# A Model of Flops*

Patrick Hummel
Yahoo!

John Morgan
UC Berkeley and Yahoo!

Phillip C. Stocken
Dartmouth College

July 10, 2011

**Abstract**

Despite surveying over 200,000 consumers, Coca-Cola's introduction of New Coke was a spectacular failure. We offer a model that can explain New Coke and other flops. A firm surveys a large number of consumers, some of whom are sincere and some of whom are strategic, and makes product decisions based on this information. When firms and consumers agree on the fraction of strategic consumers, survey information does not aggregate. When firms and consumers have differing beliefs, equilibrium is in linear strategies and information aggregates. Despite this, flops still arise. The chance of a flop increases in the noise-signal ratio, a measure of the idiosyncratic component of consumer tastes, as well as the calibration-alignment ratio, the ratio of how well calibrated the firm is with the actual fraction of strategic consumers and the degree to which firms and consumers agree on this fraction. When these ratios are large, the firm does better by eschewing surveys entirely. A firm, however, can solve the flops problem by limiting the scope for strategic consumers to influence survey outcomes.

Keywords: Flops, market surveys, heterogeneous priors, product design

JEL Codes: L15, C72, D83

*We are indebted to David Laibson, Don Lehmann, Preston McAfee, Michael Schwarz, Joel Sobel, Eric van den Steen, and Miguel Villas-Boas for helpful remarks, Dana Sisak for outstanding assistance, and seminar participants at Harvard University for comments. The second author gratefully acknowledges the financial assistance of the National Science Foundation. Please send correspondence to: John Morgan; e-mail: morgan@haas.berkeley.edu.

# 1 Introduction

In 1985, Coca-Cola CEO Robert Goizueta knew he had a problem. For the last 15 years, the market share of Coke, the company's flagship product, had been steadily declining while rival Pepsi had been gaining. Moreover, growth in the cola business had slowed dramatically with little hope of turning around in the near future. Goizueta knew he needed a game changer to reverse this momentum. Increasingly, Coke no longer fit with the tastes of a "new generation" of cola drinkers. Coca-Cola was systematic in its approach, surveying nearly 200,000 consumers to determine just the right formulation for its new and improved cola. Based on these results, the flagship cola was reformulated and marketed with great fanfare as New Coke that spring. The outcome, however, was a failure and led pundits to label the move "the marketing blunder of the century." Coke's consumer hot lines were inundated with calls from angry consumers, a class action lawsuit was filed to restore the original recipe, and consumers started hoarding and rationing cases of the original Coca-Cola. Realizing the magnitude of its error, Coca-Cola reintroduced its original formula to store shelves a mere 79 days later and subsequently withdrew New Coke from distribution in the U.S. (Coca-Cola Company, 2009). While New Coke is widely considered to be an enormous flop, similar stories abound.

The Ford Edsel, perhaps the quintessential flop, was strongly influenced by the tools of the emerging field of motivational research. Ford hired motivational researchers to interview over 800 people. The interviews were unusually broad, seeking to obtain a complete picture of consumer tastes—everything from cars to cocktails. The design of the Edsel synthesized this research to produce a car designed to propel Ford to the top of the medium-priced automobile segment that, at the time, was dominated by Pontiacs and Buicks. The result was a disaster—by the time Ford pulled the plug on the Edsel, it had lost over $250 million (about $2 billion in today's dollars). In designing a car to appeal to everyone, they produced a car that appealed to no one.[1]

Many flops result from ill-considered product extensions where a firm conducts research to identify a perceived need, matches it with a core competence, and then introduces a

[1] See Carlson (2007) for further details.

product. Some of the most famous flops matching this description include Bic disposable underwear, Cheetos lip balm, Colgate dinner entrees, and Life Savers soda (Maricano, 2009).

Despite remarkable gains in the science of market research, and the rapidly falling costs of conducting consumer surveys, firms still routinely offer products that fail spectacularly. One possible explanation is managerial hubris. Perhaps managers simply override the results of these studies to pursue pet projects that they are convinced will succeed in the market. Certainly, for some flops, this is the right explanation, but it is harder to justify for consumer product giants like Colgate-Palmolive and Frito-Lay. Another explanation is the failure to translate this research into the final product. The Pontiac Aztek, for example, suffered mainly because of design constraints imposed by General Motors to lower costs (Dederer, 2006). But in other instances, such as Crystal Pepsi, it is difficult to argue that design constraints played a major role.

In this paper, we argue that the presence of even a small fraction of strategic respondents to a survey can lead to flops. In some instances, strategic respondents destroy the information aggregation properties of large surveys. In others, surveys aggregate information, but differing views about the importance of these strategic types can lead a firm to come to erroneous conclusions about the best product design. In short, even when most consumers are sincere in representing their tastes and firms rationally account for the possibility of strategic respondents, conducting a large survey and acting on this information can lead to spectacular flops.

Some firms have eliminated market surveys from the design process. For instance, Apple famously claims to design based mainly on Steve Jobs' perception of consumer desires rather than relying on surveys (Morris, 2008). Perhaps more surprisingly, this approach has yielded a string of successes at Apple that have made it a dominant player in the technology sector. We show that, under a wide range of parameter values, the Apple strategy may be the correct one. Firms may be better off eschewing surveys entirely rather than conducting large surveys—even if the costs of conducting the survey are negligible.

While our model is explicitly about firms making product design decisions, it is also applicable in other settings. For instance, politicians often use surveys and polling to determine the appropriate policy or position to take on an issue. Our model offers a mechanism whereby such poll driven policies or positions can prove unpopular.

In our model, a firm seeks to design an ideal product—one that matches the tastes of a representative consumer. Consumer tastes share a systematic component but are idiosyncratic as well. The firm conducts an open-ended survey, which we model as a costless information transmission game, to determine tastes. While some (perhaps most) consumers simply report their tastes honestly, others are strategic and tailor their reports to try to obtain a product that matches their own tastes rather than those of the representative consumer. In a standard model in which firms and consumers share common prior beliefs over the distribution of strategic types, this is enough to destroy information aggregation in large surveys. Indeed, product choice is *unresponsive* to the survey in the unique monotone equilibrium.

When priors differ, however, information aggregation is still possible and, when the components of taste are normally distributed, equilibrium has a remarkably simple form: Strategic consumers and firms use linear strategies that are uniquely determined. Even so, differing priors can lead to perverse effects in how well the product matches consumer tastes. The match quality depends on two ratios: The first is the noise-signal ratio—the ratio of the variance of the idiosyncratic to the systematic taste component. The second is the calibration-alignment ratio—the error in the firm's prior about the fraction of strategic consumers relative to the disagreement in priors between firms and consumers. Survey performance declines in both ratios.

While it is apparent that being poorly calibrated harms the firm, the model has more subtle implications as well. In particular, a survey conducted by a well calibrated firm will still perform poorly if there is little disagreement between firms and consumers. This is true even in the limit as the firm's priors become perfectly calibrated. Idiosyncratic noise rarely plays a role in asymptotic statistics but, when survey respondents are strategic, this noise still adversely affects performance in a large survey. A key implication is that even when the firm is fairly well calibrated, there are many cases where a firm might be better served by eschewing surveys altogether. Thus, the model offers a single explanation for both the success of Apple and the failure of New Coke.

Properly constructed surveys can, however, still be tremendously useful. The key to effective survey design is to limit the scope for strategic consumers to influence survey outcomes. One way to achieve this is through a binary survey instrument—simply asking consumers for the direction in which their tastes differ relative to some status quo. A second way is to use the median of reported consumer tastes rather than the more usual sample mean. While both of these approaches have poor properties in small surveys—they sacrifice information and do not necessarily produce sincere survey responses—they are extremely effective in large surveys, producing both truthful reporting and optimal product choices. Thus, the model provides some justification for the use of Likert or other simple reporting scales rather than more involved and in-depth survey procedures.

The remainder of the paper is as follows. Section 2 sketches the model, reviews the literature, and describes our contribution. Section 3 establishes that, when all strategic players share the same prior beliefs, information aggregation from a survey is impossible and information loss is severe. Section 4 shows that, when we allow for the possibility of heterogeneous prior beliefs, information aggregation is restored. Reporting strategies are simple and intuitive, consisting of linear strategies, and in large surveys, there is a unique responsive linear equilibrium. Section 5 examines the performance of products designed using survey information. We offer conditions under which these products are likely to be flops—despite information aggregation. Indeed, for a wide range of parameter values, a firm's products would be more successful were it to forego market surveys altogether rather than conduct (and act upon) a large survey. This is the case even if it costs the firm nothing to conduct the survey. Section 6 shows how a firm can solve the flops problem by gathering less information from consumers. Section 7 concludes. The proofs to some propositions are contained in the appendix.

# 2 Model

A firm is introducing a new product. Cost considerations limit the firm to a single offering. The key decision is the type of product to introduce. Let $y \in R$ denote a product type. One can think of $y$ as summarizing all aspects of the product: its technical specifications, design features, color, taste, etc. A product will be successful if it appeals to consumer tastes, which are unknown to the firm. There is a continuum of potential consumers for the product. Each consumer $i$ has taste $t_i = \theta + \varepsilon_i$ where $\theta$ represents a systematic taste component and $\varepsilon_i$ an idiosyncratic taste component. The systematic component consists of a single draw from a normal distribution with mean zero and variance $\sigma_\theta^2$. Let $F$ denote the cumulative distribution of $\theta$ with associated density $f$. The idiosyncratic component is drawn independently for each consumer from a normal distribution with mean zero and variance $\sigma_\varepsilon^2$. Consumers only know their own taste $t_i$ and not the separate components $\theta$ and $\varepsilon_i$. The distributions of the taste components are, however, commonly known by firms and consumers. Notice that, for these taste distributions, $E[\theta|t_i] = at_i$, where

$$a \equiv \sigma_\theta^2 / \left(\sigma_\theta^2 + \sigma_\varepsilon^2\right)$$

To determine what product to offer, a firm can conduct market research by asking consumers about their tastes. This research is sufficiently inexpensive that the firm can survey an arbitrarily large number of consumers. Specifically, the firm randomly selects $n$ consumers from the population and asks them each to provide a costless report $m_i \in R$. This research can be thought of as a combination of market surveys, focus groups, taste tests, and so on. The information gleaned from these reports is then summarized by $M$, which denotes the sample mean of the reports offered.[2] The firm makes its product decision based on this information.

