5.3 The Radner Version of the ADM Economy

Arrow (1953) noted early on (in a paper not published in English until 1964) that "securities" or financial assets could, in principle, allow for the outcome of an ADM model to be replicated with far fewer markets than the archetypal Arrow-Debreu contingent commodities. These securities were special, as they were ones that paid off in only one contingency, and paid nothing in any other one, and for obvious reasons are called "**Arrow securities.**"

Radner followed this line of reasoning and imagined a market structure where instead of all trade happening at once as in the ADM setting, a small set of markets open prior to the resolution of any uncertainty. These markets allow for contingent trade in just one of the goods. Think of a world with just four physical goods—corn, wheat, alfalfa, and soybeans—in springtime, when planting is about to commence. In this economy, all eating takes places later, at harvest time. But think of the weather at harvest time as uncertain, taking one of three possible forms: sunny, cloudy, or rainy, each of which matters for the size of the harvest. After the harvest, imagine that the world ends.

In this physical setting, a Radner trading system or a Radner economy will allow only weather-contingent trade in springtime (before uncertainty is resolved) in only *one* of the physical goods—say corn. That is, participants in a Radner economy can take part in three forward markets in which they buy or sell promises to deliver or receive three goods—*corn* in rainy weather, *corn* in cloudy weather, and *corn* in sunny weather—before they know the harvest. But there would be no markets in any other goods (in our case, this just means no trade in wheat, alfalfa, or soybeans). The Radner economy then lets uncertainty over the weather resolve itself, but as soon as it does, it allows for a set of markets for immediate consumption in all goods. In our case, four markets would open: one for corn, one for wheat, one for alfalfa, and one for soybeans (and the weather would be whatever it turned out to be). These latter markets are typically called "spot" markets, because they are ones in which market participants buy and sell items for immediate consumption.

Notice that the Radner trading arrangement features fewer markets than the ADM model under uncertainty requires: instead of twelve markets (three forward markets each in corn, wheat, alfalfa, and soybeans), our current market arrangement features seven markets (three forward markets for corn in rainy, cloudy, or sunny weather, and four spot markets once the weather resolved itself). Crucially, as the number of goods and states grows, so does the difference in the number of required markets, and it grows dramatically. For instance, if there were 100 different types of crops and 20 kinds of weather, the Arrow-Debreu trading system would feature 2,000 markets, while the Radner would require just 120 markets.

More generally, under Radner trading, if there are L goods and S states, then $L + S$ markets can do the work of the L *times* S markets imagined in the time-0 ADM trading arrangement. So if there were $S = 500$ states of the world, and $L = 1{,}000$ goods as in my earlier example, then the Radner model asks that there be 1,500 markets (one forward market for delivery of a single good—corn, in our example—in each of the 500 possible states that might occur), and L markets once the uncertainty resolves (e.g., corn, wheat, and 998 other goods and services). Critically, under the Radner arrangement, "only" 500 would have to be contingent commodities. To the extent that we view (and should view—as we will see later) these types of markets as the hardest kinds of markets to arrange, as they are the ones most bedeviled by forces that induce them to fail to work well, this is good news.

In fact, as a general matter, the Radner economy will require far fewer markets to be open at any one time than the all-encompassing Arrow-Debreu model. Specifically, unlike the Arrow-Debreu model, the Radner model asks "only" that there be enough markets for the goods that one plans to consume as of that date, plus markets by which to transfer purchasing power to every possible contingency that might prevail in the *immediately following* trading session (i.e., we don't need financial markets for times further out into the future). Put this way, the analogy to insurance becomes easier to see.

A **Radner equilibrium** is therefore a particular kind of Walrasian equilibrium. It is described (as usual) by a set of prices that, when taken as given, lead all households to be able to execute their desired purchases and sales. However, this definition hides what is different about Radner equilibrium. The set of prices is really a set of *expected* prices since not all purchases take place in any one trading session, and the purchases really involve *plans* to consume and produce in a set of spot markets and in a set of future spot markets at prices *expected* to prevail then, along with choices for the one-period-ahead state-contingent claims in the one good we allow it for at each date (corn, in our example). I italicized the word "expected" to remind you yet again that while the Radner trading arrangement is more "realistic" in the sense

that it features trade occurring over time and in response to uncertainty as it occurs in our world, it asks for a great deal of forecasting power. In particular, it asks households to have tremendous "contingent" foresight for prices under various contingencies, and this should not be forgotten.

