Significance change	-1.351 (0.667)*	14 = 0.26%
		Significant sign change	1.252 (0.602)*	23 = 0.42%

Table 5: Microcredit regressions for the temptation outcome. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

single data point among the 16,561 households in Mexico determines the sign (as also discussed above in Section 2.4). To produce a statistically significant result of the opposite sign—that is, to turn Mexico’s noisy negative result into a “strong” positive result—one need remove only 15 data points, less than 0.1% of the sample. Mongolia, the smallest study in terms of sample size, is among the most robust in terms of sign changes; it takes 2% of the sample to change the sign. Producing a significant result of the opposite sign also requires more than 1% removal in the Philippines, Bosnia, Ethiopia, and Mongolia—whereas Mexico, India, and Morocco are more sensitive. We check the performance of our approximation by manually re-running the analysis with the data removed; the “Refit Estimate” column shows that the claimed reversal is always achieved in practice for these analyses.

By comparing the results of the present section with those of Sections 4.1 and 4.2, we can confirm the conclusion of Section 3.2.2, paragraph (d) that standard errors are, in general, distinct from AMIP sensitivity. Despite the fact that original estimates of Table 4 are statistically insignificant, some of these non-significant results are more AMIP-robust than some of the significant results in the Cash Transfers and Oregon Medicaid examples; consider the “Significant sign change” result in the Philippines study, for example.

We next demonstrate that the AMIP sensitivity observed in Table 4 cannot simply be ascribed to statistical insignificance. To do so, we consider a different outcome with smaller variability and show that it reveals a similar sensitivity to the profit outcome. The variable we now consider is household consumption spending on temptation goods such as alcohol, chocolate, and cigarettes, since the effect of microcredit on temptation spending was estimated by Meager (2019) with the greatest precision of all six considered outcome variables. Table 5 shows the results

of applying the AMIP to the same regression given in Eq. 26, but with temptation spending as the outcome. While somewhat more robust than the profit analyses, the difference in the approximate removal proportions in Table 5 is not large.

Finally, one might be tempted to ascribe the AMIP-non-robust results in Table 4 to outliers resulting from the heavy tails of the household profit variable (a phenomenon well-documented by Meager (2020)). However, as we discuss in Section 3.2.3, paragraph (e) above, gross error robustness is qualitatively distinct from AMIP robustness (see also the discussion of outlier trimming at the end of Section 4.2). Indeed, the more complex hierarchical model of the next section, Section 4.4, was designed precisely to accommodate the heavy tail of the household profit variable, and yet—as we will show—still exhibits a high degree of AMIP-sensitivity.

4.4 Seven RCTs of microcredit: Bayesian hierarchical tailored mixture model

In this section, we investigate a Bayesian hierarchical model, both demonstrating that even Bayesian analyses can exhibit considerable AMIP sensitivity, and showing an example of a parameter of interest for which our linear approximation performs badly. We specifically focus on a variational Bayes approximation to the tailored mixture model from Meager (2020). One might hope that any of the following aspects of the more complicated model might alleviate AMIP sensitivity: the use of hierarchical Bayesian evidence aggregation, the regularization from incorporation of priors, or the somewhat more realistic data-generating process captured in this specific tailored likelihood. Indeed, the approach of Meager (2020) was specifically motivated by the desire to capture important features of the data-generating process such as heavier tails. On the contrary, we find that the average estimated effects of microcredit remain sensitive according to the AMIP, as we did in the simpler models of Section 4.3. We also find that the linear approximation that underlies the AMIP performs poorly when attempting to decrease a particular hypervariance parameter, providing a concrete example of the limitations of our methodology, particularly for parameters near the boundary of the set of their allowable values.

4.4.1 Background

Following Meager (2020), we fit a hierarchical model (hereafter referred to as the “microcredit model”) to all the data from the seven microcredit RCTs. We model each outcome using a spike at zero and two lognormal tail distributions, one for the positive realizations of profit and one for the negative realizations. Within the model, microcredit can affect the proportion of data assigned to each of these three components as well as affecting the location and scale of the lognormal tails. There is a hierarchical shrinkage element to the model for each parameter. The hyper-

Model parameter	Original estimate	Target change	Refit estimate	Observations dropped
τ_-	0.102 (0.070)	Sign change	-0.042 (0.090)	31 = 0.09%
		Significance change	0.138 (0.071)	11 = 0.03%
		Significant sign change	-0.204 (0.106)	99 = 0.28%
τ_+	0.078 (0.033)*	Sign change	-0.021 (0.046)	74 = 0.21%
		Significance change	0.062 (0.033)	9 = 0.03%
		Significant sign change	-0.100 (0.054)	163 = 0.46%

