

Introduction

ALL YOU need to know to be a taxi driver is how to drive a car and find your way around the city, and sometimes not even that much. You pick up customers, drive them to where they want to go, get your due, and cruise off to collect yet another caller or to hunt for yet another hailer. Sometimes you negotiate the traffic, paying no attention to the passengers. Sometimes you vent your spleen on punishing taxation or corrupt local politicians to your captive audience. Occasionally you meet interesting or eccentric people, and sometimes you even act as a passenger's impromptu counselor. Then, when you have had enough after usually a long and tiring shift, you go home. "At the end of the day," as drivers might put it, the job is quite simply boring and does not take much out of you. Or does it?

In this study, we investigate how taxi drivers assess the trustworthiness of prospective passengers in two cities—Belfast, Northern Ireland, and New York—that differ in all respects except one: in both cities there are villains, such as terrorists and robbers, who pose as passengers in order to harm drivers. To appreciate just how costly mistakes in assessing prospective passengers' trustworthiness can be, consider that taxi drivers are the victims of the highest homicide rate of any occupation in the United States.¹ According to the National Institute for Occupational Safety and Health (1996), in the United States a taxi driver is sixty times more likely to be murdered on the job than the average worker. Taxi drivers are also victim to more violent assaults (184 per 1,000 workers) than any other occupation with the exception of police (306 per 1,000) and security guards (218 per 1,000) (U.S. Department of Labor 1999).² In New York City at least 250 livery cab drivers were murdered between 1990 and May 2002.

Theory and Hypotheses

Taxiing in dangerous cities dramatically intensifies the pressure on drivers' cognitive skills. Taxi drivers who work in high-crime urban areas

confront, every day and night, a series of hard trust-decision problems. Their profits depend on taking in as many passengers as possible. But typically, most passengers are unknown to them. In New York City, where the chances of carrying the same passenger more than once is low, a driver could meet as many as twenty to twenty-five new individuals on average every working day, which could add up to well over seven thousand new encounters a year.

Driving a taxi involves more than just a high number of transactions with strangers. Unlike sellers of train tickets or doughnuts on the streets, who also have many transactions with strangers, drivers are in an exposed position once they accept a passenger. They are on their own and unarmed, with one or more individuals sitting in close proximity, cruising sometimes unfriendly neighborhoods where there are few witnesses or helping hands. A study of 280 homicides of taxi drivers in the United States and Canada between 1980 and 1994 found that 75 percent of the victims were attacked by passengers inside the car (Rathbone 1994/2002).

Taxi drivers are thus under severe pressure to decide quickly, on the basis of only a little information acquired from either direct observation or the dispatcher, whether or not to accept certain passengers. Few other occupations are under a similar strain.³ Only police officers on the streets have perhaps a more demanding occupation in terms of the risks they face and the consequent importance of deciding how dangerous the individuals they meet really are. Yet law enforcers can hardly grumble since this is part of their job. Ready to face danger, they are trained to defend themselves, often armed, and able to rely on colleagues for backup. Not so the taxi drivers in hazardous cities. Ordinary men and women for whom the dangers are a nasty side effect of doing an ordinary job, they are interesting because they are more like us.

But is their ability to avoid dangerous situations purely a matter of trust? Given the fleeting nature of taxi drivers' encounters with strangers with whom they can exchange little information, one might expect taxi drivers to treat their problem as one of risk management rather than one of trust. Rather than ask, *Is this particular passenger trustworthy?* they could simply take general precautions to avoid or deter attackers, and they could think of remedies for limiting the damage in case of a bad surprise. They do take such precautions, as we shall see, but as a complement to treating their problem as one of trust rather than as an alternative to it. We devote a sizable part of our study to describing these strategies and discussing the extent to which they are adopted to limit the problem of trust.

An essential feature of our study is that we treat drivers' trust decisions in conjunction with the "sign-management" strategies of the other two protagonists of the game, the bona-fide passengers and the villains who mimic being passengers. We explore the dynamic relationships between drivers and their fares as they learn from their own experience and that of other agents and adjust their behavior accordingly. In dangerous cities, passengers must communicate that they have the capacity to pay and that they are not dangerous. For example, savvy drivers pick up passengers only at well-lit corners, not in dark alleys, and savvy passengers go to such safe places if they want to be picked up. Mimics learn and then also adopt such behavior in order to deceive drivers into picking them up. Our study explores how the three sides of this social transaction—drivers, customers, and mimics of customers—confront the challenges they face. They represent a near-perfect microcosm in which to research decisions about trust that arise in natural settings.⁴ The interactions of these players are at once simple enough to make them relatively easy to research and rich enough to contain the basic components of trust decisions.

The Basic Trust Game

The trust situation we consider is best conveyed by an example. Suppose a person extends a loan to another because she expects him to do his best to repay it, though it is clear that she would do better to refuse to make the loan if he will make no effort to repay it and it is also clear that his selfish interest is to make no effort. Then, we would say, she *trusts him to repay the loan*. The notion of trust that we adopt in this book is related to specific tasks and payoffs: the trust we have for someone in one context does not necessarily extend to other contexts (Bacharach and Gambetta 2001, 150; Hardin 2002, 9). In general, we say that a person trusts someone to do X if she acts on the expectation that he will do X when both know that two conditions obtain: if he fails to do X, she would have done better to act otherwise, and her acting in the way she does gives him a selfish reason not to do X.⁵

In any trust situation, there are no less than two players—the truster and the trustee. For the problem of trust to be a problem at all, however, the truster must be thinking about at least two possible *types* of trustee. One type of trustee is a person driven by "raw" self-interested impulses that lead him not to do X when given the opportunity. The other type is a person who has trustworthy-making qualities that motivate him to resist the pull of his raw self-interest and instead to do X (Dasgupta 1987, 53; Bacharach and Gambetta 2001, 148–52).

It is not enough to predict, say, that people will behave in a trustworthy manner if doing so is in their self-interest. This removes the problem of trust altogether by removing the tension between the two types of trustees, which is essential for being able to perceive a decision as one of trust. For if in a given situation the good trustee behaves in his self-interest, what would make the bad trustee misbehave in the same situation? Untrustworthiness would not be imaginable, and the problem of trust would disappear.⁶ Even if we assumed that people are trustworthy when the trustworthy act is in their self-interest, in order to have a problem of trust at all we still need at least two notions of self-interest: one that is, say, short-term or “raw,” untamed by any other consideration, and another that is long-term and “encapsulates” the truster’s interests as well (Hardin 2002), so that the truster is uncertain about which type of self-interest motivates the trustee in a particular situation.

The truster’s *primary trust dilemma* is thus her uncertainty over which type the trustee really is, the bad or the good type. On the one hand, to perceive that trust is a problem at all the truster must have in mind the possibility that the trustee is motivated by raw self-interest. At the same time, the truster can avoid being paralyzed by distrust only by admitting to the possibility that the truster has other qualities capable of overcoming his temptation not to do X.

