

PALTERING

Frederick Schauer¹ and Richard Zeckhauser²

To Immanuel Kant (1795, 1797) (as well as to Aristotle, St. Augustine, St. Thomas Aquinas, and countless others) lying was simply and absolutely wrong. For Kant, the morally impermissible act was the intentional assertion of a knowingly factually false proposition, and a lie's moral impermissibility derived from its incompatibility with the liar's own dignity as a human being. By contrast, others have located the wrong of the lie in its deprivation of the victim's ability to choose and thus of the victim's autonomy (Bok 1978). And still others have assessed lying's impermissibility in utilitarian terms (Bentham 1791 pg. 170). But regardless of whether it be through a focus on the liar, or on the victim, or on the well-being of society at large, we have inherited a venerable tradition which has little hesitance in condemning lying as wrong.

Although few nowadays subscribe to the Kantian view in its unalloyed form, lying is still widely considered morally objectionable. Yet despite the virtual unanimity of the view that lying is presumptively even if not absolutely wrong, it is not entirely clear which component of a lie is the principal contributor to the lie's wrongness, or whether one or more of those components is sufficient on its own. To be more precise, we commonly understand a liar to be someone who (1) intentionally utters words that he or she knows to be false, (2) where what is uttered is in fact literally false, and where utterance of the literally false words (3) produces the effect of the listener believing or being likely to believe in the truth of something that is not in fact true.³ A lie in its full glory, therefore, involves elements of intent, literal meaning, and effect. When one or more of these elements is missing, we would be hesitant to designate the activity as lying, even though it might in other ways be morally questionable or socially detrimental.

Just as a lie involves the confluence of all three of these elements, so does the absence of all of the three constitute truth-telling in its purest form -- the sincere assertion of a proposition both believed to be true and actually true, under circumstances in which the assertion either creates or reinforces a belief by the listener in a true proposition. Yet although we can thus distinguish truth-telling from lying, numerous statements are, intriguingly, neither lies nor truth-tellings. Rather, they make up a universe of morally and socially problematic statements and propositional actions in which one or more of the elements of the genuine lie is missing, but in which one or more of the elements of authentic full-bore truth-telling is missing as well. Sometimes we talk about “misstatements,” and sometimes we (or the law) criticize or punish people for “misleading” others, understanding that in such instances we are concerned with the effect on the listener more than with the moral worth of the speaker or the literal meaning of the words used.

Our goal in this paper is to explore this area of less-than-lying, and to focus in particular on the widespread practice of fudging, twisting, shading, bending, stretching, slanting, exaggerating, distorting, whitewashing, and selective reporting. Such deceptive practices are occasionally designated by the uncommon word “paltering,” which the American Heritage Dictionary defines as acting insincerely or misleadingly. Although the intended effect of a palter is the same as that of a lie, both the dictionary definition of “palter” and everyday usage of related ideas make a palter somewhat troublesome, but make it something that is not quite a full-fledged lie as well. More specifically, the palter falls short of being a lie in two important dimensions.

First, the palter may not be literally false. One of us is a furniture maker and is known by his friends to be a furniture maker. When a friend comments on the excellent workmanship of a store-bought desk in his office and he responds by saying “thank you,” he has paltered, because he has left the false impression that he made the desk himself, even though none of his words said so explicitly.⁴ Relatedly, and more commonly, paltering is somewhat more active and creative. The other of us is a tournament bridge player. When asked about his successes, he

might remark: When I played in the World Pairs Championship recently, we got to the finals, intentionally omitting mention of the fact that in pairs tournaments a sizeable fraction of the pairs – 10% in this case – reach the finals, and omitting to mention as well his more frequent worse results. Thus, the typical palter achieves its misleading effect without the use of literal falsity.

In addition to not relying on literal falsity to produce its misleading effect, the typical palter, like the ones just noted, often seems at least slightly less harmful than the typical lie. We treat palterers as sometimes unfortunate and sometimes not, but to accuse someone of paltering - even assuming he knows what the word means - falls short of calling someone a liar. Indeed, in many cases we make no charge at all against the palterer, in part because we are embarrassed to have been fooled, but in part because the Kantian legacy has left a residue of belief that the lack of literal falsity is indicative both of lesser harm and of a lesser wrong.

Although palterers often escape unscathed and even uncriticized, part of our claim in this paper is that it is nevertheless a mistake to understand palterers as being generally harmless or near harmless. Often a particular palter seems as wrongful as would a lie in the same or similar circumstances, and often its consequences are as severe. And occasionally the expected harm of a particular palter may be greater than the harm of a similar lie, *ceteris paribus*, just because palterers are more likely than liars to escape detection. And insofar as a particular palter, even if detected, is less likely to be subject to either legal or non-legal blame, the expected harm may again on occasion exceed the expected harm of an analogous lie. These factors together suggest that the use of a palter is often as reprehensible as the use of a lie, and perhaps at times more so just because choosing to palter rather than to lie, assuming equivalent intent to deceive and equivalent harm to the victim, is typically a much safer strategy.

