A careful consideration of CLARIFY: simulation-induced bias in point estimates of quantities of interest

Carlisle Rainey

Department of Political Science, Florida State University, Tallahassee, USA
Email: crainey@fsu.edu

Abstract

Some work in political methodology recommends that applied researchers obtain point estimates of quantities of interest by simulating model coefficients, transforming these simulated coefficients into simulated quantities of interest, and then averaging the simulated quantities of interest (e.g., CLARIFY). But other work advises applied researchers to directly transform coefficient estimates to estimate quantities of interest. I point out that these two approaches are not interchangeable and examine their properties. I show that the simulation approach compounds the transformation-induced bias identified by Rainey (2017), adding bias with direction and magnitude similar to the transformation-induced bias. I refer to this easily avoided additional bias as "simulation-induced bias." Even if researchers use simulation to estimate standard errors, they should directly transform maximum likelihood estimates of coefficient estimates to obtain point estimates of quantities of interest.

1. Introduction

Political scientists employ maximum likelihood (ML) to estimate a variety of statistical models. ML estimators have desirable and widely understood properties. But for many research questions, the model coefficient estimates do not directly interest the researcher. Following King *et al.* (2000), researchers often use the coefficient estimates to compute substantively meaningful quantities of interest, such as predicted probabilities, expected counts, marginal effects, and first differences. The literature offers two methods to compute point estimates for these quantities of interest.

Researchers estimate quantities of interest either by simulating quantities of interest and then averaging (e.g., King *et al.*, 2000) or by directly transforming coefficients into quantities of interest (e.g., Herron, 1999).[1] In practice, researchers' choice between these two approaches seems idiosyncratic rather than principled, depending on their preferred software package rather than any statistical criteria. Further, the methodological literature has not distinguished or compared the two approaches to estimating quantities of interest.

How does the simulation approach compare to directly transforming coefficients? Which should researchers prefer? Or are the two approaches interchangeable, as the literature seems

---

[1]CLARIFY for Stata (Tomz *et al.*, 2003), Zelig for R (Imai *et al.*, 2008; Choirat *et al.*, 2018), and clarify for R (Greifer *et al.* 2023) simulate quantities of interest and find the average. The package dynsimpie for Stata also reports the average of simulations (Philips *et al.*, 2016; Jung *et al.*, 2020). The margins command in Stata (StataCorp, 2017), the margins package in R (Leeper, 2021), and the predict() function in R for the glm (R Core Team, 2018) and polr (Venables and Ripley, 2002) classes directly transform coefficients into the quantities of interest.

to imply? Rainey (2017) shows that directly transforming coefficients into quantities of interest creates "transformation-induced" bias. I show that when researchers use the average of simulations to estimate quantities of interest, they replicate the logic of transformation-induced bias and add an additional, unnecessary bias to their estimates. I refer to this additional bias as "simulation-induced" bias. I show that simulation-induced bias occurs *in addition to* transformation-induced bias and is approximately the same size and direction. While this bias is usually small relative to the standard error, methodologists should not recommend methods that add unnecessary bias to point estimates. Instead, we should recommend methods that better adhere to the usual justifications.

2. The current practice

2.1. The plug-in estimator

First, researchers can directly transform the coefficient estimates into quantities of interest. The invariance property of ML estimators allows a researcher to find the ML estimate of a function of a parameter (i.e., a quantity of interest) by first using ML to estimate the model parameter and then applying the function to that estimate (or "plugging in the estimate") (King, 1998, 75–76; Casella and Berger, 2002, 320–321). I refer to this approach as the "plug-in" estimator. Importantly, the plug-in estimator *remains an ML estimator*. Thus, it retains all of the (desirable) properties of ML estimators (e.g., consistency, asymptotic efficiency).

As a concrete example, suppose a researcher uses ML to estimate a statistical model in which $y_i \sim f(\theta_i)$, where $i \in \{1, \ldots, N\}$ and f represents a probability distribution. The parameter θ_i connects to a design matrix X of k explanatory variables and a column of ones by a link function $g(\cdot)$, so that $g(\theta_i) = X_i\beta$, where $\beta \in \mathbb{R}^{k+1}$ represents a vector of parameters with length $k+1$. The researcher uses ML to compute estimates $\hat{\beta}^{\mathrm{mle}}$ for the parameter vector β. I denote the function that transforms model coefficients into quantities of interest as $\tau(\cdot)$. For example, if the researcher uses a logit model and focuses on the predicted probability for a specific observation X_c, then $\tau(\beta) = \mathrm{logit}^{-1}(X_c\beta) = 1/(1 + e^{-X_c\beta})$. The researcher can use the invariance property to quickly obtain a ML estimate of the predicted probability: $\hat{\tau}^{\mathrm{mle}} = \tau(\hat{\beta}^{\mathrm{mle}}) = \mathrm{logit}^{-1}(X_c\hat{\beta}^{\mathrm{mle}}) = 1/(1 + e^{-X_c\hat{\beta}^{\mathrm{mle}}})$.