The intuition for the Radner arrangement being able to deliver the ADM outcome is this: As long as I, as a market participant, can buy or arrange to receive enough corn in the rainy and sunny outcomes, I will be able to then use the spot markets to trade the corn I receive for wheat. For concreteness, just imagine that *after* the weather is decided, you can buy or sell corn and wheat from and to the WCH. Now, if I know (forecast correctly) the prices of corn and wheat under *each kind of weather* that might have occurred in the interim, then I can work out exactly how much corn I will need to deliver or have delivered if I want to buy a given amount of wheat (or corn) on the spot markets under each possible realization of weather.[6]

As I will elaborate a bit more later, an important and relatively implicit assumption made in the Radner trading arrangement is that households not be constrained in their ability to *borrow* or, more precisely, to short-sell the one commodity in which there is forward trade. In our example, the Radner arrangement requires that households be able to sell as much contingent corn (say, rainy-day corn) as they please, including *more* than they would be endowed with in the state under consideration. The understanding is that they will purchase the rest in the spot market next period (as soon as the weather is revealed). The absence of constraints on short-selling is, in turn, made possible by the implicit assumption that default is not possible—deliveries will be honored. Further below, I will revisit the problems that the absence of such "unlimited commitment" creates. For now, suffice it to say that limits on the commitment of borrowers to repay when it is perfectly feasible for them to do so will generally void the ability of Radner trading to mimic the outcome arising from trade in the full set of Arrow-Debreu contingent commodities. A lesson here is this: if we think a sequential-trading setting is likely the only realistic one we can imagine as practical, then "limited commitment" may well be a relevant barrier to attaining efficient outcomes via any sequential-trading arrangement we might imagine, including Radner's full "one-step-ahead forward markets in one good" setup.

I have just noted that households that interact through markets that open over time and in response to the resolution of uncertainty must

forecast the prices they will come to face in these various situations. However, not just any forecast will do. In keeping with our traditional notion of equilibrium as a situation in which no one is surprised by market outcomes given the particular realization of uncertainty, equilibrium under the Radner market structure will require that households forecast future prices *correctly* (see Mas-Colell, Whinston, and Green 1995, prop. 19.D.1, and Ljungqvist and Sargent 2004, ch. 8, for formal treatments). In what follows, I will describe examples of how the Radner interpretation of the Walrasian model is the one used by modern macroeconomists to organize their thinking about the "real world." For now, it is important to emphasize that once trade starts to happen over time, as certainly seems to be the case in the real world, households *must* start making forecasts of future prices—this has absolutely nothing to do with how "rational" or irrational one thinks participants in real markets are. It is a requirement that is inherent to any setting in which events unfold over time and do so in ways that are uncertain from the current perspective.

5.3.1 A Summary of Radner Trading

We have just seen that in a Radner model of the economy, one allows for the sequential trading of securities in response to the temporal unfolding of uncertainty. Consider the situation where, each time some uncertainty about the world is resolved (e.g., will there be war or not? will I lose my job or not? etc.), we allow markets to open after each such event, operate a WCH to get Walrasian prices in each of these markets, and allow households (and firms, if they want to) to reposition their entire portfolios in light of the new information under these prices. This interpretation, while still very much an extreme portrait relative to the markets one routinely sees in the real world, is certainly much more directly useful for thinking about the trade and price movements that we observe daily. At least it features trade—as opposed to merely deliveries! Moreover, it features trade in response to the arrival of new information, surely a part of why a good deal of trade actually occurs.

