Journal of Business & Economic Statistics

ISSN: 0735-0015 (Print) 1537-2707 (Online) Journal homepage: https://www.tandfonline.com/loi/ubes20

Discussion of “Transparency in Structural Research” by Isaiah Andrews, Matthew Gentzkow, and Jesse Shapiro

Stéphane Bonhomme

To cite this article: Stéphane Bonhomme (2020) Discussion of “Transparency in Structural Research” by Isaiah Andrews, Matthew Gentzkow, and Jesse Shapiro, Journal of Business & Economic Statistics, 38:4, 723-725, DOI: 10.1080/07350015.2020.1790377

To link to this article: https://doi.org/10.1080/07350015.2020.1790377

Published online: 28 Sep 2020.

Discussion of "Transparency in Structural Research" by Isaiah Andrews, Matthew Gentzkow, and Jesse Shapiro

Stéphane BONHOMME
Department of Economics, University of Chicago, Chicago, IL (*sbonhomme@uchicago.edu*)

It is a pleasure to discuss this article. In two previous articles (Andrews, Gentzkow, and Shapiro 2017, 2019) and in this one, the authors propose tools to facilitate the interpretation of structural estimation, and to better understand the mapping between model, data, and results. Their work represents an important contribution to a much needed area of research.

In this new article, they propose a definition of transparency and review a number of approaches to improve it. This includes reporting two measures they proposed previously: sensitivity (Λ) and informativeness (Δ), as well as using methods for sensitivity analysis.

I particularly like the discussion of identification, and the authors' emphasis on descriptive tools as complements to, rather than substitutes for, formal identification analyses. The article contains a number of very useful insights and methods for structural economists and other applied researchers.

My comments are related to what the authors call the *audience*; that is, the consumers and users of the research. Research is always directed to an audience: other researchers, a readership, or a decision maker. One may say that a research report is "transparent" when a member of the audience can "see through it," as one sees through a window. However, what transparency means depends greatly on who is looking through that window; in the present context, on the audience of the scientific research.

I will focus on three questions: Can a definition of transparency account for the complexity of the researcher's report (when the audience is not as sophisticated as the researcher)? Can communication through models make research more transparent (for an audience of economists)? Is transparency sufficient to ensure that structural policy predictions are credible (for an audience of policy makers)?

1. TRANSPARENCY AND COMPLEXITY

The authors' measure of transparency abstracts from the complexity of the researcher's report, and from the ability of the audience to process the information contained in the report.

Yet, the degree of sophistication of the audience matters for assessing the transparency of research. For example, consumers of research may face more severe computational constraints, lack subject-matter expertise, or have a different opportunity cost of time compared to the researcher. An important role of empirical research is to communicate scientific findings, often by producing concise summaries of the data that the audience can interpret. This role becomes increasingly important as data sets become bigger and more complex.

To see why these aspects are not currently captured by the authors' setup, consider the communication risk of Andrews and Shapiro (2020), focusing on the case where there is a single member of the audience, and abstracting from the r subscript for conciseness. Communication risk of a report $(\widehat{c},\widehat{t})$ is

$$R\left(\widehat{c},\widehat{t}\right)=\mathbb{E}\left(\min_{d(\widehat{c},\widehat{t})}\mathbb{E}\left[\left(d\left(\widehat{c},\widehat{t}\right)-c\right)^2\,\middle|\,\widehat{c},\widehat{t}\right]\right).$$

As the authors mention, $R\left(\widehat{c},\widehat{t}\right)$ is minimum when the report $\left(\widehat{c},\widehat{t}\right)$ is the entire data D. In that case, the authors' transparency index is equal to one ("full transparency").

