1

Introduction: Marginal Utility Matters

1.1 THREE VIEWS OF DECISION-MAKING

In a textbook that capped a generation of work on axiomatic preference theory, David Kreps introduced Totrep or "trade-off talking rational economic person" to ease students into the mathematical models they would need to master.[1] Totrep became a celebrity, by the standards of fictional economics personalities, and served as an exemplar of the economic agent who must choose among alternative actions. Curiously the reader never learns if Totrep can pin down the marginal trade-offs that economics is famous for. Totrep's preferences must satisfy the classical axioms of rationality that hold that all pairs of alternatives can both be judged and judged consistently, but it remains open whether Totrep can determine the marginal value of one good in terms of another.

The classical axioms of rationality do not require agents to make judgments of the form "I am willing to accept x units of good 2 for a small amount of good 1 and to give up x units of good 2 to receive the same small amount of good 1." These are the judgments that underlie the first graph drawn on the blackboards of Econ 101, the smooth indifference curve that pictures an agent's marginal trade-offs of one good for another. One of the hallmarks of economics is therefore missing from the mathematical model of Totrep's preferences.

The divide between Totrep and the intro economics classroom mirrors the grand development of neoclassical economics, where two views

[1] The textbook, Kreps (1988), explains that Michael Harrison has parental rights over Totrep.

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

of economic decision-making have dominated the intellectual landscape and divide it into two eras of roughly equal length. In the origin story of the indifference curve in the late 19th century, marginal utility played the lead: economic agents use the pleasure delivered by increments of various goods to figure out which combinations of goods will keep them at the same level of satisfaction. For the next 75 years or so, in the marginalist period of neoclassical economics, agents were accordingly modeled by utility functions with derivatives that represent the agents' marginal utilities. The smooth indifference curve is the perennial survivor of this epoch and, to this day, the smooth indifference curve makes the most sense when it is built from the ground up by agents who weigh the increments of some primordial benefit that different options can deliver.

For the last 75 years, beginning shortly after World War II, a more spare model of rationality has ruled the seminar rooms of economic theory: it requires an agent's preferences to satisfy two axioms, completeness and transitivity, and further assumptions in more specific choice contexts. The smooth indifference curve appeared difficult to defend to the best and the brightest of postwar economic theory; it was also unnecessary for the Arrow-Debreu agenda that dominated economic theory in the initial decades of this era. The existence and optimality of competitive equilibria were the crown jewels of economic theory and, for these results, marginal utilities and marginal rates of substitution are irrelevant. Since it seemed to serve no purpose, the smooth indifference curve was abandoned by those theoretically in the know.

Much of the economics profession paid little attention to the changing of the guard. The everyday models of economics continue to rest on smooth indifference curves and differentiable utility functions, and consumer optimization is still explained to undergraduates with the story that agents equate the marginal utilities of their expenditure on different goods. The transition in economic theory also passed unnoticed in the outside world. In the public imagination, economics comes down to the maxim that "everything has its price": agents will trade away anything of value if offered enough in exchange. While this saying is something of a caricature, an agent with smooth indifference curves is remarkably malleable: if after buying positive amounts of two goods, the relative price of good 1 in terms of good 2 were to rise even slightly then the agent would happily trade away some quantity of good 1. A readiness to substitute and trade goods remains a benchmark of economic orthodoxy. But this flexibility does not follow from the axioms that describe Totrep's decisions.

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

There is a third position to consider. When goods do not deliver different quantities of a common homogeneous benefit, agents may be of several minds about the trade-offs they confront. Individuals can then conclude that their options are incomparable and that they are unable to come to a preference judgment: their preferences are incomplete. Incompleteness does not imply that an agent has somehow fallen prey to irrationality; the incomparability view challenges the claim that the rational pursuit of one's interests requires an agent to form preferences. Agents must still choose, of course, even when they cannot figure out which options are best. Whether facing simple or complex choices, between apples and oranges or between detailed state-contingent plans, agents may conclude that all of their conflicting attitudes must be in agreement to approve a change over the status quo or a customary decision. Or they may resort to the safest course of action, say, the plan that makes the worst-case outcome as desirable as possible. These and other choice strategies that agents turn to when they cannot form preference judgments overshadow the pleasure calculations that economists in the 19th century, eager to apply calculus, imagined to be dominant.

When agents cannot form preferences, the options they do not know how to compare cannot be grouped into conventional indifference curves: if an agent cannot compare alternatives *a* and *b* then an improvement to *a* need not make it superior to *b*. When agents resort to safe options, it may be possible for their choices to be modeled by ordered families of indifference curves, but those curves will not display the smooth trade-offs we expect of Homo Economicus. For example, the marginal value of a good might fall discontinuously as it crosses the threshold of consumption that an agent regards as safe. In both scenarios, the smooth indifference curve disappears.

Agents who lack preference judgments cannot make arbitrary choices without jeopardizing the goals such as greater material wealth that they can identify. Sticking to the status quo is the most obvious way for agents to eliminate those dangers. Seen in this light, some of the characteristic findings of behavioral economics no longer appear as inexplicable outbreaks of irrationality. Status quo bias and kindred patterns of choice lay out exactly the decisions that individuals without preferences should take to safeguard their interests. The verdict of the economics profession is that the behavioral evidence has toppled classical rationality as a positive theory of decision-making – despite its persistence in economic theory. But if the incompleteness of preferences lies behind the manifold violations of standard choice theory, then a unified explanation of economic decision-making must go beyond the empirics of behavioral economics; a reformulation of rationality is needed.

This book will ask and answer basic questions. Which of the three views of economic decision-making is correct? Can agents always smoothly trade off disparate benefits? Why is there a discrepancy between contemporary economic theory, which has dropped the smooth indifference curve and the differentiable utility function, and the routine work of economics? Do agents obey the narrower axioms of rationality that economic theory currently backs and, if not, are they acting irrationally? How do markets perform when agents cannot make smooth trade-offs? And can government policymaking be decisive when the individuals in society are not?

It is common for economists to view the differentiable utility function as a technical convenience, not a statement of principle. In combination with the convexity of preferences, a differentiable utility allows an agent to be modeled by a system of first-order conditions, the solution of which will normally identify unique utility-maximizing demands. If instead an agent cannot form a complete set of preference judgments and thus cannot be represented by a utility function, demands are not as easy to characterize and there are multiple ways to define optimization, a morass that economists would prefer to avoid. When preference judgments are complete but utility functions fail to be differentiable, even less seems to be at stake; with some tweaks to the standard toolkit, nondifferentiable utilities can be maximized almost as easily as differentiable utilities. After going through the ritual undergraduate exercise of discovering that the demand functions for Leontief utilities appear to be well-behaved, economists mostly leave the nondifferentiable utility function behind.