2.2. The average-of-simulations estimator

Second, researchers can use the average of simulated quantities of interest as the point estimator. King *et al.* (2000) suggest the following approach:

1. *Fit the model.* Use ML to estimate the model coefficients $\hat{\beta}^{\mathrm{mle}}$ and their covariance $\hat{V}(\hat{\beta}^{\mathrm{mle}})$.
2. *Simulate the coefficients.* Simulate a large number M of coefficient vectors $\tilde{\beta}^{(i)}$, for $i \in \{1, 2, \ldots, M\}$, using $\tilde{\beta}^{(i)} \sim \mathrm{MVN}[\hat{\beta}^{\mathrm{mle}}, \hat{V}(\hat{\beta}^{\mathrm{mle}})]$, where MVN represents the multivariate normal distribution.
3. *Convert simulated coefficients into simulated quantity of interest.* Compute $\tilde{\tau}^{(i)} = \tau(\tilde{\beta}^{(i)})$ for $i \in \{1, 2, \ldots, M\}$. Most quantities of interest depend on the values of the explanatory variables. In this situation, researchers either focus on a specific observation (typically some kind of "average case") or average across all sample observations (Hanmer and Kalkan, 2013). In any case, the transformation $\tau(\cdot)$ includes this choice.[2]

---

[2]As King *et al.* (2000) note, this step might require additional simulation, to first introduce and then average over fundamental uncertainty. I ignore this additional step since (1) it is usually not necessary and (2) including it does not affect my argument.

4. *Average the simulations of the quantity of interest.* Estimate the quantity of interest using the average of the simulations of the quantity of interest, so that $\hat{\tau}^{\mathrm{avg}} = 1/M \sum_{i=1}^{M} \tilde{\tau}^{(i)}$.[3]

I refer to this as the "average-of-simulations" estimator. But what are the properties of this estimator?

While the estimates it provides are sometimes similar to well-behaved plug-in estimates, King et al. (2000) develop their method informally. Much of their theoretical argument happens quickly when they write "we draw many plausible sets of parameters from their posterior or sampling distribution" (349).[4] One might justify their method from a frequentist perspective by first thinking of their method as "informal" Bayesian posterior simulation (Gelman and Hill, 2007). This is helpful because the theory and practice of simulating from a posterior distribution are well-developed and widely understood. The Berstein–von Mises theorem (Van der Vaart, 2000, 140–146) guarantees, under relatively weak assumptions, that posterior simulations are asymptotically equivalent to the simulation procedure suggested by King, Tomz, and Wittenberg. And because the point estimator $\hat{\beta}^{\mathrm{mle}}$ (and functions of $\hat{\beta}^{\mathrm{mle}}$) are consistent, then the mean of the simulations (and the mean of functions of the simulations) are consistent as well. Therefore, one can defend $\hat{\tau}^{\mathrm{avg}}$ on the grounds that it is a consistent estimator of τ.

However, the small sample properties of the average-of-simulations remain poorly understood. Methodologists and applied researchers seem to assume that $\hat{\tau}^{\mathrm{avg}}$ is interchangeable with $\hat{\tau}^{\mathrm{mle}}$. But below, I show that the average-of-simulations algorithm compounds the transformation-induced bias described by Rainey (2017) and adds a similar bias to the estimate that I refer to as "simulation-induced bias."

3. A theory of simulation-induced bias

Before developing the theory of simulation-induced bias, I review the concept of transformation-induced bias from Rainey (2017). As Rainey (2017) shows, transforming unbiased model coefficient estimates can introduce bias into estimates of quantities of interest. Rainey (2017, 404) decomposes the bias in the estimate of the quantity of interest, which he refers to as total τ-bias, into two components: transformation-induced τ-bias and coefficient-induced τ-bias. Rainey (2017) defines these as

$$\text{total } \tau\text{-bias} = \underbrace{\mathrm{E}\big[\tau(\hat{\beta}^{\mathrm{mle}})\big] - \tau\big[\mathrm{E}(\hat{\beta}^{\mathrm{mle}})\big]}_{\text{transformation-induced}} + \overbrace{\tau\big[\mathrm{E}(\hat{\beta}^{\mathrm{mle}})\big] - \tau(\beta)}^{\text{coefficient-induced}}. \tag{1}$$

Transformation-induced τ-bias behaves systematically. The shape of the transformation $\tau(\cdot)$ determines the direction of the bias. In general, any strictly convex (concave) $\tau(\cdot)$ creates upward (downward) transformation-induced τ-bias. The direction and magnitude of the coefficient-induced τ-bias depend on the choice of $\tau(\cdot)$ and the bias in the coefficient estimates, but an unbiased estimator $\hat{\beta}^{\mathrm{mle}}$ implies the absence of coefficient-induced τ-bias.

But Rainey (2017) does not consider the average-of-simulations estimator. This raises the question: Does the average-of-simulations estimator $\hat{\tau}^{\mathrm{avg}}$ suffer the same transformation-induced bias as the plug-in estimator $\hat{\tau}^{\mathrm{mle}}$? I now turn to the average-of-simulations estimator and develop the idea of "simulation-induced bias."

---

[3]In the discussion that follows, I assume no Monte Carlo error exists in $\hat{\tau}^{\mathrm{avg}}$. In other words, I assume that M is sufficiently large so that $\hat{\tau}^{\mathrm{avg}} = \mathrm{E}[\tau(\tilde{\beta})]$, where $\tilde{\beta} \sim \mathrm{MVN}[\hat{\beta}^{\mathrm{mle}}, \hat{V}(\hat{\beta}^{\mathrm{mle}})]$.