But, as you are probably thinking, even the Radner model of trading, despite its more "modest" requirements on the number of markets that must be open prior to the resolution of uncertainty, asks way too much of traders. That is, the requirements on the richness of markets and on the ability of households to forecast that are needed to attain efficient outcomes through decentralized price-taking optimization still seem very fanciful. In a nutshell, the benefits we obtain

in terms of the reduction in the required number of forward markets is accompanied, and perhaps more than substantially "offset," by a serious forecasting requirement. And yet it is still the preferred first step in macroeconomics, even if only to serve as a benchmark against which to measure the cost of dysfunction occurring in a more "realistic" model.

5.3.2 Spot Markets and IOU Markets: Radner and How Macroeconomists Think about Market Dysfunction

To my taste, the Radner model's most profound legacy is the role it plays in helping me and my fellow macroeconomists to classify the roles played by the two different kinds of markets present in any price-based trading system. Since the Radner model features both markets for trade in goods and services that are immediately consumed and markets for the transfer of purchasing power to *future* dates and/or contingencies, it bears a fundamental, if stylized, resemblance to the market systems we observe in daily life. Specifically, it is a model where these two classes of markets and their dysfunction can then be usefully placed into two separate boxes: "spot" and "IOU," respectively.

While the term "spot market" is entirely standard, let me explain the term "IOU market." Any market that is not a pure spot market (where an item is traded for another "on the spot") is one in which an IOU has been issued. This is because *any* delay between delivery and payment necessarily entails credit on one side and an obligation on the other, and hence involves either the implicit or explicit issuance of an IOU. Financial markets—for example, markets for stocks, bonds, futures, and options—all involve (sometimes complex) bundles of fairly explicitly defined IOUs issued by one party to another. Moreover, even some seemingly spot transactions come together with IOU transactions. Notice, for example, that the purchase of a car with a service plan is really a bundle of a car with a set of IOUs issued by the car dealer to a buyer. The dealer delivers a car immediately, which is the "spot" part of the transaction, and also *promises* to repair the car at *some* future dates (e.g., warranties have time limits) and under *some* circumstances (e.g., the policy may not apply if your car gets hit by lightning). The latter is nothing but a set of IOUs.

One critical market that I will place into the IOU category is that of labor. While it is true that some do have jobs that pay essentially in a spot transaction (such as the teenager who may mow neighborhood

lawns in the summer, or a local babysitter, or a seasonal farm laborer), most other forms of trade in labor are longer-term and very much involve promises by both workers and employers. Employment is not usefully regarded for most of us as a spot market transaction. Rather, it is generally a *relationship* expected by all parties to last for at least some time (and often, an open-ended amount of time). It is one that prescribes, implicitly or explicitly, actions for employer and employee alike at various times under various contingencies. Put this way, it becomes clearer that all relationships may be viewed as the trade in (sometimes elaborate) bundles of IOUs.

5.3.2.1 Spots Are OK

It is fair to suggest that macroeconomists generally view spot markets as functioning well; this is rarely where we think the large "market failures" occur. For most consumer goods (e.g., mangoes) and producer goods (e.g., drill presses), most of the time, product quality is discernible, linear prices are the rule, sellers and buyers compete (sometimes brutally), exclusion for nonpayers is typically feasible for most purchases, and stock-outs and prolonged pile-ups of inventories of most items are decidedly rare at most retailers (and certainly across all retailers in a town on any given day).[7] In fact, spot markets tend to function quite well even in those places where the average income level is extremely low, as I have found during my annual pilgrimages to Chennai, India, with its many small retail establishments offering a vast array of products at highly competitive linear take-it-or-leave-it prices.[8]

Of course, spot markets may still sometimes fail to allow market participants to make all the exchanges they want to. As I discussed in detail in chapter 2, even leaving aside market power, public goods, and taxes, asymmetric information could throw a wrench into the efficiency of decentralized trade. In fact, the seminal paper of Akerlof (1970) first helped economists recognize the potential effects of what we have come to call adverse selection, whereby the quality of a good available for sale falls with the price it is expected by sellers to fetch—sometimes to the point of driving all sellers of high-quality goods out of the market. Akerlof's work suggested this possibility in the context of the spot market for used cars. Akerlof's work showed economists that linear prices could not be *presumed* to work efficiently. But before we grow pessimistic, it is useful to remind the reader that the First Welfare Theorem gave us only *sufficient* conditions for efficient outcomes.