Given that the expected punishment of a palter is typically minimal, it is no surprise that paltering is widespread, and although a tally would be difficult to prepare, we would conjecture that the aggregate palterers in the world do more harm than the aggregate lies precisely because they are under-punished and under-deterred as compared to pure lies, and hence are widely employed. In addition, palterers are likely to be common relative to lies because though there may

be only one way (or just a few ways) to lie about a proposition, there may be many ways to palter. Palters can abound therefore, even when effective lies are likely to be relatively rare.

In seeking to explore the dimensions of paltering, this paper is motivated by our belief that individual palters are not only often more harmful than lies, but also that paltering in the aggregate is almost certainly much more common than lying. Yet although paltering is more common than lying, the fact that it exists largely outside of the reach of the law, and indeed often outside the reach of the harshest of non-legal social sanctions, may explain why it has not been well studied. But if paltering is as harmful and as widespread as we believe, and if the typical palter leads the recipient – the palteree, if you will – to have an incorrect perception of reality, then the omnipresence of paltering may in fact be more of a social problem than many people suppose. The paper thus aims to explore the practice of paltering, alongside of and in contrast to its cousin lying, and to examine the various ways in which the two activities might be controlled.

I. Definitional Preliminaries

Bearing in mind the three-part definition of the genuine lie as involving an intent to deceive, the use of words that are literally false, and the presence of a recipient who is caused by the lie to have a misimpression of reality,⁵ we can hypothesize the relaxation of one or two of these three components. If we relax the requirement of intent, we wind up with the category of statements in which speakers say something they believe erroneously to be true, in which that belief leads them to say things that are literally false, and in which those false statements lead listeners to have a false view of some actual state of affairs. Under many accounts, for example, the statements of President Bush about Iraqi weapons of mass destruction fit this characterization, for it is unlikely that the President actually knew there were no weapons of mass destruction at the time he announced that they then existed. And although the President might be charged with making a negligent misstatement, and thus with being somewhere between slightly and highly imprudent, we would not, contemporary political hyperbole aside, normally call the unintentional misstatement a lie, no matter the degree of negligence or

recklessness. And this requirement of actual intent to deceive is consistent with the common law of fraud, at least in its purest traditional form, which requires that there be some actual intent to mislead in order to support legal liability (see Alexander & Sherwin 2003). When the actual intent to deceive is missing, there is no fraud and there is no lie.

The unintentional misstatement does not amount to fraud in law, even when the other requirements of fraud are satisfied, but there are circumstances in which the law does impose other sorts of liability for unintentional misstatements. The common law of libel, for example, made libel (and slander) strict liability torts, such that the mere utterance of a reputation-harming falsehood could support legal liability even absent knowing falsity, and indeed even absent intent to harm and even absent negligence. Similarly, various aspects of the securities, consumer protection, and food and drug laws impose liability upon the makers of false statements even where there exists no intent to deceive.

But although the nature of legal liability or moral responsibility for the unintentional utterance of a falsehood on which someone relies to his detriment is an important topic, it is not our topic here. Paltering, as we understand it, is an intentional act, even though the act is different from (although not necessarily less harmful than) lying, and paltering is typically an act intended to mislead or to defraud.⁶ Thus our concern is with a certain form of calculated deception, and not with acts that unintentionally, even if negligently or recklessly, mislead others.

Nor is our concern in this paper presented when the listener is not ultimately led to believe something that is not true. In some circumstances, listeners will understand falsehoods as just that, and then even an intent to deceive coupled with a false statement will still produce no misimpression. This is how, for example, we decode inflated letters of recommendation, or adjust for the non-representativeness of what are described on resumés as “representative” lists of publications. In other circumstances, speakers may exaggerate or distort when they believe a listener’s existing beliefs make a clear understanding or evaluation of the unadulterated truth unlikely. For example, parents may overstate the dangers of drinking, driving, smoking, or sex to

children who probably would underestimate those dangers.⁷ Such misstatements, intended to compensate for biased inferences, is a category of some interest, but not to us here, largely because such statements, by producing accuracy rather than misperception, are far from obviously morally, socially, or legally problematic.

Finally, we are not interested here in the genuinely beneficial social lie or white lie. Telling sick people that they are looking better, or saying to our spouses that their clothes are becoming, or declining an invitation by inventing a prior engagement rather than truthfully expressing distaste for the host are lies, pure and simple, but they are lies whose consequences are either beneficial or at least believed, sometimes *ex ante* and sometimes *ex post*, by the liar to be likely beneficial. Palterers may be employed in the same fashion, and may receive the same assessment. But although white lies and white palterers comprise a fascinating topic, it is again one that we are content to leave to others.