The primary trust dilemmas can be applied to the driver-passenger case. Sometimes it is passengers who wonder whether the taxi driver is of the honest sort (we discuss this case later), but in this study we have looked only at the interactions in which the driver is the *truster*. It is thus the driver’s problem to decide which trustee-customer he is observing—the trustworthy one who will pay his fare and not harm him, or the untrustworthy one who only looks like a customer but is bent on harming him in some way ranging from not paying the fare to robbery or even murder.

We can make the interaction more precise by describing it as a simple game, which we refer to as the *basic trust game* (for a full discussion of the basic trust game, see Bacharach and Gambetta 2001). The payoffs of the three types of players in a basic trust game are presented in figures I.1 and I.2; the numbers are purely illustrative. In risky areas, the driver may meet a hailer of a bad type whose payoffs are expressed in figure I.1. Such a hailer wants to be picked up only to take advantage of the driver, since for him “exploiting” is better than “behaving”—that is, 4 is better than 1. In this case, the truster-driver is facing a bad customer, and “refusing” is better than “picking up” for him—that is, a payoff 0 is better than -3 . By contrast, if the truster-driver is facing a good customer (figure I.2) for whom “behaving” is better than “exploiting” (that is, 2 is better than -4), for him the best choice is to “pick up” rather than “refuse,” since 3 is better than 0.

Figure I.1 Driver's Payoffs and Bad Customer's Payoffs

		Customer	
		Behave	Exploit
Driver	Pick up	3, 1	-3, 4
	Refuse	0, 0	0, 0

Source: Authors' compilation.

The primary trust problem for the truster-driver lies in identifying which type the hailer or caller is—a genuine customer or a villain. This dilemma epitomizes countless situations in which we find ourselves in the position of either the truster or the trustee. However, there is nothing in the game as such that suggests a solution; the game simply describes a situation of uncertainty about the true nature of a trustee with whom it is to the truster's advantage to deal only if the trustee is of a good sort. The dilemma lies in the uncertainty over the real payoffs of the trustee-customers. Once the truster knows the type of trustee he is facing, the solution of the game is trivial. But how does the truster find out, and can the trustee do anything to inform him, truthfully or otherwise, about his type?

Trust and Signs

Michael Bacharach and Diego Gambetta (2001; see also Gambetta 2005) have articulated a new theoretical framework that can answer that question, and this study was inspired by, and aims to test, some of the ideas they developed.

The theory's crucial step is to establish a link between the trust dilemma and signaling theory, a branch of rational choice theory. How-

Figure I.2 Driver's Payoffs and Good Customer's Payoffs

		Customer	
		Behave	Exploit
Driver	Pick up	-3, -2	-3, -4
	Refuse	0, 0	0, 0

Source: Authors' compilation.

ever, the particular assumptions about people's motivations that are characteristic of the rational choice approach to trust play no role in the theory we develop here.

It is not essential to our theory that trust is the product of some system of rewards and penalties which act as incentives to the trustee in repeated interactions. . . . We treat trust as the product of underlying trustworthymaking character features—one of which can be, of course, the degree of susceptibility to rewards and punishments. (Bacharach and Gambetta 2001, 149)

Our approach, which we believe to be more realistic than the rational choice approach, accommodates a wide variety of sources of trustworthiness, which include not only one's self-interest but also moral principles, social norms, and even specific dispositions that, in a given game, can make one trustworthy. The source of trustworthiness depends on which trustworthy act we expect of others in a given situation.

The first step of our theory is to hypothesize that trusters acquire, in various ways, an idea of which trust-warranting properties a trustee needs in order to be trustworthy in a given game (as well as of the obverse properties that make a trustee untrustworthy). Some of these properties—such as honesty, benevolence, a long-term horizon, a pro-social upbringing—can make a person trustworthy in many trust games, while other properties have a limited range and suffice to make someone trustworthy only in doing something in particular, as, for instance, love of children, sect membership, or, in our case, simply being busy enough or wealthy enough not to bother causing trouble to a taxi driver.⁷ The notion of trustworthiness that we use refers to properties of persons that are not necessarily applicable to all situations but can be related to specific actions. We trust someone when we trust him to do one thing we have in mind, but we may or may not trust him to do something else. People who happen to be honest can be trusted in many different instances, but in our daily interactions we often trust people to do certain things for us even if they are not honest in all respects. Tony "The Ant" Spilotro's son Vincent could trust his father, a Mafia hard man, to come home and cook him breakfast every day even after Tony separated from his wife (Pileggi 1996, 157). Someone may be trustworthy as a husband and yet an opportunist with his colleagues or the Internal Revenue Service. Trustworthiness, in this narrow sense, can be sustained not only by virtues but also by nonvirtues. Laziness, fearfulness, or a physical handicap can lead us to believe that a person will not attack or cheat us if what is required for him to be untrustworthy in a given game is energy, courage, or dexterity.

However, the presence of trustworthy-making properties is not sufficient to induce trust. Trust has two enemies, not just one: bad character and poor information. The incidence of trustworthy-making qualities that sustain good character sets an upper bound on the amount of well-placed trust in a community. Actual trust, however, may fall below that threshold because of the difficulties of communicating the presence of these qualities. As we shall see, many passengers fail to communicate their trustworthiness to drivers and to be picked up. The trouble, in particular, is that the trustee's trust-warranting properties cannot be discovered from observation. One cannot see "self-interest," "honesty," or even "identity" as such.

One may say, "I could see at once that he was an honest fellow." But it is in *signs* that one *sees* it. One observes, for instance, physiognomic features—the set of the eyes, a firm chin—and behavioral features—a steady look, relaxed shoulders—and treats them as evidence of an internal disposition. . . . Trust-warranting properties may come variably close to being observable. But, except in limiting cases, they are unobservable and signs mediate the knowledge of them. (Bacharach and Gambetta 2001, 155)

Our second step is thus to assume that a person's trustworthiness can be assessed by *evaluating observable signs of him* that the truster believes to be correlated with the unobservable trust-warranting properties.

We suppose, realistically, that the vast majority of basic trust games include some observation by trusters of features of their trustees. Any piece of observable, or otherwise perceivable, behavior, including a message, an accent, or an item of clothing, counts as a feature. It follows then that every trust-involving interaction initiated by a trustee contains an episode of the kind described here:

A truster may observe the tweed jacket and shooting stick of a woman in a departure lounge, and may infer that she is to be trusted to look after the truster's bag while he posts a last-minute letter. Or, a truster may recognize her trustee, by her face, or signature, as someone who has always repaid loans in the past, and may infer that she has now, as then, those properties of character and liquidity which make her creditworthy. Or, finally, a truster, having heard of the trust-warranting properties of a person of a certain description, may try to match that description with the person in front of her: "Is this the person picked out by the signs of identification I heard of?" (Bacharach and Gambetta 2001, 155)

Mimicry

The use of signs to solve the primary problem of trust brings up a new dilemma for trusters. "The observable features (the tweeds, the signa-

ture, the name, the accent, the club membership) are sometimes good enough signs. Sometimes, however, it is not certain that signs of trust-warranting properties are to be trusted" (156). One reason for this is that the correlation with the trust-warranting property may be weak or nonexistent—one may simply be wrong about the link between signs and properties, as in the case of prejudices. Another, more interesting reason is that there may be untrustworthy trustees who endeavor to display fraudulently the signs that can convince the truster of their trustworthiness.

