The Annals of Applied Statistics
2014, Vol. 8, No. 3, 1395–1415
DOI: 10.1214/14-AOAS762

INFERENCE FOR DEFORMATION AND INTERFERENCE IN 3D PRINTING

BY ARMAN SABBAGHI[1,*], TIRTHANKAR DASGUPTA[2,†], QIANG HUANG[3,‡] AND JIZHE ZHANG[‡]

Purdue University, Harvard University*[†] *and University of Southern California*[‡]

Additive manufacturing, or 3D printing, is a promising manufacturing technique marred by product deformation due to material solidification in the printing process. Control of printed product deformation can be achieved by a compensation plan. However, little attention has been paid to interference in compensation, which is thought to result from the inevitable discretization of a compensation plan. We investigate interference with an experiment involving the application of discretized compensation plans to cylinders. Our treatment illustrates a principled framework for detecting and modeling interference, and ultimately provides a new step toward better understanding quality control for 3D printing.

1. Interference in compensation. Additive manufacturing, or 3D printing, refers to a class of technology for the direct fabrication of physical products from 3D Computer-Aided Design (CAD) models. In contrast to material removal processes in traditional machining, the printing process adds material layer by layer. This enables direct printing of geometrically complex products without affecting building efficiency. No extra effort is necessary for molding construction or fixture tooling design, making 3D printing a promising manufacturing technique [Campbell et al. (2011), Gibson, Rosen and Stucker (2009), Hilton and Jacobs (2000), Melchels, Feijen and Grijpma (2010)]. Despite these promising features, accurate control of a product's printed dimensions remains a major bottleneck. Material solidification during layer formation leads to product deformation, or shrinkage [Wang et al. (1996)], which reduces the utility of printed products. Shrinkage control is crucial to overcome the accuracy barrier in 3D printing.

To control detailed features along the boundary of a printed product, Tong, Lehtihet and Joshi (2003) and Tong, Joshi and Lehtihet (2008) used polynomial regression models to first analyze shrinkage in different directions separately, and then compensate for product deformation by changing the original CAD accordingly. Unfortunately, their predictions are independent of the product's geometry, which is not consistent with the physical manufacturing process. Huang et al.

Received November 2013; revised April 2014.

[1]Supported by National Science Foundation Grant DGE-1144152.

[2]Supported by National Science Foundation Grant DMS-11-07004.

[3]Supported by Office of Naval Research Grant #N000141110671.

Key words and phrases. Additive manufacturing, posterior predictive checks, quality control, Rubin Causal Model, stable unit-treatment value assumption.

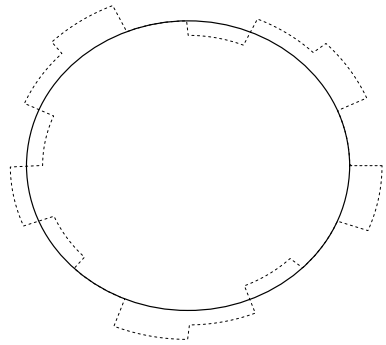

FIG. 1. *A discretized compensation plan* (*dashed line*) *to the nominal boundary* (*solid line*). *Note that compensation could be negative*.

(2014) built on this work, establishing a generic, physically consistent approach to model and predict product deformations, and to derive compensation plans. The essence of this new modeling approach is to transform in-plane geometric errors from the Cartesian coordinate system into a functional profile defined on the polar coordinate system. This representation decouples the geometric shape complexity from the deformation modeling, and a generic formulation of shape deformation can thus be achieved. The approach was developed for a stereolithography process, and in experiments achieved an improvement of one order of magnitude in reduction of deformation for cylinder products.

However, an important issue not yet addressed in the previously cited work on deformation control for 3D printing is how the application of compensation to one section of a product will affect the deformation of its neighbors. Compensation plans are always discretized according to the tolerance of the 3D printer, in the sense that sections of the CAD are altered by single amounts, for example, as in Figure 1. Furthermore, when planning an experiment to assess the effect of compensation on product deformation, it is natural to discretize the quantitative "compensation" factor into a finite number of levels, which also leads to a product having a more complex boundary. Ultimately, such changes may introduce interference between different sections of the printed product, which is defined to occur when one section's deformation depends not only on its assigned compensation, but also on compensations assigned to its neighbors [Rubin (1980)]. For example, in Figure 1, the deformation for points near the boundary of two neighboring sections should depend on compensations applied to both. By the same logic, interference becomes a practical issue when printing products with complex geometry. Therefore, to improve quality control in 3D printing, it is important to formally investigate complications introduced by the interference that results from discretization in compensation plans. We take the first step with an experiment involving a discretized compensation plan for a simple shape.

We begin in Section 2 with a review of interference, models for product deformation, and the effect of compensation. Adoption of the Rubin Causal Model

[RCM, Holland (1986)] is a significant and novel feature of our investigation, and facilitates the study of interference. Section 3.1 summarizes the basic model and analysis for deformation of cylinders given by Huang et al. (2014). Our analyses are in Sections 3.2–3.5: we first describe an experiment hypothesized to generate interference, then proceed with posterior predictive checks to demonstrate the existence of interference, and finally conclude with a model that captures interference. A statistically substantial idea in Section 3.3 is that, in experiments with distinct units of analysis and units of interpretation [Cox and Donnelly (2011), pages 18–19], the posterior distribution of model parameters, based on "benchmark" data, yields a simple assessment and inference for interference in the experiment, similar to that suggested by Sobel (2006) and Rosenbaum (2007). Analyses in Sections 3.4–3.5 demonstrate how discretized compensation plans complicate quality control through the Introduction of interference. This illustrates the fact that in complex manufacturing processes, a proper definition of experimental units and understanding of interference are critical to quality control.

2. Potential outcomes and interference.

2.1. *Experimental units and potential outcomes.* We use the general framework for product deformation given by Huang et al. [(2014), pages 3–6]. Suppose a product has intended shape ψ_0 and observed shape ψ under a 3D printing process. Deformation is informally described as the difference between ψ and ψ_0, where we can represent both either in the Cartesian coordinate system (x, y, z) or cylindrical coordinate system (r, θ, z). Cylindrical coordinates facilitate deformation modeling and are used throughout.

For illustrative purposes, we define terms for two-dimensional products (notation for three dimensions follows immediately). Quality control requires an understanding of deformation in different regions of the product that receive different amounts of compensation. We therefore define a finite number N of points on the boundary of the product, corresponding to specific angles $\theta_1, \ldots, \theta_N$, as the experimental units. The desired boundary from the CAD model is defined by the function $r_0(\theta)$, denoting the nominal radius at angle θ. We consider only one (quantitative) treatment factor, compensation to the CAD, defined as a change in the nominal radius of the CAD by x_i units at θ_i for $i = 1, \ldots, N$. Compensation is not restricted to be nonnegative. The potential radius for θ_i under compensation $\mathbf{x} = (x_1, \ldots, x_N)$ to $\theta_1, \ldots, \theta_N$ is a function of θ_i, $r_0(\cdot)$, and $\mathbf{x}$, denoted by $r(\theta_i, r_0(\cdot), \mathbf{x})$. The difference between the potential and nominal radius at θ_i defines deformation, and so

$$\Delta r\big(\theta_i, r_0(\cdot), \mathbf{x}\big) = r\big(\theta_i, r_0(\cdot), \mathbf{x}\big) - r_0(\theta_i) \tag{1}$$

is defined as our potential outcome for θ_i. Potential outcomes are viewed as fixed numbers, with randomness introduced in Section 2.3 in our general model for the potential outcomes.