Table 6: Microcredit mixture results for a selected set of model parameters. Standard errors and “significance” are based on the estimated 95% posterior credible intervals. The “Refit estimate” column shows the result of re-fitting the model removing the Approximate Most Influential Set. Stars indicate significance at the 5% level. Refits that achieved the desired change are bolded.

variances of the treatment effects are of particular interest because these capture heterogeneity in effects across studies and offer information about the transportability of results across settings.

The models in the original paper were fit via Hamiltonian Monte Carlo (HMC) with the software package Stan (Carpenter et al., 2017). It is possible to compute the Approximate Maximum Influence Perturbation for HMC, or for any Markov Chain Monte Carlo method, using the tools of Bayesian local robustness (Gustafson, 2000; Giordano et al., 2018), but the sensitivity of simulation-based estimators is beyond the scope of this paper. However, there are ways to estimate Bayesian posteriors via Z-estimators, such as with Variational Bayes (VB) techniques (Blei et al., 2017).¹⁴ Specifically, we fit the microcredit model using a variant of Automatic Differentiation Variational Inference (ADVI) described in Giordano et al. (2018, Section 5.2) (see also the original ADVI paper, Kucukelbir et al. (2017)). Since the posterior uncertainty estimates of vanilla ADVI are notoriously inaccurate, we estimated posterior uncertainty using linear response covariances, again following Giordano et al. (2018, Section 5.2).¹⁵ We verified that the posterior means and covariance estimates produced by our variational procedure and the corresponding estimates from running HMC with Stan were within reasonable agreement relative to the posterior standard deviation.

4.4.2 AMIP Sensitivity Results

We first consider the effect of microcredit on the location parameter of the positive and negative tails of profit, given respectively by the parameters τ_+ and τ_- . Roughly speaking, τ_+ and τ_- are both estimating the effect of microcredit averaged across all of the seven countries analyzed in Section 4.3. Our point estimates for τ_+ and τ_- are given by their respective VB posterior means. We used the linear response

¹⁴The Laplace approximation can also be expressed as a Z-estimator.

¹⁵When forming the Approximate Most Influential Set, we approximated the sensitivity only of the posterior means to data removal; the linear response covariances were considered fixed. However, when we report the results of re-fitting the model, we did re-calculate the linear response covariances at the new variational optimum.

Model parameter	Original estimate	Change type	Refit estimate	Prediction	Observations dropped
$\log \sigma_{\tau_-}$	-2.313	Drop 0.5% to increase	0.126	-0.811	177 = 0.50%
		Drop 0.5% to decrease	-0.151	-4.066	177 = 0.50%
$\log \sigma_{\tau_+}$	-3.100	Drop 0.5% to increase	-1.095	-1.204	177 = 0.50%
		Drop 0.5% to decrease	-1.598	-4.974	177 = 0.50%

Table 7: Results for the log posterior standard deviation estimates of the effect size distribution in the microcredit mixture model. Sign and significance are not meaningful for posterior standard deviations, so we drop 0.5% of datapoints to attempt to produce large positive and negative changes. The “Refit estimate” column shows the result of re-fitting the model removing Approximate Most Influential Set. The “Prediction” column shows the predicted change under the same perturbation.

covariance estimates to form a 95% posterior credible interval in place of confidence intervals, and consider a change “significant” if the posterior credible interval does not contain zero.

Table 6 shows the sensitivity of inference concerning τ_+ and τ_- . We see that the microcredit model’s estimates of the average effectiveness of microcredit remain highly sensitive to the removal of small percentages of the sample, despite being derived from a model that accounts for non-Gaussian data shape and is regularized by the priors. This sensitivity shows that Bayesian aggregation procedures do not necessarily produce AMIP-robust estimates.

We next examine the sensitivity of the hypervariances, which measure the variability of the effect of microcredit on these tails from country to country. Specifically, the parameters $\sigma_{\tau_+}^2$ and $\sigma_{\tau_-}^2$ represent the between-country variances of the effect of microcredit on positive and negative profit outcomes, respectively. The σ parameter can be thought of as the scale parameter analogue of the corresponding location parameter τ from Table 6. The hypervariances are of particular practical interest because they quantify how variable the effect of microcredit might be; small values of the hypervariance imply that all countries respond similarly to microcredit, whereas large values imply that one should not necessarily extrapolate the efficacy of microcredit from one country to another.