Suppose there is a sign *m*, such as looking people in the eye, that people take to be evidence of honesty. Deliberately displaying *m* is then a way of signaling that you are honest. The problem, however, is that both honest and dishonest people may have a motive for signaling *m*.

Rich people sometimes wear expensive clothes to show that they are rich; poor people sometimes wear expensive clothes to seem to be rich. Benevolent uncles smile to show they are benevolently disposed; wicked uncles smile to seem to be benevolently disposed. The deceptive instances of the strategy we call "mimicking." More precisely, a mimic of a property [*k*] is a person who does not have the property and deliberately displays *m* in order to be taken to have *k* by another. (Bacharach and Gambetta 2001, 157)

This framework is applicable also to *negative mimicry*—that is, "camouflage."

There are often signs of a trustee which are likely to be interpreted by the truster, rightly or wrongly, as indicating the lack of a trust-making quality, and so untrustworthiness. Both a trustworthy signaler who expects to be unjustly perceived if he displays such a sign *s*, or an untrustworthy one who is afraid of being detected if he does, have a reason to camouflage. That is, they can be expected to take steps *not* to show *s*. For our purposes we can therefore consider deceptive camouflaging as a special case of mimicking. The strategy of camouflaging the signs of untrustworthiness by suppressing *s* is just that of mimicking trustworthiness through displaying the notional sign "no *s*." (Bacharach and Gambetta 2001, 155n13)

We regard mimicry, both positive and negative, as a crucial component of most trust episodes in which trust is misplaced. In basic trust games, the truster reasons, as we argue he does, that there is a motive for an opportunistic trustee to mimic—that is, to emit signs of trustworthy qualities when he lacks them. Almost invariably when trust is abused, the trustee will have engaged in some act of deception aimed at making the truster think that he is the good type when this is not true. The trustee is a strategic player, not inert matter passively waiting to be appraised by the truster.

Although signs can solve the primary trust dilemma, they can also complicate the problem for the truster, who must judge whether apparent signs of trust-warranting properties are themselves to be trusted. The primary problem of trust (is this passenger a good or bad type?) is transformed into the *secondary trust dilemma*: is the sign a genuine sign of a certain trust-warranting property, or is it a mimicked sign?

Notice that the secondary trust dilemma applies also to repeated encounters in which we deal again with someone with whom we have dealt successfully in the past. In this case, the trustee enjoys a good reputation with us. Having been trustworthy in the past is his trust-warranting property in a repeated trust game. The problem is that his past trustworthiness is not directly observable; we infer it by reidentifying the person as the same person we dealt with in the past. But even the identity is not a directly observable property. We establish it through signs that uniquely belong to that person. Some of these signs, such as the face, may be resistant to mimicry, but mimics can imitate other properties, such as the voice or the signature. And this possibility gives rise to the secondary trust dilemma.

James Coleman (1990, ch. 5), who was among the first scholars to offer a sophisticated conceptualization of the notion of trust, discusses an episode taken from *The Merchant Bankers* by Joseph Wechsberg (1966) that is useful in conveying this point. The Norwegian manager of the merchant banker Hambros in London receives a call for help from a prominent ship owner in Norway who says he needs 200,000 pounds within the next half-hour or else one of his ships that was recently repaired in Amsterdam will not be released and he will lose a sizable profit. While holding the ship owner on the phone, the unfazed Hambros manager sends a telex message instructing the Amsterdam branch to pay the sum to the ship owner and have his ship released at once. A young German trainee banker who witnesses the exchange is amazed at the speed and smoothness of this transaction and asks the manager some pertinent questions: "How can you be sure that you really talked to that ship owner in Norway? It's easy to imitate a voice over the phone. How do you know he's good for two hundred thousands pounds?" (Coleman 1990, 91–93).

Coleman reports this key question—how does the manager know that he is speaking with the real ship owner rather than someone who is mimicking him?—but does not follow it up. Moreover, its importance is overlooked not only by Coleman but by most trust research, and thus the manner in which it is solved in real life has not been investigated. In the list of key features of a trust decision, Coleman includes whether or not to trust the ship owner, which is our first trust dilemma, but does not include the further essential question of whether he is *really* the ship owner. The two trust dilemmas need to be disentangled. The Hambros manager knows he can trust the ship owner to be trustworthy because

he knows him personally and knows that the ship owner has long-term interests and wants to behave well so as not to spoil his relations with the bank. In addition, the ship owner must have a special way of identifying himself to the manager, which reassures the manager of his identity. Similarly, our taxi drivers know that they can trust a *real* passenger to behave. But how do they know that a certain hailer *is* a real passenger rather than a mimicking villain?

Signaling Theory and Hypotheses

Our approach marks a departure from current theories of trust. The answers to that question in the literature vary. For some authors, trusters manage to appraise trustees' trustworthiness by a rather ineffable process. By contrast, our central hypothesis is that taxi drivers' decisions are supported by cogent reasoning. When asked, drivers often say their assessment of customers' trustworthiness is driven by "gut feelings" or a "sixth sense." Our expectation is that a logic underlies these feelings and that it consists of several cognitive steps, including an intuitive application of signaling theory. We expect that the subtlety and richness of the judgments a truster makes are due only to the complexity of the game that he is playing. The apparently ineffable nature of these judgments is an illusion that we hope to dispel by describing in detail drivers' cognitive processes.

For other, indeed for most, authors, trustworthiness is assessed by obtaining information about the real interests, constraints, and dispositions of the trustees and working out how we would behave in their position (Dasgupta 1988; Gambetta 1988, 217–24). Although this description reflects what we rationally do, it does not take into account the secondary dilemma—that is, whether and how we can trust the information we use for that reasoning. There is nothing in current approaches to say that trusters pay attention to signs; this implies, of course, that they are silent on which signs they disregard and which signs they heed.⁸ Our expectation is that drivers look for properties that are trust- or distrust-warranting ones in the game they are playing. More precisely, *they look for signs that manifest these properties*. Above all, we expect drivers to look for *reliable* signs of these properties.

The framework we use to conceptualize trust interactions has the advantage of allowing the application of signaling theory, which defines clearly what we mean by "reliable." This framework is a well-developed part of game theory, and it provides the general principles to which the information we collect about the trustees must adhere to enable us to solve the secondary trust dilemma.