The capacity of agents to trade off benefits smoothly in fact lies at the heart of conventional economics: although the Totrep axioms may omit any mention of trade-offs, the character is aptly named. But to see what trade-offs accomplish in economics, we cannot simply accept the criteria of successful model-building that the present era of economic theory has set for itself. The main results of decision and general equilibrium theory, not surprisingly, meet the tests of theoretical consistency that those traditions have laid out. In the theory of individual behavior, we instead need to examine whether rational self-interest in fact requires agents to make choices that obey the classical axioms of rationality. And we must look beyond individual optimization to the system-wide features of economic models that depend on smooth trade-offs but that general equilibrium theory has glossed over.

I will argue that the absence of smooth trade-offs leads to challenges that cannot be resolved by existing theoretical means. Before previewing this claim, let me underscore that I am not advocating a revival of the old-time religion. The smooth trade-offs and indifference curves of early neoclassical economics provided an internal theoretical coherence that the second era of economic theory has not been able to match. In terms of empirical validity, however, those assumptions and the marginal utility mechanics that lay behind them were failures and later economists have been understandably embarrassed by them. Smooth trade-offs lie at the heart of economic analysis but not of economic reality. I therefore back the third horse.

An agent that cannot pin down a marginal trade-off between goods can usually be described by a set or band of margins or supporting prices: an incremental increase in the consumption of a good will have a strictly smaller value (in terms of a comparison good) than an incremental decrease. This multiplicity of margins or valuations can be systematic, occurring not just at isolated points but at many or all consumption bundles. I will not however assume at any point in this book that agents are incapable of judging all trade-offs between goods. Agents will for example agree to part with a unit of a good when offered enough of another good in exchange. What will be missing are the marginal trade-offs and valuations that economic analysis relies on to rule over market prices and single out which government policies are optimal.

Neoclassical economics has from the outset exaggerated the importance of substitution in consumption. Economic agents do make trade-offs in consumption based in part on their preference judgments. But the magnitude of substitution may not be great enough to buffer an economy from shocks and the gains from trade that exploit differences among agents' valuations can be small. If you are seeking an explanation of the wealth of capitalist economies or of its fluctuations, substitution in consumption is not the right place to look.

The well-defined marginal rates of substitution that stem from smooth indifference curves once provided the go-to explanation of why relative prices do not move erratically through time. If instead agents are resistant to substitution and stick to particular patterns of consumption then demand will be relatively inelastic and small changes in endowments can lead relative prices to spike or plunge – a small contraction in the supply of power from the electrical grid will cause its price to jump. The neoclassical invention of smooth trade-offs assuaged these worries: the willingness of consumers to make marginal substitutions will dampen the volatility scenarios.

Once the differentiable utility function lost its standing, economists had to find an alternative argument for the stability of relative prices over time. The answer that descended from high theory was that the endowments that generate volatility are highly unlikely to occur. This account requires an economy to begin de novo at every date with a new stock of goods and therefore does not apply to societies where goods are produced. When production is present – and enough time passes for production to affect output levels – the absence of smooth trade-offs will again lead to erratic relative prices. Production can also deliver a better explanation of what curbs price volatility: firms can transfer resources across time to tamp down the price swings that unstable individual valuations can generate.

In normative economics, determinate marginal rates of substitution play an equally pivotal role: they underlie the decisiveness of the dominant concepts of economic efficiency, both social welfare maximization and Pareto efficiency. When in contrast agents' marginal valuations are ill-defined, a wide range of policy decisions will qualify as efficient. In public goods decisions, about environmental quality for example, agents consistently declare the harm done by an incremental fouling of the environment to dwarf the value gained by an incremental clean-up. A cost-benefit test will then fail to discriminate effectively: substantial intervals of environmental quality levels will pass the test. Applied welfare economics has avoided reckoning with this paralysis by ignoring, when possible, the ample evidence that agents wield bands of marginal valuations. For the practically minded economist, the way forward has been instead to employ the smallest valuations that agents report. This footwork lets the throughput of policy recommendations flow unimpeded, but that advice will be biased against public goods.

In the welfare parables of general equilibrium theory, efficiency in an exchange economy requires there to be price lines with a common slope that *support* (are tangent to) the sets of bundles that agents prefer to their own consumption. But if the smooth indifference curve is absent and is replaced by a set of margins, the discriminatory power of this requirement collapses. Economics then loses its role of showing how to fine-tune government policies. As in the case of public goods, many and sometimes every allocation will qualify as optimal and the pursuit of efficiency will therefore lead to few nonvacuous policy recommendations. If, say, an externality appears no policy response may be called for and any policy change that has even a minute impact on relative prices will usually fail to qualify as an efficiency improvement. When no policy can be dismissed as inefficient – even policies that every economist would judge to be distorting – economics becomes useless as a policy guide.

As with volatility, production rather than exchange paves the way forward. Increases in productivity, not the alignment of hypothetical indifference curves, drive the growth of social wealth. Technological change at the same time leads to sharp changes in the relative prices of factors and consumption goods and thus swings in the distribution of income. Economists tend to gloss over this conflict. The harm done by opening industries to productivity-enhancing competition either disappears into the black hole of distributional value judgments or is met with reassurances that injured parties can be made whole by carefully engineered compensation payments. Compensations accordingly became a centerpiece of how economic theory has dealt with the diverse repercussions of economic change. Under the best of circumstances, compensationism requires formidably detailed information about agents' preferences and trades. But with incomplete preferences, agents' decisions need not reveal their preferences; when agents are unable to judge and go for the safe option or the status quo, they may not view their selections as superior to their other alternatives. Discovering the information needed for compensation payments then becomes much harder.

The solution I propose provides an alternative design and rationale for policymaking that omits any mention of preferences. Compensations should give agents the opportunity to undertake the same trades they made previously; the policies that emerge then will not face any credible objections. When compensations based on ex ante trades are infeasible, policymakers can instead modulate the relative price changes that can undermine the fortunes of agents. A government moreover can constrain the relative prices facing households while still incentivizing efficient production via the prices that firms face; policies can thus both harness the efficiency gains of competition and avoid the price changes that inflict harm. This alternative approach can free welfare analysis from the apparent logjams where every policy option qualifies as efficient. Policymakers do need not acquiesce to the arbitrary programs and practices they inherit. A government need not stand by, for example, when technological change and international trade wreak harm on those caught on the losing side of dynamic comparative advantage; and the government's policy responses do not have to slow economic growth.

These economy-wide repercussions of missing preference judgments form Part II of this book.

Part I addresses individual decision-making. I begin by setting the 19th century dogma of marginal utility against the more parsimonious model of rationality that succeeded it in the mid-20th century, Robinson Crusoe comparing the gains of an extra minute gathering bananas or spearing fish against the completeness and transitivity axioms that model Totrep's preferences. Economic theory did not emerge unscathed from this transition. Utility and marginal utility not only allowed agents to pin down marginal trade-offs and thus find optimal decisions; they also showed that individuals can determine which of any pair of options is the better choice.

