

Principles for Combining Descriptive and Model-Based Analysis in Applied Microeconomics Research

Neale Mahoney[†]

Abstract

In this article, I offer guidance on how to combine descriptive and model-based empirical analysis within a paper. Drawing on examples from three recently published applied microeconomics papers, I argue that it is important to create a tight link between the descriptive analysis and the bottom-line deliverable of the model-based analysis, and I try to distill some lessons or principles for doing so. I also offer some thoughts on when a paper should start with descriptive analysis and then proceed to model-based analysis and when alternative structures may be desirable.

[†]This paper was prepared for the *Journal of Economic Perspectives* and was completed before I took leave to work in government. I thank my frequent collaborators Liran Einav and Amy Finkelstein for numerous conversations that have shaped my thinking on this topic. They deserve no blame for any faults that remain.

At a fundamental level, there is no sharp distinction between descriptive and model-based empirical analysis. For example, while it is natural to think about the partial correlations estimated by ordinary least squares regression as descriptive, interpreting these estimates requires evoking, at least implicitly, a linear model of the underlying relationship. However, it is helpful to distinguish between different types of empirical analysis, and I find the descriptive versus model-based terminology useful. I will use the term descriptive analysis to describe empirical analysis when the primary goal is to summarize patterns in the data. I will use the term model-based analysis when the goal is to estimate an economic parameter, conduct a counterfactual, or to make a statement about welfare.

In this article, I offer guidance on how to combine descriptive and model-based empirical analysis within a paper, drawing on my experience as a reader, author, and most recently a co-editor of applied microeconomics research. I will argue that it is important to construct a paper so that there is a tight link between the descriptive analysis and the bottom-line deliverable of the model-based analysis. To ground the discussion, I will begin with three recently published applied microeconomics papers: a health economics paper on prescription drug utilization, an education economics paper on school choice mechanisms, and a consumer finance paper on the pass-through of interest rates.

Drawing on examples from these papers, I will try to distill some lessons or principles. I will discuss the benefits of descriptive analysis, both for showing your identifying variation in a clear and intuitive way, and also for providing preliminary or partial evidence in support of your conclusions, even if your bottom-line conclusions require a quantitative model.¹ I will argue that you should clearly articulate the value-added of the model by explaining what you can learn from the model that cannot be learned from the descriptive analysis alone. I will also argue that you should use the descriptive analysis to guide your choices of what to model and what *not* to model. Finally, I will argue that you should choose parameters or counterfactuals

¹ I have intentionally avoided the reduced form versus structural terminology. These words have a precise meaning in certain contexts, and I do not want to risk confusion by using them imprecisely.

that are informed by the identifying variation in the data and use descriptive analysis to help the reader form a prior belief over the parameter estimates and counterfactuals that follow.

For most of this essay, I will assume that you have decided to write a paper that starts with descriptive analysis and then proceeds to model-based analysis. This is not the only – or necessarily the best – way to craft a paper. Toward the end of the essay, I will share some thoughts on when this ordering may be desirable. I will also offer a perspective on viewing research in applied microeconomics as offering a set of tradeoffs, in which the researcher needs to justify additional model-based assumptions in terms of the additional insights they deliver.

Three Running Examples

I will work with three examples, drawn from papers that span three fields within applied microeconomics: health economics, education, and consumer finance. These papers take very different approaches to the descriptive and model-based analyses. I've also chosen papers written by people I know – including a paper where I was a coauthor. The reason is that I wanted to have frank conversations with the authors about the reasons behind the choices they made. Here I provide brief summaries of the papers, focusing on the connections between the descriptive and model-based analyses that are the focus of this article.

The “Donut Hole” in Medicare Part D

The Medicare Part D program provides insurance for drug expenditures for the elderly in the U.S. However, for many years, the insurance contract had an infamous “donut hole”: consumers were subsidized by the program up to a lower level of annual expenditures and above a higher level of annual expenditures – but there was a region in the middle of the contract (the donut hole) where consumers had to pay the full cost of drugs out of pocket. Einav, Finkelstein, and Schripf (2015) study how spending on prescription drugs is influenced

by the nonlinearity in incentives created by the donut hole. They also use the variation from this nonlinearity to more broadly study consumer behavior and to estimate the impact of counterfactual insurance contracts.