In order to avoid the possibility of extrapolating to negative variances, we form a linear approximation to our variational Bayes estimates of the posterior mean of $\log \sigma$. Since $\log \sigma$ is a scale parameter measuring the variability from country to country of the effect of microcredit, its sign is not particularly meaningful, nor is it particularly interesting to ask whether its posterior credible interval contains zero. Rather, we are interested in the magnitude of $\log \sigma$. So, to investigate robustness, we use the AMIP to check the approximate maximum change achievable in either direction (increasing or decreasing the magnitude of $\log \sigma$) by removing 0.5% of the sample, about the same fraction of the data as could generate a “significant” sign change for the τ_{\pm} parameters.

The results for the hypervariances, given in Table 7, represent a useful demon-

stration of the limitations of our linear approximation. We are able to find sets of datapoints which, when dropped, produce *increases* in the hypervariances, though the our linear approximation is not nearly as accurate as in the rest of our results above. When we attempted to drop points in order to decrease the hypervariances, however, the linear approximation failed utterly; the Approximate Most Influential Set designed to produce a decrease in the hypervariances instead produced a large *increase* upon refitting.

Given that the hypervariances are constrained to be positive, our failure to produce large decreases may not be surprising. Note that the hypervariances’ posterior expectations began very small, and that decreasing them pushes the posterior of the hypervariances closer to the boundary of the admissible space. Though the log variance may in principle take arbitrarily negative values, it nevertheless appears that the model exhibits strongly non-linear dependence on the data weights for very small variances. Designing useful diagnostics for detecting and explaining such deviations from nonlinearity in complex models is an interesting avenue for future work. In the meantime, Table 7 shows the importance, when possible, of checking the accuracy of the AMIP predictions by refitting the model, and of exercising caution when using the AMIP approximation near the boundary of the parameter space.

5 Conclusion

There are different ways of quantifying the dependence between the finite-sample realization and the conclusions of statistical inference. While this dependence has become synonymous with standard errors in frequentist statistics, the notions are equivalent only under a certain paradigm that considers a hypothetical perfect random resampling exercise for the purpose of evaluating a specific parameter within a given model. This hypothetical may not capture all the data sensitivity relevant to applied social science. Much of 20th century statistics, with its focus on standard errors and sampling uncertainty, has its origins in the context of randomized agricultural trials, where the difference in yield across multiple fields is well-modeled by independent sampling variation. Contrast with trials of economic interventions to alleviate poverty, where randomly sampling individuals or communities is a challenge and interventions may be applied across very different contexts. In applied economics, statistical models are often intended to provide tractable and interpretable summaries of the impact of interventions. In doing so, models can average information across individuals in ways that may not reflect policy interests.¹⁶ For our methods to safely inform economic policy decisions, then, we need additional

¹⁶In agricultural trials, total yields are the true quantity of interest; for microcredit trials, the average treatment effect is but a convenient summary. If the average profit were to increase slightly through one individual becoming wealthy while leaving all others destitute, one could consider the intervention a failure. By contrast, if a single plant produced an entire harvest’s worth of corn, the outcome would still be desirable, if strange.

tools beyond standard errors.

In this paper, we have offered one alternative way of conceiving of and quantifying the dependence of empirical results on the sample data, beyond standard errors. Sensitivity of conclusions to data removal under our metric does not necessarily imply a problem with the sample. But the goal of inference is not to learn about the sample, but rather to learn about the population. If minor alterations to the sample can generate major changes in the inference, and we know that the environment in which we do economics is changing all the time, we ought to be less confident that we have learned something fundamental about this broader population we seek to understand, for whom we ultimately seek to make policy. We do not mean to imply that the original analysis is invalid according to classical sampling theory, and we do not recommend that researchers abandon the original full-sample results even if they are not robust according to our metric. However, reporting our metrics alongside standard errors would improve our ability to understand and interpret the findings of a given analysis.

Since AMIP analysis always indicates which data points have high (approximate) influence, our methods allow researchers not only the chance to check that the approximation worked on their own sample, but to understand what—if anything—makes these data points special. Investigating influential points may provide insight into the way in which a given inferential procedure is using the finite-sample information to generate claims about the population parameters. In addition, in cases when this sensitivity is undesirable, it may be fruitful to develop new statistical methods to ameliorate it. It seems particularly important to develop these methods in view of the actual goals and uses of economics research, rather than relying on a classical resampling paradigm that bears little resemblance to the practice of applied social science.