Once doubt was cast on marginal utility and pleasure-seeking, the larger principle that agents can order their options lost its justification. Without an explanatory psychology to fall back on, contemporary decision theory has remained silent on why an individual should satisfy the most basic axiom of rationality, the completeness assumption that individuals can form a preference judgment between any pair of options. In the face of this lacuna, the standing of completeness as a benchmark of rationality begins to wobble.

Agents find many decisions easy to judge. Everyone has favored clothes, foods, pastimes, and so forth. Agents will also readily come to preference judgments when choosing the best means to a known end – as when a worker opts for the highest-paying job. And difficult choices can sometimes be reduced to simpler alternatives that are easier to weigh. If say you compare two job options with disparate features – one offers higher salary and a longer commute – you may find the decision straightforward once you realize that the high-salary option will implicitly pay a trivial wage for your drive to work. But even in the simplest cases, you may not be able to pin down the marginal trade-offs essential to economics: you may reject a small return to a long commute but not be able to form sharper judgments.

Making matters worse, the comparisons that the agents of modern economics need to make are herculean. Jevons posited agents who faced small self-contained comparisons – how to allocate food on an ocean voyage for example – and he did not suggest that agents could compare disparate types of pleasure. The agents that live in current-day economic models, in contrast, must compare detailed state-contingent plans over a lifetime of consumption. But incomplete preferences do not have to stem from the complexity of decisions or a shortage of information. A well-informed agent facing clear alternatives might not have a best choice: there may be no bedrock of true preference that lies below.

Economists have a well-rehearsed answer to claims that agents cannot form a preference between options: make agents choose and declare that their choices reveal their preferences. I will show, however, that the "revealed preferences" that emerge from such exercises will not satisfy the classical axioms of rationality, even when agents follow decision rules that never lead to dominated outcomes. An agent in short can be rational without satisfying the axioms that supposedly characterize rationality. As a body of empirical predictions, classical rationality was therefore bound to fail, though it has taken decades of documentation for that failure to be recognized.

Our era of economics has responded to the empirical defeat of rational choice theory with a shrug: "who cares what is labeled rational, what matters is behavior." This book lays out two replies, given in embryo in this chapter, first that only the rational pursuit of self-interest can explain the apparent anomalies of real-world decision-making, and second that the appraisal of social institutions depends on a valid classification of actions as rational and irrational.

There is moreover an alternative to a divorce between rationality and behavior: characterize rationality with greater precision. When individuals face static one-shot decisions, the amendments needed are relatively minor. Instead of choosing options superior to all alternatives, agents must select undominated options. Since incomplete preferences reduce the opportunities for one decision to dominate another, decision-making then becomes easier, and indeed agents may confront an embarrassment of optima. While not a wholly new phenomenon – an agent with weakly convex indifference curves can occasionally face a budget set with more than one optimum – the multiplicity that comes with incompleteness is far-reaching. Despite this difference, the mischief that incomplete preferences can cause for the static demand for goods is limited. After all, preference theory has never been able to deliver on its promise of foundations for the downward-sloping demand function; as Becker (1962) pointed out long ago, it is easier to generate well-behaved demands from irrational behavior – specifically choices uniformly distributed on budget lines – than it is from utility maximization.

The terrain is different when agents face dynamic sequences of decisions. Individuals with incomplete preferences must then take care to avoid manipulation. The simplest way for an agent to steer clear of risks is to refuse any offer to switch to an option that the agent cannot judge relative to the pre-existing status quo, a rigidity that stands in contrast to the agents with smooth indifference curves who adapt their consumption to every relative price change. Status quo bias, the endowment effect, loss aversion – the iconic choice strategies of behavioral economics – thus emerge as validations rather than breaches of rationality. If we define agents' interests by the outcomes their decisions yield, rather than the axioms popular in economic theory, we can predict more accurately which economic behaviors will persist and which self-interest will chip away.

Incomplete preferences also resolve the puzzle of why agents so frequently fail to find a dominant option from a set of alternatives. With classical preferences, indifference is a fluke event but with incomplete preferences, an inability to judge alternatives arises systematically. In fact, once the door is open to incomplete preferences, it becomes even harder to attribute waffling to indifference: in models where agents can both be indifferent between some options and unable to form preferences for other options, indifference comes near to disappearing altogether.

The three views of decision-making adopt conflicting positions: smooth trade-offs determined by marginal utilities versus rationality axioms on preferences versus agents that cannot always come to preference judgments. The history of the contest between the first two views was written by the victors. The psychology of pleasure-seeking peddled by the early neoclassical economists appeared pointless to their mid-20th century successors and stood in the way of their scientific aspirations. Not only did the new orthodoxy hold that individual decision-making could be based on axioms of rationality rather than utility, but the smooth indifference curve appeared to be unnecessary. As I have mentioned, the features of competitive markets identified by the Arrow-Debreu model, the unifier of postwar economic theory, did not turn on marginal utilities or any of the other derivatives in the early neoclassical arsenal.[2] The labeling of neoclassical economics as marginalist was from this vantage simply a mistake. While the rear-guard defenders of utility theory put up little effective resistance, a nagging anxiety has persisted that something was lost when marginal utilities and the smooth indifference curve were dropped from the theoretical canon. One of my jobs will be to articulate this worry. We will see that the marginalist label captures part of the truth: when individuals do not substitute the satisfaction of goods at specific marginal rates, they can instead be modeled by sets of such margins.

[2] See Hahn (1961) for example.

The theory of rational choice has in turn had to face a fresh set of tribulations – the pile-up of evidence that agents do not base their decision on unified orderings that rank all of their options. Rather than carrying on its own rear-guard struggle against behavioral economics, defenders of rationality would be better served by rethinking what a theory of rational choice is supposed to accomplish.

1.2 RATIONALITY IN ECONOMIC THEORY

The three views of decision-making do agree on one important point. Each theory classifies decisions as rational or irrational, not necessarily in the narrow axiomatic sense, but with the broader meaning that some but not all decisions an agent can take will serve that agent's goals and interests. When preferences are incomplete, some decisions can still be superior to others: an agent can succeed in forming some preference judgments. Each theory also assumes that agents will not undertake an action when a preferred action is available. Without this common ground, the dialogue in this book between theoretical camps would be impossible.