In the first part of the paper, the authors use the sharp jump in the out-of-pocket price at the “kink” in the contract at the start of the donut hole to generate visually compelling descriptive evidence. They show substantial “bunching” of annual spending at the kink, which allows them to reject the null of no response to incentives. They also show that the probability of a new drug purchase decreases as spending reaches the kink, with stronger impacts in December than earlier months in the year. The anticipatory response prior to December shows that people are forward-looking, while the stronger effects in December indicate that either uncertainty or partial myopia limits the responses in earlier months.

In the second part of the paper, the authors build a dynamic model of drug utilization, which allows for a stochastic health process, price sensitivity, and (partial) forward-looking behavior. The estimated model allows the authors to go beyond the qualitative evidence on bunching and quantify the response to the nonlinear incentives created by the donut hole by comparing outcomes under the observed nonlinear contract to the ones resulting from a linear counterfactual contract. The model also allows the authors to quantify behavior in terms of economic parameters, such as a weight the consumers place on future outcomes in their decision-making. Appealingly, the model is estimated via generalized method of moments to match the bunching patterns documented in the descriptive analysis.

High School Choice in New Haven

The New Haven, Connecticut school district has offered students a mechanism for choosing between high schools since the 1990s. Such school choice mechanisms are used in many cities, and they raise some common concerns. Are students better off listing their actual choices, even knowing that certain popular schools will be oversubscribed with those listing it as a first

choice? Or should students instead play a strategic game in which they put a second or third choice at the top of their list, in the belief that they will have a better chance of actually getting into that school if they list it first? At the time of the study, New Haven had implemented a mechanism that, by allowing applicants to express the intensity of their preferences, rewarded strategic play. The tradeoff is that if students are informed and sophisticated, the ability to express intensity of preferences in the New Haven school choice mechanism can improve welfare relative to a “strategy-proof” mechanism which only rewards listing one’s true rank-order preferences. However, if students are uninformed or unsophisticated, the New Haven mechanism could lead to lower average and less equitable outcomes.

Kapor, Neilson, and Zimmerman (2020) study how the accuracy of beliefs affects the welfare from different school choice mechanisms. In the first part of the paper, the authors describe results from a survey of the school preferences of 417 students, combined with data on how students listed their school choices on their administrative application forms. They document that 32 percent of students are “revealed strategic,” in the sense that they did not list their most preferred school first in their submission to the school district. However, the authors also show that this strategic behavior is poorly informed. A descriptive analysis shows that half of the revealed strategic students are “mistakenly strategic,” in the sense that they would have been better off listing their preferred school first rather than strategically listing another school first. Based on survey responses, students often hold beliefs that differ substantially from rational expectations about the probabilities of admission: for example, they are on average highly optimistic about their admission probabilities at second-ranked schools.

In the second part of the paper, the authors build a model of high school admission applications that allows for beliefs to diverge from rational expectations as documented in the descriptive analysis. A key decision is how to model subjective beliefs. The descriptive analysis provides no evidence of strategic information acquisition, so the authors do not allow subjective beliefs to vary with preferences. Instead, motivated by the descriptive analysis, the authors allow the

wedge between subjective and rational beliefs to vary with the rank of the school chosen, priority of schools, and idiosyncratic school and individual components.

The authors use their estimated model to conduct counterfactuals that connect directly to the results from the survey. For example, the authors show that in a situation where subjective and rational beliefs diverge, a (counterfactual) strategy-proof mechanism would achieve higher welfare and improve equity. Indeed, the authors show that one needs to eliminate nearly all of the wedge between subjective beliefs and rational expectations for the New Haven mechanism, with its additional ability to express intensity of preferences, to be preferable on welfare and equity grounds. Finally, the authors show that if researchers didn't account for subjective beliefs and assumed rational expectations when estimating their model, they would have erroneously concluded that the New Haven mechanism was superior.

