Title
Adverse selection and non-take inference with coherent risk and response scoring

Journal
Journal of the Operational Research Society, 64(1)

ISSN
0160-5682

Authors
Oliver, RM
Thaker, AM

Publication Date
2013

DOI
10.1057/jors.2012.3

Peer reviewed

Journal of the Operational Research Society (2013) 64, 70–85

www.palgrave-journals.com/jors/

Adverse selection and non-take inference with coherent risk and response scoring

RM Oliver[1]* and AM Thaker[2]

[1]*University of California at Berkeley, Berkeley, CA, USA; and* [2]*InfoCentricity, Inc., Novato, CA, USA*

The authors offer a mathematical model for adverse selection by individual borrowers based on preferences for offers and the default (Bad) or non-default (Good) status of booked accounts. We define the condition for borrower risk and response when there is no adverse selection (NAS). This definition provides us with a direct comparison between the prior and posterior conditional probabilities of default by an individual borrower who Takes an offer; this allows us to obtain estimates of differential response rates for individual borrowers and the Good/Bad odds for Take, Non-Take and Accept sub-populations. Performance of different response-risk segments allows us to compare price-driven risk elasticity and price-driven response elasticity in the presence of Good or Bad adverse selections; a special case applies when the borrower's capacity to repay is not an issue. We offer limited experimental results for selected price-risk segments where action-based risk and response scores are used to estimate borrower preferences. The critical role of Non-Take inference is described.

Journal of the Operational Research Society (2013) **64,** 70–85. doi:10.1057/jors.2012.3

Published online 11 April 2012

Keywords: adverse selection; risk and response scoring; non-take inference; retail credit; risk based pricing; price-risk; price-response elasticities

Introduction

The concept of adverse selection arises in a number of financial and risk contexts where there is a belief but, unfortunately, limited experimental evidence in retail credit, that the borrower may have more relevant information about his or her own future default risk than is available to the lender at the time a loan offer is made. The presence of asymmetric or hidden information to participants in borrowing/lending negotiations is often used to explain why rational borrowers and lenders have different preferences that may not be recognized or properly assessed by the other party. Because a borrower may have personal information about conditions that make him or her more risky to the lender, he or she may not want to reveal that information and thereby jeopardize a favourable offer.

Akerlof (1970) published his well-known paper on the market for lemons in the sale of second-hand autos where the seller knows more about the quality and condition of the automobile than does the prospective buyer. Many extensions and generalizations of games between buyer and seller with asymmetric information are described in Rasmussen (1994). In credit risk, the notion that a 'Bad' borrower is more likely to respond to a high-priced offer than a 'Good' borrower may be justified on theoretical grounds or the presence of asymmetric information. Altman *et al* (1998) define Adverse Selection (with reference to insurance in health plans) as 'the tendency for sick (healthy) people to join plans at high (low) cost' and suggest that adverse selection of Goods is likely to be greater with low-priced loans for low risk borrowers. The underlying thought is that of a subpopulation of individuals seeking insurance coverage; they 'move into or out of generous or overly restrictive plans'. Ausubel (1999) defines Adverse Selection in terms of the inferior risk characteristics of the pool of customers (borrowers) who accept an offer by comparing them with the customers who accept a 'better' offer. Edelberg (2004) uses a two-period model to study the interaction between adverse selection and moral hazard and finds the counter-intuitive result that higher-risk borrowers often pay lower loan rates than lower-risk borrowers. Cressy and Toivanen (2001) appear to have been the first authors linking asymmetric information, adverse selection and loan pricing in retail credit. P&R (Phillips and Raffard, 2009) state that (Bad) Adverse Selection in a portfolio of booked loans exists if the derivative of the default (Bad) rate with respect to loan rate (price) is increasing in the loan rate. As best we understand, none of these definitions identify the individual borrowers or estimate the number of adverse selections within a booked population; Altman makes it clear that, in the context of retail credit, increased counts of Goods or Bads

*Correspondence: *RM Oliver, Operations Research, University of California at Berkeley, 260 Southampton Ave., Berkeley, CA 94707, USA.*

E-mails: oliver@ieor.Berkeley.edu; amthaker@infocentricity.com

may be the result of subsets of individuals being attracted to favourable (low price) or unfavourable (high-price) loans. Most of the analysis reported by Agarwal *et al* (2010) compares the average risk scores of borrowers who book lender offers with those that do not. Their results were based on the availability of proprietary data that made it possible to directly estimate the risk of borrowers who do not take the offers of one lender but may take a possibly different offer from another lender in the marketplace. These different ways of describing adverse selection have the common feature that those individuals who purchase insurance coverage or who take loans or purchase credit have special circumstances or knowledge about their own situation that may not be available to the insurer/lender.

Other approaches, for example Fahner (2012), study propensity scores for modeling the likelihood of historically observed treatments, actions or offers that precede observable outcomes (the propensity score is a useful device to mitigate or eliminate overt selection bias when estimating causal effects of treatments); not only are these inferences retrospective in nature, they differ from the primary objectives and models of this paper which are to predict forward-looking counts of adverse selection that result from known, controllable actions or pricing policies designed to achieve desirable performance outcomes.

The contribution that we hope to make is the structural design and analysis of probability models that quantify the likelihood of adverse selection for individual borrowers. With known risk profiles and loan pricing policies, these models might lead to improved predictions of booked portfolio default rates. While this can be done without reference to risk or response scores, their contribution is significant in that many, if not most, risk predictions and portfolio acquisition and management policies in retail credit are directly related to the use of such scores. Risk and response scores are formally defined in a later section after the probability models are developed. In principle, three scores should be considered by the lender, two of which may be offer-dependent and are referred to as action-based scores, where $\mathbf{x}$ represents a collection of behavioural, financial and demographic data relevant to the prediction of borrower performance. We use a baseline default risk score, $s_p = s_p(\mathbf{x})$, developed without any information about a lender's offer, an action-based default risk score, $s_p(r) = s_p(\mathbf{x}, r)$, and an action-based response score, $s_q(r) = s_q(\mathbf{x}, r)$ whose performance outcome is a Take or Non-Take. In this paper, higher risk scores correspond to lower probabilities of default (Bad) and higher response scores correspond to higher probabilities of response (Take).