References

- Angelucci, M. and De Giorgi, G. (2009). Indirect effects of an aid program: How do cash transfers affect ineligibles' consumption? *American Economic Review*, 99(1):486–508.
- Angelucci, M., Karlan, D., and Zinman, J. (2015). Microcredit impacts: Evidence from a randomized microcredit program placement experiment by Compartamos Banco. *American Economic Journal: Applied Economics*, 7(1):151–82.
- Attanasio, O., Augsburg, B., De Haas, R., Fitzsimons, E., and Harmgart, H. (2015). The impacts of microfinance: Evidence from joint-liability lending in Mongolia. *American Economic Journal: Applied Economics*, 7(1):90–122.
- Augsburg, B., De Haas, R., Harmgart, H., and Meghir, C. (2015). The impacts of microcredit: Evidence from Bosnia and Herzegovina. *American Economic Journal: Applied Economics*, 7(1):183–203.
- Banerjee, A., Duflo, E., Glennerster, R., and Kinnan, C. (2015). The miracle of microfinance? Evidence from a randomized evaluation. *American Economic Journal: Applied Economics*, 7(1):22–53.
- Baydin, A., Pearlmutter, B., Radul, A., and Siskind, J. (2017). Automatic differentiation in machine learning: A survey. *The Journal of Machine Learning Research*, 18(1):5595–5637.
- Blei, D., Kucukelbir, A., and McAuliffe, J. (2017). Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., Brubaker, M., Guo, J., Li, P., and Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Chatterjee, S. and Hadi, A. (1986). Influential observations, high leverage points, and outliers in linear regression. *Statistical Science*, 1(3):379–393.
- Chen, X., Tamer, E., and Torgovitsky, A. (2011). Sensitivity analysis in semiparametric likelihood models. *Cowles Foundation discussion paper*.
- Crépon, B., Devoto, F., Duflo, E., and Parienté, W. (2015). Estimating the impact of microcredit on those who take it up: Evidence from a randomized experiment in Morocco. *American Economic Journal: Applied Economics*, 7(1):123–50.
- De Bruijn, N. (1981). *Asymptotic methods in analysis*, volume 4. Courier Corporation.

- Finkelstein, A., Taubman, S., Wright, B., Bernstein, M., Gruber, J., Newhouse, J., Allen, H., Baicker, K., and Group, O. H. S. (2012). The Oregon health insurance experiment: Evidence from the first year. *The Quarterly Journal of Economics*, 127(3):1057–1106.
- Giordano, R., Broderick, T., and Jordan, M. I. (2018). Covariances, robustness and variational Bayes. *The Journal of Machine Learning Research*, 19(1):1981–2029.
- Giordano, R., Jordan, M. I., and Broderick, T. (2019a). A higher-order Swiss army infinitesimal jackknife. *arXiv preprint arXiv:1907.12116*.
- Giordano, R., Stephenson, W., Liu, R., Jordan, M. I., and Broderick, T. (2019b). A Swiss army infinitesimal jackknife. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1139–1147. PMLR.
- Gustafson, P. (2000). Local robustness in Bayesian analysis. In *Robust Bayesian Analysis*, pages 71–88. Springer.
- Hampel, F. (1974). The influence curve and its role in robust estimation. *Journal of the American Statistical Association*, 69(346):383–393.
- Hampel, F. (1986). *Robust statistics: the approach based on influence functions*, volume 196. Wiley-Interscience.
- Hansen, L. and Sargent, T. (2008). *Robustness*. Princeton University Press.
- He, X., Jurečková, J., Koenker, R., and Portnoy, S. (1990). Tail behavior of regression estimators and their breakdown points. *Econometrica: Journal of the Econometric Society*, pages 1195–1214.
- Huber, P. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability: Weather modification*, volume 5, page 221. Univ of California Press.
- Huber, P. (1981). *Robust Statistics*. John Wiley & Sons, New York.
- Karlan, D. and Zinman, J. (2011). Microcredit in theory and practice: Using randomized credit scoring for impact evaluation. *Science*, 332(6035):1278–1284.
- Krantz, S. and Parks, H. (2012). *The implicit function theorem: History, theory, and applications*. Springer Science & Business Media.
- Kucukelbir, A., Tran, D., Ranganath, R., Gelman, A., and Blei, D. (2017). Automatic differentiation variational inference. *The Journal of Machine Learning Research*, 18(1):430–474.

