

INTRODUCTION

Painting with a Broad Brush

Thinking like an Actuary

The scenario is familiar. A newspaper publishes an exposé. A government investigation issues its report. The exposé or report is about auto mechanics, for example, and reveals that many overcharge their customers, perform unnecessary repairs, and use secondhand rather than new replacement parts. Or perhaps the report is about funeral directors who prey on the bereaved, car salesmen who misrepresent the condition of used automobiles, surgeons who advise patients to have unnecessary but expensive operations, civil servants who are on the golf course rather than in their offices, stockbrokers who trade excessively with their customers' accounts in order to reap large commissions, or lawyers who commingle their clients' funds with their own. In these and countless other examples, the investigation concludes that members of the profession have persistently abused

the trust that their clients and the public have placed in them.

The reaction is virtually automatic. A day after the report appears, the professional association of automobile mechanics—or the equivalent for funeral directors, automobile dealers, surgeons, civil servants, stockbrokers, or lawyers—issues its response. The conclusion drawn by the journalists or investigators is *unfair*, the association charges. It is unfair not because the instances of abuse that were discovered did not take place. Nor is it unfair because the individuals named are honest professionals. Rather, the report is unfair because not *all* auto mechanics, funeral directors, car salesmen, surgeons, civil servants, stockbrokers, or lawyers engage in such reprehensible practices. Although *some* do, the association admits (the evidence typically being so overwhelming as to make any other response laughable), these cases are the exception and not the rule. Every profession, after all, has its bad apples, and the association is now dealing with the small minority of individuals who are giving the entire profession a bad name. In fact, the response continues, the association has just issued new rules and hired more inspectors, these measures being designed to guarantee that a few ethically lax individuals cannot succeed in tarnishing the reputation of thousands upon thousands of honest professionals.

We typically treat such self-serving protestations of unfairness with the skepticism they deserve. Yet although we are skeptical, the argument made on behalf of the impugned profession resonates with many people. Even if 40 percent of the used-car dealers turn back odometers in order to make their wares look newer than they are, there then remain 60 percent who do not. These 60 percent, however, still suffer from the stigma

created by the culpable 40 percent. Upon learning that a full 40 percent of any profession are corrupt, few of us could resist looking askance at the entire profession, including the 60 percent who have done nothing wrong. Moreover, few of us can avoid relying on such percentages as we make concrete decisions in our daily lives. As more and more cases of child molestation by priests (and also by scout leaders, baseball coaches, and day-care providers) are reported in the newspapers, many people place less trust in *all* priests, or even in all members of the clergy. As a result, the overwhelming majority of clergymen find themselves under a cloud, even though they are individuals for whom such behavior would be anathema. Or who among us would not prefer a retired federal judge to a former used-car salesman as an investment adviser? And most people think it perfectly sensible to guard their wallets and pocketbooks more carefully in Naples than in Helsinki, even as they know that the overwhelming majority of Neapolitans are every bit as honest as the overwhelming majority of Finns. Yet as we make such decisions, we may sometimes feel uneasy, for we are attributing to the many the sins of the few. Our intuitions are in conflict, because when we attribute the failings of the minority to the entire group we may sense that we are being simultaneously rational and unfair.

In stigmatizing an entire profession because of the behavior of a minority, or in engaging in similar attributions in other areas of decisionmaking, we undertake a process of *generalization*. "Painting with a broad brush," the slightly pejorative expression puts it. To refer to the practice as *stereotyping* intensifies the condemnation, and these days to describe it as *profiling* condemns more strongly still. Yet regardless of the label the structure of

the decision is the same.¹ On the basis of a characteristic of some members of a class, we reach conclusions or make decisions about the entire class. When a teenage male acquires a driver's license, his family's insurance premiums will rise dramatically, even if their son is a paragon of caution less likely to be involved in an accident than even most adults. That this particular teenage male does not possess the traits possessed by many others in the class in which he is placed is, here and elsewhere, largely beside the point.