Pass-Through of Lower Interest Rates for Banks into Increased Borrowing by Consumers

Central banks, such as the Federal Reserve, can stimulate the economy by providing banks with lower-cost capital and liquidity. The idea is that these lower costs will encourage banks to expand credit to consumers who will, in turn, increase their borrowing and spending. Agarwal, Chomsisengphet, Mahoney and Strobel (2017) argue that the impact of a reduction in banks' cost of funds on aggregate borrowing can be decomposed into the product of banks' marginal propensity to lend to borrowers and those borrowers' marginal propensity to borrow, aggregated over all borrowers in the economy. They study how frictions, such as asymmetric information, affect the pass-through of lower interest rates for banks into increased borrowing and spending by consumers. They apply this framework by estimating heterogeneous marginal propensities to borrow by consumers and marginal propensities to lend by banks in the U.S. credit card market.

In the first part of the paper, the authors directly estimate consumers' marginal propensity to borrow using quasi-experimental variation in credit limits. Banks sometimes set credit limits as

discontinuous functions of consumers' credit scores. For example, a bank might grant a \$2,000 credit limit to consumers with a credit score below 720 and a \$5,000 credit limit to consumers with a credit score of 720 or above. The authors identify 743 credit limit discontinuities in their data, located across the credit score distribution, and use these discontinuities to estimate heterogeneous marginal propensities to borrow for consumers with different credit scores.

In the second part of the paper, the authors turn to estimating banks' marginal propensity to lend to different customer groups. Estimating the marginal propensity to lend in a direct way using observed changes in banks' borrowing costs is challenging because such changes are typically correlated with shifts in the economic environment that also affect borrowing and lending decisions. The authors write down a model of optimal credit limits to show that a bank's marginal propensity to lend depends on a small number of sufficient statistics that capture the relationship between changes in lending and profits. These sufficient statistics can be estimated using the same credit limit discontinuities, allowing the authors to recover heterogeneous marginal propensities to lend to borrowers with different credit scores. The authors show that bank lending is close to the optimal level implied by the model, providing support for the modeling assumptions.

In the final part of the paper, the authors combine the model-free estimates of consumers' marginal propensity to borrow with the model-based estimates of banks' marginal propensity to lend. They then use these estimates to describe the strength of this bank lending channel and show how features of the economic environment, which influence the marginal propensity to borrow and to lend, affect the strength of this channel.

Five Principles

In this section, I discuss five principles for combining descriptive and model-based analysis, as illustrated by the three papers summarized above.

1. Show Your Variation with Descriptive Analysis.

Many applied microeconomics papers are built around an empirical approach (sometimes referred to as a research design). For this type of paper, a primary goal of the descriptive analysis is to “make the case” for the identifying variation that drives the rest of the empirical analysis. Broadly, your aim should be to explain where your variation comes from, show that it is powerful, and show that it is valid.

The right way to show the variation depends on the context. In the Medicare Part D paper, the key source of variation is the donut hole that exposes beneficiaries to increased out-of-pocket costs. The authors show that the donut hole can be characterized by a kink in the contract that maps drug spending to out-of-pocket costs. They describe and visually illustrate the donut hole in the standard insurance contract, and in the non-standard contracts that they also use in their analysis.

In the credit card paper, the key variation is the jump in credit limits at specific credit scores, which the authors take advantage of by using a regression discontinuity design. To explain and illustrate this variation, the authors provide institutional context on how bank underwriting models give rise to these type of jumps in credit limits and provide visual examples of the discontinuities in their data. They then establish the validity of these credit limit quasi-experiments by showing that other factors trend smoothly through the discontinuities and show there is no evidence of bunching above the discontinuities.

2. Use the Descriptive Analysis to Provide Preliminary Evidence.

As an author of applied microeconomics research, you should also use the descriptive analysis to provide preliminary or partial evidence for the paper’s conclusions, while recognizing that the bottom-line conclusions will require a quantitative model.

For instance, in the Medicare Part D paper, the authors show visually compelling evidence of bunching around the kink (and show that the location of this bunching moves as the kink moves across years). This evidence allows the authors to reject the null hypothesis that there is no response to incentives, but the descriptive evidence is only partial in the sense that it does not allow the authors to quantify whether the response should be considered “large” or “small” in magnitude.