In addition to differing in their tastes, consumers also differ in how they respond to the firm. A strictly positive fraction of consumers are *sincere*. These consumers report their type honestly when asked, i.e., sincere consumer $i$ reports $m_i = t_i$. The remaining consumers are strategic. These consumers choose a report $m_i$ that maximizes their payoffs recognizing the connection between their response and the product the firm ultimately offers. Specifically, suppose that strategic consumers suffer quadratic losses in the difference between the product

[2] Having the firm base its product choice on a summary statistics of the survey rather than the underlying micro-data greatly facilitates the analysis. It is also consistent with practice. Firms typically employ external research agencies to conduct marketing research. Owing to privacy concerns, their analyses tend to be generic rather than tailored to the firm (e.g., Deshpande and Zaltman, 1982; 1984; Moorman, et al., 1992). The use of classical statistics to analyze market survey data is ubiquitous (e.g., Sawyer and Peter, 1983; Lehmann, et al., 1998).

offered by the firm and their tastes. That is, the payoffs for these consumers are

$$u = -(y - t_i)^2$$

Notice that the individual report, $m_i$, does not enter directly into payoffs; it only enters to the extent it influences $y$.

The firm, however, wants a product with broad appeal and suffers quadratic losses as well.[3] Formally, suppose that the firm's payoff is

$$v = -(y - \theta)^2$$

That is, the firm wants to match its product to the tastes of the average consumer, $\theta$.

Of course, the firm does not know $\theta$ nor is it known individually by any of the consumers. If all consumers were sincere, the firm could determine $\theta$ with arbitrary precision through its research. The firm, however, cannot commit to the product it will offer as a function of the information received, $M$. Instead, the firm will select the project optimally given its posterior beliefs about $\theta$ given $M$; that is, the firm will choose the product

$$y(M) = E[\theta|M]$$

Firms and consumers may differ in their beliefs about the number of consumers who are sincere. Specifically, when a firm surveys $n$ consumers, the firm believes exactly $k_n$ of these consumers are sincere. In contrast, strategic consumers believe that exactly $h_n$ consumers are sincere for the same size survey. Thus, firms and consumers potentially have heterogeneous prior beliefs about the number of sincere consumers in the survey. We also allow for the possibility that neither set of beliefs may be correct. Since we will mainly be concerned with large surveys, it is useful to define $\delta = \lim_{n\to\infty} k_n/n$ and $\lambda = \lim_{n\to\infty} h_n/n$. Suppose further that the true number of sincere consumers in a survey of size $n$ is $s_n$ and let $\gamma = \lim_{n\to\infty} s_n/n$. Thus, $\delta$ and $\lambda$ correspond to the limiting beliefs of firms and consumers, respectively, while $\gamma$ coincides with the true fraction of sincere consumers in the limit.[4]

[3] It is straightforward to identify conditions on demand where profit maximization is equivalent to minimizing a quadratic loss function. See Alonso and Matouschek (2008) for details.

[4] In the limit of a large survey, our model is strategically equivalent to one in which a consumer is

Unlike standard models of costless information transmission, there is no commonly known inherent conflict between firms and consumers. That is, an individual consumer's "bias" is not known to the firm. Moreover, since $\varepsilon$ has zero mean, on average, the tastes of the population coincide with those of the firm. Tension in the model arises from two sources. First, while each consumer knows precisely his or her ideal product, the firm cares about the *average* consumer, and thus, even though, in expectation, preferences are aligned, for each taste realization, there is a misalignment of preferences. The second tension stems from heterogeneous prior beliefs about the fraction of sincere consumers in the survey. Firms and consumers differ in their experiences with surveys and this modeling tension is meant to reflect these differences. As we shall see, this can produce situations where the survey fully aggregates available information yet produces flops—products whose characteristics are far from the tastes of the average consumer.

**Literature Review**

Our model of consumer research contributes to the vast cheap talk literature dating back to Crawford and Sobel (1982). In this literature, there are typically one or a few experts providing information to a decision maker.[5] Tension in these models stems from the expert's bias—the difference between the expert's preferences and those of the decision maker as to the appropriate action.[6] While this bias is mainly common knowledge, a few papers study the case where it is privately known (e.g., Morgan and Stocken, 2003; Li and Madarász, 2008). In our model none of the consumers is an expert in the sense of knowing the state realization $\theta$ precisely. Thus, the firm consults with an arbitrarily large number of experts. Moreover, the exact bias of the expert ($\varepsilon$ in our model) is unknown to all parties. Finally, on average the experts, which we term consumers in our model, are unbiased. Instead, tension in our model derives from the firm anchoring on its prior beliefs over $\theta$ rather than preference misalignment.

sincere with probability $\gamma$, the firm believes a consumer is sincere with probability $\delta$, the strategic consumers believe a consumer is sincere with probability $\lambda$ and each consumer's type is an independent and identically distributed draw from these distributions.

[5]See Farrell and Rabin (1996) for an excellent survey of this literature. Grossman and Helpman (2001) offer a comprehensive perspective on costly and costless information transmission in political settings.

[6]Gordon (2010) offers a general treatment for how sender bias translates into equilibrium actions.

Our focus on large surveys also places our work in the much older literature on information aggregation. Dating back to Condorcet (1763; see Young, 1988), this literature mostly focuses on the performance of election mechanisms in aggregating preference information of voters.[7] The main finding here is that elections perform extremely well in this regard (e.g., Feddersen and Pesendorfer, 1997). Our model departs from this framework in several ways. First, the firm chooses from among a continuum of possibilities rather than making a binary choice. Second, and more importantly, the firm cannot commit to its choice ahead of time, unlike a voting rule. Instead, it collects consumer taste information and then uses it optimally given its beliefs. In that respect, our model is closer to models of polling (e.g., Meirowitz, 2005; Morgan and Stocken, 2008). Tension in these models again arises from privately known ideological bias on the part of constituents rather than anchoring on prior beliefs. These models highlight that, when the message space is binary, information aggregation is possible but that truth-telling is not generally consistent with equilibrium. A main result in our model (Proposition 3) establishes that, with heterogeneous priors, information fully aggregates with an *unrestricted* message space; however, consistent with the intuition from these earlier models, we show how the firm can still benefit from message space restrictions (Proposition 8).

We make two key behavioral departures from standard models. First, we introduce behavioral types in the form of sincere consumers. When these types first appeared in the cheap talk literature (e.g., Sobel, 1985; Benabou and Laroque, 1992), their main role was to affect the dynamics of the game. In particular, using intuition akin to Kreps and Wilson (1982) and Milgrom and Roberts (1982), strategic types imitate behavioral types to gain an advantage in the dynamic game. Our model is static, so such a role is absent. Instead, behavioral types smooth the posterior beliefs of the firm following any message. This eliminates equilibria sustained by the freedom to specify out of equilibrium beliefs (e.g., Krishna and Morgan, 2001) while allowing equilibrium reporting strategies that are continuous in the consumer's type. Continuous reporting strategies also have the property

[7] The Condorcet Jury Theorem claims that if voters behave non-strategically, then the majority in a dichotomous election is more likely than a single individual to make a correct decision and this likelihood increases as the size of the group increases (see Young, 1988).

that equilibrium information aggregation is now possible. Absent such types, equilibrium reporting strategies are typically discontinuous and lead to posteriors that partition the state space.

More recent work has studied the role of behavioral types in static settings. One branch of this literature explores the implications of the possibility of naive receivers on signaling (e.g., Ottaviani and Squintani, 2006; Kartik, Ottaviani and Squintani, 2007). Kim and Pogach (2009) also study the implications of a possibly naive sender on information transmission. Other work allows for naivete on the part of both the sender and receiver (Chen, 2011). A fundamental difference between these settings and the present paper is that they study situations where there is only a single sender whereas our main concern is with information transmission when there are a large number of senders.

Our second departure is to allow for the possibility that the firm and consumers have heterogeneous priors over the fraction of behavioral types in the population. While it has long been recognized that players might hold differing prior beliefs and that this view is not inconsistent with the assumption that players are rational (e.g., Harsanyi, 1968), there has recently been renewed interest in the effects of heterogeneous priors on player behavior (see Brunnermeier and Parker, 2005; Van den Steen, 2010). Heterogeneous priors have been shown to rationalize a variety of apparently irrational behaviors, including overconfidence in choice of actions (Van den Steen, 2004), the "pop" typical of first-day trading following an initial public offering (Morris, 1996), the winner's curse arising in private value settings (Compte, 2002), and speculative bubbles in assets prices (Scheinkman and Xiong, 2003). More recently, Van den Steen (2010) shows how relaxing the common priors assumption creates benefits to firm integration despite complete contracting.

In our model, heterogeneous priors likewise lead to a form of overconfidence—firms and consumers simultaneously believe they are obtaining an ideal product (in expectation) even though their preferences differ. Relaxing the common priors assumption in a cheap talk context changes the informational properties of equilibrium starkly: Under common priors, information loss is an inevitable consequence of equilibrium whereas with heterogeneous priors, it is not.

Finally, we contribute to the literature on survey research. Within economics, the literature largely confines itself to the problem of a monopolist conducting a survey to ascertain an uncertain demand curve (see, e.g., Manning, 1979; Venezia, 1984). The focus is on the costs and benefits of the survey rather than the sincerity of survey respondents. Within marketing, this literature recognizes the possibility that consumers might not answer sincerely when asked about issues they find sensitive, such as questions dealing with sex or health (see Fowler, 1993; Groves, 1989; Lehmann, et al., 1998; Hartman, et al., 1991). Nevertheless, the literature mostly ignores this possibility, analyzing data under the assumption that all consumers are sincere (e.g., Churchill and Iacobucci, 2005; Lehmann, et al., 1998; Sawyer and Peter, 1983). In contrast, we examine how the interaction between strategic firms and respondents impacts the informational properties of surveys and the quality of subsequent decision making.