[4]King, Tomz, and Wittenberg's claim that the draws are from the (frequentist) sampling distribution is incorrect. Take the sample mean as an example. If researchers could manipulate a sample to simulate draws from the sampling distribution, then they could compute the population mean with arbitrary precision, because the mean of the sampling distribution is the population mean. Of course, the precision of the sample mean is limited by the sample size—researchers cannot simulate from the sampling distribution.

If the transformation of coefficient estimates into an estimate of the quantity of interest is always convex (or always concave), then Jensen's inequality allows the simple statement relating $\hat{\tau}^{\text{avg}}$ and $\hat{\tau}^{\text{mle}}$ given in Lemma 1.

Lemma 1: Suppose a nondegenerate ML estimator $\hat{\beta}^{\text{mle}}$. Then any strictly convex (concave) $\tau(\cdot)$ guarantees that $\hat{\tau}^{\text{avg}}$ is strictly greater (less) than $\hat{\tau}^{\text{mle}}$.

This result is intuitive. Since I simulate using a multivariate normal distribution, $\tilde{\beta}$ has a symmetric distribution. But the distribution of $\tau(\tilde{\beta})$ is *not* symmetric. If $\tilde{\beta}$ happens to fall below the mode $\hat{\beta}^{\text{mle}}$, then $\tau(\cdot)$ pulls $\tau(\tilde{\beta})$ in toward $\hat{\tau}^{\text{mle}}$. If $\tilde{\beta}$ happens to fall above the mode $\hat{\beta}^{\text{mle}}$, then $\tau(\cdot)$ pushes $\tau(\tilde{\beta})$ away from $\hat{\tau}^{\text{mle}}$. This creates a right-skewed distribution for $\tau(\tilde{\beta})$, which pushes the average $\hat{\tau}^{\text{avg}}$ above $\hat{\tau}^{\text{mle}}$. See Appendix A for the proof.

For a convex transformation, Lemma 1 shows that $\hat{\tau}^{\text{avg}}$ is always larger than $\hat{\tau}^{\text{mle}}$. I refer to the expectation of this difference between $\hat{\tau}^{\text{avg}}$ and $\hat{\tau}^{\text{mle}}$ as "simulation-induced bias," so that

$$\text{simulation-induced } \tau\text{-bias} = \text{E}(\hat{\tau}^{\text{avg}}) - \text{E}(\hat{\tau}^{\text{mle}}).$$

Theorem 1: compares the sum of simulation- and transformation-induced τ-bias in $\hat{\tau}^{\text{avg}}$ to transformation-induced τ-bias in $\hat{\tau}^{\text{avg}}$.

Theorem 1 Suppose a nondegenerate ML estimator $\hat{\beta}^{\text{mle}}$. Then for any strictly convex or concave $\tau(\cdot)$, the sum of the simulation-induced and transformation-induced τ-bias in $\hat{\tau}^{\text{avg}}$ is strictly greater in magnitude than the transformation-induced τ-bias in $\hat{\tau}^{\text{mle}}$.

Regardless of the direction of simulation-induced and transformation-induced τ-bias, Theorem 1 shows that the magnitude of the combination in $\hat{\tau}^{\text{avg}}$ is *always* larger than the transformation-induced bias alone in $\hat{\tau}^{\text{mle}}$ for strictly convex or concave $\tau(\cdot)$. The proof follows directly from Jensen's inequality, but see Appendix A for the details.

Theorem 1 shows that $\hat{\tau}^{\text{avg}}$ compounds transformation-induced τ-bias with simulation-induced τ-bias. But is this bias substantively important? An analytical approximation provides a helpful guideline.

I approximate the simulation-induced τ-bias in $\hat{\tau}^{\text{avg}}$ as

$$\text{simulation-induced } \tau\text{-bias in } \hat{\tau}^{\text{avg}} = \underbrace{\left(\text{E}(\hat{\tau}^{\text{avg}}) - \tau\big[\text{E}(\hat{\beta}^{\text{mle}})\big]\right)}_{\text{t.i. } \tau\text{-bias in } \hat{\tau}^{\text{avg}}} - \underbrace{\left(\text{E}(\hat{\tau}^{\text{mle}}) - \tau\big[\text{E}(\hat{\beta}^{\text{mle}})\big]\right)}_{\text{t.i. } \tau\text{-bias in } \hat{\tau}^{\text{mle}}}$$

$$= \text{E}(\hat{\tau}^{\text{avg}}) - \text{E}(\hat{\tau}^{\text{mle}})$$

$$= \text{E}(\hat{\tau}^{\text{avg}} - \hat{\tau}^{\text{mle}})$$

$$= \text{E}\big(\text{E}\big[\tau(\tilde{\beta})\big] - \tau(\hat{\beta}^{\text{mle}})\big)$$

$$= \text{E}\left(\underbrace{\text{E}\big[\tau(\tilde{\beta})\big] - \tau\big[\text{E}(\tilde{\beta})\big]}_{\substack{\text{approximated in Equation (1),} \\ \text{p. 405, of Rainey (2017)}}}\right) \qquad (2)$$

$$\approx \text{E}\left[\frac{1}{2}\sum_{r=1}^{k+1}\sum_{s=1}^{k+1} H_{rs}(\hat{\beta}^{\text{mle}})\hat{V}_{rs}(\hat{\beta}^{\text{mle}})\right],$$

where the remaining expectation occurs with respect to $\hat{\beta}^{\mathrm{mle}}$, $H(\hat{\beta}^{\mathrm{mle}})$ represents the Hessian matrix of second derivatives of $\tau(\cdot)$ at the point $\hat{\beta}^{\mathrm{mle}}$ and, conveniently, $\hat{V}(\hat{\beta}^{\mathrm{mle}})$ represents the estimated covariance matrix for $\hat{\beta}^{\mathrm{mle}}$.