In the school choice paper, the authors provide evidence that students are “revealed strategic” in how they list schools but are simultaneously “mistakenly strategic” in the sense that they would sometimes have been better off if they had listed schools in a non-strategic way. This indicates that mistakes may be important, but without further modelling assumptions, it cannot fully establish the quantitative importance of these mistakes.

Choosing how much and exactly what descriptive evidence to show is a balancing act. Weak or irrelevant descriptive evidence is a waste of time and can create problematic first impressions. At the same time, some readers may find the basic descriptive evidence more credible than model-based results and you do not want to shortchange these readers. Getting feedback in seminars and conferences is useful for striking the appropriate balance.

3. Use the Descriptive Analysis to Guide Choices of What to Model – and Not Model.

Another key function of the descriptive analysis is to guide and support modeling choices. In the Medicare Part D paper, the authors show that consumers respond to the donut hole before the end of the year, but to a lesser extent than their response at year’s end. These facts motivate the specification of a model where consumers are forward-looking, but potentially not fully so. In the school choice paper, the authors present descriptive evidence that suggests that mistakes are the result of poor information. Based on survey responses, students often hold beliefs that differ substantially from rational expectations admission probabilities: one example, as noted, is that they are on average highly optimistic about their admission probabilities at

second-ranked schools. This motivates the decision to model mistakes as arising from mistaken beliefs, as opposed to another mechanism.

In my view, a signal benefit of a paper that starts with descriptive analysis and then presents the model is that you can use the descriptive evidence to justify what *not* to model. In this way, your modelling choices can be more transparent and less arbitrary, without sacrificing the ability to capture key features of the environment. For instance, in the school choice paper, a natural consideration is whether people engage in strategic information acquisition. That is, acquire more or better information about schools in their consideration set. In the descriptive analysis, the authors do not find that students have better information about the schools they are considering. Rather than falling into the trap of extending the model because of convention or because an extension would be “cool,” the descriptive analysis provides the authors with evidence to justify their decision *not* to model strategic information acquisition – so that they can focus on what matters in their setting.

4. Clearly Articulate the Value-Added of the Model.

As mentioned at the start, I think it is useful to think about the model as offering the reader a tradeoff: If the reader is willing to accept the assumptions embedded in the model, then you can deliver additional and more economically relevant results.

In the Medicare Part D paper, the authors use the limitations of the descriptive analysis to motivate the model. In particular, they describe how the descriptive evidence on bunching allows them to *qualitatively* establish that there is a response to incentives but does not allow them to *quantify* the magnitude of this response. To gauge the economic magnitude of the response, and to gain a deeper understanding of partially forward-looking consumer behavior, they need to know how people would have behaved under a counterfactual linear contract without a donut hole. Because of the dynamic nature of behavior, estimating such counterfactual behavior requires a model.

In the credit card paper, the authors can recover consumers' marginal propensities to borrow in a model-free way using the credit limit discontinuities. However, recovering banks' marginal propensities to lend from time series data is difficult because shifts in banks' cost of funds – which are the result of policy actions by the Central Bank – often occur precisely when the economic environment is rapidly changing.² This motivates their model-based approach, in which they use a small number of sufficient statistics to pin down the lending propensities. They argue that the assumptions underlying this model-based approach – that bank lending responds optimally to changes in the cost of funds and that they can measure the incentives faced by banks – are reasonable in their setting.

The bottom line is that you need to engage in an act of persuasion and sell the model to the reader. To do so, you want to clearly articulate that the value added of the model is high, in that it delivers considerably more insight than the descriptive analysis alone.

5. Choose Parameters of Interest and Counterfactuals That Are Informed by Your Variation.

Having specified and estimated a model, the final part of many papers discusses parameter estimates or conducts counterfactuals. The goal here is to deliver analysis that is more economically relevant than what could have been learned from the descriptive analysis alone – but is still informed by the data. Both of these are important. To get the reader to accept stronger assumptions, you need to be able to offer more economically relevant outcomes. At the same time, the results will be more credible if there is a tight link between the underlying variation presented in the descriptive work and the parameters or counterfactuals delivered by the model.

² For example, there was a large drop in U.S. banks' cost of funds in the fall of 2008, when in response to the financial crisis the policy interest rate of the Federal Reserve (federal funds rate) was set to near-zero. However, this was exactly the period when lenders and borrowers were updating their expectations about the economy, making it hard to separate out the effects of the drop in the cost of funds.