The two most frequent defenses of rationality argue that it provides reasonably accurate predictions of behavior and that competing theories that predict irrational actions, even if correct at a moment in time, are prone to failure as agents learn how to better serve their interests. Without passing judgment on these arguments, let me propose a third. Even if agents are bent on acting irrationally and disregard all attempts at persuasion, the determination of which actions serve their interests and which do not sets a research agenda. With a classification in hand, we can check empirically whether agents follow rational courses of action and how resistant to change their irrationalities are. When we come across an irrationality – an individual whose preferences are reversed by an immaterial redescription of options, a firm that refuses to adopt a technology that makes more money – does it become less likely over time? If not then an analysis of the irrationality should look for the constraints, both institutional and psychological, that prevent agents from correcting their errors. The most routine question in economics asks: Why has some apparent profit opportunity gone unexploited? Why for example don't perpetually overbooked restaurants charge more? Why did the British steel industry in the late 19th century stick with outmoded technologies?[3] The next

[3] See Kahneman et al. (1986), Becker (1991), and Karni and Levin (1994) on the first example and Temin (1966) and McCloskey (1973) on the second.

stage of research investigates whether new conventions and institutions emerge that can remedy the irrationalities, either with or without the cooperation of the irrational agents. For example, dynamically inconsistent agents may be unable to stop themselves from spending wealth that they had earlier decided to save for the distant future. Do savings options then appear that limit the opportunities for agents to thwart the plans laid down at earlier dates by their more prudent selves?[4]

This third argument will be more relevant for this book than the first two. I will argue that an absence of preference judgments will lead agents into behaviors that look irrational from the perspective of the classical axioms of rationality, for example, status quo bias and willingness-to-accept/willingness-to-pay disparities. Although these behaviors lead to intransitive revealed preferences, they in fact shield agents from harm when they have no preference judgments to guide them. So, if my argument is correct that an inability to compare options does not hurt agents or otherwise qualify as irrational then the agenda of economics should shift. Economists should stop trying to find out whether status quo bias will fade away when agents are given time for reflection: as an expression of rationality, there is no reason why it should. The soundness of this program however turns on the accuracy of my characterization of rationality.

A fourth and final argument in favor of rationality is that it brings diagnostic clarity. Economics analyzes the flaws of behavior and social institutions by first postulating an ideal world free of frictions, where the rational expectations and actions of agents lead to efficient outcomes, and then examining the effects of adding a candidate distortion to this hypothetical world. If, say, you want to argue that a pollution externality leads to inefficiency, your case will be more difficult if you assume in addition that agents act irrationally. With two sources of trouble, assignment of blame for the damage done can be ambiguous. The two malefactors can also offset each other: if agents irrationally underutilize a pollutant (from the point of view of their self-interest), they could unintentionally eliminate the inefficiency normally triggered by the externality. The same principles apply to individual decision-making. If you want to show that violations of independence in the theory of choice under uncertainty can expose agents to manipulation then you should not also assume that preferences are intransitive: those agents will already be manipulable.

4 Laibson (1997).

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

But for this diagnostic method to work, the ideal world must portray the interests of agents accurately: otherwise the count of distortions will not be correct. So, once we relinquish the fiction that agents can always rank options and judge marginal trade-offs, the dangers and inefficiencies of economic life need to be recast. To analyze the world's many potential hazards – externalities, intransitive choices, missing markets, the incompleteness of contracts, asymmetries of information – one should not suppose that indifference curves are smooth or that preferences satisfy the classical axioms of rationality. For example, the explanation in Chapter 4 for why a rational agent must have transitive preferences does not and should not assume that agents have complete preferences: if it did then we would leave open the possibility that the drawbacks of intransitivity were an artifact of completeness. Conversely, a demonstration that incompleteness does no harm should not presuppose that agents violate those axioms such as the transitivity of preferences (as opposed to the transitivity of revealed preferences) that do serve their interests. To take a different example, I will argue in Chapter 8 that classical welfare economics offers workable advice only if preferences are complete; a more accurate characterization of the ideal frictionless world will thus expose the ineffectiveness of traditional policy advice.

These uses of rationality mark the difference between economics and those social sciences, such as psychology, that do without norms of behavior. An economic understanding of individuals and social institutions must go beyond how agents act and pose a pertinent set of counterfactuals. When agents violate a putative norm of rational conduct, we must ask if the agents would be better off if they instead complied with the norm, and if so, what is preventing them from doing so.

"The combined assumptions of maximizing behavior, market equilibrium, and stable preferences, used relentlessly and unflinchingly, form the heart of the economic approach. . . ." These words of Gary Becker from 1976 leave the meaning of "maximizing" ambiguous. Does maximizing mean that, in a properly constructed economic model, agents pursue their interests? Or does it mean that each agent's goals are sufficiently unified that they can be assimilated into a single utility function? And if an agent maximizes a utility function, must that function display smooth trade-offs, as the utility functions in all of Becker's models do? If only the answer to the first of these questions is "yes" then this book hews to Becker's philosophy. The shading though will admittedly look different: preferences will be displaced from their position at the heart of economics and production becomes central.

1.3 THE IDEOLOGY AND SCOPE OF TRADE-OFFS

When economists advise governments or run universities, they are eager to point out that decisions present trade-offs. Redistributions of income bring tax distortions and hence inefficiency. Lower tuition comes at the cost of increased faculty teaching loads. The lockdown that saves lives from an infectious disease will choke off economic activity. This fixation – that trade-offs are the be-all and end-all of governance – provides economics with its own distinctive ideology. Just as battlefield casualties force a medic into triage decisions, life confronts us with hard choices. Real economists are ready to tackle them.

The difficult decisions the world serves up are resolved in economics by the differentiable preference trade-offs of the indifference curve. Thus armed, the agents of economic models can identify a small class of solutions to the problems they face. This resolution made perfect sense in the 19th century when the differential calculus provided the lingua franca of science. Enough time has passed, however, to allow economics to acknowledge that agents do not always know how to compare and judge, both on the margin and overall. Deciding whether to accept this fact may be the ultimate hard choice for an economist.

Smooth preference trade-offs will be absent when different goods appeal to an agent but the agent cannot identify a specific rate at which increments of those goods can be substituted without benefit or loss. These lapses from the neoclassical ideal come in two varieties. In the first, an agent views some trade-offs of consumption as neither beneficial nor harmful. Since these appraisals moreover will not depend on the exact terms of the trade-off, they cannot be explained away as cases of indifference. For instance, a one unit sacrifice of good 1 might require at least x additional units of good 2 to leave an agent at least as well off but a gain of less than x units of good 2 might not harm the agent. As we will see, incompleteness provides the only convincing account of these judgments. In the second, an agent might have classically rational preferences but the agent's indifference curves are kinked. A one unit sacrifice of good 1 might again require x units of good 2 as compensation but an x unit sacrifice of good 2 might not be remedied by a one unit gain of good 1. In both cases, an agent will display an interval of margins or supporting prices at some or all consumption options.

