Introduction

At the 1927 meetings of the American Economic Association, Paul Douglas presented a paper entitled "A Theory of Production," which he had coauthored with Charles Cobb. The paper proposed the now familiar Cobb–Douglas function as a general mathematical representation of the relationship between the amounts of capital and labor employed in the US manufacturing sector and the quantity of output produced by that sector. The paper's innovation, however, was not the function itself, as this functional form had been previously proposed by Knut Wicksell and others, but the use of the function as the basis of a statistical procedure for estimating the parameters of this relationship.[1] It is this procedure, a linear regression of the log of a measure of the output of some production activity on the logs of measures of inputs used in the activity, that I call in this book "the Cobb–Douglas regression." In a broader sense, the paper's innovation was the idea motivating and underlying the particular linear regression used by Cobb and Douglas: that a stable, quantifiable relationship between the inputs to and outputs of production processes existed and could be discovered through regression analysis, and that knowledge of this relationship would help to answer important questions of economic theory and policy.

Paul Douglas's innovation of 1927 was a very successful one. The production function regression, in which a measure of output is regressed on measures of inputs, has become a general-purpose empirical tool in economics. Production function regressions are today seen as a means of answering a wide range of questions in a number of fields of economics. They are estimated with cross section, time series, and panel data sets,

[1] Lloyd (2001) discusses previous theoretical uses of the Cobb–Douglas production function to describe the relationship between inputs and outputs.

using variables measured at the level of the individual production process, the firm, the industry, and the national economy. Douglas's idea that such regressions yield meaningful and useful information is widely accepted, and even the original Cobb–Douglas form of the regression continues to be used by economists. This book tells the story of the introduction and diffusion of the Cobb–Douglas regression.

I have two purposes in writing this book. The first is to provide a detailed and accurate historical case study of the introduction and diffusion of an important empirical research technique in economics. Though a history of a single empirical technique may seem very narrow in focus, the life story of the Cobb–Douglas regression has the potential to shed light on several aspects of the history of twentieth-century economics. First, it provides a valuable new perspective on the development of empirical research in economics in the twentieth century, much as did Mary Morgan's (1990) case studies of the estimation of supply and demand curves and of empirical business cycle research in her *History of Econometrics*. Also, it will become clear that the history of the Cobb–Douglas regression intersects with the histories of several research areas and policy debates in economics. I have focused in particular on the role of the Cobb–Douglas regression in the field of agricultural economics and in the history of empirical research into economic growth. However, Douglas's idea that the parameters of something called a production function could be estimated using regression analysis has come to be central to other areas of economics as well, including, since the 1970s, research and policy advice related to education, and, since the 1990s, the empirical analysis of oligopoly behavior and its implications for antitrust law.

As a historical case study, the book focuses on the period from the mid-1920s, when Douglas was doing the preliminary statistical work that led to the 1927 paper, to the late 1960s, when Douglas's idea had come to be widely accepted by economists, and the Constant Elasticity of Substitution (CES) production function was emerging as the first popular alternative to the Cobb–Douglas function for use in empirical production function research. The main source material for this history is the published economic literature. I review and analyze the articles, books, reviews, comments and rejoinders in which the regression was used, praised, criticized, or defended. The decision to focus on published economic literature is warranted, I would argue, because economists in their roles as researchers were the initial and crucial gatekeepers determining whether the Cobb–Douglas regression would indeed come to be a widely used

research tool. The diffusion of the technique could only occur if and when a significant number of economists decided that it produced trustworthy information relevant to questions they believed to be important. Economists' decisions to accept or reject the Cobb–Douglas regression and the results it produced were made largely on the basis of technical considerations related to the specialized disciplinary knowledge of their field, and the published literature is the central source for one who wishes to understand these considerations.

