7

Development Traps I: Complementarities

7.1. Introduction

We've now studied some of the classical theories of economic growth and their implications. Some implications, such as the predicted positive link between investment rates and growth rates, are borne out by the data. Others are not. For instance, there is a lot less convergence than the simple Solow model would have us believe, though there are pockets of countries that have clustered close together. Variations of the Solow model, which allow for differences in the accumulation of physical *and* human capital, do match the observed diversity in per-capita incomes more closely.

But I want us to go further. It is easy enough to point to variations in rates of physical and human capital accumulation, but the extra step to branding these as *root causes* of underdevelopment is perilously tempting. We should avoid taking that step. The low accumulation of capital — mirrored in savings rates or educational choices — is just as much a symptom of development as a cause. To the extent that it intertwines with the development process (or the lack of it), we cannot rely on these observations as *explanations*. Of course, we can condition on these variables all we want, and triumphantly point to what's left as evidence for (conditional) convergence, but after a point, that isn't saying very much.

Do we really base a theory of underdevelopment on the presumption that individuals in different parts of the world have some *intrinsic* difference in their willingness — or ability — to save, to work, or to procreate? If you do believe that, you could argue that such-and-such country is underdeveloped because it is populated by people who habitually save very little, or are lazy, or reproduce like rabbits. Or perhaps more subtly, you could claim that there is something wrong with the "cultures" of these countries, leading to low savings, rapid demographics, laziness or corruption.

Somehow, that does not seem intellectually right. We would like to have a theory which — while not belittling or downplaying the role of social, cultural and political factors — does not simply stop there. After all, low incomes can provoke, in turn, low savings rates, because a larger fraction of income is needed for current consumption. The same is true of demographics — underdevelopment might be a cause of high population growth rates, just as high population growth rates themselves retard the development process. Two-way causation may be fundamental to our lives.

That is, a more substantive view of economic development is that much of it possesses "chicken-and-egg" features, in which underdevelopment is seen not as a failure of some fundamental economic parameters, or socio-cultural values, but as an interactive outcome that hangs together, perhaps precipitated by a belief system or by historical conditions. The very same fundamentals may result in multiple, self-sustaining outcomes. There are three reasons for taking this approach very seriously.

[1] It leads to a theory, or a set of theories, in which economic "convergence" (of incomes, wealth, levels of well-being) across countries is not to be automatically had. Yes, there are factors that we can "condition" on, but once we start accepting the possibility that often, these factors *cannot* be kept equal, the notion of conditional convergence starts losing its relevance and a broader view of development must be sought. Convergence may or may not take center-stage in that broader view.

[2] It does not rely on "fundamental" differences across peoples or cultures. This is not to say that culture — like conditional convergence — does not play a role. It does. But I also take the view that culture, or religion, along with several other economic, social and political institutions, are all part of some broader interactive theory in which "first cause" is to be found — if at all — in prior histories or the vagaries of geography. That isn't to say that the earlier ahistorical theories are wrong, but that they explain matters at one level, and that to understand more we must go deeper, by not necessarily regarding as exogenous what these models regard as exogenous.

[3] It does *not* concede to the nihilistic dictum that "anything can happen." Only certain configurations of variables can lock together in the theories that we now go on to study. But it does take one step back from naive determinism. It argues that there could be a certain degree of openness about observed outcomes, a gap that might be filled by historical accounts largely based on hindsight, or a consideration of small accidents that magnified and hardened over time, and in some cases we may have to admit that we don't know how a particular configuration came into being, but be content with our ability to rule *out* a large number of other possibilities, though not all.

[4] It implies a very different role for government policy. Specifically, the theory will place greater value on one-time, or temporary, interventions than theories that are based on so-called fundamentals. For instance, imagine that it is some deep cultural, country-specific trait that keeps corruption high; then, while a suitably punitive policy will lower corruption, there is no telling when that policy can be taken away, or indeed, if it can be taken away at all. For in the absence of the policy, the theory would tell us that society would revert to its familiar corrupt ways. In contrast, a theory that fully recognizes the two-way influences that run between development and corruption would promote a policy that pushes the interactive cycle into a new equilibrium. Once that happens, *the policy can be removed*. In this sense, the new theory has very different implications from its traditional counterpart.

Very broadly, then, we are interested in seeing how historical forces and expectations shape the overall economic pattern displayed by a country or region. You may not have attended numerous conferences on economic development at the time of reading this book, but I have (and perhaps you soon will). Many are the occasions when such questions are raised, and equally numerous are the occasions when the panelist or speaker weightily invokes the complex, deep role played by culture, society, and "historical forces." Indeed, the thorny issues of development are difficult, and to understand them we must understand all the complexities of society, in its many forms. On the other hand, simply *asserting* how complicated matters are does not take anybody very far (although I have often been surprised by the frequency with which such pronouncements are described as "insightful," as "going beyond the narrow boundaries of our discipline," and so on).

The fact of the matter is that people the world over are intrinsically the same: they are all human beings, with the same hopes and desires. Why do cultures and economies, born from the same intrinsic material, perform so differently? The purpose of this chapter (and the next) is to discuss such issues. Two themes will be central. The first, that receives relative emphasis in this chapter, can be described as *multiple equilibrium*: the notion that at any given moment, society could be locked into one of a variety of outcomes, depending on individual *beliefs* about the actions of others: their choice of technological platform, their holding of (or flight from) a currency, their notion of what's fashionable, their proclivity to be corrupt or not, and a plethora of other situations that we discuss below. The second theme is closely related but conceptually distinct: it is the notion that historical events — early technologies, social norms, or institutions — might cast their long shadow into the present, locking today's society into different outcomes. We might call this a theory of multiple *steady states*: while for each historical condition, there might be a unique equilibrium, those equilibrium paths are capable of steering society into distinct long-run outcomes that are deeply dependent on distant histories.

Any particular example will have shades of both multiple equilibrium and history-dependence, as we shall see. In any case, these themes break with the idea of convergence from the very start, and instead begin to entertain the possibility that different outcomes could co-exist within the same structure.

7.2. Complementarities and Multiplicity

A central feature in this chapter is the idea of a *complementarity*: a situation in which the taking of some action by a larger number of individuals encourages even more individuals to adopt that action. We begin by developing this notion.

7.2.1. Qwerty. The top left row of letters on your computer keyboard begins with the familiar sequence $q, w, e, r, t, y, \ldots$. Have you ever wondered just how this particular sequence came into being? David (1985) offers a fascinating account of typewriter history, observing that the QWERTY arrangement sprang largely from a need to avoid typebar tangles. The original typewriters were mechanical, of course. You hit the desired key and a long lever bearing the imprint of this key would rise up, much like a striking cobra, and hit the typewriter ribbon. If you hit two keys at once or two keys in very quick succession, the two levers would jam. I still have memories of my father in the office of our Calcutta home, reaching into the well of his Remington typewriter to deal with tangled keys, eyes squinted against the ever-present cigarette smoke.

The QWERTY layout emerged, in part, to reduce the frequency of such jams, *by slowing down the speed of typing*. In no way was it the best keyboard layout for an electronic typewriter or a computer. Indeed, the Dvorak system, introduced in 1932, was an alternative that repeatedly won speed-typing contests. Why, then, despite its obvious inefficiency today, does QWERTY still rule the roost?

To understand this, it is important to realize that typing skills were part of the intricate network of business and industry. Typists came from typing schools and were almost exclusively hired by firms.[82] Given that all firms were hiring QWERTY-trained typists, it made little sense for any *one* of these firms to invest in, say, Dvorak-style keyboards, and train their typists accordingly; the costs involved in retraining were simply too high. Yet, if *all* firms and typing schools (by an stupendous act of coordination) had adopted a different system, the efficiency gains may have been significant. We have, then, a self-fulfilling situation that is difficult for any *one* entity to get out of, because the return to each entity depends on what every other entity is doing. You simply cannot ask the question, "QWERTY or Dvorak for you?" in a vacuum. Your answer depends on how others have answered the question.

The QWERTY example illustrates a *complementarity*. The number of willing adopters depends on how many others adopt, or are expected to adopt. Depending on the initial history (in this case the existence of jam-prone mechanical typewriters), multiple adoption configurations are possible. We live in a QWERTY world. Perhaps there is a parallel universe in which electronic keyboards were invented before typewriters, there is no risk of jamming keys, and everybody does Dvorak instead of QWERTY.

7.2.2. Traffic. I live in lower Manhattan, and there are two ways of escaping the city (by car, that is) if you want to head upstate for a calm weekend. You could head east to the FDR Drive and go up the eastern edge of Manhattan, or you could take the West Side Highway (WSH) and drive up the western edge.

Figure 7.1 displays the total number of weekend travelers as the width of the horizontal axis. Count the number of people traveling on the FDR from left to right and the number on the WSH from right to left. The greater the number on each route, the higher the travel cost on that route (congestion), so the lower is the net payoff of that route. An increase in the number of users inflicts "negative externalities" on all users. These "payoff curves" are depicted in the diagram.