A signaling game is said to have a "sorting" or "separating" equilibrium if, in equilibrium, a good type emits one kind of signal and a bad

type emits a different kind. No mimicry is possible in this case. The conditions for a signaling game to have a sorting equilibrium are well known. A genuine customer always emits a signal s , and a nongenuine one never emits s , if two conditions hold:

1. The benefit to a genuine customer of being treated as a genuine customer exceeds the cost for him of emitting s .
2. The benefit to a nongenuine customer of being treated as a genuine customer is less than the cost for him of emitting s .

We call these the “can” and “cannot” conditions. If the truster can see or hear the trustee, as in most basic trust games naturally occurring, the “can” condition is easily met since the trustee can communicate almost costlessly with the truster. Sorting therefore depends essentially on the “cannot” condition—on whether among these signals are some that nongenuine customers cannot afford to use. No poisoner seeks to demonstrate his honesty by drinking from the poisoned chalice. The best situation from a taxi driver’s perspective is when a genuine customer can afford to emit s and a nongenuine customer cannot.

Our general prediction is therefore that *drivers screen passengers looking for reliable signs of trust- or distrust-warranting properties*, in the sense that they look for signs that are too costly for a mimic to fake but affordable for the genuine article, given the benefit that each can expect in the situation. In other words, drivers will not be either erratic in the signs they watch for or easily satisfied by cheaply mimickable ones.

Some customers display signs that are not just unaffordable to the mimic but outright impossible to fake. In the economics jargon, they have an infinite cost. In ordinary parlance, they are quite simply constraints on mimics’ options. These signs, such as the face or DNA, are uniquely attached to the bearer and impossible for the mimic to reproduce, regardless of payoffs. Constraints are perfectly separating. In many real-world cases, however, we are not so lucky and find that most signals at our disposal are not perfectly separating. There are many signs that to some degree are more costly for a mimic (because of, say, penalties for their abuse) but are still emitted by a small number of mimics. Most of the time someone who looks like a Hasidic Jew will be a Hasidic Jew, but occasionally he may be a Palestinian suicide bomber. When the signal is only partly contaminated, it remains credible enough for a minority of mimics to gain from using it and for the majority of honest signalers to keep using it. In such an equilibrium, called “semi-sorting,” at least some truth can be transmitted. The prospects for even a semi-sorting signal are worse, however, in dangerous areas. The higher the base-rate probability of encountering a villain, the more credible, or separating, a signal needs to be to be effective. Sometimes, we further predict, no single sign

suffices and *drivers look for clusters of signs* that, if pointing in the same direction, may together come close to discriminating the good from the bad passengers. The cost of mimicry grows as the number of signs to be manipulated in order to persuade also grows.

In most practical circumstances, good customers display, as a matter of course, signs that are correlated with being a good customer. If one is middle-aged, white, and female, in face-to-face encounters one automatically reveals those features, and at no cost. If drivers know that there is a negligible probability that a woman will turn out to be a robber, they will look for signs of gender that the hailer displays. If persuaded that the hailer is a woman, they will decide to pick her up. These "cues" are "there anyway," and the genuine possessor of the right properties need take no action to manifest them. Indeed, she need hardly be aware of the cues' effect. Moreover, a feature of cues is that while costless for true possessors, they are very costly to produce and display for those who are not true possessors of the property, so they amply meet both the "can" and "cannot" conditions. If the properties to which cues refer are thought to make their owner trustworthy, the taxi business can proceed unimpeded, and trust is not perceived as much of a problem.

The trust game that taxi drivers play, however, is not always so smoothly resolved. It becomes more challenging, and thus potentially more revealing of agents' reasoning, under two conditions. The first condition is either an absence of cues or uninformative cues. For instance, in Northern Ireland religious affiliation often cannot be established just by looking at people. Most young males look alike, whether Catholic or Protestant. This is problematic for Catholic drivers, who fear becoming the target of a sectarian attack by picking up a Protestant man. In these cases, we hypothesize, *drivers "probe" further or customers volunteer additional information of a kind that a mimic would find hard to fake*. Because these information-enhancing activities are costly, the driver and the customer are more strategically aware.

Under the second condition, customers are genuine but cannot avoid displaying the "wrong" cues. For instance, young black males in New York or groups of young men in Belfast who are bona-fide customers face higher costs of signaling to drivers that they are bona fide, since their manifest features are associated with negative properties. To be picked up, we hypothesize, *they need either to display further reassuring signals or to invite probing* if they are to offset the effect of these negative cues. For example, they may stress their reassuring features: if the driver is Catholic, a passenger may choose to display or stress his own "Catholic" features. Or an African American passenger can send a white friend to hail a cab, as black English soccer players have had to do in London at night (Thorpe 1994).⁹ They may even equip themselves in advance of hailing a taxi with reassuring objects to display. For instance,

the black academic Lawrence Otis Graham buys and waves the *Wall Street Journal* to make drivers stop (Graham 1995).¹⁰ We expect, however, that it will not always be possible for customers to find signals of this kind that meet the “can” condition as well as the “cannot” condition, and even if they are genuine, they may end up having to walk home.

At the same time, we expect to observe drivers *behaving rationally and not being persuaded by easy-to-mimic positive signs*. We expect that drivers discount signs such as a smile or a necklace with a cross or a T-shirt inscribed with “Make Love Not War” and that they watch instead for signs that are cheap to observe but relatively hard to mimic. These could include a physical disability, a familiar face, or the “goodness” of the establishment from which the customer is emerging.

Notice, however, that there are some persuasive signals an honest customer can transmit only if the driver decides to stop and probe and not if he drives away. We therefore also expect that honest customers engage in signaling actions that are insufficient to reassure a rational driver fully—for a mimic too would find these actions cheap—but are nonetheless sufficient to make him stop. For example, a customer could remove or veil all the signs that are easy to remove or veil, such as holding his leather jacket on his arm to display the absence of concealed weapons or covering up a tattoo by wearing long sleeves. We also expect to find that bona-fide passengers with negative cues will induce drivers to stop not only by hiding their indicators of untrustworthiness but also, if that is not enough, by displaying some of the same “fraudulent” signs used by mimics. They will deceive, of course, but only in order to tell the truth: once the driver stops, they are then in a better position to convey the “separating” signals necessary to reassure him fully as to their true type.

By testing these hypotheses, we should be able to establish the validity of our general claims—namely, that drivers and customers conform to the signaling theory of trust by means of a wealth of detailed instances of suitable behavior that could hardly be cogently explained by other theories of trust.

So far we have mentioned signs and signals interchangeably without drawing a distinction between them, but there is one, and an important one at that. Signals are any observable features of an agent that are *intentionally* displayed by a signaler for the purpose of raising the probability that the receiver assigns to a certain state of affairs. A sign, by contrast, can be anything pertaining to a person that is perceptible and can modify our beliefs about that person while not being displayed with the intention of achieving that result. Signs, however, are dormant potential signals. They are the raw material of signals to which signalers who are aware of playing a strategic game can choose to resort. The basic form of a sign-signal transformation is a signaler taking steps to display a preexisting sign. One trigger of this transformation is the bearer’s realization of the meaning of

certain signs in the eyes of an observer. An individual may be unaware that his accent is informing others of a quality of his until an observer acts in a way that makes him aware.