Not every pathology of preference theory indicates an inability to say what one good is worth in terms of another. A telling case is provided by lexicography, where an agent considers one good to be so superior to a second good that any increase in the consumption of the former, no

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

matter how small, is superior to any increase in the latter, no matter how large. For example, let the bundle of two goods (x_1, x_2) be preferred to (y_1, y_2) if and only if $x_1 > y_1$ or $(x_1 = y_1$ and $x_2 \geq y_2)$. Although lexicography is the bad boy of utility theory – the preference is complete and transitive and yet has no utility representation[5] – a lexicographic agent *can* say what one good is worth in terms of another: there is no amount of good 2 the agent would accept for any amount of good 1. The market demand of such an agent, moreover, is identical to the demand of an agent with the differentiable utility $u(x_1, x_2) = x_1$ when prices are strictly positive.[6] A world of such lexicographic agents moreover would create no ripples; firms would not produce good 2 and it would cease to exist. Lexicography thus amounts to a decided trade-off. It is the undecided trade-offs that spell trouble for conventional economics.

You may walk away unconvinced, despite my efforts, that agents display ranges of marginal valuations rather than the single margins of smooth indifference curves and that such preferences pass the test of rationality. Intellectual debates on fundamentals are always prone to cycles of objection and rejoinder. But you ought to know what turns on your allegiance to the smooth textbook indifference curve. With smooth trade-offs, market prices will display less volatility and classical welfare economics will be able to generate usable policy recommendations. Your allegiance however will also mean that you must live with a stubborn discrepancy between your predictions and the reality of individual choice behavior. You will have trouble explaining why agents hold persistently to the status quo and why they so often conclude that their options include no dominant alternative.

Since the link between smooth trade-offs and their economic consequences will often be the implicit subject matter, many of the arguments in this book have a contingent or hypothetical character; the causal mechanisms at play are as important as the facts on the ground. For example, the goal of Chapter 7 is not that pricing is volatile, the point is that the best explanations of the incidence and logic of volatility rest on production rather than preferences. If you are skeptical of volatility, you should know what arguments best rationalize your position. For a second example, my assumption throughout the book that preferences (as opposed to revealed preferences) are transitive is not an empirical claim: it is made to

[5] Debreu (1954).

[6] This point is due to Richter (1966) though he gives this agent a more complicated utility function that fails to be differentiable.

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

isolate without any extraneous distortions the impact of incomparability and incompleteness.

1.4 A THEORETICAL RUN-THROUGH

The two varieties of preferences without smooth trade-offs mentioned in the previous section will be presented in two models, local incomparability and safety bias. They provide the centerpieces of Part I and appear in Chapters 3 and 5. In the chapter in between, you will find the case for the rationality of agents who are not always capable of comparing alternatives. The book's preference applications in Part I and economy-wide analyses of pricing and policymaking in Part II are built on this foundation.

The run-through below will outline the two main models and their connections to pricing and policymaking – which might otherwise be difficult to see. For readers interested in practical applications, the run-through also offers a theory-light path: read the synopsis below possibly joined to a sightseeing tour of Chapters 3 and 5, take a sample of the preference applications in Sections 4.3 and 6.3, and then jump to Part II. I have omitted from the run-through several topics, including the book's positive proposals for policymaking in Chapter 9, when the run-through provides a sufficient bridge.

A Puzzle

Jamie and Pat are deciding on a movie. Jamie says "You should choose since I am indifferent." Pat makes the same plea. Are these claims believable? Are the parties indifferent, as they claim, or are they waffling, unable to rank the available alternatives?

Classical consumer preferences rule out incompleteness, the inability of an agent to form a preference between some pairs of options. So if we infer from Jamie and Pat's conversation that neither holds a strict preference between some pair of movies, then the only remaining classical possibility is that they are indifferent. But indifference is a highly unlikely event. If a preference satisfies the textbook assumptions of continuity and increasingness, then the pairs of bundles that are indifferent form a low-dimensional subset of the space of all pairs of bundles: each indifference curve will be a "thin" subset of the positive orthant, a line (usually assumed to be convex) when there are just two goods. As a practical matter therefore we should rarely if ever see agents who are indifferent between alternatives. But since we see agents like our moviegoers all the time, the classical model cannot be correct.

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

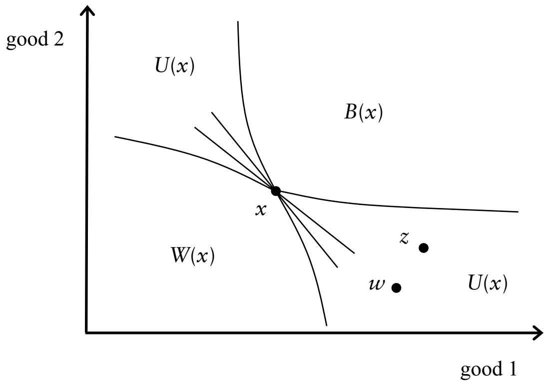

FIGURE 1.1 Locally incomparable preferences

Local Incomparability

To address this puzzle and other problems with the classical model, the agents most frequently encountered in this book will have a form of incomplete preferences where they cannot decide how much a small increment of a good is worth in terms of another good. Given any bundle x, the set of bundles one of our agents regards as weakly better than x, which I label $B(x)$, will display a kink at x. Equivalently, multiple price lines will *support* $B(x)$ at x, that is, $B(x)$ lies to the northeast of multiple price lines through x. The agent thus has an interval or band of margins. See Figure 1.1. Remarkably enough, kinks of this sort are intimately connected to incompleteness. If $B(x)$ displays a kink at x greater than some minimal size for every bundle x then under mild conditions the agent's preferences must be incomplete. So some bundles y will not land in either $B(x)$ or in $W(x)$, the bundles weakly worse than x, but in a third category $U(x)$ of bundles unranked relative to x.

Preferences that fit the pattern of Figure 1.1 – the presence for each x of multiple price lines that support $B(x)$ and that, in the vicinity of x, are contained in $U(x)$ – will be called *locally incomparable*. For such preferences, the $B(x)$ sets will frequently overlap as x varies, illustrated in Figure 1.2. And due to their incompleteness, locally incomparable preferences never have utility representations: a utility would place any bundle y in either $B(x)$ or $W(x)$ or both.

One way that locally incomparable preferences can arise is from *unanimity aggregation* where an individual holds several candidate preference relations to be reasonable and commits to a ranking of alternatives only when all of the candidates agree. For example, in Figure 1.1, each of the two segments of the boundary of $B(x)$ could be a portion of an

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

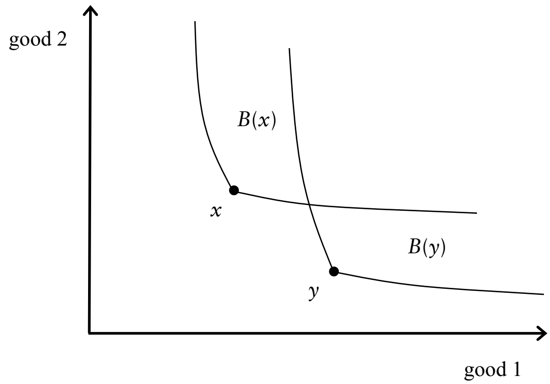

FIGURE 1.2 The overlap of better-than sets

indifference curve of a fully orthodox candidate preference: a bundle y then lies in $B(x)$ only if the two candidates unanimously back y over x.