We denote the occurrence of a default as a Bad (B) and the occurrence of non-default as a Good (G). Goods and Bads are often used to denote other measures of performance such as late payment or fraud. We also denote the 'acceptance of a lender's offer by a borrower' as the outcome of a random event with two possible outcomes, one of which is that the borrower 'Takes' (T) the offer; the other possible outcome is 'Not-Take' (N). A Bad corresponds to default, a trapped state, not delinquency, which is a transient state. One measure that is often used in the literature to indicate the extent of adverse selection is the separation between the cumulative distributions of baseline (late payment) risk scores for Take and Non-Takes for different price offer segments (see Figure 1). We believe this measure is unsatisfactory, as differences in the score distributions do not quantify preferences or expected Bad counts posterior to the 'Take' event. Nothing is said about the internal composition or arrangement of the Goods and Bads within the Take or Non -Take populations nor is it obvious how one should discriminate between increased preferences of all borrowers for less expensive offers which are offset or exaggerated by differential preferences from Good or Bad borrowers.

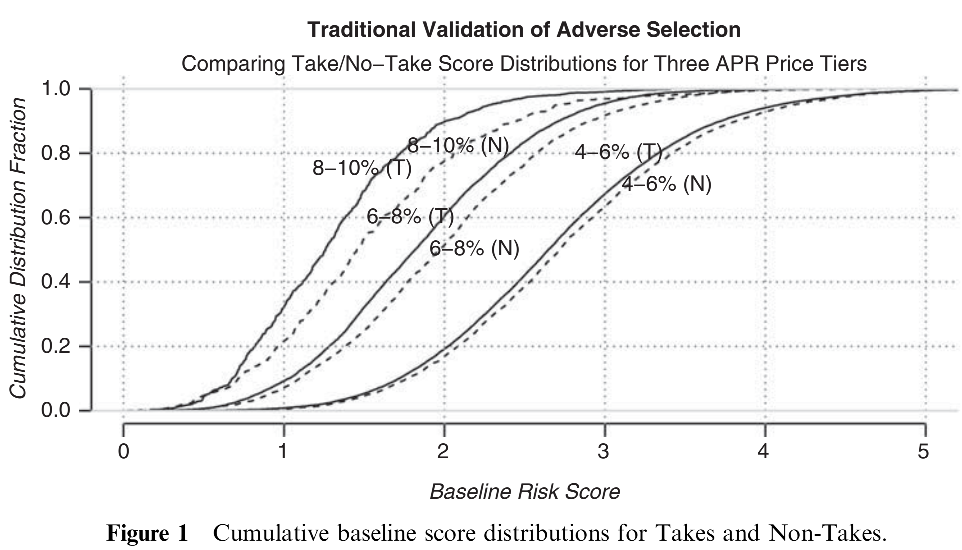

Figure 1 Cumulative baseline score distributions for Takes and Non-Takes.

Thomas (2009) defines Bad adverse selection in terms of the increased likelihood that booked borrowers at all loan rates and a given score will be Bad when compared with the prior prediction for all borrowers at the same score but without any condition or loan rate. P&R (2010) use a definition that incorporates '… price changes with the same credit score …'. We know of no study where there is an attempt to identify and count adverse selections from the risk and price composition of booked and unbooked borrowers.

In our experience many of the descriptions and explanations of adverse selection are unclear for at least four reasons: first, adverse selection is not formulated in terms of counts that differ from the counts expected when there is no adverse selection. Second, scores that are used to make risk predictions or set pricing policies may not reference appropriate data or performance outcomes at the time price offers are being made by the lender. If a lender offers a borrower a loan rate based on *delinquency scores* and then observes that the *number of defaults* (distinct from delinquencies) at a later time differs substantially from delinquency predictions, one should not conclude that adverse selection explains the difference. Third, there must be clarity in the definition of the performance variables in order to disentangle adverse selections from ordinary preferences and the aversion of most borrowers (both Good and Bad) to higher prices. Finally, symmetry suggests there should be an allowance for positive ('adverse') selection by Goods as well as negative adverse selection by Bads.

An example of bad adverse selection

Table 1 illustrates an example that occurs when an offer is made to a sub-population of 1500 risk-acceptable borrowers. Cell entries are either observed counts or expectations derived from risk scores. The columns correspond to Good/Bad (Non-Default/Default) outcomes with the final column representing the observed number of Takes and Non-Takes. By contrast, the bottom row corresponds to borrowers who are acceptable for loan offers and the middle rows correspond to borrower Takes and Non-Takes. The lightly shaded grey cells in the second and third columns contain the expected number of counts based on estimates of risk assigned to each borrower. Typically, scores are calculated and available before responses to offers are known and are based on relevant (predictive) characteristics of individual borrowers. The PopOdds for the predicted number of Goods and Bads among Takes, Non-Takes and Accepts in the second and third column is 7.5:1. Cells with dark shades of grey correspond to actual counts of responses and non-responses as well as Good/Bad performance of those who book. It is worth noting that the observed counts of Bads in the Take group is approximately 40% higher than the predicted number; this results in a lower Good/Bad odds of 5:1 for the Take sub-population. It appears to be an indication of Bad adverse selection even when there is no comparison with the change in counts when higher or lower priced offers are made to borrowers with similar risk assessments. Performance of Non-Takes who might have accepted offers with other lenders is generally unknown to the lender but it is possible that the Good/Bad count for Non-Takes, if they could be observed, would be different from that of the Take sub-population. Although it is not obvious how one should estimate cell counts in the Non-Take and Accept rows after the Good and Bad counts among Takes have been observed, once the uncertain count in the number of Bads among Non-Takes is specified as Z_1, one must maintain conservation in all unshaded cells by expressing counts in terms of the observed responses and Good/Bad counts among Takes. In this case the odds of Good in the sub-populations of Take (T), Non-Take (N) and Accepts (A) are given by

$$o_T = \frac{250}{50} = 5,\ o_N = \frac{1200 - Z_1}{Z_1},\ o_A = \frac{1450 - Z_1}{50 + Z_1}.$$

