1

Introduction: Is Macroeconomics Useful?

Microeconomics is what economists know about. But macroeconomics is what they want *to know about. That's what makes it so interesting.*

Benjamin Friedman (quoted in Fisher, 1993)

In 1936, in the middle of the worst economic outcome the modern world had ever seen, John Maynard Keynes published a book entitled *The General Theory of Employment, Interest and Money*. The purpose of the book was to try to explain how a market economy could find itself persistently stuck in a situation where, despite ample industrial capacity and willing workers, unemployment and deprivation were widespread. Pretty much no one reads Keynes anymore (especially economists), but the *General Theory* did leave one enduring legacy: It gave rise to macroeconomics as a separate branch of economics, in which the units of study are economic aggregates – things like the overall levels of production, employment, and inflation in an economy.

One other piece of the Keynesian legacy was not so enduring. This was the implicit assumption that in order to explain the behavior of macroeconomic aggregates, it is necessary to use a method of analysis that is distinct from the microeconomic strategy of analyzing the world in terms of optimizing households and firms whose actions are (somehow) coordinated so as to yield an equilibrium in which supply and demand are equal. That disconnect between macroeconomics and microeconomics was viewed by some economists (not all) as a bit of an embarrassment, but it was usually justified on the grounds that pathologies like business cycles and involuntary unemployment simply didn't seem like "equilibrium" phenomena in the microeconomic sense. But, in 1979, Lucas and Sargent published a broadside that argued this latter feature of macroeconomics was wrongheaded

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

and – as far as informing economic policy went – counterproductive.[1] Lucas and Sargent maintained that macroeconomic phenomena *were* equilibrium phenomena and could be studied in a Walrasian general equilibrium framework like the one developed by Arrow, Debreu, and others. Today, variants of the Lucas–Sargent approach permeate mainstream macroeconomic theory, including "new-Keynesian" economics and other attempts to construct optimizing, equilibrium models that can produce "Keynesian" results.

In this chapter, I discuss reasons why the existing general equilibrium framework of microeconomics isn't suitable for analyzing macroeconomic questions (or microeconomic questions, for that matter). I also consider the separate question of whether and to what extent it makes sense to take an aggregative approach to studying the economy, and the difficulties we face in choosing to do so.

1.1 Existence versus Stability (and Why the Latter Matters More)

It is one of those ironies of history that almost at the same time that Lucas and Sargent were laying out their case to macroeconomists, *micro*economic theorists were starting to have serious doubts about the predictive content and overall usefulness of general equilibrium theory. The theory's predictive content had been called into question by a set of papers published between 1973 and 1976 that demonstrated that the assumptions needed to secure the existence of a Walrasian general equilibrium were not enough to tie down the aggregate properties of the economy in any meaningful way.[2] In particular, stronger – and therefore special – assumptions were needed to demonstrate that an equilibrium was unique; moreover, no general comparative statics results could be achieved (it wasn't even possible to unambiguously demonstrate that increasing the amount of a good in an economy would lower its price).[3]

[1] Like most pieces of agitprop, the Lucas and Sargent paper vastly overstated the deficiencies of the old order. And, as in most revolutionary movements, counterrevolutionary activity was dealt with harshly – witness Lucas's (1994) choleric reaction to Ball and Mankiw (1994).

[2] The papers were Sonnenschein (1972, 1973), Mantel (1974, 1976), and Debreu (1974). In a nutshell, what they showed was that summing individual demands led to an aggregate (excess) demand relation that was continuous, didn't exhibit money illusion, and obeyed Walras's law – but that was it. Conversely, any function that had these properties could actually occur as the excess demand function of an economy.

[3] More recent attempts to salvage general equilibrium theory from the Sonnenschein–Mantel–Debreu wreckage have met with little success; see Brown and Matzkin (1996), Brown and Shannon (2000), and Nachbar (2002) for some representative examples.

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

However, there was an even more problematic issue with general equilibrium theory that was starting to become apparent around this time: While the theory could prove that an equilibrium existed, it couldn't show that it would be *stable*, in the sense that mechanisms were present that would return the economy to equilibrium if it happened to be moved away from it, or that would bring the economy to an equilibrium (not necessarily Walrasian) if it didn't start off there. For a macroeconomist, stability seems far more relevant than existence: A basic question in macroeconomics is whether the economy will recover "on its own" after a recession or whether any self-correcting tendencies are too weak to be relied on (or too unreliable to depend on). Without a demonstration of stability and a theory of the out-of-equilibrium processes that deliver it, that question can't be answered with any generality. Relatedly, a belief that the economy will eventually return to a state of full employment as long as prices and wages are given enough time to adjust receives no justification from the existence proofs of general equilibrium theory.