Now, let me make a simple modification to this setup. As the authors point out, $R\left(\widehat{c},\widehat{t}\right)=\mathbb{E}\left[\mathrm{var}\left(c\,|\,\widehat{c},\widehat{t}\right)\right]$ is equal to the within-$\left(\widehat{c},\widehat{t}\right)$ variance of c. This quantity can alternatively be written as the minimum Bayes risk over decision rules that are functions of the report; that is,

$$R\left(\widehat{c},\widehat{t}\right)=\min_{d(\cdot)}\mathbb{E}\left[\left(d\left(\widehat{c},\widehat{t}\right)-c\right)^2\right],$$

where the minimization is now over decision rules. Suppose that only a limited set of decision rules is available to the audience. Denote this set as $\mathcal{D}_m$, where m is the dimension of the report. The audience has a set $\mathcal{D}_m$ of decision rules for each m, and the researcher chooses m, as well as the m-dimensional vector $\left(\widehat{c},\widehat{t}\right)$ that she reports to the audience. The previous expression motivates defining a *constrained* version of communication risk of $\left(\widehat{c},\widehat{t}\right)$ as

$$R^{\mathrm{c}}\left(\widehat{c},\widehat{t}\right)=\min_{d(\cdot)\in\mathcal{D}_m}\mathbb{E}\left[\left(d\left(\widehat{c},\widehat{t}\right)-c\right)^2\right].$$

When the set of decision rules available to the audience is restricted, it may no longer be optimal for the researcher to report the data D. To see this in a simple example, consider a case where $D=(D_1,\ldots,D_K)'$, $\mathbb{E}(c\,|\,D)=D'\beta_0$, and the set of decision rules for reports $X\in\mathbb{R}^m$ is

$$\mathcal{D}_m=\left\{d(\cdot)\,:\,d(X)=X'\beta\,,\,\|\beta\|_0\leq 1\right\},$$

Journal of Business & Economic Statistics
October 2020, Vol. 38, No. 4
DOI: 10.1080/07350015.2020.1790377

where $\|\beta\|_0$ denotes the number of nonzero coefficients in β. A rationale for considering this class may be that the audience understands bivariate relationships between c and any scalar component of the report well, but is not sophisticated enough to analyze multivariate relationships.[1]

In this setting, reporting $D'\beta_0$ to the audience is strictly better than reporting D in general. Indeed, when the report is D the audience's constrained optimal decision given D depends on a single element D_k, for some $k = 1, \ldots, K$, whereas the unconstrained optimal decision $\mathbb{E}(c \mid D) = D'\beta_0$ is a linear combination of all elements $D_1, \ldots, D_K$. Hence, this simple modification of the setup provides a rationale for the researcher to communicate specific features of the data to the audience.

Note that in this example there are multiple optimal reports. For instance, reporting both $D'\beta_0$ and D together also achieves the minimum value of R^c. It may be desirable to penalize R^c for the complexity of the set of decision rules. In the present setting, the dimension m of the report is a natural measure of the complexity of $\mathcal{D}_m$. For example, the cost, for the audience, of having to examine m bivariate relationships may scale linearly with m. Such a complexity penalty will induce the researcher to report more parsimonious summaries of the data.

There are different routes one could follow to extend the authors' framework to account for a not-fully sophisticated audience. The modifications I have discussed here are simply meant to start thinking about how this might be done.

2. TRANSPARENCY THROUGH MODELS

Models are a common communication technology among economists. Moreover, "auxiliary" models are often used in combination with structural models, to convey the mechanism and study identification in simple settings, and to estimate the parameters of the structural model by indirect inference. While the authors focus on situations where the report contains a vector $\widehat{s}$ of descriptive statistics, discussing multiple models and reporting estimates based on them can also make structural analysis more transparent.

My goal here is to discuss the usefulness of combining structural models with reduced-form (or "quasi-structural") models, and how this combination can improve the transparency of empirical work. To do so I will focus on life-cycle models of consumption, saving and income, which are the topic of the structural analysis in Gourinchas and Parker (2002) that the authors use as an example.