If Z_1 were equal to 200 then the odds of all three groups would be 5:1, that is no adverse selection. Smaller values of Z_1 lead to larger values of Good/Bad odds in the Accept population, and greater disparity between the odds of the Take and Non-Take populations.

Definition of adverse selection and NAS

Our characterization of the prediction-decision problem for assessing and making loan offers is shown in the extensive form influence diagrams of Figure 2 (a, b). The origination of loans is a two-stage decision problem where a baseline risk score provides the lender a risk assessment of those potential borrowers that are deemed to be acceptable for offers. A second, and later, decision by the

Table 1 Predicted and observed cell counts for Goods/Bads, Takes/Non-takes

	E [#Goods]	*E [#Bads]*	*Observed #Goods*	*Observed #Bads*	*Totals*
Takes (T)	264.7	35.3	250	50	300
Not takes (N)	1058.8	141.2	$1200-Z_1$	Z_1	1200
Total accepts (A)	1323.5	176.5	$1450-Z_1$	$50+Z_1$	1500

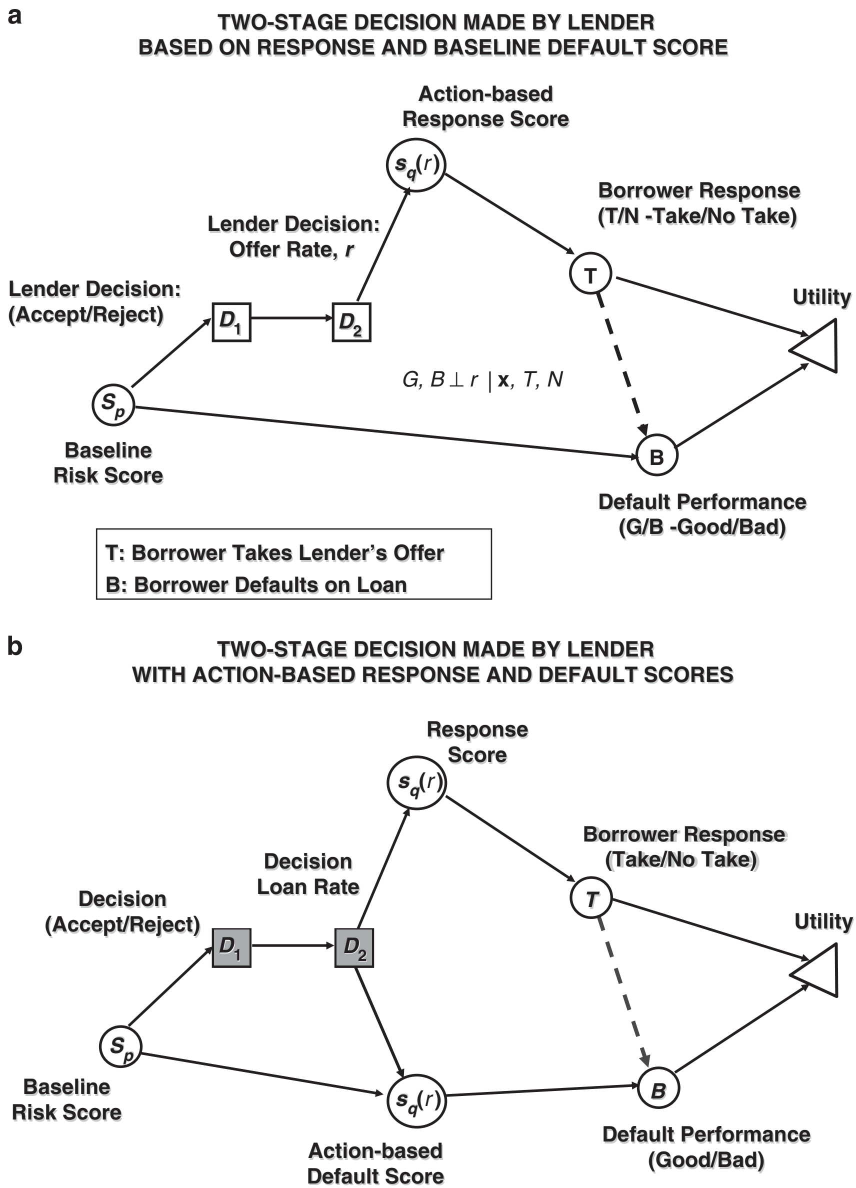

Figure 2 (a) Offer decisions, responses and defaults with baseline risk score; (b) Offer decisions, responses and defaults with action-based scores.

lender, possibly as the result of an optimization procedure based on business objectives and tradeoffs, matches the individual borrower with an offer; at that point in time the lender may also use an action-based response score to predict borrower preferences and likely response and risk outcomes.