The phenomenon of the teenage male whose acquisition of a driver's license causes his family's insurance premiums to skyrocket reminds us that decisionmaking by generalization is the stock in trade of the insurance industry. Indeed, the insurance industry has its own name for this kind of decisionmaking. To be an *actuary* is to be a specialist in generalization, and actuaries engage in a form of decisionmaking that is sometimes called *actuarial*. Actuaries guide insurance companies in making decisions about large categories (teenage males living in northern New Jersey) that have the effect of attributing to the entire category certain characteristics (carelessness in driving) that are probabilistically indicated by membership in the category, but that still may not be possessed by a particular member of the category (this *particular* teenage male living in northern New Jersey).

Occasionally the actuarial generalizations of the insurance industry become controversial. One example is the use of generalizations about the comparative safety of different neighborhoods as a basis for setting the rates for homeowners' insurance or determining the willingness of a bank to give a mortgage. Normally these generalizations attract little attention, but they become controversial when they have a racially disproportionate

effect, raising mortgage and insurance costs in predominantly black neighborhoods more than in largely white ones.² Similarly, there are protests at the prospects of requiring gay men to pay more for health insurance because as a class they have a statistically greater likelihood of contracting the HIV virus, and there is anger at the thought that people whose genetic makeup makes them more likely to suffer from cancer or heart disease would have to pay higher premiums for life insurance.³

Yet although examples like these generate controversy, in the vast majority of instances the actuarial behavior of the insurance industry is accepted simply as a fact of life. As a resident of Massachusetts, which has more than its share of bad drivers, I pay higher premiums for automobile insurance than I would were I a fulltime resident of Vermont. And this is so not only because the prevalence of bad drivers in Massachusetts increases the likelihood of my being involved in an accident that is not my fault. Rather, the principal reason my insurance premiums are so high is that I am actuarially saddled with the driving habits of my fellow Massachusetts residents, even though these are habits I do not share. And if as a safe driver I buy a high-performance sports car, my insurance premiums will again jump, even if the sports car has no greater monetary value than the staid station wagon I now own, and even if the sports car is no more attractive to thieves. When I buy a sports car I become actuarially encumbered with the risk-taking driving attributes of most sports car owners, attributes that again I do not share, and indeed that if given the opportunity I could prove to my insurance company I do not share. But the insurance company is stunningly uninterested in providing me the opportunity to demonstrate that the generalizations about Massachusetts

drivers and sports car owners do not apply to me. I wish it were otherwise, but in the spectrum of life's injustices, both of these seem rather far down the list.

Insurance companies, of course, are not the only ones who operate actuarially. We all do, with far more frequency than most people typically acknowledge. We operate actuarially when we choose airlines on the basis of their records for safety, on-time performance, or not losing checked luggage. We operate actuarially when we associate personal characteristics such as a shaved head, a tattoo, and black clothing with behavioral characteristics, such as racist beliefs and a propensity to violence, that the personal characteristics seem probabilistically but not inexorably to indicate.⁴ And we operate actuarially when we indulge in ordinarily harmless stereotyping by nationality, as in describing Italians as emotional and Scots as dour, even while recognizing that more pernicious stereotyping by race, by gender, by sexual orientation, and by ethnicity remains widespread. Still, once we see that insurance companies are in the business of systematic stereotyping, that employers stereotype when they assume that certain characteristics (good grades from a prestigious university) will predict successful job performance, that police detectives focus on suspects by aggregating stereotypes, and that most of us stereotype in much of our daily lives, we cannot so easily dismiss the practice of stereotyping—or profiling—as necessarily morally wrong. Nor can we dismiss as irrational the practice of generalizing even when the generalization produces errors in particular cases. This book tries to sort out when painting with a broad brush is desirable and when it is not, recognizing, for the time being, that the word “desirable” encompasses a wide range of considera-

tions of prudence, efficiency, and, most importantly, morality.

Generalizations Good and Bad

At the outset it is important to draw a distinction between those generalizations that have no statistical or factual basis and those that do. Following the statisticians, we can label the generalizations that are devoid of such empirical foundations as *spurious*. And although spurious generalizations will crop up at various points in this book, our primary concern will be not with the use of spurious generalizations, but rather with the use and misuse of those generalizations that do have a sound statistical basis. We can label the generalizations with a sound statistical basis *nonspurious*, but what it is for a generalization to have a sound statistical basis is a tricky question deserving some preliminary exploration.