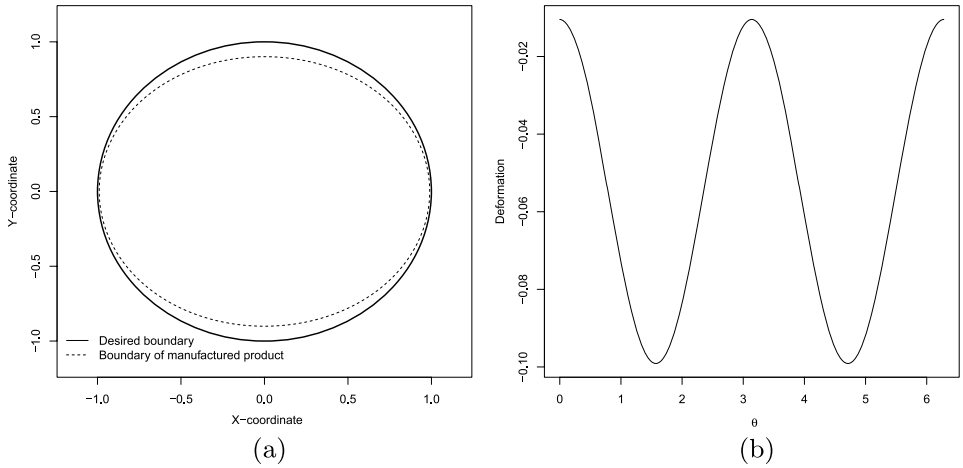

FIG. 2. (a) *Ideal shape* (*solid line*) *versus the actual shape* (*dashed line*). (b) *Visualization of shrinkage.*

This definition of the potential outcome is convenient for visualizing shrinkage. For example, suppose the desired shape of the product is the solid line, and the manufactured product when $\mathbf{x} = \mathbf{0} = (0, \ldots, 0)$ is the dashed line, in Figure 2(a). Plotting the deformation at each angle yields a visualization amenable to analysis [Figure 2(b)]. Orientation is fixed: we match the coordinate axes of the printed product with those of the CAD model.

2.2. *Interference.* A unit θ_i is said to be affected by interference if

$$\Delta r\big(\theta_i, r_0(\cdot), \mathbf{x}\big) \neq \Delta r\big(\theta_i, r_0(\cdot), \mathbf{x}'\big)$$

for at least one pair of distinct treatment vectors $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^N$ with $x_i = x_i'$ [Rubin (1980)]. If there is no interference, then $\Delta r(\theta_i, r_0(\cdot), \mathbf{x})$ is a function of $\mathbf{x}$ only via the component x_i. As the experimental units reside on a connected boundary, the deformation of one unit may depend on compensations assigned to its neighbors when the compensation plan is discretized. Perhaps less plausible, but equally serious, is the possible leakage of assigned compensations across units. These considerations explain the presence of the vector $\mathbf{x}$, containing compensations for all units, in the potential outcome notation (1). Practically, accommodations made for interference should reduce bias in compensation plans for complex products and improve quality control.

2.3. *General deformation model.* Following Huang et al. [(2014), pages 6–8], our potential outcome model under compensation plan $\mathbf{x} = \mathbf{0}$ is decomposed into three components:

$$\Delta r\big(\theta_i, r_0(\cdot), \mathbf{0}\big) = f_1\big(r_0(\cdot)\big) + f_2\big(\theta_i, r_0(\cdot), \mathbf{0}\big) + \varepsilon_i. \tag{2}$$

Function $f_1(r_0(\cdot))$ represents average deformation of a given nominal shape $r_0(\cdot)$ independent of location θ_i, and $f_2(\theta_i, r_0(\cdot), \mathbf{0})$ is the additional location-dependent deformation, geometrically and physically related to the CAD model. We can also interpret $f_1(\cdot)$ as a low-order component and $f_2(\cdot, \cdot, \mathbf{0})$ as a high-order component of deformation. The ε_i are random variables representing high-frequency components that add on to the main trend, with expectation $\mathbb{E}(\varepsilon_i) = 0$ and $\operatorname{Var}(\varepsilon_i) < \infty$ for all $i = 1, \ldots, N$.

Figure 2 demonstrates model (2). In this example, $r_0(\theta) = r_0$, so $f_1(\cdot)$ is a function of r_0, and $f_2(0, r_0, \mathbf{0}) = f_2(2\pi, r_0, \mathbf{0})$. Decomposition of deformation into lower and higher order terms yields

$$\Delta r(\theta_i, r_0, \mathbf{0}) = c_{r_0} + \sum_k \{a_{r_0,k} \cos(k\theta_i) + b_{r_0,k} \sin(k\theta_i)\} + \varepsilon_i, \tag{3}$$

where $f_1(r_0) = c_{r_0}$, and $\{a_{r_0,k}, b_{r_0,k}\}$ are coefficients of a Fourier series expansion of $f_2(\cdot, \cdot, \mathbf{0})$. The $\{a_{r_0,k}, b_{r_0,k}\}$ terms with large k represent the product's surface roughness, which is not of primary interest.

2.4. *General compensation and interference models.* Under the polar coordinate system, a compensation of x_i units at θ_i can be thought of as an extension of the product's radius by x_i units in that direction. Bearing this in mind, we first follow Huang et al. [(2014), page 8] to extend (2) to accommodate compensations, and then build upon this to give an extension that can help capture interference resulting from discretized compensation plans.

Let $r(\theta_i, r_0(\cdot), (x_i, \ldots, x_i)) = r(\theta_i, r_0(\cdot), x_i\mathbf{1})$ denote the potential radius for θ_i under compensation of x_i units to all points. Compensation $x_i\mathbf{1}$ is equivalent, in terms of the final manufactured product, as if a CAD model with nominal radius $r_0(\cdot) + x_i$ and compensation $\mathbf{0}$ was initially submitted to the 3D printer. Then

$$\begin{aligned} r(\theta_i, r_0(\cdot), x_i\mathbf{1}) - \{r_0(\theta_i) + x_i\} &= r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}) - \{r_0(\theta_i) + x_i\} \\ &= \Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}), \end{aligned} \tag{4}$$

where $\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0})$ follows the same form as (2), abbreviated as

$$\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}) = \mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0})\} + \varepsilon_i. \tag{5}$$

Consequently, the potential outcome for θ_i is

$$\begin{aligned} \Delta r(\theta_i, r_0(\cdot), x_i\mathbf{1}) &= r(\theta_i, r_0(\cdot), x_i\mathbf{1}) - r_0(\theta_i) \\ &= r(\theta_i, r_0(\cdot), x_i\mathbf{1}) - \{r_0(\theta_i) + x_i\} + x_i \\ &= \Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0}) + x_i \\ &= \mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0})\} + x_i + \varepsilon_i. \end{aligned} \tag{6}$$

The last two steps follow from (4) and (5), respectively. If x_i is small relative to $r_0(\theta_i)$, then (6) can be approximated using the first and second terms of the Taylor expansion of $\mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x_i, \mathbf{0})\}$ at $r_0(\theta_i)$:

$$\begin{aligned}\Delta r\big(\theta_i, r_0(\cdot), x_i\mathbf{1}\big) &\approx \mathbb{E}\big\{\Delta r\big(\theta_i, r_0(\cdot), \mathbf{0}\big)\big\} \\ &\quad + (x_i - 0)\bigg[\frac{d}{dx}\mathbb{E}\big\{\Delta r\big(\theta_i, r_0(\cdot) + x, \mathbf{0}\big)\big\}\bigg]_{x=0} + x_i + \varepsilon_i \\ &= \Delta r\big(\theta_i, r_0(\cdot), \mathbf{0}\big) + \big\{1 + h\big(\theta_i, r_0(\cdot), \mathbf{0}\big)\big\}x_i,\end{aligned} \tag{7}$$

where $h(\theta_i, r_0(\cdot), \mathbf{0}) = [d/dx\mathbb{E}\{\Delta r(\theta_i, r_0(\cdot) + x, \mathbf{0})\}]_{x=0}$. Under a specified parametric model for the potential outcomes, this Taylor expansion is performed conditional on the model parameters. When there is no interference,

$$\Delta r\big(\theta_i, r_0(\cdot), \mathbf{x}\big) = \Delta r\big(\theta_i, r_0(\cdot), x_i\mathbf{1}\big)$$

for any $\mathbf{x} \in \mathbb{R}^N$, and so (7) is a model for compensation effects in this case.

We can generalize this model to incorporate interference in a simple manner for a compensation plan $\mathbf{x}$ with different units assigned different compensations. As all units are connected on the boundary of the product, unit θ_i's treatment effect will change due to interference from its neighbors, so that θ_i will deform not just according to its assigned compensation x_i, but instead according to a compensation $g_i(\mathbf{x})$. Thus, we generalize (7) to

$$\Delta r\big(\theta_i, r_0(\cdot), \mathbf{x}\big) \approx \Delta r\big(\theta_i, r_0(\cdot), \mathbf{0}\big) + \big\{1 + h\big(\theta_i, r_0(\cdot), \mathbf{0}\big)\big\}g_i(\mathbf{x}), \tag{8}$$

where the *effective treatment* $g_i(\mathbf{x})$ is a function of x_i and assigned compensations for neighbors of θ_i (with the definition of neighboring units naturally dependent on the specific product), hence potentially a function of the entire vector $\mathbf{x}$. Allowing the treatment effect for θ_i to depend on treatments assigned to its neighboring units effectively incorporates interference in a meaningful manner, as will be seen in the analysis of our experiment.