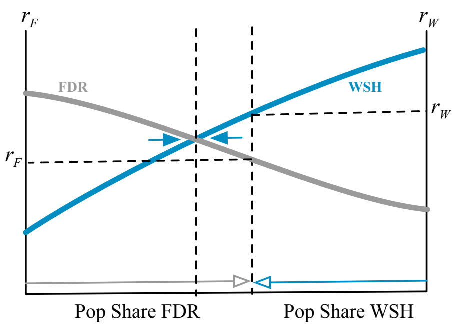

Figure 7.1. Anti-Complementarities: FDR and WSH.

A particular allocation of travelers is described by the point C, with attendant travel returns of r_F (for FDR) and r_W (for the WSH). It is reasonable to expect that this allocation will "drift" to the left in the diagram, as more travelers join the higher-return WSH. The system heads towards D, as travelers divide up the two routes to equalize returns. *At D*, the returns are equalized across the two highways. History can have no effect on the equilibrium dispersion of traffic on these routes. It wouldn't matter if in one universe the FDR was built first, and in another the WSH was built first. In both universes, traffic would gravitate towards the same "equilibrium." For

[82]We are considering times that precede the growth of the market for personal typewriters.

instance, in the first universe, the return to travel on FDR would be a (see diagram) to begin with, while the newly built WSH would have a higher return of b. But of course, traffic will begin to switch to the new route, which drives up its cost and brings down the cost of the old route. Ultimately, the commuters will use the highways in a way so that returns are equalized across routes. This solution ultimately occurs regardless of which route was created first.

7.2.3. Qwerty Revisited. But replace FDR by Qwerty and the West Side Highway by Dvorak, and the story is dramatically different. To see why, study Figure 7.2. The diagram looks very similar to Figure 7.1 at first glance, but look again and you'll see it isn't. The return to each technology is now an *increasing* function of the number of existing users, the opposite of the traffic example. The return in a Qwerty-only society is a. The introduction of Dvorak (with no takers just yet) will come at the lower return b. And there *will* be no takers, because no individual has the incentive to switch: $a > b$ after all. The same is true in a parallel universe where Dvorak comes first: there will be no takers for Qwerty. History matters.

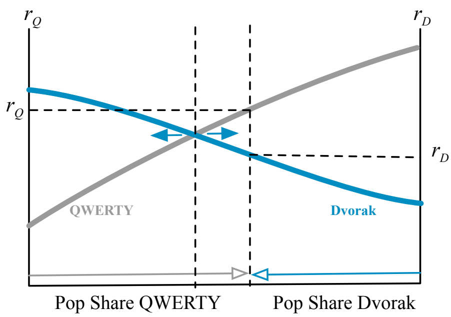

Figure 7.2. Complementarities: Qwerty and Dvorak.

What about the traffic equilibrium point D? It's still there, but now it is "unstable." It is true that if by some magical fluke, Qwerty and Dvorak were "allocated" to users *exactly* as demanded by D, no one would have an incentive to switch. But this point is as unstable as a pencil teetering on its tip: the slightest move right or left will spiral the system further away to one of the "edge" allocations. For instance, eyeball the point C and notice the tendency to move away from D, not towards it.

Finally, note that the all-Dvorak return may well be higher than the all-Qwerty return. (It is, in our diagram.) All-Dvorak may be a better social outcome compared to all-Qwerty. But that does not prevent societal lock-in, if Qwerty is already there by the advantage of historical precedence.

7.2.4. (Many) Other Examples. *Platforms.* Several websites work off complementarities, attempting to lock individuals in by promising access to a wide range of *other* individuals who are also locked in. Facebook, AirBnB, or Tinder work off exactly this principle. Sites that have accrued a lot of users are attractive to new users precisely because of the complementarity: the very size of their membership creates variety and acts as an insulating device against rival incursions.

Infrastructure. Suppose that a public transportation system has a fixed overhead cost, plus a variable cost per user. To break even, users must individually pay the variable cost, and their share of the collective fixed cost. If the number of users is small, the fixed cost cannot be spread very efficiently. That will deter use, thus completing the vicious circle of a small number of users. Conversely, if there are many users, the fixed cost will be defrayed among many people: that, too, could be self-justifying, leading to a "virtuous circle" of high usage.

Finance. A well-diversified, “thick" financial market is attractive to individuals who are seeking relatively safe returns on their assets. At the same time, that larger inflow of money is more conducive to market diversification: market makers can open up new areas of financial investment. So diversification begets diversification, precisely because there are complementarities in financial market participation.

Capital Deepening. Greater roundaboutness in production increases the productivity of capital, the scale of aggregate production, and in this way the final demand for individual machine varieties. This may in turn justify the greater roundaboutness of production.

Social Capital. High rural-urban migration can destroy social capital back in rural areas. In turn, that destruction can increase the pace of rural-urban migration.

Discrimination. Individuals or groups discriminated against may not invest in human capital, perpetuating that discrimination.

Currency Crises. Apart from the fundamentals of holding or selling a currency, there is a strong incentive to sell if other individuals are selling. This forms the basis of a class of currency-crisis theories based on complementarities.

Corruption. Consider the decision to evade taxes. Suppose that every tax return with evaded taxes emits a suspicious signal, and auditors examine a fraction of those returns. If a country has a limited workforce of auditors, then there can be multiple equilibria involving different levels of evasion. In one equilibrium, evasions are low, and most suspicious returns can therefore be audited. That deters evasion, closing the cycle. In another equilibrium, evasion is widespread and only a small percentage of returns can be audited, thereby encouraging widespread evasion in response.

Speeding. A similar phenomenon can be seen with another form of “corruption": speeding on highways. If you drive, you will have surely noticed that traffic forms into local clusters. Some clusters obey the speed limit, because a single deviant is likely to be picked up by the police. Other clusters are traveling at high speed, in the hope that the police will have to pick one of many speeders, and it won’t be them. The same driver can behave very differently as she encounters different clusters on her journey.

Social Norms. Social norms provide a leading example of complementaries. Individuals want to conform to what other individuals see as proper behavior. In different societies and settings, the very same individual can behave in very different ways depending on how *others* behave. Examples include the blowing of car horns, throwing garbage on the streets, spitting in public, or jumping the queue.

Monkeying With The Rupee

There’s a story about how to catch a monkey. You use a jar or an empty coconut shell and fill it with peanuts. Monkey approaches, reaches into jar and clenches its greedy little paw around the peanuts. But it can’t pull its full hand out, and what’s more, it won’t let the peanuts go. End of monkey.

How governments in developing countries have wished for similar success when it comes to foreign investment! How India or Indonesia or Brazil or many others have hoped that the hard currency could come monkeying in, and then stay, forever enraptured by the goodies that emerging markets have to offer! Alas, it’s never worked that way. What flows in can flow out. Foreign capital flow is a two-edged sword, and both edges can be very sharp indeed.

The basic economics of this is pretty simple. Imagine a huge stock of hard-currency-denominated investible funds, forever sloshing around in search of the best returns. For a developing country, the urge to tap into these funds is immense. Ideally, it would like those funds to appear as hard, irreversible investment. But that isn't the way the peanuts are shelled. Much — most? — of the investment will appear in softer form: emerging market stocks, or for the not so choosy, emerging market funds, or for the still less choosy, emerging government bonds. And so it was that India started on the Great Upward Path: money pouring into its coffers from abroad, accompanying tariff and quota liberalization then permitting easy purchase of foreign goods without a huge depreciation in the rupee, the outward drain being more than easily matched by the inward flow.

But all of this stuff, apart from the hard investment, is reversible. Much of the short or medium-return to portfolio investments is prey to severe herding. Scenario 1: money comes into an economy, stock prices climb, the currency stays strong, rates of return are high. Scenario 2: money flees, the stock market tanks, the currency nosedives, returns fall. Now listen to the one sentence that explains (almost) everything: *both Scenarios 1 and 2 can coexist in the same economy with the same fundamentals*. Expectations can drive enormous regime changes.

But then, what drives the transition from one regime to another? Often, though not always, the answer is that there is no answer. Or at the very least, there isn't an answer which in any way can predict this abrupt transition in any deterministic fashion. Markets almost always react long before the fundamentals necessitate those reactions. For instance, a developing-country government might have a large amount of debt denominated in hard currency. Perhaps the citizenry gets too used to the inflow of hard currency and ratchets up its lifestyle, so that the country runs a current account deficit. Or perhaps there is a war or an internal conflict, or a debate regarding economic policy. Perhaps a few Dr. Dooms issue gloomy forecasts. All of this is true (to varying degrees) of India. The country may be perfectly solvent nonetheless, but the specter of possible future insolvency can precipitate a crisis today as the slush money is sucked out. Rome may not have been built in a day, but currency holdings often appear to be: what goes up can come down very fast indeed, without any necessary fundamental justification.

