

NBER WORKING PAPER SERIES

INFERENCE ON WINNERS

Isaiah Andrews
Toru Kitagawa
Adam McCloskey

Working Paper 25456
<http://www.nber.org/papers/w25456>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138
January 2019

We would like to thank Tim Armstrong, Stephane Bonhomme, Gregory Cox, Aureo de Paula, Patrick Kline, Hannes Leeb, Anna Mikusheva, Mikkel Plagborg-Moller, Jose Luis Montiel Olea, Jack Porter, Adam Rosen, Frank Schoerfheide, Jesse Shapiro, and participants at numerous seminars and conferences for helpful comments. Andrews gratefully acknowledges financial support from the NSF under grant number 1654234. Kitagawa gratefully acknowledges financial support from the ESRC through the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001) and the European Research Council (Starting grant No. 715940). Initial version posted May 10, 2018. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Isaiah Andrews, Toru Kitagawa, and Adam McCloskey. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Inference on Winners
Isaiah Andrews, Toru Kitagawa, and Adam McCloskey
NBER Working Paper No. 25456
January 2019
JEL No. C12,C13

ABSTRACT

Many empirical questions can be cast as inference on a parameter selected through optimization. For example, researchers may be interested in the effectiveness of the best policy found in a randomized trial, or the best-performing investment strategy based on historical data. Such settings give rise to a winner's curse, where conventional estimates are biased and conventional confidence intervals are unreliable. This paper develops optimal confidence sets and median-unbiased estimators that are valid conditional on the parameter selected and so overcome this winner's curse. If one requires validity only on average over target parameters that might have been selected, we develop hybrid procedures that combine conditional and projection confidence sets to offer further performance gains relative to existing alternatives.

Isaiah Andrews
Department of Economics
Harvard University
Littauer M18
Cambridge, MA 02138
and NBER
iandrews@fas.harvard.edu

Adam McCloskey
Department of Economics
University of Colorado at Boulder
256 UCB
Boulder, CO 80309
adam.mccloskey@colorado.edu

Toru Kitagawa
Department of Economics
University College London
Gower Street
London WC1E 6BT
United Kingdom
and Institute For Fiscal Studies - IFS
t.kitagawa@ucl.ac.uk

A data appendix is available at <http://www.nber.org/data-appendix/w25456>

Inference on Winners*

Isaiah Andrews[†] Toru Kitagawa[‡] Adam McCloskey[§]

December 31, 2018

Abstract

Many empirical questions can be cast as inference on a parameter selected through optimization. For example, researchers may be interested in the effectiveness of the best policy found in a randomized trial, or the best-performing investment strategy based on historical data. Such settings give rise to a winner’s curse, where conventional estimates are biased and conventional confidence intervals are unreliable. This paper develops optimal confidence sets and median-unbiased estimators that are valid conditional on the parameter selected and so overcome this winner’s curse. If one requires validity only on average over target parameters that might have been selected, we develop hybrid procedures that combine conditional and projection confidence sets to offer further performance gains relative to existing alternatives.

KEYWORDS: WINNER’S CURSE, SELECTIVE INFERENCE

JEL CODES: C12, C13

1 Introduction

A wide range of empirical questions involve inference on target parameters selected through optimization over a finite set. In a randomized trial considering multiple treatments, for

*We would like to thank Tim Armstrong, Stéphane Bonhomme, Gregory Cox, Áureo de Paula, Patrick Kline, Hannes Leeb, Anna Mikusheva, Mikkel Plagborg-Møller, José Luis Montiel Olea, Jack Porter, Adam Rosen, Frank Schoerfheide, Jesse Shapiro, and participants at numerous seminars and conferences for helpful comments. Andrews gratefully acknowledges financial support from the NSF under grant number 1654234. Kitagawa gratefully acknowledges financial support from the ESRC through the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001) and the European Research Council (Starting grant No. 715940). Initial version posted May 10, 2018.

[†]Department of Economics, Harvard University, iandrews@fas.harvard.edu

[‡]CeMMAP and Department of Economics, University College London, t.kitagawa@ucl.ac.uk

[§]Department of Economics, University of Colorado, adam.mccloskey@colorado.edu

instance, one might want to learn about the true average effect of the treatment that performed best in the experiment. In finance, one might want to learn about the expected return of the trading strategy that performed best in a backtest. Perhaps less obviously, in threshold regression or tipping point models, researchers first estimate the location of a threshold by minimizing the sum of squared residuals and then seek to estimate the magnitude of the discontinuity taking the estimated threshold as given.

Estimators that do not account for data-driven selection of the target parameters can be badly biased, and conventional t-statistic-based confidence intervals may severely under-cover. To illustrate the problem, consider inference on the true average effect of the treatment that performed best in a randomized trial.¹ Since it ignores the data-driven selection of the treatment of interest, the conventional estimate for this average effect will be biased upwards. Similarly, the conventional confidence interval will under-cover, particularly when the number of treatments considered is large. This gives rise to a form of winner’s curse, where follow-up trials will be systematically disappointing relative to what we would expect based on conventional estimates and confidence sets. This form of winner’s curse has previously been discussed in contexts including genome-wide association studies (e.g. Zhong and Prentice, 2009; Ferguson et al., 2013) and online A/B tests (Lee and Shen, 2018).

This paper develops estimators and confidence sets that eliminate these biases and inference failures. There are two distinct perspectives from which to consider bias and coverage. The first conditions on the target parameter selected, for example on the identity of the best-performing treatment, while the second is unconditional and averages over possible target parameters. As we discuss in the next section, conditional validity is more demanding but may be desirable in some settings, for example when one wants to ensure validity conditional on the recommendation made to a policy maker. Both perspectives differ from inference on the effectiveness of the “true” best treatment, as in e.g. Chernozhukov et al. (2013) and Rai (2018), in that we consider inference on the

¹Such a scenario seems to be empirically relevant, as a number of recently published randomized trials in economics either were designed with the intent of recommending a policy or represent a direct collaboration with a policy maker. For example, Khan et al. (2016) assesses how incentives for property tax collectors affect tax revenues in Pakistan, Banerjee et al. (2018) evaluates the efficacy of providing information cards to potential recipients of Indonesia’s *Raskin* programme, and Duflo et al. (2018) collaborates with the Gujarat Pollution Control Board (an Indian regulator tasked with monitoring industrial emissions in the state) to evaluate how more frequent but randomized inspection of plants performs relative to discretionary inspection. Baird et al. (2016) finds that deworming Kenyan children had substantial beneficial effects on their health and labor market outcomes into adulthood, and Björkman Nyqvist and Jayachandran (2017) finds that providing parenting classes to Ugandan mothers has a greater impact on child outcomes than targeting these classes at fathers.

effectiveness of the (observed) best-performing treatment in the experiment rather than the (unobserved) best-performing treatment in the population.²

Considering first conditional inference, we derive optimal unbiased and equal-tailed confidence sets. Our results build on the rapidly growing literature on selective inference (e.g. Harris et al. (2016); Lee et al. (2016); Tian and Taylor (2016); Fithian et al. (2017)), which derives optimal conditional confidence sets in a range of other settings. We further observe that the results of Pfanzagl (1994) imply optimal median-unbiased estimators for conditional settings, which does not appear to have been previously noted in the selective inference literature. Hence, for settings where conditional validity is desired, we propose optimal inference procedures that eliminate the winner’s curse noted above. We further show that in cases where this winner’s curse does not arise (for instance because one treatment considered is vastly better than the others) our conditional procedures coincide with conventional ones. Hence, our corrections do not sacrifice efficiency in such cases.

A common alternative remedy for the biases we consider is sample splitting. In settings with independent observations, choosing the target parameter using the first part of the data and constructing estimates and confidence sets using the second part ensures unbiasedness of estimates and validity of conventional confidence sets conditional on the target parameter. Such conventional split-sample procedures can have undesirable properties, however. In particular, the target parameter is generally more variable than if constructed using the full data. Moreover, since only the second part of the data is used for inference, Fithian et al. (2017) show that conventional split-sample procedures are inadmissible within the class of procedures with the same target parameter. Motivated by this result, in the supplement to the paper we develop computationally tractable confidence sets and estimators that dominate conventional sample-splitting.

We next turn to unconditional inference. One approach to constructing unconditional confidence sets is projection, applied in various forms and settings by e.g. Romano and Wolf (2005), Berk et al. (2013), and Kitagawa and Tetenov (2018a). To obtain a projection confidence set, we form a simultaneous confidence band for all potential target parameters and take the implied set of values for the target parameter of interest. The resulting confidence sets have correct unconditional coverage but, unlike our conditional intervals, are wider than conventional confidence sets even when the latter are valid. On the other hand, we find in simulations that projection intervals outperform conditional intervals in

²See Dawid (1994) for an early discussion of this distinction, and an argument in favor of inference on the best-performing treatment in the experiment.

cases where there is substantial randomness in the target parameter, e.g. when there is not a clear best treatment.

Since neither conditional nor projection intervals perform well in all cases, we introduce hybrid confidence sets that combine conditioning and projection. These maintain most of the good performance of our conditional confidence intervals in cases for which the winner’s curse does not arise but are subsets of (conservative) projection intervals by construction, limiting their maximal under-performance relative to projection confidence sets. We also introduce hybrid estimators that allow a controlled degree of bias while limiting the deviation from the conventional estimator.

We derive our main results in the context of a finite-sample normal model with an unknown mean vector and a known covariance matrix. This model can be viewed as an asymptotic approximation to non-normal finite sample problems where the optimal policy may not be obvious from the data. To formalize this connection, in the supplement to the paper we show that the procedures we derive are uniformly asymptotically valid over large classes of data-generating processes.

Since we are not aware of any other full-sample procedures that ensure validity conditional on the target parameter, our simulations focus on unconditional performance. The simulation designs are based on an empirical welfare maximization application from Kitagawa and Tetenov (2018b) and a threshold regression application from Card et al. (2008). In both settings, we find that while our conditional procedures exhibit good unconditional performance in cases where the objective function determining the target parameter has a well-separated optimum, their unconditional performance can be poor in other cases. By contrast, our hybrid procedures perform quite well: hybrid confidence sets are shorter than the previously available alternative (projection intervals) in all specifications, and are shorter than conditional intervals in all but the well-separated case (where they are nearly the same). Hybrid estimators eliminate nearly all the bias of conventional estimators, and are less dispersed than our exactly median unbiased estimators. These results show that while optimal conditional performance is attainable, conditional validity can come at the cost of unconditional performance. By combining conditional and projection approaches, our hybrid procedures yield better performance than either and offer a substantial improvement over existing alternatives.

While most of our simulation results focus on comparing our full-sample conditional and hybrid approaches to existing full-sample alternatives, Card et al. (2008) originally conducted inference based on a conventional split-sample approach. Hence, our simulations

based on Card et al. (2008) also compare conventional sample splitting procedures to our improved split-sample ones. We similarly find substantial performance improvements in these split-sample settings.

In this paper we focus on frequentist inference, and in particular on ensuring coverage and controlling bias under all parameter values. If one instead takes a Bayesian perspective then, as discussed by e.g. Dawid (1994), the selection issue does not arise since Bayesian inference conditions on the data and thus on any form of data-driven selection. One way to interpret this point is that e.g. the Bayes posterior median is median unbiased for the true parameter value under the prior. As highlighted by Dawid (1994), however, this property hinges crucially on the specification of the prior. If we consider frequentist performance in cases where the data are generated in a manner inconsistent with the prior, Bayes procedures may have large biases. In settings where we observe independent estimates for a large number of different parameters and are willing to assume that these parameters are drawn from some common unknown distribution, we can avoid this issue by adopting an empirical Bayes approach and estimating the prior (see Efron, 2011; Ferguson et al., 2013). Many settings, including our empirical welfare and threshold regression examples, lack this structure however, rendering this approach inapplicable.

It is important to emphasize that we take the rule for selecting the target parameter as given. In policy-evaluation contexts, for example, our goal is to evaluate the effectiveness of recommended policies taking the rule for selecting a recommendation as given, rather than to improve the rule. There are a number of reasons why valid confidence sets and median-unbiased estimates are of interest in such settings. One might be interested in understanding the true effectiveness of a selected policy for scientific reasons. Alternatively, one might want to assess uncertainty about the effect of a new policy for forecasting and risk management purposes. Finally, after a policy has been implemented or a follow-up trial conducted, one may want to test whether observed differences in efficacy can be explained solely by the winner’s curse.

This paper is related to the literature on tests of superior predictive performance (e.g. White (2000); Hansen (2005); Romano and Wolf (2005)). This literature studies the problem of testing whether some strategy or policy beats a benchmark, while we consider the complementary question of inference on the effectiveness of the estimated “best” policy. Our conditional inference results combine naturally with the results of this literature, allowing one to condition inference on e.g. rejecting the null hypothesis that no policy outperforms a benchmark.

As mentioned above, our results are also closely related to the growing literature on selective inference. Fithian et al. (2017) describe a general conditioning approach applicable to a wide range of settings, while a rapidly growing literature including e.g. Harris et al. (2016); Lee et al. (2016); Tian and Taylor (2016) works out the details of this approach for a range of settings. Likewise, our analysis of conditional confidence sets examines the implications of the conditional approach in our setting. Our results are also related to the growing literature on unconditional post-selection inference, including e.g. Berk et al. (2013); Bachoc et al. (2017, 2018); Kuchibhotla et al. (2018). This literature considers analogs of our projection confidence sets for inference following model selection.

Beyond the new settings considered, we make two main theoretical contributions relative to the selective and post-selection inference literatures. First, when one only requires unconditional validity, we propose the class of hybrid inference and estimation procedures. We find that hybrid procedures offer large gains in unconditional performance relative both to conditional procedures and to existing unconditional alternatives. Second, for settings where conditional inference is desired, we observe that the same structure used to develop optimal conditional confidence sets also allows construction of optimal quantile unbiased estimators using the results of Pfanzagl (1994).³

In the next section, we begin by introducing the problem we consider and the techniques we propose in the context of a stylized example. Section 3 introduces the normal model in which we develop our main results, and shows how it arises as an asymptotic approximation to empirical welfare maximization and threshold regression examples. Section 4 develops our optimal conditional procedures, discusses their properties, and compares them to sample splitting. Section 5 introduces projection confidence intervals and our hybrid procedures. Finally, Sections 6 and 7 report results for simulations calibrated to empirical welfare maximization and threshold regression applications, respectively. The supplement to the paper collects proofs and other supporting material for the results in the main text, derives a computationally tractable split-sample approach that dominates conventional split-sample inference, shows that the finite sample results developed in the main text translate to uniform asymptotic results over large classes of data generating processes, and provides additional simulation results.

³Our asymptotic results are also novel relative to the literature. In particular, Tibshirani et al. (2018) establish uniform asymptotic validity for conditional confidence sets based on similar ideas to ours, but only under particular local sequences. We impose an analogous restriction for some of our asymptotic results but not others. See the supplement for details and further discussion.

2 A Stylized Example

We begin by illustrating the problem we consider, along with the solutions we propose, in a stylized example based on Manski (2004). In the treatment choice problem of Manski (2004) a treatment rule assigns treatments to subjects based on observable characteristics. Given a social welfare criterion and (quasi-)experimental data, Kitagawa and Tetenov (2018b) propose what they call empirical welfare maximization (EWM), which selects the treatment rule that maximizes the sample analog of the social welfare criterion over a class of candidate rules.

For simplicity suppose there are only two candidate policies: θ_1 corresponding to “treat everyone” and θ_2 corresponding to “treat no one.” Suppose further that our social welfare function is the average of an outcome variable Y . If we have a sample of independent observations $i \in \{1, \dots, n\}$ from a randomized trial where a binary treatment $D_i \in \{0, 1\}$ is randomly assigned to subjects with $Pr\{D_i = 1\} = d$, then as in Kitagawa and Tetenov (2018b) the scaled empirical welfare under (θ_1, θ_2) is

$$(X_n(\theta_1), X_n(\theta_2)) = \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{D_i Y_i}{d}, \frac{1}{\sqrt{n}} \sum_{i=1}^n \frac{(1-D_i) Y_i}{1-d} \right).$$

EWM selects the rule $\hat{\theta} = \operatorname{argmax}_{\theta \in \{\theta_1, \theta_2\}} X_n(\theta)$.⁴

Kitagawa and Tetenov (2018b) show that the welfare from the policy selected by EWM converges to the optimal social welfare at the minimax optimal rate, providing a strong argument for this approach. Even after choosing a policy, we may want estimates and confidence intervals for its implied social welfare in order to learn about the size of the policy impact and communicate with stakeholders. For a fixed policy θ , the empirical welfare $X_n(\theta)$ is unbiased for the true (scaled) social welfare $\mu_n(\theta)$ under the corresponding policy.⁵ By contrast, the empirical welfare of the estimated optimal policy $X_n(\hat{\theta})$ is biased upwards relative to the true social welfare $\mu_n(\hat{\theta})$ since we are more likely to select a given policy when the empirical welfare over-estimates the true welfare. Likewise, confidence sets for $\mu_n(\hat{\theta})$ that ignore estimation of θ may cover $\mu_n(\hat{\theta})$ less often than we intend. This is a form of winner’s curse: estimation error leads us to over-predict the benefits of our chosen policy and to misstate our uncertainty about its effectiveness.

⁴If the summands are instead weighted by sample propensity scores, we obtain Manski’s conditional empirical success rule and the asymptotically optimal rules of Hirano and Porter (2009) with a symmetric loss.

⁵ $X_n(\theta)$ is exactly mean-unbiased and asymptotically median-unbiased.