# 3 Homogeneous Priors

Compared with standard cheap talk models, our model makes two behavioral departures: First, we assume that a fraction of consumers are sincere and merely report their tastes honestly when asked. Second, we assume that firms and strategic consumers differ in their views about the fraction of sincere consumers in the survey population. In this section, we admit only the first of these departures and explore the implications of the sincere consumers on the quality of market research.

Straightforward intuition would suggest that, the larger the fraction of sincere consumers, the greater the informational benefit to the firm. Clearly, if all consumers were sincere, the firm could infer the systematic taste component, $\theta$, arbitrarily precisely—that is, information aggregates.[8] When none of the consumers are sincere, it cannot. Thus, one might surmise that the quality of information increases with the fraction of sincere consumers. Indeed, this would be the case if the behavior of strategic consumers remained fixed; however, their behavior adjusts to the fraction of sincere types in the population. We show that, even if the fraction of strategic consumers is small, their influence on information aggregation is large.

[8] Formally, we say that information aggregates if, in equilibrium, there exists an invertible function mapping $M$ onto $\theta$.

Specifically, when there are strategic consumers, there is no equilibrium in which information aggregates. Formally,

**Proposition 1** *Under homogeneous priors, there does not exist an equilibrium for a large survey in which information aggregates.*

Proposition 1 reveals a fundamental tension in the model. While, on average, strategic consumers share the same preferences as the firm, individually they do not. Thus, each strategic consumer sends a report so that, in expectation, she gets her ideal product, $t_i$. At the same time, the firm uses its research to design a product that converges in probability to the tastes of the mean consumer, $\theta$. But, as Proposition 1 shows, these two objectives are mutually inconsistent when information aggregates.

Of course, the fact that information does not fully aggregate does not imply that information loss is necessarily severe. Indeed, in the Crawford and Sobel (1982) framework, the misalignment of preferences between sender and receiver ensures that some information is lost, but the welfare effects of this loss are minimal as preferences become aligned. As we will show below, this same intuition does not hold when the fraction of sincere consumers is large. Our next result shows that, even in the limit as the fraction of sincere consumers goes to one, information loss remains severe.

To formalize this idea, we study equilibria that are well-behaved in the following sense: Fix an equilibrium and let $y(\theta)$ and $M(\theta)$ denote the equilibrium product choice and sample mean in state $\theta$ in the limit as the number of consumers surveyed goes to infinity. We say that an equilibrium is *monotone* if $y(\theta)$ and $M(\theta)$ are (weakly) monotone in $\theta$. Monotone equilibria capture the following idea: When the systematic component of tastes shifts upward, both the firm and the typical consumer prefer a higher product type to a lower one. Monotone equilibria capture this idea but are weaker—if product types were *decreasing* in $\theta$, this too would satisfy monotonicity but make little economic sense. Technically, the restriction to monotone strategies rules out perverse strategies on the part of the strategic types leading to firm strategies where $y(\theta)$ is infinitely oscillatory in $\theta$

With this restriction in mind, we can now offer a precise statement about the severity of information loss under homogeneous priors:

**Proposition 2** *In any monotone equilibrium, the firm's product choice is not responsive to the survey; that is, $y(M) = 0$ for all $M$.*

Proposition 2 is, essentially, an unravelling result. The impossibility of simultaneously satisfying the objectives of the firm and strategic consumers leads to an "arms race" whereby strategic consumers send ever more extreme messages which, in turn, leads the firm to ever more discount the survey results. This process equilibrates only when the firm is totally unresponsive to survey results. Moreover, the restriction to monotone equilibria is stronger than what is needed for the proof. It suffices that $M$ and $y$ are monotone for sufficiently large (small) values of $\theta$.

# 4 Heterogeneous Priors

The combination of sincere and strategic consumers in the common priors model proves to be a volatile mixture leading to severe information loss. In this section, we show that, by relaxing the assumption of common prior beliefs, information aggregation is restored. There exists an equilibrium in which information aggregates, and moreover, it takes on a simple and intuitive form: both reports and products are chosen using linear strategies. That is, strategic consumers use a reporting strategy $\beta t_i$, where $\beta$ is a parameter, while the firm selects the product using the strategy $\phi M$ where, again, $\phi$ is a parameter.

First, consider the case where the product is *responsive* to the reports of consumers, i.e., where $\phi \neq 0$ and finite. Recall that the firm chooses product $y$ such that

$$y(M) = E[\theta|M]$$

Given $\delta$ and $\beta$, the firm believes that $M$ converges in probability to $\delta\theta + (1-\delta)\beta\theta$. The firm should then choose $\phi$ such that

$$\phi(\delta\theta + (1-\delta)\beta\theta) = \theta$$

which yields

$$\phi = \frac{1}{\delta + (1-\delta)\beta} \tag{1}$$

Now consider strategic consumer $i$. She should choose $m_i$ such that

$$\begin{aligned} t_i &= E\left[y\left(M\right)|t_i\right] \\ &= E\left[\phi M|t_i\right] \end{aligned}$$

Given her beliefs about the populace of respondents, it follows that

$$t_i = \phi \frac{h_n E\left[t_j|t_i\right] + \left(n - h_n - 1\right)\beta E\left[t_j|t_i\right] + m_i}{n}$$

for $j \neq i$. Then

$$m_i = \frac{n t_i}{\phi} - h_n E\left[t_j|t_i\right] - \left(n - h_n - 1\right)\beta E\left[t_j|t_i\right]$$

Substituting in the strategic respondent's strategy $m_i = \beta t_i$ and rewriting yields

$$\frac{\beta t_i}{n} = \frac{t_i}{\phi} - \frac{h_n}{n} E\left[t_j|t_i\right] - \frac{\left(n - h_n - 1\right)}{n}\beta E\left[t_j|t_i\right]$$

In a large survey, we then have

$$0 = \frac{t_i}{\phi} - \lambda E\left[t_j|t_i\right] - \left(1 - \lambda\right)\beta E\left[t_j|t_i\right]$$

Recall that $E\left[t_j|t_i\right] = E\left[\theta|t_i\right] = a t_i$. Solving for $\beta$, we have

$$\beta = \left(\frac{1}{\phi} - \lambda a\right)\frac{1}{\left(1 - \lambda\right) a} \tag{2}$$

Equilibrium then consists of simultaneously solving equations (1) and (2), which yields:

**Proposition 3** *Suppose that priors are heterogeneous, i.e., $\delta \neq \lambda$. Then the following comprises the unique responsive linear equilibrium in a large survey.*

$$\begin{aligned} m_i &= \beta t_i \\ y &= \phi M \end{aligned}$$

*where $\beta = \frac{(\delta - \lambda a)}{a(1-\lambda)-(1-\delta)}$ and $\phi = \frac{(a(1-\lambda)-(1-\delta))}{a(\delta-\lambda)}$.*

*Furthermore, in this equilibrium, information aggregates. Specifically, $M$ converges in probability to the invertible function $\psi\left(\theta\right) = \left(\gamma + \left(1 - \gamma\right)\beta\right)\theta$.*

Information aggregates under heterogeneous priors because, when firms and consumers have differing beliefs about the fraction of sincere consumers in the population, they can simultaneously believe that their objectives are being met. In particular, given beliefs $\delta$, the firm is choosing its product to match with what it perceives to be the average type in the population, $\theta$. At the same time, given beliefs $\lambda$, strategic consumers respond to the survey confident that they will get their most preferred product in expectation. The combination of normally distributed tastes and quadratic payoffs leads to simple linear strategies.

Proposition 3 showed that there was a unique responsive linear equilibrium. Unfortunately, there is also a linear equilibrium that is unresponsive. This is the linear analog to a familiar "babbling equilibrium" one often encounters in models of costless information transmission. Specifically,

**Proposition 4** *The following comprises an unresponsive linear equilibrium:*

$$\begin{aligned} m_i &= \beta t_i \\ y &= 0 \end{aligned}$$

*where* $\beta = -k_n / (n - k_n)$.

While Proposition 4 shows that an unresponsive linear equilibrium exists, it is not robust to trembles—small perturbations in the equilibrium strategies. Formally, we say that a linear equilibrium is *tremble-proof* if small perturbations in the firm's strategy lead to small perturbations in strategic consumers' reporting strategies and vice-versa. With this definition in hand, we can now show:

**Proposition 5** *In large surveys, there is a unique tremble-proof linear equilibrium corresponding to the strategies contained in Proposition 3.*

Despite the apparent complexity of the model, equilibrium can be strikingly simple—both consumers and the firm play linear strategies and information aggregates in a large survey. Moreover, the equilibrium is unique in the class of linear strategies (using either the refinement of responsiveness or tremble-proofness). This is in sharp contrast to standard models of strategic information transmission in which information often does not aggregate, strategies are typically partitional, and equilibrium multiplicity is commonly a problem. One might then be tempted to conclude that, with heterogeneous priors, a firm's decision about whether to conduct a large survey or to forego surveying has a clear answer—in favor of the large survey. As we will show in the next section, the presence of heterogeneous priors makes this determination less obvious than it first appears.

# 5 Flops

In this section, we investigate the quality of product designs when a firm conducts a large survey. Earlier we saw that information aggregated in a linear equilibrium, and thus, one might have some measure of confidence that large surveys would produce "correct" product designs. Our main result of this section is to show that, instead of producing good designs, under some conditions acting on a large survey is likely to lead to a product that is a flop. Put differently, information aggregation does not inoculate a firm from producing a design that is ill-suited to the average consumer.

We first consider what would seem to be an obvious question: Should a firm undertake a large survey or no survey whatsoever in designing its product? Recall that when a firm designs based on a large survey, its expected payoffs are

$$V = -\left(\phi M - \theta\right)^2$$

and since $M$ converges in probability to $\left(\gamma + \left(1 - \gamma\right)\beta\right)\theta$, then, using the expressions for $\phi$ and $\beta$ given in Proposition 3, it follows that for large $n$, the firm's payoffs converge in probability to

$$\begin{aligned} V &= -\left(\frac{\left(\gamma - \lambda\right)a + \delta - \gamma}{\left(\delta - \lambda\right)a} - 1\right)^2 \theta^2 \\ &= -\left(\frac{\delta - \gamma}{\delta - \lambda} \times \frac{\sigma_\varepsilon^2}{\sigma_\theta^2}\right)^2 \theta^2 \end{aligned} \qquad (3)$$

Notice that the firm's expected payoffs depend on a pair of ratios. We term the first expression, $\left|\left(\delta - \gamma\right) / \left(\delta - \lambda\right)\right|$, the *calibration-alignment ratio*. The numerator reflects how well calibrated the firm's beliefs are about the fraction of strategic consumers in the population. The denominator reflects the alignment of priors between the firm and strategic consumers.