- Leamer, E. (1984). Global sensitivity results for generalized least squares estimates. *Journal of the American Statistical Association*, 79(388):867–870.
- Leamer, E. (1985). Sensitivity analyses would help. *The American Economic Review*, 75(3):308–313.
- Maclaurin, D., Duvenaud, D., and Adams, R. (2015). Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, volume 238.
- Masten, M. and Poirier, A. (2020). Inference on breakdown frontiers. *Quantitative Economics*, 11(1):41–111.
- Meager, R. (2019). Understanding the average impact of microcredit expansions: A Bayesian hierarchical analysis of seven randomized experiments. *American Economic Journal: Applied Economics*, 11(1):57–91.
- Meager, R. (2020). Aggregating distributional treatment effects: A Bayesian hierarchical analysis of the microcredit literature. *LSE working paper*.
- Mosteller, F. and Tukey, J. (1977). *Data Analysis and Regression: A Second Course In Statistics*. Pearson, USA.
- Reeds, J. (1976). *On the definition of von Mises functionals*. PhD thesis, Statistics, Harvard University.
- Rigollet, P. (Spring 2015). Course 18.s997 High Dimensional Statistics. Massachusetts Institute of Technology: MIT OpenCourseWare, <https://ocw.mit.edu/>. Accessed: 2021-08-24.
- Saltelli, A. (2004). Global sensitivity analysis: An introduction. In *Proc. 4th International Conference on Sensitivity Analysis of Model Output (SAMO '04)*, pages 27–43.
- Serfling, R. (2009). *Approximation theorems of mathematical statistics*, volume 162. John Wiley & Sons.
- Sobol, I. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1-3):271–280.
- Stefanski, L. and Boos, D. (2002). The calculus of M-estimation. *The American Statistician*, 56(1):29–38.
- Tarozzi, A., Desai, J., and Johnson, K. (2015). The impacts of microcredit: Evidence from Ethiopia. *American Economic Journal: Applied Economics*, 7(1):54–89.
- Van der Vaart, A. (2000). *Asymptotic statistics*, volume 3. Cambridge University Press.

- Wilson, A., Kasy, M., and Mackey, L. (2020). Approximate cross-validation: Guarantees for model assessment and selection. In *International Conference on Artificial Intelligence and Statistics*, pages 4530–4540. PMLR.
- Young, A. (2019). Consistency without inference: Instrumental variables in practical application. <http://personal.lse.ac.uk/YoungA/CWOI.pdf>. Accessed: 2020-11-27.
- Yu, B. (2013). Stability. *Bernoulli*, 19(4):1484–1500.

Appendix A Detailed Proofs

Lemma 1. Let χ_1, \dots, χ_N be real-valued scalars with $\frac{1}{N} \sum_{n=1}^N \chi_n^2 = 1$. Then $\max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N |\vec{w}_n - 1| \chi_n \leq \sqrt{\alpha}$.

Proof. Without loss of generality, let the χ_n be unique (if they are not, add an arbitrarily small amount of jitter to break ties), and let $q_{1-\alpha}$ denote their $[(1-\alpha)N]$ -th largest value. The maximum $\max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N |\vec{w}_n - 1| \chi_n$ is achieved at \vec{w} which sets to zero the weights of all $\{n : \chi_n \geq q_{1-\alpha}\}$, so

$$\max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N |\vec{w}_n - 1| \chi_n = \frac{1}{N} \sum_{n=1}^N \mathbb{I}(\chi_n \geq q_{1-\alpha}) \chi_n.$$

Let \hat{F}_χ denote the empirical distribution on χ_n conditional on the data d_n , and note that $q_{1-\alpha}$ is fixed in \hat{F}_χ . Applying Cauchy-Schwartz to the preceding display with the distribution \hat{F}_χ gives

$$\frac{1}{N} \sum_{n=1}^N \mathbb{I}(\chi_n \geq q_{1-\alpha}) \chi_n \leq \sqrt{\frac{1}{N} \sum_{n=1}^N \mathbb{I}(\chi_n \geq q_{1-\alpha})^2} \sqrt{\frac{1}{N} \sum_{n=1}^N \chi_n^2} = \sqrt{\frac{\lfloor N\alpha \rfloor}{N}} \leq \sqrt{\alpha},$$

since $\frac{1}{N} \sum_{n=1}^N \chi_n^2 = 1$ and at most $\lfloor \alpha N \rfloor$ points are greater than $q_{1-\alpha}$. \square

The following lemma satisfies Condition 1 of Giordano et al. (2019b).