To simplify the analysis and develop corrected inference procedures, we turn to asymptotic approximations. Under mild conditions the central limit theorem implies that our estimates of social welfare are asymptotically normal:

$$\begin{pmatrix} X_n(\theta_1) - \mu_n(\theta_1) \\ X_n(\theta_2) - \mu_n(\theta_2) \end{pmatrix} \Rightarrow N\left(0, \begin{pmatrix} \Sigma(\theta_1) & \Sigma(\theta_1, \theta_2) \\ \Sigma(\theta_1, \theta_2) & \Sigma(\theta_2) \end{pmatrix}\right), \quad (1)$$

where the asymptotic variance Σ can be consistently estimated while the scaled social welfare μ_n cannot be. To simplify the analysis, for this section only we assume that $\Sigma(\theta_1, \theta_2) = 0$.⁶ Motivated by (1), we abstract from approximation error and assume that we observe

$$\begin{pmatrix} X(\theta_1) \\ X(\theta_2) \end{pmatrix} \sim N\left(\begin{pmatrix} \mu(\theta_1) \\ \mu(\theta_2) \end{pmatrix}, \begin{pmatrix} \Sigma(\theta_1) & 0 \\ 0 & \Sigma(\theta_2) \end{pmatrix}\right)$$

for $\Sigma(\theta_1)$ and $\Sigma(\theta_2)$ known, and that $\hat{\theta} = \operatorname{argmax}_{\theta \in \Theta} X(\theta)$ with $\Theta = \{\theta_1, \theta_2\}$.

As discussed above, $X(\hat{\theta})$ is biased upwards as an estimator of $\mu(\hat{\theta})$. This bias arises both conditional on $\hat{\theta}$ and unconditionally. To see this note that $\hat{\theta} = \theta_1$ if $X(\theta_1) > X(\theta_2)$, where ties occur with probability zero. Conditional on $\hat{\theta} = \theta_1$ and $X(\theta_2)$, $X(\theta_1)$ follows a normal distribution truncated below at $X(\theta_2)$. Since this holds for all $X(\theta_2)$, $X(\theta_1)$ has positive median bias conditional on $\hat{\theta} = \theta_1$.⁷

$$Pr_\mu\left\{X(\hat{\theta}) \geq \mu(\hat{\theta}) \mid \hat{\theta} = \theta_1\right\} > \frac{1}{2} \text{ for all } \mu. \quad (2)$$

Since the same argument holds for $\hat{\theta} = \theta_2$, $\hat{\theta}$ is likewise biased upwards unconditionally:

$$Pr_\mu\left\{X(\hat{\theta}) \geq \mu(\hat{\theta})\right\} > \frac{1}{2} \text{ for all } \mu. \quad (3)$$

Note that (3) differs from (2) in that the target parameter is random. Unsurprisingly given this bias, the conventional confidence set which adds and subtracts a quantile of the standard normal distribution times the standard error need not have correct coverage.

To illustrate these issues, Figure 1 plots the coverage of conventional confidence sets, as well as the median bias of conventional estimates, in an example with $\Sigma(\theta_1) = \Sigma(\theta_2) = 1$. For comparison we also consider cases with ten and fifty policies, $|\Theta| = 10$ and $|\Theta| = 50$, where

⁶One can show that $\Sigma(\theta_1, \theta_2) = -\mu(\theta_1)\mu(\theta_2)$, so this restriction arises naturally if one models μ as shrinking with the sample size to keep it on the same order as sampling uncertainty: $\mu_n = \frac{1}{\sqrt{n}}\mu^*$.

⁷It also has positive mean bias, but we focus on median bias for consistency with our later results.

we again set $\Sigma(\theta)=1$ for all θ and for ease of reporting assume that all the policies other than the first are equally effective: $\mu(\theta_2)=\mu(\theta_3)=\dots=\mu(\theta_{-1})$. The first panel of Figure 1 shows that while the conventional confidence set has reasonable coverage when there are only two policies, its coverage can fall substantially when $|\Theta|=10$ or $|\Theta|=50$.⁸ The second panel shows that the median bias of the conventional estimator $\hat{\mu}=X(\hat{\theta})$, measured as the deviation of the exceedance probability $Pr_{\mu}\{X(\hat{\theta})\geq\mu(\hat{\theta})\}$ from $\frac{1}{2}$, can be quite large. The third panel shows that the same is true when we measure bias as the median of $X(\hat{\theta})-\mu(\hat{\theta})$. In all cases we find that performance is worse when we consider a larger number of policies, as is natural since a larger number of policies allows more scope for selection.

Our results correct these biases. Returning to the case with $|\Theta|=2$ for simplicity, let $F_{TN}(x(\theta_1);\mu(\theta_1),x(\theta_2))$ denote the (truncated normal) distribution function for $X(\theta_1)$ truncated below at $x(\theta_2)$ when the true social welfare for θ_1 is $\mu(\theta_1)$. For fixed $x(\theta_1)>x(\theta_2)$ this function is strictly decreasing in $\mu(\theta_1)$, and for $\hat{\mu}_{\alpha}$ that solves $F_{TN}(X(\theta_1);\hat{\mu}_{\alpha},X(\theta_2))=1-\alpha$, Proposition 1 below shows that

$$Pr_{\mu}\left\{\hat{\mu}_{\alpha}\geq\mu(\hat{\theta})|\hat{\theta}=\theta_1\right\}=\alpha \text{ for all } \mu.$$

Hence, $\hat{\mu}_{\alpha}$ is α -quantile unbiased for $\mu(\hat{\theta})$ conditional on $\hat{\theta}=\theta_1$, and the analogous statement holds conditional on $\hat{\theta}=\theta_2$. Indeed, Proposition 1 shows that $\hat{\mu}_{\alpha}$ is the optimal α -quantile unbiased estimator conditional on $\hat{\theta}$.

Using this result, we can eliminate the biases discussed above. The estimator $\hat{\mu}_{1/2}$ is median unbiased and the equal-tailed confidence interval $CS_{ET}=[\hat{\mu}_{\alpha/2},\hat{\mu}_{1-\alpha/2}]$ has conditional coverage $1-\alpha$, where we say that a confidence set CS has conditional coverage $1-\alpha$ if

$$Pr\left\{\mu(\hat{\theta})\in CS|\hat{\theta}=\theta_j\right\}\geq 1-\alpha \text{ for } j\in\{1,2\} \text{ and all } \mu. \quad (4)$$

While the equal-tailed confidence interval is easy to compute, there are other confidence sets available in this setting. As in Lehmann and Scheffé (1955) and Fithian et al. (2017) it is possible to construct a uniformly most accurate unbiased (UMAU) confidence set, CS_U , conditional on $\hat{\theta}$. To construct CS_U , we collect the parameter values not rejected by a uniformly most powerful unbiased test conditional on $\hat{\theta}$. While straightforward to implement, the exact form of this test is somewhat involved and so is deferred to Section 4 below. The equal-tailed confidence set CS_{ET} is not unbiased, so there is not a clear ranking between CS_{ET} and CS_U .

⁸For example, these could correspond to cases where we consider “treat no one” along with nine or forty nine different treatment assignment rules, respectively.




Figure 1: Performance of conventional procedures in examples with 2, 10, and 50 policies.

The law of iterated expectations implies that CS_{ET} and CS_U have unconditional coverage $1-\alpha$ as well:

$$Pr_\mu\left\{\mu(\hat{\theta}) \in CS\right\} \geq 1-\alpha \text{ for all } \mu. \quad (5)$$

Unconditional coverage is easier to attain, so relaxing the coverage requirement from (4) to (5) may yield tighter confidence sets in some cases. Conditional and unconditional coverage requirements address different questions, however, and which is more appropriate depends on the problem at hand. In the EWM problem, for instance, a policy maker who is told the recommended policy $\hat{\theta}$ along with a confidence interval may want the confidence interval to be valid conditional on the recommendation, which is precisely the conditional coverage requirement (4). In particular, this ensures that if one considers repeated instances in which EWM recommends a particular course of action (e.g. departure from the status quo), reported confidence sets will in fact cover the true effects a fraction $1-\alpha$ of the time. On the other hand, if we only want to ensure that our confidence sets cover the true value with probability at least $1-\alpha$ on average across the distribution of recommendations, it suffices to impose the unconditional requirement (5).

We are unaware of alternative procedures that ensure conditional coverage (4).⁹ For unconditional coverage (5), however, Kitagawa and Tetenov (2018a) propose an unconditional confidence set based on projecting a simultaneous confidence band for μ to obtain a confidence set for $\mu(\hat{\theta})$. In particular, let c_α denote the $1-\alpha$ quantile of $\max_j |\xi_j|$ for $\xi = (\xi_1, \xi_2)' \sim N(0, I_2)$ a two-dimensional standard normal random vector. If we define CS_P as

$$CS_P = \left[Y(\hat{\theta}) - c_\alpha \sqrt{\Sigma(\hat{\theta})}, Y(\hat{\theta}) + c_\alpha \sqrt{\Sigma(\hat{\theta})} \right],$$

this set has correct unconditional coverage (5).

Figure 2 plots the median (unconditional) length of 95% confidence sets CS_{ET} , CS_U , and CS_P , along with the conventional confidence set, again in cases with $|\Theta| \in \{2, 10, 50\}$. We focus on median length, rather than mean length, because the results for Kivaranovic and Leeb (2018) imply that both CS_{ET} and CS_U have infinite expected length.¹⁰ As Figure 2 illustrates, the median lengths of CS_{ET} and CS_U are shorter than the (nonrandom) length of CS_P when $|\mu(\theta_1) - \mu(\theta_{-1})|$ exceeds four, and converges to the length of the

⁹As noted in the introduction and further discussed in Section 4.3 below, split-sample confidence intervals also have conditional coverage but change the definition of $\hat{\theta}$.

¹⁰While Kivaranovic and Leeb (2018) do not consider the behavior of unbiased confidence sets, one can show that the expected length of the level $1-\alpha$ unbiased confidence set is bounded below by that of the level $1-2\alpha$ equal-tailed confidence set.

conventional interval as $|\mu(\theta_1) - \mu(\theta_{-1})|$ tends to infinity. When $|\mu(\theta_1) - \mu(\theta_{-1})|$ is small, on the other hand, CS_{ET} and CS_U can be substantially wider than CS_P . Both features become more pronounced as we increase the number of policies considered, and are still more pronounced for higher quantiles of the length distribution. To illustrate, Figure 3 plots the 95th percentile of the distribution of length in the case with $|\Theta| = 50$ policies, while results for other quantiles and specifications are reported in Section E of the supplement.

In Figure 4 we plot the median absolute error $Med_\mu(|\hat{\mu} - \mu(\hat{\theta})|)$ for different estimators, and find that the median-unbiased estimator likewise exhibits larger median absolute error than the conventional estimator $X(\hat{\theta})$ when $|\mu(\theta_1) - \mu(\theta_{-1})|$ is small.¹¹ This feature is again more pronounced as we increase the number of policies considered, or if we consider higher quantiles as in Section E of the supplement.

Recall that CS_U is the optimal unbiased confidence set, while the endpoints of CS_{ET} are optimal quantile unbiased estimators. So long as we impose correct conditional coverage (4) and unbiasedness, there is therefore no scope to improve unconditional performance. If we instead require only correct unconditional coverage (5), improved performance is possible.

To improve performance, we consider hybrid confidence sets CS_{ET}^H and CS_U^H . As detailed in Section 5.2 below, these confidence sets are constructed analogously to CS_{ET} and CS_U , but further condition on the event that the true social welfare falls in the level $1 - \beta$ projection interval CS_P^β for $\beta < \alpha$. This ensures that the hybrid confidence sets are never longer than the level $1 - \beta$ projection interval, and so both limits the performance deterioration when $|\mu(\theta_1) - \mu(\theta_{-1})|$ is small and ensures that the expected length of hybrid confidence sets is always finite. These hybrid confidence sets have correct unconditional coverage (5), but do not in general have correct conditional coverage (4). By relaxing the conditional coverage requirement, however, we obtain major improvements in unconditional performance, as illustrated in Figure 2. In particular, we see that in the cases with 10 and 50 policies, the hybrid confidence sets have shorter median length than the unconditional interval CS_P for all parameter values considered. The gains relative to conditional confidence sets are large for many parameter values, and are still more pronounced for higher quantiles of the length distribution, as in Figure 3 and Section E of the supplement. In Figure 4 we report results for a hybrid estimation procedure based on a similar approach (detailed in Section 5.3 below), and again find substantial performance improvements.

The improved unconditional performance of the hybrid confidence sets is achieved by

¹¹The proof of Proposition 1 of Kivaranovic and Leeb (2018) implies that the mean absolute error of the median unbiased estimator is infinite.

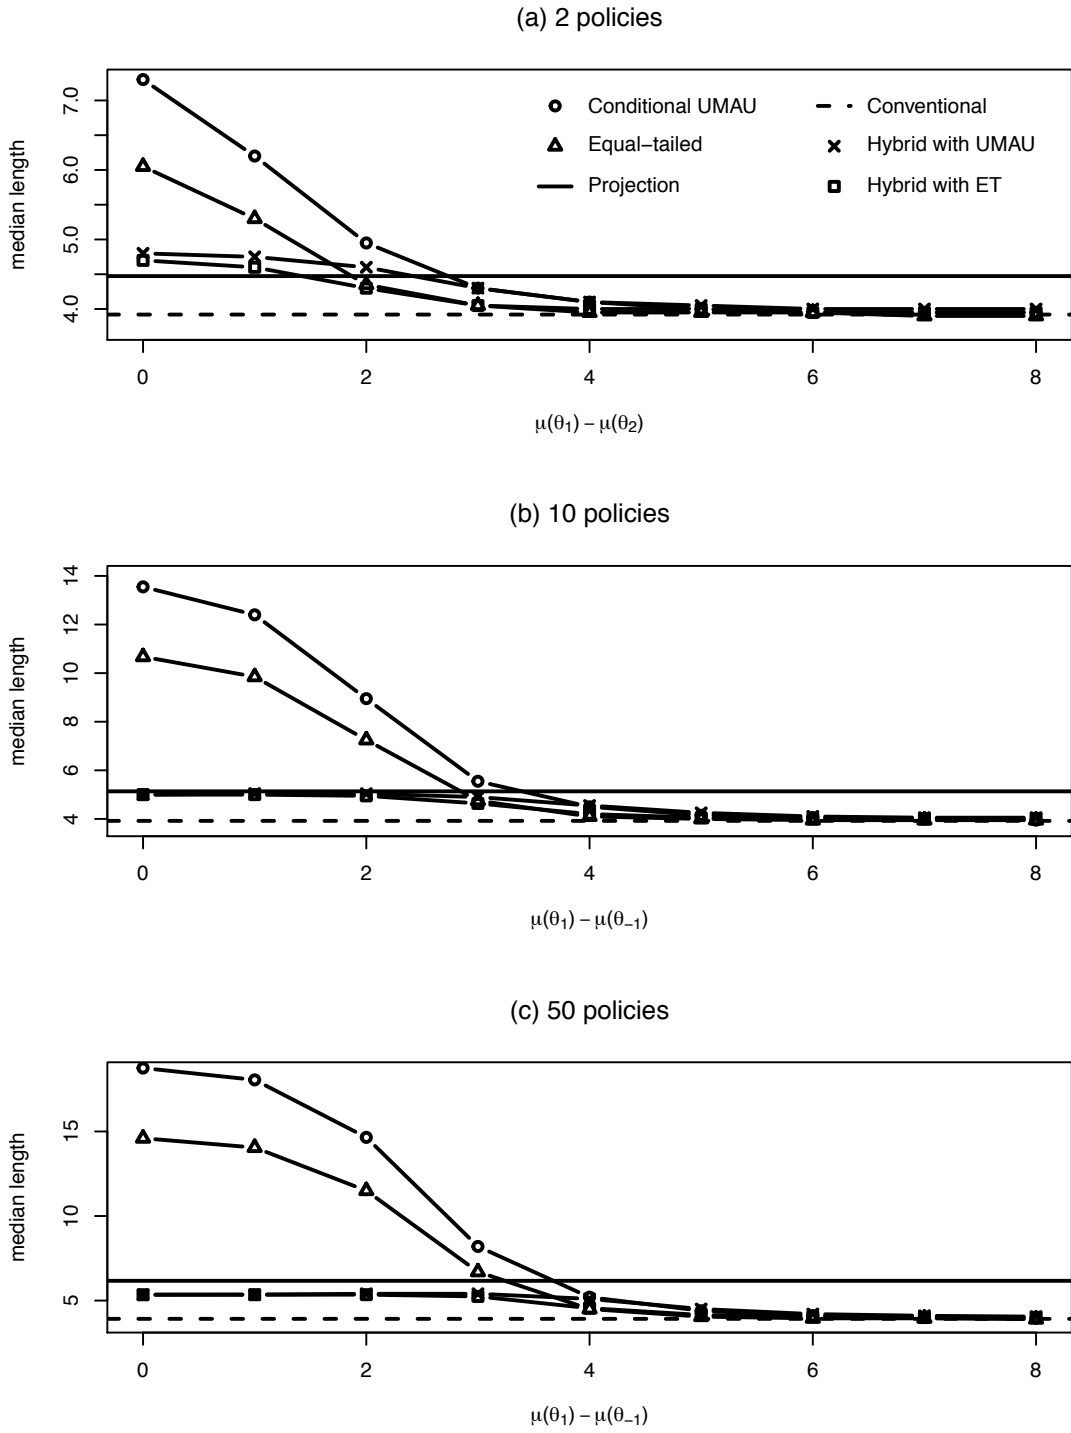


Figure 2: Median length of confidence sets for $\mu(\hat{\theta})$ in cases with 2, 10, and 50 policies.