We term the second expression, $\sigma_\varepsilon^2/\sigma_\theta^2$, the *noise-signal ratio*. It is the ratio of the variance in the idiosyncratic component of consumer tastes compared to the variance of the systematic component. Finally, equation (3) reveals that, despite information aggregation, the product design does not converge to the tastes of the average consumer, $\theta$. Otherwise, this expression would be zero. Rather, it reveals that, unless there is no variance in the idiosyncratic component of consumers' tastes or the firm's priors are perfectly calibrated with the truth, the firm's product designs meet the needs of the market only imperfectly.

Now suppose that the firm eschews surveys entirely. Lacking any additional information on which to make its design decision, the firm opts for a design that matches its prior beliefs about $\theta$, i.e., $y = E[\theta] = 0$. In this case, for every realization of $\theta$, such a firm suffers losses of

$$V_0 = -\theta^2 \tag{4}$$

Comparing equations (4) to (3), it may be readily seen that performing a large survey is superior if and only if

$$\left|\frac{\delta - \gamma}{\delta - \lambda}\right| \times \frac{\sigma_\varepsilon^2}{\sigma_\theta^2} < 1 \tag{5}$$

Thus, we have established that:

**Proposition 6** *A firm benefits from a large survey if and only if the magnitudes of the calibration-alignment ratio and the noise-signal ratio are sufficiently small. Formally, a large survey is better than no survey if and only if equation* (5) *holds.*

Proposition 6 reveals that, even when information aggregates in a survey, firm management may be better off ignoring this information entirely and designing based on gut feel. This is more likely to be the case for products that have a large idiosyncratic taste component or in settings in which firms are relatively less sure about the fraction of strategic consumers. In technology industries where the product is something new and difficult for consumers to envision ahead of time, one might expect that the noise-signal ratio will be large. For instance, in designing his high-tech horseless carriage (car), Henry Ford famously quipped that "If I had asked my customers what they wanted, they would have said a faster horse."[9] Likewise, for the design of intermediate inputs in the supply chain, survey respon-

[9] As quoted by Kelley and Littman (2005).

dents will be other downstream businesses. In this case, both the firm and the "consumers" probably share similar beliefs about the fraction of strategic types in the population and hence the calibration-alignment ratio is likely to be high. In either case, a firm might be better off not surveying at all.

When do product designs turn out to be flops? To examine this question, we say that a product is a *flop* if it is more than $z > 0$ standard deviations away from the tastes of the average consumer, $\theta$. The probability that a design is a flop following a large survey is then given by

$$\Pr[flop] = \Pr[|\phi M - \theta| > z\sigma_\theta]$$

and, since $M$ converges in probability to $(\gamma + (1-\gamma)\beta)\theta$, it then follows (after some algebra) that

$$\Pr[flop] = 2\Phi\left(-\left|\frac{\delta-\lambda}{\gamma-\delta}\right|\frac{\sigma_\theta^2}{\sigma_\varepsilon^2}z\right) \tag{6}$$

where $\Phi(\cdot)$ is a cumulative distribution function of a standard normal random variable. Notice that the chance of a flop again depends on the calibration-alignment ratio, the noise-signal ratio and $z$, the tolerance for missing the ideal product, $\theta$.

When the firm eschews surveys entirely, the probability of a flop is simply

$$\Pr[flop] = 2\Phi(-z) \tag{7}$$

Now, the comparison between equations (6) and (7) is straightforward. Once again, a firm reduces the chance of a flop by conducting a large survey (relative to not surveying) if and only if $\left|\frac{\delta-\lambda}{\gamma-\delta}\right|\frac{\sigma_\theta^2}{\sigma_\varepsilon^2} > 1$, which is equivalent to the condition in (5) derived under payoff maximization. Thus, we have shown:

**Proposition 7** *Choosing whether to conduct a large survey so as to maximize the firm's expected payoff is equivalent to choosing whether to conduct a large survey to reduce the probability of a flop.*

In other words, payoff maximization and flop minimization lead to the same decision as to when to conduct a large survey. Proposition 7 implies that, when the firm wishes to maximize payoffs, it can delegate the product design decision to a product manager who might have very different concerns—say avoiding a failed product—and still achieve this objective. Figure 1 sketches the frontier delineating the decision whether to survey. This frontier is simply a rectangular hyperbola where $|(\delta - \gamma) / (\delta - \lambda)| \times \sigma_\varepsilon^2/\sigma_\theta^2 = 1$. The locus described by $|(\delta - \gamma) / (\delta - \lambda)| \times \sigma_\varepsilon^2/\sigma_\theta^2 = \kappa$, for arbitrary constant $\kappa > 0$, traces out an "iso-flop" line—it describes the locus of calibration-alignment and noise-signal ratios where the probability of a flop is identical. Level curves lying to the northeast of the one depicted in the figure are associated with greater chances of a flop while those to the southwest are associated with reduced chances of a flop.

We have emphasized the critical role played by the calibration-alignment and noise-signal ratios in whether a design turns out to be a flop or not. It is useful to define $\rho$ to be the product of these two ratios. Formally, $\rho = |(\delta - \gamma) / (\delta - \lambda)| \times \sigma_\varepsilon^2/\sigma_\theta^2$. If we let the flop threshold be $z = 1/2$ (i.e., a product is deemed a flop if it falls outside of a one standard deviation band around $\theta$), then we can compute the probability of a flop as a function of $\rho$. Figure 2 depicts this plot. The figure highlights that, unless $\rho$ is very low, the chance of flops is substantial. Indeed, even when $\rho$ is only 0.75, the majority of products are flops. At the point of indifference between surveying or not, the chance of a flop is about 62 percent. Surveying when $\rho$ is higher than this indifference point further exacerbates the flops problem. The main point of the figure is that, unless tastes are highly correlated (i.e., $\sigma_\varepsilon^2$ is small) and the firm is well calibrated (i.e., $\delta$ is close to $\gamma$), the probability of producing a flop is surprisingly high despite information aggregation.

Conventional wisdom suggests that large surveys are most important when tastes exhibit large idiosyncratic variation (e.g., Lehmann, et al., 1998). Under the standard view, the size of the survey helps to wash out this variation thus revealing the systematic component of tastes. Our results suggest the opposite conclusion. When idiosyncratic taste variation is large, surveys perform poorly and flops are likely. Indeed, in the limit, as $\sigma_\varepsilon^2 \to \infty$, a flop results almost surely. Intuitively, larger idiosyncratic taste variation creates greater tension between strategic respondents and the firm. Since the firm gives little weight to individual survey responses (relative to its prior beliefs), strategic respondents move increasingly aggressively to shift the firm's product design to their ideal. The result of this tension is that, in the limit, strategic respondents effectively jam the signals of sincere respondents and the firm becomes hypersensitive to small changes in $M$ and responds accordingly. Since the firm is not perfectly calibrated, this hypersensitivity magnifies the firm's errors in design, inevitably resulting in a flop.

Our results also highlight the special role played by assuming priors are homogeneous. Under this assumption, one comes to the pessimistic conclusion that surveys are never helpful. When priors are not homogeneous, there is scope for information revelation in large surveys. Moreover, the greater the disagreement in the priors of the firm and the strategic types, the more scope there is for surveys to be helpful. One might be tempted to equate disagreement in the sense of differences between $\lambda$ and $\delta$ with misalignment of preferences. In most models of strategic information transmission, such misalignment is harmful to information transmission (e.g., Crawford and Sobel, 1982). Here, disagreement is helpful in that it reduces the incentives for strategic types to distort and for firms to respond to these distortions. Indeed, our results show that, all else equal, when disagreement is large, flops are unlikely.

**Outsized Role of Strategic Consumers**

Clearly, the flops problem is caused by distortions in the reports of strategic consumers; that is, on their "outsized role" in determining survey results. One might have the intuition that strategic consumers wield this influence by choosing either very large or very small survey responses, (i.e., large equilibrium values of $\beta$). Were this the case, then one possible solution to the flops problem would be to constrain the space of reports through upper and lower bounds. The next example shows that even when strategic consumers only modestly exaggerate their tastes, aggregate survey results can be sufficiently distorted that flops are a significant problem; thus capping the reporting space is not a panacea.

**Example 1** *Suppose $\delta = 1$, $\lambda = 3/4$, $\gamma = 7/10$, $\sigma_\theta^2 = 2$, and $\sigma_\varepsilon^2 = 1$. Strategic types exaggerate modestly, sending a report $m = 3t$. The chance of a product that falls outside of a one standard deviation band around $\theta$ is 40.5 percent in a large survey, so flops are a significant possibility.*

The relatively poor calibration of the firm in Example 1 is obviously helpful in generating the result. A more subtle intuition might be that, when the firm is well-calibrated, extreme exaggeration is required to produce flops. Our next example shows that exaggeration increases only modestly and flops arises fairly frequently even when the firm is closely calibrated to the true fraction of strategic consumers.

**Example 2** *Suppose $\delta = 1$, $\lambda = \gamma = 0.95$, $\sigma_\theta^2 = 5$, and $\sigma_\varepsilon^2 = 2$. Strategic types send a report $m = 9t$. The chance of a product that falls outside of a one standard deviation band around $\theta$ is 21.1 percent in a large survey.*

Again, it is the miscalibration rather than the exaggeration by strategic consumers that is at the heart of the flops problem. One might think that the problem essentially disappears when the firm is arbitrarily well-calibrated. Our next result shows that this is not the case. Recall that the chance of a flop depends on the calibration-alignment *ratio* rather than solely on the calibration level itself. With this in mind, the next example shows that, even when nearly all consumers are sincere, the chance of flops remains stubbornly high.