For instance, in the Medicare Part D paper, the main counterfactual is the effect of removing the kink. This comparison is clearly economically relevant: it is the natural benchmark to gauge the effect of the kink and it was a frequently discussed – and eventually implemented – policy reform. Since the descriptive analysis shows bunching, it is closely connected to the variation in the data.

In the school choice paper, the focus of the model is to incorporate inaccurate beliefs – and the resulting mistakes – into a state-of-the-art school choice model. With model-based estimates of inaccurate beliefs in hand, the authors can then examine the effect of a counterfactual strategy-proof mechanism – and examine the effects of correcting beliefs holding the choice mechanism fixed. The counterfactual mechanism with correct beliefs helps quantify the cost of mistaken beliefs that is identified in the descriptive analysis, while the strategy-proof mechanism shows the benefits of a practical solution to the problem of inaccurate information. Indeed, the New Haven schools have now, with the researchers' help, rolled out a version of a strategy-proof mechanism.

The credit card paper uses the model and evidence from the quasi-experiments to recover banks' marginal propensities to lend. The marginal propensities to lend, combined with the directly estimated marginal propensities to borrow, allow the researchers to recover the pass-through of changes to banks' cost of funds. The heterogeneous estimates of banks' marginal propensities to lend are closely connected to the prior descriptive analysis, using the same quasi-experiments that are used in the model-free analysis to estimate consumers' marginal propensities to borrow.

More generally, it's important to emphasize that counterfactuals or discussion of economic parameters shouldn't be an afterthought, completed at the eleventh hour before a presentation or submission deadline. Choosing counterfactuals that provide economically relevant insights that go beyond what you could learn from the descriptive analysis but are still informed by your data – that use but don't abuse your model – requires careful thought and consideration. Don't sell yourself short.

Data-then-Model or Model-then-Data?

For most of this essay, I've taken as given that an applied microeconomics research paper should start with descriptive analysis and then proceed to model-based analysis. However, an obvious meta-question is whether a data-then-model or model-then-data ordering is preferable.

Choosing how to structure a paper can be difficult – and I don't think there is always a right choice. Editors and authors sometimes disagree about the appropriate ordering, and my coauthors and I have sometimes switched the ordering during the course of a project. There are also more complex organizational structures – such as the use of an illustrative toy model, descriptive analysis, and then a richer econometric model – that I will not delve into here. With these caveats in mind, here are some thoughts that can help inform this decision.

It can be preferable to lead with a model when you need a model to guide decisions on what data to collect. Consider a field experiment where you collect your own survey data. For such a project, you would ideally use model-based reasoning to guide your decisions on what questions to ask in your survey. When writing the paper, it may be useful to present the model first to help motivate and justify the survey design.

Similarly, it can make sense to start with the model when the data is non-standard and you need the model to provide guidance on what sort of basic data analysis to conduct. For instance, if you have social network data, it may be hard to summarize the structure of the social network before introducing a model that can help define measures of network structure.

It can also be advisable to start with the model when the conceptual idea imbedded in the model is the main contribution of the paper. For instance, if your paper is proposing a new economic mechanism, then it is natural to first present the model that lays out this mechanism, and then present the data analysis that allows you to quantify its importance.

Conversely, it can be useful to lead with the data analysis if you want to use facts in the data to guide the modeling choices. For instance, in the Medicare Part D paper, the decisions of what to emphasize in the model – price sensitivity, uncertainty, forward-looking but not perfectly forward-looking behavior – are motivated by facts uncovered in the descriptive analysis. Similarly, in the school choice paper, the decision to write down a model with inaccurate beliefs – along with the specific decisions on how to model the wedge between subjective beliefs and rational expectations – would have been hard to motivate without the preceding descriptive analysis.

Another reason to use a data-then-model structure is that it keeps more readers engaged with your paper for longer. For better or for worse, I suspect that readers of applied microeconomics research are more likely to be “turned off” by a model than by descriptive analysis. So, if you lead with your model, you may lose some readers fairly early in the paper; whereas if you start with the descriptive analysis, you’re more likely to retain your readers for at least some of your findings, even if you still lose them when you get to the model section.