Consider Figure 2 when the Good/Bad outcomes of performance are conditionally independent of borrower response (Take/Not-Take). To mimic a conditional independence structure think of a coin-tossing experiment in which one number and one outcome are stamped on each of two coins associated with a single borrower: the numbers on the coins are the prior probabilities of a Take and Bad, respectively. Independent tosses of each biased coin are made and the observed outcomes T or N and G or B are then stamped on the respective coins alongside the probability numbers. The G or B designation identified on any coin associated with an N on the paired coin is covered with masking tape so that an external observer, say the lender, cannot discern the default/non-default status of Non-Takes. In the case of T coins the G or B stamp on the paired coin is revealed to the lender.

When Adverse Selection is known to be present, the Good/Bad outcomes of performance depend on Take/Not-Take outcomes of borrower response, including the special case where the only risk score used is the baseline score mentioned above. Most risk control departments in banks and lending agencies are concerned about the implications of higher risk levels of those who respond to offers while marketing departments focus on response rates; thus there is a natural trade-off and tension that occurs in the decision-making process. As mentioned earlier, limited experimental evidence suggests that the Bad rate for those who Take an offer is higher than it would be for those who do not; alternatively, the odds of a Good in the Take sub-population is lower than it would be for the Non-Take sub-population among Accepts. Although there is general agreement about the existence of adverse selection there does not appear to be agreement on a precise definition on how it should be measured. To mimic the conditional dependence structure in Figure 2(b), think of the earlier coin-tossing experiment having the same two biased coins associated with a single borrower. The first coin toss determines whether we get the outcome T or N. If it is a T the probability on the second coin is changed and then tossed with the result of a G or B outcome revealed to the lender. In the case of an N, the probability on the second coin is also changed; it is tossed but the G or B outcome is *not revealed* to the lender whereas the G or B outcome on the paired coin that results from a T outcome is always revealed to the lender.

Clearly, Figure 2(a) contains one path where loan rate indirectly influences the Good/Bad outcomes whereas in Figure 2(b) we find there are two possible paths, one direct, one indirect, for the influence of loan rate on the Good/Bad outcomes. This feature appears to offer a sensible way to distinguish between direct and indirect adverse selection if that is important to the lender. In the discussions that follow we assume that the terms of the loan are entirely captured by the loan rate, r (or the premium over the risk-free rate) although the ideas put forth in this paper, can be generalized to the case where $\mathbf{r}$ is itself a vector. We use $\mathbf{x}$, a vector, to denote behavioural, financial, demographic characteristics and other relevant data to define the unconditional probability of a Bad in the Accept population as

$$p(B|\mathbf{x}, r) = \Pr\{Bad|data\,\mathbf{x},\ offer\,r\},\quad \mathbf{x} \in \chi, r \in \mathcal{R}, \tag{1}$$

where $\mathbf{x} \in \chi$, $r \in \mathcal{R}$ denotes sets of predictive and loan offer rate data available to the lender. For the prospective borrower the appropriate data might easily include privileged or private information that is not in χ. The probability of a Bad conditional on a Take is

$$p(B|T, \mathbf{x}, r) = \Pr\{Bad|Take, data\,\mathbf{x},\ offer\,r\}. \tag{2}$$

To obtain the marginal probability in (1), Bads among Non-Takes (N) must also be considered because

$$p(B|\mathbf{x}, r) = p(B|T, \mathbf{x}, r)p(T|\mathbf{x}, r) + p(B|N, \mathbf{x}, r)p(N|\mathbf{x}, r) \tag{3}$$

always holds. In this paper the existence of adverse selection for a borrower is an inequality, which states that the conditional default probability given a Take is larger than the unconditional prior default probability for a prospective borrower identified by $(\mathbf{x}, r)$,

$$p(B|T, \mathbf{x}, r) > p(B|\mathbf{x}, r). \tag{4}$$

Thomas (2009, p 171) uses a similar but different inequality involving risk scores where the right-hand side is conditioned on s, not $\mathbf{x}$, and excludes any condition on r. The many-to-one inequality $p(B|T, s, r) > p(B|s)$ for a fixed s and different values of r compares the posterior risk assessment of loan rate and score on borrowers who book with prior assessments made absent loan rate and booking information. It is not obvious that multiple borrowers, having the same s and different loan rates, will simultaneously meet or fail the stated inequality. Because of (3) and (4), coherence requires that

$$p(G|N, \mathbf{x}, r) > p(G|\mathbf{x}, r). \tag{5}$$

Let us now consider response rates of borrowers, obviously of great interest to lenders and organizations offering different credit products. Retail Credit experts, such as Gerbino and Rosenberger (personal e-mail communication), have suggested that in the presence of adverse selection a Bad borrower is more likely to take the offer than a randomly selected member of the Accept population. In contrast to (4), we have

$$q_B(\mathbf{x}, r) \triangleq p(T|B, \mathbf{x}, r) > p(T|\mathbf{x}, r) \triangleq q(\mathbf{x}, r). \tag{6}$$

The probability statements in (4) and (6) are equivalent because Bayes' Rule tells us that the connection between the conditional probability of a Bad given a Take and the conditional probability of a Take by a Bad is equality of the Bayes' factors (Good, 1961):

$$\frac{p(B|T, \mathbf{x}, r)}{p(B|\mathbf{x}, r)} = \frac{p(T|B, \mathbf{x}, r)}{p(T|\mathbf{x}, r)} \tag{7}$$

Another way to state this equality is that if we observe a higher default rate among the bookings than was expected from accepted borrower applicants then we must have a higher booking rate from those that default, i.e. evidence of a Take is informative.