The stability problem has not been solved – and probably never will be – though that fact no longer seems to vex microeconomists overly much. Perhaps the best run at the problem in the context of standard general equilibrium theory was made forty years ago by Fisher (1983), who didn't even find his own (partial) solution all that convincing. The technical problem involves finding a realistic trading process that will act as a Lyapunov function (intuitively, a Lyapunov function is a function that "squeezes" the state of a system toward its equilibrium point; it is used in mathematics to demonstrate the stability of differential equation systems). One reasonably realistic candidate might be a process in which the set of trading opportunities that are perceived as profitable at disequilibrium prices becomes smaller and smaller over time as agents arbitrage them away. The reason Fisher concluded the problem is likely insoluble is that the ability of real-world agents to act on new perceived opportunities for arbitrage – including those that turn out to be incorrect – makes stability impossible to demonstrate without additional strong (and unrealistic) assumptions.

Despite its truly fundamental importance, equilibrium stability receives virtually no attention from microtheorists nowadays. As evidence, Mas-Colell, Whinston, and Green's (1995) exhaustive survey of microeconomic theory devotes only seven of its 980-plus pages to the question of stability (and that mostly to the uninteresting and largely irrelevant concept of so-called tâtonnement stability), but ultimately asserts that the topic is not central to economics (p. 620). That claim is hard to take seriously, though: If you can't show that an economy will converge to a particular

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

equilibrium – and do so relatively rapidly – then it's difficult to argue that such an equilibrium has any real-world relevance, or any claim to priority as an object of study. For instance, comparative statics exercises seem rather pointless if one can't argue that the economy will actually tend to move to its new equilibrium point; likewise, concepts like Friedman's (1968) "natural rate" of unemployment – which is explicitly associated with a Walrasian equilibrium – have no meaningful content.

As Fisher (1983, 2011) also points out, various interesting complications arise once we take the stability problem seriously. Many of these are also extremely relevant to macroeconomics.

- If households and firms find themselves in an economy that is away from equilibrium, they will see that prices can change and will also likely realize that their plans to buy and sell might not come to fruition. Perceived constraints on sales in product and labor markets can, in turn, give rise to Keynesian-style (and decidedly non-Walrasian) underemployment equilibria.[4]
- A claim that the efficiency properties of competitive equilibrium can be enjoyed by instituting "market reforms" is specious, since there's no guarantee that such an equilibrium will actually result (and other bad things could happen along the way).
- Once trading out of equilibrium occurs, phenomena such as path dependence and hysteresis can easily arise. Consider the Edgeworth box diagram in Figure 1.1, which shows a two-person exchange economy with a unique Walrasian equilibrium at point A (given by the intersection of the offer curves ω_1 and ω_2) that is supported by prices P_A from the initial endowment E. Say that prices instead start out at P_B. Points along P_B are not Walrasian equilibria (there is excess demand for the x-axis good and excess supply of the y-axis good). However, mutually beneficial trades are possible, and if they are allowed to do so at the disequilibrium price, then any subsequent equilibrium can easily differ from A – even (especially) if relative prices are restored to P_A.
- Finally, the assumption that markets clear or that any disequilibrium states are resolved quickly enough that the economy will always be at a rational expectations equilibrium is simply that – an assumption.

[4] Interestingly, in situations like these Walras's law will only hold in an expectational form – see Fisher (1983, chapter 7).