By contrast, our interest is in the category in which the second criterion for a lie is relaxed while the first and the third are retained. That is, we are interested in the category of statements in which the speaker intends for the listener to have a misimpression, and in which the listener does wind up with a misimpression as a result of the speaker's statement, but in which the connection between the speaker's intent to deceive and the listener's state of having a misimpression is not the literal falsehood, as in the true lie, but something short of literal or exact falsity. This is the category of the successful palter, and this category will be our primary focus in the balance of this paper.

II The Varieties of Paltering

Examples of paltering are widespread. Sometimes paltering seems to occur when people take advantage of vague language, as when a used-car dealer describes a car as a "cream-puff," when a real estate agent describes a location as "prestige" or "highly desirable," when a stock broker describes a new public offering as "hot," or when a restaurant announces on the menu that its own signature pastrami and cheese sandwich is "famous." But since almost all of the

consumers of such loose exaggerations are well aware of the practice,⁸ and consequently accept the vagueness of the language used, it is not clear that vague language by itself, in most circumstances, even qualifies as mild paltering.

True examples of paltering, therefore, are the ones that occur when something the palterer says (or does not say) or does (or does not do) is intended to leave the impression that a specific state of affairs obtains, and when the recipient as a result of the palterer's actions or inactions believes that this specific state of affairs obtains, but when in fact the reality is quite different and perhaps just the opposite. Often the recipient's misimpression is a consequence of failing to correct a wrong impression. People who physically resemble a famous person may be treated better than they would otherwise expect and, knowing what is happening, may simply let the misimpression stand. So too with people who share names with the famous.⁹ And examples like this only scratch the surface, with people often taking advantage of the countless misimpressions that are all around us, and that are often held by those with whom we deal.

Much more common, however, is creating a wrong impression through deliberate action. Ph.D.s will often make restaurant or hotel reservations as "Dr. So-and-so," hoping in the process to lead the establishment to believe that they are (typically wealthy) physicians and not (typically non-wealthy) academics. Advertisers wishing to draw consumers' attention to the contents of an envelope will frequently put a government warning about tampering with the mail on the outside of the envelope, and will at the same time omit a return address, thus intentionally attempting to create the misimpression that the envelope contains an official letter from a government agency. People will often refer to a famous person by his or her first name, attempting to create the impression of close friendship. The Internal Revenue Service is alleged deliberately to select the period immediately preceding the April 15 tax filing deadline as the time to initiate tax fraud criminal prosecutions and to send out routine press releases about audit practices, presumably hoping in the process to lead taxpayers to have a subjective probability of audit and of criminal prosecution that is considerably higher than the actual objective probability of those occurrences. And politicians will often take advantage of the availability heuristic (Tversky & Kahneman

1974; Tversky & Kahneman 1973) by presenting extreme and unrepresentative examples of various problems and benefits -- discussing Willy Horton, for example -- hoping thereby to lead their listeners into a predictably mistaken generalization. These are but a small number of examples, but even such a short list should be sufficient to establish that paltering – the deliberate attempt to create a misimpression in someone by means other than by uttering a literal falsehood – is as widespread as it is interesting.

III The Political Economy of Paltering

Most people would much prefer to palter than to lie. Perhaps this is due to their upbringing, and perhaps the preference has other causes. But whatever the cause, it is hardly clear why paltering should be thought preferable to lying, or deemed less reprehensible. The law is of course concerned with problems of proof, and so much of law's tolerance for paltering is likely less a function of paltering not being a social wrong and more a function of the difficulty of proving that it has occurred. The current state of the law aside, however, once we understand paltering as involving the same intent as lying and the same effect as lying, and lacking only the exact mechanism of literal falsehood, it is hard to understand why the law, except for the obvious problems of proof involved, would be less concerned with paltering than lying.

Indeed, if we forget about the law for a moment, and just think more generally about the political economy of paltering and lying, we might suspect that paltering is in some sense worse than lying.¹⁰ Because lies involve literal untruths, they are easy, or at least easier than palters, to identify, whether legally or through loss of reputation. And because lies are easier to identify with some certainty, they are also easier to punish when identified. Conversely, because palters are harder to identify, there is a considerable incentive for those who wish to deceive others to turn to paltering rather than to lying. First, the personal discomfort from paltering is likely to be less than that from lying. After all, no one talks about galvanic skin responses attending the clever misrepresentation. Second, the same problems of proof that may lead the law to be comparatively unconcerned with paltering may make paltering easier to get away with than

lying, the law aside. Thus, those who intend to deceive will have multiple incentives to palter rather than lie. But if the harm of the palter is no less than the harm of the lie, then the very fact that it is safer to engage in one rather than another equally harmful act would suggest that it is the safer but equally harmful act – the palter – that is likely to become the greater social problem.