Figure 3: 95th percentile of length of confidence sets for $\mu(\hat{\theta})$ in case with 50 policies.

requiring only unconditional, rather than conditional, coverage. To illustrate, Figure 5 plots the conditional coverage given $\hat{\theta} = \theta_1$ in the case with two policies. As expected, the conditional intervals have correct conditional coverage, while coverage distortions appear for the hybrid and projection intervals when $\mu(\theta_1) \ll \mu(\theta_2)$. In this case $\hat{\theta} = \theta_2$ with high probability but the data will nonetheless sometimes realize $\hat{\theta} = \theta_1$. Conditional on this event, $X(\theta_1)$ will be far away from $\mu(\theta_1)$ with high probability, so projection and hybrid confidence sets under-cover.

3 Setting

This section introduces our general setting, which extends the stylized example of the previous section in several directions. We assume that we observe normal random vectors $(X(\theta)', Y(\theta)')'$ for $\theta \in \Theta$ where Θ is a finite set, $X(\theta) \in \mathbb{R}^{d_X}$, and $Y(\theta) \in \mathbb{R}$. In particular, for $\Theta = \{\theta_1, \dots, \theta_{|\Theta|}\}$, let $X = (X(\theta_1)', \dots, X(\theta_{|\Theta|}')')'$ and $Y = (Y(\theta_1), \dots, Y(\theta_{|\Theta|}))'$. Then

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim N(\mu, \Sigma) \quad (6)$$

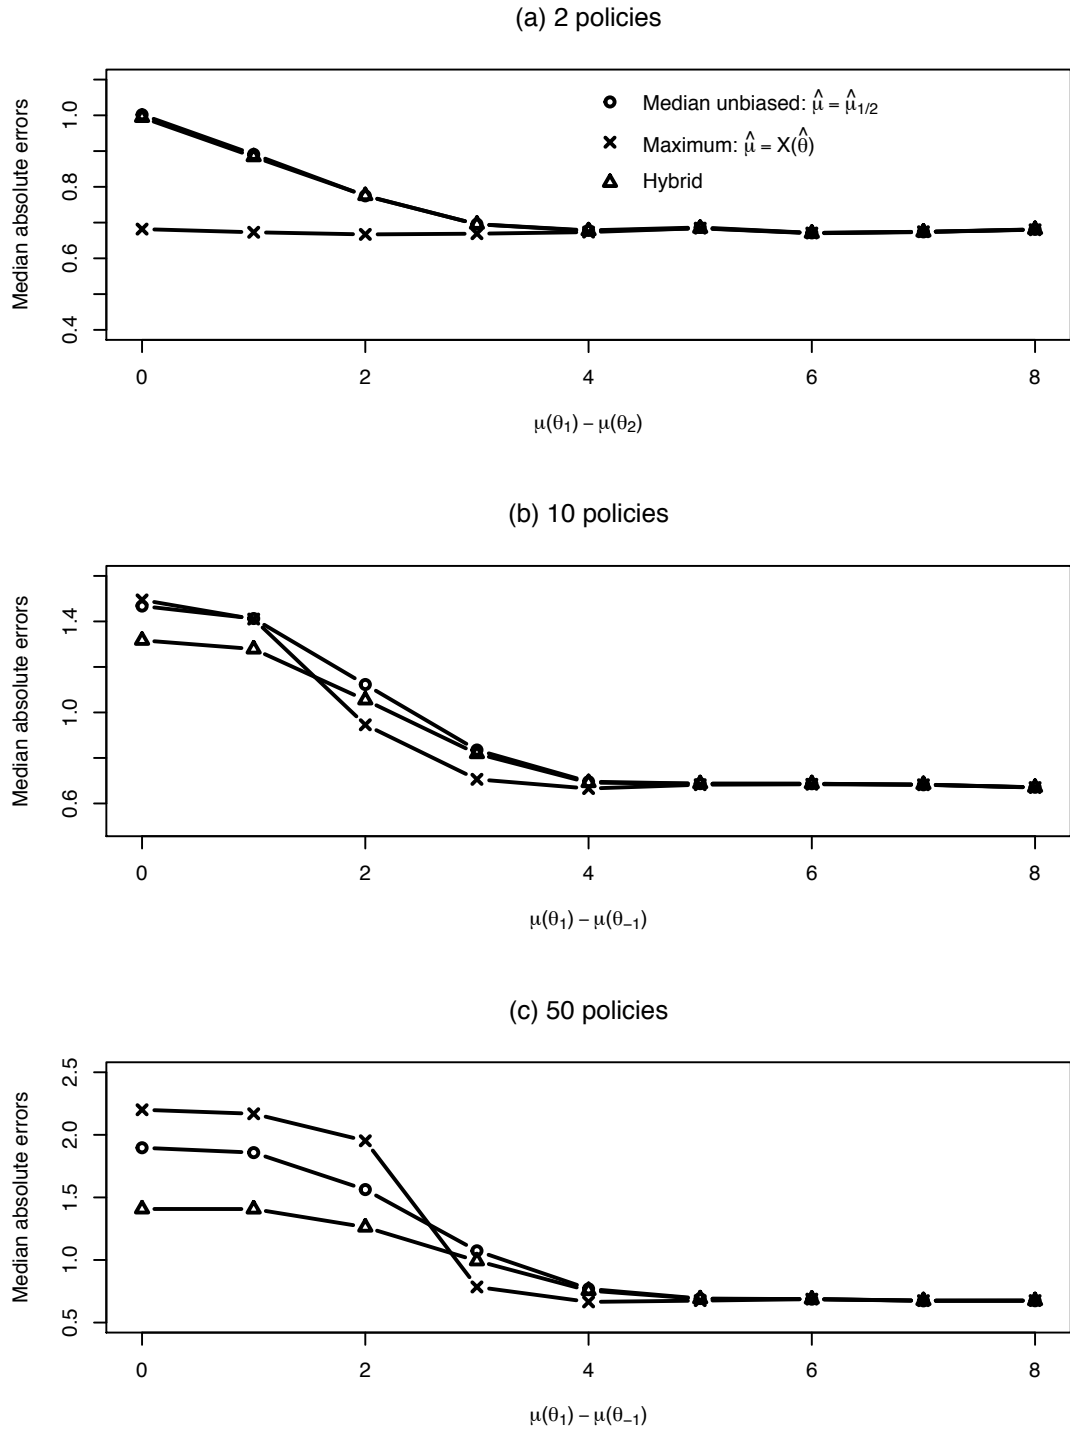


Figure 4: Median absolute error of estimators of $\mu(\hat{\theta})$ in cases with 2, 10, and 50 policies.




Figure 5: Coverage conditional on $\hat{\theta} = \theta_1$ in case with two policies.

for

$$E \left[\begin{pmatrix} X(\theta) \\ Y(\theta) \end{pmatrix} \right] = \mu(\theta) = \begin{pmatrix} \mu_X(\theta) \\ \mu_Y(\theta) \end{pmatrix},$$

$$\Sigma(\theta, \tilde{\theta}) = \begin{pmatrix} \Sigma_X(\theta, \tilde{\theta}) & \Sigma_{XY}(\theta, \tilde{\theta}) \\ \Sigma_{YX}(\theta, \tilde{\theta}) & \Sigma_Y(\theta, \tilde{\theta}) \end{pmatrix} = Cov \left(\begin{pmatrix} X(\theta) \\ Y(\theta) \end{pmatrix}, \begin{pmatrix} X(\tilde{\theta}) \\ Y(\tilde{\theta}) \end{pmatrix} \right).$$

We assume that Σ is known, while μ is unknown and unrestricted unless noted otherwise. For brevity of notation, we abbreviate $\Sigma(\theta, \theta)$ to $\Sigma(\theta)$. We will show that this model arises naturally as an asymptotic approximation. We assume throughout that $\Sigma_Y(\theta) > 0$ for all $\theta \in \Theta$, since the inference problem we study is trivial when $\Sigma_Y(\theta) = 0$.

We are interested in inference on $\mu_Y(\hat{\theta})$, where $\hat{\theta}$ is determined based on X . We define $\hat{\theta}$ through either the *level maximization* problem where (for $d_X = 1$)

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} X(\theta), \quad (7)$$

or the *norm maximization* problem where (for $d_X \geq 1$)

$$\hat{\theta} = \underset{\theta \in \Theta}{\operatorname{argmax}} \|X(\theta)\|, \quad (8)$$

with $\|\cdot\|$ denoting the Euclidean norm.¹² We will again be interested in constructing confidence sets for $\mu_Y(\hat{\theta})$ that are valid either conditional on the value of $\hat{\theta}$ or unconditionally, as well as median-unbiased estimates. We may also want to condition on some additional event $\hat{\gamma}=\tilde{\gamma}$, for $\hat{\gamma}=\gamma(X)$ a function of X which takes values in the finite set Γ . In such cases, we aim to construct confidence sets for $\mu_Y(\hat{\theta})$ that are valid conditional on the pair $(\hat{\theta},\hat{\gamma})$. Examples of such additional conditioning events are discussed below.

In the remainder of this section, we show how this class of problems arises in examples and discuss the choice between conditional and unconditional confidence sets in each case. We first revisit the EWM problem in a more general setting and show that it gives rise to the level maximization problem (7) asymptotically. We then discuss threshold regression models and show that they reduce to the norm maximization problem (8) asymptotically. We also briefly discuss other examples giving rise to level and norm maximization problems, and note that finite sample results for level and norm maximization in the normal model (6) translate to uniform asymptotic results over large classes of models.

Empirical Welfare Maximization As in the last section, we aim to select a welfare-maximizing treatment rule from a set of policies Θ in the EWM problem of Kitagawa and Tetenov (2018b). Let us assume that we have a sample of independent observations $i \in \{1, \dots, n\}$ from a randomized trial where treatment is randomly assigned conditional on observables C_i with $Pr\{D_i=1|C_i\}=d(C_i)$. We consider policies that assign units to treatment based on the observables, where rule θ assigns i to treatment if and only if $C_i \in \mathcal{C}_\theta$. The scaled empirical welfare under policy θ is¹³

$$X_n(\theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left(\frac{Y_i D_i}{d(C_i)} 1_{\{C_i \in \mathcal{C}_\theta\}} + \frac{Y_i(1-D_i)}{1-d(C_i)} 1_{\{C_i \notin \mathcal{C}_\theta\}} \right).$$

EWM again selects the policy that maximizes empirical welfare: $\hat{\theta}_n = \arg\max_{\theta \in \Theta} X_n(\theta)$.

The definition of Y_n in this setting depends on the object of interest. We may be interested in the overall social welfare, in which case we can define $Y_n = X_n$. Alternatively we could be interested in social welfare relative to the baseline of no treatment, in which case we can define $Y_n(\theta)$ as the difference in scaled empirical welfare between policy θ and

¹²For simplicity of notation we will assume $\hat{\theta}$ is unique almost surely unless noted otherwise. Our conditional analysis does not rely on this assumption, however: see footnote 20 below.

¹³Kitagawa and Tetenov (2018b) primarily consider welfare relative to the baseline of no treatment, which yields the same optimal policy.

the policy that treats no one, which we denote by $\theta=0$:

$$Y_n(\theta) = X_n(\theta) - X_n(0) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left[\frac{Y_i D_i}{d(C_i)} - \frac{Y_i(1-D_i)}{1-d(C_i)} \right] 1\{C_i \in \mathcal{C}_\theta\}.$$

Likewise, we might be interested in the social welfare for a particular subgroup defined by the observables, say \mathcal{S} , in which case we can take

$$Y_n(\theta) = \frac{\sqrt{n} \sum_{i=1}^n \left(\frac{Y_i D_i}{d(C_i)} 1\{C_i \in \mathcal{S} \cap \mathcal{C}_\theta\} + \frac{Y_i(1-D_i)}{1-d(C_i)} 1\{C_i \in \mathcal{S} \setminus \mathcal{C}_\theta\} \right)}{\sum_{i=1}^n 1\{C_i \in \mathcal{S}\}}.$$

For $\mu_{X,n}$ and $\mu_{Y,n}$ the true scaled social welfare corresponding to X_n and Y_n ,

$$\begin{pmatrix} X_n - \mu_{X,n} \\ Y_n - \mu_{Y,n} \end{pmatrix} \Rightarrow N(0, \Sigma) \quad (9)$$

under mild conditions, where the covariance Σ will depend on the data generating process and the definition of Y_n but is consistently estimable. By contrast, the scaling of X_n and Y_n means that $\mu_{X,n}$ and $\mu_{Y,n}$ are not consistently estimable. As in the last section, this suggests the asymptotic problem where we observe normal random vectors (X, Y) as in (6) with Σ known and $\hat{\theta}$ defined as in (7), the level maximization problem.¹⁴

As argued in the last section, if a policy maker is given a recommended policy $\hat{\theta}$ as well as a confidence set for $\mu_Y(\hat{\theta})$, it is natural to require that the confidence set be valid conditional on the recommendation. It may also be natural to condition on additional variables. For example, if a recommendation is made only when we reject the null hypothesis that no policy in Θ improves outcomes over the base case of no treatment, $H_0: \max_{\theta \in \Theta} \mu(\theta) \leq \mu(0)$, then it is also natural to condition inference on this rejection.¹⁵ To cover this case we can define $\hat{\gamma} = \gamma(X)$ as a dummy for rejection of H_0 . If on the other hand we care only about performance on average across a range of recommendations, we need only impose unconditional coverage. \triangle

The level maximization problem arises in a number of other settings as well. For example, selecting the “best” policy from a collection considered in A/B tests is closely

¹⁴Under mild regularity conditions, (9) also holds in settings where the empirical welfare involves estimated propensity scores and/or estimated outcome regressions, e.g., the hybrid procedures of Kitagawa and Tetenov (2018b) and the doubly robust welfare estimators of Athey and Wager (2018).

¹⁵In the case of $|\Theta|=2$, conditioning on this rejection can be interpreted as conditioning on the event that the decision criterion of Tetenov (2012) supports the same policy.

related to EWM. Further afield, the literature on tests of superior predictive performance (c.f. White (2000); Hansen (2005); Romano and Wolf (2005)) considers the problem of testing whether some trading strategies or forecasting rules amongst a candidate set beat a benchmark. If we define $X_n = Y_n$ as the vector of performance measures for different strategies, X_n is asymptotically normal under mild conditions (see e.g. Romano and Wolf (2005)). If one wants to form a confidence set for the performance of the “best” strategy based on X_n (perhaps also conditioning on the result of a test for superior performance), this reduces to our level maximization problem asymptotically.

Another example comes from Bhattacharya (2009) and Graham et al. (2014), who consider the problem of optimally matching individuals to maximize peer effects. For X_n again a scaled objective function, the results of Bhattacharya (2009) show that his problem reduces to level maximization asymptotically when one considers a finite set of assignments. More broadly, any time we consider M -estimation with a finite parameter space and are interested in the value of the population objective or some other function at the estimated optimal value, this falls into our level maximization framework under mild conditions.

We next discuss an example of threshold regression estimation, showing that it gives rise to our norm-maximization problem asymptotically.

Threshold Regression Estimation Suppose we observe data on an outcome Y_i , a threshold regressor Q_i and a k -dimensional vector of regressors C_i for $i \in \{1, \dots, n\}$. We assume there is a linear but potentially regressor-dependent relationship between Y_i and C_i :

$$Y_i = C_i'(\beta + \varphi_n(Q_i)) + U_i, \quad (10)$$

where $Q_i \in \mathbb{R}$ and the residuals U_i are orthogonal to Q_i and C_i . Similarly to Elliott and Müller (2014) and Wang (2018), the function $\varphi_n : \mathbb{R} \rightarrow \mathbb{R}^k$ determines the value of the regressor-dependent coefficient $\beta + \varphi_n(Q_i)$. This model nests the traditional threshold regression model (see e.g. Hansen (2000) and references therein) by taking

$$\varphi_n(Q_i) = 1(Q_i > \theta)\delta, \quad (11)$$

where $\theta \in \mathbb{R}$ is the “true” threshold. The threshold model (11) is often used as a parsimonious approximation to a more general linear regression model with regressor-dependent coefficients. For example, Card et al. (2008) use the threshold model to approximate a theoretical model with smoothly-varying regressor-dependent coefficients. See also the

motivations for this model discussed in Hansen (1997, 2000).

Since the threshold regression model is widely used in practice, we consider a researcher who fits the model (11). To allow the possibility of misspecification, however, we assume only that the data is generated by (10). To provide a good asymptotic approximation to finite sample behavior, we follow Elliott and Müller (2007, 2014) and Wang (2018) and model parameter instability as on the same order as sampling uncertainty, with $\varphi_n(Q_i) = \frac{1}{\sqrt{n}}g(Q_i)$ for a fixed function g . We further assume that

$$\frac{1}{n} \sum_{i=1}^n C_i C_i' 1(Q_i \leq \theta) \rightarrow_p \Sigma_C(\theta), \quad \frac{1}{n} \sum_{i=1}^n C_i C_i' g(Q_i) 1(Q_i \leq \theta) \rightarrow_p \Sigma_{Cg}(\theta), \quad (12)$$

and

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n C_i U_i 1(Q_i \leq \theta) \Rightarrow G(\theta), \quad (13)$$

all uniformly in $\theta \in \mathbb{R}$. Here $\Sigma_C : \mathbb{R} \rightarrow \mathbb{R}^{k \times k}$ is a consistently-estimable matrix-valued function and $\Sigma_C(\theta)$ is full rank for all θ in the interior of the support of Q_i , $\Sigma_{Cg} : \mathbb{R} \rightarrow \mathbb{R}^k$ is a vector-valued function, and $G(\cdot)$ is a k -dimensional mean zero Gaussian process with a consistently estimable covariance function that is positive definite when evaluated at points in the interior of the support of Q_i . Conditions (12) and (13) are analogous to Conditions 1(ii) and 1(iv) of Elliott and Müller (2007) for structural break models in a time-series setting. See Wang (2018) for sufficient conditions that give rise to (12) and (13).