The interactions between mimics, taxi drivers, and passengers tread the fine line between signs and signals. The mimics with villainous intentions are always aware that they are signaling, albeit fraudulently. They veil or hide their bad signs and forge and imitate the signs that a normal passenger would show. Sometimes an act of deceitful mimicry consists precisely of making a deceitful signal look like an innocent sign that a nonvillainous customer would carry without thinking. Taxi drivers can be taken in, but, as we hypothesize, they are generally on guard, screening first and then probing, if not persuaded, by encouraging passengers to display more convincing signals. By contrast, prospective passengers who are good types and carry no negative cues automatically persuade drivers of their goodness and are unaware that any signaling is going on. They become aware of their signs only when they expect to be or are refused by drivers. They may then try to discern what it is that puts taxi drivers off before resorting to conscious signaling that either masks their negative signs or displays new ones that offset the others.

Despite the difference, the reliability of signs and signals is governed by the same condition of signaling theory: signs, just like signals, transmit discriminating information to drivers to the extent to which they conform to the “cannot” condition.

The Acquisition of Background Knowledge

Signaling theory is an abstract tool, but some strong assumptions about the players’ background knowledge are crucial to the solutions it offers. Players need to know the size of both the benefits and the costs—for both possessors and nonpossessors of the trust-warranting properties—of emitting a given signal, and the truster needs to know the base-rate probability that a signaler has them. Players must also know a range of signs, their association with the relevant properties, and how others interpret them. If religiosity is known to make a person trustworthy, the players need to know which signs express it. They must know, for instance, that skullcaps are donned by male orthodox Jews, roughly how difficult it is for someone else to find and wear one convincingly, and that by convention women do not wear them. So although we hypothesize that players use strategies with the same abstract properties wherever they are, *the actual parameters of their strategies are very particular to their situation*. Like the devil, the knowledge required to play a trust game lies in the details.

Much background knowledge is learned simply as a by-product of living and observing life in a certain area, either by direct experience or through information sources that one trusts. A focused effort may be

needed, however, to learn this knowledge, or to teach others like themselves, when the area of operation is complex, the agents are new to it, or new situations and new kinds of informative episodes have arisen. This should be true, for instance, in cities in which many drivers are immigrants or in which passengers include not only residents but also transient people of diverse kinds. Even an experienced taxi driver from Bogotá has much to learn before he can feel comfortable in New York City. Mimics too have an interest in observing the behavior of both drivers and the genuine passengers on whom they plan to model their actions.

Although we do not make any specific hypotheses about knowledge acquisition and diffusion, our research pays descriptive attention to how real-world players acquire and update the knowledge needed for playing the trust game. Taxi drivers may receive formal training and be socialized by their organizations. They talk to other drivers, engage in queries with dispatchers, compare notes, and warn each other; older, more experienced drivers socialize younger, more naive drivers into the rules of the game and what to look for when deciding whether to pick up a passenger. Similarly, both mimics and customers experience success and failure and adjust their behavior accordingly. They read the news about incidents involving taxi drivers and learn whom to fear and how not to look. Mimics try to look tame and have to learn the salient features of “tameness”; so too do genuine passengers who have learned that something about them displays the wrong kind of sign and makes drivers avoid them.

Whereas signaling theory assumes that players have rational beliefs, we do not rule out the possibility that a certain amount of inaccurate or stereotypical knowledge permeates drivers’ “street-level epistemology,” as Russell Hardin (1993) puts it. A distinctive feature of our research is that we investigate empirically the actual judgment and decision-making processes that taxi drivers, passengers, and mimics use in such situations. In other words, rather than presume full rationality, our study explicates the actual content of the “street-level epistemologies” of the respective players about such decisions and works out the extent to which nonrational reasoning affects judgment and choice. For instance, in Northern Ireland there is a widespread and inaccurate belief that one can identify Catholics by the fact that their eyes are set closer to each other. In the United States some drivers may make decisions driven by crude racism. Or subtler dynamics involving, for instance, availability biases may be at work. For example, in the aftermath of an incident drivers may for some time become excessively prudent.

Thus, our approach does not rule out the possibility that the beliefs that enter into drivers’ applications of signaling theory may sometime be wrong. On the one hand, we seek to verify a series of signal-theoretic predictions for drivers’ and customers’ behavior. On the other hand, we do not exclude the influence on the decisionmakers’ judgments of

“extrarational” factors, such as their preexisting stereotypes. In particular, we expect drivers to be more rational in assessing positive signs than they are in assessing negative ones. Because of the well-known asymmetry between trust and distrust, false negative beliefs about the import of signs are likely to be more durable than false positive ones. Drivers do not pick up people displaying the former, so they do not get counter-evidence directly, and if all drivers hold the wrong beliefs, they do not even get it indirectly by communicating with each other.

In conclusion, our understanding of the steps that one needs to take to establish a trustee’s trustworthiness can be summarized in terms of three ways in which one can err and make the wrong judgment. These mistakes are hierarchically ordered in that if one commits the first mistake, the second will be irrelevant, and so on.

The first mistake is to look for the wrong trust- or distrust-warranting property—that is, a property unrelated to trustworthiness. This is the case when one entertains *unfounded beliefs*. One may think that Latino women are more dangerous than white ones, or one may infer from the property of one member of a group that all members of that group have that property, but neither belief is true. Or one may have the wrong causal theory. For example, one may believe that voting Democratic or being a Scorpio makes one more dangerous.

The second mistake is also due to wrong beliefs, in this case about the signs involved. One variety of this mistake is wrongly picking a sign of a trust-warranting property that is not in fact a sign of the property. A driver may be interested, for instance, in whether a fare is Catholic and may further believe that he can identify Catholics by their narrow-set eyes, which is an *unfounded stereotype*. Or he may misinterpret a sign that has multiple causes by picking the wrong cause: reading illness-related hand tremors, for instance, as the result of nervousness brought on by bad intentions, or interpreting dark glasses as an attempt to avoid eye contact because of nasty intentions rather than an attempt to hide grief.

The third kind of mistake—and the crucial one in terms of our theory—is to be persuaded of a passenger’s trustworthiness by a signal that, if the signaler is honest, is truly connected with the trust-warranting property the driver is looking for but that can be easily mimicked by a dishonest signaler. Assuming that the qualities drivers look for are truly associated with trustworthiness or untrustworthiness and that the signs they observe are truly signs of those qualities, there is still the question of how reliable the signs are.