3. Experimental design and analysis for interference.

3.1. *Compensation model for cylinders.* Huang et al. [(2014), page 12] constructed four cylinders with $r_0 = 0.5, 1, 2$, and 3 inches, and used $N_{0.5} = 749$, $N_1 = 707$, $N_2 = 700$, and $N_3 = 721$ equally-spaced units from each. Based on the logic in Section 2.3, they fitted

$$\Delta r(\theta_i, r_0, \mathbf{0}) = x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i) + \varepsilon_i \tag{9}$$

to the data, with $\varepsilon_i \sim \mathrm{N}(0, \sigma^2)$ independently, and parameters α, β, a, b, x_0, and σ independent of r_0. Specifically, for the cylinder, the location-independent term is thought to be proportional to r_0, so that with overexposure of x_0 units it would be of the form $x_0 + \alpha(r_0 + x_0)$. Furthermore, the location-dependent term is

TABLE 1
Summary of 1000 *posterior draws of parameters after a burn-in of* 500 *when no compensation is applied. This is drawn from Table* 5 *in Huang et al.* (2014). *Effective sample size is abbreviated as ESS throughout*

	Mean	**SD**	**Median**	**95% credible interval**	**ESS**
α	-1.34×10^{-2}	1.6×10^{-4}	-1.34×10^{-2}	$(-1.37, -1.31) \times 10^{-2}$	8198
β	5.7×10^{-3}	3.1×10^{-5}	5.71×10^{-3}	$(5.65, 5.8) \times 10^{-3}$	9522
a	8.61×10^{-1}	7.33×10^{-3}	8.61×10^{-1}	$(8.47, 8.75) \times 10^{-1}$	8223
b	1.13	5.46×10^{-3}	1.13	(1.12, 1.14)	9424
x_0	8.79×10^{-3}	1.5×10^{-4}	8.79×10^{-3}	$(8.5, 9.07) \times 10^{-3}$	8211
σ	8.7×10^{-4}	1.18×10^{-5}	8.7×10^{-4}	$(8.5, 8.9) \times 10^{-4}$	9384

thought to be a harmonic function of θ_i, and also proportional to r_0, of the form $\beta(r_0 + x_0)\cos(2\theta_i)$ with overexposure. Independent errors are used throughout because the focus is on a correct specification of the mean trend in deformation (Appendix A contains a discussion on this point). Huang et al. (2014) specified

$$a \sim \mathrm{N}(1, 2^2), \qquad b \sim \mathrm{N}(1, 1^2), \qquad \log(x_0) \sim \mathrm{N}(0, 1^2)$$

and placed flat priors on α, β, and $\log(\sigma)$, with all parameters independent a priori. Posterior draws of the parameters were obtained by Hamiltonian Monte Carlo [HMC, Duane et al. (1987)] and are summarized in Table 1, with convergence diagnostics discussed in Appendix B. A simple comparison of the posterior predictive distribution of product deformation to the observed data [Huang et al. (2014), page 19] demonstrates the good fit, and so we proceed with this specification and parameter inferences to design and analyze an experiment for interference.

Substituting $\Delta r(\theta_i, r_0, \mathbf{0})$ from (9) into the general model (6), we have

$$\begin{aligned}(10)\qquad &\Delta r(\theta_i, r_0, x_i\mathbf{1}) \\ &\quad = x_0 + x_i + \alpha(r_0 + x_0 + x_i)^a + \beta(r_0 + x_0 + x_i)^b \cos(2\theta_i) + \varepsilon_i.\end{aligned}$$

The Taylor expansion at $r_0 + x_0$, as in (7), yields the model

$$\begin{aligned}(11)\qquad &\Delta r(\theta_i, r_0, x_i\mathbf{1}) \\ &\quad = x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i) \\ &\qquad + \{1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1}\cos(2\theta_i)\}x_i + \varepsilon_i.\end{aligned}$$

We incorporate interference for a plan $\mathbf{x}$ with different units assigned different compensations by changing x_i in the right side of (11) to $g_i(\mathbf{x})$, with the functional form of $g_i(\mathbf{x})$ derived by exploratory means in Section 3.3.

3.2. *Experimental design for interference.* Under a discretized compensation plan, the boundary of a product is divided into sections, with all points in one section assigned the same compensation. In the terminology of Cox and Donnelly [(2011), pages 18–19], these sections constitute units of analysis, and individual angles are units of interpretation. We expect interference for angles near neighboring sections. Interference should be substantial for a large difference in neighboring compensations, and negligible otherwise.

This reasoning led to the following restricted Latin square design to study interference. We apply compensations to four cylinders of radius 0.5, 1, 2, and 3 inches, with each cylinder divided into 16 equal-sized sections of $\pi/8$ radians. One unit of compensation is 0.004, 0.008, 0.016, and 0.03 inch for each respective cylinder, and there are only four possible levels of compensation, $-1, 0, +1$, and $+2$ units. Two blocking factors are considered. The first is the quadrant and the second is the "symmetry group" consisting of $\pi/8$-radian sections that are reflections about the coordinate axes from each other. Symmetric sections form a meaningful block: if compensation x is applied to all units, then we have from (11) that for $0 \le \theta \le \pi/2$,

$$\begin{aligned}\mathbb{E}\{\Delta r(\theta, r_0, x\mathbf{1})|\alpha, \beta, a, b, x_0, \sigma\} &= \mathbb{E}\{\Delta r(\pi - \theta, r_0, x\mathbf{1})|\alpha, \beta, a, b, x_0, \sigma\} \\ &= \mathbb{E}\{\Delta r(\pi + \theta, r_0, x\mathbf{1})|\alpha, \beta, a, b, x_0, \sigma\} \\ &= \mathbb{E}\{\Delta r(2\pi - \theta, r_0, x\mathbf{1})|\alpha, \beta, a, b, x_0, \sigma\},\end{aligned}$$

suggesting a need to control for this symmetry in the experiment. Thus, for each product, we conceive of the 16 sections as a 4×4 table, with symmetry groups forming the column blocking factor and quadrants the row blocking factor. Based on prior concerns about the possible severity of interference and resulting scope of inference from our model (7), the set of possible designs was restricted to Latin squares (each compensation level occurs only once in any quadrant and symmetry group), where the absolute difference in assigned treatments between two neighboring sections does not exceed two levels of compensation. Each product was randomly assigned one design from this set, with no further restriction that all the products have the same design.

Our restricted Latin square design forms a discretized compensation plan that blocks on two factors suggested by the previous deformation model, and remains model-robust to a certain extent. The chosen experimental designs are in Figure 3, and observed deformations for the manufactured products are in Figure 4. There are $N_{0.5} = 6159$, $N_1 = 6022$, $N_2 = 6206$, and $N_3 = 6056$ equally spaced angles considered for the four cylinders.

3.3. *Assessing the structure of interference.* Our first task is to assess which units have negligible interference in the experiment. To do so, we use the suggestions of Sobel (2006) and Rosenbaum (2007), who describe when interest exists in comparing a treatment assignment $\mathbf{x}$ to a baseline.