This approximation appears similar to the approximation for the transformation-induced τ-bias, which (adjusting notation slightly) Rainey (2017, 405, Equation (1)) computes as

$$\text{t.i. } \tau\text{-bias} \approx \frac{1}{2} \sum_{r=1}^{k+1} \sum_{s=1}^{k+1} H_{rs}\big[\mathrm{E}(\hat{\beta}^{\mathrm{mle}})\big] V_{rs}(\hat{\beta}^{\mathrm{mle}}), \tag{3}$$

where $H[\mathrm{E}(\hat{\beta}^{\mathrm{mle}})]$ represents the Hessian matrix of second derivatives of $\tau(\cdot)$ at the point $\mathrm{E}(\hat{\beta}^{\mathrm{mle}})$ and $V(\hat{\beta}^{\mathrm{mle}})$ represents the covariance matrix of the sampling distribution of $\hat{\beta}^{\mathrm{mle}}$.

When one compares Equations (2) and (3), they yet again compare the *expectation of a function* to the *function of the expectation*. Therefore, Equations (2) and (3) are not exactly equal. But, as a rough guideline, one should expect them to be similar. And to the extent that the two are similar, the additional simulation-induced τ-bias in $\hat{\tau}^{\mathrm{avg}}$ is about the same as the transformation-induced τ-bias in $\hat{\tau}^{\mathrm{mle}}$.

Because of the similarity between Equations (2) and (3), the simulation-induced τ-bias becomes large under the conditions identified by Rainey (2017) as leading to large transformation-induced τ-bias: when the non-linearity in the transformation $\tau(\cdot)$ is severe and when the standard errors of $\hat{\beta}^{\mathrm{mle}}$ are large. While the transformation-induced τ-bias vanishes as the number of observations grows large, it can be substantively meaningful for the sample sizes commonly encountered in social science research. In standard modeling situations, Rainey (2017) demonstrates that transformation-induced bias can (1) be larger than the bias in the estimates of the model coefficients and (2) shrink to zero more slowly as the sample size increases. By extension, the same claims hold for simulation-induced bias (which, again, appears *in addition to* transformation-induced bias).

4. The intuition of simulation-induced bias

To develop the intuition for the theoretical results above, I examine a stylized example with simulations, an alternative analytical approach, and an empirical example.

4.1. Using a drastic, convex transformation: $\tau(\mu) = \mu^2$

To develop an intuition for the simulation-induced τ-bias in $\hat{\tau}^{\mathrm{avg}}$, consider the simple (unrealistic, but heuristically useful) scenario in which $y_i \sim \mathrm{N}(0, 1)$, for $i \in \{1, 2, \ldots, n = 100\}$, and the researcher wishes to estimate μ^2. Suppose that the researcher knows that the variance equals one but does not know that the mean μ equals zero. The researcher uses the unbiased ML estimator $\hat{\mu}^{\mathrm{mle}} = \sum_{i=1}^{n} y_i/n$ of μ, but ultimately cares about the quantity of interest $\tau(\mu) = \mu^2$. The researcher can use the plug-in estimator $\hat{\tau}^{\mathrm{mle}} = (\hat{\mu}^{\mathrm{mle}})^2$ of $\tau(\mu)$. Alternatively, the researcher can use the average-of-simulations estimator, estimating $\tau(\mu)$ as $\hat{\tau}^{\mathrm{avg}} = 1/M \sum_{i=1}^{M} \tau(\tilde{\mu}^{(i)})$, where $\tilde{\mu}^{(i)} \sim \mathrm{N}(\hat{\mu}^{\mathrm{mle}}, 1/\sqrt{n})$ for $i \in \{1, 2, \ldots, M\}$.

The true value of the quantity of interest is $\tau(0) = 0^2 = 0$. However, the ML estimator $\hat{\tau}^{\mathrm{mle}} = (\hat{\mu}^{\mathrm{mle}})^2$ equals zero if and only if $\hat{\mu}^{\mathrm{mle}} = 0$. Otherwise, $\hat{\tau}^{\mathrm{mle}} > 0$. Since $\hat{\mu}^{\mathrm{mle}}$ almost surely differs from zero, $\hat{\tau}^{\mathrm{mle}}$ is biased upward.

Moreover, even if $\hat{\mu}^{\mathrm{mle}} = 0$, $\tilde{\mu}^{(i)}$ almost surely differs from zero. If $\tilde{\mu}^{(i)} \neq 0$, then $(\tilde{\mu}^{(i)})^2 > 0$. Thus, $\hat{\mu}^{\mathrm{avg}}$ is almost surely larger than the true value $\tau(\mu) = 0$ even when $\hat{\mu} = 0$.