We note that (7) *always* holds and lays the foundation for defining No Adverse Selection (*NAS*) as the equality of prior (before Take) and posterior (after Take) probabilities,

$$\textbf{NAS}: \quad \frac{p(B|T, \mathbf{x}, r)}{p(B|\mathbf{x}, r)} = \frac{p(T|B, \mathbf{x}, r)}{p(T|\mathbf{x}, r)} = 1. \tag{8}$$

This condition corresponds to Figure 2 where the arc between T and B nodes can be removed. Obviously, G can be substituted for B in the definition of NAS. We define Bad Adverse Selection (**BAS**) in (4) as the case where both Bayes' factors in (7) are greater than one, that is

$$\mathbf{BAS}: \quad \frac{p(B|T,\mathbf{x},r)}{p(B|\mathbf{x},r)} = \frac{p(T|B,\mathbf{x},r)}{p(T|\mathbf{x},r)} > 1. \tag{9}$$

Thus, a borrower who Takes an offer has a higher probability of default than before the Take event, an equivalent statement is that the Take probability conditioned on Bad is higher than it is for the entire group of prospects. We label adverse selection for Goods as **GPS** (Good positive selection) when the inequality is reversed and a Good is attracted to an under-priced loan or insurance policy. Notice, also, that (8) and (9) hold even when borrowers have the capacity to pay and the prior probability of default is independent of the loan rate.

Neither (8) nor (9) requires a comparison of different loan rates where one offer is in some sense inferior to another, *only that risk and response outcomes are conditionally dependent so that posterior probabilities of Good/Bad and Take/Not Take outcomes differ from their priors.* By contrast, Ausubel (1999) states that there is experimental evidence of adverse selection from inferior offers and P&R (2010) require an examination of the derivatives of borrower default or, as a minimum, the comparison of different price offers:

$$p(B|T,\mathbf{x},r') > p(B|T,\mathbf{x},r), \quad r' > r. \tag{10}$$

In our view the difficulty with using (10) in a definition of adverse selection is that because of the price change, it becomes difficult to distinguish between two different effects that may be taking place simultaneously: the offer at the lower rate r is more attractive and preferred by most borrowers to one at a higher rate. At the same time the fractional Good/Bad composition may be changing so that it is not obvious how to disentangle ordinary preferences from adverse selections.

If we are certain that NAS holds in Table 1 then the PopOdds of Takes, Non-Takes and Accepts are equal and the counts would have to be adjusted for each Z_1 as shown in Table 2.

For purposes of discussion let us assume we know there are Bad adverse selections in the count of Takes in the first row. For a given value of Z_1 we would have to reduce the number of Bad adverse selections by Y_1, to obtain NAS; thus Y_1, the number of Bad adverse selections, is the difference between the count of observed Bads when adverse selection is present and the number of Bads among Takes under NAS. To maintain conservation we must also adjust three other cells by an internal reallocation of Bad and Total responses; note that counts in some cells depend only on Z_1, some only on Y_1, but only one cell has a count which depends on both and equals $Z_1 + Y_1$. In this simple example there is BAS when $0 \leqslant Y_1 \leqslant 50$.

Table 2 Counts of Goods/Bads, Takes/Not takes under NAS

	Goods	*Bads*	*Totals*
Takes (T)	250	$50-Y_1$	$300-Y_1$
Not takes (N)	$1200-Z_1$	Z_1+Y_1	$1200+Y_1$
Total accepts (A)	$1450-Z_1$	$50+Z_1$	1500

From the definition in (8),

$$o_A = \frac{1450 - Z_1}{50 + Z_1} = \frac{1200 - Z_1}{Z_1 + Y_1} = \frac{250}{50 - Y_1}. \tag{11}$$

Solving for Y_1, in terms of Z_1 yields the count of Bads in the Take row that would result in NAS for each Z_1,

$$Y_1(Z_1) = \frac{50o_A - 250}{o_A} = \frac{60,000 - 300Z_1}{1450 - Z_1}. \tag{12}$$

Conservation of counts must always be satisfied in Table 2. When BAS is thought to be present, NAS could only have been achieved when the total number of responses is less than or equal to the booked accounts; the size of the reduction depends on our initial estimate for Z_1. Clearly, Non-Take inference is an essential requirement for the estimation of NAS. In our case we want to infer Bads among Non-Takes, rather than the traditional Reject Inference which is to infer the Bads that might have resulted had the lender accepted members of the Reject sub-population.

When adverse selection is present the marginal and differential (conditional) Take rates are given by

$$q = \Pr\{T\} = \frac{300}{1500} = 0.20,$$
$$q_G \triangleq \Pr\{T|G\} = \frac{250}{1450 - Z_1},$$
$$q_B \triangleq \Pr\{T|B\} = \frac{50}{50 + Z_1}$$

but with NAS the Take rate for Bads and the marginal Take rate are equal and depend on both Y_1, and Z_1:

$$q = \Pr\{T\} = \Pr\{T|B\} = \Pr\{T|G\} = \frac{50 - Y_1(Z_1)}{50 + Z_1} = \frac{300 - Y_1(Z_1)}{1500}$$

As an example consider the case when $Z_1 = 200$ and $Y_1 = 0$ the PopOdds of all rows equals 5:1 and we have NAS. The Take rates of Goods and Bads are equal to one another and to the marginal rate, $q = 0.20$. On the other hand, when $o_A = 10$, we find from (11) and (12) that $Z_1 = 86.4$, $Y_1 = 25$ which means that adverse selection accounts for about half the number of observed Bads in the Take population. With adverse selection the fraction of

Bads among Takes is $1/6=0.167$ but is less among Non-Takes because

$$p\{B|N\} = \frac{Z_1}{1200} = \frac{86.4}{1200} = 0.072. \quad (13)$$

In this case, conditional Take rates for the Good and Bad sub-populations differ greatly from one another and are given by

$$q_G = \frac{250}{1363.6} = 0.183, \quad q_B = \frac{50}{136.4} = 0.367 \quad (14)$$

Higher priced offer

Assume that a new higher priced offer is made to the same population. If the loan rate in Tables 1 and 2 was r and the increase in the loan price is Δr, the new rate is $r+\Delta r$. It is found that the observed Good/Bad counts are 160 and 40 as shown in Table 3 so that the Take count is reduced to 200 and the number of Non-Takes is increased to 1300. The observed counts in top row and right-most column in Table 3 correspond to the case where Y terms are zero and there are no entries in the unshaded cells. Subscript 2 denotes effects of the higher priced offer. These new counts are due to two distinct effects: the reduced appeal of the higher priced offer to all borrowers in combination with what appears to be a disproportionate increase in the fraction of observed defaults (Bads), that is the adverse selects.