The standard threshold estimator $\hat{\theta}_n$ chooses θ to minimize the sum of squared residuals in an OLS regression of Y_i on C_i and $\mathbf{1}(Q_i > \theta)C_i$ across a finite grid of thresholds Θ .¹⁶ For

$$X_n(\theta) = \begin{pmatrix} (\sum_{i=1}^n C_i C_i' 1(Q_i \leq \theta))^{-\frac{1}{2}} (\sum_{i=1}^n C_i \eta_i 1(Q_i \leq \theta)) \\ (\sum_{i=1}^n C_i C_i' 1(Q_i > \theta))^{-\frac{1}{2}} (\sum_{i=1}^n C_i \eta_i 1(Q_i > \theta)) \end{pmatrix},$$

with $\eta_i \equiv U_i + n^{-1/2} C_i' g(Q_i)$, arguments analogous to those in the proof of Proposition 1 in Elliott and Müller (2007) imply that $\hat{\theta}_n = \arg\max_{\theta \in \Theta} \|X_n(\theta)\| + o_p(1)$, where $o_p(1)$ is an asymptotically negligible term. Hence, $\hat{\theta}_n$ is asymptotically equivalent to the solution to a particular norm-maximization problem (8).

Suppose we are interested in the approximate change in the j th parameter $\delta_j = e_j' \delta$,

¹⁶Note that finiteness of Θ is without loss of generality if Q_i is finitely-supported, but that we otherwise limit attention to a finite collection of thresholds.

where e_j is the j^{th} standard basis vector.¹⁷ In practice it is common to estimate δ by least squares imposing the estimated threshold $\hat{\theta}_n$. When the threshold regression model (11) is misspecified, however, there is neither a “true” threshold θ nor a “true” change coefficient δ . Instead, the population regression coefficient $\delta(\theta)$ imposing threshold θ depends on θ . Thus, for threshold θ , the coefficient of interest is $\delta_j(\theta)$. Denote the OLS estimate imposing threshold θ by $\hat{\delta}_j(\theta)$ and define $Y_n(\theta) = \sqrt{n}\hat{\delta}_j(\theta)$. If we define $\mu_{Y,n}(\theta) = \sqrt{n}\delta_j(\theta)$ as the scaled coefficient of interest and $\mu_{X,n}(\theta)$ as the population analog of $X_n(\theta)$, Section B.2 of the supplement shows that

$$\begin{pmatrix} X_n(\theta) - \mu_{X,n}(\theta) \\ Y_n(\theta) - \mu_{Y,n}(\theta) \end{pmatrix} \Rightarrow N(0, \Sigma(\theta)) \quad (14)$$

uniformly over a parameter space Θ contained in the interior of the support of Q_i , where the covariance matrix $\Sigma(\theta)$ is consistently estimable but $\mu_{X,n}(\theta)$ and $\mu_{Y,n}(\theta)$ are not. As before, this suggests the asymptotic problem (6) where we now define $\hat{\theta}$ through norm maximization (8).

Since the estimated threshold $\hat{\theta}$ is random and the parameter of interest $\delta_j(\theta)$ depends on θ , it is important to account for this randomness in our inference procedures. In particular, it may be appealing to condition inference on the estimated threshold $\hat{\theta}$, since we only seek to conduct inference on $\delta_j(\hat{\theta})$ when $\hat{\theta} = \tilde{\theta}$. It may also be natural to condition inference on additional variables. For example, if we report a confidence set for the change coefficient $\delta_j(\hat{\theta})$ only when we reject the null hypothesis of parameter constancy, $H_0: \varphi_n(\theta) = 0$ for all θ , it is natural to condition inference on this rejection. As above, this can be accomplished by defining $\hat{\gamma} = \gamma(X)$ as a dummy for rejection of H_0 , and conditioning inference on $(\hat{\theta}, \hat{\gamma})$. Even if we only desire coverage of $\delta_j(\hat{\theta})$ on average over the distribution of $\hat{\theta}$, and so prefer to consider unconditional confidence sets, accounting for the randomness of $\hat{\theta}$ remains important. If on the other hand we are confident that the threshold model is correctly specified, so that (11) holds in the data, it will typically be more appealing to focus on inference for the “true” parameters as in Elliott and Müller (2014) and Wang (2018). \triangle

An analogous analysis applies to estimation and inference in the traditional structural break model (see e.g. Hansen (2001) and Perron (2006) and references therein) under local asymptotics as in Elliott and Müller (2007, 2014). Moreover, while our discussion of threshold regression estimation focuses on the linear model (10), Elliott and Müller

¹⁷By changing the definition of Y_n below, our results likewise apply to the pre-change parameters β_j and the post-change parameters $\beta_j + \delta_j$, amongst other possible objects of interest.

(2014) show that structural break estimation in nonlinear models with time-varying parameters gives rise to the same asymptotic problem. Hence, our results apply in that setting as well. Likewise, Wang (2018) shows that the same asymptotic problem arises in nonlinear threshold models.¹⁸ Further afield, one could generalize our approach to consider norm-minimization rather than norm-maximization, and so derive results for GMM-type problems with finite parameter spaces.

Uniform Asymptotic Validity We have shown that the empirical welfare maximization and threshold regression problems asymptotically resemble level and norm maximization based on the finite-sample normal model (6). Section D of the supplement builds on this connection and shows that if we consider classes of data generating processes such that (X_n, Y_n) are uniformly well-approximated by the normal model (6), we have a uniformly consistent estimator $\hat{\Sigma}_n$ for Σ , and Σ satisfies mild regularity conditions, our finite-sample results in the normal model (6) translate to uniform asymptotic results. These uniformity results apply to level maximization settings without any restrictions on the behavior of $(\mu_{X,n}, \mu_{Y,n})$. In norm maximization settings, by contrast, we limit attention to $(\mu_{X,n}, \mu_{Y,n})$ lying in bounded sets, since this is the context for which the asymptotic results of Elliott and Müller (2007, 2014) and Wang (2018) imply an asymptotic norm-maximization representation.¹⁹

4 Conditional Inference

This section develops conditional inference procedures for our general setting. We seek confidence sets with correct coverage conditional on $(\hat{\theta}, \hat{\gamma})$,

$$Pr_{\mu} \left\{ \mu_Y(\hat{\theta}) \in CS | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma} \right\} \geq 1 - \alpha \text{ for all } \tilde{\theta} \in \Theta, \tilde{\gamma} \in \Gamma, \text{ and all } \mu. \quad (15)$$

As in the stylized example of Section 2, we consider both equal-tailed and uniformly most accurate unbiased confidence sets.²⁰ We also derive optimal conditionally α -quantile-unbiased estimators, which for $\alpha \in (0, 1)$ satisfy

$$Pr_{\mu} \left\{ \hat{\mu}_{\alpha} \geq \mu_Y(\hat{\theta}) | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma} \right\} = \alpha \text{ for all } \tilde{\theta} \in \Theta, \tilde{\gamma} \in \Gamma, \text{ and all } \mu. \quad (16)$$

¹⁸In a manuscript circulated after the initial public version of this paper, Hyun et al. (2018) consider the related problem of conditional inference for changepoint detection, but the changepoint estimation methods they consider cannot be cast as norm-maximization, so their results do not overlap with ours.

¹⁹If one instead considers cases where $(\mu_{X,n}, \mu_{Y,n})$ diverges, as occurs for example in threshold regression with non-vanishing parameter instability, the problem reduces to level-maximization asymptotically.

²⁰If $\hat{\theta}$ is not unique we change the conditioning event $\hat{\theta} = \tilde{\theta}$ to $\tilde{\theta} \in \arg\max X(\theta)$ or $\tilde{\theta} \in \arg\max \|X(\theta)\|$ for the level and norm maximization problems, respectively.

Our conditional procedures depend on the conditioning events of interest. We analyze these conditioning events for our general level and norm maximization settings, and illustrate them in our EWM and threshold regression examples. We then discuss conventional sample splitting as an alternative conditional approach and briefly discuss the construction of dominating procedures. Finally, we show that our conditional procedures converge to conventional ones when $Pr_\mu\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\} \rightarrow 1$ so the latter are valid.

4.1 Optimal Conditional Inference

Since $\hat{\theta}$ and $\hat{\gamma}$ are functions of X , we can re-write the conditioning event in terms of the sample space of X as $\{X: \hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\} = \mathcal{X}(\tilde{\theta}, \tilde{\gamma})$. Thus, for conditional inference we are interested in the distribution of (X, Y) conditional on $X \in \mathcal{X}(\tilde{\theta}, \tilde{\gamma})$. Our results below imply that under mild conditions, the elements of Y other than $Y(\tilde{\theta})$ do not help in constructing a quantile-unbiased estimate or unbiased confidence set for $\mu_Y(\hat{\theta})$ conditional on $X \in \mathcal{X}(\tilde{\theta}, \tilde{\gamma})$. Hence, we limit attention to the conditional distribution of $(X, Y(\tilde{\theta}))$ given $X \in \mathcal{X}(\tilde{\theta}, \tilde{\gamma})$.

Since $(X, Y(\tilde{\theta}))$ is jointly normal unconditionally, it has a multivariate truncated normal distribution conditional on $X \in \mathcal{X}(\tilde{\theta}, \tilde{\gamma})$. Correlation between X and $Y(\tilde{\theta})$ implies that the conditional distribution of $Y(\tilde{\theta})$ depends on both the parameter of interest $\mu_Y(\hat{\theta})$ and μ_X . To eliminate dependence on the nuisance parameter μ_X , we condition on a sufficient statistic. Without truncation and for any fixed $\mu_Y(\tilde{\theta})$, a minimal sufficient statistic for μ_X is

$$Z_{\tilde{\theta}} = X - \left(\Sigma_{XY}(\cdot, \tilde{\theta}) / \Sigma_Y(\tilde{\theta}) \right) Y(\tilde{\theta}), \quad (17)$$

where we use $\Sigma_{XY}(\cdot, \tilde{\theta})$ to denote $Cov(X, Y(\tilde{\theta}))$. $Z_{\tilde{\theta}}$ corresponds to the part of X that is (unconditionally) orthogonal to $Y(\tilde{\theta})$ which, since $(X, Y(\tilde{\theta}))$ are jointly normal, means that $Z_{\tilde{\theta}}$ and $Y(\tilde{\theta})$ are independent. Truncation breaks this independence, but $Z_{\tilde{\theta}}$ remains minimal sufficient for μ_X . The conditional distribution of $Y(\hat{\theta})$ given $\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}, Z_{\tilde{\theta}}=z\}$ is truncated normal:

$$Y(\hat{\theta}) | \hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}, Z=z \sim \xi | \xi \in \mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z), \quad (18)$$

where $\xi \sim N(\mu_Y(\tilde{\theta}), \Sigma_Y(\tilde{\theta}))$ is normally distributed and

$$\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z) = \left\{ y : z + \left(\Sigma_{XY}(\cdot, \tilde{\theta}) / \Sigma_Y(\tilde{\theta}) \right) y \in \mathcal{X}(\tilde{\theta}, \tilde{\gamma}) \right\} \quad (19)$$

is the set of values for $Y(\tilde{\theta})$ such that the implied X falls in $\mathcal{X}(\tilde{\theta}, \tilde{\gamma})$ given $Z_{\tilde{\theta}}=z$. Thus, conditional on $\hat{\theta}=\tilde{\theta}$, $\hat{\gamma}=\tilde{\gamma}$, and $Z_{\tilde{\theta}}=z$, $Y(\hat{\theta})$ follows a one-dimensional truncated normal

distribution with truncation set $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z)$.

Using this result, it is straightforward to construct quantile-unbiased estimators for $\mu_Y(\hat{\theta})$. Let $F_{TN}(y; \mu_Y(\tilde{\theta}), \tilde{\theta}, \tilde{\gamma}, z)$ denote the distribution function for the truncated normal distribution (18). This distribution function is strictly decreasing in $\mu_Y(\tilde{\theta})$. Define $\hat{\mu}_\alpha$ as the unique solution to

$$F_{TN}(Y(\hat{\theta}); \hat{\mu}_\alpha, \tilde{\theta}, \tilde{\gamma}, Z_{\tilde{\theta}}) = 1 - \alpha. \quad (20)$$

Proposition 1 below shows that $\hat{\mu}_\alpha$ is conditionally α -quantile-unbiased in the sense of (16), so $\hat{\mu}_{\frac{1}{2}}$ is median-unbiased while the equal-tailed interval $CS_{ET} = [\hat{\mu}_{\alpha/2}, \hat{\mu}_{1-\alpha/2}]$ has conditional coverage $1 - \alpha$. Moreover, results in Pfanzagl (1979) and Pfanzagl (1994) on quantile-unbiased estimation in exponential families imply that $\hat{\mu}_\alpha$ is optimal in the class of quantile-unbiased estimators.

To establish optimality, we add the following assumption:

Assumption 1

If $\Sigma = \text{Cov}((X', Y')')$ has full rank, then the parameter space for μ is open and convex. Otherwise, there exists some μ^ such that the parameter space for μ is an open convex subset of $\{\mu^* + \Sigma^{\frac{1}{2}}v : v \in \mathbb{R}^{\dim(X, Y)}\}$ where $\Sigma^{\frac{1}{2}}$ is the symmetric square root of Σ .*

This assumption requires that the parameter space for μ be sufficiently rich.²¹ When Σ is degenerate (for example when X and Y are perfectly correlated as in the EWM example with $X = Y$), this assumption further implies that (X, Y) have the same support for all values of μ . This rules out cases in which some a pair of parameter values μ_1, μ_2 can be perfectly distinguished based on the data. Under this assumption, $\hat{\mu}_\alpha$ is an optimal quantile-unbiased estimator.

Proposition 1

Let $\hat{\mu}_\alpha$ be the unique solution of (20). $\hat{\mu}_\alpha$ is conditionally α -quantile-unbiased in the sense of (16). If Assumption 1 holds, then $\hat{\mu}_\alpha$ is the uniformly most concentrated α -quantile-unbiased estimator in that for any other conditionally α -quantile-unbiased estimator $\hat{\mu}_\alpha^$ and any loss function $L(d, \mu_Y(\tilde{\theta}))$ that attains its minimum at $d = \mu_Y(\tilde{\theta})$ and is quasiconvex in d for all $\mu_Y(\tilde{\theta})$,*

$$E_\mu \left[L\left(\hat{\mu}_\alpha, \mu_Y(\tilde{\theta})\right) | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma} \right] \leq E_\mu \left[L\left(\hat{\mu}_\alpha^*, \mu_Y(\tilde{\theta})\right) | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma} \right]$$

²¹The assumption that the parameter space is open can be relaxed at the cost of complicating the statements below.

for all μ and all $\tilde{\theta} \in \Theta$, $\tilde{\gamma} \in \Gamma$.

Proposition 1 shows that $\hat{\mu}_\alpha$ is optimal in the strong sense that it has lower risk (expected loss) than any other quantile-unbiased estimator for a large class of loss functions.

Rather than considering equal-tailed intervals, we can alternatively consider unbiased confidence sets. Following Lehmann and Romano (2005), we say that a level $1-\alpha$ two-sided confidence set CS is unbiased if its probability of covering any given false parameter value is bounded above by $1-\alpha$. Likewise, a one sided lower (upper) confidence set is unbiased if its probability of covering a false parameter value above (below) the true value is bounded above by $1-\alpha$. Using the duality between tests and confidence sets, a level $1-\alpha$ confidence set CS is unbiased if and only if $\phi(\mu_{Y,0}) = 1\{\mu_{Y,0} \notin CS\}$ is an unbiased test for the corresponding family of hypotheses.²² The results of Lehmann and Scheffé (1955) applied in our setting imply that optimal unbiased tests conditional on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}\}$ are the same as optimal unbiased tests conditional on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}, Z_{\tilde{\theta}} = z_{\tilde{\theta}}\}$. These optimal tests take a simple form.

Define a size α test of the two-sided hypothesis $H_0: \mu_Y(\tilde{\theta}) = \mu_{Y,0}$ as

$$\phi_{TS,\alpha}(\mu_{Y,0}) = 1\left\{Y(\tilde{\theta}) \notin [c_l(Z_{\tilde{\theta}}), c_u(Z_{\tilde{\theta}})]\right\} \quad (21)$$

where $c_l(z)$, $c_u(z)$ solve

$$Pr\{\zeta \in [c_l(z), c_u(z)]\} = 1-\alpha, \quad E[\zeta 1\{\zeta \in [c_l(z), c_u(z)]\}] = (1-\alpha)E[\zeta]$$

for ζ that follows a truncated normal distribution

$$\zeta \sim \xi | \xi \in \mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z), \quad \xi \sim N(\mu_{Y,0}, \Sigma_Y(\tilde{\theta})).$$

Likewise, define a size α test of the one-sided hypothesis $H_0: \mu_Y(\tilde{\theta}) \geq \mu_{Y,0}$ as

$$\phi_{OS-, \alpha}(\mu_{Y,0}) = 1\left\{F_{TN}(Y(\tilde{\theta}); \mu_{Y,0}, \tilde{\theta}, \tilde{\gamma}, z) \leq \alpha\right\} \quad (22)$$

and a test of $H_0: \mu_Y(\tilde{\theta}) \leq \mu_{Y,0}$ as

$$\phi_{OS+, \alpha}(\mu_{Y,0}) = 1\left\{F_{TN}(Y(\tilde{\theta}); \mu_{Y,0}, \tilde{\theta}, \tilde{\gamma}, z) \geq 1-\alpha\right\}. \quad (23)$$

²²That is, $H_0: \mu_Y(\tilde{\theta}) = \mu_{Y,0}$ for a two-sided confidence set, $H_0: \mu_Y(\tilde{\theta}) \geq \mu_{Y,0}$ for a lower confidence set and $H_0: \mu_Y(\tilde{\theta}) \leq \mu_{Y,0}$ for an upper confidence set.