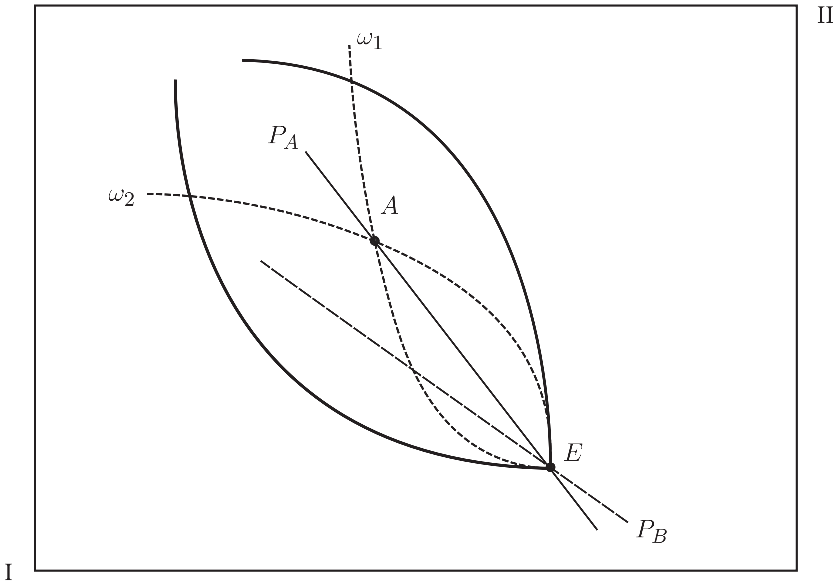

Figure 1.1 Path dependence under disequilibrium trading.

1.2 Microfoundations and Aggregation

If Lucas and Sargent's call to model macroeconomic phenomena using the tools of general equilibrium theory has so little to recommend it, what about their prescription that macroeconomic theory must be grounded in the optimizing behavior of households and firms? Here, again, we quickly run into trouble. Since the 1950s, economists have become increasingly aware that there is no reason to expect that individual-level behavior will show up in any recognizable form in the aggregate unless additional highly restrictive assumptions are made.[5] Worse still, even the simple *time-series* properties of individual-level variables are unlikely to be preserved once the data are aggregated up. And while there do exist ways to deal with this problem – some good, most bad – even the sounder methods are not especially useful from a practical perspective.

Start with one of the bad solutions: the assumption that the overall economy can be described in terms of an "average" household or firm, a so-called *representative agent*. Even today, a remarkable number of macroeconomists

[5] Even without being acquainted with this earlier literature, the Sonnenschein–Mantel–Debreu results should give us a hint that such troubles are lurking in the background.

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

use representative-agent models for both theoretical and empirical research despite there being absolutely no reasonable justification for doing so (a fact that has been well-known for years). As Kirman (1992) details:

- There is no way to formally justify that a combination of individual maximizing agents will itself act like a single maximizer – or, put more starkly, "[t]here is simply no direct relation between individual and collective behavior" (p. 118).
- Using a representative-agent framework to model the effect of a policy change is apt to be misleading: Once a policy setting (or some other feature of the model) is altered, the representative agent can end up responding in a manner that is different to what we would find by aggregating each individual's response to the policy (or other) change. Hence, one thing that makes microfoundations useful in principle – namely, the idea that by explicitly grounding the behavioral responses of the economy's members in optimizing behavior that reflects their "tastes and technology," we can obtain a structural model that is suitable for policy evaluation – is actually absent from representative-agent models.
- In the case of a representative consumer, the "preferences" of the representative agent need not match the preferences of the individuals in the economy, even if each individual happens to prefer the same thing. That is, situations can arise in which the representative agent prefers A to B, even though every individual prefers B to A.

Similarly, on the production side, the conditions needed in order to talk about a "representative firm" (in the sense of an aggregate production relation) are sufficiently stringent that they are unlikely to hold for any economy at any point in time (Fisher, 1993), a point that we will discuss at length in Chapter 4.

In a fundamental sense, it is strange that a macroeconomist would even *want* to entertain the notion of using a representative- or single-agent approach for modelling or estimation purposes. Many of the most interesting decisions in the economy – for example, those pertaining to the choice of how much to produce, hire, consume, and invest – likely reflect behavioral and informational feedbacks that are generated by the myriad interactions that take place among individual agents and groups. As a result, we might expect at least some macroeconomic phenomena to be *emergent*, manifesting behaviors and dynamics that, because they are the result of these interactions, cannot be predicted simply by scaling up or extrapolating the behavior of a single (isolated) agent. Put more simply,

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

we probably aren't going to be able to capture the dynamics of groups of agents very well if we restrict ourselves to considering a single representative agent. And studying the dynamics of groups of agents is pretty much what macroeconomics is all about.