This comparative propensity to palter rather than to lie is likely exacerbated by the way in which many – probably most – people have somewhat of a hard time telling a straight (and non-white) lie. Whether this reluctance is caused by hard-wired moral sensibilities, or, more likely, by socially reinforced condemnation of lying as such (George Washington refused to lie about the cherry tree, we were taught,¹¹ but we were not taught what he might have done had there been paltering options available¹²), it seems plain that most of us have or have developed an internal “reject” button that makes it hard for us to lie. But this internal reject button exists within a psyche that also resists taking actions that are to our detriment, and so the palter often emerges as the self-interested but internally palatable alternative to lying.

Although there are thus incentives for people to palter, it is also likely that the cost of identifying a palter is higher than that of identifying a pure lie, in part because many palters, like President Clinton’s statement that “I did not have sexual relations with that woman,” have the element of deniability.¹³ That is, the palterer can, if exposed, often claim, unlike the liar, to have been misunderstood. Because careful wordings tend to be misheard often in direct hearing, and even more often in second-hand accounts, the palterer gets the benefit of the likely understanding coupled with a defense based on the exact and careful wording actually used.

Deniability, combined with the absence of a plainly demonstrable falsehood, thus makes it more costly to identify a palter than to identify a lie. And so if it is cheap to palter and expensive to identify a palter, and beneficial to palter if undiscovered, we can expect the practice to be widespread. And thus we find ourselves with a practice that is often beneficial to those who engage in it, difficult to identify, and hard to penalize through the law and outside the law, all of which combine to make the practice likely to be common. Thus, the palter presents almost all of the same harms as the lie, but because the lie is easier to identify, easier to penalize with

and without the law, and is subject to internal controls on its use,¹⁴ it may turn out, surprisingly perhaps, that lying is less common and less of a problem, in the aggregate, than paltering.

IV Penalizing Paltering

Although neither the law of perjury nor the traditional law of fraud penalizes paltering, the situation is changing. The securities laws, for example, penalize “material” omissions,¹⁵ and in modern times civil penalties for deceit or fraudulent misrepresentation cover conduct as well as words, and encompass a wide range of non-disclosures, passive acts, half-truths, and evasions (see Restatement (Second) of Torts §550 (1977) (literally true statements intentionally creating misimpressions); §527 (partial truths); see generally, Keeton, Dobbs, Keeton, & Owen 1984 pg. 725-738). So too with crimes of larceny, where nowadays the crime of obtaining property by false pretenses is sufficiently broad as to include a variety of ways in which those with larcenous motives may induce false beliefs in their victims (see Model Penal Code §223.3).

Yet although the law is broadening, it still remains a narrow and (comparatively) rarely used weapon against paltering. The law indeed has its place, but it often plays a subordinate role in the control of anti-social behavior, especially as compared to the role played by social norms, informal sanctions, and various other extra-legal remedies. And because of this, we turn, at least here, primarily to the question of non-legal remedies for paltering, and to the crucial role played by reputation in explaining the dynamics of paltering.

In considering these remedies, a number of interrelated dynamics come into play when we are considering paltering. First, the palterer is likely to be better off where paltering is more effective. Moreover, the effectiveness of a palter will depend both on how difficult it is to distinguish a palter from the truth as well as on the percentage of palterers. When distinguishing palterers from the truth is difficult, and when the number of palterers is small, paltering is likely to be more effective.

Consider the question of frequency. For some activities, like driving in excess of the speed limit, the incentives to break the law vary positively with the prevalence of the activity,

because the prevalence reduces the probability of enforcement, and at times even of detection. The driver on the Massachusetts Turnpike who drives at 78 miles per hour (the speed limit is 65) along with scores of others is far less likely to be apprehended than the driver who drives at 78 when everyone else is going no faster than 66. If all motorists were identical, and if prevalence increased attractiveness, there would then be two equilibria, one in which no one exceeded the speed limit and another where everyone did. But for other activities, like murder, the incentive to commit is largely independent of the prevalence of the activity. Whether there are more or less murders is largely irrelevant to the individual murderer.

Unlike speeding, where increased frequency increases the incentive to participate, however, and unlike murder, where the incentive is independent of frequency, paltering is an activity where the practice becomes less attractive as more people participate. As used-car dealers and rug sellers have come to lament, paltering is far more effective when only a small number of people engage in it. And indeed the limiting case is the one in which the behavior is so widely expected (and accepted) that no one is deceived and we do not have paltering at all. But even short of this limiting case, the more common is paltering, the harder it is for the individual palterer, since recipients will set higher standards for accepting some statement as true. A palter will thus be most effective in an environment in which palters are rare, and hence where recipients are especially likely to be trusting. An individual palterer, therefore, has an interest in there not being much paltering.

An additional problem, however, is the way in which this dynamic limits the effectiveness of sanctions. Sanctions, when effective, reduce the incidence of the behavior they are directed against, but this reduced incidence, at least in the case of paltering,¹⁶ will, as we have just seen, increase the incentive to palter for an individual palterer. Thus, the less paltering there is, the more a potential palterer will have an incentive to engage in it, thus limiting the effectiveness of any scheme of sanctions, and suggesting that there is likely to be an irreducible positive level of paltering.