Incompleteness and local incomparability can account for Jamie and Pat's interchange. A bundle is classically defined to be indifferent to x if it is contained in both $B(x)$ and $W(x)$. In Figure 1.1, x itself is the only such bundle but the bundles in $U(x)$, which are unranked relative to x, are plentiful. If in addition $B(x)$ and $W(x)$ change continuously as x varies then incompleteness is robust: when y lies in the interior of $U(x)$, an absence of preference will also hold for pairs near x and y. So if we reinterpret our couple's claims of indifference as an inability to form preferences then their actions, though they violate the classical model of consumer theory, can readily be explained.

Little turns on how the bundles in $U(x)$ are labeled, and we could instead declare each bundle in $U(x)$ to be indifferent to x. The drawback of this redescription is that the agent's preferences and specifically the agent's indifference relation will then fail to be transitive. For example, if z has slightly more of both goods than w and both lie in $U(x)$ then w will be indifferent to x and hence weakly preferred to x, x will be indifferent and hence weakly preferred to z, and yet z will be strictly preferred to w (assuming the preference is increasing). Letting $\succsim$ and $\succ$ denote weak and strict preference, the intransitivity $w \succsim x \succsim z \succ w$ has appeared. Since with this new preference we would face the chore of disentangling such intransitivities of indifference from the more significant intransitivities of strict preference, I will stick to the standard definition of indifference. The standard definition should also reassure you that the overlap of $B(x)$ sets in Figure 1.2 does not imply that the underlying preference is intransitive;

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

since the boundary of $B(x)$ need not contain any points indifferent to x besides x itself, an overlap does not lead to indifference curves that cross. Indeed incompleteness nearly snuffs out indifference: it turns out that x can be indifferent to only a negligible fraction of points on the boundary of $B(x)$.

Why is it important to keep track of whether preferences are transitive? Decades of tradition in economics have tied rational self-interest to preferences that satisfy the two classical axioms of rationality, completeness and transitivity. This unfortunate fusion, never backed up by rigorous argument, has misled generations of economists, students and teachers alike.

Failures of transitivity can inflict harm by leading agents through sequences of decisions that leave them with options worse than what they began with. The money pump – where an agent with the preferences $x \succ y \succ z \succ x$ will happily pay a small sum in each round of an endless sequence of exchanges of a worse for a better option – is only the most famous and extreme case of the damage intransitivity can wreak.

Incompleteness on the other hand does no harm if the preference judgments that agents can make satisfy transitivity. Agents can protect themselves from damaging sequences of trades simply by declining to swap options that they do not know how to order. As long as an agent does accept any offer of an alternative that is strictly preferred to the agent's current holding, the agent will not be led to an option worse than what the agent could have reached with a different trading strategy.

These conclusions bring to light the rational self-interest behind one of the enduring regularities of behavioral economics, sticking to the status quo. In Figure 1.1, z has slightly more of both goods than w but neither is ranked relative to x. A strategy of exchanging unranked options at will could lead the agent from z to x and then to w. Status quo bias will stop the agent from being ensnared by vetoing the first exchange.

The rationality logic behind status quo bias is visible however only if we treat choice and preference as different concepts. If the agent in Figure 1.1 follows the status-quo bias strategy of making exchanges only when offered a preferred alternative, the agent will sometimes choose x over z and sometimes choose z over x: it will depend on which is the status quo. The agent will similarly vacillate between x and w. On the revealed preference view that choice and preference are different names for the same phenomenon, an intransitivity has appeared: the agent has revealed a strict preference for z over w and indifference between w and x and between x and z. The same intransitivity arose when we labeled the options in $U(x)$ as indifferent to x and it underlies the skepticism of economists that agents will persist in status quo bias. But if we distinguish between preference and choice then we can see that the intransitivity that stems from status quo bias, far from inflicting harm, protects agents from destructive chains of decisions.

When preference is defined by an agent's judgments of well-being, self-interest will require those judgments to be transitive but not complete: an agent facing some pair of options might not be able to say which is better. And the intransitivity that accompanies status quo bias shows that a choice definition of preference will not rescue the completeness-cum-transitivity theory of rationality. An agent's revealed preferences must be complete – if x and z are unranked then x when it is the status quo will be revealed preferred to z – but they will fail to be transitive.

A locally incomparable agent exemplifies the status quo bias that accompanies incomplete preferences. Such an agent will have a band of marginal valuations, illustrated in the multiple price lines of Figure 1.1 that support $B(x)$. The steepest of these lines indicates that when facing a unit loss of good 1 the agent will require a large amount of good 2 as compensation, the flattest of the lines indicates that the agent will sacrifice only a small amount of good 2 for a unit gain of good 1. This differential response to gains and losses goes under several names, including loss aversion and the willingness-to-accept/willingness-to-pay disparity, and the differential also serves as a common definition of status quo bias. None of these classic patterns of behavioral economics need to be posited as ad hoc psychological facts: they stem from the demands of rationality when agents cannot judge trade-offs.

The agents with smooth indifference curves that appear in orthodox consumer theory, in contrast, will not display status quo bias or any of its cognates. If to begin such an agent has purchased positive amounts of all goods then he or she will necessarily want to make further trades if prior to consumption relative prices were to shift even slightly. Incompleteness therefore scores a second empirical success over the textbook indifference curve, beyond its better explanation of how frequently agents cannot strictly rank their options.

Rooting status quo bias in incomplete preferences also lets us dispense with the awkward psychological asymmetry where the pain of losses inherently cuts deeper than the pleasure of gains. Suppose that unanimity aggregation lies behind Figure 1.1 with the two halves of the boundary of $B(x)$ representing portions of the indifference curves of two candidate preferences. The simplest way for this agent to unanimity aggregate is to

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

see if a change relative to x benefits the candidate preference that is most pessimistic about the change: if this candidate preference approves the change then so will the other candidate. So if the change is a loss of good 1 and a gain of good 2, the agent can let the candidate with the steeper indifference curve make the decision and if the change is a gain of good 1 and a loss of good 2, then the candidate with the flatter indifference curve can make the call. This delegation amounts to an optimization technique, not an impulsive reaction to losses.