Why, then, does the representative-agent approach continue to pervade macroeconomics? Very likely, some simply don't realize how shaky the justification for this framework actually is, while some others probably don't much care. I also expect that very few people have ever had a paper rejected from a mainstream journal for assuming a representative agent – just as how in the 1970s, no procurement manager ever lost their job for buying IBM. Additionally, these frameworks are very tractable relative to most alternatives. That said, someone of a critical mindset can claim with strong justification that in essentially every context, the predictions that are derived from representative-agent models are at best meaningless and at worst completely misleading. And at a minimum, such models have no claim to being "microfounded" in any serious way.[6]

In recent years, there has been an attempt to move away from a strict representative-agent approach by allowing for some form of heterogeneity among agents (typically consumers). For example, the so-called heterogeneous agent new-Keynesian models try to address a shortcoming associated with the transmission mechanism of monetary policy in standard new-Keynesian models. In the standard model, monetary policy affects the real economy through an intertemporal substitution channel: Changes in real interest rates cause a representative household to either postpone its current consumption or bring forward its future consumption. The problem is that empirical estimates of the sensitivity of consumption to interest rate changes find it to be quite low, which implies that this channel can't really be consistent with monetary policy having a large effect on aggregate demand. The proposed solution involves introducing what amounts to a Keynesian multiplier mechanism by adding households whose consumption is closely tied to their current income; as a result, even a modest response of spending to a change in the return on liquid assets can yield sizeable overall effects on consumption.[7]

[6] As a colleague of mine once put it, "I have no problem with microfoundations. But in what way is $\max E_t \sum_{i=0}^{T} \beta^i U(c_{t+i})$ microfounded?"

[7] See Kaplan, Moll, and Violante (2018) for a canonical example. This description omits a number of subtleties associated with these models; we will consider them – along with their shortcomings – in greater detail in Chapter 7.

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

A second example is provided by Angeletos and Lian (2022), which attempts to explain why shifts in aggregate demand can drive business cycles. In an economy with a representative consumer (and full information), a shock to aggregate demand should not have a "multiplier" effect – that is, an effect on output larger than what is implied by the shock itself – because permanent income is unchanged. In order to generate a larger effect, Angeletos and Lian assume that consumers are imperfectly informed about the state of the aggregate economy, and so misperceive an aggregate demand shock as an idiosyncratic shock to their (permanent) income (the formal setup uses the sort of "islands" economy that Lucas, 1972 employed in order to ensure that individual misperceptions didn't cancel out in the aggregate). The resulting reduction in consumption and aggregate demand induces additional pessimism about permanent income, a further decline in demand, and so on, thereby amplifying the effect of the initial shock.

Although these models do incorporate heterogeneity of a sort, the rationale for doing so has nothing to do with the fact that representative-agent models are theoretically suspect, but rather reflects an attempt to deal with the problems that arise when consumer behavior is modelled using a permanent-income framework with rational expectations. Put differently, approaches like these do highlight the fragility of conclusions derived from representative-agent frameworks, and if one insists on modelling household behavior as the outcome of an intertemporal optimization problem, then modifications like these will be needed in order to obtain halfway plausible results. But these analyses still work within a market-clearing framework where economic fluctuations are essentially equilibrium phenomena. And they still assume agents who solve complicated optimization and information-acquisition problems that are not all that likely to provide realistic descriptions of individual behavior (at least, they have never been shown to do so).[8]

While using representative agents to model the behavior of economic aggregates has no justification and shouldn't be done, we unfortunately don't have workable alternatives that *can* be usefully employed in a practical setting. On the theoretical side, at least two alternatives have been proposed; although both are interesting and explicitly model interactions among heterogeneous agents, neither has found much applicability to policy analysis.

[8] We will return to some of these themes in Chapter 3. It also isn't clear that a model that relies on such a contrived setup as an islands economy can seriously be labelled "microfounded" – as Solow (1983) once pointed out, no one has ever discovered such an island (or even found a message in a bottle that came from one). Much the same criticism can be levelled against models that rely on contrivances like Dixit–Stiglitz aggregation.

The first alternative, *agent-based modelling*, constructs model economies with large numbers of heterogenous agents who follow specified behavioral rules (including rules governing how their expectations are formulated) and who interact with each other in different ways and in different venues, such as making transactions in different markets.[9] These model economies are simulated with a computer, and these simulations can generate complex aggregate dynamics despite assuming relatively simple individual-level decision rules.