**Example 3** *Suppose $\delta = 1$ and strategic consumers are well-calibrated, i.e., $\lambda = \gamma < 1$. In the limit as consumers become sincere, the probability of a flop becomes*

$$\lim_{\gamma \to 1} 2\Phi \left( - \left| \frac{1-\gamma}{\gamma - 1} \right| \frac{\sigma_\theta^2}{\sigma_\varepsilon^2} z \right) = 2\Phi \left( -\sigma_\theta^2 / \sigma_\varepsilon^2 z \right) > 0$$

*and is independent of $\gamma$ along the limiting path.*

The key insight from this example is that the probability of a flop depends on the noise-signal ratio regardless of how well the firm is calibrated. Thus, the presence of even a small fraction of strategic consumers fundamentally changes the quality of firm decisions compared to the case where such consumers are absent.

# 6 Avoiding Flops

The influence of strategic consumers stems, in part, from the fact that the firm is observing only a summary statistic of the survey data, the sample mean, rather than the data itself. For normally distributed tastes, the sample mean is a sufficient statistic when all consumers are sincere, but not when there are strategic consumers. Typically, one needs to turn to scoring rules or other exotic mechanisms to induce truthful revelation and thereby avoid flops. Here we show that no such complexity is required.

Instead, a firm need only adjust its choice of summary statistics in interpreting the results of survey data. Absent strategic respondents, there are many other summary statistics that the firm might use to infer tastes and, indeed, in a large survey the firm ought to be indifferent among these alternatives. In this section, we show that, in the presence of strategic respondents, the seemingly innocuous choice of summary statistics has profound effects on incentives for truthful reporting and hence on the informational properties of surveys. We consider two such statistics and show how, merely through the judicious use of summary statistics, the firm can solve the flops problem.

**Binary Surveys**

Suppose that the firm continues to solicit open-ended responses to its survey and that consumers still offer costless messages in $R$; however, now the firm bases its designs only on the statistic, $\mu$, the fraction of responses where $m \geq 0$. In the limit, this is also a sufficient statistic for the firm to deduce the state in a large sample model of sincere consumers with normally distributed tastes. Here, we will show that the incentive properties of this statistic are vastly superior to using the more familiar sample mean as the basis for decision making.

**Proposition 8** *When the firm uses a $\mu$ statistic as the basis for its decisions, then truthful revelation is an equilibrium in a large survey. The $\mu$ statistic solves the flops problem.*

The proof is by construction. First, suppose that the firm acts as though all surveyed consumers were truthful. Under this hypothesis, in a large survey, $\mu$ will converge in probability to

$$\mu \to 1 - \Phi\left(-\frac{\theta}{\sigma_\varepsilon}\right)$$

Hence, in equilibrium, the firm's product will converge to

$$\begin{aligned} y(\mu) &= -\sigma_\varepsilon \Phi^{-1}(1-\mu) \\ &= \theta \end{aligned}$$

in a large survey.

It remains to show that strategic consumers can do no better than respond truthfully. We provide an intuitive argument below and relegate the formal argument, which is more involved, to the appendix. A strategic individual faces a choice between reporting $m \geq 0$ and $m < 0$. The particular value of $m$ in each region is irrelevant. Given taste $t_i$, a strategic consumer expects $\theta$ to take on the value $at_i$; thus, in a large survey, when $t_i > 0$, a strategic consumer expects that the design will converge to $at_i < t_i$. Notice that the equilibrium design is "lower" than her ideal design, so she strictly prefers to report $m = t_i$ rather than any value of $m < 0$ since reporting $m < 0$ will only lead to a design that is further from her ideal in expectation. Similarly, when $t_i < 0$, the equilibrium design is "higher" than her ideal design, so she strictly prefers to report $m = t_i$ rather than any value of $m \geq 0$ since reporting $m \geq 0$ will only lead to a design that is further from her ideal in expectation.

Therefore, truth-telling is incentive compatible for strategic consumers and the firm's expectations about consumer reports are confirmed. Thus, we have identified a fully revealing equilibrium, regardless of calibration, alignment, or the noise-signal ratio.

Proposition 8 points out the crucial connection between the data processing strategy of the firm and incentives for the strategic consumers. The key to solving the flops problem is for the firm to focus on the fraction of consumers sending positive messages. A simple way a firm might gather this data is by using blunter survey instruments, such as Likert scales. In terms of the model, were the firm to survey consumers using only a simple binary survey instrument, then the equilibrium we constructed in Proposition 8 would remain valid and the firm could end up making better decisions than if it used a more comprehensive or open-ended survey instrument. In short, when some consumers are strategic, the firm can learn more by asking less in a large survey.

**Sample Median**

The binary survey instrument works by not allowing strategic individuals to exaggerate their type, but at the cost of coarsening the information gleaned from surveys when there are only a small number of respondents. A different way one might solve the "exaggeration problem" is by trimming the summary statistic so as to be less sensitive to outliers. Utilizing the median, rather than the sample mean, represents the most stringent trimming strategy available to the firm. To analyze this problem, it is helpful to discretize the action spaces of the players and allowing the "grid" over the set of possible actions to become arbitrarily fine. Specifically, suppose that the action spaces of firms and consumers are constrained to a grid ranging from $-l$ to $l$. Let $g > 0$ denote the fineness of the grid. There are $2w + 1$ actions available where $w = l/g$. Thus, the set of actions is given by $m \in \{-l, -l+g, -l+2g, .., -3g, -2g, -g, 0, g, 2g, 3g, ...l-2g, l-g, l\}$ and similarly for $y$. Note that as $g$ becomes arbitrarily small and $l$ becomes arbitrarily large, the action space becomes arbitrarily close to the action space in the continuum model.

Consider a survey consisting of an odd number of individuals (so the median is unambiguously defined) and where the firm bases its product choice on the median report.[10] We will show that truth-telling comprises an equilibrium in a large survey utilizing this technology. Since the message space and the type space no longer coincide, we need to define truth-telling. We will say that a *truthful* response is one where the message that is closest to the consumer's taste is selected. As above, we will assume that sincere consumers are truthful and consider the incentives of strategic consumers. A truth-telling equilibrium is an equilibrium where all consumers engage in truthful strategies and the firm is choosing a product optimally under the assumption of truthful responses.

**Proposition 9** *For any grid with fineness $g$ and length $2l$,when the firm selects its product on the basis of the median report, truth-telling is an equilibrium for a sufficiently large survey.*

Why does trimming using the median lead to truth-telling? The report of a strategic type is relevant only to the extent that it is the median report and thus, in determining whether to distort a report, a strategic consumer conditions on that event. In a small survey, the firm will still anchor on its prior beliefs and therefore there is still some benefit from limited exaggeration; however, as the survey size increases, less weight is placed on the firm's prior belief and therefore there is little incentive to distort conditional on being the median report. The discrete action space ensures that there is a sufficiently large survey whereby this slight misalignment of preferences between the firm and the strategic consumer ceases to matter and reporting becomes truthful.

[10]The results also hold if there is an even number of respondents.

Now one might conjecture that the same effect will also hold when the sample mean is the summary statistic. The key difference is that, when the median statistic is used, strategic consumers only affect the product choice conditional on being the median report, whereas when the sample mean is used, their reports always have a small effect on the product choice. Thus, under the median, strategic types have a small chance of making a large change in the product based on their report, while with the sample mean, they have a large (probability one) chance of making a small change in the product. In the former instance, truth-telling is optimal, while in the latter instance, it is not since the product choice does not necessarily coincide with the tastes of the strategic consumer.

Since Proposition 9 holds for any grid with fineness $g$ and length $2l$, it then follows that, letting the length become arbitrarily long (i.e., $l \to \infty$) and the gaps arbitrarily small (i.e., $g \to 0$), truth-telling remains an equilibrium for large surveys along this sequence. Moreover, as $g \to 0$ and $l \to \infty$, the action space converges to the countably infinite analog of the continuum action model. While this is not a proof that truth-telling or near truth-telling is an equilibrium in the continuum model, it is suggestive.

Finally, while Proposition 9 was derived under the assumption that the systematic and idiosyncratic components of consumer tastes are normally distributed with known mean and variance, the result also holds under more relaxed assumptions. In particular, so long as the idiosyncratic taste component is symmetrically distributed around zero, then the systematic taste component can be drawn from any distribution and truth-telling will still hold. Thus, the firm need only have minimal knowledge of the data generating process giving rise to consumer tastes for flops to be avoided using this statistic. Moreover, note that truth-telling is an equilibrium for a sufficiently large sample irrespective of the firm's beliefs about the fraction of sincere and strategic consumers, and in particular, irrespective of whether the firm is well calibrated.

# 7 Conclusion

The science of market research has progressed rapidly since the time of the Ford Edsel. Now, at relatively modest cost, firms can (and do) conduct large surveys of consumers' tastes. This data is critical to final product offerings. Yet, despite these gains, products that prove to be flops are commonplace—as any visit to a local movie theater will attest.

The problem, we argue, is that some consumers are sophisticated in responding to surveys—they answer survey questions strategically rather than honestly. As a consequence, conducting large surveys by no means guarantees the avoidance of flops. The effects of these strategic consumers on survey data depend crucially on the beliefs held by consumers and firms about their prevalence. When beliefs are common, the presence of even a small fraction of strategic respondents makes information aggregation impossible. In contrast, when beliefs differ, large surveys can still aggregate information. Moreover, equilibrium reporting strategies are simple and intuitive—there is a unique, responsive equilibrium where reporting strategies are linear in tastes. The key intuition is that heterogeneous prior beliefs allow strategic consumers and firms to simultaneously believe that their objectives are being met. That is, all parties believe that the survey will produce their ideal product even though views as to what this product is differ.