When the Passenger Is the Truster

We should also briefly mention the trust game that is the reverse of the one we consider in this study: when the passenger is the truster and

worries about the type and intentions of the driver-trustee. This is a serious problem in many cities of the world where passengers are cheated, robbed, or raped by “taxi drivers.” For example, the London Metropolitan Police report that in 2001, 233 women were sexually assaulted by mimic-rapists who passed themselves off as minicab drivers (the London equivalent of Belfast private-hire taxis and New York livery cabs). That amounted to 24 percent of all rapes and serious sexual assaults in London in that year.¹¹ Mimicking can be of two kinds: a villain may mimic being a taxi driver when he is not by using a plausible car and showing up in the rich pickup points at the right time of day or night. Razaq Assadullah was jailed for eight years for raping a twenty-eight-year-old secretary on her way home after a night out with friends. The victim thought she was safe because rather than pick one of the drivers touting for business—the riskiest way of getting a ride—she had gone to the minicab office to get a car. But her rapist was working with a false ID, a fake driving license, and no insurance. The victim was also with a friend, but after Assadullah dropped off her friend, he drove her to a deserted street, climbed into the back of the car, and raped her (Weale 2003). Alternatively, genuine taxi drivers may switch to rapist mode. In this case, they mimic not being a taxi driver generally but being one on *that* occasion. The latter case is not so infrequent; the *Observer* reported on November 17, 2002, that “one police operation revealed that half the mini cab drivers [in London] stopped at roadside checks had criminal records, including convictions for serious sex and assault offenses.”

Originally we thought of investigating this side of the problem too, since it obviously needs to be studied, but decided against it. We needed to limit the size and scope of our fieldwork, and we could not have used the same sources to study both games. Also, drivers meet the problem of trustworthiness far more frequently than passengers do. One reason is the mere numerical property of this market: a passenger will get far fewer rides than a driver will supply, so the risk of any one passenger meeting a bad driver is, *a priori* at least, lower than the risk of any one driver meeting a bad customer.

A second, and more important, reason is that in cities where the taxi industry is regulated—such as both of the cities in this study—the problem of mimicking taxi drivers is limited. In particular this is so because the industry has developed standardized insignias that make it easier for passengers not just to identify a taxi in the traffic and distinguish it from other cars but also to distinguish a real taxi from a mimic taxi. In addition, personalized identifiers displayed in cars make it easy to catch a driver if he misbehaves. An aspiring mimic cannot easily pretend to be a genuine yellow cab, for instance. He would have to steal one, but once other drivers know the number of the medallion, the thief would not be able to go far in the city before he would be very quickly spotted.

His face would not match that of the driver in the photo card displayed in the cab, and passengers would quickly become suspicious. And one cannot easily set up a phony phone line pretending to be a livery car company just for the sake of taking advantage of passengers. One can put up a sign to pretend to be a livery cab or a minicab belonging to a nonexistent firm, but the chances of success will be lower. People will be familiar with the name and emblems of genuine companies and thus less inclined to trust taxis displaying unknown identifiers. The risk of encountering mimic-taxis goes up in cities in which most cars are unmarked and anyone can act as a taxi driver. The actual number of incidents may, however, be lower in the latter case, for passengers, anticipating the higher risk, may more often choose other means of transport or screen drivers' identities more carefully.

That is the case in Derry, the second largest city in Northern Ireland, which we briefly investigated. Over two-thirds of the 1,500 taxis that operate there are now illegal, according to Eamon O'Donnell, development worker for North West Taxi Proprietors Ltd. (NWTP), which aims to represent and regulate the taxi industry in Derry (personal communication, June 21, 2001). NWTP has proposed a "code of practice" that sets out obligations to be placed on all members of the taxi industry in the city, including depot owners as well as drivers, and a "bill of rights" for taxicab passengers entitling them to, among other things, a safe and secure journey. The NWTP proposal, which has not been accepted yet, includes the introduction of a unique logo to be shared by all firms. Within the logo, they plan to insert two numbers, one to identify the firm and the other the driver. The logo would then be either painted on the side of the car or posted on the windshield. Their suggestion would be that passengers only enter cars displaying this logo.

We found that one of the two main reasons to regulate the industry in that way is precisely that of reassuring customers who, according to our interviewee, use taxis less than they would for fear of being attacked by phony drivers. (The other reason is to limit the competition.) This proposal indicates an awareness of the importance of introducing hard-to-fake and easy-to-observe signs to limit villains' ability to mimic a taxi or a taxi driver, and it shows the importance of standardization as a means to achieve that.

Sources and Method

The core data for this study are of an ethnographic kind. The data were collected in both New York City and Belfast using partially structured interviews and participant observation with drivers, dispatchers, and passengers. The interviews were taped and transcribed; in analyzing the transcriptions, we focused on respondents' accounts of the signs they

look for or display and on the reasoning behind their actions. Additional data were collected from newspaper reports, taxi industry reports, statistical sources, and the academic literature.¹² We should stress that the two cities we selected, Belfast and New York, are both notorious for the risks they pose for drivers but differ with respect to the kinds of such risks. We expect that the strategies brought to bear on trust games will differ in interesting ways as well. The choice of cities was intended to provide us with a vivid contrast in the extent to which trust games depend on particular local conditions and the knowledge derived from them and, at the same time, to enable us to make a preliminary test of whether the principles underlying agents' "street-level epistemologies" (Hardin 1993) are, as we anticipate, the same everywhere.

We have no special attachment to the research method we chose. Its main drawback is, of course, that it does not allow a proper scientific test of our hypotheses, for we cannot control and measure the variables relevant to trust decisions. However, given the complexity of the cognitive steps involved in trust decisions, there are serious obstacles in designing controlled studies and employing subjects like taxi drivers. The designs ought to maintain a sufficiently strong incentive-compatibility and avoid narrowing or trivializing the stimuli so as to make the results relevant for real-world decisions.

We considered two alternative designs to test whether taxi drivers conform, in their acceptance decisions, to the predictions of signaling theory, notably that, where would-be fares may be dangerous, the cheaper it is to fake a sign of being a safe fare, the lower the acceptance rate. The two designs were:

1. Presenting drivers with photographs or "vignettes" of would-be fares and asking whether they would accept such a fare
2. Presenting drivers with a conventionally designed isomorph of the real-life trust game—for example, a game with money payoffs having the same strategic structure as the real-life encounter

In the end we decided against both designs. The first design suffers from two drawbacks: it is not "incentive-compatible," and more importantly, the display gives only partial information about the situation and is therefore bound to prompt questions about other, contextual factors that form part of the standard decision process for real drivers. For example, drivers might ask, "Where do I see this person? What time is it? Is there a football match in town?" These questions cover many dimensions, so if we supplied the information, we would then lack statistical power to disentangle their effects. If, conversely, we left them unspecified, we would have no control over what is really going on in the minds of the

subjects. We were no longer convinced that we could get clear confirming evidence of appropriate sensitivity to the costs of mimicry—not because drivers do not have this sensitivity, but because of the inevitable noisiness of the data that we would get from directing subjects' attention to real-world scenarios.