7.3. The Complementarity Map

The QWERTY-Dvorak story and every other example that I've given above can be captured cleanly by a very simple device, which I will call the *complementarity map*. Study Figure 7.3 in what follows. Each panel in that Figure plots on its horizontal axis the fraction n of people who are *expected* to take a particular action (e.g., using QWERTY, throwing garbage, speeding, fleeing from a currency, using a particular online platform, and so on). On the vertical axis, we plot the fraction $x(n)$ of people who — in response to n — *want* to take that action. That's the map we're after.

As a concrete example, consider the decision to be corrupt; say, to take bribes. Imagine that everyone has a mental threshold: they will take bribes if the chances of being caught are low enough. Those chances, in turn, depend negatively on the fraction of corrupt individuals in the society (perhaps because policing is costly). So the greater the fraction of people n *expected* to be corrupt, the greater is the fraction of people $x(n)$ who will *want* to be corrupt. That is a complementarity, and indeed, the map that connects "expectation" n to "response" $x(n)$ is upward-sloping. typically starting

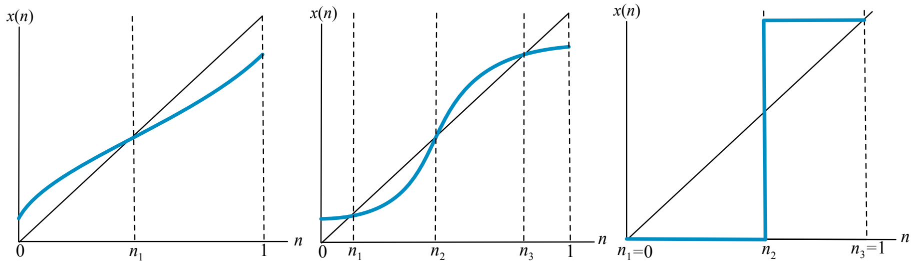

(a) Single Equilibrium **(b)** Multiple Equilibria **(c)** Corner Equilibria

Figure 7.3. The Complementarity Map.

above the “diagonal line" $x(n) = n$ and ending below it; see Figure 7.3.[83] That map will have to cut (or touch) the diagonal $x(n) = n$. Each contact of the map with the diagonal — that is, each n with $x(n) = n$ — represents a potential equilibrium. Values of n with $x(n) \neq n$ cannot constitute such an equilibrium because the expectation of the fraction of corrupt people (n) does not match its realized value (which is $x(n)$).

How many equilibria will there be? In general, there's no telling. There could be just one, as Panel A of the Figure shows. Or there could be several. In the latter case, we have multiple equilibria, as shown in Panel B by the three points n_1, n_2 and n_3.[84] Two of these — n_1 and n_3 — are “stable," and the third (n_2) is not. To see this, perturb the equilibrium n_1 by slightly increasing n. The mapping “responds" by telling us that the number who want to take the action at this higher expected n falls short of that expectation: the map lies below the diagonal. So n_1 acts as a gravitational attractor: the system falls back towards it again (much like the dynamics around the steady state of the Solow model). The same is true of perturbations to the left of n_1. Similar observations hold in the neighborhood of the equilibrium n_3. Once “in the orbit" of these equilibria, the number of corrupt individuals (to continue with our example) converges inexorably to one of the two levels n_1 and n_3.

But this isn't true of the intermediate intersection n_2. That is also an equilibrium in the sense that expectations are self-fulfilling, but it is unstable. As mentioned before, this equilibrium like the proverbial pencil teetering on its tip, and a slight perturbation will send the system away from it. To its right, the system accumulates mass as it goes towards n_3. And to its left, the system collapses to n_1. (You should carefully verify these claims.) In short, “stable equilibria" are given by those intersections of the complementarity map that cut the diagonal “from above," as in the case of n_1 and n_3. And initial conditions matter: if the system starts to the left of n_2, it will fall back towards n_1, and if it starts to the right of n_2, it will careen towards n_3. The history of where we start plays an important, though perhaps not exclusive role.

[83]While we're about it, let's pay some attention to the end-points. When $n = 0$, so that no one is expected to be corrupt, $x(n)$ could still be positive, if there are some individuals who would engage in corrupt activitiies anyway. Likewise, when $n = 1$ — everyone is expected to be corrupt — $x(n)$ is typically strictly less than 1, because there will be people who value honesty no matter what.

[84]There could be even more crossings, but three is enough to make all the points we need to.

You should check that the QWERTY-Dvorak example also fits into this discussion. To do so, just use n and $x(n)$ for the fraction of users choosing QWERTY. Verify that you can translate the QWERTY example perfectly into Panel C of Figure 7.3, where $n_1 = 0$ stands for zero-QWERTY adoption and $n_3 = 1$ stands for full-QWERTY adoption.

Finally, the complementarity map also works just fine for situations without complementarities! In those cases it will typically be *downward*-sloping. To see this, map the highway example into our diagram: just use n and $x(n)$ for the fraction of individuals using FDR, but notice that the resulting map *must be* decreasing. Notice how a decreasing complementarity map typically makes for a single equilibrium: a declining function can have at most one intersection with the diagonal.[85] A complementarity map could even switch slopes. We will see an example of this in the next section.

7.4. Underdevelopment as an Equilibrium Trap

Pervasive complementarities can lead to a "low-level equilibrium trap," while at the same time there is another, better equilibrium, if only all agents could appropriately *coordinate* their actions to reach it. The genesis of this view lies in the classic work of economists such as Allyn Young (1928), Paul Rosenstein-Rodan (1943), Ragnar Nurkse (1961) and Albert Hirschman (1969), research that wasn't part of the mainstream for many decades until some of it was revived in modern form.[86] According to this view, economic underdevelopment is the outcome of a massive *coordination failure*, in which several investments do not occur because other complementary investments are not made, and these latter investments are not forthcoming because the former are missing. Low investment all around can be an equilibrium, in the sense that no individual investor will want to deviate from that state of affairs, and lose still more money.

A general failure of investment can reduce national income, thereby leading to low demand for all products, and depressing the return to any one investment. This was a point that deeply concerned economists such as Rosenstein-Rodan and Nurkse. An individual firm expands if it finds that expansion is privately profitable. But its actions have economy-wide implications that are both related to private profitability but also transcend it. The former occurs because greater profits are paid out to owners of the firm and thereby become part of national income. The latter occurs because payments to inputs (such as labor) *also* become part of national income, but to our firm, those payments represent *costs*, not profits. Overall, a firm interested in private profits may or may not act in the interests of the economy as a whole.

7.4.1. National Income and Market Size. To expand on the effects of low national income, think about market sizes for firms.[87] Sophisticated intermediate inputs that are highly productive at large scale may well be jettisoned when market size is small. That jettisoning can lead to low income and small market size, completing a vicious circle. Take construction, for instance. In developing countries, this is a pretty labor-intensive activity. The area is cleared by hand, rubble is removed in small baskets carried by

[85] For the careful reader: this is true of *single-dimensional* maps. With multi-dimensional actions, there may indeed be multiple equilibria, but those equilibria won't be "ranked" by low to high intensities of some action, as they are for increasing complementarity maps.

[86] See, for example, Murphy, Shleifer, and Vishny (1989), Matsuyama (1991), Krugman (1991), Rodríguez-Clare (1996), Ciccone and Matsuyama (1996), or Baland and Francois (1996).

[87] See Ciccone and Matsuyama (1996) for an explicit account, and also Dixit and Stiglitz (1977), Grossman and Helpman (1991), Romer (1990), and Rodríguez-Clare (1996).

hand, cement is often mixed at the site and carried by hand, and walls are put up brick by brick. In industrialized economies, each of these tasks has been automated: cranes are used for clearing and prefabricated walls are erected at the site. Each automated instrument, in turn, is produced through a complicated activity: think about how cranes and prefabricated walls are themselves produced. These sophisticated inputs can be extremely costly to use at small scale, but their absence leads to lower productivity and national income, and thereby back to small scale via low demand. The cycle is complete. But the other side of the coin is a virtuous circle. A large market increases the demand for sophisticated intermediates goods, thereby generating high income, high demand, and large markets, thus completing the circle.

7.4.2. Linkages. Or perhaps the externalities imposed by one investment on another are more direct. An expansion of, say, the steel industry might have direct input linkages to, say, coal via the availability of mining equipment. These linkages were particularly emphasized in the writings of Albert Hirschman. The outputs of some sectors often serve as inputs for other sectors. When investment is low, the consequent low-output regime propagates "forwards," generating input scarcity for other sectors and inhibiting investment there. The converse interaction is also relevant. If a sector that uses my output as an input is depressed, then the shock waves propagate "backwards." I will not invest in my sector, not because I have a shortage of inputs, but because I expect that there will not be enough demand for my products.