Proposition 2

If Assumption 1 holds, $\phi_{TS,\alpha}$, $\phi_{OS-, \alpha}$, and $\phi_{OS+, \alpha}$ are uniformly most powerful unbiased size α tests of their respective null hypotheses conditional on $\hat{\theta} = \tilde{\theta}$ and $\hat{\gamma} = \tilde{\gamma}$.

To form uniformly most accurate unbiased confidence sets we collect the values not rejected by these tests. The two-sided uniformly most accurate unbiased confidence set is $CS_U = \{\mu_{Y,0} : \phi_{TS,\alpha}(\mu_{Y,0}) = 0\}$. CS_U is unbiased and has conditional coverage $1 - \alpha$ by construction. Likewise, we can form lower and upper one-sided uniformly most accurate unbiased confidence intervals as $CS_{U,-} = \{\mu_{Y,0} : \phi_{OS-, \alpha}(\mu_{Y,0}) = 0\} = (-\infty, \hat{\mu}_{1-\alpha}]$, and $CS_{U,+} = \{\mu_{Y,0} : \phi_{OS+, \alpha}(\mu_{Y,0}) = 0\} = [\hat{\mu}_\alpha, \infty)$, respectively. Hence, we can view CS_{ET} as the intersection of level $1 - \frac{\alpha}{2}$ uniformly most accurate unbiased upper and lower confidence intervals. Unfortunately, no such simplification is generally available for CS_U , though Lemma 5.5.1 of Lehmann and Romano (2005) guarantees that this set is an interval.

4.2 Conditioning Sets

Thus far we have left the conditioning events $\mathcal{X}(\tilde{\theta}, \tilde{\gamma})$ and $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z)$ abstract. To implement our conditional procedures, however, we need tractable representations of $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z)$. We first derive the form of this conditioning event for the level maximization problem (7) and the norm maximization problem (8) without additional conditioning variables $\hat{\gamma}$. We then discuss the effect of adding conditioning variables and illustrate in our examples.

In level maximization problems without additional conditioning variables, we are interested in inference conditional on $X \in \mathcal{X}(\tilde{\theta})$ for $\mathcal{X}(\tilde{\theta}) = \left\{ X : X(\tilde{\theta}) = \max_{\theta \in \Theta} X(\theta) \right\}$. The following result, based on Lemma 5.1 of Lee et al. (2016), derives $\mathcal{Y}(\tilde{\theta}, z)$ in this setting.

Proposition 3

Let $\Sigma_{XY}(\tilde{\theta}) = \text{Cov}(X(\tilde{\theta}), Y(\tilde{\theta}))$. Define

$$\mathcal{L}(\tilde{\theta}, Z_{\tilde{\theta}}) = \max_{\theta \in \Theta : \Sigma_{XY}(\tilde{\theta}) > \Sigma_{XY}(\tilde{\theta}, \theta)} \frac{\Sigma_Y(\tilde{\theta}) (Z_{\tilde{\theta}}(\theta) - Z_{\tilde{\theta}}(\tilde{\theta}))}{\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, \theta)},$$

$$\mathcal{U}(\tilde{\theta}, Z_{\tilde{\theta}}) = \min_{\theta \in \Theta : \Sigma_{XY}(\tilde{\theta}) < \Sigma_{XY}(\tilde{\theta}, \theta)} \frac{\Sigma_Y(\tilde{\theta}) (Z_{\tilde{\theta}}(\theta) - Z_{\tilde{\theta}}(\tilde{\theta}))}{\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, \theta)},$$

and

$$\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) = \min_{\theta \in \Theta : \Sigma_{XY}(\tilde{\theta}) = \Sigma_{XY}(\tilde{\theta}, \theta)} - (Z_{\tilde{\theta}}(\theta) - Z_{\tilde{\theta}}(\tilde{\theta})).$$

If $\mathcal{V}(\tilde{\theta}, z) \geq 0$, then $\mathcal{Y}(\tilde{\theta}, z) = [\mathcal{L}(\tilde{\theta}, z), \mathcal{U}(\tilde{\theta}, z)]$. If $\mathcal{V}(\tilde{\theta}, z) < 0$, then $\mathcal{Y}(\tilde{\theta}, z) = \emptyset$.

Thus, the conditioning event $\mathcal{V}(\tilde{\theta}, z)$ is an interval bounded above and below by easy-to-calculate functions of z . While we must have $\mathcal{V}(\tilde{\theta}, z) \geq 0$ for this interval to be non-empty, $Pr_{\mu} \left\{ \mathcal{V}(\hat{\theta}, Z_{\hat{\theta}}) < 0 \right\} = 0$ for all μ so this constraint holds almost surely when we consider the value $\hat{\theta}$ observed in the data. Hence, in applications we can safely ignore this constraint and calculate only $\mathcal{L}(\hat{\theta}, Z_{\hat{\theta}})$ and $\mathcal{U}(\hat{\theta}, Z_{\hat{\theta}})$.

The norm maximization conditioning event is $\mathcal{X}(\tilde{\theta}) = \left\{ X : \|X(\tilde{\theta})\| = \max_{\theta \in \Theta} \|X(\theta)\| \right\}$. This conditioning event involves nonlinear constraints so the results of Lee et al. (2016) do not apply. The expression for $\mathcal{V}(\tilde{\theta}, z)$ is more involved, but remains easy to calculate.

Proposition 4

Define

$$\begin{aligned} A(\tilde{\theta}, \theta) &= \Sigma_Y(\tilde{\theta})^{-2} \sum_{i=1}^{d_X} \left[\Sigma_{XY,i}(\tilde{\theta})^2 - \Sigma_{XY,i}(\theta, \tilde{\theta})^2 \right], \\ B_Z(\tilde{\theta}, \theta) &= 2\Sigma_Y(\tilde{\theta})^{-1} \sum_{i=1}^{d_X} \left[\Sigma_{XY,i}(\tilde{\theta}) Z_{\tilde{\theta},i}(\tilde{\theta}) - \Sigma_{XY,i}(\theta, \tilde{\theta}) Z_{\tilde{\theta},i}(\theta) \right], \\ C_Z(\tilde{\theta}, \theta) &= \sum_{i=1}^{d_X} \left[Z_{\tilde{\theta},i}(\tilde{\theta})^2 - Z_{\tilde{\theta},i}(\theta)^2 \right]. \end{aligned}$$

For

$$\begin{aligned} D_Z(\tilde{\theta}, \theta) &= B_Z(\tilde{\theta}, \theta)^2 - 4A(\tilde{\theta}, \theta)C_Z(\tilde{\theta}, \theta), \quad H_Z(\tilde{\theta}, \theta) = \frac{-C_Z(\tilde{\theta}, \theta)}{B_Z(\tilde{\theta}, \theta)}, \\ G_Z(\tilde{\theta}, \theta) &= \frac{-B_Z(\tilde{\theta}, \theta) - \sqrt{D_Z(\tilde{\theta}, \theta)}}{2A(\tilde{\theta}, \theta)}, \quad \text{and} \quad K_Z(\tilde{\theta}, \theta) = \frac{-B_Z(\tilde{\theta}, \theta) + \sqrt{D_Z(\tilde{\theta}, \theta)}}{2A(\tilde{\theta}, \theta)}, \end{aligned}$$

define

$$\begin{aligned} \ell_Z^1(\tilde{\theta}) &= \max \left\{ \max_{\theta \in \Theta: A(\tilde{\theta}, \theta) < 0, D_Z(\tilde{\theta}, \theta) \geq 0} G_Z(\tilde{\theta}, \theta), \max_{\theta \in \Theta: A(\tilde{\theta}, \theta) = 0, B_Z(\tilde{\theta}, \theta) > 0} H_Z(\tilde{\theta}, \theta) \right\}, \\ \ell_Z^2(\tilde{\theta}, \theta) &= \max \left\{ \max_{\theta \in \Theta: A(\tilde{\theta}, \theta) < 0, D_Z(\tilde{\theta}, \theta) \geq 0} G_Z(\tilde{\theta}, \theta), \max_{\theta \in \Theta: A(\tilde{\theta}, \theta) = 0, B_Z(\tilde{\theta}, \theta) > 0} H_Z(\tilde{\theta}, \theta), G_Z(\tilde{\theta}, \theta) \right\}, \\ u_Z^1(\tilde{\theta}, \theta) &= \min \left\{ \min_{\theta \in \Theta: A(\tilde{\theta}, \theta) < 0, D_Z(\tilde{\theta}, \theta) \geq 0} K_Z(\tilde{\theta}, \theta), \min_{\theta \in \Theta: A(\tilde{\theta}, \theta) = 0, B_Z(\tilde{\theta}, \theta) < 0} H_Z(\tilde{\theta}, \theta), K_Z(\tilde{\theta}, \theta) \right\}, \\ u_Z^2(\tilde{\theta}) &= \min \left\{ \min_{\theta \in \Theta: A(\tilde{\theta}, \theta) < 0, D_Z(\tilde{\theta}, \theta) \geq 0} K_Z(\tilde{\theta}, \theta), \min_{\theta \in \Theta: A(\tilde{\theta}, \theta) = 0, B_Z(\tilde{\theta}, \theta) < 0} H_Z(\tilde{\theta}, \theta) \right\}, \end{aligned}$$

and

$$\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) = \min_{\theta \in \Theta: A(\tilde{\theta}, \theta) = B_Z(\tilde{\theta}, \theta) = 0 \text{ or } D_Z(\tilde{\theta}, \theta) < 0} C_Z(\tilde{\theta}, \theta).$$

If $\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) \geq 0$ then

$$\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) = \bigcap_{\theta \in \Theta: A(\tilde{\theta}, \theta) > 0, D_Z(\tilde{\theta}, \theta) \geq 0} \left[\ell_Z^1(\tilde{\theta}), u_Z^1(\tilde{\theta}, \theta) \right] \cup \left[\ell_Z^2(\tilde{\theta}, \theta), u_Z^2(\tilde{\theta}) \right].$$

If $\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) < 0$, then $\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) = \emptyset$.

While the expression for $\mathcal{V}(\tilde{\theta}, z)$ in this setting is long, it is easy to calculate in practice and can be expressed as a finite union of intervals using DeMorgan's laws. As before, $Pr_{\mu} \left\{ \mathcal{V}(\hat{\theta}, Z_{\hat{\theta}}) < 0 \right\} = 0$ for all μ so we can ignore this constraint in applications.

Our derivations have so far assumed we have no additional conditioning variables $\hat{\gamma}$. If we also condition on $\hat{\gamma} = \tilde{\gamma}$, then for $\mathcal{X}_{\gamma}(\tilde{\gamma}) = \{X : \gamma(X) = \tilde{\gamma}\}$, we can write $\mathcal{X}(\tilde{\theta}, \tilde{\gamma}) = \mathcal{X}(\tilde{\theta}) \cap \mathcal{X}_{\gamma}(\tilde{\gamma})$. Likewise, for $\mathcal{Y}_{\gamma}(\tilde{\gamma}, z)$ defined analogously to (19), $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z) = \mathcal{Y}(\tilde{\theta}, z) \cap \mathcal{Y}_{\gamma}(\tilde{\gamma}, z)$. The form of $\mathcal{X}_{\gamma}(\tilde{\gamma})$ and $\mathcal{Y}_{\gamma}(\tilde{\gamma}, z)$ depends on the conditioning variables $\hat{\gamma}$ considered. To illustrate we next discuss the effect of conditioning on the outcomes of pretests in our EWM and threshold regression examples.

Empirical Welfare Maximization (continued) Suppose that we report estimates and confidence sets for welfare only if the improvement in empirical welfare from the estimated optimal policy over a baseline policy $\theta = 0$ exceeds a threshold c , i.e. $X(\hat{\theta}) - X(0) \geq c$. For instance, we might report results only when the test of White (2000) rejects the null that no policy has performance exceeding the baseline, $H_0 : \max_{\theta \in \Theta} \mu_X(\theta) \leq \mu_X(0)$. This implies that we report results only if $X(\hat{\theta}) - X(0) \geq c$ for c a critical value depending on Σ . We can set $\gamma(X) = 1 \left\{ X(\hat{\theta}) - X(0) \geq c \right\}$ and it is natural to condition inference on $\hat{\gamma} = 1$.

Assuming $\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0) > 0$ for simplicity, the conditioning event in this setting is $\mathcal{X}_{\gamma}(1) = \left\{ X : X(\hat{\theta}) - X(0) \geq c \right\}$ and one can show that

$$\mathcal{Y}_{\gamma}(1, Z_{\tilde{\theta}}) = \left\{ y : y \geq \frac{\Sigma_Y(\tilde{\theta}) \left(c - Z_{\tilde{\theta}}(\tilde{\theta}) + Z_{\tilde{\theta}}(0) \right)}{\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0)} \right\}.$$

See Section B.1 of the supplement for details, as well as expressions for other values of $\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0)$. In the present case, provided $\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}) \geq 0$, $\mathcal{Y}(\tilde{\theta}, 1, Z_{\tilde{\theta}}) =$

$[\mathcal{L}^*(\tilde{\theta}, Z_{\tilde{\theta}}), \mathcal{U}(\tilde{\theta}, Z_{\tilde{\theta}})]$, where $\mathcal{U}(\tilde{\theta}, Z_{\tilde{\theta}})$ is the upper bound derived in Proposition 3 while

$$\mathcal{L}^*(\tilde{\theta}, Z_{\tilde{\theta}}) = \max \left\{ \mathcal{L}(\tilde{\theta}, Z_{\tilde{\theta}}), \frac{\Sigma_Y(\tilde{\theta}) \left(c - Z_{\tilde{\theta}}(\tilde{\theta}) + Z_{\tilde{\theta}}(0) \right)}{\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0)} \right\},$$

for $\mathcal{L}(\tilde{\theta}, Z_{\tilde{\theta}})$ defined as in Proposition 3. Hence, when $\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0) > 0$, conditioning on $\hat{\gamma} = 1$ simply modifies the lower bound $\mathcal{L}(\tilde{\theta}, Z_{\tilde{\theta}})$. Likewise, when $\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0) < 0$ or $\Sigma_{XY}(\tilde{\theta}) - \Sigma_{XY}(\tilde{\theta}, 0) = 0$, conditioning on $\hat{\gamma} = 1$ modifies $\mathcal{U}(\tilde{\theta}, Z_{\tilde{\theta}})$ and $\mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}})$, respectively. \triangle

Threshold Regression Estimation (continued) Suppose that we report estimates and confidence sets for the change parameter $\delta_j(\hat{\theta})$ only if we reject the null hypothesis of no threshold, $H_0: \delta(\theta) = 0$ for all $\theta \in \Theta$. Suppose, in particular, that we test this hypothesis with the sup-Wald test of Andrews (1993). Analogous results to those shown in Elliott and Müller (2014) provide that in our setting, such a test rejects asymptotically if and only if $\|X(\hat{\theta})\| > c$ for a critical value c that depends on Σ . We can set $\gamma(X) = 1 \left\{ \|X(\hat{\theta})\| > c \right\}$ and it is again natural to condition inference on $\hat{\gamma} = 1$.

In this setting $\mathcal{X}_\gamma(1) = \left\{ X : \|X(\hat{\theta})\| > c \right\}$. As before, the expressions for the conditioning sets are involved but straightforward to compute. In particular, for $\bar{\mathcal{V}}(Z_{\tilde{\theta}})$, $\bar{\mathcal{L}}(Z_{\tilde{\theta}})$, and $\bar{\mathcal{U}}(Z_{\tilde{\theta}})$ defined in Section B.2 of the supplement, if $\bar{\mathcal{V}}(Z_{\tilde{\theta}}) \geq 0$ then $\mathcal{Y}_\gamma(1, Z_{\tilde{\theta}}) = (\bar{\mathcal{L}}(Z_{\tilde{\theta}}), \bar{\mathcal{U}}(Z_{\tilde{\theta}}))^c$, where \mathcal{S}^c denotes the complement of a generic set \mathcal{S} . Thus,

$$\mathcal{Y}(\tilde{\theta}, 1, Z_{\tilde{\theta}}) = (\bar{\mathcal{L}}(Z_{\tilde{\theta}}), \bar{\mathcal{U}}(Z_{\tilde{\theta}}))^c \cap \bigcap_{\theta \in \Theta: A(\tilde{\theta}, \theta) > 0, D_Z(\tilde{\theta}, \theta) \geq 0} \left[\ell_Z^1(\tilde{\theta}, u_Z^1(\tilde{\theta}, \theta)) \cup \left[\ell_Z^2(\tilde{\theta}, \theta), u_Z^2(\tilde{\theta}) \right] \right]$$

when $\min \left\{ \mathcal{V}(\tilde{\theta}, Z_{\tilde{\theta}}), \bar{\mathcal{V}}(Z_{\tilde{\theta}}) \right\} \geq 0$. Details and expressions under other realizations of $\bar{\mathcal{V}}(Z_{\tilde{\theta}})$ can be found in Section B.2 of the supplement. \triangle

As these example illustrate, it is straightforward to incorporate additional conditioning variables $\hat{\gamma}$ in both the level and norm maximization problems provided one can characterize the set $\mathcal{Y}_\gamma(\tilde{\gamma}, z)$. While such characterizations are easy to obtain in many cases, they depend on the conditioning variable considered and must be derived on a case-by-case basis.

4.3 Comparison to Sample Splitting

A common remedy in practice for the problems we study is to split the sample. If we have iid observations and select $\hat{\theta}^1$ based on the first half of the data, conventional estimates and confidence intervals for $\mu_Y(\hat{\theta}^1)$ that use only the second half of the data will be

(conditionally) valid. Hence, it is natural to ask how our conditioning approach compares to this conventional sample splitting approach.