This dynamic works the other way as well, however, and thus it might initially appear

that paltering is to some extent a self-enforcing, self-limiting, and self-correcting activity. The more paltering there is, the higher the cost (greater probability of detection, more elaborate palterers necessary to produce the same effect, etc.) of engaging in it, and thus paltering might appear to some people to be a problem that can, at least to a significant degree, take care of itself. It will not go away, but it may be inherently self-limiting, and thus not be in need of external sanctions in order to keep its incidence to manageable proportions.

Yet although at the extremes such self-limitation may be expected to occur, short of the extremes there is a familiar coordination problem, a problem exacerbated because it is often the case that paltering cuts across pre-existing social or professional groupings. Across such groupings, and to a lesser extent within the groupings, individual palterers will have an interest in there being less paltering so that their own palterers will be more effective. The best palterer of all is the first in some arena. But this dynamic affects everyone, at least on the assumption that there is a bit of the palterer in each of us. Without external enforcement of some sort, therefore, no one will have an incentive not to palter, but everyone will have an incentive to have others not palter, and it is just in this situation when outside help is most needed. If this outside help takes the form of punishment, the palterers who are punished will clearly be worse off and the world will be better off. And even if there is a cost to punishment, it will typically be the case that the loss to the palterer will be greater than the loss to the punisher, which is why the incentives encourage gossip and other forms of reputation-damaging sanctions against palterers.

Before turning to an analysis of the externalities that have been hinted at among players, we should identify the three reasons why a society – comprised of senders and receivers, with many players in both roles – should want fewer palterers in its midst. First, the senders do not like palterers because it makes it harder to have themselves believed. Second, receivers do not like palterers because it makes them more likely to be deceived. And third, if there are instruments that ameliorate the effects of paltering, either by enabling senders to demonstrate that their statements are whole truths, or receivers to discern the truthful essence of a statement, there will be costs to employ them. And such costs will be higher the more palterers there are.

In assessing various approaches to this outside help, which in this context would consist largely of calibrated social sanctions, we need to focus on two kinds of errors. And in doing so we can first put ourselves in the shoes of an individual confronting someone who makes a statement that may or may not be a palter. In this situation, to borrow the language of statistics and decision theory, the Type II error will be in believing someone who is paltering. And the Type I error will be in not believing someone who in fact is telling the whole truth. Obviously there are degrees of truth and untruth, and degrees of believing and non-believing, but we will stick to this simplified form of presentation.

Now imagine we have a world, call it World A, in which there are one hundred people delivering one message every day. And to simplify further, imagine that in World A there is another and distinct group of one hundred people receiving one message every day. Now assume that in World A there is a five percent chance that any given message is a palter, perhaps because all of the message deliverers palter five percent of the time, or because ninety-five percent never palter and five percent always palter. But whatever the cause, any individual message is five percent likely to be a palter. The task of the message recipient is to make a decision under these conditions about whether to believe a given message.

Assume that if the recipient accepts all messages at face value, he will make five percent Type II errors and no Type I errors. And if the recipient scrupulously sorts all of the messages in order to exclude those that could conceivably be palters, assume that he would screen out all of the palters but also ten percent of the truthful messages. So then he would make no Type II errors and 9.5% ($95\% \times 10\%$) Type I errors.¹⁷

Now let us move to another society, which we will call World B. In World B, assume that ten percent of the statements are palters, and only ninety percent are truthful. A recipient who accepts all statements in World B will thus make ten percent Type II errors. And if the recipient rejects all even slightly suspicious messages, he will make 9.0% ($90\% \times 10\%$) Type I errors. Thus, the recipient's possible tradeoff rate between Type II and Type I errors is greater in World B than it was in World A. Thus, he should be more skeptical, and disbelieve more

messages. Using this more formal analysis, therefore, we can see why it is that the honest used car salesman, like the midnight stroller in a neighborhood plagued with burglaries, is at a particular disadvantage. Ultimately, of course, the enhanced doubting by receivers in World B would feed back to affect the behavior of senders.

Recall P.T. Barnum's statement that "There's a sucker born every minute." Assuming that Barnum was right (we are confident that his estimate was actually on the low side), he was smart enough to recognize that a plentiful supply of suckers would make being a huckster more attractive. Barnum's implicit model was that there were two types of receivers, suckers and sophisticates, where sophisticates knew how to avoid being taken in by a palter. Barnum, who was particularly smart in these matters, was presumably interested in the ratio of suckers to hucksters, hoping for a higher ratio. So if the supply of suckers increased, the returns to hucksters would increase. But this in turn would stimulate the supply of hucksters, leading to a fall off in business, and eventually an equilibrium would be reached.