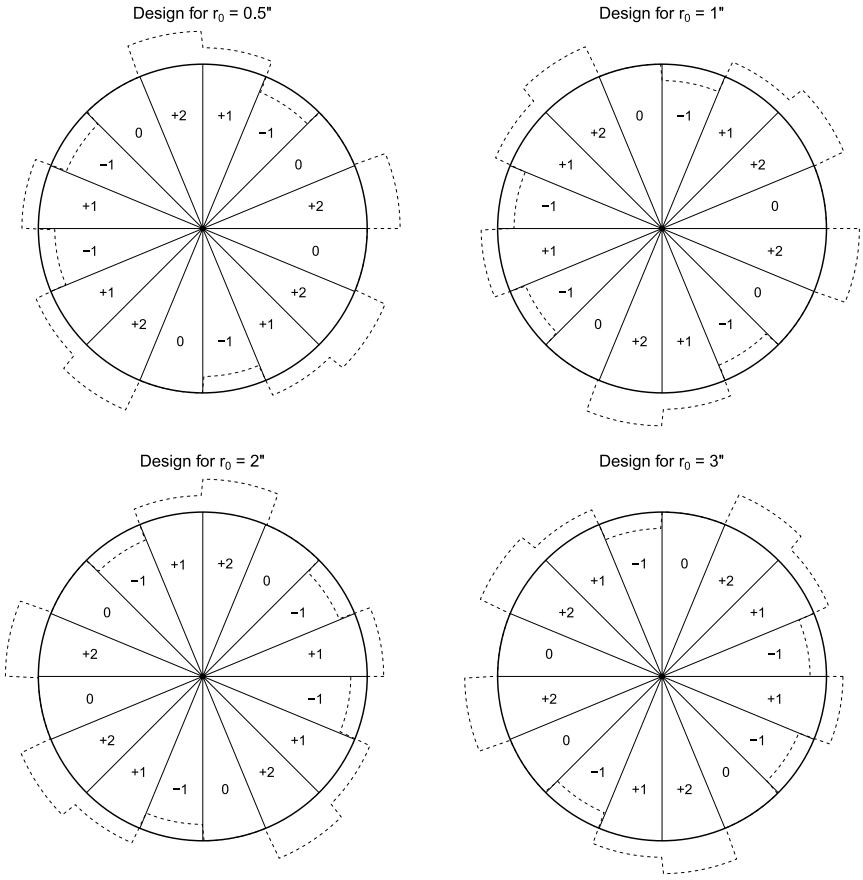

FIG. 3. *Experimental designs. Dashed lines represent assigned compensations.*

We have in Section 3.1 data on cylinders that receive no compensation (denoted by $\mathbf{D}_n$) and a model (9) that provides a good fit. Furthermore, we have a hypothesized model (11) for compensation effects when interference is negligible, which is a function of parameters in (9). If the manufacturing process is in control, posterior inferences based on $\mathbf{D}_n$ then yield, by (11), predictions for the experiment. In the absence of any other information, units in the experiment with observed deformations deviating strongly from their predictions can be argued to have substantial interference. After all, if θ_i has negligible interference under assignment $\mathbf{x} = (x_1, \ldots, x_N)$, then

$$\Delta r(\theta_i, r_0, \mathbf{x}) = \Delta r\big(\theta_i, r_0, (x_i, \ldots, x_i)\big) = \Delta r(\theta_i, r_0, x_i \mathbf{1}).$$

This suggests the following procedure to assess interference:

(1) Calculate the posterior distribution of the parameters conditional on $\mathbf{D}_n$, denoted by $\pi(\alpha, \beta, a, b, x_0, \sigma | \mathbf{D}_n)$.

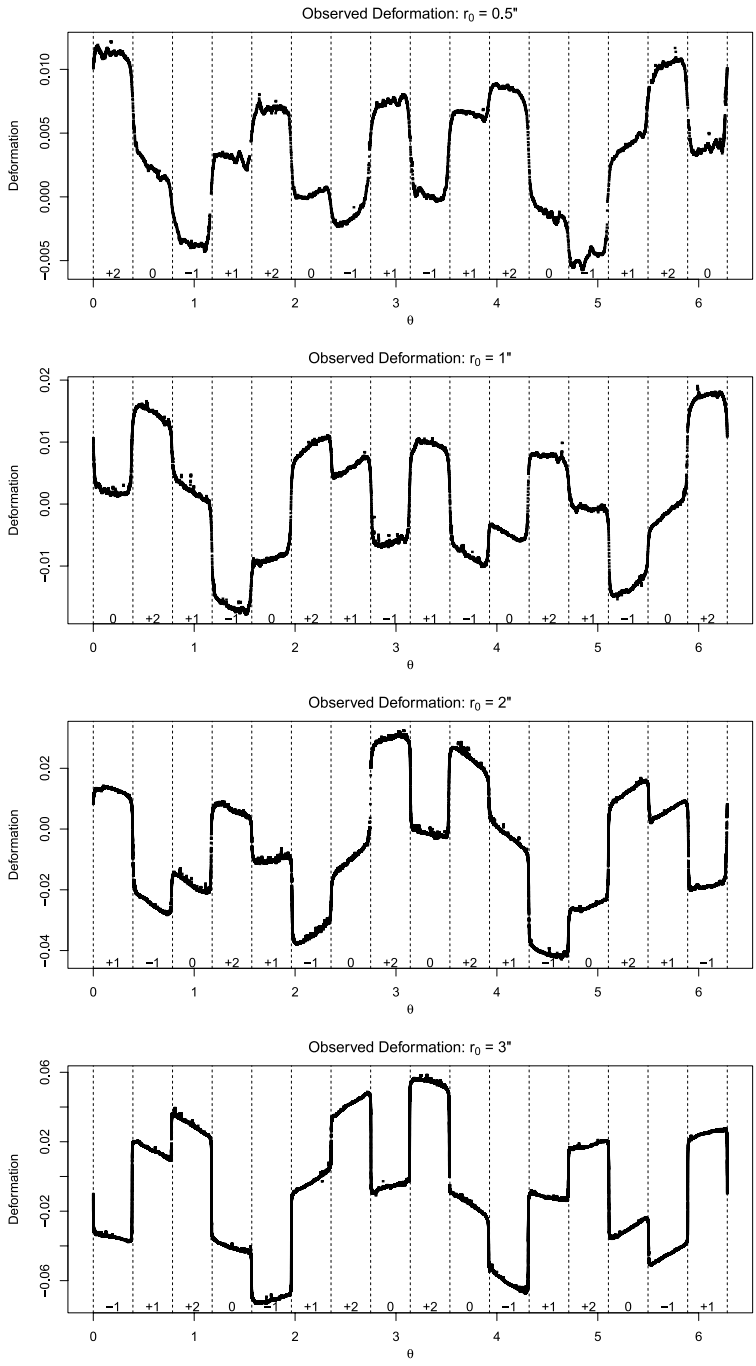

FIG. 4. *Observed deformations in the experiment. Dashed lines represent sections, and numbers at the bottom of each represent assigned compensations.*

(2) For every angle in the four cylinders, form the posterior predictive distribution of the potential outcome corresponding to the observed treatment assignment (Figure 3) using model (11) and $\pi(\alpha, \beta, a, b, x_0, \sigma|\mathbf{D}_n)$.

(3) Compare the posterior predictive distributions to the observed deformations in the experiment.

- If a unit's observed outcome falls within the 99% central posterior predictive interval and follows the posterior predictive mean trend, it is deemed to have negligible interference.
- Otherwise, we conclude that the unit has substantial interference.

This procedure is similar to the construction of control charts [Box, Luceño and Paniagua-Quiñones (2009)]. When an observed outcome lies outside the 99% central posterior predictive interval, we suspect existence of a special cause. As the entire product is manufactured simultaneously, we believe that the only reasonable assignable cause is interference.

We implemented this procedure and observed that approximately 70%–80% of units, primarily in the central regions of sections, have negligible interference (Appendix C). This is clearly seen with another graph that assesses effective treatments, which we proceed to describe.

Taking expectations in (11), the treatment effectively received by θ_i is

$$\begin{aligned}&\big(\mathbb{E}\{\Delta r(\theta_i, r_0, \mathbf{x})|\alpha, \beta, a, b, x_0, \sigma\} - x_0\\ &\quad - \alpha(r_0 + x_0)^a - \beta(r_0 + x_0)^b \cos(2\theta_i)\big)\\ &\quad /\big(1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1}\cos(2\theta_i)\big).\end{aligned} \tag{12}$$

We gauge $g_i(\mathbf{x})$ by plugging observed data from the experiment and posterior draws of the parameters based on $\mathbf{D}_n$ into (12). These discrepancy measure [Meng (1994)] calculations, summarized in Figure 5, again suggest that central angles in each section have negligible interference: estimates of their effective treatments correspond to their assigned treatments. There is a slight discrepancy between assigned treatments and inferred effective treatments for some central angles, but this is likely due to different parameter values for the two data sets. Of more importance is the observation that the effective treatment of a boundary angle θ_i is a weighted average of the treatment assigned to its section, $x_{i,M}$, and its nearest neighboring section, $x_{i,NM}$, with the weights a function of the distances (in radians) between θ_i and the midpoint angle of its section, $\theta_{i,M}$, and the midpoint angle of its nearest neighboring section, $\theta_{i,NM}$. All these observations correspond to the intuition that interference should be substantial near section boundaries.