For ease of exposition, in this section we focus on even sample splits. Asymptotically, such splits yield a pair of independent and identically distributed normal draws (X^1, Y^1) and (X^2, Y^2) , both of which follow (6), albeit with a different scaling for (μ, Σ) than in the full-sample case.²³ Sample splitting procedures calculate $\hat{\theta}^1$ as in (7) and (8) for level and norm maximization, respectively, replacing X by X^1 . Inference on $\mu_Y(\hat{\theta}^1)$ is then conducted using (X^2, Y^2) . In particular, the conventional 95% sample-splitting confidence interval for $\mu_Y(\hat{\theta}^1)$,

$$\left[Y^2(\hat{\theta}^1) - 1.96\sqrt{\Sigma_Y(\hat{\theta}^1)}, Y^2(\hat{\theta}^1) + 1.96\sqrt{\Sigma_Y(\hat{\theta}^1)} \right],$$

has correct (conditional) coverage and $Y^2(\hat{\theta}^1)$ is a median-unbiased estimator for $\mu_Y(\hat{\theta}^1)$.

While conventional sample splitting resolves the inference problem, this comes at a cost. First, $\hat{\theta}^1$ is based on less data than in the full-sample case, which is unappealing since a policy recommendation estimated with a smaller sample size leads to a lower expected welfare (see, e.g., Theorems 2.1 and 2.2 in Kitagawa and Tetenov (2018b)). Moreover, even after conditioning on $\hat{\theta}^1$, the full-sample average $\frac{1}{2}(X^1, Y^1) + \frac{1}{2}(X^2, Y^2)$ remains a minimal sufficient statistic for μ . Hence, using only (X^2, Y^2) for inference sacrifices information.

Fithian et al. (2017) formalize this point and show that conventional sample splitting tests (and thus confidence sets) are inadmissible.²⁴ Motivated by this result, in Section C of the supplement we derive optimal confidence sets and estimates that are valid conditional on $\hat{\theta}^1$. These optimal split-sample procedures involve truncated normal distributions which are difficult to compute, however, so we also propose computationally straightforward alternatives. These alternatives dominate conventional split-sample methods, but are in turn dominated by the (computationally intractable) optimal split-sample procedures. Nevertheless, these computationally straightforward alternative procedures dominate their conventional counterparts by a substantial margin in simulations calibrated to Card et al. (2008) and reported in Section 7.

Splitting the sample changes the target parameter from $\mu_Y(\hat{\theta})$ to $\mu_Y(\hat{\theta}^1)$, so split-sample

²³Section C of the supplement considers cases with general sample splits and describes the scaling for (μ, Σ) . Intuitively, the scope for improvement over conventional split-sample inference is increasing in the fraction of the data used to construct X_1 .

²⁴Corollary 1 of Fithian et al. (2017) applied in our setting shows that for any sample splitting test based on Y^2 , there exists a test that uses the full data and has weakly higher power against all alternatives and strictly higher power against some alternatives.

approaches are not directly comparable to our full-sample conditioning approach developed above. Nonetheless, while conventional sample splitting methods are dominated, calculating $\hat{\theta}^1$ based on only part of the data may increase the amount of information available for inference and so allow tighter confidence intervals. Thus, depending on how we weight noisier choices of θ against more precise inference on $\mu_Y(\hat{\theta})$, it may be helpful to split the sample and use a procedure that dominates conventional split-sample inference. See Tian and Taylor (2016) and Tian et al. (2016) for related discussions.

4.4 Behavior When $Pr_\mu\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\}$ is Large

As discussed in Section 2, if we ignore selection and compute the conventional (or “naive”) estimator $\hat{\mu}_N=Y(\hat{\theta})$ and the conventional confidence set

$$CS_N = \left[Y(\hat{\theta}) - c_{\alpha/2, N} \sqrt{\Sigma_Y(\hat{\theta})}, Y(\hat{\theta}) + c_{\alpha/2, N} \sqrt{\Sigma_Y(\hat{\theta})} \right] \quad (24)$$

where $c_{\alpha, N}$ is the $1 - \alpha$ -quantile of the standard normal distribution, $\hat{\mu}_N$ is biased and CS_N has incorrect coverage conditional on $\hat{\theta} = \tilde{\theta}$, $\hat{\gamma} = \tilde{\gamma}$. These biases are mild when $Pr_\mu\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\}$ is close to one, however, since in this case the conditional distribution is close to the unconditional one. Intuitively, $Pr_\mu\{\hat{\theta}=\tilde{\theta}\}$ is close to one for some $\tilde{\theta}$ when $\mu_X(\theta)$ or $\|\mu_X(\theta)\|$ has a well-separated maximum in the level and norm maximization problems, respectively. This section shows that our procedures converge to conventional ones in this case.

In particular, suppose first that for some sequence of values $\mu_{Y, m}$ and $z_{\tilde{\theta}, m}$ the probability that $\hat{\theta}=\tilde{\theta}$ and $\hat{\gamma}=\tilde{\gamma}$, conditional on $Z_{\tilde{\theta}}=z_{\tilde{\theta}, m}$, converges to one as $m \rightarrow \infty$. Then our conditional confidence sets and estimates converge to the usual confidence sets and estimates.

Lemma 1

Consider any sequence of values $\mu_{Y, m}$ and $z_{\tilde{\theta}, m}$ such that $Pr_{\mu_{Y, m}}\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma} | Z_{\tilde{\theta}}=z_{\tilde{\theta}, m}\} \rightarrow 1$. Then under $\mu_{Y, m}$, conditional on $\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}, Z_{\tilde{\theta}}=z_{\tilde{\theta}, m}\}$ we have $CS_U \rightarrow_p CS_N$, $CS_{ET} \rightarrow_p CS_N$, and $\hat{\mu}_{\frac{1}{2}} \rightarrow_p Y(\tilde{\theta})$, where for confidence sets \rightarrow_p denotes convergence in probability of the endpoints.

Lemma 1 discusses probabilities conditional on $Z_{\tilde{\theta}}$. If we consider a sequence of values μ_m such that $Pr_{\mu_m}\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\} \rightarrow_p 1$, the same result holds when conditioning only on $\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\}$ and unconditionally.

Proposition 5

Consider any sequence of values μ_m such that $Pr_{\mu_m}\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\} \rightarrow 1$. Then under μ_m ,

we have $CS_U \rightarrow_p CS_N$, $CS_{ET} \rightarrow_p CS_N$, and $\hat{\mu}_{\frac{1}{2}} \rightarrow_p Y(\tilde{\theta})$ both conditional on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}\}$ and unconditionally.

These results provide an additional argument for using our procedures: they remain valid when conventional procedures fail, but coincide with conventional procedures when the latter are valid. On the other hand, as we saw in Section 2, there are cases where our conditional procedures have poor unconditional performance.

5 Unconditional Inference

Rather than requiring validity conditional on $(\hat{\theta}, \hat{\gamma})$ we can instead require coverage only on average, yielding the unconditional coverage requirement

$$Pr\left\{\mu(\hat{\theta}) \in CS\right\} \geq 1 - \alpha \text{ for all } \mu. \quad (25)$$

All confidence sets with correct conditional coverage in the sense of (15) also have correct unconditional coverage provided $\hat{\theta}$ is unique with probability one.

Proposition 6

Suppose that $\hat{\theta}$ is unique with probability one for all μ . Then any confidence set CS with correct conditional coverage (15) also has correct unconditional coverage (25).

Uniqueness of $\hat{\theta}$ implies that the conditioning events $\mathcal{X}(\tilde{\theta}, \tilde{\gamma})$ partition the support of X with measure zero overlap. The result then follows from the law of iterated expectations.

A sufficient condition for almost sure uniqueness of $\hat{\theta}$ is that Σ_X has full rank. A weaker sufficient condition is given in the next lemma. Cox (2018) gives sufficient conditions for uniqueness of a global optimum in a much wider class of problems.

Lemma 2

Suppose that for all $\theta, \tilde{\theta} \in \Theta$ such that $\theta \neq \tilde{\theta}$, either $Var\left(X(\theta)|X(\tilde{\theta})\right) \neq 0$ or $Var\left(X(\tilde{\theta})|X(\theta)\right) \neq 0$. Then $\hat{\theta}$ is unique with probability one for all μ .

While the conditional confidence sets derived in the last section are unconditionally valid, unconditional coverage is less demanding than conditional coverage. Hence, if we are only concerned with unconditional coverage, relaxing the coverage requirement from (15) to (25) may allow us to obtain shorter confidence sets in some settings.

In this section we explore the benefits of such a relaxation. We begin by introducing unconditional confidence sets based on projections of simultaneous confidence bands for

μ . We then introduce hybrid confidence sets that combine projection confidence sets with conditioning arguments. We do not know of estimators for $\mu_Y(\hat{\theta})$ that are unconditionally α -quantile-unbiased but not conditionally unbiased, but introduce hybrid estimators which substantially reduce variability at the cost of permitting a small unconditional bias.

5.1 Projection Confidence Sets

One approach to obtain an unconditional confidence set for $\mu_Y(\hat{\theta})$ is to start with a joint confidence set for μ and project on the dimension corresponding to $\hat{\theta}$. This approach was used by Kitagawa and Tetenov (2018a) for inference in EWM, and by Romano and Wolf (2005) in the context of multiple testing. This approach has also been used in a large and growing statistics literature on post-selection inference including e.g. Berk et al. (2013), Bachoc et al. (2017), Kuchibhotla et al. (2018), and Bachoc et al. (2018). Laber and Murphy (2011) consider a variant of projection for inference on the generalization error of an estimated classifier, obtaining a smaller critical value via a first-stage pretest with a divergent critical value.

To formally describe the projection approach, let c_α denote the $1 - \alpha$ quantile of $\max_{\theta} |\xi(\theta)| / \sqrt{\Sigma_Y(\theta)}$ for $\xi \sim N(0, \Sigma_Y)$. If we define

$$CS_\mu = \left\{ \mu : |Y(\theta) - \mu_Y(\theta)| \leq c_\alpha \sqrt{\Sigma_Y(\theta)} \text{ for all } \theta \in \Theta \right\},$$

then CS_μ is a level $1 - \alpha$ confidence set for μ .²⁵ If we then define

$$CS_P = \left\{ \tilde{\mu}_Y(\hat{\theta}) : \exists \mu \in CS_\mu \text{ such that } \mu_Y(\hat{\theta}) = \tilde{\mu}_Y(\hat{\theta}) \right\} = \left[Y(\hat{\theta}) - c_\alpha \sqrt{\Sigma_Y(\hat{\theta})}, Y(\hat{\theta}) + c_\alpha \sqrt{\Sigma_Y(\hat{\theta})} \right]$$

as the projection of CS_μ on the parameter space for $\mu_Y(\hat{\theta})$, then since $\mu \in CS_\mu$ implies $\mu_Y(\hat{\theta}) \in CS_P$, CS_P satisfies the unconditional coverage requirement (25). As noted in Section 2, however, CS_P does not generally have correct conditional coverage.

The width of the confidence set CS_P depends on the variance $\Sigma_Y(\hat{\theta})$ but does not otherwise depend on the data. To account for the randomness of $\hat{\theta}$, the critical value c_α is larger than the conventional two-sided normal critical value. This means that CS_P will be conservative in cases where $\hat{\theta}$ takes a given value $\tilde{\theta}$ with high probability. To improve performance in this case, we next consider hybrid confidence sets.

²⁵Note that we consider a studentized confidence band that adjusts the width based on $\Sigma_Y(\hat{\theta})$, while Kitagawa and Tetenov (2018a) consider an unstudentized band. Romano and Wolf (2005) argue for studentization in a closely related problem.

5.2 Hybrid Confidence Sets

As shown in Section 2, conditional and projection confidence sets each have good unconditional performance in some cases, but neither is fully satisfactory. Hybrid confidence sets combine these procedures to obtain good performance over a wide range of parameter values.

Hybrid confidence sets are constructed to be subsets of the level $1-\beta$ projection confidence set CS_P^β for $0 \leq \beta < \alpha$. A hybrid confidence set collects the values $\mu_{Y,0} \in CS_P^\beta$ not rejected by a hybrid test. Like our conditional tests, hybrid tests of $H_0: \mu_Y(\tilde{\theta}) = \mu_{Y,0}$ condition on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}\}$, but they further condition on the event that the null value is contained in the projection confidence set, i.e. $\mu_{Y,0} \in CS_P^\beta$. This changes the conditioning event to

$$\mathcal{Y}^H(\tilde{\theta}, \tilde{\gamma}, \mu_{Y,0}, z) = \mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, z) \cap \left[\mu_{Y,0} - c_\beta \sqrt{\Sigma_Y(\tilde{\theta})}, \mu_{Y,0} + c_\beta \sqrt{\Sigma_Y(\tilde{\theta})} \right]$$

for c_β as defined in Section 5.1.

Similarly to our conditional confidence sets, we construct hybrid confidence sets by inverting both equal-tailed and uniformly most powerful unbiased hybrid tests. To construct the equal-tailed test, we define $\phi_{OS-, \alpha}^H$ and $\phi_{OS+, \alpha}^H$ analogously to $\phi_{OS-, \alpha}$ and $\phi_{OS+, \alpha}$ in (22) and (23), respectively, using the conditioning event $\mathcal{Y}^H(\tilde{\theta}, \tilde{\gamma}, \mu_{Y,0}, Z_{\tilde{\theta}})$ rather than $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, Z_{\tilde{\theta}})$. The equal-tailed hybrid test of $H_0: \mu_Y(\tilde{\theta}) = \mu_{Y,0}$ is

$$\phi_{ET, \alpha}^H(\mu_{Y,0}) = \max\{\phi_{OS-, \alpha/2}^H(\mu_{Y,0}), \phi_{OS+, \alpha/2}^H(\mu_{Y,0})\},$$

which rejects if either of the upper or lower size $\alpha/2$ one-sided tests rejects. The level $1-\alpha$ equal-tailed hybrid confidence set is $CS_{ET}^H = \left\{ \mu_{Y,0} \in CS_P^\beta : \phi_{ET, \frac{\alpha-\beta}{1-\beta}}^H(\mu_{Y,0}) = 0 \right\}$, which collects the set of values in CS_P^β which are not rejected by $\phi_{ET, \frac{\alpha-\beta}{1-\beta}}^H$.

To form a hybrid confidence set based on inverting unbiased tests, we likewise define $\phi_{TS, \alpha}^H$ analogously to $\phi_{TS, \alpha}$ in (21), using the conditioning event $\mathcal{Y}^H(\tilde{\theta}, \tilde{\gamma}, \mu_{Y,0}, Z_{\tilde{\theta}})$ rather than $\mathcal{Y}(\tilde{\theta}, \tilde{\gamma}, Z_{\tilde{\theta}})$. By the results of Proposition 2, we know that $\phi_{TS, \alpha}^H(\mu_{Y,0})$ is the uniformly most powerful level α unbiased test of $H_0: \mu_Y(\tilde{\theta}) = \mu_{Y,0}$ conditional on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}, \mu_{Y,0} \in CS_P^\beta\}$. The corresponding level $1-\alpha$ confidence set is then $CS_U^H = \left\{ \mu_{Y,0} \in CS_P^\beta : \phi_{U, \frac{\alpha-\beta}{1-\beta}}^H(\mu_{Y,0}) = 0 \right\}$.

For $\beta=0$ the hybrid confidence sets coincide with the conditional confidence sets CS_{ET} and CS_U . For $\beta>0$ on the other hand, the hybrid confidence sets are contained in CS_P^β and the level of hybrid tests that condition on $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}, \mu_{Y,0} \in CS_P^\beta\}$ are correspondingly

adjusted to $\frac{\alpha-\beta}{1-\beta}$. This adjustment is necessary because the true value $\mu_Y(\hat{\theta})$ sometimes falls outside CS_P^β , and if we do not account for this our hybrid confidence sets may under-cover. With this adjustment, however, hybrid confidence sets have coverage at least $1-\alpha$ both conditionally and unconditionally.

Proposition 7

The hybrid confidence sets CS_{ET}^H and CS_U^H have conditional coverage $\frac{1-\alpha}{1-\beta}$:

$$Pr_\mu \left\{ \mu(\tilde{\theta}) \in CS_{ET}^H | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}, \mu_Y(\tilde{\theta}) \in CS_P^\beta \right\} = \frac{1-\alpha}{1-\beta},$$

$$Pr_\mu \left\{ \mu(\tilde{\theta}) \in CS_U^H | \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}, \mu_Y(\tilde{\theta}) \in CS_P^\beta \right\} = \frac{1-\alpha}{1-\beta},$$

for all $\tilde{\theta} \in \Theta$, $\tilde{\gamma} \in \Gamma$, and all μ . Moreover, provided $\hat{\theta}$ is unique with probability one for all μ , both confidence sets have unconditional coverage between $1-\alpha$ and $\frac{1-\alpha}{1-\beta} \leq 1-\alpha+\beta$:

$$\inf_\mu Pr_\mu \left\{ \mu(\hat{\theta}) \in CS_{ET}^H \right\} \geq 1-\alpha, \quad \sup_\mu Pr_\mu \left\{ \mu(\hat{\theta}) \in CS_{ET}^H \right\} \leq \frac{1-\alpha}{1-\beta},$$

$$\inf_\mu Pr_\mu \left\{ \mu(\hat{\theta}) \in CS_U^H \right\} \geq 1-\alpha, \quad \sup_\mu Pr_\mu \left\{ \mu(\hat{\theta}) \in CS_U^H \right\} \leq \frac{1-\alpha}{1-\beta}.$$

Hybrid confidence sets strike a balance between the conditional and projection approaches. The maximal length of hybrid confidence sets is bounded above by the length of CS_P^β . For small β , hybrid confidence sets will be close to conditional confidence sets and thus to the conventional confidence set when $\{\hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma}\}$ with high probability. However, for $\beta > 0$, hybrid confidence sets do not fully converge to conventional confidence intervals as $Pr_\mu \left\{ \hat{\theta} = \tilde{\theta}, \hat{\gamma} = \tilde{\gamma} \right\} \rightarrow 1$.²⁶ Nevertheless, in our simulations we find the performance of the hybrid and conditional approaches to be quite similar in these well-separated cases.