My perspective on this material is in the main that of a historian of econometrics, or more accurately, a historian of the use of statistical evidence by economists. The economists in my narrative who sought to estimate production functions faced a variety of challenges, which in many cases were specific forms of generic challenges that hindered empirical research in economics in the mid-twentieth century. The procedure required measures of inputs and outputs, measures that often had to be constructed from imperfect and incomplete statistical data. Throughout the period I examine, linear regression was the statistical method used to estimate production functions, but the decision to use linear regression necessitated a number of subsidiary decisions, such as the form to be taken by the estimating equation. Theory gave, at best, uncertain guidance to the economist making these decisions; this, combined with the quality of the available data, ensured that any decision made was vulnerable to criticism. Further, the interpretation of the results of the regression required an assessment of problems arising from inaccuracies in the construction of input and output measures, or failures in the statistical assumptions required for linear regression to produce reliable estimates. I give particular attention to describing these challenges, the strategies used by researchers to deal with them, and the roles played by statistical and economic theory in both the development and the criticism of those strategies.

Although Paul Douglas is the most important person in the story of the Cobb–Douglas regression, readers will find that this book is not, even in part, a biography of Douglas. Douglas's 1971 autobiography, *In the Fullness of Time*, reveals him to be a fascinating individual – intelligent, energetic, ambitious, and a complex man who had a highly successful career as an academic economist, culminating in his election as president of the American Economic Association, and in politics, including three terms as a US senator. A good biography of Douglas would be most welcome, but this book is about the life of his most enduring creation, not the life of the creator. I only refer to biographical material on Douglas

when I believe it provides insights into the choices he made in his production function research.

Douglas's 1927 paper with Cobb (published in 1928) marked the beginning of a twenty-year research program in which Douglas, working with various collaborators, estimated the Cobb–Douglas regression using a variety of data sets. Part I of the book describes that research and reactions to it by various groups of critics, both friendly and unfriendly. The research program can be divided into two phases: the first, which included Douglas's 1934 book, *The Theory of Wages*, in which he estimated the regression with annual time-series data from the manufacturing sector as a whole, and the second, from about 1937 to 1947, in which he used a cross-section version of the regression with input and output data aggregated to the level of individual industries.

The original time-series regression of the 1928 Cobb–Douglas paper arose out of Douglas's prior research interests and represented cutting-edge work in empirical economics. Between 1928 and the publication of *The Theory of Wages*, Douglas came to understand and represent his work with the regression as a contribution to a larger effort by sophisticated empirical economists to build a body of empirical knowledge on the "valuable theoretical scaffolding" of "the neoclassical school" (Douglas 1934, xii), or, as he sometimes put it, to make economics "a more progressive science."

From *The Theory of Wages* onward, one finds two consistent and strongly argued claims in Douglas's production studies: The Cobb–Douglas regression could produce results that were economically meaningful and important, and their meaning was closely related to the marginal productivity theory of distribution. In the cross-section studies of the late 1930s and early 1940s, however, Douglas and his coauthors took shifting and sometimes inconsistent positions on a number of subsidiary questions related to these themes. These included questions of inference and statistical procedure, such as "Under what circumstances does a least squares regression based on the Cobb–Douglas production function produce meaningful information?" or "How does one judge whether a particular application of the Cobb–Douglas regression to a data set has produced useful estimates of economically important quantities and relationships?" as well as questions related to the theoretical and practical significance of the results, such as "Could the results be used to test the marginal productivity theory? To detect the presence or absence of competitive conditions in some part of the economy? To provide quantitative measures of key theoretical relationships identified by the marginal productivity theory?"

There are several reasons for this ambiguity. One is that after 1935, Douglas worked with coauthors and allowed them to write their own defenses and interpretations of the Cobb–Douglas regression procedure into jointly authored publications. Another is that numerous innovative statistical practices were being introduced into the economic literature during the period, and Douglas wanted to keep his research program at the cutting edge of empirical economics – indeed, this was one reason for his eagerness to work with coauthors. A third reason was the existence in the profession of more than one conceptualization of the marginal productivity theory. And it could simply be said that Douglas was in the process of learning about the technique he had introduced to the profession in 1927, developing, with the help of his coauthors and critics, a firmer grasp on its potential and its limitations.

Douglas began his first successful campaign for the US Senate in 1948, effectively ending his research career. However, applications of Douglas's regression multiplied over the following two decades – a 1963 survey article in a leading economics journal cited more than thirty published books and articles in which estimated Cobb–Douglas regressions could be found (Walters 1963a). Criticisms of the procedure also continued to appear, some making new points and some essentially reiterating points made by earlier critics.