I illustrate this fact clearly by repeatedly simulating y and computing $\hat{\tau}^{\mathrm{mle}}$ and $\hat{\tau}^{\mathrm{avg}}$. Figure 1 shows the first four of 10,000 total simulations. The figure shows how the unbiased estimate $\hat{\mu}^{\mathrm{mle}}$ is translated into $\hat{\tau}^{\mathrm{mle}}$ and $\hat{\tau}^{\mathrm{avg}}$.
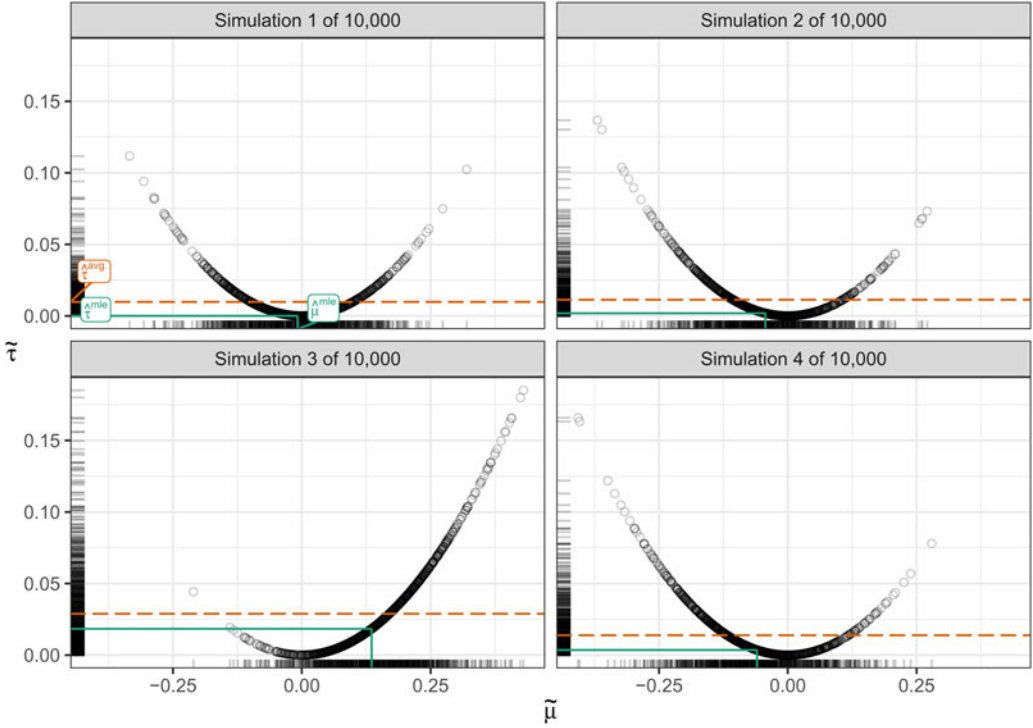
Figure 1. The first four Monte Carlo simulations of $\hat{\mu}$. These four panels illustrate the relationship between $\hat{\tau}^{\text{mle}}$ and $\hat{\tau}^{\text{avg}}$ described by Lemma 1 and Theorem 1.

First, to find $\hat{\tau}^{\text{avg}}$, I complete three steps: (1) simulate $\tilde{\mu}^{(i)} \sim \text{N}(\hat{\mu}^{\text{mle}}, 1/10)$ for $i \in \{1, 2, \ldots, M = 1{,}000\}$, (2) calculate $\tilde{\tau}^{(i)} = \tau(\tilde{\mu}^{(i)})$, and (3) calculate $\hat{\tau}^{\text{avg}} = 1/M \sum_{i=1}^{M} \tilde{\tau}^{(i)}$. The rug plot along the horizontal axis shows the distribution of $\tilde{\mu}$. The hollow points in Figure 1 shows the transformation of each point $\tilde{\mu}^{(i)}$ into $\tilde{\tau}^{(i)}$. The rug plot along the vertical axis shows the distribution of $\tilde{\tau}$. Focus on the top-left panel of Figure 1. Notice that $\hat{\mu}^{\text{mle}}$ estimates the true value $\mu = 0$ quite well. However, after simulating $\tilde{\mu}$ and transforming $\tilde{\mu}$ into $\tilde{\tau}$, the $\tilde{\tau}$s fall far from the true value $\tau(0) = 0$. The dashed orange line shows the average of $\tilde{\tau}$. Notice that although $\hat{\mu}^{\text{mle}}$ is unusually close to the truth $\mu = 0$ in this sample, $\hat{\tau}^{\text{avg}}$ is much too large.

Second, to find $\hat{\tau}^{\text{mle}}$, I transform $\hat{\mu}^{\text{mle}}$ directly using $\hat{\tau}^{\text{mle}} = (\hat{\mu}^{\text{mle}})^2$. The solid green lines show this transformation. The convex transformation $\tau(\cdot)$ has the effect of lengthening the right tail of the distribution of $\tilde{\tau}$, pulling the average well above the mode. This provides the basic intuition for Lemma 1.

The remaining panels of Figure 1 repeat this process with three more random samples. Each sample presents a similar story—the convex transformation stretches the distribution of $\tilde{\tau}$ to the right, which pulls $\hat{\tau}^{\text{avg}}$ above $\hat{\tau}^{\text{mle}}$.