While hybrid confidence sets combine the conditional and projection approaches, they can yield overall performance more appealing than either. In Section 2 we found that hybrid confidence sets had a shorter median length for many parameter values than did either the conditional or projection approaches used in isolation. Our simulation results in Sections 6 and 7 below provide further evidence of outperformance in realistic settings.

²⁶Indeed, one can directly choose β to yield a given maximal power loss for the hybrid tests relative to conditional tests in the well-separated case. Such a choice of β will depend on Σ , however. For simplicity we instead use $\beta = \alpha/10$ in our simulations. Romano et al. (2014) and McCloskey (2017) find this choice to perform well in two different settings when using a Bonferroni correction.

It is worth contrasting our hybrid approach with more conventional Bonferroni corrections as in e.g. Romano et al. (2014); McCloskey (2017). A simple Bonferroni approach for our setting intersects a level $1-\beta$ projection confidence interval CS_P^β with a level $1-\alpha+\beta$ conditional interval that conditions only on $\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\}$. Bonferroni intervals differ from our hybrid approach in two respects. First, they use a level $1-\alpha+\beta$ conditional confidence interval, while the hybrid approach uses a level $\frac{1-\alpha}{1-\beta}$ conditional interval, where $\frac{1-\alpha}{1-\beta} \leq 1-\alpha+\beta$. Second, the conditional interval used by the Bonferroni approach does not condition on $\mu_Y(\tilde{\theta}) \in CS_P^\beta$, while that used by the hybrid approach does. Consequently, hybrid confidence sets never contains the endpoints of CS_P^β , while the same is not true of Bonferroni intervals.

5.3 Hybrid Estimators

The simulation results of Section 2 showed that our median-unbiased estimator can sometimes be much more dispersed than the conventional estimator $\hat{\mu}=Y(\hat{\theta})$. While we do not know of an alternative approach to construct exactly median-unbiased estimators in our setting, a version of our hybrid approach yields estimators that control both median bias and dispersion relative to $\hat{\mu}=Y(\hat{\theta})$.

To construct hybrid estimators we again condition on both $\{\hat{\theta}=\tilde{\theta}, \hat{\gamma}=\tilde{\gamma}\}$ and $\mu_Y(\tilde{\theta}) \in CS_P^\beta$. Conditional on these events and $Z_{\tilde{\theta}}=z$, we know that $Y(\tilde{\theta})$ again lies in $\mathcal{Y}^H(\tilde{\theta}, \tilde{\gamma}, \mu_Y(\tilde{\theta}), z)$. Let $F_{TN}^H(y; \mu_Y(\tilde{\theta}), \tilde{\theta}, \tilde{\gamma}, z)$ denote the conditional distribution function of $Y(\tilde{\theta})$, and define $\hat{\mu}_\alpha^H$ to solve $F_{TN}^H(Y(\hat{\theta}); \hat{\mu}_\alpha^H, \hat{\theta}, \hat{\gamma}, Z_{\hat{\theta}}) = 1-\alpha$.

Proposition 8

For $\alpha \in (0,1)$, $\hat{\mu}_\alpha^H$ is unique and $\hat{\mu}_\alpha^H \in CS_P^\beta$. If $\hat{\theta}$ is unique almost surely for all μ , $\hat{\mu}_\alpha^H$ is α -quantile-unbiased conditional on $\mu_Y(\hat{\theta}) \in CS_P^\beta$:

$$Pr_\mu \left\{ \hat{\mu}_\alpha^H \geq \mu_Y(\hat{\theta}) \mid \mu_Y(\hat{\theta}) \in CS_P^\beta \right\} = \alpha \text{ for all } \mu.$$

Proposition 8 implies several notable properties for the hybrid estimator. First, since $Pr_\mu \left\{ \mu_Y(\hat{\theta}) \in CS_P^\beta \right\} \geq 1-\beta$ by construction, one can show that

$$\left| Pr_\mu \left\{ \hat{\mu}_\alpha^H \geq \mu_Y(\hat{\theta}) \right\} - \alpha \right| \leq \beta \cdot \max\{\alpha, 1-\alpha\} \text{ for all } \mu.$$

This implies that the absolute median bias of $\hat{\mu}_{\frac{1}{2}}^H$ (measured as the deviation of the exceedance probability from $1/2$) is bounded above by $\beta/2$. On the other hand, since $\hat{\mu}_{\frac{1}{2}}^H \in CS_P^\beta$ we have $\left| \hat{\mu}_{\frac{1}{2}}^H - Y(\hat{\theta}) \right| \leq c_\beta \sqrt{\Sigma_Y(\tilde{\theta})}$, so the difference between $\hat{\mu}_{\frac{1}{2}}^H$ and the conventional esti-

mator $Y(\hat{\theta})$ is bounded above by half the width of CS_P^β . As β varies, the hybrid estimator interpolates between the median-unbiased estimator $\hat{\mu}_{\frac{1}{2}}$ and the conventional estimator $Y(\hat{\theta})$.

6 Simulations: Empirical Welfare Maximization

Our first set of simulations considers the EWM setting introduced in Section 3. We calibrate our simulations to experimental data from the National Job Training Partnership Act (JTPA) Study, which was previously used by Kitagawa and Tetenov (2018b) to study empirical welfare maximization. For a detailed description of the study see Bloom et al. (1997).

We have data on $n = 11,204$ individuals i and the treatment D_i is binary; $D_i = 1$ indicates assignment to a job training program and $D_i = 0$ indicates non-assignment. The probability of assignment is constant: $d(c) = \Pr(D_i = 1 | C_i = c) = 2/3$. We consider rules that allocate treatment based on years of education C_i . In the data, C takes integer values ranging from 6 to 18 years. As in Section 3, rule θ assigns i to treatment if and only if $C_i \in \mathcal{C}_\theta$.

We consider two classes of policies. The first, which we call threshold policies, treat all individuals with fewer than θ years of education: $\mathcal{C}_\theta = \{C : C \leq \theta\}$. The second, which we call interval policies, treat all individuals with between θ_l and θ_u years of education: $\mathcal{C}_\theta = \{C : \theta_l \leq C \leq \theta_u\}$, where a policy θ consists of a (θ_l, θ_u) pair. The total number of policies $|\Theta|$ is equal to 13 and 91 for the threshold and interval cases, respectively. We define $X_n(\theta)$ as a scaled estimate for the increase in income from policy θ relative to the baseline of no treatment. For Y_i individual income measured in hundreds of thousands of dollars,

$$X_n(\theta) = \frac{1}{\sqrt{n}} \sum_{i=1}^n \left(\frac{Y_i D_i}{d(C_i)} 1_{\{C_i \in \mathcal{C}_\theta\}} - \frac{Y_i (1 - D_i)}{1 - d(C_i)} 1_{\{C_i \notin \mathcal{C}_\theta\}} \right),$$

and we consider inference on the average increase in income, so $Y_n = X_n$.

For our simulations, we focus on the asymptotic problem and draw normal vectors X with known variance Σ_X equal to a (consistent) estimate for the asymptotic variance of X_n based on the JTPA data and take $\hat{\theta} = \arg\max_\theta X(\theta)$. The object of interest is thus $\mu_X(\hat{\theta})$. The mean vector $\mu_{X,n}$ of X_n is not consistently estimable due to the \sqrt{n} scaling, so we consider three specifications for the mean μ_X of X . Specification (i) sets $\mu_X = 0$, so all policies yield the same welfare as the baseline of no treatment. Specification (ii) sets $\mu_X = (0, -10^5, \dots, -10^5)$, so one policy is vastly more effective than the others. Finally, specification (iii) sets $\mu_X = X_n$ for X_n calculated in the JTPA data. Intuitively, we expect that specification (i) will be unfavorable to conditional confidence sets since in Section 2 these performed poorly when all policies were equally effective. Specification (ii) should

be favorable to conditional confidence sets since in this case $\hat{\theta}$ selects one policy with high probability, and the results of Section 4.4 apply. Finally, specification (iii) is calibrated to the data and it is not obvious which approaches will perform well in this setting.

To the best of our knowledge our conditional confidence sets are the only known procedures available with correct conditional coverage given $\hat{\theta}$. Hence, we focus on unconditional performance and compare the conditional confidence sets CS_{ET} and CS_U and the hybrid confidence sets CS_{ET}^H and CS_U^H to the projection confidence set CS_P . The conditional and hybrid confidence sets are novel to this paper, but (unstudentized) projection confidence sets were previously considered for this problem by Kitagawa and Tetenov (2018a). We take $\alpha=0.05$ in all cases and so consider 95% confidence sets. For hybrid confidence sets we set $\beta=\alpha/10=.005$. All reported results are based on 10^4 simulation draws.

Table 1 reports the unconditional coverage $Pr_{\mu}\{\mu_X(\hat{\theta}) \in CS\}$ of all five confidence sets, along with the conventional confidence set CS_N as in (24). As expected, all confidence sets other than CS_N have correct coverage in all settings considered. The conditional confidence sets are exact, with coverage equal to 95% up to simulation error. By contrast, hybrid confidence sets tend to be slightly conservative, and projection confidence sets are often quite conservative, with coverage close to one when we consider interval policies.

Table 1: Unconditional Coverage Probability

DGP	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H	CS_P	CS_N
Class of Threshold Policies						
(i)	0.949	0.950	0.952	0.953	0.986	0.922
(ii)	0.952	0.952	0.956	0.956	0.991	0.952
(iii)	0.95	0.95	0.955	0.955	0.992	0.952
Class of Interval Policies						
(i)	0.952	0.949	0.956	0.953	0.992	0.837
(ii)	0.95	0.951	0.954	0.954	0.998	0.950
(iii)	0.951	0.95	0.954	0.955	0.998	0.948

We next compare the length of confidence sets. Projection confidence sets were proposed in the previous literature and their length is proportional to the standard error $\sqrt{\Sigma_X(\hat{\theta})}$ for the welfare of the estimated optimal policy. Hence, CS_P provides a natural benchmark against which to compare the length of our new confidence sets. In Table 2 we compare our new confidence sets to this benchmark in two ways, first reporting the median lengths of CS_{ET} , CS_U , CS_{ET}^H , and CS_U^H relative to CS_P (that is, the ratio of the median of their lengths), and then reporting the fraction of simulation draws for which

our new confidence sets are longer than CS_P .

Focusing first on specification (i) for which $\mu_X=0$, we see that conditional confidence sets are longer than CS_P according to both measures in the threshold and interval policy specifications. Hence, as expected, this case is unfavorable to these confidence sets. By contrast, our hybrid confidence sets are shorter than the projection sets both in median length and in the substantial majority of simulation draws. Turning next to specification (ii) for which μ_X has a well-separated maximum, we see that, as expected, conditional confidence sets are much shorter than projection confidence sets. Hybrid confidence sets perform nearly as well. Finally in specification (iii) for which μ_X is calibrated to the data, we see that the performance of the conditional sets is between its performance in cases (i) and (ii), and that hybrid confidence sets again perform best.

Overall, these simulation results favor the hybrid confidence sets relative to both the conditional and projection sets. The benefits of hybrid confidence sets are still more pronounced if we consider higher quantiles of the length distribution, reported in Section F of the supplement. We do not find a strong advantage for either CS_{ET}^H or CS_U^H , though when the two differ CS_{ET}^H typically performs better. Since CS_{ET}^H is also typically easier to calculate, these simulation results suggest using CS_{ET}^H in this setting.

Table 2: Length of Confidence Sets Relative to CS_P in EWM Simulations

DGP	Median Length Relative to CS_P				Probability Longer than CS_P			
	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H
Class of Threshold Policies								
(i)	1.17	1.27	0.63	0.64	0.71	0.80	0.04	0.35
(ii)	0.75	0.75	0.76	0.76	0	0	0	0
(iii)	0.84	0.93	0.84	0.89	0.33	0.43	0	0.19
Class of Interval Policies								
(i)	1.54	1.65	0.77	0.76	0.79	0.88	0	0
(ii)	0.63	0.64	0.65	0.65	0	0	0	0
(iii)	0.78	0.88	0.76	0.81	0.32	0.42	0	0

We next consider the properties of our point estimators. The initial columns of Table 3 report the simulated median bias of our median unbiased estimator $\hat{\mu}_{\frac{1}{2}}$, our hybrid estimator $\hat{\mu}_{\frac{1}{2}}^H$, and the conventional estimator $X(\hat{\theta})$, measured both as the difference in the exceedance probability from $\frac{1}{2}$ and as the median studentized estimation error. The hybrid estimator is quite close to being median unbiased. By contrast, the conventional estimator exhibits substantial bias when μ_X does not have a well-separated maximum.

The final three columns of Table 3 report the median absolute studentized error for the estimators considered. These results show that the median unbiased estimator $\hat{\mu}_{\frac{1}{2}}$ has a larger median absolute error than the conventional estimator $X(\hat{\theta})$ in all designs except the well-separated case (ii), where all three estimators perform similarly. The hybrid estimator $\hat{\mu}_{\frac{1}{2}}^H$ likewise has a larger median absolute error than the conventional estimator. Additional results reported in Section F of the supplement show that the hybrid estimator substantially outperforms the median unbiased estimator when one considers higher quantiles of absolute error.

Table 3: Bias and Median Absolute Error of Point Estimators

DGP	$Pr_{\mu}\left\{\hat{\mu} > \mu_X(\hat{\theta})\right\} - \frac{1}{2}$			$Med_{\mu}\left(\frac{\hat{\mu} - \mu_X(\hat{\theta})}{\sqrt{\Sigma_X(\hat{\theta})}}\right)$			$Med_{\mu}\left(\frac{ \hat{\mu} - \mu_X(\hat{\theta}) }{\sqrt{\Sigma_X(\hat{\theta})}}\right)$		
	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$X(\hat{\theta})$	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$X(\hat{\theta})$	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$X(\hat{\theta})$
Class of Threshold Policies									
(i)	-0.007	-0.007	0.391	-0.02	-0.02	0.82	1.11	1.10	0.88
(ii)	-0.001	0.001	0.001	0	0	0	0.67	0.67	0.67
(iii)	-0.001	-0.001	0.104	0	0	0.25	0.80	0.79	0.67
Class of Interval Policies									
(i)	0	0.003	0.5	0	0.02	1.3	1.42	1.39	1.30
(ii)	-0.002	0.001	0.001	0	0	0	0.65	0.65	0.66
(iii)	0	0.001	0.148	0	0	0.35	0.86	0.86	0.69

The results of this section confirm our theoretical findings. Conditional confidence sets and estimators perform well when the optimal policy is well-separated but can otherwise underperform existing alternatives. Hybrid confidence sets outperform existing alternatives in all cases, nearly matching conditional confidence sets in well-separated cases while maintaining much better performance in other settings. Finally, hybrid estimators eliminate almost all median bias while obtaining a substantially smaller median absolute error than the exact median-unbiased estimator. Hence, we find strong evidence favoring our hybrid confidence sets relative to the available alternatives and evidence favoring our hybrid estimators if bias reduction is desired.

7 Simulations: Tipping Point Estimation

Our second set of simulation results is based on the tipping point model of Card et al. (2008), a leading application of the threshold regression model discussed throughout this paper as a running example. Card et al. (2008) study the evolution of neighborhood

composition as a function of minority population share. In particular, for Y_i the normalized change in the white population of census tract i between 1980 and 1990, C_i a vector of controls, and Q_i the minority share in 1980, Card et al. (2008) consider the specification

$$Y_i = \beta + C_i' \alpha + \delta 1\{Q_i > \theta\} + U_i,$$

which allows the white population share to change discontinuously when the minority share exceeds some threshold θ . They then fit this model, including the break point θ , by least squares. See Card et al. (2008) for details on the data and motivation. We consider data from Chicago and Los Angeles with $n=1,820$ and $n=2,035$ observations, respectively, estimating the model separately in each city.²⁷

Results in Wang (2018) show that if we model the coefficient δ as on the same order as sampling uncertainty, this threshold regression model satisfies the high-level conditions (12)–(13) we introduced in Section 3. Hence, we can immediately apply our results for the norm-maximization problem to the present setting. Specifically, we define X_n as discussed in Section 3 and $\hat{\theta}_n$ is again asymptotically equivalent to the solution to a norm-maximization problem $\arg\max_{\theta \in \Theta} \|X(\theta)\|$.²⁸ We define $Y_n(\theta) = \sqrt{n}\hat{\delta}(\theta)$ to be proportional to the estimated change coefficient imposing tipping point θ , so we again consider the problem of inference on the change coefficient while acknowledging randomness in the estimated threshold.

Our simulations draw normal random vectors (X, Y) from the limiting normal model derived in Section 3. This model depends on the function Σ_C and the covariance function of G in Section 3 which we (consistently) estimate from the Card et al. (2008) data. It also depends on the function $\Sigma_{cg}(\cdot)$. Since this is not consistently estimable, we consider three specifications. Specification (i) assumes there is no coefficient change, corresponding to $\delta=0$. Specification (ii) assumes that there is a single large change, setting $\delta=-100\%$ and taking the true threshold to equal the estimate in the Card et al. (2008) data. Finally, specification (iii) calibrates $\Sigma_{cg}(\cdot)$ to the data, corresponding to the analog of model (10) where the intercept term in the regression may depend arbitrarily upon a neighborhood's minority share. This specification implies that the break model is misspecified but as discussed above, our approach remains applicable in this case, unlike the results of Wang

²⁷We focus on these cities following Wang (2017), a previous version of Wang (2018), since Card et al. (2008) note that their tipping point estimation method appears more appropriate for larger cities.