Part II of the book is concerned with the diffusion of the Cobb–Douglas regression, and begins with a chapter that discusses three important moments in the story of the Cobb–Douglas regression after 1947. The first is the appearance in Lawrence Klein's (1953) econometrics textbook of an econometric model that included a Cobb–Douglas production function. The second is E. H. Phelps Brown's 1957 article criticizing Douglas's applications of the regression. The third is the introduction of the CES production function in 1961 by Kenneth Arrow, Hollis Chenery, Bagicha Minhas, and Robert Solow.

However, the bulk of my discussion of the postwar (and post-Douglas) spread of Douglas's regression takes the form of two case studies: (i) the adoption of the Cobb–Douglas regression by agricultural economists and (ii) the story of how the Cobb–Douglas regression came to be used as a tool for the measurement and explanation of economic growth and technical change. In both cases, I begin by describing the preexisting situation in the area of research into which the regression was introduced. Each had its own set of motivating questions, theoretical framework(s), cache of relevant data, and associated empirical challenges. In both cases, a small group of researchers introduced the Cobb–Douglas regression as a means

of answering one or more of the important questions in the field. As these innovators described how the regression could be applied to the questions they cared about, some of the existing claims and counterclaims about what Douglas's versions of the regression could or could not accomplish, and under what circumstances, became irrelevant, while the new applications led to new challenges to, and new modifications of, Douglas's regression procedure. In each of the two fields, the regression, and Douglas's broader idea about the fruitfulness of using regression analysis to estimate input-output relationships, gained a foothold, creating yet another genus of the Cobb–Douglas regression (with associated species) that sidestepped at least some of the criticisms leveled at Douglas's original studies. In the process, each field also became home to a new constituency with a vested interest in defending the value and legitimacy of its version of the regression.

My second purpose in writing this book is to explore possible answers to two questions raised by the story of the Cobb–Douglas regression: First, what explains the success of the Cobb–Douglas regression, that is, what factors contributed to its gradual transformation from an innovative and controversial statistical procedure to a widely accepted general-purpose tool in empirical economics? As a historical matter, the acceptance of the Cobb–Douglas regression was the crucial first step in the process by which economists would embrace and more fully implement Douglas's innovative idea of using regression analysis to estimate production relations. Further, the factors that contributed to economists' acceptance of the Cobb–Douglas regression shaped in important ways the more general strategies and techniques that economists would eventually develop in pursuing the many research questions embodied in Douglas's idea. So, explaining the success of the Cobb–Douglas regression goes a long way towards explaining the success of the regression approach to estimating production relations.

The second and related question is whether the analysis of factors that contributed to the success of the Cobb–Douglas regression offers more general insights into the factors that determine the success or failure of innovative research techniques in the social and perhaps natural sciences. In posing this second question, I have in mind a variety of types of innovative research tools, techniques, or procedures: new statistical techniques, like the Cobb–Douglas regression, but also new experimental methods, new modeling strategies, new conceptual frameworks, and so on. I regard an innovation as having become "successful" when it comes to be used by a large number of researchers, and a significant number of

researchers who do not actually use the technique believe that it can contribute to the advance of knowledge, and regard as credible the knowledge it has produced. This broad consensus on the value of the tool should persist for a while – decades, not years – and would be evidenced by a steady and substantial stream of peer reviewed or otherwise professionally approved published research in which the tool is used.[2] Few economists would dispute that the production function regression was, and remains, a success in this sense.[3]

I reflect on the first question, that is, what factors led to the success of the Cobb–Douglas regression, at various points in the book. Some of these reflections are brief and quite speculative, but there are four potential reasons for the success of Douglas's regression that I elaborate in more detail in a final chapter. A first reason is the decision by Douglas to link his procedure, and his interpretation of the estimates it produced, to fundamental concepts of the neoclassical approach to economics, which was destined to grow in influence over the course of the twentieth century. A second is what I call the adaptability of the Cobb–Douglas regression. As I have mentioned, one could discern from Douglas's papers several "versions" of the Cobb–Douglas regression. There were the two "actual" versions (the time series and the industry-level cross section), but Douglas endorsed the idea that other versions should be developed as well (firm level, or time series within a single industry, or cross section with aggregate data from different states or countries, or any of the above with alternative functional form assumptions). As the Cobb–Douglas regression came to be adopted by other researchers in other fields of economics, for purposes other than those originally envisioned by Douglas, even more versions