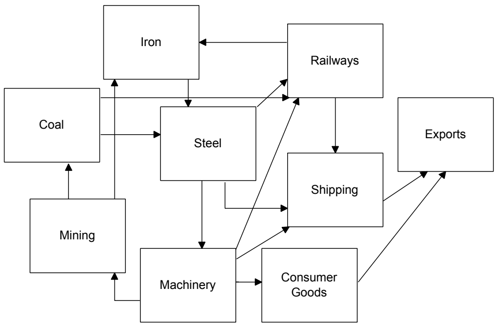

Figure 7.4. Linkages and Coordination.

Figure 7.4 illustrates a network of interaction between various sectors. Only a small number of industries are involved in this diagram, but you see the idea. The arrows are there to suggest that one industry might facilitate the development of another by easing the conditions of production in the latter industry. These are the "forward linkages," but you can think of the opposite directions as potential conduits for "backward linkages." Think of each box as being either "switched on" (high investment) or "switched off" (low investment). Thus two different configurations can come about in two otherwise identical economies. In one, all activity is depressed — each sector is switched off. In another, the boxes are switched on — each sector is in a high-investment state. Both configurations may well be self-reinforcing. Because the gain to investing depends positively on investments made by others, we therefore have a complementarity. From this perspective, the possible multiplicity of equilibria should not be at all surprising.

7.4.3. Investments With Complementarities. To develop this last argument a bit further, suppose that an economy has many firms. Index these firms by the label x on the interval $[0, 1]$. Suppose that each firm x can either use a simple, low-profit method of production, or it can invest and actively scale up its production using a powerful and productive technology. There's a catch, though: while the new technology is more productive and there are potential profits to be made, it's costly to implement. There is a setup cost for firm x; call it $s(x)$. Write the firm labels so that $s(x)$ is increasing; then low-index sectors have lower setup costs and are more amenable to investment. Notice that firm x's *net* profit from investment is given by

$$\pi = R - s(x),$$

where R is the present value of gross profits or returns accruing to each firm, subtracting its variable outlays on labor and other inputs. The important externality is that R could well vary with the fraction n of investing firms in the economy, and in the case of particular interest to us, $R = R(n)$ is *increasing* in n.

If we think of the low-profit method of production as yielding a return of zero, or equivalently if we think of $R(n)$ as the *extra* returns from investment over and above those from the low-profit method, then every firm with index x such that

$$\pi(x,n) = R(n) - s(x) \geqslant 0 \tag{7.1}$$

will invest, under the expectation that n firms in the economy are investing. That allows us to build a complementarity map which runs through the following causal arrows: $n \mapsto R(n) \mapsto x(n)$, where (remembering that firms have been arranged in order of increasing setup cost), $x(n)$ is the *largest index* in the economy for which equation (7.1) holds. Figure 7.5 does just this. Any value of n determines the expected return $R(n)$ for each potential firm. By comparing these with setup costs, we can determine the share of firms that actually want to invest, which is $x(n)$. When $R(n)$ increases with n, $x(n)$ must increase as well, thereby giving rise to a complementarity, as displayed in the Figure.

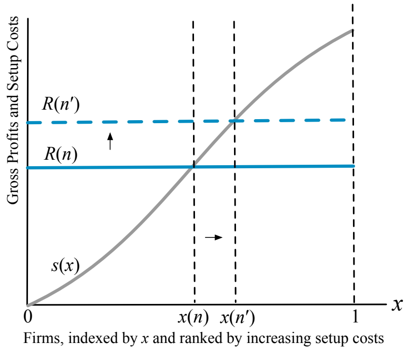

Figure 7.5. Investment Complementarities.

I want you to note very carefully that x increases with n because in the story that interests us, more investing firms *heightens* the return to all investments, much as a rising tide raises all boats. Such linkages are typically salient *across* industries, as higher production in one industry generates propitious conditions for higher production in another. That contrasts sharply with the logic that we are used to from an introductory microeconomics class. There, more firms generate less returns for any one firm because the playing field within the industry gets crowded. That logic isn't wrong, of course, but it is only one effect that's particularly suited to firm competition *within* an industry. In that setting, $x(n)$ would decrease as n increases. The situation is analogous to our FDR-WSH construction, and is typically used to show how free entry determines the number of active firms within a single industry.

Let's return to our case of interest, then, in which economy-wide investment generates better opportunities for all firms. The problem is that no individual firm fully internalizes these global effects. That is why it is possible to have multiple equilibria, just as in our other examples. For instance, when the share n_1 invests, it generates enough economy-wide effects for precisely n_1 firms to invest, which is graphically captured by the function $x(n)$ intersecting the 45° line at the point n_1. If some more firms are expected to invest, that will indeed induce greater investment all around in response — if $x(n)$ is an increasing function — but that response may not be enough to justify those heightened expectations to begin with. That is captured by the fact that $x(n)$ falls below the 45° line just to the right of n_1, and it makes n_1 a "locally stable" outcome.

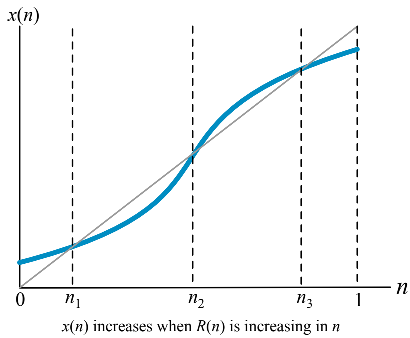

Figure 7.6. Investment Complementarities.

But just as in our discussion in Section 7.3, there could be other locally stable outcomes. In Figure 7.6, the left panel shows another equilibrium intersection n_3, with a larger amount of investment. Like n_1, this too is robust to small perturbations of investment in its neighborhood. But no firm can unilaterally precipitate a move from n_1 to n_3 on its own. Indeed, the entirety of our discussion on complementarity maps in Section 7.3 applies to this situation: the possibilities of multiple locally stable equilibria, as well as the intermediate, unstable tipping point n_2 that must be crossed if we are to move from one locally stable outcome to another. A "big push" — in the form of a large but temporary infusion of investment that tips the system past the locally unstable point n_2 — may be called for. This is the view of development policy that's confronted and debated in the work of Rosenstein-Rodan, Nurkse, Hirschman and others.

7.4.4. The Private and the Social. The act of investment has private payoffs for the firm (namely, profits), but also affects the fortunes of other firms. As already discussed, a firm's actions could have the serendipitous by-product of raising national income, or generating direct demands for the outputs of others, or making possible the expansion of other firms by relieving input shortages. For instance, if there is widespread unemployment of labor, a firm's investment might contribute to the lowering of that unemployment and so benefit other firms (not to mention the benefit to labor itself). To the firm, however, it's all lopsided: having to pay for that extra labor is a *cost*, and the socially valuable investment may well not be forthcoming.

This approach is to be contrasted with a viewpoint closely associated with Adam Smith. In his monumental *The Wealth of Nations*, Smith argued that we need not rely on firm benevolence to guarantee market efficiency. Market *prices* and attendant *private* interests will be enough to avoid those failures. Prices serve as signposts, directing investment to socially efficient ends. How can we reconcile the two views?[88]

The answer is that *all* externalities, and especially the externalities from investment that we are emphasizing now, can be viewed as stemming from the absence of, or imperfection in, some appropriate set of markets. An entrepreneur who builds a mass transportation system might raises real estate prices of nearby locations, creating a (positive) externality. It wouldn't be an externality if there were a market in which potential real estate beneficiaries could compensate the entrepreneur. If there were

[88]To keep the discussion simple, I have abstracted from some more subtle points. In particular, Smith's arguments, which lie at the core of the so-called "First Fundamental Theorem of Welfare Economics" — that every competitive equilibrium is socially efficient — do *not* imply that there is just one equilibrium. There could be many, so multiplicity is not ruled out by Smith. But no two equilibria can be unambiguously ranked, in the sense that all economic agents are better off in one equilibrium than in the other.

such a market, the externality would be completely internalized and the system would be built if and only if it were socially beneficial.

But this way of thinking is not always useful or realistic. Think of traffic congestion on a busy street. People caught in the traffic are affected to different degrees. Imagine a competitive market system with market-determined compensation rates to be paid from one driver to another for all imaginable contingencies (one if you are on your way to a hospital to have a baby, one if you are late for a meeting, one if you are out for a relaxed drive, and so on). If such markets existed and all contingencies were immediately verifiable, then every street would hum efficiently. But this is a vision of markets taken to the point of absurdity.[89] Several examples in this chapter fall into this category, and more they involve social interactions of this sort, the less likely it is that markets will be available to cover the external effects of every such interaction.

On the other hand, forcing ourselves to confront the "why-isn't-there-a-market" counterview can sometimes leads to useful insights. Recall Figure 7.4, with its inter-sectoral linkages. Where, precisely, is the externality? The figure suggests, for instance, that the steel industry cannot expand because there is a limited market for products that use steel. These limited markets go from product to product, possibly feeding back on the steel industry (see the figure for some of these feedback loops). How does this notion of a "limited market" transmit itself to a producer?