Safety Bias

When agents cannot judge trade-offs, they can turn to many ways of making decisions – unanimity aggregation is just one possibility. An agent facing a difficult choice can take a cautious approach by letting the decision be settled by a worst-case assessment of each option. To this end, the agent can deploy a set of *welfare functions*, each of which assesses all of the agent's options. The worst case for an option x will be given by the smallest of these welfare assessments and the agent when facing a pair of options will choose the option with the greater worst-case assessment; these rankings of the pairs in turn define an entire preference relation. When one of the agent's welfare functions fails to provide the worst-case assessment for either of two options x and y, then that function is irrelevant to the agent's preference between x and y. So, in contrast to unanimity aggregation, the present preferences do not require unanimous agreement.

When the agent's assessments all agree on the welfare level of an option x then x guarantees that welfare level: x is *safe*. Since different movements away from x can affect which welfare function delivers the worst-case assessment, the agent's better-than set $B(x)$ will be kinked at x: a set of margins or price lines will support $B(x)$ at x. In contrast to local incomparability, however, the multiplicity of supporting prices occurs only at the safe consumptions rather than at every consumption. The agent's decisions will therefore be biased in favor of safe options.

In intertemporal choice, for example, an agent who is unsure how to trade-off consumption at different dates might consider, for each time t, a welfare function that equals time-t consumption x_t to provide one of the reasonable assessments. The minimum coordinate of a bundle $x = (x_1, \ldots, x_T)$ would then be the agent's worst-case assessment of x and the agent when choosing between two options would select the bundle whose minimum coordinate is greater. The safe consumptions would then consist of the constant bundles – the 45° line when there are two periods.

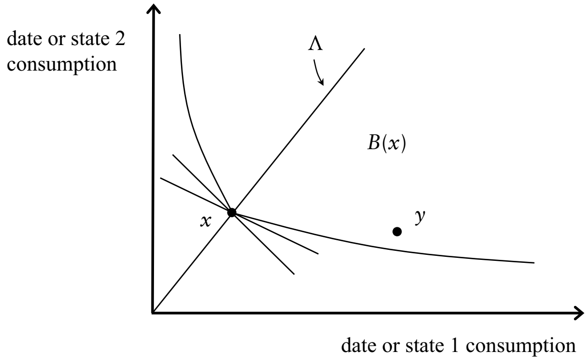

FIGURE 1.3 Safety bias

To see with two periods that a kink in $B(x)$ occurs when $x_1 = x_2$, fix date 1 consumption at $\bar{x}_1$ and increase x_2 from 0. The worst-case assessment then begins at 0, increases with x_2, and plateaus at $x_2 = \bar{x}_1$ when the welfare function that gives the worst-case assessment switches from x_2 to x_1. A kink in $B(x)$ therefore occurs at $(\bar{x}_1, \bar{x}_1)$.

This Leontief-style example is extreme for clarity's sake. An agent could well regard a ray of time-varying consumption bundles to be safe – social conventions might convince the agent that consumption must increase over time to qualify as safe. An agent would also be likely to accept some intertemporal trade-offs and prefer a nonsafe y over a safe x when the gains that come with y at some dates sufficiently outstrip the losses at other dates. These possibilities are pictured for the two-period case in Figure 1.3, where the preferences arise from two welfare functions and Λ is a ray of safe bundles. Each of the two parts of the lower boundary of $B(x)$ is a portion of an indifference curve of one of the two welfare functions and, since the two parts meet at a point on Λ, the two welfare functions assign the same welfare number to the two indifference curves. As in the Leontief case, kinks occur along the ray because as a path of consumption crosses the ray the welfare function that gives the worst-case assessment switches.

In choice under uncertainty, an agent might be unsure which subjective probability distribution to employ. Each distribution in some set of distributions the agent deems reasonable would then define an expected utility function that can serve as a welfare function and the worst-case assessment of an option x would be the minimum of these expected utilities when evaluated at x. Consumption constant across states of the world

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

will then form a safe pattern; with two states, the ray Λ of safe bundles in Figure 1.3 would be the 45° line.

For both intertemporal and uncertain choice, multiple rays of safe consumption are also possible.

An agent with rays of safe bundles along which the agent's indifference curves are kinked and supported by multiple prices will be called *safety biased* – whether or not the psychological backstory of making the worst-case assessment as beneficial as possible holds. Although an inability to make preference judgments is the most prominent source of safety bias, there can be other origins, for example, complementarities in consumption. The left and right shoes that motivate Leontief utilities provide a textbook example though the vertical and horizontal stretches of Leontief indifference curves will not hold in any important case of safety bias in this book. One goal will be to show that complementaries are not as innocuous as economists have been trained to believe.

Welfare Economics and Volatility

Local incomparability and safety bias will drive the economy-wide applications in Part II of the book.

In welfare economics, the omnipresent kinks of local incomparability imply that an agent has a range of valuations of one good in terms of another. So if, for example, good 1 in Figure 1.1 is a public good and good 2 is the agent's private consumption, the kink in $B(x)$ implies that the agent would sacrifice only a little private consumption to scale up the public good a notch and would demand a lot of private consumption to take it down a notch. When these gaps are summed across agents, there will be a large band of aggregate marginal valuations of the public good. So if the marginal cost of the public good in terms of private consumption falls into this band, classical welfare criteria will offer only limited help to a government deciding how much of the public good to construct: both large and small quantities will qualify as efficient.

A similar policy paralysis besets the economy as a whole. Due to the second welfare theorem, if an allocation of goods $x = (x^1, \dots, x^I)$ among I individuals is Pareto efficient then a common price vector will support each agent i's better-than set at x^i. Under local incomparability however a band of vectors can typically serve as supporting prices. If moreover this band moves continuously as the allocation changes and the band has maximum dimension, then for allocations near x each agent i's better-than set will typically continue to be supported at i's new allocation by a common price vector. The first welfare theorem therefore implies that

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

the new allocation will also qualify as efficient. Similarly, when an allocation x is Pareto efficient and the economy subsequently undergoes a small change to the modeling environment, perhaps a small externality, the price vector that initially supported the agents' better-than sets at x will usually continue to do so; the change will therefore call for no policy response. The discriminatory power of economic policy advice thus leaches away.

In early neoclassical economics, the differentiable utility function underlay the stability of prices: small exogenous changes to endowments or technology will typically lead to a small response in the marginal rates of substitution that hold in equilibrium and thus in equilibrium relative prices. With safety-biased preferences, on the other hand, the kinks in indifference curves undermine this argument: small endowment changes in the vicinity of a kink will now force large changes in supporting price lines. Relative prices will consequently be volatile. Modern general equilibrium theory has countered that it is exceedingly unlikely that endowments land at just the points that allow agents to consume in the neighborhood of indifference-curve kinks. The kinks however are attractors. In their presence, an intertemporal investment in a future consumption good will earn a future utility return that fails to vary smoothly as a function of current consumption forgone: the rate of return will fall discontinuously as the production of the future consumption good places agents along their safe rays. Investment can therefore propel an economy to the endowment points where agents do consume at indifference-curve kinks. The volatility scenarios will then play out: small endowment changes in the future will drive substantial movements in relative prices.