Even when information aggregates, flops remain a distinct possibility. We show that the chance of a flop depends on two ratios: the noise-signal ratio, which measures the relative amount of idiosyncratic taste variation among consumers, and the calibration-alignment ratio, which measures the accuracy of the firm's views about the fraction of strategic consumers relative to the degree to which priors are aligned. Absent strategic consumers, the noise-signal ratio would have no bearing on the quality of a large survey in that idiosyncratic taste variation "washes out" of the analysis. With sophisticated consumers, a higher noise-signal ratio leads to larger errors in the firm's assessment of the tastes of the average consumer and, as a consequence, more flops—even in the limit. The calibration-alignment ratio reveals that firms are better served when their beliefs as to the fraction of strategic consumers are well-calibrated with the true fraction and worse off when consumers share the firm's beliefs. Because the chance of a flop depends on the calibration-alignment *ratio,* flops still arise even as a firm becomes arbitrarily well-calibrated. When these ratios are large and the firm uses the mean of the survey responses to make a decision, a firm is better off eschewing surveys entirely rather than relying on large surveys. This is true even if surveys are costless to conduct. In our view, this helps to explain the remarkable success of Apple, whose product designs are mainly informed by Steve Jobs' perceptions of consumer tastes rather than survey evidence (Morris, 2008).

Our results highlight the hazards of using open-ended surveys and standard summary measures such as the sample mean to determine product choices. This is not to say that surveys cannot play an important role in product design provided that they are well designed. We offer two simple variations in survey design that solve the flops problem. Offering a binary survey where individuals can only respond as to the direction of their tastes relative to the status quo completely eliminates the incentives for misrepresentation. Alternatively, extreme trimming of outliers through the use of the median report can also align incentives in a large survey. Both designs capture enough information so that the firm can make optimal product choices. By asking for or using less information in its survey, the firm ends up obtaining more and better information. Thus, our model suggests an incentive rationale for the widespread use of Likert scales and trimming in survey design.

More broadly, while many firms embrace surveys as a central input in their decision making (e.g., Deshpande and Zaltman, 1982; 1984), this paper highlights the importance of careful survey design and analysis of the data. Typically, survey designers worry about choosing questions to extract as much information as possible without worrying too much about incentives to misrepresent. Likewise, statistical techniques for analyzing survey data rarely correct for strategic misrepresentation. But consumers are not mere passive data points nor draws from some unbiased urn. Rather, they have their own preferences, their own agendas, and will pursue these if given the opportunity. Seemingly well-meaning but naive procedures can easily give rise to flops while procedures that seem less effective at extracting information can, in fact, be powerful tools for making correct decisions.

We have couched our results in terms of surveys conducted by firms to guide product offerings, but our model is also relevant to political decision making. Polling data increasingly guides the formulation of policy. If anything, the problem of strategic survey respondents is even more pronounced when one is dealing with emotionally charged issues such as abortion, immigration, and the environment. Our results offer a cautionary tale for policy-makers seeking to govern or design policies merely to reflect the "will of the people" as measured by survey data. Product flops are costly to firms, but policy flops can have far graver consequences—not just to individual politicians but to society at large.

# A Appendix

**Proof of Proposition 1:**

Suppose to the contrary that information aggregates. Assume without loss of generality that consumer 1 is a strategic consumer. Her expected payoff from reporting a message $m_1$ is given by

$$E_\theta \left[ - \left( y \left( M \right) - t_1 \right)^2 | t_1 \right]$$

Therefore, for finite $n$, the consumer sends a message $m_1$ satisfying

$$\max_{m_1} - \int_{-\infty}^{\infty} \left( y \left( \frac{m_1 + \sum_{i=2}^{n} m_i}{n} \right) - t_1 \right)^2 dH \left( \sum_{i=1}^{n} m_i | t_1 \right)$$

where $H \left( \sum_{i=1}^{n} m_i | t_1 \right)$ is the cumulative distribution function of the sum of the consumers' messages conditional on $t_1$. Differentiating with respect to $m_1$ and using the first order condition yields

$$\int_{-\infty}^{\infty} \left( y \left( \frac{m_1 + \sum_{i=2}^{n} m_i}{n} \right) - t_1 \right) y' \left( \frac{m_1 + \sum_{i=2}^{n} m_i}{n} \right) dH \left( \sum_{i=1}^{n} m_i | t_1 \right) = 0$$

for the equilibrium $m_1$. Since information aggregates, in the limit as $n \to \infty$, $\sum_{i=1}^{n} m_i / n$ converges to an invertible function $M(\theta)$, and furthermore, $y \left( M \left( \theta \right) \right) = \theta$. Thus, as $n \to \infty$, this first order condition can be expressed as

$$\int_{-\infty}^{\infty} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta = 0$$

where $f \left( \theta | t_1 \right)$ is the probability density function of $\theta$ conditional on $t_1$. For this condition to hold, it must be the case that $\int_{x}^{\infty} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta$ exists and is finite for all $x$. Equivalently, one can write the first order condition as

$$\int_{-\infty}^{t_1} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta + \int_{t_1}^{\infty} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta = 0$$

Suppose that $y' \left( M \left( \theta \right) \right) > 0$. Then, taking limits as $t_1 \to \infty$, yields

$$\begin{aligned}
& \lim_{t_1 \to \infty} \int_{-\infty}^{t_1} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta + \lim_{t_1 \to \infty} \int_{t_1}^{\infty} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta \\
= & \lim_{t_1 \to \infty} \int_{-\infty}^{t_1} \left( \theta - t_1 \right) y' \left( M \left( \theta \right) \right) f \left( \theta | t_1 \right) d\theta \\
< & \ 0
\end{aligned}$$

The equality follows from the following observation: since $\int_x^\infty (\theta - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta$ exists and is finite for all $x$, it follows that $\int_{t_1}^\infty (\theta - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta$ exists and is finite for all $t_1$. The integrand in this expression decreases at an exponential rate in $t_1$ as $t_1 \to \infty$, implying the integral decreases at an exponential rate in $t_1$. Therefore, $\lim_{t_1\to\infty} \int_{t_1}^\infty (\theta - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta = 0$.

The case where $y'(M(\theta)) < 0$ is analogous. We thus conclude that for a sufficiently large $t_1$, a strategic consumer's first order condition cannot be satisfied. Accordingly, with homogeneous priors, information aggregation is impossible. ■

**Proof of Proposition 2:**

Suppose to the contrary that $y(M) \neq 0$. Consider the strategic consumer's first order condition, which can be written as

$$\int_{-\infty}^{\infty} (E[\theta|M(\theta)] - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta = 0$$

Fix $b$ such that $b \in (a, 1)$. Taking limits as $t_1 \to \infty$ yields

$$\begin{aligned} & \lim_{t_1\to\infty} \int_{-\infty}^{bt_1} (E[\theta|M(\theta)] - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta \\ & + \lim_{t_1\to\infty} \int_{bt_1}^{\infty} (E[\theta|M(\theta)] - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta \\ = \; & \lim_{t_1\to\infty} \int_{-\infty}^{bt_1} (E[\theta|M(\theta)] - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta \end{aligned}$$

where the equality follows because the integrand in this expression decreases at an exponential rate in $t_1$ as $t_1 \to \infty$, implying the integral in the second term decreases at an exponential rate in $t_1$.

Note $E[\theta|M(\theta)] \leq E[\theta|\theta \geq bt_1]$ for all $\theta \leq bt_1$. Also note $E[\theta|\theta \geq bt_1] < t_1$ for sufficiently large $t_1$. Therefore, for sufficiently large $t_1$, $E[\theta|M(\theta)] < t_1$ for all $\theta \leq bt_1$. Further, observe that because $y(\theta)$ is monotonic in $\theta$ and $M(\theta)$ is monotonic in $\theta$, $y'(M(\theta))$ has only a single sign. Accordingly, when $y'(M(\theta)) \geq 0$, with a strict equality occurring over some positive measure of $\theta$, it follows that

$$\lim_{t_1\to\infty} \int_{-\infty}^{bt_1} (E[\theta|M(\theta)] - t_1)\, y'(M(\theta))\, f(\theta|t_1)\, d\theta < 0$$

Because the first order condition is not satisfied, we conclude $y(M) = 0$. The case where $y'(M(\theta)) \leq 0$ is analogous. ■

**Proof of Proposition 4:**

Since the firm's product does not vary with $M$, the strategic respondent's strategy is consistent with equilibrium play. It remains to show that there is no profitable deviation for the firm. Suppose, without loss of generality, that the firm believes that consumers $i = 1, 2, ..., k$ are sincere while the remainder are strategic; in this proof, the subscript $n$ on the variable $k$ has been dropped to streamline the exposition. Given the firm's beliefs,

$$M = \frac{t_1 + ... + t_k}{n} + \frac{\beta(t_{k+1} + ... + \ t_n)}{n}$$

which can be written as

$$M = \frac{k\theta}{n} + \frac{\varepsilon_1 + ... + \varepsilon_k}{n} + \frac{\beta(n-k)\theta}{n} + \frac{\beta(\varepsilon_{k+1} + ... + \ \varepsilon_n)}{n}$$

Using the fact that $\beta = -k/(n-k)$ yields

$$M = \frac{\varepsilon_1 + ... + \varepsilon_k}{n} + \frac{\beta(\varepsilon_{k+1} + ... + \ \varepsilon_n)}{n}$$

Taking expectations yields

$$E[\theta|M] = E[\theta] = 0$$

Thus, $y = 0$ is optimal. ■

**Proof of Proposition 5:**

It may be readily verified that tremble-proofness of the strategies in Proposition 3 follows from the fact that the best response functions given in equations (1) and (2) are continuous in the other players' strategies. To see that the strategies in Proposition 4 are not tremble-proof, consider a perturbation $\phi = \eta \neq 0$. The incentive constraint for the strategic respondent $i$ is to choose $m$ such that

$$t_i = E[\phi M | t_i]$$

Given the strategic respondent $i$'s beliefs about the populace of respondents, it follows that

$$t_i = \phi\left(\frac{h_n E[t_j|t_i] + (n - h_n - 1) E[\beta t_j|t_i] + m_i}{n}\right)$$

Then

$$m_i = \frac{nt_i}{\phi} - h_n E[t_j|t_i] - (n - h_n - 1) E[\beta t_j|t_i]$$

Substituting in the strategic respondent's strategy $m_i = \beta t_i$ yields

$$\beta t_i = \frac{nt_i}{\phi} - h_n E[t_j|t_i] - (n - h_n - 1) \beta E[t_j|t_i]$$

which can be rewritten as

$$\beta t_i = \frac{nt_i}{\phi} - h_n a t_i - (n - h_n - 1) \beta a t_i$$

Solving for $\beta$ yields

$$\beta = \left(\frac{n}{\phi} - h_n a\right)\left(\frac{1}{1 + (n - h_n - 1) a}\right)$$