I repeat this process 10,000 total times to produce 10,000 estimates of $\hat{\mu}^{\text{mle}}$, $\hat{\tau}^{\text{mle}}$, and $\hat{\tau}^{\text{avg}}$. Figure 2 shows the density plots for the 10,000 estimates (i.e., the sampling distributions of $\hat{\mu}^{\text{mle}}$, $\hat{\tau}^{\text{mle}}$, and $\hat{\tau}^{\text{avg}}$). As I show analytically, $\hat{\mu}^{\text{mle}}$ is unbiased with a standard error of $\sigma/\sqrt{n} = 1/\sqrt{100} = 1/10$. Both $\hat{\tau}^{\text{mle}}$ and $\hat{\tau}^{\text{avg}}$ are biased upward, but, as Theorem 1 suggests, $\hat{\tau}^{\text{avg}}$ has more bias than $\hat{\tau}^{\text{mle}}$. And as the approximation suggests, $\hat{\tau}^{\text{avg}}$ has about twice the bias of $\hat{\tau}^{\text{mle}}$. Indeed, the exact bias of each estimator is easy to compute for this simple example. Appendix B shows that the biases are $1/n = 1/100$ and $2/n = 2/100$ in this example.

Figure 2. The sampling distributions of $\hat{\mu}^{\text{mle}}$, $\hat{\tau}^{\text{mle}}$, and $\hat{\tau}^{\text{avg}}$.

4.2. Using the law of iterated expectations

One can also develop the argument analytically via the law of iterated expectations. It helps to alter the notation slightly, making two implicit dependencies explicit. I explain each change below and use the alternate, more expansive notation only in this section.

The law of iterated expectations states that $E_Y(E_{X|Y}(X \mid Y)) = E_X(X)$, where X and Y represent random variables. The three expectations occur with respect to three different distributions: E_Y denotes the expectation with respect to the marginal distribution of Y, $E_{X|Y}$ denotes the expectation with respect to the conditional distribution of $X|$ Y, and E_X denotes the expectation with respect to the marginal distribution of X.

Outside of this section, the reader should understand that the distribution of $\tilde{\beta}$ depends on $\hat{\beta}^{\text{mle}}$ and could be written as $\tilde{\beta} \mid \hat{\beta}^{\text{mle}}$. To remain consistent with previous work, especially Herron (1999) and King *et al.* (2000), I use $\tilde{\beta}$ to represent $\tilde{\beta} \mid \hat{\beta}^{\text{mle}}$. The definition of $\tilde{\beta}$ makes this usage clear. In this section only, I use $\tilde{\beta} \mid \hat{\beta}^{\text{mle}}$ to represent the conditional distribution of $\tilde{\beta}$ and use $\underline{\tilde{\beta}}$ to represent the <u>unconditional</u> distribution of $\tilde{\beta}$. Intuitively, one might imagine (1) generating a data set y, (2) estimating $\hat{\beta}^{\text{mle}}$, and (3) simulating $\tilde{\beta} \mid \hat{\beta}^{\text{mle}}$. If I perform steps (1) and (2) just once, but step (3) repeatedly, I generate a sample from the conditional distribution $\tilde{\beta} \mid \hat{\beta}^{\text{mle}}$. If I perform steps (1), (2), and (3) repeatedly, then I generate a sample from the <u>unconditional</u> distribution $\underline{\tilde{\beta}}$. The unconditional distribution helps us understand the nature of the simulation-induced τ-bias.

Applying the law of iterated expectations, I obtain $E_{\underline{\tilde{\beta}}}(\underline{\tilde{\beta}}) = E_{\hat{\beta}^{\text{mle}}}(E_{\tilde{\beta}|\hat{\beta}^{\text{mle}}}(\tilde{\beta} \mid \hat{\beta}^{\text{mle}}))$. The three identities below connect the three key quantities from Theorem 1 to three versions of $E_{\hat{\beta}^{\text{mle}}}(E_{\tilde{\beta}|\hat{\beta}^{\text{mle}}}(\tilde{\beta} \mid \hat{\beta}^{\text{mle}}))$, with the transformation $\tau(\cdot)$ applied at different points.

$$\tau\left[\underset{\hat{\beta}^{\text{mle}}}{E}\left(\underset{\tilde{\beta}|\hat{\beta}^{\text{mle}}}{E}\left(\tilde{\beta} \mid \hat{\beta}^{\text{mle}}\right)\right)\right] = \tau\left[\underset{\underline{\tilde{\beta}}}{E}\left(\underline{\tilde{\beta}}\right)\right] = \tau\left[\underset{\hat{\beta}^{\text{mle}}}{E}\left(\hat{\beta}^{\text{mle}}\right)\right], \tag{4}$$

$$\underset{\hat{\beta}^{\text{mle}}}{E}\left(\tau\left[\underset{\tilde{\beta}|\hat{\beta}^{\text{mle}}}{E}\left(\tilde{\beta} \mid \hat{\beta}^{\text{mle}}\right)\right]\right) = \underset{\hat{\beta}^{\text{mle}}}{E}\left(\tau\left[\hat{\beta}^{\text{mle}}\right]\right) = \underset{\hat{\beta}^{\text{mle}}}{E}\left(\hat{\tau}^{\text{mle}}\right), \text{ and } \quad \longleftarrow \quad \text{Switch } \tau \text{ and an E once.}$$