Let us say that a generalization characteristically takes the form of “*x*’s are *y*.” In this schematic form, *x* is some identifiable attribute or description, and *y* is an act, property, or behavioral propensity that the attribute or description is thought to indicate or predict.⁵ “Scots are dour.” “Ferraris are fast.” “Golden Retrievers are good with children.” “Swiss watches are well made.” “Airline employees are rude.” “Nonsmokers live longer than smokers.” “Blondes have more fun.” “German wines are sweet.”

Sometimes these generalizations are *universal*, meaning that *all* of the *x*’s are *y*. In many cases universal generalizations are universal as a matter of definition. Philosophers often use the example “All bachelors are unmarried.” Here the generalization is necessarily uni-

versal, because it is part of the definition of a bachelor that a bachelor be unmarried. If someone is married, he simply cannot be a bachelor. Thus the generalization "All bachelors are unmarried" is a universal generalization just because all the *x*'s—the bachelors—must by definition be *y*—unmarried.

Sometimes, however, universal generalizations are universal not for definitional reasons but for empirical ones. Some fact about the world and not about the word makes the generalization universal. For instance, it used to be believed that all swans were white. Even before the people who made the generalization that all swans were white realized that there were black swans (as there are in Australia, but nowhere else in the world, except in zoos or as Australian transplants), they probably recognized that something could have been a swan and have been a color other than white. It is just that there happened to be no nonwhite swans, or so they thought. So until black swans were discovered, "All swans are white" was a universal generalization; but the possibility remained open, even then, that there could be black swans, unlike the possibility, nonexistent by definition, that there could be married bachelors. Most people believe that the generalization "All human beings are less than nine feet tall" is universal, but the universality is again empirical and not definitional. Nothing in the defining characteristics of humanity requires that human beings be less than nine feet tall, although it is a contingent fact about the world that all of them are.⁶ I would be astonished if I were to observe a person taller than nine feet, even more than people were astonished when they first encountered black swans, but my observation will not cause me to doubt that the nine-footer is a human being.

Universal generalizations, whether the source of the universality be definitional or empirical, tend to interest philosophers, but most of the generalizations that the rest of us employ and encounter on a daily basis are not universal. "Swiss cheese has holes." "Italians are demonstrative." "Philosophers are clever." "Volvos are reliable." "You get what you pay for" (inexpensive goods are poorly made). In using and understanding such statements, we recognize that they apply only to most cases, and although we could make things more precise by ensuring that we always prefaced each of the foregoing with "all," "most," "many," "a disproportionate number of," or some similar qualification, it is noteworthy that we rarely do so. Life is short, and time even shorter, and our language is consequently replete with unstated qualifications. As a result, it is an entrenched feature of linguistic usage that statements like "Swiss cheese has holes," "Volvos are reliable," and countless others are taken to be properly used even when not prefaced with the word "all," and even when it would be a mistake to do so. We know that some Swiss cheese has no holes, that some Italians are non-demonstrative, that some philosophers are dim, that some Volvos are lemons, and that some inexpensive products are well made. Nevertheless, the existence of particular cases in which the generalization does not hold rarely causes us to refrain from using the generalization, provided that the generalization holds true for most cases. As long as we remain satisfied that most Italians are demonstrative, it is an accepted part of everyday talk to assert that "Italians are demonstrative," even as we recognize that more than a few Italians are not.

Indeed, we often generalize, and are understood to have generalized appropriately, even when less than a

majority of the x 's are y . When we say, "Bulldogs have bad hips," we do not necessarily mean that even a majority of bulldogs have bad hips. What we mean is that bulldogs have a higher percentage of hip problems than the category of dogs as a whole, and a higher percentage of hip problems than do other breeds. We would understand as sound the claim that bulldogs have bad hips even if only 30 percent of all bulldogs had bad hips, so long as the percentage of hip problems for dogs as a whole was, say, 5 percent, and so long as the percentage of hip problems for other breeds was lower than the percentage for bulldogs. And when on the basis of accident statistics or personal experience we actuarially castigate a group (males under the age of twenty-five, for example) as "bad drivers," we often mean only that the group has a noticeably higher percentage of bad drivers than some other group, or than the group of all drivers. To say that young men are reckless drivers is not to say that the majority of young men are reckless drivers, but only that the percentage of reckless drivers among young men is noticeably higher than the percentage of bad drivers in some comparison group (young women, for example) or in some background reference group (all drivers, for example). So when soccer aficionados conclude that English soccer fans are violent, they believe their conclusion is justified on the basis of there being a higher incidence of violence among English soccer fans than among their counterparts in other countries, even as they recognize that the number of individuals contributing to the higher incidence still falls far short of a majority.