3.4. *A simple interference model.* We first alter (11) to

$$\begin{aligned}&\Delta r(\theta_i, r_0, \mathbf{x})\\ &\quad = x_0 + \alpha(r_0 + x_0)^a + \beta(r_0 + x_0)^b \cos(2\theta_i)\\ &\qquad + \big\{1 + a\alpha(r_0 + x_0)^{a-1} + b\beta(r_0 + x_0)^{b-1}\cos(2\theta_i)\big\}g_i(\mathbf{x}) + \varepsilon_i,\end{aligned} \tag{13}$$

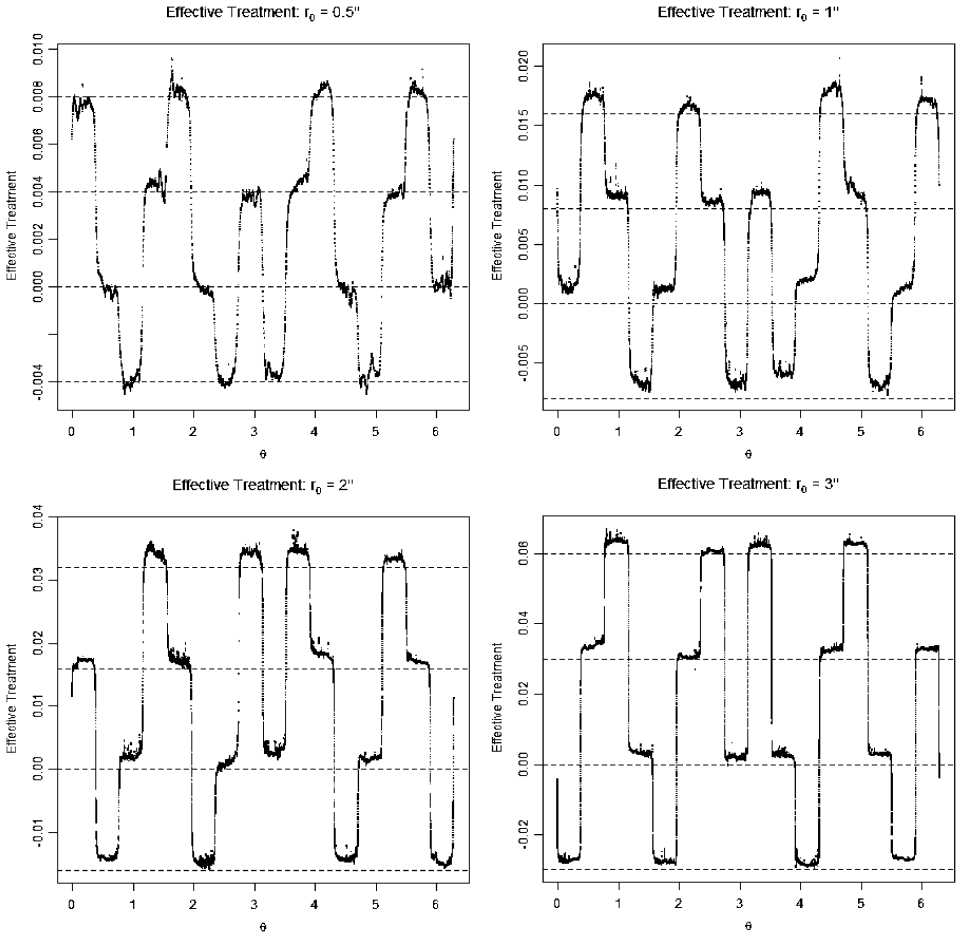

FIG. 5. *Gauging effective treatment* $g_i(\mathbf{x})$ *using* (12). *Four horizontal lines in each subfigure denote the possible compensations, and dots denote estimates of treatments that units effectively received in the experiment.*

where

$$
\begin{aligned}
g_i(\mathbf{x}) = {} & \left\{1+\exp\left(-\lambda_{r_0}|\theta_i-\theta_{i,NM}|+\lambda_{r_0}|\theta_i-\theta_{i,M}|\right)\right\}^{-1} x_{i,M} \\
& + \left\{1+\exp\left(\lambda_{r_0}|\theta_i-\theta_{i,NM}|-\lambda_{r_0}|\theta_i-\theta_{i,M}|\right)\right\}^{-1} x_{i,NM},
\end{aligned} \tag{14}
$$

with $\theta_{i,M}, \theta_{i,NM}$ denoting midpoint angles for the $\pi/8$-radian sections containing and neighboring nearest to θ_i, respectively, and $x_{i,M}, x_{i,NM}$ compensations assigned to these sections. Effective treatment $g_i(\mathbf{x})$ is a weighted average of the unit's assigned treatment $x_i = x_{i,M}$ and the treatment $x_{i,NM}$ assigned to its nearest neighboring section. Although the form of the weights is chosen for computational convenience, we recognize that (14) belongs to a class of models agreeing

TABLE 2
Summary of posterior draws for simple interference model

	Mean	SD	Median	95% credible interval	ESS
α	-1.06×10^{-2}	1.53×10^{-4}	-1.06×10^{-2}	$(-1.09, -1.03) \times 10^{-2}$	8078
β	5.79×10^{-3}	3.69×10^{-5}	5.79×10^{-3}	$(5.72, 5.86) \times 10^{-3}$	8237
a	9.5×10^{-1}	9.46×10^{-3}	9.5×10^{-1}	$(9.31, 9.69) \times 10^{-1}$	8150
b	1.12	6.64×10^{-3}	1.12	(1.0, 1.13)	8504
x_0	7.1×10^{-3}	1.43×10^{-4}	7.1×10^{-3}	$(6.82, 7.39) \times 10^{-3}$	8404
σ	3.14×10^{-3}	1.36×10^{-5}	3.14×10^{-3}	$(3.11, 3.17) \times 10^{-3}$	8924
$\lambda_{0.5}$	32.66	2.05	32.62	(28.69, 36.76)	8686
λ_1	48.24	2	48.12	(44.5, 52.6)	8666
λ_2	76.83	1.78	76.78	(73.42, 80.44)	8770
λ_3	86.08	0.83	86.06	(84.49, 87.68)	8385

with prior subject-matter knowledge that interference may be negligible if the implemented compensation plan is sufficiently "continuous," in the sense that the theoretical compensation plan is a continuous function of θ and the tolerance of the 3D printer is sufficiently fine so that discretization of compensation is negligible (Appendix D).

We fit the model in (13) and (14), having 10 total parameters, to the experiment data. The prior specification remains the same, with $\log(\lambda_{r_0}) \sim \mathrm{N}(0, 4^2)$ independently a priori for $r_0 = 0.5, 1, 2$, and 3 inches. A HMC algorithm was used to obtain 1000 draws from the joint posterior distribution after a burn-in of 500, and these are summarized in Table 2.

This model provides a good fit for the 0.5 and 1 inch cylinders, but not the others. As an example, in Figure 6(a) the posterior mean trend does not correctly capture the observed transition across sections for the 3 inch cylinder. The problem appears to reside in (14). This specification implies that effective treatments of units $\theta_i = k\pi/8$ for $k \in \mathbb{Z}_{>0}$ are equal-weighted averages of compensations applied to units $k\pi/8 \pm \pi/16$. To assess the validity of this implication, we use the posterior distribution of the parameters to calculate, for each θ_i, the inferred effective treatment in (12). An example of these calculations, Figure 6(b), shows that the inferred effective treatment for $\theta_i = \pi$ is nearly 0.06 inch, the compensation applied to the right-side section. Thus, specification (14) is invalidated by the experiment.