$$\tag{5}$$

$$\underset{\hat{\beta}^{\text{mle}}}{E}\left(\underset{\tilde{\beta}|\hat{\beta}^{\text{mle}}}{E}\left(\tau\left[\tilde{\beta} \mid \hat{\beta}^{\text{mle}}\right]\right)\right) = \underset{\underline{\tilde{\beta}}}{E}\left(\tau\left[\underline{\tilde{\beta}}\right]\right) = \underset{\underline{\tilde{\beta}}}{E}\left(\hat{\tau}^{\text{avg}}\right). \quad \longleftarrow \quad \text{Switch } \tau \text{ and an E again.}$$

$$\tag{6}$$

If I subtract Equation (5) from Equation (4), I obtain the transformation-induced τ-bias in $\hat{\tau}^{\text{mle}}$ (see Equation (1) for the definition of transformation-induced τ-bias). To move from Equation

(4) to (5) I must swap $\tau(\cdot)$ with an expectation once. This implies that, if $\tau(\cdot)$ is convex, Equation (5) must be greater than Equation (4). This, in turn, implies that the bias is positive.

To obtain the τ-bias in $\hat{\tau}^{avg}$ I must subtract Equation (6) from (4). But to move from Equation (4) to (6) I must swap $\tau(\cdot)$ with an expectation *twice*. Again, if $\tau(\cdot)$ is convex, then Equation (6) must be greater than Equation (4). However, because one should expect $\hat{\beta}^{mle}$ and $\tilde{\beta} \mid \hat{\beta}^{mle}$ to have similar distributions, one should expect the additional swap to roughly double the bias in $\hat{\tau}^{avg}$ compared to $\hat{\tau}^{mle}$. This additional swap creates the additional, simulation-induced τ-bias.

4.3. Using a re-analysis of Holland (2015)

Holland (2015) presents a nuanced theory that describes the conditions under which politicians choose to enforce laws and supports the theoretical argument with a rich variety of evidence. In particular, she elaborates on the *electoral* incentives of politicians to enforce laws. I borrow three Poisson regressions and hypotheses about a single explanatory variable to illustrate how the plug-in estimates can differ from the average-of-simulations estimate.

Holland writes:

> My first hypothesis is that enforcement operations drop off with the fraction of poor residents in an electoral district. So district poverty should be a negative and significant predictor of enforcement, but only in politically decentralized cities [Lima and Santiago]. Poverty should have no relationship with enforcement in politically centralized cities [Bogota] once one controls for the number of vendors.

I use Holland's hypothesis and data to illustrate the behavior of the average-of-simulations and plug-in estimators. I refit Model 1 from Table 2 in Holland (2015) for each city. I then use each model to compute the percent increase in the enforcement operations for each district in the city if the percent of the district in the lower class dropped by half. For example, in the Villa Maria El Triunfo district in Lima, 84 percent of the district is in the lower class. If this dropped to 42 percent, then the average-of-simulations estimate suggests that the number of enforcement operations would increase by about 284 percent (from about 5 to 20). The plug-in estimate, on the other hand, suggests an increase of 264 percent (from about 5 to 17). The plug-in estimate, then, is about 7 percent smaller than the average-of-simulations estimate—a small, but noticeable shrinkage.

Figure 3 shows how the estimates change (usually shrink) for all districts when I switch from the average-of-simulations estimates to plug-in estimates. Table 1 presents the details for the

Figure 3. This figure compares the average-of-simulations estimates with the plug-in estimates using three Poisson regression models from Holland (2015). The quantity of interest is the percent increase in the enforcement operations when the percent of a district in the lower class drops by half. The arrows show how the estimates change when I switch from the average-of-simulations to the plug-in estimate.

Table 1. This table presents the details for the districts labeled in Figure 3

City	District	Average of simulations			ML estimate			
		% Change[a]	from[b]	to[c]	% Change	from	to	shrinkage[d]
	Usme	9%	5.5	5.8	7%	5.4	5.8	16%
	Cuidad Bolivar	8%	6.0	6.4	7%	5.9	6.3	15%
	Bosa	8%	7.1	7.4	7%	6.9	7.4	15%
	San Cristobal	7%	12.6	13.4	6%	12.5	13.3	14%
Bogota	Santa Fe	6%	27.3	29.0	5%	26.6	28.0	11%
	Villa Maria El Triunfo	284%	5.3	19.5	264%	4.7	17.1	7%
	Villa El Salvador	231%	7.3	23.5	217%	6.8	21.4	6%
	Ventanilla	192%	8.4	23.4	182%	8.2	23.0	5%
	Lurin	176%	6.9	17.4	168%	6.4	17.1	5%
Lima	Chacalacayo	167%	6.7	16.4	159%	6.2	16.1	5%
	La Pintana	642%	1.4	4.0	301%	0.8	3.4	53%
	Cerro Navia	545%	1.5	4.2	272%	1.0	3.6	50%
	Lo Espejo	517%	1.4	4.3	263%	1.0	3.5	49%
	Renca	451%	1.3	4.1	241%	1.0	3.4	47%
Santiago	San Ramon	451%	1.2	4.0	241%	1.0	3.3	47%