Although there are several varieties of linguistically and empirically plausible generalizations, and although the use of nonuniversal generalizations is an entrenched feature of our linguistic and decisional lives, it still re-

mains important to distinguish those nonuniversal generalizations that have a sound statistical basis from those that do not. Generalizations are of course statistically sound when they are universal, as in "Bachelors are unmarried" and "Human beings are less than nine feet tall." And generalizations are usually statistically sound when they accurately portray the traits of a majority of the members of the class, as in "Swiss cheese has holes." Moreover, as the example of bulldogs with bad hips demonstrates, generalizations are statistically sound whenever the generalization accurately portrays the members of the class as having a greater prevalence of some trait than has some larger class of which the group is ordinarily taken to be a part, even though the trait appears in less than a majority of the members of both groups. What makes "Bulldogs have bad hips" an example of a statistically sound generalization, despite the fact that a majority of bulldogs do not have bad hips, is that knowing that a dog was a bulldog would improve one's ability to predict hip problems relative to whatever information one had other than the information that the dog was a bulldog. As long as the probability of a dog's having hip problems given that the dog is a bulldog is greater than the probability of a dog's having hip problems given no information about the breed of dog, we can say that the trait of being a bulldog is *relevant*, and we can say that generalizing from that trait meets the minimum threshold of statistical (or actuarial) soundness.

The bulldog example illustrates not only that a generalization can be statistically sound despite the fact that the generalization applies only to a minority of members of the class, but also that a key feature of a sound generalization is its comparative dimension. Even if a general-

ization accurately represents the traits of even a majority of some group, we would think it misleading to use the generalization if the trait was no more present in the group we were talking about than it was in any other group. Suppose that by some measure of honesty—Would you turn in to the police an envelope of cash you found on the street even if you were absolutely certain that no one would ever know if you kept it for yourself?—60 percent of all human beings were honest. And suppose that by the same measure it turns out, not surprisingly, that 60 percent of Swedes were honest. Under these circumstances it seems misleading to say, “Swedes are honest,” even though a majority of Swedes *are* honest. Only if Swedes were more honest than some other group would there be a point in making this statement. So although “Swiss cheese has holes” is accurate for most pieces of Swiss cheese, the significance of the statement lies in the fact that Swiss cheese generally has holes and most other kinds of cheese generally do not. If all cheese had holes, saying “Swiss cheese has holes” would be as odd as saying, in most contexts, that “Volvos have four wheels.”

Yet even though generalizations can be statistically sound in a number of different ways, there *are* statistically unsound generalizations. Take the generalization “Capricorns are self-confident.” If one is a Capricorn purely by virtue of having been born between December 22 and January 19, there is clearly a recognizable category of Capricorns. There is also a recognizable category of self-confidence, and we all understand that self-confidence is a trait that some people have more of than do others. “Capricorns are self-confident” is consequently a statistically sound generalization if all Capricorns are self-confident, if the trait of self-confidence appears in

Capricorns to a greater degree than it appears in the population at large, or if self-confidence appears in Capricorns to a greater degree than it appears in members of the population born under other astrological signs. To give this generalization the greatest chance of achieving statistical soundness, we can focus just on the possibility that Capricorns possess the trait of self-confidence to a greater degree than the population at large, independent of the absolute percentage of Capricorns who are self-confident. Thus, if Capricorns possess the trait of self-confidence in a higher proportion than the trait exists in the entire population, then the generalization that Capricorns are self-confident will be statistically sound, and we will have better information by which to predict whether a person is self-confident if we know that the person is a Capricorn than if we do not have that information.