If switching from being an honest sender to being a huckster were costless, the ultimate returns to both suckers and hucksters would be the same as before the suckers were born, and nothing would have changed. Thus Barnum was obviously interested in a world in which there were many suckers and few hucksters, and his interest in their being fewer hucksters was not just about being able to divide up the pie in fewer sections. Rather, it was primarily in his interest in raising the returns of huckstering, which would come from there being fewer hucksters to make suckers, or sophisticates who could not costlessly detect hucksterism, wary. So as long as the ratio of suckers to hucksters remained high, Barnum would do very well. And what is interesting is that this is a world in which society would also do well. A world in which there are a small number of hucksters preying on a large basically trusting population that is rarely fleeced is, under reasonable empirical assumptions, probably a better world than one in which a large number of hucksters seeking to fleece an equivalent number of wised-up potential suckers. Thus a world with fewer hucksters, even if it is a world with fewer older but wiser suckers, would be a world with few actual fleecings. The task, then, is to devise strategies for getting to this better

world and avoiding the worse one. And the goal is certainly to avoid the worst world, a world in which there are a large number of hucksters serially fleecing the same suckers, suckers who rarely learned from their own bad experiences.¹⁸

Chief among the strategies for getting to this better world is likely to be one focusing on reputation. Although the law may have a role to play, and although direct condemnation – “You dirty palterer!” – can occasionally be valuable, we believe that reputational mechanisms that spread condemnation, and hence raise the expected costs of paltering, are more likely to be effective. To be caught in a serious palter may not hurt much, but to be caught in half a dozen may truly tarnish a reputation. The businessman of yesteryear, who could establish that “my word is my bond,” had a reputation that protected against both lying and paltering. The more a society rewards those who scrupulously avoid harmful palters, and are known as straight shooters, the more important it will be for people to avoid being known as even minor palterers.

We do not think it wise for reputational enforcement to rely (even if it could) on brittle barriers, in which reputations would remain good if there were no harmful palters, and would turn bad if but one were detected. Reputational enforcement needs to be more flexible than this. That is partly a function of the fact that all of us would probably like to palter a bit, and also partly of the fact that, given current standards, all of us probably palter from time to time. But more important, it is frequently difficult to determine whether a statement is a palter. There are Type I errors (rejecting a truthful statement as a palter), as well as Type II errors (responding to a palter as if it were the truth), and optimal detection will lead to some Type I errors. Thus, sending someone to the gallows for a single detected palter will be undesirable. But over the longer term, particularly if perceived palters to one individual are passed on to others, it should be easy to distinguish the palterer from the truth teller.

The need for measured drops in reputation is also partly a function of the fact that, as hinted above, the optimal level of paltering may not be zero. Just as there is a line between the palter and the lie, and between the palter and the truth, so too is there a line between the harmful palter and the beneficial white lie, or, to keep the parallelism, between the harmful palter and the

white palter. If there were no paltering, especially in a world of uncertainty about the line between the palters that are harmful and those that are beneficial, there might be too much blunt and harmful truth, and thus the optimal level of paltering, although likely low, is also unlikely to be zero.

The task then is to imagine a reputational mechanism that recognizes that people will be deterred from paltering not primarily by the inner voice that tells them not to, because there are also inner voices telling us that sometimes it is good for us to palter even if it is not good for society that we do so. Indeed, some of us get a sneaky but good feeling from fooling others, partly because a palter requires more ingenuity than a lie. Moreover, there is still another voice telling us that there are some socially beneficial palters and that this might just be one of them. Under these circumstances, the fear that we will be regarded poorly by others, or the fear that others will not deal with us, can prove to be a much more significant regulatory mechanism than self-policing.

How might such a regulatory mechanism work? First, it must separate and sanction most severely the intentional palters, leaving the negligent or innocent palters for different forms of punishment. So let us focus on the palters that we most want to limit, the self-interested intentional palters that are intended to hurt others, and that will in fact hurt others. Even for these, the reputational penalty cannot simply be one that imposes large penalties in a small number of instances. Although such strategies seem appealing in theory, low enforcement/high penalty regulatory strategies tend to work least well when the high penalties are socially unacceptable (see Schauer 2003 pg. 311-328). And high penalties are socially unacceptable in this context for a variety of reasons. One is that there is a bit of the palterer in each of us, and the line between self-interested palterers and harmless ones is not easy to draw. Is it wrong, for example, and, if so, how wrong, to enhance our own athletic accomplishments in order to impress a member of the opposite sex? People tend to be comfortable with imposing heavy penalties for crimes that they themselves cannot imagine committing, but when it is a socially harmful activity that they can see themselves doing – driving under the influence is the classic

example of a socially harmful activity that is under-penalized because of the “there but for the grace of God” phenomenon – the willingness to punish heavily is severely weakened. Moreover, it is considered bad form to mistakenly accuse someone of being less than honest. This is well-known to students who ask for extensions on final papers for reasons of computer failure, death of a grandparent, and the like. Although the students know that the story is false (these, by the way, are typically lies and not palters), and knows that the teacher knows that the story is likely false, the student also knows that the penalties to the accuser for making a false accusation of dishonesty are high. Few teachers want to accuse the student of dishonesty only to be met by a death certificate for the student’s grandmother, or the dated repair bill for the computer, or the verified record of a genuine medical emergency. So under conditions of uncertainty, even suspicious conditions, we often do not accuse when an accusation is merited, and for similar reasons we often are unwilling to impose severe reputational penalties even when they appear to be justified.