Another posterior predictive check helps clarify the problem. From (14),

$$g_i(\mathbf{x}) = w_i x_{i,M} + (1 - w_i) x_{i,NM},$$

and so

$$w_i = \frac{g_i(\mathbf{x}) - x_{i,NM}}{x_{i,M} - x_{i,NM}}, \tag{15}$$

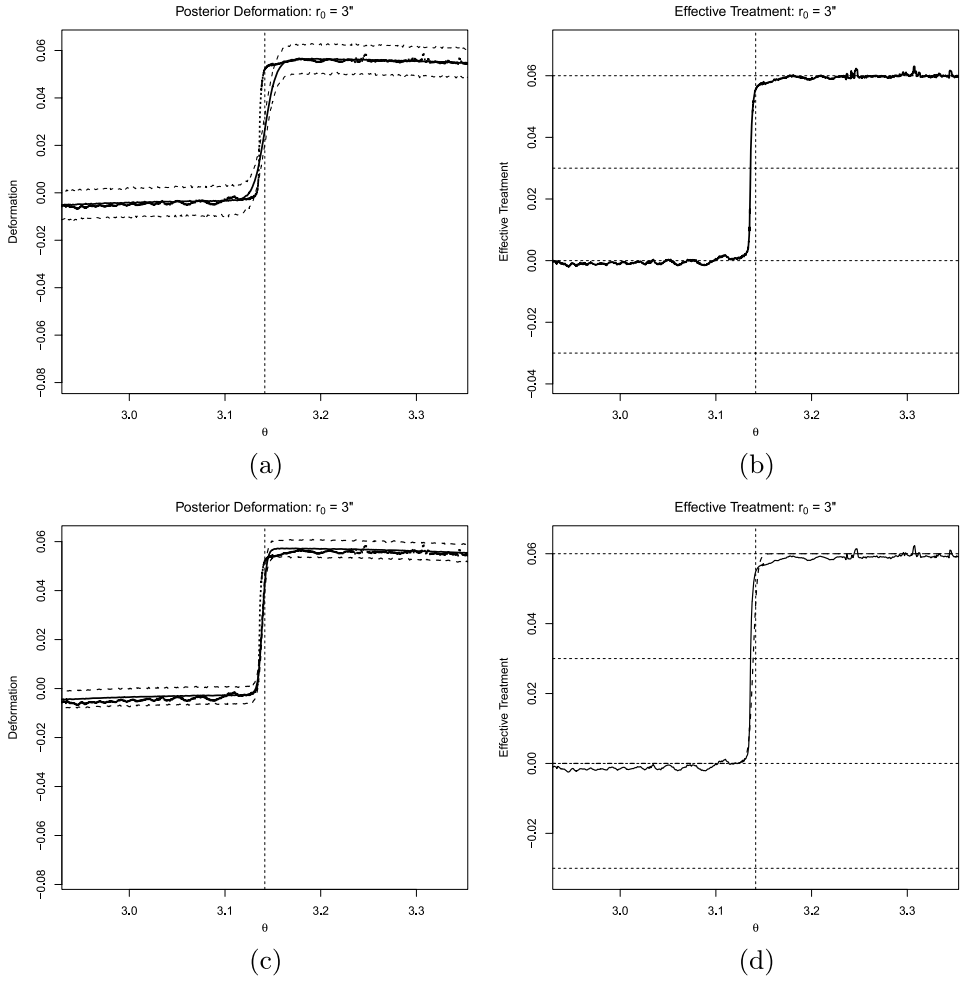

FIG. 6. (a) *An example of the type of erroneous predictions made by model* (13), (14) *for the* 3 *inch cylinder. The vertical line is drawn at* $\theta = \pi$, *marking the boundary between two sections. Units to the left of this line were given* 0 *compensation, and units to the right were given* $+2$ *compensation. The posterior mean trend is represented by the solid line, and posterior quantiles are represented by dashed lines. Observed data are denoted by dots.* (b) *Corresponding inferred effective treatment for* $15\pi/16 \le \theta \le 17\pi/16$. (c) *Refined posterior predictions for* $r_0 = 3$ *inches,* $15\pi/16 \le \theta \le 17\pi/16$. (d) *Comparing inferred effective treatments* (*solid line*) *with refined effective treatment model* (*dashed line*) *for the* 3 *inch cylinder.*

which is well defined because $x_{i,M} \neq x_{i,NM}$ in this experiment. Plugging in the inferred effective treatments, calculated from (12), into (15), we then diagnose how to modify (14) to better model interference in the experiment.

This calculation was made for all cylinders, and the results for $r_0 = 3$ inches are summarized in Figure 7 as an example. Rows in this figure show the weights

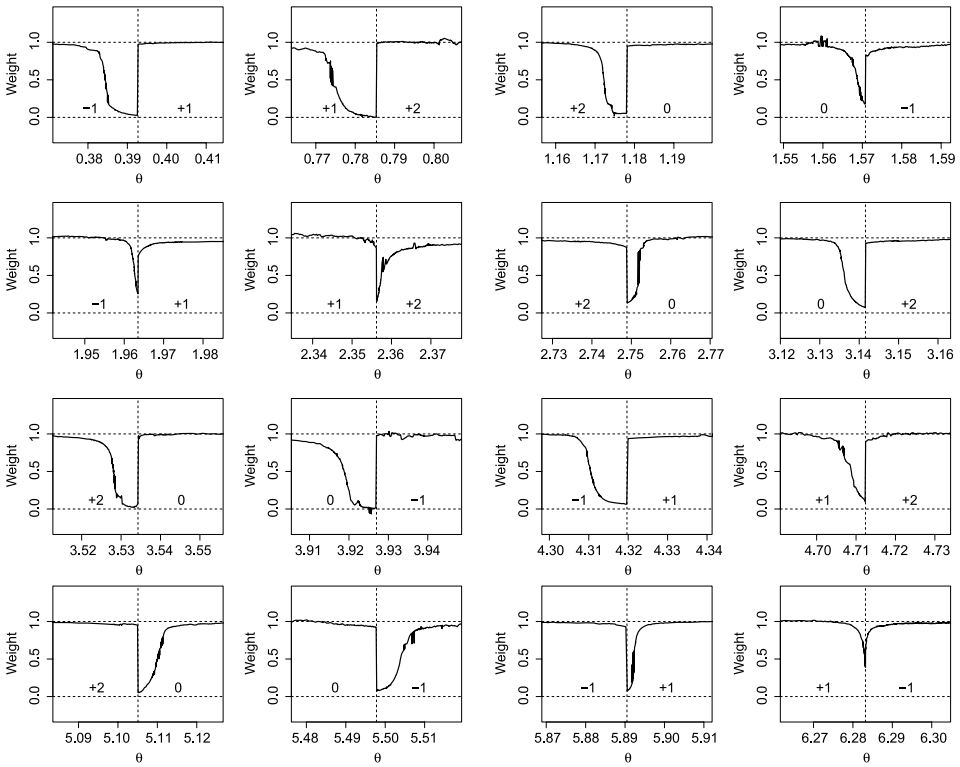

FIG. 7. *Inferring weights w_i in the interference model for the $r_0 = 3$ inch cylinder, using effective treatments calculated from equation (12), based on the posterior distribution of parameters from Section 3.4 and equation (15). Vertical lines represent $\theta = k\pi/8$ for $k = 1, \ldots, 16$, and numbers at the bottom of each subfigure represent assigned compensations.*

for each quadrant, and we focus on their behavior in neighborhoods of integral multiples of $\pi/8$. Neither the decay in the weights [represented by λ_{r_0} in (14)] nor the weight for integral multiples of $\pi/8$ remain constant across sections. In fact, these figures suggest that λ_{r_0} is a function of $\theta_{i,M}, \theta_{i,NM}$, and that a location term is required. They also demonstrate a possible, subtle quadrant effect and, as our experiment blocks on this factor, we are better able to use these posterior predictive checks to refine our simple interference model and capture this unexpected deformation pattern.

3.5. *A refined interference model.* Our refined effective treatment model is of the same form as (14), with λ_{r_0} replaced by $\lambda_{r_0}(\theta_{i,M}, \theta_{i,NM})$, and $|\theta_i - \theta_{i,M}|$, $|\theta_i - \theta_{i,NM}|$ replaced by $|\theta_i - \theta_{i,M} - \delta_{r_0}(\theta_{i,M}, \theta_{i,NM})|$, $|\theta_i - \theta_{i,NM} - \delta_{r_0}(\theta_{i,M}, \theta_{i,NM})|$, respectively. Here, $\delta_{r_0}(\theta_{i,M}, \theta_{i,NM})$ represent location shifts across sections suggested by the previous posterior predictive checks.

Our specific model is

$$\delta_{r_0}(\theta_{i,M}, \theta_{i,NM}) = \delta_{r_0,0} + \sum_{k=1}^{3} \{\delta^c_{r_0,k} \cos(k\theta_{i,B}) + \delta^s_{r_0,k} \sin(k\theta_{i,B})\}, \tag{16}$$

$$\begin{aligned}\lambda_{r_0}(\theta_{i,M}, \theta_{i,NM}) &= \mathbb{I}\big(|x_{i,M} - x_{i,NM}| = 1\big)\lambda_{r_0,1} \\ &\quad + \mathbb{I}\big(|x_{i,M} - x_{i,NM}| = 2\big)\lambda_{r_0,2},\end{aligned} \tag{17}$$

where $\theta_{i,B} = (\theta_{i,M} + \theta_{i,NM})/2$ and $|x_{i,M} - x_{i,NM}|$ is measured in absolute units of compensation. From Figure 7 and the fact that

$$\delta_{r_0}(\theta_{i,M}, \theta_{i,NM}) = \delta_{r_0}(\theta_{i,M} + 2\pi, \theta_{i,NM} + 2\pi),$$

location shifts should be modeled using harmonic functions.