There is only one answer to this question. It is that the producer fears that an expansion of output will lead to a fall in the price of her good. This argument includes as a special case the complete inability to find a buyer, which is just another way of saying that extra quantities cannot be sold at *any* positive price. That lowered price is, of course, a positive externality for a potential buyer of the product, but the producer does not internalize the fortunes of the buyer, *nor does the market make him do so*. Similarly, if a producer feels that she cannot expand production easily because the price of her *inputs* will go up, seh may not carry out the expansion even though the higher price is good for the supplier of the input).[90]

It is very important to understand that such answers make sense only if the market system is not competitive in the usual textbook sense: *that there is a price at which one can buy or sell unlimited quantities of a commodity*. If a producer (or a buyer) is negligible compared to the overall size of the market, her actions cannot significantly affect the price nor can her output go unsold at the prevailing market price. The externalities arising from backward and forward linkages would be fully internalized by the market price, and our previous arguments would not apply.

We must, therefore, examine whether or not it is reasonable to drop the competitive markets hypothesis in some situations. Does it make sense to assume that a seller can sell "unlimited" quantities of some good at any given price? Perhaps the answer is yes (at least as an approximation) when the economic agents involved are small relative to the overall size of the sector. For instance, with frictionless international trade, the competitive markets hypothesis may be compelling: entrepreneurs in any one country

[89] One reason why it is absurd is that contingencies are *not* verifiable, so that the required system of payments cannot generally be implemented. On a more subtle note, the very fact that payments must be made may contaminate the social interactions in the first place.

[90] These externalities are referred to as *pecuniary* externalities, as opposed to technological externalities (congestion, pollution and the like). See Scitovsky (1954) for a discussion.

may be "small" relative to the scale of the world economy. But if international trade is not frictionless, the considerations raised here may be issues of first-order importance.

In short, none of this is to say that Adam Smith and his followers were *logically* wrong, but rather to emphasize that the concept of a competitive price at which an individual can scale up all she wants to buy, sell, borrow or invest, is often an untenable ask. Moreover, even a *market* for an activity may not even exist in the first place, depending on the particular activity and externality under consideration. Such lacunae — if we can call them that — can really flip things around. For instance, suppose that there is involuntary unemployment of labor. (We'd all agree that this is not an uncommon phenomenon.) A firm's investment might then contribute to the lowering of that unemployment and so benefit other firms, not mention the unemployed workers. That would be a socially cool thing to do. But it is the other way around: having to pay for that extra labor is a *cost* to the firm. In the fairytale world of competitive, efficient markets, that involuntary unemployment wouldn't exist to begin with.

7.4.5. Breaking Bad Equilibria. Multiple equilibria or steady states arising from complementarities have a distinctive feature: *temporary* policies can have *permanent* effects. For instance, suppose that we currently located in the high-corruption equilibrium n_3 in Panel B of Figure 7.3, and that a temporary hike in fines or a large anti-corruption drive pushes the system temporarily to left of n_2. Then *even if* the anti-corruption drive is suspended or completely stopped, the system can move of its own accord to the low-corruption equilibrium. You would not see this, say, in the Solow model of economic growth: for instance, if a temporary boost to the savings rate is removed, the system will begin to revert to the old steady state.

Temporary changes can also have bad effects. Suppose that n represents the fraction of people in politics that take recourse to defending or promoting fake news. It is entirely reasonable that if more politicians take this action, then more will want to do so — perhaps the punishment for lying through one's teeth gets diluted as a jaded society gets used to these tactics. In short, there are complementarities. A "low fake-news" equilibrium n_1 can be utterly destroyed by the temporary actions of a rogue politician and his cronies who have no respect for the facts. The system can be tipped over to the right of n_2, and finally head into the bad equilibrium with everyone (or a majority) spouting fake news, described by n_3.

It is very much the same with the investment complementarities described in Section 7.4.3. Whether or not a "bad equilibrium" arises depends on the expectations that each entrepreneur holds about the others. These expectations, in turn, are driven by past history. A region that is historically stagnant may well be expected to stay that way, whereas another region that has been historically active may continue to flourish. At the same time, there may be nothing that is intrinsically different between the two regions. What is the role of policy, then?

Rosenstein-Rodan, writing near the end of the Second World War in 1943, provided an intellectual foundation for the Marshall Plan, which poured many billions of dollars into the reconstruction of post-War Europe (with one eye — or perhaps both — firmly on the so-called Communist threat). He introduced the idea of a *big push*, a policy that would simultaneously raise investment in many different sectors of the economy. Such a policy has two important features. First, it requires a massive (presumably public) investment in many different sectors of the economy at once. Second, an idea of the *quantitative* allocation of the investment across different sectors of the economy is necessary. For instance, from the perspective of the linkages argument illustrated by Figure 7.4, it would be necessary to know the correct mix of investments, otherwise there would be, say, too much investment in shipping and too little in coal.

This is a tall order on both counts. The investment is presumably immense. For most developing countries, this would require an infusion of foreign aid that is probably way too large. In post-war Europe, the Marshall Plan did provide a huge infusion of funds, and it was in the context of such possibilities, reparations in particular, that Rosenstein-Rodan made his original argument. Significantly more problematic than the size of the required investment are the informational requirements of such an exercise. It is simply not practical for the government to have the information to parcel out the investment to different sectors or for the government to have the information to run each such sector.

There is another objection to this policy that is more fundamental than either of the considerations described in the preceding text. It is that the policy does not exploit the fact that the desired outcome *is also an equilibrium*. Put another way, if incentives are provided carefully and selectively, we can partly rely on the market to correct this coordination failure.

Albert Hirschman (1958) made influential use of this idea. His proposal was simple but profound: instead of following the big push, which is akin to a policy of "balanced growth," follow a deliberate policy of *un*balanced growth. That is, selectively promote the development of certain key sectors in the economy, and as the linkages generated by these key sectors make their presence felt, the market will respond to the unbalanced situation by making the other investments spontaneously.

Such key or *leading sectors* could be chosen on the basis of number and strength of intersectoral linkages that they possess. The more or more intense those links, the greater the candidacy of such sectors to be leading sectors. The character of the link, whether it is forward or backward, also assumes some importance. Forward linkages are essentially facilitatory: they increase the viability of some other sector from the point of view of production, from the supply side as it were. Backward linkages increase the *demand* for the product of another sector; the flow of information to the benefiting firm may be sharper.

Finally, it is important to look at the "intrinsic profitability" of each sector. Suppose the export sector has far more linkages than the development of highways, peso for peso. Is this an argument for the government to make a large-scale investment in the export sector at the expense of highways? Not necessarily, and the reason for this is that export production is more likely to be spontaneously taken up by the private sector, because there are profits to be made. It is harder (though not impossible) to turn highway construction and maintenance into a profitable venture, especially in poor countries where charging profit-making tolls might be politically or economically infeasible. In this case, *the government maximizes the chances of overcoming coordination failure by investing in the least profitable activity*, provided of course that such activities have linkages as well. Thus a leading sector need not be intrinsically profitable, but it must spur other sectors that are.[91]

7.5. Expectations and History

So far, we've viewed multiple equilibria as reinforcing cycles of beliefs and actions that lock together. In this final section, we review this notion of equilibrium, and begin the transition to theories of history-dependence.

Figure 9.4 shows us two sectors, Old and New. Each person can choose to "live" in either sector. Just for simplicity, I've drawn the diagram so that the payoff from living in Old doesn't depend on how many live there: it's fixed at the height of the horizontal line. In New, there are complementarities: the return r_{new} to living in New depends positively on the number of individuals living in New.[92] Many models of complementarities can be neatly captured by this little paradigm: for instance, "joining the revolution," "selling a currency," or "adopting a particular technology" can all be thought of as "choosing to live" in a particular sector. Panel (a) captures the obviously interesting case: when no one is in New, its rate of return is smaller than in Old, and when everyone is in New, its rate of return exceeds that in Old.

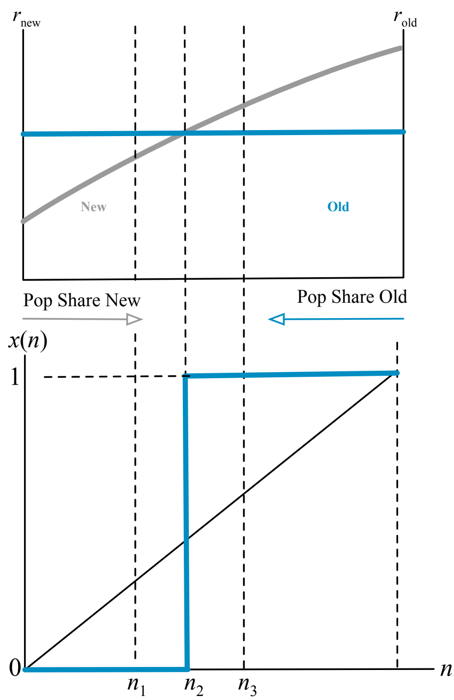

Figure 7.7. History versus Expectations

History is captured by the initial allocation of individuals across the two sectors. If that allocation lies to the left of n_2, at which the two returns are the same — say at n_1 — individuals will progressively gravitate from New to Old. This describes the failure of a new sector when there is not enough mass to keep the momentum going. Matters end with everybody in Old and nobody in New. Of course, had there been sufficient critical mass, matters would have been different. For instance, suppose that history initially put us at n_3, where the net rate of return to being in New is positive. Then there is an accelerating tempo of change, as people switch over to the new technology (or some new product or way of life, or join a revolution or sell the currency, whatever).