In a standard life-cycle model of consumption and saving, Hall (1978) and Hall and Mishkin (1982) considered the special case where households' utility is quadratic. This specification is restrictive, since in particular it does not account for precautionary saving behavior. Yet, under quadratic utility there is a simple relationship between consumption growth and shocks to income. Specifically, in a model where the income level is the sum of a permanent component (a random walk $\eta_{it} = \eta_{i,t-1} + V_{it}$) and a transitory component (a white noise ε_{it}), Hall and Mishkin showed (1982) that consumption growth is given by

$$C_{it} - C_{i,t-1} = \beta_t V_{it} + \gamma_t \varepsilon_{it}, \tag{1}$$

which is linear in the permanent and transitory shocks to income. (1) holds exactly in the structural model with quadratic utility, when interest rate and discount rate coincide. In that case $\beta_t = 1$, and γ_t is a function of the interest rate and the time horizon.

The crucial point is that Equation (1) has a "life of its own," in the sense that it is empirically useful irrespective of whether the researcher is willing to assume that preferences are quadratic. For example, (1) can be directly estimated using panel data to document the degree of insurance to income fluctuations that households possess. This relationship also provides moments that can be used to study identification or estimate the parameters of a structural model. This approach has been very influential in the literature, especially through the development of log-linearized approximations under more general utility classes (Carroll 1997; Blundell, Pistaferri, and Preston 2008).

A drawback of specifications based on quadratic utility or log-linearized approximations is that they do not generally nest structural models. A different strategy is to work directly from the nonlinear reduced form; that is, to estimate the nonlinear policy rules, laws of motion of state variables, and payoff equations. Arellano, Blundell, and Bonhomme (2017) applied such an approach to income and consumption dynamics. In their setup, consumption is a nonlinear function of assets, permanent income, and transitory income, of the form

$$C_{it} = g_t(A_{it}, \eta_{it}, \varepsilon_{it}, \nu_{it}), \tag{2}$$

where A_{it} are assets, g_t is a nonlinear, unspecified function, and ν_{it} are unobserved taste shocks.

The approach based on the nonlinear reduced form emphasizes two key features of the model: heterogeneity and exogeneity. In the present setting, heterogeneity in (2) is represented by the observed and latent state variables, that is, A_{it} and $(\eta_{it}, \varepsilon_{it}, \nu_{it})$, respectively. In turn, exogeneity assumptions take the form of conditional independence given state variables. For example, when the taste shocks ν_{it} are independent of all other current and past state variables, consumption at time t is independent of $\eta_{i,t-1}$ conditional on $(A_{it}, \eta_{it}, \varepsilon_{it})$. Such dynamic restrictions are helpful to ensure that the nonlinear reduced-form model is identified despite the presence of the latent income components η_{it} and ε_{it}. However, except for heterogeneity and exogeneity, (2) does not restrict other features of the structural model such as the functional forms of preferences and expectations.

Communicating estimates of nonlinear reduced forms such as (2) can be useful to improve the transparency of structural work. Indeed, derivative effects such as marginal propensities to consume, and impulse responses to income shocks, can be inferred from the reduced-form equations without the need to recover the underlying structural primitives. Such empirical quantities are often of interest by themselves, and they provide natural "targets" for structural estimation.

However, an important limitation of nonlinear reduced forms is that they cannot be directly used for counterfactual prediction in general. This leads me to the last point in this discussion.

[1]The ℓ^0 "norm" $\|\cdot\|_0$ is related to the ℓ^1 norm used in the Lasso. ℓ^1-penalization is being used in many fields. An economic application of it is the theory of "sparse bounded rationality" in Gabaix (2014).

3. TRANSPARENCY OF POLICY PREDICTIONS

Perhaps the main goal of structural economics is to produce predictions of the effects of policies that have not yet been implemented. To discuss transparency in this context, for concreteness I will focus on the evaluation of a conditional cash transfer program in Mexico by Todd and Wolpin (2006, 2008), which is also the empirical context studied by Attanasio, Meghir, and Santiago (2012) that the authors use as an example.

In their 2006 article, Todd and Wolpin specified and estimated a dynamic structural model, and used the estimated model for counterfactual prediction. In their subsequent 2008 article, they showed that certain counterfactual questions can be answered without the entire structure of the model. Specifically, they explained how a matching estimator can be used to predict the effect of a subsidy on school attendance for a subsample of children, even in the absence of subsidy variation in the sample. Their analysis exploits variation in children's wage and households' income, which under the model is sufficient to compensate for the lack of variation in the subsidy.