Substituting in for $\phi = \eta$ gives

$$\beta = \left(\frac{n}{\eta} - h_n a\right)\left(\frac{1}{1 + (n - h_n - 1) a}\right).$$

Observe that as $\eta \to 0$, $\beta \to$sign$(\eta)\,\infty$, which does not equal

$$\beta = -\frac{k}{n-k}$$

Hence, we conclude that the equilibrium is not tremble-proof. ■

**Proof of Proposition 8:**

Define $f(\theta|t)$ to be the probability density function corresponding to the posterior beliefs of a consumer with type $t$; these beliefs are normally distributed with mean $at$ and variance $a\sigma_\varepsilon^2$. Under the putative truth-telling equilibrium, the consumer believes that the product design converges to $\theta$. Suppose that $t < 0$. If the consumer changes her message from $m < 0$ to $m \geq 0$, this will change the mean product design by some positive amount. The difference in consumer's payoff from reporting truthfully compared to a deviation where $m > 0$ is

$$u(t, m < 0) - u(t, m > 0) = \int_{-\infty}^{\infty} -(\theta - t)^2 f(\theta|t)\, d\theta - \int_{-\infty}^{\infty} -(\theta + q(\theta, n) - t)^2 f(\theta|t)\, d\theta$$

where $q(\theta, n)$ denotes the change in the product as a consequence of a consumer changing her message. Below, we will show that this expression is positive; i.e. the consumer does not benefit from a deviation.

We first show that for large $n$,

$$q(\theta, n) = \frac{\sigma_\varepsilon}{n\Phi'\left(\frac{\theta}{\sigma_\varepsilon}\right)}$$

where $\Phi$ is the cumulative distribution function of a standard normal random variable and $\Phi'$ its associated density. To see this, note that when a consumer changes her message and reports $m \geq 0$, the fraction of positive messages changes from $\mu$ to $\mu + 1/n$, where $\mu$ is the fraction of positive messages reported when all consumers are truthful. When all consumers are truthful in equilibrium, the firm chooses a project that equals $y = \sigma_\varepsilon \Phi^{-1}(\mu)$. In contrast, when a consumer changes her message and reports $m \geq 0$, the firm chooses a project equal to $y' = \sigma_\varepsilon \Phi^{-1}(\mu + 1/n)$. Thus, when a consumer changes her message, it changes the product by an amount equal to $\sigma_\varepsilon \left(\Phi^{-1}(\mu + 1/n) - \Phi^{-1}(\mu)\right)$. For large $n$, the change in the firm's product converges to

$$\begin{aligned} q(\theta, n) &= \frac{\sigma_\varepsilon}{n} \frac{d}{d\mu}\left(\Phi^{-1}(\mu)\right) \\ &= \frac{\sigma_\varepsilon}{n\Phi'\left(\Phi^{-1}(\mu)\right)} \\ &= \frac{\sigma_\varepsilon}{n\Phi'\left(\frac{\theta}{\sigma_\varepsilon}\right)} \end{aligned}$$

It now follows that the change in the consumer's utility is given by

$$\begin{aligned} & \int_{-\infty}^{\infty} \left(-(\theta - t)^2 + (\theta + q(\theta, n) - t)^2\right) f(\theta|t)\, d\theta \\ = & \int_{-\infty}^{\infty} \left(q(\theta, n)^2 + 2q(\theta, n) \times (\theta - t)\right) f(\theta|t)\, d\theta \\ = & \int_{-\infty}^{\infty} \left(\left(\frac{\sigma_\varepsilon}{n\Phi'\left(\frac{\theta}{\sigma_\varepsilon}\right)}\right)^2 + 2\frac{\sigma_\varepsilon}{n\Phi'\left(\frac{\theta}{\sigma_\varepsilon}\right)}(\theta - t)\right) \frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}} \exp\left(-\frac{(\theta - at)^2}{2a\sigma_\varepsilon^2}\right) d\theta \\ = & \int_{-\infty}^{\infty} \left(\frac{\sigma_\varepsilon}{n\left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\theta}{\sigma_\varepsilon}\right)^2\right)\right)}\right)^2 \frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}} \exp\left(-\frac{(\theta - at)^2}{2a\sigma_\varepsilon^2}\right) d\theta \qquad (8) \\ & + \int_{-\infty}^{\infty} \frac{2\sigma_\varepsilon(\theta - t)}{n\left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\theta}{\sigma_\varepsilon}\right)^2\right)\right)} \frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}} \exp\left(-\frac{(\theta - at)^2}{2a\sigma_\varepsilon^2}\right) d\theta \end{aligned}$$

where the second equality follows after substituting in for $q$ and the third equality follows from expanding the expression.

Consider each term separately. Observe that the first term in (8) can be written as

$$\int_{-\infty}^{\infty}\left(\left(\frac{\sigma_\varepsilon}{n\left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\theta}{\sigma_\varepsilon}\right)^2\right)\right)}\right)^2\frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}}\exp\left(-\frac{(\theta-at)^2}{2a\sigma_\varepsilon^2}\right)\right)d\theta$$

$$= \int_{-\infty}^{\infty}\left(\frac{\sqrt{2\pi}\sqrt{\sigma_\varepsilon^2}}{n^2\sqrt{a}}\exp\left(\frac{1}{\sigma_\varepsilon^2}\left(1-\frac{1}{2a}\right)\theta^2+\frac{\theta}{\sigma_\varepsilon^2}t-\frac{1}{2}\frac{a}{\sigma_\varepsilon^2}t^2\right)\right)d\theta$$

which converges to infinity when $a > 1/2$ because $(1-1/2a) > 0$ and is positive when $a \leq 1/2$.

Observe that the second term in (8) can be written as

$$\int_{-\infty}^{\infty}2\frac{\sigma_\varepsilon}{n\left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\theta}{\sigma_\varepsilon}\right)^2\right)\right)}(\theta-t)\frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}}\exp\left(-\frac{(\theta-at)^2}{2a\sigma_\varepsilon^2}\right)d\theta$$

$$= \int_{-\infty}^{\infty}\frac{2}{n}\frac{\theta-t}{\sqrt{a}}\exp\left(\frac{\theta^2a-(\theta-at)^2}{2a\sigma_\varepsilon^2}\right)d\theta$$

$$= \frac{2}{n\sqrt{a}}\exp\left(\frac{t^2a^2}{2(1-a)\sigma_\varepsilon^2}\right)\int_{-\infty}^{\infty}(\theta-t)\exp\left(\frac{-(1-a)\left(\theta-\frac{at}{1-a}\right)^2}{2a\sigma_\varepsilon^2}\right)d\theta$$

$$= \frac{2\sqrt{2\pi}}{n}\sqrt{\frac{\sigma_\varepsilon^2}{1-a}}\exp\left(\frac{t^2a^2}{2(1-a)\sigma_\varepsilon^2}\right)\int_{-\infty}^{\infty}(\theta-t)\frac{1}{\sqrt{2\pi}\sqrt{\frac{a\sigma_\varepsilon^2}{1-a}}}\exp\left(\frac{-(1-a)\left(\theta-\frac{at}{1-a}\right)^2}{2a\sigma_\varepsilon^2}\right)d\theta$$

$$= t\frac{2\sqrt{2\pi}}{n}\sqrt{\frac{\sigma_\varepsilon^2}{1-a}}\exp\left(\frac{t^2a^2}{2(1-a)\sigma_\varepsilon^2}\right)\left(\frac{a}{1-a}-1\right)$$

Thus, the expression

$$\int_{-\infty}^{\infty}\frac{2\sigma_\varepsilon(\theta-t)}{n\left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{1}{2}\left(\frac{\theta}{\sigma_\varepsilon}\right)^2\right)\right)}\frac{1}{\sqrt{2\pi}\sqrt{a\sigma_\varepsilon^2}}\exp\left(-\frac{(\theta-at)^2}{2a\sigma_\varepsilon^2}\right)d\theta$$

is finite and bounded when $a > 1/2$, is zero when $a = 1/2$, and is positive when $a < 1/2$.

Combining the two terms we observe that the sum of the terms converges to infinity when $a > 1/2$ and is positive when $a \leq 1/2$. This indicates that regardless of whether $a > 1/2$ or $a \leq 1/2$, a consumer with $t < 0$ will not choose to deviate and send a positive message. Analogous arguments apply when $t > 0$. ■

**Proof of Proposition 9:**

We show that for any $g$ and $l$, there is some sufficiently large $n$ such that it is an equilibrium for the firm to choose a product equal to the median of the sample reports and the strategic consumers to send a message equal to the point on the grid closest to their type. First note that if the strategic consumers always send a message equal to the point on the grid closest to their type, then in the limit as the number of survey responses becomes arbitrarily large, the probability the median of the sample reports equals the point on the grid that is closest to $\theta$ becomes arbitrarily close to 1, and the best response for the firm is to choose a product that is equal to the median of the sample reports. Also note that if the firm always chooses a product equal to the median of the sample reports, then the best response for the strategic consumers is to send a message equal to the point on the grid closest to their type, $m^*$. If a strategic consumer deviated by instead sending a message $m^* + jg$ for some positive integer $j$ (assuming such a deviation is feasible), this change would only change the sample median in cases where this median was in $[m^*, m^* + jg)$ before, in which case the deviation could only lead to a greater median than before. Since the firm is choosing a product equal to the sample median, this would lead to a product that is further from the strategic consumer's taste than before. This would not be a profitable deviation. Similar reasoning shows that a strategic consumer cannot profitably deviate by sending a message smaller than $m^*$. Thus, it is a best response for a strategic consumer to send a message equal to the point on the grid that is closest to her type. ■

# References

[1] Alonso, R., and N. Matouschek. "Optimal Delegation." *Review of Economic Studies*, Vol. 75, No. 1 (2008): 259–293.

[2] Benabou, R., and G. Laroque. "Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility." *Quarterly Journal of Economics*, Vol. 107 (1992): 921-958.