The second design, by contrast, lacks ecological validity and fails to tap into the specialized knowledge and “gut feelings” of taxi drivers for their real environment. Yet our central interest is in determining how adept taxi drivers are at making these decisions. Offering money payments would give incentives but change the frame (partly because the strategic situation is presented neutrally). These well-known difficulties about this experimental methodology, which is standard in experimental economics, apply here with special force because of our interest in whether signaling theory works in a particular form of real-life interactions.

Thus, the design of future research ought to tap into the decision-making know-how of taxi drivers on the ground while at the same time not stimulating the full complexity of their real-life heuristics; this complexity is too great for effective hypothesis-testing. Since we have not as yet clearly devised how to steer a middle course between this Scylla and Charybdis, we decided to limit ourselves to the ethnographic study. The fieldwork had the advantage of allowing us to observe the agents under their natural and often stringent constraint; in the lab no one is going to come after you if you get the wrong answer. It also allowed us to collect a wealth of vivid details on the “semiotics” of trust that no experimental or narrowly focused research could have yielded. We hope that our fieldwork, whose features we now describe, will inspire and guide further experimental and survey research.

Belfast

Much of the research work in Belfast went as planned; it took twenty weeks, from mid-April 2001 to the end of August 2001. First, we carried out an extensive search of five Northern Ireland newspapers over the past thirty years, looking for accounts of incidents in which taxi drivers had been involved as either victims or perpetrators of attacks since the so-called political “Troubles” began. We found forty-nine such reported incidents.

Next we interviewed forty-five taxi drivers in Belfast, including four who also worked as dispatchers and one person who worked only as a dispatcher. The bulk of the interviews were done with drivers from twelve private-hire firms from across the city. We also interviewed drivers who drove public-hire taxis based at Belfast City Airport and the City Center. In total in Belfast, we interviewed nineteen Catholic and twenty-six Protestant drivers.¹³

We contacted the drivers in a variety of ways. First of all, Heather Hamill reestablished contacts with three former taxi drivers she had met during previous fieldwork research in West Belfast. Our interviews with the former drivers proved to be a valuable pilot for this study by helping us to formulate our interview questions. These drivers also introduced us to a number of their colleagues, some of whom agreed to participate in the study.

To broaden the sample and contact drivers who worked in other areas, we put up posters and left flyers in taxi depots, but we did not get even one driver from advertising in this way. So we simply phoned for taxis. In a sense, we too mimicked an ordinary customer. We chose safe pickup and drop-off locations within each geographical area, such as a library or college. During the journey we explained who we really were and what we were doing and asked the driver if we could interview him, either there and then or at a later date. This strategy bore varying results. The first time we tried it, the driver's opening question, after asking Heather her destination, was, "So what do you do then?" To which she replied, "Well, it's funny that you should ask. . . ." He subsequently introduced us to his two brothers-in-law, who were also taxi drivers.

The novelty of our study impressed a number of our friends in Belfast, who, as they took taxis in the course of their daily business, took it upon themselves to advertise our work to the perhaps unlucky drivers who picked them up. Our sample thus snowballed outward as our interviewees passed us on to their friends and colleagues.

Other encounters were less positive. Some cabbies paranoically thought that we had targeted them specifically and were undercover police or government employees (for the Department of Health and Social Security, DHSS) checking on whether they were working while claiming unemployment benefits. Others thought we were from the Child Support Agency and had been sent by their ex-wife. Mark Jordan, our straight research assistant, had the additional problem of having his invitation to participate taken for a gay pickup line.

Because of Belfast's sectarian divisions (see table I.1), we needed a neutral location where we could carry out prearranged interviews. Fortunately, a local charity offered us a room in its offices located in a religiously neutral area near the City Center of Belfast. On other occasions, we went to taxi depots and interviewed the drivers as we sat beside them in their parked cars. The cars offered the drivers both privacy and the comfort of being in their own environment, and thus we were often seen jumping out of one taxi and into another as the next interviewee pulled up alongside.

We sat in the dispatch office of five different taxi companies and observed the dispatcher and the interaction between the drivers; we also drove around with five drivers while they were working. Heather Hamill's presence in the taxi prompted ironic comments from cus-

Table I.1 Belfast Interviewee Codes

Characteristics of Interviewees	Codes
Job or role	
Driver	—
Customer	Cust
Dispatcher	D
Bouncer	B
Religion	
Catholic	C
Protestant	P
Type of taxi driven	
Public hire	PH
Private hire	—
Area in which driver, dispatcher, or bouncer worked	
City Center	CC
West Belfast	WB
East Belfast	EB
North Belfast	NB
Gender	
Male	—
Female	F

Source: Authors' compilation.

tomers, particularly from late-night customers, many of whom thought she was the driver's untrusting wife or girlfriend. We were not sure whether this assumption was due to Belfast taxi drivers' reputations in matters of money or matrimony, or perhaps both.

We then interviewed customers and mimics: forty ordinary customers, ten "bad-looking" customers, and ten runners and three robbers who posed as good customers.¹⁴ We also arranged a group discussion with eight of the "bad-looking" customers. Finally, we took a large number of taxi rides posing as customers in a variety of locations and times of day and night.

The Belfast drivers' fondness for, and skill in, storytelling provided us with very rich data. Acting on information from one of our interviewees, we learned that staff from the Department of Environment (DOE) enforcement department, which polices the taxi industry, were posing as passengers in order to catch unlicensed "pirate" drivers and private-hire drivers illegally picking up fares off the street. We met with the head of the DOE's enforcement team to learn more about the strategies of these unexpected mimics.

We also interviewed the coordinator of North West Taxi Proprietors Ltd., who drove a taxi in Derry; five members of the Department of Vehicle Licensing Northern Ireland (DVLNI), the local government department that regulates the taxi industry; and four bouncers, or door-men, an occupation that faces screening problems similar to those faced by drivers.

New York

By contrast, not much went as planned in the New York City fieldwork. We worked harder and longer but gained less information of the type and quality we wanted.¹⁵ We focused on livery cab drivers rather than yellow cabs because the former, as we make clear in chapter 7, take much greater risks.

Despite our impeccable academic credentials—in addition to the University of Oxford and the Russell Sage Foundation, Heather Hamill also obtained a very useful attachment to Columbia University—we met with uncooperative and suspicious responses from all quarters. We cannot really say how much this response was intensified by the gloom that the terrorist attacks had caused, or whether perhaps this response was compounded by the anthrax scare that occurred during the first few months of our fieldwork.

In addition, the fieldwork suffered from specific shortcomings. First, we did not have informal contacts in New York, as we had among taxi drivers in Belfast. The people we knew were not the type to have personal contacts with taxi drivers—the total population and the social segregation in New York are much greater than they are in Belfast. We also encountered language barriers in the predominantly Spanish-speaking world of livery drivers. We found it difficult to make ourselves understood by the drivers we encountered when we phoned up a base and requested a taxi in the hope of recruiting a driver for an interview, as we had done in Belfast.