Thus, the imposition of reputational penalties is impeded by the way in which imposition of such punishment can cost the punisher as well as the punishee, by the fact that the optimal level of paltering is above zero, and by the fact that our universal desire to palter now and then makes us squeamish about heavy penalties. The task, therefore, is to devise a system of social or reputational sanctions for paltering that takes account of the unavailability of the theoretically efficient high penalty/low enforcement devices that might otherwise seem optimal. So if we are restricted to lower penalty enforcement mechanisms, we are restricted as well, unless we are unwilling to be effective, to low penalty/high enforcement devices.

Such devices are hardly rare. Indeed, the typical parking ticket is a classic example, even if we assume (counterfactually in some locations) that the law enforcement goal is to minimize the activity rather than to maximize revenue. In the context of paltering, therefore, the task is to imagine a set of reputational sanctions that is the equivalent of the parking ticket. That is, they should be easy to give out, but not so costly that they would not be given. Some of these may sound like gossip, but gossip, appropriately regulated itself by second-order social norms, can be

a useful first-order social norm in the control of anti-social behavior. Gossip can be a powerful social deterrent, and one of the useful things about gossip is that an individual item of gossip will rarely have a serious negative effect, but that multiple items will transform gossip into large-scale reputational penalties. Some jurisdictions raise the marginal cost of parking tickets as they accumulate in a year, just to discourage those who think of parking fines as basically a convenience that saves money on average, or just represents a beneficial tradeoff of ease-of-parking against money. Such an escalation scheme thus separates the true scofflaws from the person who has accidentally overstayed his meter contribution. So too, one item of gossip about a palterer may do little, two items do some but not too much, and three produce a widespread reputation for dishonesty. Such an approach is quite consistent with rational decision-making, given the occurrence of multiple events each of which would be low probability given good behavior (see Schauer & Zeckhauser 1996). Similarly, an extremely harmful palter would incur a large reputational penalty. Thus, if reputations spread fluidly, then appropriately directed and regulated gossip about palterers may succeed in imposing large deterrents on serial palterers, but only small deterrents and penalties on those who occasionally palter.¹⁹ If properly calibrated, this might well achieve the optimal result of aligning the societal interest in limiting large-scale paltering with the individual tendency in almost everyone to occasionally fall prey to the temptation to palter.

References

- Alexander, Larry and Sherwin, Emily. 2003. Deception in Morality and Law. *Law and Philosophy* 22: 393-450.
- Aquinas, Thomas. *Summa Theologiae* 41 (2a2ae 110): 3.
- Aristotle. *Nicomachean Ethics* Bk. IV (1127a): 28-30.
- Augustine. 1952. Against Lying. In *Treatises on Various Subjects*, edited by R.J. Deferrari. New York: Fathers of the Church.
- Bentham, Jeremy. 1931. *The Theory of Legislation*, edited by C.K.Ogden. New York: Harcourt, Brace and Co.
- Bok, Sissela. 1978. *Lying: Moral Choice in Public and Private Life*. New York: Pantheon Books.
- Kant, Immanuel. 1898. On a Supposed Right to Tell Lies from Benevolent Motives. In *Kant's Critique of Practical Reason and Other Works on the Theory of Ethics*, edited by Thomas K. Abbott. London: Longmans, Green, and Co.
- Kant, Immanuel. 1996. *The Metaphysics of Morals*, edited by Mary Gregor. Cambridge: Cambridge University Press.
- Keeton, W. Page, Dobbs, Dan B., Keeton, Robert E. and Owen, David G. 1984. *Prosser and Keeton on the Law of Torts*. 5 ed. St. Paul, MN: West Publishing Co.

- Resnick, Paul, and Zeckhauser, Richard. 2002. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In *The Economics of the Internet and E-Commerce*, edited by M.R. Baye, 127-157. Amsterdam: Elsevier Science.
- Schauer, Frederick. 2003. *Profiles, Probabilities, and Stereotypes*. Cambridge: Harvard University Press.
- Schauer, Frederick and Zeckhauser, Richard. 1996. On the Degree of Confidence for Adverse Decisions. *Journal of Legal Studies* 25: 27-52.
- Tversky, Aaron, and Kahneman, Daniel. 1973. Availability: A Heuristic for Judging Frequency and Probability. *Cognitive Psychology* 5: 207-232.
- Tversky, Aaron, and Kahneman, Daniel. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 211: 1124-1130.