Let us assume, for the moment, that the default risks of the borrowers in the Accept population are unaffected by the change in the offer rate with a PopOdds for Accepts given by o_A used in Table 2—this would correspond to the case where all borrowers in the Accept group are unaffected by the higher loan rate and have the capacity to pay. When we allocate the count of Goods and Bads in the Non-Take row in Table 3 (as we did in Table 2) we find that the count of Bads among Takes is always 10 larger for the higher priced offer; furthermore, the difference in the count of Bad adverse selections between the high priced and low priced offer can be positive or negative.

Using the same calculations as in (11), the Bads and the number of adverse selections among Takes for both offers are shown in Table 4 for different values of the Accept PopOdds, o_A.

Let us assume that the PopOdds for Accepts is $o_A=10$, and that the loan rate of offer 1 is 6% and offer 2 is 7%, that is an increase of 100 basis points. A discrete version of price-response elasticity (negative of the traditional elasticity used in the economics literature) for the conditional Take rate of Goods and Bads can be easily calculated. The low price offer Take probabilities are obtained from (14) and the high price offer Take probabilities can be calculated from the second row of Table 4. Thus we obtain the Good/Bad elasticities:

$$\begin{aligned}\varepsilon_B^{(T)} &= \frac{\Delta q_B}{\Delta r}\frac{r}{q_B} = -6\frac{0.073}{0.367} = -1.20 \\ > \varepsilon_G^{(T)} &= \frac{\Delta q_G}{\Delta r}\frac{r}{q_G} = -6\frac{0.066}{0.183} = -2.16. \quad (15)\end{aligned}$$

In this example the price response elasticity for Bads is larger than for Goods which means that the percentage decrease in response rates for Goods is greater (in absolute value) than for Bads. The marginal price-response elasticity is therefore

$$\begin{aligned}\varepsilon^{(T)} &= \frac{\Delta q}{\Delta r}\frac{r}{q} = \varepsilon_G^{(T)} p(G|T) + \varepsilon_B^{(T)} p(B|T) \\ &= -2.16(0.833) - 1.20(0.167) = -2.00.\end{aligned}$$

Differential elasticities

To show how differential Take rates are influenced by the preferences of borrowers and the presence of adverse selection, we define a price-risk elasticity in addition to the traditional and well-known response ('price-volume') elasticities for Takes. Even though it is slightly more

Table 3 Counts of Goods/Bads, Takes/Not Takes without adverse selection in higher priced offer

	Goods	*Bads*	*Totals*
Takes	160	$40-Y_2$	$200-Y_2$
Not takes	$1300-Z_2$	Z_2+Y_2	$1300+Y_2$
Total accepts	$1460-Z_2$	$40+Z_2$	1500

Table 4 Adverse selection counts for low and high priced offer with different accept PopOdds

Accept PopOdds, o_A	*Non-take bads, Z_1*	*Adverse selection, Y_1*	*Non-take bads, Z_2*	*Adverse selection, Y_1*	*Change Y_2-Y_1*
20:1	21.4	37.5	31.4	32	−5.5
10:1	86.4	25	96.4	24	−1
9:1	100	22.2	110	22.2	0
8:1	116.7	18.8	126.7	20	1.2
7:1	137.5	14.3	147.5	17.1	2.8
6:1	164.3	8.3	174.3	13.3	5
5:1	200	0	210	8	8
4:1	250	−12.5	260	0	12.5

complicated, it is important to use a notation that adheres to the standard convention for conditional dependence and independence. *Price-risk elasticity* is the percentage change in probability of default as a function of small percentage changes in price or loan rate. We denote the conditional *price-response elasticity* of Takes or Non-Takes (superscript T or N) with a B or G subscript. Because most of these elasticities cannot be measured directly there is a need for Non-Take inference to establish the magnitude of differential Take rates in distinct Good/Bad sub-populations. The conditional price-response elasticities are:

$$\varepsilon_B^{(T)} \triangleq \frac{\partial q_B(\mathbf{x},r)}{\partial r}\frac{r}{q_B(\mathbf{x},r)},$$
$$\varepsilon_G^{(T)} \triangleq \frac{\partial q_G(\mathbf{x},r)}{\partial r}\frac{r}{q_G(\mathbf{x},r)}, \tag{16}$$

with a similar definition for Non-Takes. The unconditional (marginal) price-response elasticity for Takes is

$$\varepsilon^{(T)} \triangleq \frac{\partial p(T|\mathbf{x},r)}{\partial r}\frac{r}{p(T|\mathbf{x},r)} = \frac{\partial q}{\partial r}\frac{r}{q(\mathbf{x},r)} = \frac{\partial \ln q(\mathbf{x},r)}{\partial \ln r}. \tag{17}$$

With this definition, price-response elasticity is non-positive for Takes and non-negative for Non-Takes. Marginal *price-risk elasticity* is defined as:

$$\delta^{(B)} \triangleq \frac{\partial p(B|\mathbf{x},r)}{\partial r}\frac{r}{p(B|\mathbf{x},r)} = \frac{\partial \ln p(B|\mathbf{x},r)}{\partial \ln r}. \tag{18}$$