1.5 CONVENTIONS, DEFINITIONS, AND AVOIDING MATH

Mathematical models supply the alpha and omega of this book and many Greek letters in between, but the synopsis above and the preference applications in Sections 4.3 and 6.3 allow a math-avoiding route to Chapters 8 through 10 on policymaking, where the models are more concrete and self-contained. Discovering the intersections of mathematics and economic content, where a formal model quietly imports a vision of human nature, is one of the pleasures of economic theory. But not necessarily on every occasion.

The initial sections of most chapters provide informal summaries of the models to follow and, in the second half of the book, earlier-in-the-chapter partial equilibrium treatments are lighter going that later-in-the-chapter general equilibrium models. When a section or displayed example bears obvious mathematical warning signs, you may drive by. If you are not well-versed in general equilibrium theory, portions of Chapter 7 invite skimming: I could not in this case shepherd the math into skippable chunks.

Many proofs will be found in the Appendix rather than the body of the text.

The items below lay out the notation used in this book. While most of these definitions will be repeated along the way, this list can serve as a backup.

In models with goods, there will when possible be L goods and I agents, the agents will consume nonnegative bundles, subscripts will denote goods, and superscripts will denote individuals. So agent i's consumption of goods will normally be a vector of L real numbers $x^i = (x^i_1, \ldots, x^i_L) \geq 0$. Both L and I will always be finite.

For any two vectors of n real numbers x and y, $x \leq y$ holds if $x_i \leq y_i$ for all i, $x < y$ holds if $x \leq y$ and $x \neq y$, and $x \ll y$ holds if $x_i < y_i$ for all i. Greater than inequalities are defined analogously.

Preferences will often be expressed as binary relations. A binary relation $\succsim$ on a set X is *complete* if, for all $x, y \in X$, either $x \succsim y$ or $y \succsim x$ (or both); *transitive* if, for all $x, y, z \in X$, $x \succsim y$ and $y \succsim z$ imply $x \succsim z$; *asymmetric* if $x \succsim y$ implies not $y \succsim x$; *reflexive* if, for all $x \in X$, $x \succsim x$; incomplete when $\succsim$ is not complete; and *classically rational* if $\succsim$ is complete and transitive. When a binary relation $\succsim$ is incomplete and neither $x \succsim y$ nor $y \succsim x$ holds then x and y are $\succsim$*-unranked* or $\succsim$*-unrelated* and I sometimes then write $x \perp y$. If either $x \succsim y$ or $y \succsim x$ (or both) then x and y are $\succsim$*-related*.

Given a binary relation $\succsim$ on X that is interpreted as a preference, the *strict preference* $\succ$ is defined by $x \succ y$ if and only if $x \succsim y$ and not $y \succsim x$, the *indifference relation* $\sim$ is defined by $x \sim y$ if and only if $x \succsim y$ and $y \succsim x$. If neither $x \succ y$ nor $y \succ x$ hold then (in Chapter 6 only) I will write $x \sim^* y$, which need not imply $x \sim y$.

A preference relation $\succsim$ on X is *represented by a utility function u* from X to $\mathbb{R}$ if, for all $x, y \in X$, $x \succsim y$ if and only if $u(x) \geq u(y)$.

When the set of consumption possibilities consists of bundles of goods, $X = \mathbb{R}^L_+$, I will normally assume that preferences and utility functions are increasing, though utilities will on occasion be weakly increasing to admit convenient illustrations. Increasingness can usually be weakened considerably, for example to local nonsatiation. Given $X = \mathbb{R}^L_+$, a preference relation $\succsim$ is *increasing* if, for all $x, y \in X$, $x > y$ implies $x \succ y$; *weakly*

https://doi.org/10.1017/9781009340731.002 Published online by Cambridge University Press

increasing if, for all $x, y \in X$, $x \gg y$ implies $x \succ y$; *convex* if, for all $x \in X$, $B(x) = \{y \in X : y \succsim x\}$ is a convex set and *strictly convex* if $B(x)$ is a strictly convex set, that is, for all $y \neq z$ and $0 < \alpha < 1$, $\alpha y + (1-\alpha)z$ lies in the interior of X; *locally nonsatiated*, if for all $x \in X$ and all $\varepsilon > 0$, there is a $y \in X$ such that $\|y - x\| < \varepsilon$ and $y \succ x$; and *smooth* if $\succsim$ can be represented by a differentiable function u.[7] A function u from X to $\mathbb{R}$ is *increasing* if, for all $x, y \in X$, $x > y$ implies $u(x) > u(y)$ and is *weakly increasing* if, for all $x, y \in X$, $x \gg y$ implies $u(x) > u(y)$.

If x and y are two vectors of n real numbers, then $x \cdot y$ denotes the dot or inner product $\sum_{i=1}^{n} x_i y_i$. Unless clarity demands an exception, I will omit the bounds in sums: so $\sum_i x_i$ rather than $\sum_{i=1}^{n} x_i$.

The set $B(x)$ above, sometimes denoted $B_{\succsim}(x)$, is the *better-than set* of the binary relation $\succsim$ at x. A $p \in \mathbb{R}^L$ *supports* $B(x)$ at x if $p \cdot (z - x) \geq 0$ for all $z \in B(x)$.

Let f be a function from $\mathbb{R}_+^L$ (or some other convex set) to $\mathbb{R}$. If f is differentiable then $D_{x_i} f(x)$ will denote the partial derivative of f with respect to its ith argument evaluated at $x = (x_1, \ldots, x_n)$ and $Df(x)$ will be the vector of these partial derivatives. When $L = 1$ (the domain of f is $\mathbb{R}_+$) then $D_+ f(x)$ or $D^+ f(x)$ will denote the right derivative of f evaluated at x, and $D_- f(x)$ or $D^- f(x)$ will denote the left derivative. The function f is *concave* if, for all x and y in $\mathbb{R}_+^L$ and all $0 \leq \alpha \leq 1$, $f(\alpha x + (1-\alpha)y) \geq \alpha f(x) + (1-\alpha) f(y)$. If $\geq$ is replaced by $\leq$ in the last inequality then f is *convex*. The function f is *strictly concave* if, for all $x \neq y$ and all $0 < \alpha < 1$, $f(\alpha x + (1-\alpha)y) > \alpha f(x) + (1-\alpha) f(y)$. If $>$ is replaced by $<$ in the last inequality then f is *strictly convex*.

[7] Debreu (1972) provides a condition on preferences that implies the existence of such a representation.