This model provides a better fit. Comparing Figure 6(c), which displays posterior predictions from the refined model (based on one chain of posterior draws using a standard random walk Metropolis algorithm), with the previous model's predictions in Figure 6(a), we immediately see that the refined model better captures the posterior predictive mean trend. Similar improvements exist for the other sections and cylinders. We also compare the original inferred effective treatments obtained from (12) with the refined model in Figure 6(d) and again observe that the new model better captures interference.

3.6. *Summary of the experimental design and analysis*. Three key ingredients relating to the data, model, and experimental design have made our series of analyses possible, and are relevant and useful across a wide variety of disciplines. First is the availability of benchmark data, for example, every unit on the cylinder receiving zero compensation. Second is the potential outcomes model (11) for compensation effects when there is no interference, defined in terms of a fixed number of parameters that do not depend on the compensation plan $\mathbf{x}$. These two enable calculation of the posterior predictive distribution of potential outcomes under the assumption of negligible interference. The final ingredient is the explicit distinction between units of analysis and units of interpretation in our design, which provides the means to assess and model interference in the experiment. Comparing observed outcomes from the experiment to posterior predictions allows one to infer the structure of interference, which can be validated by further experimentation.

These considerations suggest that our methodology can be generalized and applied to other experimental situations with units residing on connected surfaces. In general, when experimenting with units on a connected surface, a principled and step-by-step analysis using the three ingredients above, as illustrated in this paper, can ultimately shed more light on the substantive question of interest.

4. Conclusion: Ignoring interference inhibits improvements. To manufacture 3D printed products satisfying dimensional accuracy demands, it is important to address the problem of interference in a principled manner. Huang et al. (2014) recognized that continuous compensation plans implemented on printers with a sufficiently fine tolerance can effectively control a product's printed dimensions without inducing additional complications through interference. Their models for product deformation motivated our experiment that introduces interference through the application of a discretized compensation plan to the boundary of a cylinder. Combining this experiment's data with inferences based on data for which every unit received no compensation led to an assessment of interference in terms of how units' effective treatments differed from that physically assigned. Further analyses effectively modeled interference in the experiment.

It is important to note that the refined interference model's location and scale terms (16), (17) are a function of the compensation plan. For example, reflecting the assigned compensations across the y axis would accordingly change the location shifts. The implication of this and all our previous observations for manufacturing is that severely discretized compensation plans introduce interference, and, if this fact is ignored, then quality control of 3D printed products will be hindered, especially for geometrically complex products relevant in real-life manufacturing.

Many research challenges and opportunities for both statistics and additive manufacturing remain to be addressed. Perhaps the most important is experimental design in the presence of interference. For example, when focus is on the construction of specific classes of products (e.g., complicated gear structures), optimum designs can lead to precise estimates of model parameters, hence improved compensation plans and control of deformation. An important and subtle statistical issue that then arises is how the structure of interference changes as a function of the compensation plan derived from the experimental design. Instead of being a weighted average of the treatment applied to its section and nearest neighboring section, the derived compensation plan may cause a unit's effective treatment to be a weighted average of treatments applied to other sections as well, with weights depending on the absolute difference in applied compensations. Knowledge of the relationship between compensation plans derived from specific experimental designs and interference is necessary to improve quality control in general, and therefore is an important issue to address for 3D printing.

APPENDIX A: CORRELATION IN ε

In all our analyses, we assumed the ε_i were independent. As pointed out by a referee, when units reside on a constrained boundary, independence of error terms is generally unrealistic. However, we believe that our specific context helps justify this simplifying assumption for several reasons.

First, the major objective driving our work on 3D printing is compensation for product deformation. To derive compensation plans, it is important to accurately

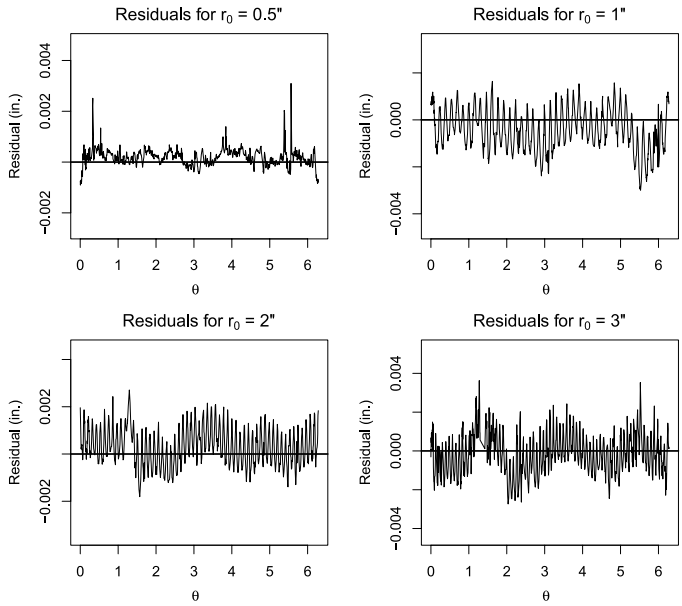

FIG. 8. *Residuals for the model fit in Section 3.1. Here, the residual is defined as the difference between the observed deformation and the posterior mean of deformation for each angle* θ_i.

specify the mean trend in deformation. Although incorporating correlation may change parameter estimates that govern the mean trend, we do not believe that modeling the correlation in errors will substantially help us compensate for printed product deformations. This is something we intend to address further in our future work.

Second, there is a factor that may further confound the potential benefits of including correlated errors in our model: the resolution of the CAD model. To illustrate, consider the model fit in Section 3.1. We display the residual plots in Figure 8. All residuals are (in absolute value) less than 1% of the nominal radius for $r_0 = 0.5$ inch and at most approximately 0.1% of the nominal radius for $r_0 = 1, 2, 3$ inches, supporting our claim that we have accurately modeled the mean trend in deformation for these products. However, we note that for $r_0 = 1, 2, 3$ inches, there is substantial negative correlation in residuals between adjacent units, with the residuals following a high-frequency harmonic trend. There is a simple explanation for this phenomenon. Our first manufactured products were $r_0 = 1, 2, 3$ inches, and the CAD models for these products had low resolution. Low resolution in the CAD model yields the high-frequency pattern in the residual plots. The next product we constructed was $r_0 = 0.5$ inch, and its CAD model had higher resolution than that previously used, which helped to remove this high-frequency pattern. Minor trends appear to exist in this particular plot, and an ACF plot formally reveals significant autocorrelations. Accordingly, we observe that the correlation in residuals is a function of the resolution of the initial CAD model. In consideration of

our current data and our primary objective to accurately capture the mean trend in deformation, we use independent ε_i throughout. We intend to pursue this issue further in our future work, for example, in the direction of Colosimo, Semeraro and Pacella (2008).

Furthermore, as pointed out by the Associate Editor, correlations in residuals for more complicated products may be accounted for by modeling the interference between units, which is precisely the focus of this manuscript.

APPENDIX B: MCMC CONVERGENCE DIAGNOSTICS

Convergence of our MCMC algorithms was gauged by analysis of ACF and trace plots, and effective sample size (ESS) and Gelman and Rubin [(1992), GR] statistics, which were calculated using 10 independent chains of 1000 draws after a burn-in of 500. In Sections 3.1 and 3.4, the ESS were all above 8000 (the maximum is 10,000), and the GR statistics were all 1.

APPENDIX C: ASSESSING INTERFERENCE

The results of the first procedure described in Section 3.3 are displayed in Figure 9: bold lines represent posterior means, dashed lines quantiles forming the 99% central posterior intervals, and dots the observed outcomes in the experiment, with separate figures for each nominal radius and compensation. For example, the graph in the first row and column of Figure 9 contains the observed data for angles in the 0.5 inch radius cylinder that received -1 compensation. This figure also contains the posterior predictive mean and 99% intervals for all angles under the assumption that -1 compensation was applied uniformly to the cylinder. Although only four sections of the cylinder received this compensation in the experiment, forming this distribution makes the posterior predictive mean trend transparent, and so helps identify when a unit's observed outcome deviates strongly from its prediction.

APPENDIX D: NOTE ON A CLASS OF INTERFERENCE MODELS

Compensation is applied in practice by discretizing the plan at a finite number of points, according to some tolerance specified by the size (in radians) for each section or, alternatively, the maximum value of $|\theta_{i,M} - \theta_{i,NM}|$.