Lemma 2. Let $W_\alpha^* := \{\vec{1} + t(\vec{w} - \vec{1}) : \vec{w} \in W_\alpha, t \in [0, 1]\}$ Under Assumption 3,

$$\begin{aligned} \max_{\vec{w} \in W_\alpha^*} \sup_{\theta \in \Omega_\theta} \left\| \frac{1}{N} \sum_{n=1}^N (\vec{w}_n - 1) G(\theta, d_n) \right\|_1 &\leq \sqrt{D} C_{gh} \sqrt{\alpha} \quad \text{and} \\ \max_{\vec{w} \in W_\alpha^*} \sup_{\theta \in \Omega_\theta} \left\| \frac{1}{N} \sum_{n=1}^N (\vec{w}_n - 1) H(\theta, d_n) \right\|_1 &\leq \sqrt{D} C_{gh} \sqrt{\alpha}. \end{aligned}$$

Proof. We prove the result for $G(\theta, d_n)$; the proof for $H(\theta, d_n)$ follows analogously. By the triangle inequality and the relationship between $\|\cdot\|_2$ and $\|\cdot\|_1$,

$$\left\| \frac{1}{N} \sum_{n=1}^N (\vec{w}_n - 1) G(\theta, d_n) \right\|_1 \leq \sqrt{D} C_{gh} \frac{1}{N} \sum_{n=1}^N |\vec{w}_n - 1| \frac{\|G(\theta, d_n)\|_2}{C_{gh}}.$$

Apply Lemma 1 with $\chi_n := \frac{\|G(\theta, d_n)\|_2}{C_{gh}}$ to control the maximum of the sum over W_α . Finally, the results extends to W_α^* since

$$\max_{\vec{w} \in W_\alpha^*} \frac{1}{N} \sum_{n=1}^N |\vec{w}_n - 1| \chi_n = \max_{t \in [0, 1]} \max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N |t(\vec{w}_n - 1)| \chi_n = \max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N |(\vec{w}_n - 1)| \chi_n.$$

\square

We need the following lemma to extend the result of Giordano et al. (2019b), Theorem 1 to smooth functions.

Lemma 3. *Let Assumptions 3 and 4 hold. For sufficiently small α , there exists a constant $C_b < \infty$ such that, for any $a \in \mathbb{R}^N$,*

$$\max_{\vec{w} \in W_\alpha^*} \left\| \left(\left. \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \right|_{\vec{w}} - \left. \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \right|_{\vec{1}} \right) a \right\|_2 \leq C_b \frac{\|a\|_2}{\sqrt{N}} \sqrt{\alpha}.$$

Proof. As in the proof of Theorem 1, for the remainder of the proof assume that $\alpha \leq \frac{\Delta^2}{DC_{gh}^2}$, and observe that Assumptions 1-5 and Condition 1 of Giordano et al. (2019b) are satisfied. For the duration of this proof, define the shorthand notation

$$H(\vec{w}) := \frac{1}{N} \sum_{n=1}^N \vec{w}_n H(\hat{\theta}(\vec{w}), d_n) \quad \text{and} \quad G(\vec{w}) := \frac{1}{N} \sum_{n=1}^N a_n G(\hat{\theta}(\vec{w}), d_n).$$

Then, by the indicated results from Giordano et al. (2019b),

$$\begin{aligned} & \left\| \left(\left. \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \right|_{\vec{w}} - \left. \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \right|_{\vec{1}} \right) a \right\|_2 \\ &= \left\| -H(\vec{w})^{-1} G(\vec{w}) + H(\vec{1})^{-1} G(\vec{1}) \right\|_2 \quad (\text{Proposition 4}) \\ &\leq \left\| -(H(\vec{w})^{-1} - H(\vec{1})^{-1}) G(\vec{w}) \right\|_2 + \left\| H(\vec{1})^{-1} (G(\vec{1}) - G(\vec{w})) \right\|_2 \\ &\leq \left\| -(H(\vec{w})^{-1} - H(\vec{1})^{-1}) G(\vec{w}) \right\|_2 + C_{op} \delta. \quad (\text{Condition 1, Assumption 2}) \end{aligned}$$