What has so far limited the applicability of this approach in a policy setting is the difficulty in demonstrating that the models' predictions are robust and that the models themselves faithfully capture some feature of real-world economic dynamics. For the former, there is always a sneaking suspicion that a particular result depends on the specific decision rules agents are assumed to follow (and apparently this is difficult to check in large models – see Dawid and Gatti, 2018, p. 70). For the latter, the fact that a model can generate business cycles or financial crises, or roughly match the dynamics of (or correlations between) macroeconomic aggregates like real GDP or inflation might not be viewed as sufficient validation of the model (though in fairness, it's not clear that any other type of theoretical model is tested more rigorously). Similarly, the fact that these models' principal strength is that they can generate emergent dynamics means almost by definition that it will be difficult to describe and assess the causal mechanisms that are at play in a particular model, and especially how they depend on the way that individual-level behavior and interactions are modelled.

A second theoretical alternative is to take a page from physics – specifically, statistical physics – and model aggregate phenomena in terms of the statistical distributions of outcomes that are generated by large numbers of interacting agents. This approach, which is developed at length by Aoki (1996, 2002) and Aoki and Yoshikawa (2007), is intriguing – there is an intuitive appeal in thinking about a macroeconomic equilibrium as an inherently stochastic object that emerges from the bottom up, rather than as the result of feeding stochastic shocks into a system of (typically linear) equations that supposedly capture the average responses of different classes of agents.[10] However, the methodology has not really caught on. One likely

[9] An early example of this sort of approach in a macroeconomic context is described in chapter 9 of Nelson and Winter (1982); see Dawid and Gatti (2018) for a relatively up-to-date overview of some macroeconomic applications of agent-based modelling.

[10] In a preface to Aoki and Yoshikawa (2007), Yoshikawa draws a connection between this approach and Tobin's (1972) conception of a "stochastic macro-equilibrium" in which "random intersectoral shocks keep individual [markets] in diverse states of disequilibrium"

reason is that the mathematics used are unfamiliar to most economists (after all, physics envy doesn't imply physics training). On a deeper level, though, many of the theoretical models derived using this approach are closer to proofs of concept – that is, demonstrations that one can get models of this sort to generate phenomena such as endogenous cyclical movements in aggregate production. Such exercises seem reminiscent of Slutzky's (1937) famous result that moving sums of random variables can yield time series whose fluctuations *look* like business cycles – interesting and suggestive, but hard to know what to make of.[11]

On the empirical side, an approach to modelling the behavior of aggregates in a way that tries to correctly capture the effects of individual heterogeneity involves starting from models that are fit to microlevel data. For example, if we have panel data on household income, expenditures, and other characteristics, we can fit a consumer demand system and allow its parameters to depend on observed household attributes. The resulting estimates can be summed or (equivalently) averaged in order to obtain "correct" aggregate relations. Generally, these relations will be different from what we would get by evaluating the demand system using the aggregates themselves (for example, average income or the fraction of households with a particular attribute): Intuitively, any sort of nonlinearity in the specification, including interactions between household attributes and variables like household income, will drive wedges between the various relations.[12]

In practical terms, the existence of these sorts of issues implies that we will not be able to recover individual-level behavioral relationships using aggregate data; similarly, we should not expect the restrictions on individual behavior that are implied by microtheory to be applicable to an aggregate model, or to be apparent in aggregate data. In fact, it turns out that even the *time-series* properties of individual-level data will not generally carry

but "the perpetual flux of particular markets produces fairly definite aggregate outcomes." (The approach is also very appealing to those whose vision of macro theory was shaped by Asimov, 1951.)

11 A related literature that goes by the name "econophysics" has achieved somewhat wider acceptance. However, most of this work has been focused on describing and understanding the behavior of financial markets – not the broader economy – using tools derived from statistical physics.

12 To give two trivial examples, the log of an average is not the same as the average of a log, while the presence of a nonzero covariance between a characteristic x and income y means that the average of xy will not equal the average of x times the average of y. See Stoker (1993) and Blundell and Stoker (2007) for useful summaries of the issues involved (the latter reference also includes applications to intertemporal consumption modelling and to models of labor supply).

through to the aggregate level – for example, if household-level behavior implies that variables like consumption and wealth will be cointegrated for each household, it need not be the case that the aggregate analogs of these series will be cointegrated as well.[13] Finally, it seems likely that equations fit to aggregate data will tend to be unstable if the composition of the population changes over time, which is one reason (among many) why we should not expect empirical macroeconomic relationships to work very well over long periods of time.[14]