Now here is a different story that works entirely on the basis of expectations. Begin again at n_1. Now imagine that for some reason *everybody* believes that *everybody*

[91]This simple observation also suggests that the role of government is often to take up intrinsically unprofitable activities and cautions us to not make blanket statements such as "governments are incapable of covering their costs" without studying the sectors involved.

[92]For detailed analysis of similar models, see, for example, Krugman (1991a), Matsuyama (1991), Acemoglu (1997) and Adserà and Ray (1998).

else will be in New tomorrow. Never mind where this belief came from. Only note that *if* this belief is genuinely held by someone, she must also believe that New is the sector to be in because the return there is higher. Consequently, she will gravitate to New. However, if everybody thinks the same way, everybody will be in New, and the seemingly absurd belief is completely justified!

From this second perspective, it seems that history plays no role at all. *Irrespective* of initial conditions, there are only two self-justifying outcomes that are possible—everyone in Old or everyone in New—and that *both* outcomes are always possible, depending only on expectations.[93]

Of course, there is a serious problem with the analysis of the previous paragraph. It ignores the fact that sectoral shifts take *time*. It *is* possible that a swarm of people gravitate from Old to New, but it does not follow that the rate of return to being in that sector *instantaneously* adjusts to the new level as displayed by the diagram. While the diagram could accurate as a long- or even a medium-run description of returns, it cannot be a good description of what happens *immediately*.

Once we recognize this time factor, our belief-based argument begins to unravel.[94] Even with the expectation that everybody else will be in New, our individual will not move there just yet. She will *wait* until the rate of return in that sector climbs above the rate of return that she is currently getting in Old. But if our person thinks this way, others will as well, so people will postpone the move, and that destroys the self-justifying nature of the expectation that everyone else will indeed move there tomorrow. History makes a triumphant reentry: now *expectations* play no role and initial conditions once again must determine the entire outcome.

The essential point of the foregoing argument is that people will delay their decisions beyond what they expect others to do, so as to avoid the time lag that occurs in the build-up of returns. This tendency to delay a potentially costly move can trap the economy. This is precisely what occurred in the situations described by Rosenstein-Rodan and Hirschman: everybody wants everybody else to go first. There is no solution to *that* little problem.

I draw two lessons from this section. The first is that expectations and history are not that different, at least conceptually. Once we introduce some delay into the adjustment of rates of return, the self-fulfilling leaps of expectations become unlikely, turning those expectations into something akin to physically sluggish objects, which move through time just as other physical objects — such as capital or income — might move. A unified theory that gives pride of place to initial allocations in determining final outcomes can then apply to both "real" variables such as capital, and social-pyschological variables such as expectations.

The second point implicit in this discussion is that we can also tease out a richer theory of when expectations have more or less "historical weight." Take fashion, for instance. One characteristic of fashion is that the initial users might be thought of as "mavericks": they don't care that they are different; indeed, they *like* the idea of being different. Hence, there is no cost and possibly a gain to "going first" to New: their payoffs are different from the ones depicted in Figure 9.4 and incorporate the pleasures

[93]As usual, there is also a third allocation n_2 at which the rates of return in the two sectors are exactly equalized, but it is unstable as already discussed.

[94]See Chamley and Gale (1994) and Adserà and Ray (1998).

of being non-conformist. Once this happens, the conformist followers can follow in due course. Just how many mavericks are needed to upset a going equilibrium will also depend on how quickly or slowly the returns adjust: with fashion, they presumably move fast, so a relatively small number of mavericks will do the trick. That gives us a theory of "fashion cycles": mavericks will continually blaze the trail from one equilibrium to another, with the conformists following them.[95] History sits lightly on fashionable outcomes, because there are both conformists and non-conformists among the *fashionistas*.

On the other hand, in the case of a development trap, or technological lock-in, "going first" — or nonconformism — means taking economic losses. Unless there are entrepreneurs who deliberately run against the economic tide, because of overoptimism or social vision, perhaps, or just plain arrogance, we cannot break an equilibrium trap in the same way as old fashions are broken. Sometimes there are such entrepreneurs, but in general, history will weigh heavily on economic outcomes.

Finally, and even with sluggishly moving variables, expectations do play a role when there is some advantage to going first. One such advantage comes from the costs of entry into a new sector. It may be that earlier entrants get the better jobs, for instance, because of more years of experience. Alternatively, if we take these sectors literally as two geographical regions, then it may be that congestion increases the costs of moving later, perhaps because of difficulties in finding suitable housing. In such cases, it is possible that economic agents trade off the lower *current* returns from moving now with the higher costs of moving in the future. This might spark a process of "migration" from Old to New even though Old is currently more profitable. Sustaining this migration requires the belief that others will migrate as well. In this sense, expectations might overcome history to some degree.[96]

7.6. A Complementarity? Fertility Decline in Bangladesh

Like many countries, Bangladesh has gone through a sustained period of fertility decline, with birth rates slowing as death rates fell across the country. As death rates dropped from over 40 per thousand in the 1940s, down to well under 10 per thousand, birth rates have correspondingly followed, albeit with a lag — from around 50 per thousand to around 20 per thousand. This is just one instance of the famous *demographic transition*, a subject we shall take up in detail in Chapter 14.

Such a change is due to many factors, of course, but it is hard to ignore the massive change in *social* attitudes that must have taken place over this period, particularly with respect to the use of contraception. As in many countries, with both high and low income, contraceptive use among couples — and open discussion of such use — can be a tricky business. Might a complementarity have been at work, increasing the ease of contraceptive use among any couple as it became more widespread and therefore socially acceptable? Munshi and Myaux (2006) argue that the answer is in the affirmative:

> "Most societies have traditionally put norms into place to regulate fertility. When the economic environment changes, individuals gradually learn . . . about the new reproductive equilibrium that will emerge in their

[95] On fashion cycles, see, for example, Pesendorfer (1995).

[96] This is the argument made in Adserà and Ray (1997).

> community. This characterization of the fertility transition as a process of changing social norms is applied to rural Bangladesh, where norms are organized at the level of the religious group and interactions rarely cross religious boundaries ... Local changes in reproductive behavior occur independently across religious groups despite the fact that all individuals in the village have access to the same family planning inputs.

Their study concentrates on the Comilla district of Bangladesh, specifically a cluster of 70 villages in the Matlab *thana*, a sub-district of Comilla. These villages were covered under an intensive program in 1978, called the Maternal Child Health–Family Planning (MCH-FP) project. A Community Health Worker visited each village once every fortnight, and contraceptives were provided free of cost. Over a period of just ten years (1983–93), the total fertility rate, measured as the total number of children that a woman is expected to have over her reproductive years,[97] dropped from an unsustainably high level of 4.5 to a modest 2.9. Over the same period, contraceptive use among couples rose from 40% to 63%. These are solid changes, for which the program is certainly to be credited in part, though it must be also remembered that such changes were afoot all over the country. In any case, we are not evaluating the success of the MCH-FP project here, though the determination with which the Community Health Workers (all women) went about their task is laudable, given the strong resistance they faced from large sections of the village community.[98] Our goal here simply is to spot the complementarity — did contraceptive use *in the community* have an impact on use at the level of the couple, over and above any family planning program?

I want to show you that this is not an easy empirical question to answer. Imagine, for instance, the following sort of regression equation:

$$y_{it} = A + \gamma y_{i,t-1} + \beta x_{t-1}^{v(i)} + \eta Z_{it} + C_t^{v(i)} + \epsilon_{it}, \tag{7.2}$$

where t is the period of observation (say annual or semi-annual), y_i is 0-1 contraceptive use by individual i,[99] $v(i)$ is i's village of residence, x^v is village-level lagged contraceptive use, Z_i is a vector of individual-specific characteristics (such as age), and C_t^v is a village- *and* time-specific dummy variable. Our specific interest is in the term β, for if that were positive, it would suggest a complementarity. But how is β to be identified, given the presence of C_t^v, which is meant to be a proxy term for everything that we do not observe in specific villages over specific periods? The answer is simple: it *cannot* be identified, because all the co-movements between individual-level and average village-level use could be driven by (unobserved) variations in program intensities or implementation that vary both across villages and over time.[100] There is no way in

[97] The total fertility rate in any year is defined as the total number of children that a woman would expected to have, using the age-specific fertility rates in that year.