Suppose compensation plan $x(\theta)$ is a continuous function of θ, and define

$$w_i = \frac{h(|\theta_i - \theta_{i,M}|)}{h(|\theta_i - \theta_{i,M}|) + h(|\theta_i - \theta_{i,NM}|)},$$

with $h : \mathbb{R} \to \mathbb{R}_{>0}$ a monotonically decreasing continuous function, and

$$g_i(\mathbf{x}) = w_i x_{i,M} + (1 - w_i) x_{i,NM}.$$

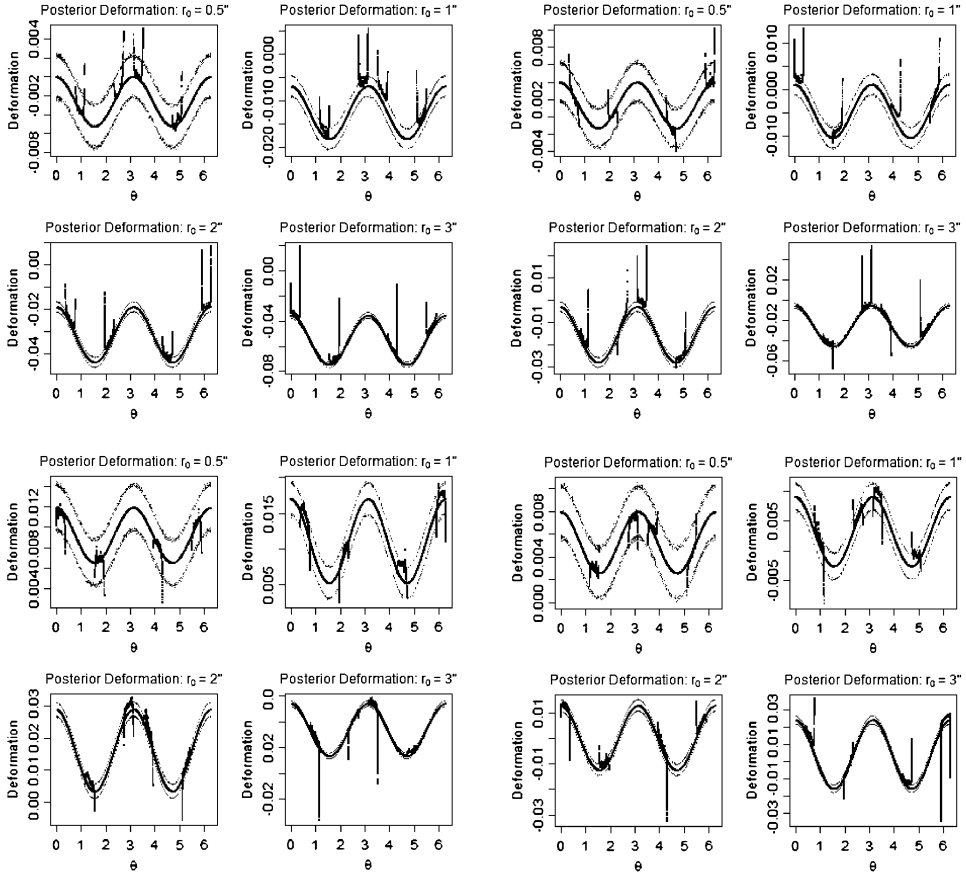

FIG. 9. *Assessing interference in the experiment based on posterior inferences drawn from the no-compensation data. Clockwise from top left*: *predictions for units that received* $-1, 0, +1$, *and* $+2$ *compensation.*

Then for the cylinders considered in our experiment, $g_i(\mathbf{x}) \to x_i$ as $|\theta_{i,M} - \theta_{i,NM}| \to 0$. This is because $|x_{i,M} - x_{i,NM}| \to 0$ as $|\theta_{i,M} - \theta_{i,NM}| \to 0$, and

$$0 \leq |\theta_i - \theta_{i,NM}| - |\theta_i - \theta_{i,M}| \leq |\theta_{i,M} - \theta_{i,NM}|.$$

Acknowledgements. We are grateful to Xiao-Li Meng, Joseph Blitzstein, David Watson, Matthew Plumlee, the Editor, Associate Editor, and a referee for their valuable comments, which improved this paper.

REFERENCES

BOX, G. E. P., LUCEÑO, A. and PANIAGUA-QUIÑONES, M. D. C. (2009). *Statistical Control by Monitoring and Adjustment*, 2nd ed. Wiley, Hoboken, NJ. MR2509873

CAMPBELL, T., WILLIAMS, C., IVANOVA, O. and GARRETT, B. (2011). *Could 3D Printing Change the World? Technologies, Potential, and Implications of Additive Manufacturing*. Atlantic Council, Washington, DC.

COLOSIMO, B. M., SEMERARO, Q. and PACELLA, M. (2008). Statistical process control for geometric specifications: On the monitoring of roundness profiles. *J. Qual. Technol.* **40** 1–18.

COX, D. R. and DONNELLY, C. A. (2011). *Principles of Applied Statistics*. Cambridge Univ. Press, Cambridge. MR2817147

DUANE, S., KENNEDY, A., PENDLETON, B. J. and ROWETH, D. (1987). Hybrid Monte Carlo. *Phys. Lett. B* **195** 216–222.

GELMAN, A. and RUBIN, D. (1992). Inference from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–472.

GIBSON, I., ROSEN, D. and STUCKER, B. (2009). *Additive Manufacturing Technologies*: *Rapid Prototyping to Direct Digital Manufacturing*. Springer, Berlin.

HILTON, P. and JACOBS, P. (2000). *Rapid Tooling*: *Technologies and Industrial Applications*. CRC Press, Boca Raton.

HOLLAND, P. W. (1986). Statistics and causal inference. *J. Amer. Statist. Assoc.* **81** 945–970. MR0867618

HUANG, Q., ZHANG, J., SABBAGHI, A. and DASGUPTA, T. (2014). Optimal offline compensation of shape shrinkage for 3D printing processes. *IIE Transactions on Quality and Reliability*. To appear.

MELCHELS, F., FEIJEN, J. and GRIJPMA, D. (2010). A review on stereolithography and its applications in biomedical engineering. *Biomaterials* **31** 6121–6130.

MENG, X.-L. (1994). Posterior predictive p-values. *Ann. Statist.* **22** 1142–1160. MR1311969

ROSENBAUM, P. R. (2007). Interference between units in randomized experiments. *J. Amer. Statist. Assoc.* **102** 191–200. MR2345537

RUBIN, D. (1980). Comment on "Randomization analysis of experimental data: The Fisher randomization test," by D. Basu. *J. Amer. Statist. Assoc.* **75** 575–595.

SOBEL, M. E. (2006). What do randomized studies of housing mobility demonstrate?: Causal inference in the face of interference. *J. Amer. Statist. Assoc.* **101** 1398–1407. MR2307573

TONG, K., JOSHI, S. and LEHTIHET, E. (2008). Error compensation for fused deposition modeling (FDM) machine by correcting slice files. *Rapid Prototyping J.* **14** 4–14.

TONG, K., LEHTIHET, E. and JOSHI, S. (2003). Parametric error modeling and software error compensation for rapid prototyping. *Rapid Prototyping J.* **9** 301–313.

WANG, W., CHEAH, C., FUH, J. and LU, L. (1996). Influence of process parameters on stereolithography part shrinkage. *Mater. Des.* **17** 205–213.

A. SABBAGHI
DEPARTMENT OF STATISTICS
PURDUE UNIVERSITY
250 N. UNIVERSITY STREET
WEST LAFAYETTE, INDIANA 47907
USA
E-MAIL: armansabbaghi.stat@gmail.com

T. DASGUPTA
DEPARTMENT OF STATISTICS
HARVARD UNIVERSITY
1 OXFORD STREET, 7TH FL.
CAMBRIDGE, MASSACHUSETTS 02138
USA
E-MAIL: dasgupta@stat.harvard.edu

Q. HUANG
J. ZHANG
DANIEL J. EPSTEIN DEPARTMENT OF INDUSTRIAL
AND SYSTEMS ENGINEERING
UNIVERSITY OF SOUTHERN CALIFORNIA
3715 MCCLINTOCK AVENUE
LOS ANGELES, CALIFORNIA 90089
USA
E-MAIL: qiang.huang@usc.edu
jizhezha@usc.edu