Being unable to recover individual-level behavior from aggregate data is only an issue, of course, to the extent that we *want* to recover such relations; if we do, then we should probably be looking to microlevel data in the first place. (That said, enough problems attend the use of microlevel data when trying to estimate behavioral relationships that such exercises are often less concerned with theoretical purity than they are with obtaining well-fitting and tractable specifications, especially where capturing the effects of heterogeneity is concerned.) But it is also the case that in nearly every practically relevant application, it will simply not be feasible to correctly aggregate empirical equations that are based on microlevel data: In the United States, the main source of household-level consumption data, the Consumer Expenditure Survey, is not really intended to be used to fit demand systems (instead, its main purpose is to compute weights for the consumer price index); moreover, its data are only available with a considerable lag. There is also no source of readily available, firm-level producer data in the United States that would be suitable for estimating microlevel production relationships.

Where does all this leave us in terms of the sorts of aggregate measures – real GDP and its components, their corresponding price measures, and so on – that are actually produced by US statistical agencies?

In general, the conditions that need to be met in order to ensure that these aggregates will summarize microlevel production and demand relationships in a sensible and well-behaved way turn out to be so stringent that they are almost certainly never met by any real-world economy – and statistical agencies do not approach the problem in this way. Instead, these aggregate series are defined so that they will have certain desirable and intuitively reasonable

[13] See Forni and Lippi (1997) for a detailed discussion of this topic.

[14] In addition, we will typically not be able to use estimates from microlevel empirical work in the context of an aggregate model – for example, as a way to "calibrate" a macromodel's parameters (see Browning, Hansen, and Heckman, 1999, for an early discussion of this topic).

"index number" properties.[15] To give one example, price and quantity indexes for total consumption use a formula that ensures that nominal consumption – which, being a value denominated in dollars, is reasonably straightforward to define and compute – will equal the quantity index when it is divided by the price index, and vice versa. While measures like these can be given a tenuous grounding in choice theory, such a grounding will only apply to a single consumer (or a representative agent); once we are dealing with the economy as a whole, any such theoretical justification vanishes.

Similarly, a measure of aggregate production like real GDP shouldn't be interpreted as the output of some economy-spanning firm that uses a neoclassical production technology to make a single homogenous good. Instead, real GDP is better thought of as an index that starts with the change in the total dollar value of the goods and services produced in the United States for final demand and then tries as best it can to remove the portion of this change that is attributable to changes in the individual prices of these goods and services.[16] The resulting quantity index (we hope) provides a reasonable gauge of the change in the overall level of real activity across different points in time that is both useful in its own right and that can be usefully related to other aggregates, such as the unemployment rate or economywide employment. Likewise, aggregate price indexes should allow us to make useful statements regarding the broad direction and magnitude of economywide price changes.

From a purely statistical perspective, using aggregate data can carry one significant advantage: Under certain circumstances, aggregating individual observations will help to reveal common "macrolevel" influences. We can see how this might occur with the following extremely stylized example (which is taken from Forni and Lippi, 1997, chapter 1). Assume that an individual-level variable, x_t^i, is the sum of two components:

$$x_t^i = X_t^i + \xi_t^i. \tag{1.1}$$

Here, X_t^i denotes the effect of macroeconomic shocks, while ξ_t^i is an individual-specific term. (For instance, if x_t^i is an individual's income, X_t^i would reflect the dependence of their income on economywide conditions, while ξ_t^i could be something like a pay raise or a bonus that the individual receives in period t.) These individual-specific shocks are orthogonal across individuals, as well as being orthogonal to any of the "macro" terms.

[15] We will consider some of these topics in more detail in Appendix B.

[16] For the sticklers, here "final demand" is meant to include inventory investment.

Purely for illustrative purposes (and to obtain a simple expression), let's assume that for any two individuals i and j, their respective macro terms X_t^i and X_t^j have a constant correlation equal to ρ (the fact that these terms would have *some* correlation isn't too hard to believe – they are related to economywide shocks, after all – though the assumption that this correlation is the same for everyone at every point in time is a rather special one). Let's also normalize things so that the variance of X_t^i is the same for everyone (and equal to 1), and that the macro term explains a fraction R^2 of the variability of x_t^i, which we accomplish by setting the variance of ξ_t^i equal to $(1-R^2)/R^2$.