Endnotes

¹Frank Stanton Professor of the First Amendment, John F. Kennedy School of Government, Harvard University.

²Frank Plumpton Ramsey Professor of Political Economy, John F. Kennedy School of Government, Harvard University.

³The effect is typically produced by the victim's belief in the sincerity of the liar's assertion, and so we might include this belief by the victim as a separate component of what constitutes a lie.

⁴Or consider the situation, one that most academics have experienced, in which a smiling and grateful student announces that he has been offered the job or place in graduate school for which he thinks you have recommended him, when in fact you have still not gotten around to writing the letter. When you respond by saying that you are happy for him, or even "you're welcome," you have paltered, even though you have not lied.

⁵Causation would be less important to a Kantian liar-focused account of the wrong of lying, but we believe it relatively non-controversial under most other accounts to maintain that a lie derives part of its wrongfulness from the actual way in which it causes a recipient either to come to false belief, or to be reinforced in a false belief.

⁶When houses are sold in Cambridge, Massachusetts, it is now the norm that the seller is asked to respond to a questionnaire. One question is whether the house has leaked recently. But consider a response of "A few years ago, during that very rainy summer, there was mild seepage into the basement, perhaps a cup of water overall." If the roof has been leaking since then, however, this would be a palter. The potential buyer will not be suspicious, but leaving the

question blank would raise suspicions, and stating explicitly that there were no roof leaks would be actionable.

⁷We acknowledge, but nevertheless bracket, the agency problem in the example, such that the parent is trying to get the child to make the decision that the parent would prefer, and not necessarily the decision that even a well-informed child would make.

⁸And especially so when the role of the person making the statement – car dealers, carpet salesmen, and so on – causes recipients of the messages to have their antennae up.

⁹Robert Solo was a well-regarded economist at Michigan State University who undoubtedly had many temptations to use palters of omission when people were impressed upon meeting him because of the similarity of his name to that of the considerably more famous Robert Solow of MIT.

¹⁰ See Thomas Lutz, this volume, who also addresses the social impact of misleading or false actions and words. For example, he observes that having his then girlfriend see through his crocodile tears was important to their relationship.

¹¹And it may be that the very teaching of this tale itself involves some paltering. Teachers and parents suspect the story may be apocryphal, but they resolve the uncertainty in the way that assists the lesson they wish to impart rather than in the way they believe most likely true.

¹²The tree might have had a disease, for example, a disease that young George could have accurately described in order to avoid the fact that, disease notwithstanding, he chopped down the tree just for fun.

¹³Clinton obviously miscalculated on this occasion, partly because he may have overestimated by orders of magnitude the number of people who would treat oral sex and sexual relations as different categories. Even so, Clinton still had the implausible but not completely impossible defense that he thought of the two as falling in significantly different categories, and therein lies the core of the idea of deniability.

¹⁴Among these harms are various externalities, especially the way in which both paltering and lying tend to hurt people who do not engage in the practice. The more untruths there are out there, the harder it is for the truth-teller to be believed, as any honest used-car dealer or carpet salesman will tell you. Honest people must then either suffer the consequences of not being believed, or spend additional resources – hire intermediaries, or offer special warranties, or pay for inspections – to be believed. Yet their counterparts in professions where liars and palterers are scarce face no such costs.

¹⁵Section 12(2) of the Securities Act of 1933, for example, prohibits not only making an “untrue statement of a material fact,” but also the “omi[ssion of] a material fact necessary in order make the statements, in light of the circumstances in which they were made, not misleading . . .”

¹⁶And so too, for closely related reasons, for burglary. An increase in burglary in a neighborhood will increase the watchfulness of neighbors, the prevalence of alarm systems, and the frequency of police patrols. Thus the more burglary there is the harder it is for the individual burglar, and the less burglary that exists the easier it will be for an individual burglar to burgle undetected.

¹⁷Of course, the tradeoff rate is likely to increase with the intensity of rejection. That is because

there is almost certainly heterogeneity among messages. Thus, some palterers are relatively easy to detect, implying that eliminating them would entail few Type I errors. As the costs in Type I errors to avoid Type II errors becomes greater. This implies that the possibility curve showing the possible values for the two types of errors bulges towards (is convex to) the origin.

¹⁸And that is why, popular conception aside, retail establishments are very interested in finding out the names and identifying information for people who have declared personal bankruptcy. Such people may be poor credit risks, but interests rates can take care of this. What is more important is that the list of those who have declared personal bankruptcy is a list of those especially likely to be vulnerable to sales pitches and especially likely to spend beyond their means.

¹⁹The success of eBay in promoting trust is due in large part to its reputation system, where buyers are asked to rate sellers after each transaction, and reputations are posted (see Resnick and Zeckhauser 2002).