Then,

$$\begin{aligned} & \left\| (H(\vec{w})^{-1} - H(\vec{1})^{-1}) G(\vec{w}) \right\|_2 \\ &= \left\| H(\vec{w})^{-1} (H(\vec{1}) - H(\vec{w})) H(\vec{1})^{-1} G(\vec{w}) \right\|_2 \\ &\leq 2C_{op}^2 \left\| (H(\vec{1}) - H(\vec{w})) G(\vec{w}) \right\|_2 \quad (\text{Assumption 2, Lemma 6}) \\ &\leq 2C_{op}^2 \sqrt{D} (1 + DC_w L_h C_{op}) \delta \|G(\vec{w})\|_2 \quad (\text{Lemma 5, Matrix norms}) \\ &= 2C_{op}^2 \sqrt{D} (1 + DC_w L_h C_{op}) \delta \left\| \frac{1}{N} \sum_{n=1}^N a_n G(\hat{\theta}(\vec{w}), d_n) \right\|_2 \\ &\leq 2C_{op}^2 \sqrt{D} (1 + DC_w L_h C_{op}) \delta C_{gh} \frac{\|a\|_2}{\sqrt{N}}. \quad (\text{Assumption 3, Cauchy-Schwartz}) \end{aligned}$$

Combining, and using our Lemma 2 to give $\delta = \sqrt{D} C_{gh} \sqrt{\alpha}$, gives the desired result. \square

Proof of Theorem 1. For the duration of the proof, define the linear approximation $\hat{\theta}^{\text{lin}}(\vec{w}) := \hat{\theta} + \left. \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \right|_{\vec{1}} (\vec{w} - \vec{1})$. Assumption 3 is equivalent to Assumptions 1-4 of Giordano et al. (2019b), and Lemma 2 satisfies Condition 1 of Giordano et al. (2019b) with $\delta = \sqrt{D} C_{gh} \sqrt{\alpha}$. Assumption 5 of Giordano et al. (2019b) is satisfied for W_α with $C_w = 1$. Define, as in Giordano et al. (2019b), $C_{IJ} := 1 + DL_h C_{op}$ and $\Delta := \min \left\{ \Delta_\theta C_{op}^{-1}, \frac{1}{2} C_{op}^{-1} C_{IJ}^{-1} \right\}$. So Lemma 3 and Theorem 1 of Giordano et al.

(2019b) give, respectively, that

$$\max_{\vec{w} \in W_\alpha^*} \left\| \hat{\theta}(\vec{w}) - \hat{\theta} \right\|_2 \leq C_{op} \sqrt{DC_{gh}} \sqrt{\alpha} \quad \text{and} \quad (27)$$

$$\alpha \leq \frac{\Delta^2}{DC_{gh}^2} \quad \Rightarrow \quad \max_{\vec{w} \in W_\alpha^*} \left\| \hat{\theta}^{\text{lin}}(\vec{w}) - \hat{\theta}(\vec{w}) \right\|_2 \leq 2C_{op}^2 C_{IJ} DC_{gh}^2 \alpha. \quad (28)$$

For the remainder of the proof assume that $\alpha \leq \frac{\Delta^2}{DC_{gh}^2}$ so that Eq. 28 applies.

For any $\vec{w} \in W_\alpha$, define $\omega(t) := \vec{1} + t(\vec{w} - \vec{1}) \in W_\alpha^*$. By the fundamental theorem of calculus,

$$\phi(\hat{\theta}(\vec{w}), \vec{w}) - \hat{\phi} = \int_0^1 \frac{d\phi(\omega(t))}{dt} \Big|_t dt = \int_0^1 \left(\frac{d\phi(\omega(t))}{dt} \Big|_t - \frac{d\phi(\omega(t))}{dt} \Big|_1 \right) dt + \frac{d\phi(\omega(t))}{dt} \Big|_1. \quad (29)$$

where, by the chain rule,

$$\frac{d\phi(\omega(t))}{dt} \Big|_t = \frac{\partial \phi(\theta, \omega(t))}{\partial \theta^T} \Big|_{\hat{\theta}(\omega(t))} \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\omega(t)} (\vec{w} - \vec{1}) + \frac{\partial \phi(\hat{\theta}(\omega(t)), \vec{w})}{\partial \vec{w}^T} \Big|_{\omega(t)} (\vec{w} - \vec{1}).$$

It will be useful to adopt a specific “big O” notation for the remainder of the proof, by which we mean the following. If we write $x = O(\sqrt{\alpha})$ for some quantity x , we mean that there exists a constant C , available as a closed-form function of constants defined in Assumptions 3 and 4, such that $x \leq C\sqrt{\alpha}$ for all $\alpha \leq \frac{\Delta^2}{DC_{gh}^2}$. An analogous notation meaning is given to $x = O(\alpha)$. This “big O” notation can be manipulated in the usual ways (De Bruijn, 1981).