Yet even taking the version of the claim that would be easiest to demonstrate, it turns out that the claim implicit in this generalization is almost certainly untrue. There is no evidence supporting the belief that the attribute of being a Capricorn is predictive of self-confidence, and in the absence of any such evidence it is highly likely that the relationship between being a Capricorn and being self-confident is spurious. In the language of the law of evidence it is *irrelevant*. Yet regardless of what we call it the conclusion is the same—we have no reason to believe that the trait of self-confidence appears more commonly in Capricorns than in anyone else.⁷ If we want to know whether a person is self-confident, being told that the person is a Capricorn helps us not at all. In such a case, we can—and should—say that the generalization is irrelevant, that it is spurious, and that it is statistically unsound.

Another illustration of a spurious generalization is provided by the fortunately lost art of phrenology. Phrenology flourished in the middle of the 1800s, and the many people who believed in it believed that the contours of one's skull provided relevant evidence about a person's personality and natural talents. Because certain traits were fostered in certain parts of the brain, phrenologists and their followers believed, determining the size and location of external cranial features corresponding to those brain locations could provide valuable information about various mental abilities and traits of personality. One's acquisitiveness, self-esteem, combativeness, and facility with language, for example, were all thought to be reflected on a person's cranial terrain. As we know, however, it turned out that there was no evidence that any of phrenology's premises were actually true.⁸ Phrenology is no longer practiced, chiefly because generalizations about personality and ability based on cranial features were exposed as simply spurious. It was not that phrenological generalizations, like most generalizations, were not universal. It is that these generalizations (like astrological generalizations, which still, curiously, persist) had no basis whatsoever in even probabilistic empirical fact.

Astrology and phrenology do not, unfortunately, exhaust the universe of commonly employed but statistically unsound generalizations. In Japan, and more recently in China, it is widely believed that one can generalize about the personalities of people on the basis of their blood types. Consequently, people sometimes are not hired for jobs for which they are otherwise qualified because of blood types, despite the fact that there is no more evidence of the relationship between blood type and personality than there is of the relationship between astrological sign or skull terrain and personality.⁹ Closer

to home, consider the often-believed generalization that gay men lack physical courage. Again, assume that we can distinguish gay men from their heterosexual counterparts. And assume as well that there is some measure of physical courage, such as the willingness to risk death for a cause one believes in, as with a soldier who does not break ranks and run even under direct fire from the enemy. The generalization that gay men lack physical courage would then be statistically sound insofar as gay men possessed the proclivity to physical courage in a smaller proportion than it was possessed by all men, or all people, or in a smaller proportion than it was possessed by heterosexual men. But once again there appears to be not a shred of evidence to support this proposition. The generalization that gay men lack physical courage, like the generalizations that Capricorns are self-confident, that people with bumps on the left front side of their skull are lacking in self-esteem, and that people with type O blood are loyal, is statistically unsound. The trait around which the generalization is constructed—sexual orientation—turns out to be completely irrelevant to the characteristic—courage—that we may often be trying to predict.

When applied to groups of people, statistically unsound generalizations are often castigated as *prejudices*. In common parlance, to have a prejudice is to have an unsubstantiated belief about a person because of his or her membership in a group. But this way of describing what it is to be prejudiced, no matter how common, is ambiguous. Some people might say that to have an unsubstantiated belief about a person is to subscribe to a statistically unsound generalization, such as believing that gay men are cowardly or that Scorpios lack self-confidence. Other people, however, often say that it is

“prejudiced” not only to rely on a spurious generalization, but also to apply a statistically sound but non-universal generalization to a particular member of the group. Consider, for example, the case of people who have been incarcerated for crimes and served their sentences—ex-convicts, in the common vernacular. We know that many ex-convicts proceed to live exemplary lives, never again committing a crime. Yet the percentage of ex-convicts who commit subsequent crimes is much greater than the percentage of the population at large who commit crimes,¹⁰ and thus someone who was “prejudiced” against ex-convicts would not necessarily be someone operating on a mere superstition. On the contrary, a “prejudice” against ex-convicts might be the consequence of acting on a nonuniversal generalization that actually does rest on a sound statistical foundation. To be prejudiced against ex-convicts is quite different from being prejudiced against Scorpios, and the fact that the same word is often used to apply to both forms of acting on a generalization only compounds the confusion.