²⁸While Card et al. (2008) optimize over all possible tipping points between 5% and 60%, consistent with our theoretical results we limit attention to a finite set of thresholds. In particular, we consider 100 evenly-spaced quantiles of the minority share, and then further restrict attention to thresholds between 5% and 60%. We also tried several other discretization schemes and found very similar results in all cases.

(2018). Indeed, Card et al. (2008) acknowledge that the tipping point model only approximates their underlying theoretical model of neighborhood ethnic composition, so misspecification seems likely in this setting.

We again focus on the unconditional performance of our proposed procedures along with existing alternatives. All reported results are based on 10^4 simulation draws. Table 4 reports coverage for the confidence sets CS_{ET} , CS_U , CS_{ET}^H , CS_U^H , and CS_P , along with the conventional confidence set CS_N . As for the simulations calibrated to the EWM application, we see that all confidence sets other than CS_N have correct coverage, CS_P often over-covers, the conditional confidence sets have exact coverage and the hybrid confidence sets exhibit minimal over-coverage. In this application, the conventional confidence set CS_N severely under-covers for some simulation designs.

Table 4: Unconditional Coverage Probability

DGP	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H	CS_P	CS_N
Chicago Data Calibration						
(i)	0.948	0.95	0.949	0.949	0.95	0.750
(ii)	0.951	0.95	0.956	0.955	0.994	0.951
(iii)	0.947	0.946	0.951	0.951	0.990	0.934
Los Angeles Data Calibration						
(i)	0.949	0.948	0.949	0.948	0.95	0.615
(ii)	0.952	0.952	0.956	0.956	0.996	0.952
(iii)	0.951	0.949	0.955	0.954	0.996	0.95

Table 5 compares the lengths of our confidence sets to that of CS_P . For each confidence set we again report both median length relative to CS_P and the frequency with which the confidence set is longer than CS_P . Here we see that the conditional confidence sets can be relatively long, while the hybrid confidence sets provide marked performance improvements across the specifications considered. Similarly to the simulation exercises of the previous section, the benefits of the hybrid confidence sets can become even more pronounced at different length quantiles. See Section G of the supplemental appendix. Remarkably, neither of the hybrid confidence sets is longer than CS_P in any simulation draw across all specifications examined. The overall message is similar to that of the previous section: hybrid confidence sets possess clear advantages for unconditional inference and CS_{ET}^H seems to be the most compelling option, especially given its computational simplicity.

Finally, we consider the properties of our point estimators. The initial columns of Table 6 report median bias measured both with the deviation of the exceedance proba-

Table 5: Length of Confidence Sets Relative to CS_P in Tipping Point Simulations

	Median Length Relative to CS_P				Probability Longer than CS_P			
	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H	CS_{ET}	CS_U	CS_{ET}^H	CS_U^H
Chicago Data Calibration								
(i)	1.33	1.38	0.94	0.94	0.83	0.89	0	0
(ii)	0.72	0.72	0.74	0.74	0	0	0	0
(iii)	0.82	0.93	0.82	0.87	0.35	0.44	0	0
Los Angeles Data Calibration								
(i)	1.26	1.29	0.86	0.85	0.58	0.62	0	0
(ii)	0.68	0.68	0.69	0.69	0	0	0	0
(iii)	0.68	0.70	0.70	0.72	0.15	0.19	0	0

bility from $\frac{1}{2}$ and with the studentized median estimation error. We again see that $\hat{\mu}_{\frac{1}{2}}$ is median-unbiased (up to simulation error) and that $\hat{\mu}_{\frac{1}{2}}^H$ exhibits minimal median bias. By contrast, in specification (i) the conventional estimator $Y(\hat{\theta})$ has substantial median bias as measured by the studentized median estimation error, though very little as measured by the exceedance probability. This latter feature reflects the fact that the density of $Y(\hat{\theta}) - \mu_Y(\hat{\theta})$ has very little mass near zero in this specification.

Turning to median absolute studentized error, we see that all estimators perform similarly when the series has a single large break. By contrast, the median unbiased estimator $\hat{\mu}_{\frac{1}{2}}$ performs better than the conventional estimator $Y(\hat{\theta})$ in specification (i) (no break) but performs worse in specification (iii). The hybrid estimator is weakly better than the unbiased estimator in all cases, with performance gains in case (i) and equal performance in the other two cases. Again, the performance gains are more pronounced if one considers higher quantiles of the absolute error distribution, as reported in Section G of the supplement.

7.1 Split-Sample Procedures

While we have so far compared the performance of our conditional and hybrid procedures to the projection confidence set CS_P and conventional estimator $Y(\hat{\theta})$, Card et al. (2008) instead adopt a sample-splitting approach, using two-thirds of the data to select the break-date and a third of the data for inference. In this section we compare the performance of this conventional split-sample procedure to that of the implementable split-sample alternative developed in Section C of the supplement. We consider the same calibrations to the Card et al. (2008) data as above and choose the sample split as in Card et al. (2008).

Table 7 compares the conventional split-sample confidence set CS_{SS} and estimator $Y^2(\hat{\theta}^1)$ used by Card et al. (2008) to our (equal-tailed) alternative split-sample confidence

Table 6: Bias and Median Absolute Error in Tipping Point Simulations

	$Pr_\mu\left\{\hat{\mu} > \mu_Y(\hat{\theta})\right\} - \frac{1}{2}$			$Med_\mu\left(\frac{\hat{\mu} - \mu_Y(\hat{\theta})}{\sqrt{\Sigma_Y(\hat{\theta})}}\right)$			$Med_\mu\left(\left \frac{\hat{\mu} - \mu_Y(\hat{\theta})}{\sqrt{\Sigma_Y(\hat{\theta})}}\right \right)$		
	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$Y(\hat{\theta})$	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$Y(\hat{\theta})$	$\hat{\mu}_{\frac{1}{2}}$	$\hat{\mu}_{\frac{1}{2}}^H$	$Y(\hat{\theta})$
Chicago Data Calibration									
(i)	0	0	0.01	-0.01	0.01	0.64	1.51	1.38	1.52
(ii)	-0.01	-0.01	-0.01	-0.03	-0.03	-0.03	0.66	0.66	0.66
(iii)	-0.01	-0.01	-0.15	-0.03	-0.03	-0.37	0.83	0.83	0.71
Los Angeles Data Calibration									
(i)	0	0	0	0	0	-0.8	1.38	1.29	1.80
(ii)	0	0	0	0.01	0.01	0.01	0.67	0.67	0.67
(iii)	0	0	0.006	0	-0.01	-0.016	0.74	0.74	0.68

set CS_{SS}^A and median-unbiased estimator $\hat{\mu}_{\frac{1}{2},SS}^A$. See Section C of the supplement for definitions. These results clearly reflect the dominance of our alternative split-sample procedures, with substantial performance improvements for both confidence sets and estimators across all calibrations. These improvements are largest in the well-separated case (ii), but are nearly as large in the data-calibrated case (iii). Section G of the supplement provides ratios of the 5th, 25th, 50th, 75th and 95th quantiles of the lengths of CS_{SS}^A relative to the those of CS_{SS} as well as the quantiles of $|\hat{\mu} - \mu_Y(\hat{\theta}^1)| / \sqrt{\Sigma_Y(\hat{\theta}^1)}$ for $\hat{\mu} = \hat{\mu}_{\frac{1}{2},SS}^A$ and $\hat{\mu} = Y^2(\hat{\theta}^1)$. There, our new split-sample procedures can be seen to dominate the conventional ones across all quantiles and simulation designs considered, often by very wide margins.

Table 7: Performance Measures of Split-Sample Procedures

	Median Length Relative to CS_{SS}	$Med_\mu\left(\frac{ \hat{\mu}-\mu_Y(\hat{\theta}^1) }{\sqrt{\Sigma_Y(\hat{\theta}^1)}}\right)$	
DGP	CS_{SS}^A	$\hat{\mu}_{\frac{1}{2},SS}^A$	$Y^2(\hat{\theta}^1)$
Chicago Data Calibration			
(i)	0.83	0.57	0.67
(ii)	0.58	0.38	0.66
(iii)	0.64	0.44	0.67
Los Angeles Data Calibration			
(i)	0.78	0.55	0.69
(ii)	0.58	0.39	0.67
(iii)	0.59	0.42	0.68

8 Conclusion

This paper considers a form of the winner’s curse that arises when we select a target parameter for inference based on optimization. We propose confidence sets and quantile unbiased estimators for the target parameter that are optimal conditional on its selection. We hence recommend our conditional inference procedures when it is appropriate to remove uncertainty about the choice of target parameters from inferential statements. These conditionally valid procedures are also unconditionally valid, but we find that they sometimes have unappealing (unconditional) performance relative to existing alternatives. If one is satisfied with correct unconditional coverage and (in the case of estimation) a small, controlled degree of bias, we propose hybrid inference and estimation procedures which combine conditioning with projection confidence sets. Examining performance in simulations calibrated to empirical welfare maximization and tipping point applications, we find that our hybrid approach performs well in both cases.

Our results suggest a range of opportunities for future work. First, rather than considering inference on $\mu_Y(\hat{\theta})$, under suitable assumptions one could build on our results to forecast $Y(\hat{\theta})$. Alternatively, while conditional and projection confidence sets have antecedents in the literature on inference after model selection, including in Berk et al. (2013) and Fithian et al. (2017), there is no analog of our hybrid approach in this literature. Our very positive simulation results for the hybrid approach in the present setting suggest that this approach might yield appealing performance in a range of post-selection-inference settings. Even if a fully conditional approach is desired in the post-selection problem, as in Fithian et al. (2017), one could consider the analog of our optimal median-unbiased estimates that condition on the selected model. Finally, the problem of estimating the value of a dynamic treatment rule (c.f. Chakraborty and Murphy, 2014; Han, 2018) is closely related to our level-maximization setting, so it seems likely that our results could prove to be useful there as well.

References

- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61(4):821–856.
- Athey, S. and Wager, S. (2018). Efficient policy learning. *arXiv preprint*. arXiv 1702.0289.
- Bachoc, F., Leeb, H., and Potscher, B. (2018). Valid confidence intervals for post-model-selection predictors. *Annals of Statistics (forthcoming)*.

- Bachoc, F., Preinerstorfer, D., and Steinberger, L. (2017). Uniformly valid confidence intervals post-model-selection. Unpublished Manuscript.
- Baird, S., Hicks, J. H., Kremer, M., and Miguel, E. (2016). Worms at work: Long-run impacts of a child health investment. *The Quarterly Journal of Economics*, 131(4):1637–1680.
- Banerjee, A., Hanna, R., Kyle, J., Olken, B. A., and Sumarto, S. (2018). Tangible information and citizen empowerment: Identification cards and food subsidy programs in Indonesia. *Journal of Political Economy*, 126(2):451–491.
- Berk, R., Brown, L., Buja, A., Zhang, K., and Zhao, L. (2013). Valid post-selection inference. *Annals of Statistics*, 41(2):802–831.
- Bhattacharya, D. (2009). Inferring optimal peer assignment from experimental data. *Journal of the American Statistical Association*, 104:486–500.
- Björkman Nyqvist, M. and Jayachandran, S. (2017). Mothers care more, but fathers decide: Educating parents about child health in Uganda. *American Economic Review*, 107(5):496–500.
- Bloom, H. S., Orr, L. L., Bell, S. H., Cave, G., Doolittle, F., Lin, W., and Bos, J. M. (1997). The benefits and costs of jtpa title ii-a programs: Key findings from the national job training partnership act study. *Journal of Human Resources*, 32(3):549–576.
- Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *Quarterly Journal of Economics*, 123:177–216.
- Chakraborty, B. and Murphy, S. (2014). Dynamic treatment regimes. *Annual Review of Statistics and Its Applications*, 1:447–64.
- Chernozhukov, V., Lee, S., and Rosen, A. (2013). Intersection bounds: Estimation and inference. *Econometrica*, 81(2):667–737.
- Cox, G. (2018). Almost sure uniqueness of a global minimum without convexity. Unpublished Manuscript.
- Dawid, A. P. (1994). Selection paradoxes in bayesian inference. *Institute of Mathematical Statistics Lecture Notes - Monograph Series*, 24:211–220.

- Duflo, E., Greenstone, M., Pande, R., and Ryan, N. (2018). The value of regulatory discretion: Estimates from environmental inspections in India. *Econometrica*, 86(6):2123–2160.
- Efron, B. (2011). Tweedie’s formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614.
- Elliott, G. and Müller, U. K. (2007). Confidence sets for the date of a single break in linear time series regressions. *Journal of Econometrics*, 141(2):1196–1218.
- Elliott, G. and Müller, U. K. (2014). Pre and post break parameter inference. *Journal of Econometrics*, 180:141–157.
- Ferguson, J. P., Cho, J. H., Yang, C., and Zhao, H. (2013). Empirical bayes correction for the winner’s curse in genetic association studies. *Genetic Epidemiology*, 37(1):60–68.
- Fithian, W., Sun, D., and Taylor, J. (2017). Optimal inference after model selection. *arXiv*.
- Graham, B. S., Imbens, G. W., and Ridder, G. (2014). Complementarity and aggregate implications of assortative matching: A nonparametric analysis. *Quantitative Economics*, 5:29–66.
- Han, S. (2018). Identification in nonparametric models for dynamic treatment effects. Unpublished Manuscript.
- Hansen, B. E. (1997). Inference in TAR models. *Studies in Nonlinear Dynamics and Econometrics*, 2:1–14.
- Hansen, B. E. (2000). Sample splitting and threshold estimation. *Econometrica*, 68:575–603.
- Hansen, B. E. (2001). The new econometrics of structural change: Dating breaks in u.s. labor productivity. *Journal of Economic Perspectives*, 15(4):117–128.
- Hansen, P. R. (2005). A test for superior predictive ability. *Journal of Business and Economic Statistics*, 23(4):365–380.
- Harris, X. T., Panigrahi, S., Markovic, J., Bi, N., and Taylor, J. (2016). Selective sampling after solving a convex problem. *arXiv*.
- Hirano, K. and Porter, J. R. (2009). Asymptotics for statistical treatment rules. *Econometrica*, 77:1683–1701.

- Hyun, S., Lin, K., G'Sell, M., and Tibshirani, R. J. (2018). Post-selection inference for changepoint detection algorithms with application to copy number variation data. Unpublished Manuscript.
- Khan, A. Q., Khwaja, A. I., and Olken, B. A. (2016). Tax farming redux: Experimental evidence on performance pay for tax collectors. *The Quarterly Journal of Economics*, 131(1):219–271.
- Kitagawa, T. and Tetenov, A. (2018a). Supplement to “who should be treated? empirical welfare maximization methods for treatment choice”. *Econometrica*.
- Kitagawa, T. and Tetenov, A. (2018b). Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica*, 86(2):591–616.
- Kivaranovic, D. and Leeb, H. (2018). Expected length of post-model-selection confidence intervals conditional on polyhedral constraints. Unpublished Manuscript.
- Kuchibhotla, A., Brown, L., Buja, A., George, E., and Zhao, L. (2018). A model free perspective for linear regression: Uniform-in-model bounds on post selection inference. Unpublished Manuscript.
- Laber, E. and Murphy, S. (2011). Adaptive confidence intervals for the test error in classification. *Journal of the American Statistical Association*, 106(495):904–913.
- Lee, J. D., Sun, D. L., Sun, Y., and Taylor, J. E. (2016). Exact post-selection inference, with application to the LASSO. *Annals of Statistics*, 44:907–927.
- Lee, M. R. and Shen, M. (2018). Winner’s curse: Bias estimation for total effects of features in online controlled experiments. In *KDD*.
- Lehmann, E. and Scheffé, H. (1955). Completeness, similar regions, and unbiased estimation: Part ii. *Sankhyā: The Indian Journal of Statistics*, 15(3):219–236.
- Lehmann, E. L. and Romano, J. P. (2005). *Testing Statistical Hypotheses*. Springer, third ed. edition.
- Manski, C. F. (2004). Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246.

- McCloskey, A. (2017). Bonferroni-based size-correction for nonstandard testing problems. *Journal of Econometrics*, 200(1):17–35.
- Perron, P. (2006). Dealing with structural breaks. In *Palgrave Handbook of Econometrics*, volume 1: Econometric Theory, pages 278–352. Palgrave.
- Pfanzagl, J. (1979). On optimal median unbiased estimators in the presence of nuisance parameters. *Annals of Statistics*, 7(1):187–193.
- Pfanzagl, J. (1994). *Parametric Statistical Theory*. De Gruyter.
- Rai, Y. (2018). Statistical inference for treatment assignment policies. Unpublished Manuscript.
- Romano, J. P., Shaikh, A., and Wolf, M. (2014). A practical two-step method for testing moment inequalities. *Econometrica*, 82(5):1979–2002.
- Romano, J. P. and Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282.
- Tetenov, A. (2012). Statistical treatment choice based on asymmetric minimax regret criteria. *Journal of Econometrics*, 166(1):157–165.
- Tian, X., Bi, N., and Taylor, T. (2016). MAGIC: a general, powerful and tractable method for selective inference. *arXiv*.
- Tian, X. and Taylor, J. (2016). Selective inference with a randomized response. *Annals of Statistics* (forthcoming).
- Tibshirani, R. J., Rinaldo, A., Tibshirani, R., and Wasserman, L. (2018). Uniform asymptotic inference and the bootstrap after model selection. *The Annals of Statistics*, 46(3):1255–1287.
- Wang, Y. (2017). Inference in the threshold model. Unpublished Manuscript.
- Wang, Y. (2018). Inference in the change point model with time-varying moments. Unpublished Manuscript.
- White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5):1097–1126.

Zhong, H. and Prentice, R. (2009). Correcting “winner’s curse” in odds ratios from genomewide association findings for major complex human diseases. *Genetic Epidemiology*, 34(1):78–91.