[3] Brunnermeier, M. K., and J. A. Parker. "Optimal Expectations." *American Economic Review*, Vol. 95, No. 4 (September 2005): 1092-1118.

[4] Carlson, P. "The Flop Heard Around the World." *The Washington Post* (September 4, 2007). Available at http://www.washingtonpost.com/wp-dyn/content/article/2007/09/03/AR2007090301419.html.

[5] Chen, Y. "Perturbed communication games with honest sender and naive receivers." *Journal of Economic Theory*, Vol. 146 (2011): 401-424.

[6] Churchill, Jr., G. A., and D. Iacobucci. *Marketing Research: Methodological Foundations*, 9th ed., Mason, OH: Thomson Learning (2005).

[7] Compte, O. "The Winner's Curse with Independent Private Values." Working Paper, CERAS-ENPC, 2002.

[8] Coca-Cola Company. "Coke Lore: The Real Story of New Coke." (2009). Available at http://www.thecoca-colacompany.com/heritage/cokelore_newcoke.html.

[9] Crawford, V. P., and J. Sobel. "Strategic Information Transmission." *Econometrica*, Vol. 50, No. 6 (1982): 1431-51.

[10] Dederer, A. "In Defense Of: The Pontiac Aztek." *The Truth About Cars* (October 25, 2006). Available at http://www.thetruthaboutcars.com/in-praise-of-the-pontiac-aztek/.

[11] Deshpande, R., and G. Zaltman. "Factors Affecting the Use of Market Research Information: A Path Analysis." *Journal of Marketing Research*, Vol. 19, No. 1 (1982): 14-31.

[12] Deshpande, R., and G. Zaltman. “A Comparison of Factors Affecting Researcher and Manager Perceptions of Market Research Use.” *Journal of Marketing Research*, Vol. 21, No. 1 (1984): 32-38.

[13] Farrell, J., and M. Rabin. “Cheap Talk.” *Journal of Economic Perspectives,* Vol. 10 (Summer 1996): 103-118.

[14] Feddersen, T., and W. Pesendorfer. “Voting Behavior and Information Aggregation in Elections with Private Information.” *Econometrica*, Vol. 65, No. 5 (1997): 1029-1058.

[15] Fowler, F. J. *Survey Research Methods.* 2nd ed. Newbury Park, CA: Sage (1993).

[16] Gordon, S. “On Infinite Cheap Talk Equilibria.” Working paper, Université de Montréal (August 2010).

[17] Grossman, G. M., and E. Helpman. *Special Interest Politics.* Cambridge, MA: The MIT Press (2001).

[18] Groves, R. M. *Survey Errors and Survey Costs.* New York: Wiley (1989).

[19] Harsanyi, J. S. “Games with Incomplete Information Played by ‘Bayesian’ Players, I-III, Part III. The Basic Probability Distribution of the Game.” *Management Science*, Vol. 14, No. 7 (1968): 486-502.

[20] Hartman, R. S., M. J. Doane, and C.-K. Woo. “Consumer Rationality and the Status Quo.” *Quarterly Journal of Economics*, Vol. 106, No. 1 (1991): 141-162.

[21] Kartik, N., M. Ottaviani, and F. Squintani. “Credulity, lies, and costly talk.” *Journal of Economic Theory*, Vol. 134 (2007): 93-116.

[22] Kelley, T., and J. Littman. *The ten faces of innovation: IDEO’s strategies for beating the devil’s advocate & driving creativity throughout your organization.* New York: Currency/Doubleday (2005).

[23] Kim, K.-M., and J. Pogach. “Honesty vs. White Lies.” Working paper, University of Pennsylvania (May 2009).

[24] Kreps, D., and R. Wilson. "Reputation and Imperfect Information." *Journal of Economic Theory*, Vol. 27 (1982): 253-279.

[25] Krishna, V., and J. Morgan. "Asymmetric Information and Legislative Rules: Some Amendments." *American Political Science Review*, Vol. 95 (2001): 435-52.

[26] Lehmann, D. R., S. Gupta, and J. H. Steckel. *Marketing Research.* Reading, MA: Addison-Wesley (1998).

[27] Li, M., and K. Madarász. "When Mandatory Disclosure Hurts: Expert Advice and Conflicting Interests." *Journal of Economic Theory*, Vol. 139, No. 1 (March 2008): 47-74.

[28] Manning, R. "Market Research by a Monopolist: A Bayesian Analysis." *Economica*, New Series, Vol. 46, No. 183 (1979): 301-306.

[29] Maricano, M. "Cheetos Lip Balm & More Bizarre Brand Extensions." *Wall Street Journal* (July 15, 2009). Available at http://online.wsj.com/article/SB124690095610001477.html.

[30] Meirowitz, A. "Polling Games and Information Revelation in the Downsian Framework." *Games and Economic Behavior*, Vol. 51 (2005): 464-489.

[31] Milgrom, P., and J. Roberts. "Limit Pricing and Entry under Incomplete Information: An Equilibrium Analysis." *Econometrica,* Vol. 50, No. 2 (1982): 443-459

[32] Morgan, J., and P. C. Stocken. "An Analysis of Stock Recommendations." *RAND Journal of Economics*, Vol. 34, No. 1 (2003): 183-203.

[33] Morgan, J., and P. C. Stocken. "Information Aggregation in Polls." *American Economic Review*, Vol. 98, No. 3 (June 2008): 864-896.

[34] Moorman, C., G. Zaltman, and R. Deshpande. "Relationships between Providers and Users of Market Research: The Dynamics of Trust within and between Organizations." *Journal of Marketing Research*, Vol. 29, No. 3 (1992): 314-328.

[35] Morris, B. "Steve Jobs speaks out." CNNMoney.com (March 7, 2008). Available at http://money.cnn.com/galleries/2008/fortune/0803/gallery.jobsqna.fortune/index.html

[36] Morris, S. "Speculative Investor Behavior and Learning." *Quarterly Journal of Economics*, Vol. 111, No. 4 (1996): 1111-1133.

[37] Ottaviani, M., and F. Squintani. "Naive Audience and Communication Bias." *International Journal of Game Theory*, Vol. 35, No. 1 (December, 2006): 129-150.

[38] Sawyer, A. G., and J. P. Peter. "The Significance of Statistical Significance Tests in Marketing Research." *Journal of Marketing Research*, Vol. 20, No. 2 (1983): 122-133.

[39] Scheinkman, J. A., and W. Xiong. "Overconfidence and Speculative Bubbles." *Journal of Political Economy*, Vol. 111, No. 6 (2003): 1183-1219.

[40] Sobel, J. "A Theory of Credibility." *Review of Economic Studies*, Vol. 52 (1985): 557-573.

[41] Van den Steen, E. "Rational Overoptimism (and Other Biases)." *American Economic Review*, Vol. 94, No. 4 (2004): 1141-1151.

[42] Van den Steen, E. "Interpersonal Authority in a Theory of the Firm." *American Economic Review*, Vol. 100, No. 1 (2010): 466-490.

[43] Venezia, I. "Optimal Investments in Market Research." *European Journal of Operational Research*, Vol. 18, No. 2 (1984): 198-207.

[44] Young, H. "Condorcet's Theory of Voting." *American Political Science Review,* Vol. 82, No. 4 (1988): 1231-1244.

**Figure 1: Level curve of Probability of a Flop as a function of Calibration-Alignment ratio and Noise-Signal ratio**

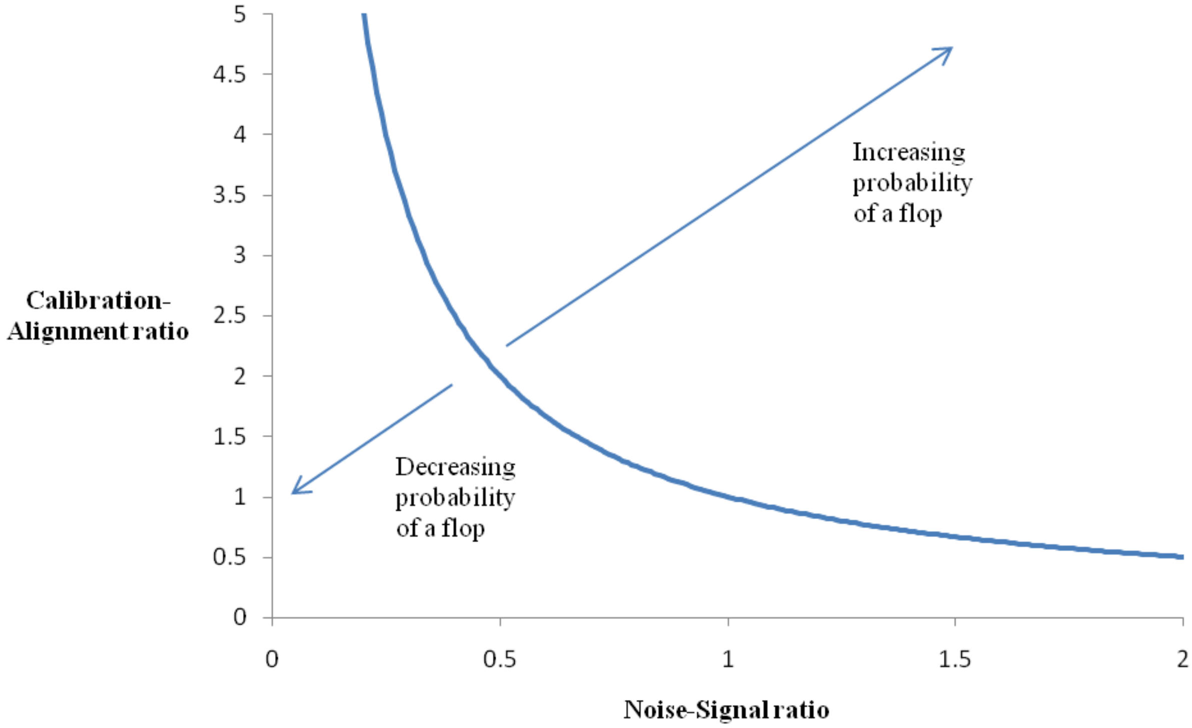

**Figure 2: Probability of a Flop as a function of $\rho$, the product of the Calibration-Alignment ratio and Noise-Signal ratio**