A third appeal of the data-then-model ordering is that it is often a better reflection of the research process. Based on experience and conversations with colleagues, my sense is that many applied microeconomics researchers conduct extensive descriptive analysis before undertaking the effort of specifying and estimating a structural model. While papers should not be written as a chronology of the research process, ordering the paper in the same way in which the research was done often comes across as more natural.

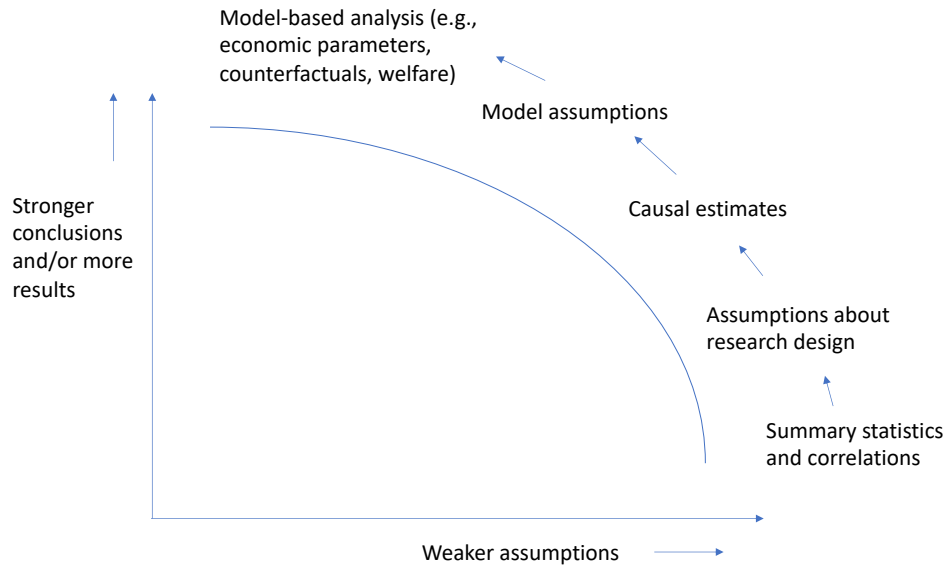
Taking a step back, a metaphor I find useful is the exploration of a decision tree. In constructing a paper, it is smart to lead with the analysis that most quickly and efficiently prunes branches from this tree. If there is an overwhelming number of possible branches of data analysis, it may be more natural to start with the model to guide which branches to explore. If there is a rich set of models that could be plausible, it may be more useful to start with the data analysis to narrow the scope of the modeling exercise.

Concluding Thoughts

It is useful to think about data-then-model papers as tracing out a frontier that trades off the strength of the assumptions for more economically relevant results, as shown in Figure 1. At each stage in the paper, you are offering the reader a deal: if you accept some additional assumptions, then I will provide you with additional results. If the reader is willing to accept assumptions about the validity of the empirical approach, you can offer causal estimates. If the reader is willing to accept additional assumptions about the economic environment, you can deliver additional results in terms of economic parameters, counterfactuals, or welfare.

Economist-readers understand tradeoffs, and my sense is that they will be more likely to accept model-based assumptions if the paper is structured in a way such that they know they are getting something in return. In addition, economists have highly heterogeneous preferences about the kinds of model-based assumptions with which they are comfortable. This type of structure allows the reader to situate themselves at the point on this frontier that best matches their preferences – and allows the reader to “get off the train” at the point where they are no longer comfortable with the tradeoff being offered.

Figure 1: Tracing out the frontier between the strength of the assumptions and more economically relevant results



References

Agarwal, Sumit, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel. "Do Banks Pass through Credit Expansions to Consumers Who want to Borrow?" *The Quarterly Journal of Economics* 133, no. 1 (2018): 129-190.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D." *The Quarterly Journal of Economics* 130, no. 2 (2015): 841-899.

Kapor, Adam J., Christopher A. Neilson, and Seth D. Zimmerman. "Heterogeneous Beliefs and School Choice Mechanisms." *American Economic Review* 110, no. 5 (2020): 1274-1315.