[98] Initial responses were extremely hostile, as the health workers were viewed as violating *purdah* and additionally espousing policy that went against existing norms of free reproduction. These norms, it should be added, were predominantly local, as neither Hindu nor Muslim scripture prohibit contraception. It must be remembered that rural Bangladesh was (and in several parts still is) a conservative society, though at the same time we see dynamic and progressive changes in Bangladesh today in all walks of life.

[99] The "individual" refers to a woman, and contraceptive use by her includes all forms that the couple might use.

[100] This is to be distinguished from including only individual- or village-fixed effects that do not vary over time, along with time controls that do not vary across villages. If we had reason to believe that this were all, we could identify β. But there is no reason to believe that such fixed effects are enough — for instance, there could be village-specific time trends.

this specification to separate out a genuine causal effect running from x to y — our sought-after complementarity — from, say, the joint impact of the Community Health Worker in a particular village over different periods of time (this latter effect would be subsumed in the C_t^v dummies).

Munshi and Myaux's workaround is to exploit the idea that while both Hindus and Muslims co-exist in the Matlab villages, their *within-group* interactions are likely to be stronger than their *cross-group* interactions. This is not the place to wonder why that's the case, or even to bemoan that fact, but to accept that this is a plausible hypothesis and remain laser-focused on the question of complementarity. But how might the asymmetry of group interactions help to answer our question? The argument is that *if* the interaction hypothesis has teeth, then we would expect village-level *Muslim* contraceptive use to affect a *Muslim* couple's contraceptive use, but less so (or perhaps not at all) the contraceptive use of a Hindu couple. (And vice versa, I hasten to add.) Specifically, think now of two regressions. For a Muslim individual i, consider

$$y_{it} = A_m + \gamma_m y_{i,t-1} + \beta_{mm} x_{t-1}^{v(i),m} + \beta_{mh} x_{t-1}^{v(i),h} + \eta_m Z_{it} + C_t^{v(i),m} + \epsilon_{it} \tag{7.3}$$

where the m and h labels should be self-explanatory. (And we can write down a similar equation for a Hindu individual.) The hypothesis that we now work with is that cross-religious influences are minimal, so that we expect β_{mm} to be positive, but we expect β_{mh} to equal zero.

But why don't we just have the same problem as we had before? Why can't the village-level, religion-level *and* time-specific dummies still confound all that is there to be explained, just as they did in equation (11.18)? The answer is that a confounding process of this sort would have to *simultaneously* account for not just the positive β_{mm} and β_{hh} but they would *also* have to account for the zero cross-group terms β_{mh} and β_{hm}. That is a tall order which would require, in particular, the two sets of effects $C_t^{v,m}$ and $C_t^{v,h}$ to be uncorrelated with each other. That means that what is happening under the MCH-FP program in a specific village (and at a specific time) for Hindus, is uncorrelated with the same program in the same village, and at the same time, for Muslims! Logically speaking, this juxtaposition of events is possible — stranger things have happened in our universe. But at some point we must also give in to common sense. In the words of Munshi and Myaux, "it is difficult to imagine that potential omitted variables such as program effects or economic change are completely uncorrelated across religious groups within the same village."

Table 7.1 summarizes the key findings. The dependent variable is individual contraceptive use, which is a binary decision in any period, and the independent variables are lagged contraceptive use in the same village, *separately* neasured among those belonging to the individual's own religion group and the cross-religious group. The third independent variable controls for the individual's own lagged use of contraception in the previous period. The coefficients for the pwn-group effect are positive and strongly significant, suggesting that village use does influence individual use. Remarkably, the cross-group effects are insignificant, suggesting that community effects flow through the same religious group, but not through the cross-religious group.

The use of religious groupings serves two roles. The first is of intrinsic interest: if there is a complementarity, is it mediated by cultural or religious similarity. Table 7.1 strongly indicates that the answer to that question is in the affirmative. And the

	All villages		Minority > 5%		Minority > 15%	
	Muslims	Hindus	Muslims	Hindus	Muslims	Hindus
Lagged contraception	0.217	0.161	0.193	0.169	0.207	0.168
(own-group)	(0.013)	(0.014)	(0.016)	(0.017)	(0.018)	(0.020)
Lagged contraception	0.008	0.009	0.007	0.024 -	0.001	0.019
(cross-group)	(0.006)	(0.007)	(0.011)	(0.016)	(0.013)	(0.024)
Lagged contraception	0.698	0.712	0.704	0.710	0.706	0.717
(self)	(0.003)	(0.005)	(0.004)	(0.005)	(0.004)	(0.006)
R^2	0.513	0.559	0.520	0.558	0.521	0.565
No. Observations	139,875	43,101	79,927	29,771	49,730	20,756

Table 7.1. The contraception decision for Hindu and Muslim couples, viewed as a linear probability function of the listed variables. Each sub-column marked "Muslim" contains the Muslim regression, and those marked "Hindu" the Hindu regressions. Standard errors robust to heteroskedasticity and correlated residuals within each village-6-month-period are in parentheses. Source: Munshi and Myaux (2006, Table 3).

second role, as already discussed, is statistical: by showing that the cross-group effects are zero while the own-group effects are positive, it suggests — again strongly — that these outcomes are not driven by some unobserved effect that is specific to village, religion and time. For if they were, those terms would have to be uncorrelated within the same village, which is not a logical impossibility, but does appear to fly in the face of common sense.

These findings complete our brief description of the Munshi-Myaux study. It lends some credence to the norm-based fertility transition model. The community prevalence of contraceptive use appears to reinforce individual choices, suggesting a complementarity in social norms. Finally, a program like the MCH-FP does not need to last forever, if this kind of theory is correct. Once individual and social norms lock together in the new equilibrium, the program can be removed without the entire process unwinding of its own accord.

Appendix: Joining the Revolution

On December 17, 2010, a 26 year old Tunisian street vendor named Tarek al-Tayeb Mohamed Bouazizi set himself on fire as a protest against the harassment that he suffered at the hands of municipal officials. This single act served as a catalyst for a huge wave of political demonstrations in Tunisia, leading to a change of the Presidency in January 2011, and setting the stage for the cascade of demonstrations and protests across North Africa and the Middle East, known today as the Arab Spring.

It is, of course, absurd to imagine that a single man's act, brave and sad though it may have been, was fundamentally responsible for the outpouring of discontent, for which there were systematic and deep reasons throughout the region. Rather, what it was was a signal around which individuals could rally in a coordinated way (aided by the widespread use of social media), thereby propelling societies from seemingly peaceful configurations to a state of utter revolution. My goal is to show you how you can make sense of some of this by using the ideas of complementarities and multiple equilibrium developed in this chapter.

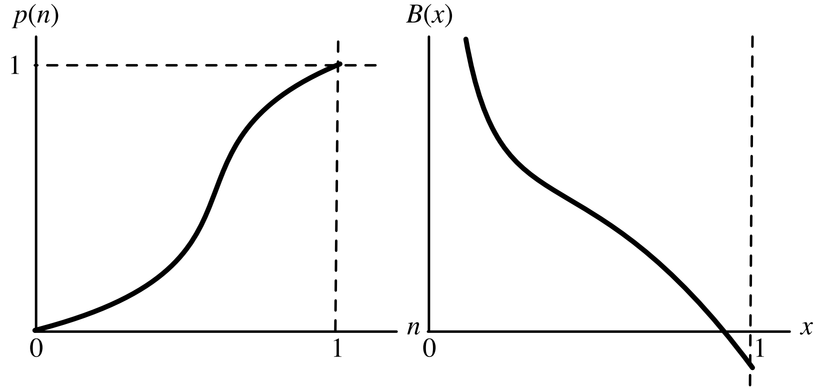

Figure 7.8. The Success Probability and Benefit Functions.

And when I say "some of this," I mean just that: there are so many complexities about the dynamics of revolution that we just do not understand. But that shouldn't intimidate you into thinking that you can't understand *any* of it. (That's a common defeatist approach adopted by many who are scared of formal reasoning and therefore hide behind their scorn of simple yet insightful modeling devices.)

So, let's think of a revolution as an equilibrium in a (potentially) multiple-equilibrium context. Begin by supposing that there are lots of people, all arrayed on the interval $[0,1]$, and that a person is just indexed by her location on that line ("person x" being the person at "location x" in $[0,1]$). Each person just has one decision to make, to join the revolution, or not.

Joining the revolution comes with a potential benefit, and a potential cost. The benefit is that if the revolution succeeds, an individual who joined can say that "she was there," that "she helped topple the regime," and this has some value B (there may be real payoffs as well, such as political or social recognition, or economic rewards). The cost is that if the revolution fails, then there will be a price to pay, a loss L (incarceration or worse). Success is an uncertain prospect, and we assume that it depends on the number of people n who join the revolution, which we denote by n. Write the probability of success as some increasing function $p(n)$.