The people who apply the word “prejudice” to negative views about the class of ex-convicts are people who understand that the word “prejudice” is pejorative. They are unlikely to describe as a prejudice the belief that accountants are cautious, even though some accountants plainly are not.¹¹ But when such people describe as “prejudices” the statistically sound generalization about the propensity toward crime of ex-convicts, it is because they believe it is usually wrong to prejudge people even on the basis of statistically sound group characteristics. For such people it is wrong to condemn or distrust all ex-convicts on the basis of group characteristics not held by all members of the group. As long as there are some or many ex-convicts who will not commit further crimes,

being suspicious of *all* ex-convicts is wrong, so the argument goes, even if it is true that ex-convicts as a group are considerably more likely to commit crimes than the population at large. Because the word “prejudice” is pejorative, therefore, it is often used to refer not only to statistically unsound generalizations, but also to the inappropriate use of even statistically sound generalizations. But the fact that both of these quite different meanings of “prejudice” are widespread only serves to underscore the confusion.

Much the same can be said about the word “stereotype,” which has an even more negative connotation. As with “prejudice,” it seems wrong to use the word “stereotype” to refer to those cases in which a generalization is both empirically and morally appropriate.¹² But if stereotyping is wrong, and if condemnation is implicit in using the word “stereotype,” we have the same ambiguity about whether stereotyping is wrong only when the stereotype lacks any statistical foundation, or whether it is wrong also when stereotypes that do have a sound statistical foundation are used to make decisions about entire classes, including the members of the class to whom the stereotype does not individually apply.¹³

All of this suggests that there are serious and unresolved definitional issues that surround the topic of generalization, and suggests as well that there is a great deal of loose language being thrown around, especially when we are generalizing about the characteristics of human beings on the basis of identifiable attributes. And although there is little point to thinking that we can answer difficult problems about the use of nonuniversal generalizations simply by defining our terms, we can make a start in avoiding confusion by recognizing some of the definitional issues before we get much further into

the analysis. We should recognize, for example, that it may be misleading to describe both generalizations about the courage of gay men and generalizations about the recidivist tendencies of ex-convicts as prejudices, and that it may be misleading in a different way to describe generalizations about the self-confidence of Capricorns and generalizations about the meticulousness of accountants as stereotypes.

The reason we cannot resolve by definitional fiat the difficult issues surrounding the use of generalizations is that the relationship among universality, statistical soundness, and morality is exactly the matter at issue. That the terms in common usage, such as "prejudice," "stereotype," and "profiling," are ambiguous on the conceptual, empirical, and moral issues is not a solution, but a signal that we have a problem in need of analysis. This definitional detour, and the distinction it emphasizes between spurious and nonspurious generalizations, highlight the primary concern that pervades this book.

This primary concern, a concern that the definitional ambiguity between statistically sound and statistically unsound generalizations illuminates, is about the appropriate (and inappropriate) uses of statistically sound but nonuniversal generalizations. The problem of the statistically unsound generalization is not, to be sure, an unimportant one. Indeed, we will return to it repeatedly in considering when, if at all, it is appropriate to make decisions on the basis of even statistically sound generalizations. It is possible, after all, that a reluctance to employ even statistically sound generalizations will helpfully decrease the reliance on statistically unsound and dangerous generalizations. Nevertheless, our primary focus here will be on the statistically sound generalization, the generalization constructed around a trait or property

that is relevant for the purpose at hand. In focusing on the statistically sound but nonuniversal generalization, we will examine a widespread feature of decisionmaking, for the problem of whether, and when, to rely on such statistically sound but nonuniversal generalizations pervades our decisionmaking experience. As one usage of the word "prejudice" indicates, many people believe it wrong to make individual decisions on the basis of nonuniversal group characteristics even if the group attributions have a solid statistical grounding. From this perspective it is simply always wrong to make actuarial decisions. All human beings—teenage males who drive cars, ex-convicts, used-car salesmen, Scots, accountants, and everyone else—deserve to be treated as individuals and not simply as members of a group, so the argument goes, and actuarial decisions about human beings are in most instances morally wrong. But although this belief in the wrongfulness of reliance on even statistically sound but nonuniversal generalizations is widespread, it still may not be correct. Indeed, it may not even be plausible. In any event, exploring the legitimacy of relying on statistically sound but nonuniversal generalizations in making decisions is the theme that connects all the specific topics addressed in this book.

Is Particularism a Moral Imperative?