[2] Using this standard of success, it is not hard to also find examples of unsuccessful attempts to introduce innovative research tools. Three cases of innovative empirical tools/approaches introduced into twentieth-century economics research that were not successful by my definition would be the measures of business cycle behavior developed by Arthur Burns and Wesley Mitchell (1946, discussed in Morgan [1990]); Ragnar Frisch's confluence analysis (Hendry and Morgan 1989); and "periodogram analysis" (Cargill 1974).

[3] Some citation-based evidence of this success can be offered as well. A 2003 article proposing an improved approach to production function estimation, "Estimating Production Functions Using Inputs to Control for Unobservables" (Levinsohn and Petrin 2003) has been cited over 1,000 times. A subsequent related article, "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables" (Wooldridge 2009) has been cited 160 times. An empirical article demonstrating a "production function approach" to measuring the effect of health on economic growth (Bloom, Canning, and Sevilla 2003) has been cited over 250 times. By way of comparison, the median article published in one of the five most prestigious journals in economics in 2007 had received 50 cites by 2015 (Hamermesh 2018).

were born. I believe that the fact that many versions and interpretations of Douglas's regression existed, both actual and potential, contributed to the growth in its use by economists. Criticisms that were devastating to one version of the regression were often irrelevant to others. Rejection of the results that the Douglas group was actually publishing did not constitute rejection of the broader idea represented by Douglas's program: One could still believe that if one used the right type of data, rather than the type that Douglas had used, a regression of the log of output on logged measures of inputs could yield useful information, although perhaps not the information that Douglas claimed it would yield.

Another explanation of the success of Douglas's innovation involves a discussion of Douglas's rhetoric of persuasion. It will become clear that Douglas was, at times, the model of the objective scientist: open and frank about the limitations of his method, about the weaknesses of his data, about the potential biases in the results he had produced, and about the wide gap between the conditions his regression procedure assumed and the conditions under which the data he used were actually generated. But at other times, he adopted the tone and rhetoric of an advocate. When responding to critics, their weaker criticisms were identified and rebutted but strong ones ignored or twisted into a form more easy to dismiss. When Douglas summarized his past work, he often stressed or exaggerated positive results and glossed over or ignored negative ones. Even particular results that he identified as problematic when originally published would sometimes be represented in subsequent accounts as part of the body of evidence that supported the legitimacy of the procedure.

More important, perhaps, than Douglas's ability to persuade was his audience's willingness to be persuaded. Simply put, economists really wanted to believe in Douglas's idea, if not his regression. Douglas's production function research faced harsh, detailed, and technically competent criticisms from the very beginning, some of which called into question not just Douglas's particular applications of his regression but also the general feasibility of estimating neoclassical production functions. At the same time, early articles that use or endorse the future potential of Douglas's regression seem to show an *eagerness to believe*. My argument is not that there was simply some critical mass of economists who were so carried away by the potential of the Douglas's method that they were willing to ignore criticisms or accept deficient defenses of the method. Instead, I think it is important that many economists were sufficiently excited by the potential of the method that, rather than abandoning it in the face of warranted criticism, they were willing to continue to teach it and use it

while working towards improvements in data and statistical methodology that would rescue it.

To put this argument in a more concrete form, I focus on two economists, Earl Heady and Zvi Griliches, who responded to the new research tool in a similar way. (Heady's work with the regression is described in detail in Chapter 5, and Griliches's in Chapter 6). Both men became early and frequent users of the technique, in research that displayed a solid understanding of how it worked as a research tool as well as the various problems associated with its application. More importantly, they communicated, by words and example, the attitude that despite its weaknesses, the technique was potentially very valuable, that the best way to realize that potential was to continue using the technique while working to address the weaknesses, and that even while this process of improvement was going on, the technique was still able to contribute to knowledge. I propose that the efforts of these men, or more precisely, the manner in which they advocated for the new technique, was an important factor in its eventual success.