[a]Quantity of interest; percent change in enforcement operations when the percent in the lower class drops by half.
[b]Enforcement operations when the percent in the lower class equals its observed value.
[c] Enforcement operations when the percent in the lower class equals half its observed value.
[d] Shrinkage in the quantity of interest due to switching from the average of simulations to the ML estimator.

labeled cases in Figure 3. In Bogota, the estimate shrinks by 11 percent in Sante Fe and 16 percent in Usme. In Lima, the estimate shrinks by 5 percent in Chacalacayo and 7 percent Villa Maria El Triunfo. The shrinkage is much larger in Santiago, where the standard errors for the coefficient estimates are much larger. The estimate shrinks by about 47 percent in San Ramon and 53 percent in La Pintana. The median shrinkage is 6 percent in Bogota, 2 percent in Lima, and 37 percent in Santiago. For many districts in Santiago, the average-of-simulations estimate is about *twice* the plugin estimate. These estimates clearly show that the average of simulations and the ML estimate can differ meaningfully in actual analyses.

5. Conclusion

Many social scientists turn to King *et al.* (2000) for advice on interpreting, summarizing, and presenting empirical results. The authors improved empirical research by highlighting the importance of substantively meaningful quantities of interest. I agree with King and Zeng's (2006) summary of the literature: "[w]hether such effects are calculated via analytical derivation or what is now the more common approach of statistical simulation, political scientists have made much progress in learning how to make sophisticated methods speak directly to their substantive research questions" (132).

Researchers estimate quantities of interest either by averaging simulated quantities of interest (e.g., CLARIFY in Stata, Zelig in R, or clarify in R) or using the invariance property of ML estimators (e.g., margins in Stata and R). In practice, researchers' choice between these two estimators seems idiosyncratic rather than principled, depending on their preferred software package rather than any statistical criteria. The methodological literature recommends both, but has not distinguished or compared the two approaches to estimating quantities of interest.

When researchers use the average of simulations (King *et al.*, 2000) to estimate quantities of interest, they replicate the logic of transformation-induced bias (Rainey, 2017) and add simulation-induced bias to the estimates. This additional bias is roughly the same magnitude and direction as transformation-induced bias and occurs *in addition to* transformation-induced bias. While the additional bias is usually small relative to the standard error, methodologists should not recommend methods that add unnecessary bias to point estimates. Instead, we should

recommend methods that better adhere to the usual evaluative standards. Even if we recommend using simulation to estimate standard errors, we should not recommend averaging simulations to obtain the point estimate. Instead, researchers should directly transform ML estimates of coefficients to obtain ML estimates of the quantities of interest. The resulting point estimates inherit the desirable properties of ML estimators and avoid unnecessary simulation-induced bias.

References

Casella G and Berger RL (2002) *Statistical Inference*, 2nd Edn. Pacific Grove, CA: Duxbury.

Choirat C, Honaker J, Imai K, King G and Lau O (2018) *Zelig: Everyone's Statistical Software*. Version 5.0-12, http://zeligproject.org/

Gelman A and Hill J (2007) *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge: Cambridge University Press.

Greifer N, Worthington S, Iacus S and King G (2023) *clarify: Simulation-Based Inference for Regression Models*. https://iqss.github.io/clarify/

Hanmer MJ and Kalkan KO (2013) Behind the curve: clarifying the best approach to calculating predicted probabilities and marginal effects from limited dependent variable models. *American Journal of Political Science* 57, 263–277.

Herron MC (1999) Postestimation uncertainty in limited dependent variable models. *Political Analysis* 8, 83–98.

Holland AC (2015) The distributive politics of enforcement. *American Journal of Political Science* 59, 357–371.

Imai K, King G and Lau O (2008) Toward a common framework for statistical analysis and development. *Journal of Computational and Graphical Statistics* 17, 892–913.

Jung YS, Souza FDS, Philips AQ, Rutherford A and Whitten GD (2020) A command to estimate and interpret models of dynamic compositional dependent variables: new features for dynsimpie. *The Stata Journal* 20, 584–603.

King G (1998) *Unifying Political Methodology: The Likelihood Theory of Statistical Inference*. Ann Arbor: Michigan University Press.

King G and Zeng L (2006) The dangers of extreme counterfactuals. *Political Analysis* 14, 131–159.

King G, Tomz M and Wittenberg J (2000) Making the most of statistical analyses: improving interpretation and presentation. *American Journal of Political Science* 44, 341–355.

Leeper TJ (2021) Margins: marginal effects for model objects. R package version 0.3.26.

Philips AQ, Rutherford A and Whitten GD (2016) dynsimpie: A command to examine dynamic compositional dependent variables. *The Stata Journal* 16, 662–677.

Rainey C (2017) Transformation-induced bias: unbiased coefficients do not imply unbiased quantities of interest. *Political Analysis* 25, 402–409.

R Core Team (2018) *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

StataCorp (2017) *Stata 15 Base Reference Manual*. College Station, TX: Stata Press.

Tomz M, Wittenberg J and King G (2003) Clarify: software for interpreting and presenting statistical results. *Journal of Statistical Software* 8, 1–30.

Van der Vaart AW (2000) *Asymptotic Statistics*. Cambridge: Cambridge University Press.

Venables WN and Ripley BD (2002) *Modern Applied Statistics with S*, 4th Edn. New York: Springer.