Market participants could, and indeed do, augment (and sometimes replace) linear prices with a variety of other contractual features, such as warranties or promises of free auto servicing, etc. As a result, if these promises can be expected to be honored, even spot markets plagued by asymmetry of information on product quality may work well. This is a quantitative question, and one that is now getting more attention because economists have the computational and game-theoretic tools to analyze such cases.

5.3.2.2 IOUs, Maybe Not So Much?

It is probably fair to say that if macroeconomists disagree on the extent of market dysfunction, it is most often in their assessment of the performance of IOU markets. There is a good reason for this: asymmetric information problems are likely worse in IOU markets, and on top of it, there is a second class of problems—those created by limits on parties' commitment to act in the future as they now promise to do. While I've already argued that imperfect commitment plays a role in creating difficulties for centralized systems, we now see that it can create problems in decentralized settings as well. By the end of this book, I hope to have persuaded you that, on balance, limited commitment is the central impediment to allowing societies to attain "good outcomes."

A reason that limited commitment is the main form of sand in the gears of economic life is that even the occasional problems one sees in spot markets may have their roots in IOU-market dysfunction. Think of the market for used cars. While bad cars may lurk among the good ones on any used-car lot, one may not be doomed to a crapshoot when buying a used car. A warranty may be just what is needed to separate the good from the bad and the ugly. But what's a warranty? It's a bundle of IOUs. So the extent to which a warranty on a car can overcome a buyer's fears that it is a lemon depends on the buyer's faith that the issuer of such promises will make good later on. Barring this, there is less reason to be confident that the spot market for cars will work well.

Thus, while limits to commitment are generally irrelevant (almost by definition) in impeding spot transactions, in IOU markets asymmetric information and limited commitment can *interact* to further worsen matters. I will discuss some models that study these impediments to trade later, but for now, let me expand on the kinds of

problems IOU markets can present. Let's think of the decision to buy health insurance. Insurance purchases are an example that can feature the problems of both asymmetric information and limited commitment. As already broached in chapter 2, these are economists' two "usual suspects" in creating problems for decentralized trade. In the insurance context, asymmetric information can cause problems in the following manner: If I know more about my own condition than the insurance company does, the insurance company should worry that I will lie about my weight and my cigarette addiction. The insurance company cannot simply raise premiums; if they did, the relatively healthy might drop out, leaving the pool filled with even more overweight smokers. Of course, they might offer me a high-deductible plan, but by definition this is *incomplete* insurance: when something bad happens, it *will* cost me out of pocket. Insurance as we experience it is emphatically not a Radner economy contingent claim.

Now, enter limited commitment. By definition, this is a problem that is confined to IOU markets. In the insurance context, one certainly hears of people being dropped by an insurer the moment they become severely ill. Of course, this does not always happen, but it does happen to some. At the time I make a decision to purchase insurance, what is relevant is my *assessment* of the likelihood that I will be unceremoniously dropped at some critical juncture, perhaps on a technicality in the fine print of the contract. If we all share an assessment of this risk, notice that the adverse selection may get worse. All else being equal, potential buyers now worry that they will be dropped. If that's the case, prices will have to remain "high" to make up for the fact that the relatively healthy face even less incentive to buy in. In the example of a car service plan cited above, the problem is similar: the buyer wonders if the servicer will honor his promises. And, knowing this, the seller may not be able to offer a highly comprehensive service plan: the buyer won't believe that he'll make good on the promise. In turn, the seller may lack a credible way to show the buyer he believes in the car he's selling. In the end, then, what could have been a mutually beneficial trade might not occur.

Even worse, trade might not occur simply because buyers are pessimistic in a way that leads to outcomes that *fail to disconfirm* their pessimism. An example might be a setting where all would-be buyers of insurance think that no insurer will ever make good on a policy. In this world, no one would ever buy insurance, and as a result, no insurer would ever get the chance to prove anyone wrong in their