Just as price-response elasticities have been used to measure the change in borrower preferences for percentage increases in rate, risk elasticities can be thought of as the percentage change in 'Goodness' or 'Badness' for percentage increases in rate. While the marginal price-risk elasticity refers to the default rate of members of the Accept population, the conditional price-risk elasticities represent defaults within Take and Non-Take groups. The roles of superscripts and subscripts are reversed:

$$\delta_T^{(B)} \triangleq \frac{\partial p(B|T,\mathbf{x},r)}{\partial r}\frac{r}{p(B|T,\mathbf{x},r)},$$
$$\delta_N^{(B)} \triangleq \frac{\partial p(B|N,\mathbf{x},r)}{\partial r}\frac{r}{p(B|N,\mathbf{x},r)}. \tag{19}$$

Price-risk elasticity for Takes is positive with BAS and negative with GPS. Because of the equality of Bayes' factors in (7), unequal differential Take rates for Goods and Bads exist if and only if there are unequal differential Good/Bad default rates for Take and Non-Take sub-populations. From symmetry arguments it is easy to see that there are a total of six risk and six response elasticities for Takes, Non-Takes and Accepts,

$$\mathbf{E} = \begin{bmatrix} \varepsilon_G^{(T)} & \varepsilon_B^{(T)} & \varepsilon^{(T)} \\ \varepsilon_G^{(N)} & \varepsilon_G^{(N)} & \varepsilon_G^{(N)} \end{bmatrix}, \quad \mathbf{\Delta} = \begin{bmatrix} \delta_T^{(G)} & \delta_T^{(B)} \\ \delta_N^{(G)} & \delta_N^{(B)} \\ \delta^{(G)} & \delta^{(B)} \end{bmatrix}, \tag{20}$$

where superscripts in $\mathbf{E}$ are associated with rows and subscripts with columns, the opposite being true with $\mathbf{\Delta}$. Note, also, that the final column in the former and the bottom row in the latter refer to the marginal elasticities (without subscripts). The vector of price-response elasticities for Takes corresponds to the top row of $\mathbf{E}$ and the price-risk elasticities for Bads correspond to the rightmost column of the $\mathbf{\Delta}$ matrix in (20). One case of special interest occurs when (18) is zero which means that the bottom row of $\mathbf{\Delta}$ is zero.

In much of what follows and in most of the practical risk and credit scoring applications that we are familiar with, default predictions do not explicitly incorporate offer or rate terms which means that the lenders are not concerned about the borrower's ability or 'capacity to pay'. When the capacity to pay is not an issue the prior probability of default for Accepts is independent of the loan rate; this represents a special case that yields simplified formulas for risk and response elasticities. Some models in the literature assume the presence of Bad adverse selection when action-based default scores include a negative term proportional to the loan offer rate, r. In our models, adverse selection only depends on whether the added information from a Take or Non-Take does or does not change the posterior default score or probability; for this reason, the use of a baseline score which is conditionally independent of offer terms, r, may nevertheless lead to BAS (GPS) when potential borrowers are disproportionately attracted to certain offers.

Conservation of price-risk and price-response elasticities

To capture the relationship between Take and Default rate elasticities we again use Bayes' Rule. The posterior conditional probability of a Bad given a Take is given by

$$p(B|T,\mathbf{x},r) = \frac{p(T|B,\mathbf{x},r)}{p(T|\mathbf{x},r)}p(B|\mathbf{x},r) = \frac{q_B(\mathbf{x},r)}{q(\mathbf{x},r)}p(B|\mathbf{x},r). \tag{21}$$

If the Take probability of the Bad sub-population is the same as the unconditional response rate, the last ratio equals one, the probability of a Bad in the Take population is equal to the probability of a Bad for the population of Accepts and we have NAS. It is a straightforward calculation to express the partial derivative of the posterior probability of a Bad in (21) with respect to the loan rate as

$$\frac{\partial p(B|T,\mathbf{x},r)}{\partial r}=\frac{\partial}{\partial r}\frac{q_B(\mathbf{x},r)p(B|\mathbf{x},r)}{p(T|\mathbf{x},r)}$$
$$=\frac{1}{p(T|\mathbf{x},r)}\left(p(B|\mathbf{x},r)\frac{\partial q_B(\mathbf{x},r)}{\partial r}+q_B(\mathbf{x},r)\frac{\partial p(B|\mathbf{x},r)}{\partial r}-\frac{q_B(\mathbf{x},r)p(B|\mathbf{x},r)}{p(T|\mathbf{x},r)}\frac{\partial p(T|\mathbf{x},r)}{\partial r}\right). \quad (22)$$

The middle term inside the large parenthesis is proportional to the change in the Bad probability of accepted borrowers. Although this term may, in some cases, be unaffected by changes in the loan rate, in general, its presence and influence must be assessed. By factoring out common terms we can write (22) in terms of the difference between price-response elasticity for all respondents and the conditional price-risk elasticity for Bads as well as a term which is proportional to the ratio of the conditional to the unconditional Bad probability.

By factoring out the three terms on the rhs of (21) from (22), we obtain the equivalent expression

$$\frac{\partial p(B|T,\mathbf{x},r)}{\partial r}=\frac{1}{r}\frac{q_B(\mathbf{x},r)p(B|\mathbf{x},r)}{p(T|\mathbf{x},r)}\left(\varepsilon_B^{(T)}-\varepsilon^{(T)}+\frac{\partial p(B|\mathbf{x},r)}{\partial r}\frac{r}{p(B|\mathbf{x},r)}\right)$$
$$=\frac{p(B|T,\mathbf{x},r)}{r}\left(\varepsilon_B^{(T)}-\varepsilon^{(T)}+\frac{\partial p(B|\mathbf{x},r)}{\partial r}\frac{r}{p(B|\mathbf{x},r)}\right). \quad (23)$$

On dividing both sides by the factor to the left of the large parenthesis, we obtain a ‘no free lunch’ conservation equation for the deviations in conditional elasticities from their marginals, namely

$$\delta_T^{(B)}-\delta^{(B)}=\varepsilon_B^{(T)}-\varepsilon^{(T)}. \quad (24)$$

The risk exchange equation in (24) always holds and is independent of whether the borrower does or does not have the capacity to pay. The left-hand side is the deviation of the conditional price-risk elasticity for Bads among Takes from its marginal value; this always equals the deviation in the conditional price-response elasticity for Takes among Bads from its marginal. By symmetry it follows that there are four Take/Non-Take price-risk counterparts to the traditional price-response elasticities; each conservation equation corresponds to the paired differences between conditional risk and response elasticities.