Now let's define an aggregate variable $\mathbf{X}_t$ (say, total income) as the sum of the individual (income) terms x_t^i:

$$\mathbf{X}_t = \sum_{i=1}^{n} x_t^i = \sum_{i=1}^{n} (X_t^i + \xi_t^i), \tag{1.2}$$

where we assume there are n individuals. Under the various assumptions, the variance of the sum of the macro terms will be $n+n(n-1)\rho$ and the variance of the sum of the individual-level terms will be $n(1-R^2)/R^2$. Hence, in a sample of n individuals, the fraction of the variance of the *aggregate* variable $\mathbf{X}_t$ that will be explained by the macro term (call this R_n^2) will be

$$R_n^2 = \frac{1+(n-1)\rho}{(n-1)\rho + (1/R^2)}. \tag{1.3}$$

What this means is that even if the macro term explains very little of the variability of the individual variables x_t^i – say, $R^2 = 0.01$ – and even if the macro effects are not very correlated across individuals (say $\rho = 0.01$), we won't need to aggregate over too many individuals in order to have the macro term explain a reasonably large fraction of the aggregate variable (for this example, 10,000 individuals would yield an R_n^2 equal to 0.50).

It's important not to make too much of this result – for various reasons, including the unrealistic nature of the example, things are unlikely to be quite this neat in the real world.[17] And, of course, none of this solves the basic problem that aggregate variables will not typically behave as individual-level variables writ large. But the example does hold out some hope that an atheoretical, basically empirical approach might end up

[17] In particular, the effective population size n is not always large, as many aggregate statistics are based on surveys with limited sample sizes. In addition, sampling variability can interact with some commonly used index number formulas in a way that prevents it from washing out at the aggregate level.

capturing informative common movements in the aggregate data that we can associate with macroeconomic shocks.[18]

1.3 Toward a Practical Macroeconomics

The issues raised in this chapter have two implications for the role that microeconomics should play in macroeconomics, neither of which is very constructive.

- First, the lack of any convincing theoretical demonstration of the stability of a Walrasian general equilibrium implies that a market-clearing general equilibrium is neither a relevant nor an interesting object of study, especially when we want to consider dynamic responses of the economy (such as those that occur over the course of a business cycle).
- Second, it is extremely unlikely that the aggregate data that we actually have access to will reflect recognizable theoretical or empirical microeconomic relations. As a result, the type of microfounded model that dominates mainstream macroeconomic thinking will provide no useful predictions about macroeconomic processes or outcomes, and no useful guidance regarding what sorts of empirical macroeconomic relationships are likely to be well-specified or stable over time.

Put more plainly, there is no especially good reason to use microeconomic theory to explain or predict the changes in economic aggregates that we actually observe, or even as a framework for modelling macroeconomic phenomena: Mainstream microeconomic theory simply isn't up to the task, even if we are willing to suspend disbelief and entertain the notion that people's behavior in the economic realm can be well described with the tools of that theory.

[18] This notion would seem to be inconsistent with Gabaix (2011), who argues that the distribution of firms is so fat-tailed that idiosyncratic productivity shocks among a small number of large firms account for a large fraction of macroeconomic fluctuations. However, the productivity shocks Gabaix measures are actually sales shocks, and many of the specific examples he gives reflect changes in demand, not firms' ability to produce more or less efficiently. And even for big firms, demand (and therefore sales) will be determined by a large number of agents, some of whom will be in other countries. In addition, using (net) sales ignores inventory investment, which – as we'll see in Chapter 3 – appears to be a major contributor to business cycles and which is (if not imported) part of *some* firm's output. (Also of note is that Gabaix's key theoretical derivation, which he uses to justify his empirical approach, assumes competitive conditions even though his focus is on extremely large firms in what are no doubt highly concentrated industries.)

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

What, then, might be a practical alternative? The empirical approach that we will consider in much of this book involves treating macroeconomic aggregates *on their own terms*, in the hopes of discerning relationships among these variables – essentially, statistical regularities or stylized facts – that are reasonably robust and well-specified in a time-series sense. We will then attempt to come up with plausible interpretations for these empirical relationships (as well as assessing how well mainstream macroeconomic theory is able to do so).