Whenever available, such strategies that rely on nonlinear reduced-form implications of the structural model, instead of requiring its full structure, are useful tools to make the analysis more transparent.

However, many policy-relevant counterfactual predictions use the detailed specification of the dynamic structural model, including preferences, expectations, and their functional forms. In this context, ensuring that research is transparent may not be sufficient for a structural prediction to be successful. While transparency helps interpreting the mapping between assumptions and results, the counterfactual c of interest often lies far from the available sample, and any descriptive statistic or empirical model may not help predict c very much. Even a perfectly transparent account of the assumptions and how they map to the results may not be very informative about how credible a particular counterfactual exercise is.

An alternative approach is to evaluate the predictions of the structural model out of sample. Todd and Wolpin (2006) performed such a "validation" exercise by taking advantage of a randomized control trial in Mexico. Specifically, they estimated their structural model using control villages only, and then used treated villages as a "hold-out" sample for evaluation. Intuitively, the out-of-sample evaluation provides an indirect check about the assumptions of the model, to the extent that those matter for the sort of counterfactual questions that the model will be used for. When interested in counterfactual predictions, such validation exercises based on hold-out samples are valuable complements to the strategies that the authors discuss in the article.

4. CONCLUSION

Questions related to misspecification and sensitivity analysis are central to applied research, and the three articles that the authors have written on these topics make important progress. In particular, this third opus provides useful guidance to structural empirical researchers on how to make their work more transparent.

Based on this discussion, I suggest three lines for further work: refining the definition and measure of transparency to account for the complexity of research reports, studying examples where $\widehat{s}$ is an estimated nonlinear empirical model, and formalizing the usefulness of structural validation exercises. There are important open questions related to this last point, such as what information (if any) to hold out when estimating a structural model, and how to assess the power of out-of-sample validation tests. I expect that the authors' work will motivate more research on these issues.

Acknowledgments

I thank the authors, Manuel Arellano, and Lars Hansen for helpful comments.

[Received April 2020. Revised June 2020.]

References

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017), "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 132, 1553–1592. [723]

——— (2019), "On the Informativeness of Descriptive Statistics for Structural Estimates," NBER Working Paper No. 25217. [723]

Andrews, I., and Shapiro, J. M. (2020), "A Model of Scientific Communication," Working Paper. [723]

Arellano, M., Blundell, R., and Bonhomme, S. (2017), "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85, 693–734. [724]

Attanasio, O. P., Meghir, C., and Santiago, A. (2012), "Education Choices in Mexico: Using a Structural Model and a Randomized Experiment to Evaluate PROGRESA," *Review of Economic Studies*, 79, 37–66. [725]

Blundell, R., Pistaferri, L., and Preston, I. (2008), "Consumption Inequality and Partial Insurance," *American Economic Review*, 98, 1887–1921. [724]

Carroll, C. D. (1997), "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55. [724]

Gabaix, X. (2014), "A Sparsity-Based Model of Bounded Rationality," *Quarterly Journal of Economics*, 129, 1661–1710. [724]

Gourinchas, P.-O., and Parker, J. A. (2002), "Consumption Over the Life Cycle," *Econometrica*, 70, 47–89. [724]

Hall, R. E. (1978), "Stochastic Implications of the Life Cycle Permanent Income Hypothesis: Theory and Evidence," *Journal of Political Economy*, 86, 971–987. [724]

Hall, R. E., and Mishkin, F. S. (1982), "The Sensitivity of Consumption to Transitory Income: Estimates From Panel Data on Households," *Econometrica*, 50, 461–481. [724]

Todd, P. E., and Wolpin, K. I. (2006), "Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility," *American Economic Review*, 96, 1384–1417. [725]

——— (2008), "Ex Ante Evaluation of Social Programs," *Annales d'Economie et de Statistique*, 263–291. [725]