To begin with, by definition of W_α , we have $\max_{\vec{w} \in W_\alpha} \frac{1}{N} \sum_{n=1}^N (\vec{w}_n - 1)^2 = \frac{\lfloor \alpha N \rfloor}{N} \leq \alpha$, so $\max_{\vec{w} \in W_\alpha} \left\| (\vec{w} - \vec{1}) / \sqrt{N} \right\|_2 \leq \sqrt{\alpha}$.

Next, observe that Eqs. 27 and 28 together imply that

$$\max_{\vec{w} \in W_\alpha^*} \left\| \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\vec{1}} (\vec{w} - \vec{1}) \right\|_2 \leq \max_{\vec{w} \in W_\alpha^*} \left\| \hat{\theta}^{\text{lin}}(\vec{w}) - \hat{\theta}(\vec{w}) \right\|_2 + \max_{\vec{w} \in W_\alpha^*} \left\| \hat{\theta}(\vec{w}) - \hat{\theta} \right\|_2 = O(\sqrt{\alpha}).$$

By Lemma 3 below, we have that

$$\max_{t \in [0,1]} \max_{\vec{w} \in W_\alpha^*} \left\| \left(\frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\omega(t)} - \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\vec{1}} \right) (\vec{w} - \vec{1}) \right\|_2 \leq C_b \frac{\left\| \vec{w} - \vec{1} \right\|_2}{\sqrt{N}} \sqrt{\alpha} = O(\alpha).$$

Combining the previous two displays gives, by the triangle inequality, that

$$\max_{t \in [0,1]} \max_{\vec{w} \in W_\alpha^*} \left\| \frac{d\hat{\theta}(\vec{w})}{d\vec{w}^T} \Big|_{\omega(t)} (\vec{w} - \vec{1}) \right\|_2 = O(\sqrt{\alpha}).$$

Finally, by the Lipschitz property of the partial derivatives in Assumption 4, we

have that

$$\begin{aligned} \max_{t \in [0,1]} \left\| \frac{\partial \phi(\theta, \omega(t))}{\partial \theta} \Big|_{\hat{\theta}(\omega(t))} - \frac{\partial \phi(\theta, \vec{1})}{\partial \theta} \Big|_{\hat{\theta}} \right\|_2 &= O(\sqrt{\alpha}) \quad \text{and} \\ \max_{t \in [0,1]} \sqrt{N} \left\| \frac{\partial \phi(\hat{\theta}(\omega(t)), \vec{w})}{\partial \vec{w}} \Big|_{\omega(t)} - \frac{\partial \phi(\hat{\theta}, \vec{w})}{\partial \vec{w}} \Big|_{\vec{1}} \right\|_2 &= O(\sqrt{\alpha}). \end{aligned}$$

Again, the triangle inequality with the boundedness of the partial derivatives of ϕ at $\vec{w} = \vec{1}$ implies

$$\max_{t \in [0,1]} \left\| \frac{\partial \phi(\theta, \omega(t))}{\partial \theta} \Big|_{\hat{\theta}(\omega(t))} \right\|_2 \quad \text{and} \quad \max_{t \in [0,1]} \sqrt{N} \left\| \frac{\partial \phi(\hat{\theta}(\omega(t)), \vec{w})}{\partial \vec{w}} \Big|_{\omega(t)} \right\|_2 = O(\sqrt{\alpha}).$$

Combining the above results gives that

$$\max_{t \in [0,1]} \left\| \frac{d\phi(\omega(t))}{dt} \Big|_t \right\|_2 = O(\sqrt{\alpha}) \quad \text{and} \quad \max_{t \in [0,1]} \left\| \frac{d\phi(\omega(t))}{dt} \Big|_t - \frac{d\phi(\omega(t))}{dt} \Big|_1 \right\|_2 = O(\alpha),$$

from which the desired conclusion follows by Eq. 29. \square