There are pitfalls, of course, in taking such an approach. When you stop to think about it, it seems hard to believe that *any* sort of empirical macroeconomic regularity would exist in an economy as geographically spread out and complex as ours, let alone one that could be used for forecasting or to predict the likely consequences of a policy change. And as we will see, many of these relations are in fact unstable over time or are able to explain a relatively small fraction of the variability we observe in the data.[19] Moreover, the evidence for a particular explanation or interpretation of an empirical finding will rarely be dispositive, which means that a large number of warning labels will need to be affixed to such an explanation if its purpose is to inform policy.

Unavoidably, therefore, an undertaking of this sort will be more art than (pseudo-)science. But even though macroeconomics probably never will be a science – we have too few relevant observations with which to permit sensible inductive reasoning, and no agreed-on standard of evidence with which to assess proposed hypotheses – it does share the four goals of any science; namely, explanation, understanding, prediction, and control. And because our society contains policy institutions that seek to affect macroeconomic outcomes in a deliberate way, we need to try to achieve these four goals as best we can.

1.4 Digression: Does Aggregation Save *Micro*economics?

The Sonnenschein–Mantel–Debreu theorems demonstrate that the structure that is given to individual-level behavior by standard assumptions about consumer preferences is largely washed away by aggregation. But can aggregation itself yield properties that would allow certain key microeconomic

[19] One important practical use of a statistical macroeconomic model is to identify when residuals are starting to emerge, as these can indicate that a consequential change in the economic environment is starting to take place.

https://doi.org/10.1017/9781009465779.002 Published online by Cambridge University Press

propositions to hold? In a remarkable attack on the problem, Hildenbrand (1993) argues that the answer to this question is a qualified "yes."[20]

Hildenbrand asks what features of average behavior and of the data would permit us to claim that the "law of market demand" – very loosely, the notion that prices and demand move in opposite directions – will hold.[21] We can see the intuition behind Hildenbrand's argument as follows.[22]

Start from the Marshallian demand functions $x(p, y)$, where y is income. Taking the derivative with respect to the price vector and adding and subtracting the term $D_y\, xx'$ (the prime denotes a transpose) yields:

$$D_p\, x(p, y) = D_p\, x(p, y) + D_y\, xx' - D_y\, xx'. \tag{1.4}$$

Note that the first two terms give the Slutsky matrix. If we take the average of this equation over all consumers and assume that the average Slutsky matrix is negative semidefinite, then the law of market demand will hold if we can show (empirically) that $D_y\, xx'$ is positive definite on average.

The average value of xx' can be thought of as a measure of dispersion or "spread," since the average of the outer product of the demand vectors will equal the variance of the outer product plus the outer products of the mean demands. What Hildenbrand examines empirically is whether this spread is an increasing function of income; he does so by using household expenditure surveys to compute various measures of dispersion for different pairs of commodities at different income levels. He concludes that there is evidence for increasing spread in the data, which suggests that the law of market demand might well be an aggregate property of the data.

Why is this answer only a *qualified* yes? Several issues are associated with Hildenbrand's analysis, including the use of broad commodity aggregates like "housing," "food," and "transport" (which can only be justified under relatively strict assumptions), as well as a blurred distinction between income and expenditure (which puts the negative semidefiniteness of the Slutsky matrix on shakier ground, as income can be saved rather than spent). In addition, the experiment that would actually need to be done in order to assess the sign of $D_y xx'$ would involve giving each consumer slightly more

[20] A very distant relative of Hildenbrand's argument appears in Hicks (1956, VII.7).

[21] Strictly, the law of market demand states that the vector of price changes and the vector of quantity changes "point" in opposite directions; that is, for two different price vectors p and q, aggregate demand $F(\cdot)$ satisfies $(p - q) \cdot \big(F(p) - F(q)\big) < 0$. This property in turn ensures the uniqueness of a Walrasian equilibrium.

[22] This discussion is taken from Lewbel (1994).

income and then seeing whether the average spread of their demands rose; because that is not possible to do, Hildenbrand instead looks at demands for households with higher or lower income levels. (He calls the assumption that the two experiments would yield similar results "metonymy.") Finally – and most importantly – even if we accept Hildenbrand's conclusion regarding market demand, it still has nothing to say about Fisher's stability question. So while shifting our focus to the behavior of aggregates isn't quite enough to restore one's faith in the relevance of general equilibrium theory, it does at least let us assert with mild conviction that market demand curves slope down – at least in a dark room from a distance.