This is the simplest bare-bones model that we can write down, but it leads to some interesting variations. For instance, the benefits and costs of joining may be different for different individuals. An easy way to introduce heterogeneity is to think of the benefits as person-dependent (you can do the same with L as well). In this variation, think of B not as the same across people but as a function $B(x)$, where $B(x)$ is the success benefit accruing to a person with index x. Without losing any generality, we can arrange people on $[0,1]$ so that the benefits decrease as the index increases. So the revolutionary firebrands are on the left, where they belong, with $B(x)$ being large when x has a small index, while people close to an index of 1 have the lowest values of B, which may even be negative. Figure 7.8 depicts some typical shapes for the functions $p(n)$ and $B(x)$.

I want us to construct the *complementarity map*: the function described in the main chapter that links the number of people expected to join, to the number of people who *want* to join (under that expectation). To do so, note that if n people are expected to join, then a person x will want to join if his expected payoff from doing so is positive; that is, if

$$p(n)B(x) - [1 - p(n)]L > 0.$$

Moving terms around, we can write this a bit more carefully: *join* if

$$B(x) > \frac{1 - p(n)}{p(n)} L, \tag{7.4}$$

don't join if the opposite inequality holds, and be indifferent if equality holds.[101] That gives us the required complementarity map as introduced in Section 7.3: for each "anticipated n," throw in all the people for whom (7.4) holds, and any or all of the people for whom (7.4) holds with equality. That mapping tells us how many people *want* to join assuming that n people are *expected* to join.

In an equilibrium, just as we equate supply with demand, the expected and actual numbers must match. Put another way, equilibria are given by the intersections of the complementarity map with the diagonal or 45^0 line; see Panel B of Figure 7.3. That panel shows a general mapping which may wiggle and intersect the diagonal one or more times, but what we can be sure of is that such a map cannot ever decline: if more people are expected to join the revolution, then the probability of success climbs, and so at least as many people as before want to join. As already noted, this nondecreasing nature of the mapping is fundamental to a situation of complementarities.

Now let's look at some special cases.

Case 1. Multiplicity in Homogeneous Societies. Suppose that benefits and losses are the same for everyone. Then the joining condition (7.4) reduces to

$$B > \frac{1 - p(n)}{p(n)} L. \tag{7.5}$$

Observe that the right-hand side progressively decreases as n goes up. In fact, if we assume that $p(0) = 0$ (no participation implies a zero chance of success), and $p(1) = 1$ (full participation ensures success), then the right-hand side of (7.5) declines all the way from infinity to zero. Whenever n is low enough, then, no one will want to participate, and once n crosses the threshold such that (7.5) begins to hold, *everyone* will want to participate. So the complementarity mapping looks just as in Panel A of Figure 7.9. The very same society exhibits an equilibrium with utter quiescence ($n_1^* = 0$) and one with full-throated participation ($n_3^* = 1$)!

Case 2. No Multiplicity. We noted in Section 7.3 that complementarities need not always imply multiplicity. We illustrate that here. Suppose that individuals differ in their perceptions of revolutionary benefits, and the function $B(x)$ has the special form

$$B(x) = \frac{1 - x}{x}.$$

The idea is that $B(x)$ is extremely high for some individuals (take x close to zero), while for others revolution has almost zero participatory value (take x close to 1). We

[101] Careful reader: we are assuming that everyone is risk-neutral; the analysis is very similar and just a bit more notationally challenged if people are risk-averse.

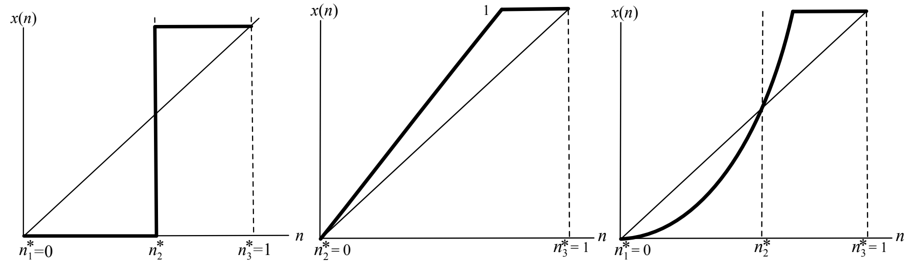

(a) Case 1 **(b)** Case 2 **(c)** Case 3

Figure 7.9. COMPLEMENTARITIES IN JOINING THE REVOLUTION.

also choose a special form for $p(n)$:

$$p(n) = n,$$

which says that the probability of success is the same as the fraction of people who participate.

Now for every n, the value of the complementarity map is the largest possible solution to the joining condition (7.4), so that if $x(n)$ denotes this value, it solves the equation

$$\frac{1 - x(n)}{x(n)} = \frac{1 - n}{n} L,$$

provided that such a solution exists with $x(n) \leqslant 1$. Otherwise (7.5) holds with strict inequality throughout, and $x(n) = 1$. Solving for $x(n)$ as a function of n, we see (after some manipulation) that

$$x(n) \text{ is the maximum of 1 and } \frac{n}{n + (1 - n)L} \tag{7.6}$$

The mapping in (7.6) isn't very hard to figure out. When the loss parameter L equals exactly 1, it reduces to something very simple: $x(n) = n$ for every n. For any smaller loss parameter, the right-hand side of (7.6) exceeds n, so that $x(n) > n$ except at the points $n = 0$ and $n = 1$. For any larger loss parameter, the right-hand side of (7.6) falls below n, so that $x(n) < n$ whenever $n > 0$.

We therefore have a unique equilibrium for every value of L, unless L happens to be exactly 1 which is a knife-edge outcome.[102] When $L < 1$, a revolution is "inevitable": benefits are always high relative to costs, and the only stable equilibrium involves $n = 1$. Panel B in Figure 7.9 illustrates this subcase. When $L > 1$, so that costs are now high relative to benefits, the participants are never numerous enough, and the revolution can never happen. There are still complementarities, of course: the complementarity map slopes upwards, but benefits fall off as our person index becomes larger, and that slows down the new joiners at a pace fast enough so that complementarities are never strong enough for multiple equilibrium. (We haven't pictured this case graphically but you should draw it for yourself.)

[102]In this knife-edge case, every joining rate is actually an equilibrium, but it really isn't interesting as the parameters really have to match in an entirely coincidental way.

I've shown you this case to drive home the point that complementarities may be *necessary* for multiple equilibria of this sort, but they are not always sufficient. Too little complementarity will imply that the mapping will never build up enough steam to cross the diagonal. But too "enthusiastic" a response will imply that the mapping with *always* be above the diagonal. These are precisely the two situations that occur in this example, and so we always have uniqueness.

Case 3. Multiplicity, Sometimes. In our last example, we keep the same benefit function as in the previous case, but we change the probability function so that it is more sensitive to a larger number of people. That is, we keep $B(x) = (1 - x)/x$ but now assume that

$$p(n) = n^2.$$

In this case, using (7.5) again, the complementarity mapping is given by

$$\frac{1 - x(n)}{x(n)} = \frac{1 - n^2}{n^2}L,$$

subject as before to the upper bound $x(n) = 1$, so that

$$x(n) \text{ is the maximum of 1 and } \frac{n^2}{n^2 + (1 - n^2)L} \tag{7.7}$$

As in the previous example, it is useful to distinguish between $L > 1$ and $L < 1$. *At* $L = 1$, the mapping in (7.7) reduces to $x(n) = n^2$, with only one stable equilibrium: at 0. And the same is true for all higher losses $L > 1$: $x(n) < n$ throughout and the revolution can never get off the ground. This case resembles the unique equilibrium of the previous example.

With $L < 1$, matters are different. If you remember, in the previous example this was associated with $x(n) > n$ throughout. But that wont happen here, and the reason it won't happen is that the probability of success is relatively unaffected by small values of n: the function $p(n) = n^2$ "builds up" at high values of n, but is (relatively) dormant at low values. Therefore, at low expectations, $x(n)$ falls short of n and there is indeed a stable equilibrium outcome at zero. On the other hand, $x(n)$ crosses the diagonal and exceeds n for higher values of n, the exact value of that threshold depending on the value of L. There is another stable equilibrium with full participation. This subcase with multiplicity is very similar to our first example. Panel C of Figure 7.9 illustrates.

A Final Remark. Our example, with its stark coexistence of quiescence and revolutionary fervor, is perhaps the starkest model of the Arab Spring. Think of a disaffected society, but one in which individuals fear reprisals in the event of protest (this is L). As long as everyone feels that fewer than n_2^* protesters will be on the streets, *no one* will be on the streets, and from the outside that society will look comparatively peaceful. But once an incident carries sufficient symbolic potency, such as the death of Bouazizi, the expectation of widespread protest can serve as its own justification: the system tips to the right of n_2^*, and from there on, to the revolutionary equilibrium n_3^*. Bouazizi's death had no intrinsic impact on the parameters of society, but his death was not in vain.