After a false start, we decided to go through more formal channels, starting with the Taxi and Limousine Commission (TLC), the regulating body of the taxi industry in New York. That agency, however, was not at all cooperative and showed little interest in the study. We were unable to interview any TLC employees and thus could get neither their views nor as much background data as was desirable. We made tactful, if relentless, requests for interviews until the last week of the fieldwork, to no avail. We were referred from person to person, promises went unfulfilled, and phone calls were not returned.

We turned to the New York State Federation of Taxi Drivers (NYSFTD), the largest union representing livery drivers. Although it took some time to make contact with this group, the strategy proved to be a

good one, and their help was invaluable. We interviewed forty-five livery drivers at NYSFTD offices in the Bronx, Queens, and Brooklyn. The bulk of the interviews were done with drivers from livery companies. We also interviewed two dispatchers and five drivers who drove yellow taxis based in Manhattan.

If Belfast drivers tended toward logorrhoea, the New Yorkers veered toward aphasia; they were less articulate and prone to provide monosyllabic answers. The deficiency of their responses was compounded by the language problem. Although we requested English-speaking drivers, a proportion did not speak English well enough to understand our questions or answer them fully, and our Spanish was rudimentary. The NYSFTD kindly provided us with a translator, who helped immensely, but the involvement of a translator prevented the free-flowing conversational style that elicited many of the rich details we obtained in Belfast, and many of the finer points were lost in translation. (We used the translator for seventeen of the interviews with drivers.) Furthermore, even with the help of NYSFTD staff, some of the drivers did not really understand our interests and were nervous and suspicious. Some doubted our intentions, wondering if we worked for the government or a newspaper and were keen to confirm that they had legal status in the United States. Although we interviewed forty-five Hispanic drivers, only thirty-two of the interviews were of sufficient quality to be usefully analyzed.¹⁶

We then turned our attention to drivers who worked in northern Manhattan. These drivers, who were mostly West African or African American, did not belong to the NYSFTD or any other union, so we had no official channels by which we could gain access to them. Instead, we decamped to the Metro North station at 125th Street and Lexington Avenue in Harlem, where we approached drivers waiting for passengers leaving the station. Although we had a varied reception, we were able to interview eighteen drivers using this strategy.

In addition to interviewing sixty-three drivers, we interviewed ten ordinary customers, twenty "bad-looking" customers, and six NYSFTD employees. We also took a large number of taxi rides posing as customers in a variety of locations and times of day and night.

The one aspect of our research in New York that did go better than in Belfast was the collection of material from newspaper archives. We did an extensive search of the *New York Times* and the *New York Post* since 1990, looking for accounts of incidents in which taxi drivers (both livery and yellow cab) had been involved as either victims or perpetrators of attacks and also looking for accounts of changes in policy and legislation that had affected the taxi industry. This yielded a great deal more information, both in quantity and quality, than our newspaper search in Belfast. We collected 128 cases of drivers murdered or violently attacked

and 59 cases of individuals accused of attacking a taxi driver. We also supplemented our data by subscribing to an Internet newsgroup to which taxi drivers contribute from the United States, the United Kingdom, and Australia.

Ideally, we would have liked to interview mimics of the worst sort—namely, those who were or had been arrested for attacking drivers, with either sectarian (Belfast) or robbery (New York) motives. In the time scale and with the resources available to us, we were unable to track these people down. Most of the information we have on mimics is derived from other sources, either the drivers themselves or the accounts of various incidents that we collected through our newspaper search.¹⁷

Anonymity

In return for the opportunity to tape-record the interviews, we promised the drivers that we would not reveal their identity. We removed any reference to specific taxi firms in Belfast or taxi bases in New York, and we gave each driver a code. In Belfast (see table I.1), the code provides information about the drivers' religion, the type of taxis they drove, the areas they drove in, and their gender. The private-hire drivers are coded first by religion, followed by area in which they worked, their gender (if female), and then an identifying number. "C-WBF3" tells us that the driver was Catholic, worked in West Belfast, was female, and was interview number 3. Public-hire drivers are coded by religion and a number. "P-PH5" describes a Protestant driver who drove a public-hire taxi and was interview number 5.

Customers are coded by religion, identified as a customer, gender-coded if female, and assigned a number. "P-Cust2" tells us that the respondent was Protestant, a customer, and a male and that his interview number was 2.

Dispatchers are coded by religion, identified as a dispatcher, coded by the area in which they work and their gender (if female), and given a number. "C-WBDF2" tells us that the respondent was a female Catholic who worked as a dispatcher in West Belfast and that her interview number was 2.

Bouncers are coded by religion, identified as a bouncer, coded by gender (if female), and given a number. "P-BF6" refers to a female Protestant bouncer whose interview number was 6.

In New York (see table I.2), drivers are coded by area in which they worked, gender (if female), and type of taxi they drove, and they are given a number. "BX7" is the code for a male driver who worked in the Bronx and was given the interview number 7. Customers are coded by skin color, identified as a customer, coded by gender (if female), and given a number. "BCustF1" refers to a black, female customer who was

Table I.2 New York Interviewee Codes

Characteristics of Interviewees	Codes
Job or role	
Driver	—
Customer	Cust
Dispatcher	D
Skin color	
Black	B
White	W
Hispanic	S
Type of taxi driven	
Yellow cab	YC
Livery cab	—
Area in which driver or dispatcher worked	
Bronx	BX
Brooklyn	BR
Queens	Q
North Manhattan	NM
Gender	
Male	—
Female	F

Source: Authors' compilation.

interview number 1. Dispatchers are coded by area in which they worked, identified as a dispatcher, coded by gender (if female), and given a number. "BRD2" is the code for a male dispatcher who worked in Brooklyn and was interview number 2.

Outline of the Book

The book is divided into two parts, one for each city, beginning with Belfast. Within each part we follow the same sequence of chapters. In the first chapter of each part, we describe the city and the toll that taxiing has taken on drivers. We give an account of the origins and motivations of the drivers and of the general dangers they face when driving in their city. Finally, we describe how they become knowledgeable about the city in which they work and how they exchange information about the dangers they face. In the next chapters (chapters 2 and 7), we give an account of the various types of mimics who pose as bona-fide passengers in order to harm the drivers in some way and describe the strategies that mimics use to persuade drivers to pick them up. In chapters 3 and 8, we describe the vast array of general precautions, deterrents, and

remedial actions that drivers take to protect their safety. In chapters 4 and 9, we discuss the various ways in which drivers screen passengers, the signs they consider, how they reason about signs, and the choices they make. In chapters 5 and 10, we look at how drivers further probe passengers, either before picking them up or after they are in the car. We also look at the problem of probing from the point of view of bona-fide customers who display negative cues and want to persuade drivers to pick them up. In the conclusion, we discuss our findings from the research, both expected and unexpected.