The goal of this book is thus to examine, and qualifiedly to defend, the practice of painting with a broad brush. In defending painting with a broad brush, I challenge an increasingly pervasive theme in contemporary thought. Sometimes this theme is described as "particularist," for the particular is the opposite of the general. Thus to make decisions on the basis of the characteristics of par-

ticular events or particular individuals, rather than on the basis of the characteristics of the groups or classes of which the particulars may be members, is often thought to be a moral imperative. Indeed it is often thought to define the concept of justice, and justice has long been thought necessarily to reside in particulars. When the philosopher Onora O'Neill maintains that "theorists of justice assume that interpretation will lead to differentiated action in differing cases," she describes an attitude not dissimilar to that endorsed by Thomas Hardy in *Tess of the D'Urbervilles* when he characterized Clare, Tess's husband, as having made the "mistake" of "allowing himself to be influenced by general principles to the disregard of the particular instance."¹⁴ And perhaps most blunt was William Blake, who opined that "to generalize is to be an idiot. To particularize is the alone distinction of merit. General knowledges are those knowledges that idiots possess."¹⁵ For O'Neill, for Hardy, for Blake, and for many others, maximum particularity is a characteristic of both justice and wisdom, and reliance on nonparticular categories or principles is at best a necessary evil, at worst an injustice, and all too often a demonstration of stupidity. Indeed, Blake's characterization of the generalizer as an idiot will be important to bear in mind, if only to make clear that the target is not one of straw. In countless walks of life, getting to the particulars of *this* situation or case or person is put forth as the ideal, with generalization seen as at best the lazy alternative. Lawyers celebrate deciding "one case at a time" and rail against the ways in which rigid rules interfere with the exercise of common sense in individual situations.¹⁶ Philosophers extol the virtues of particularism in making moral judgments.¹⁷ Feminist theorists often describe a qualified particularism as preferable to a traditionally

male belief in the superiority of generalization and abstraction.¹⁸ And psychologists frequently maintain that we should look at each person as an individual, with "going beyond categorization" being "the right thing to do in almost every case."¹⁹

Yet the distinguished recent and not-so-recent provenance of the particularist position may still not make the position right. Is reliance on generalizations simply the *modus operandi* of the lazy or the stupid? Perhaps so, but perhaps instead generalization is more inevitable and even more desirable than Blake and others have thought. Taking that possibility seriously is what joins the diverse topic considered in the following pages.

Because all the features of a particular event or decisionmaking occasion are commonly described as the *context* of that event or occasion, it is also common to refer to particularist decisionmaking as contextual. Looking at something in context is assumed to be good, and saying that someone has described a decision or a statement *out of context* is rarely a compliment. Nevertheless, part of what I seek to defend in this book is acontextual decisionmaking, decisionmaking that may deliberately ignore some of the context of the decision and some of the particular features of events or individuals, even features that may themselves be relevant to the decision to be made. As with the actuarial decisionmaking practices of the insurance industry, nonparticularist or acontextual decisionmaking is rarely defended. The fact that we have a battery of readily available pejoratives—*stereotyping, profiling, painting with a broad brush, ignoring the context, Procrustean, one size fits all*—shows just how strong the pull to particularism is. Yet on closer analysis it may appear that there is more to be said (and not only by the insurance companies) for actuarial decisionmaking than is

commonly supposed. However often painting with a broad brush may be denied and condemned, it may still turn out on closer inspection to be an often inevitable and frequently desirable dimension of our decisionmaking lives.

To put it most starkly, and to set aside for now some of the definitional questions discussed above, part of my goal here is to defend the practice of stereotyping or profiling. But this, as will become apparent in later chapters, is not to defend the practice of stereotyping or profiling by, say, race, or gender, or sexual orientation. In most cases involving generalizing on the basis of race, gender, and sexual orientation, the stereotyping is simply statistically unsound, as with the stereotypes that blacks are lazy, that women are bad drivers, or that gay men lack courage. Not surprisingly, nothing in this book is directed toward the defense of statistically or empirically unsupportable generalizations. More importantly, however, we will see why even many statistically supportable generalizations—women have less upper-body strength than men—often ought not, for compelling moral reasons, to be the basis for decisions that affect people's lives. Just why this is so will occupy much of chapters 5 and 6. Yet a large part of the motivation for this book is the impression that we have moved too quickly from condemning stereotyping or profiling on the basis of characteristics such as race, gender, and sexual orientation to believing in the moral impermissibility of all stereotyping and all profiling. Perhaps people have too hastily taken the wrong of racial stereotyping, for example, as simply one instance of the wrong of stereotyping itself. It is this belief—that stereotyping and profiling and generalizing are *always* wrong—that this book will

examine from numerous different angles, and it is this belief that in important respects this book will challenge.

Implicit in any defense of generality in decisionmaking is a defense of error. To make coarse-grained decisions is to make decisions in some particular cases that are less optimal than the decisions we might make in those cases if we focused only on the features of that case, and if we focused on all the relevant features of that case. Making such coarse-grained decisions—painting with a broad brush—is thus to make mistakes. Defending the making of mistakes is a formidable task. It may even be unwise. Yet as much of the philosophical and legal literature on rule-utilitarianism in particular and rules in general has shown,²⁰ creating a decisionmaking procedure predicted at the outset to make some number of errors will often lead to fewer errors in the long run than will creating a decisionmaking procedure that in theory produces no errors but that in practice produces many. Even more importantly, however, avoiding mistakes is often not the only goal. A decisionmaking procedure that makes more mistakes in the long run may serve other important goals and embody other important values.²¹ Just as we refrain from pointing out all the errors that our loved ones make, and refrain even when pointing out these errors would reduce the total number of errors made, so too might we wish at times to adopt, in the service of accuracy-independent values, decision procedures that can be expected to make some number of mistakes, and even more mistakes of some kind that would be made by other plausible decision procedures. It is often the case that generality in decisionmaking is just such a procedure, and so while part of this book's goal is to demonstrate that such a procedure often produces

fewer errors in the long run, another and possibly more difficult task is to show that the use of generalizations in decisionmaking—stereotyping—may sometimes be desirable even when it produces more and not fewer errors in the long run.

Although the analysis that follows is unavoidably a moral one, it is not the plan of this book to take off the shelf a moral framework or methodology and then apply it to the problem of generality. And it is certainly not the goal of this book to presuppose a strictly utilitarian or consequentialist focus on decisionmaking *efficiency* or *economy*. To show that it is often efficient to ignore individual differences is hardly a difficult task. Much harder, and much more important, is confronting the question whether ignoring individual differences is right, just, or fair even when it is efficient. Is ignoring individual differences unjust only when it is inefficient, and then only *because* it is inefficient? Perhaps efficiency does exhaust justice; but here we proceed on the assumption that it does not, and thus proceed on the assumption that there are important questions that must be asked about generality in decisionmaking that cannot be subsumed under efficiency-based or even more broadly utilitarian frameworks.

Instead of selecting a moral standpoint at the outset, we will explore a diverse collection of topics in which the problem of generality looms large. In examining the intuitions and arguments that might apply to these topics, we can generate a series of smaller moral propositions that might then be part of a larger moral framework. Rather than applying a moral framework to the problem of generality, therefore, this book proceeds on the assumption that more-particular solutions to the problem

of generality might help compose ~~that~~ larger moral framework, rather than the other way around. Perhaps ironically, therefore, the approach here to the problem of generality will have a strongly particularistic flavor. We will examine problems as diverse as those of basing insurance rates and other social consequences on genetic differences; controlling vicious dogs and other dangerous animals on the basis of breed characteristics; identifying the targets of customs inspections, tax audits, airport searches, and police stops by reference to sets of characteristics commonly referred to as “profiles”; allowing statistics to be used as evidence in civil and criminal trials; imposing mandatory retirement for airline pilots and other forms of discrimination on the basis of age; taking gender differences into account in military and related contexts; criminalizing behavior not because it is intrinsically wrong, but because it indicates the likelihood that wrongful behavior has occurred; prohibiting the use of traditional methods of food production; and imposing sentences on those convicted of crimes not on the basis of whether the punishment fits the crime or fits the criminal, but instead on the basis of a predetermined guideline. From these and many other varied instances there will, it is to be hoped, emerge a greater understanding of the problem of generality in decisionmaking, a more sensitive appreciation of the relationship between generality and justice, and, most importantly, an increased awareness of the circumstances under which generality in decisionmaking is appropriate and the circumstances under which it is not.