Expressions similar to (22) can be derived for the rate of change of the Good probability where G replaces B in the conditional statements.

$$\frac{\partial p(G|T,\mathbf{x},r)}{\partial r}\frac{r}{p(G|T,\mathbf{x},r)}-\frac{\partial p(G|\mathbf{x},r)}{\partial r}\frac{r}{p(G|\mathbf{x},r)}=\varepsilon_B^{(T)}-\varepsilon^{(T)}. \quad (25)$$

Because the conditional Good/Bad probabilities must sum to one, their partial derivatives sum to zero which means that one can convert the derivatives for Good probabilities to equivalent expressions involving Bad probabilities. Adding (25) to (23) and simplifying terms yields the result that the price-risk elasticity for Takes consists of two terms, one being proportional to the difference in the differential price-response elasticities, the other being proportional to the price-risk elasticity for Accepts,

$$\delta_T^{(B)}=\frac{\partial p(B|T,\mathbf{x},r)}{\partial r}\frac{r}{p(B|T,\mathbf{x},r)}$$
$$=(\varepsilon_B^{(T)}-\varepsilon_G^{(T)})p(G|T,\mathbf{x},r)+\delta^{(B)}\left(\frac{p(G|T,\mathbf{x},r)}{p(G|\mathbf{x},r)}\right). \quad (26)$$

If borrowers have the capacity to pay, the price-risk profiles of Accepts is unaffected by the loan rate; thus, the rightmost term in (26) vanishes and we obtain an equation that is similar, but not identical, in structure to one obtained earlier by P&R (2010).

If borrowers do not have the capacity to pay and the action-based default score is explicitly dependent on the loan rate, there is an additional contribution from loan rate changes to the marginal default probabilities in the lender's Accept population. If there is no adverse selection and the posterior to prior probability of default equals one we recover a special case that expresses risk elasticity deviations in terms of the deviations in response elasticities.

Bad adverse selection when borrowers have the capacity to pay

In many practical applications of credit risk decisions, the lender only uses a baseline score such as a Bureau score or a proprietary internal score based on past performance and relevant predictors and does not find the need to include the loan rate as a predictor of Good/Bad.

This coincides with the assumption that there is no ‘Capacity’ effect, that is the loan rate offer does not, in and of itself, directly influence the probability of default even though there may be an indirect linkage to loan rate resulting from the Take and Non-Take preferences of borrowers. It is interesting that this takes us back to one of the three components of the original 'Three C's Rule' referring to Character, Capacity and Collateral (Lewis, 1992), where, traditionally, decisions were based on making judgments about the ability and capacity of the borrower to repay loans as promised.

As we have already mentioned there is an important special case when a change in the loan rate does not affect the probability of a Bad or Good in the population of borrowers acceptable to lenders; in such cases prices only influence the internal reallocation of Goods and Bads within the Take and Non-Take sub-populations but not the overall risk of the Accept population. Derivatives of the marginal Bad rate and price-risk elasticity (without a condition on Take (T) or Non-Take (N)) can now be set equal to zero:

$$\delta^{(B)} = \delta^{(B)}(\mathbf{x}, r) = \frac{\partial p(B|\mathbf{x}, r)}{\partial r} \frac{r}{p(B|\mathbf{x}, r)} = 0. \qquad (27)$$

Even though this simplifying condition does not always hold we emphasize that adverse selection can still result from the internal reallocation of Goods and Bads within Takes. In such cases the final term on the right-hand side of (26) vanishes so that the rate of change of the risk of default (conditional Bad probability for Takes) is inversely proportional to the loan rate and directly proportional to the product of default/non-default probabilities and the difference between the differential price-response elasticities:

$$\frac{\partial p(B|T, \mathbf{x}, r)}{\partial r} = (\varepsilon_B^{(T)}(\mathbf{x}, r) - \varepsilon_G^{(T)}(\mathbf{x}, r)) \times \frac{p(B|T, \mathbf{x}, r)p(G|T, \mathbf{x}, r)}{r}. \qquad (28)$$

The condition is again on the Take group, not the population of all borrowers as a whole. Assume that the probability of a Good is positive. An equivalent statement to (28) is that price-risk elasticity for Takes is proportional to the product of the Good probability and the difference in the conditional price-response elasticities is now:

$$\begin{aligned}&\delta_T(\mathbf{x}, r)\\ &= (\varepsilon_B^{(T)}(\mathbf{x}, r) - \varepsilon_G^{(T)}(\mathbf{x}, r))p(G|T, \mathbf{x}, r) > 0\\ &\quad \textit{iff} \quad \varepsilon_B^{(T)}(\mathbf{x}, r) > \varepsilon_G^{(T)}(\mathbf{x}, r). \qquad (29)\end{aligned}$$

Thus the positivity of (29) is the result of greater price-response sensitivity for Bads than for Goods and is zero only when conditional elasticities are equal and there is no Adverse Selection—note that if the probability of a Good among the Takes is zero, we must have either BAS or NAS. If the inequality is reversed the statement holds for Goods. Furthermore, the price-response elasticity for the group of borrowers who take the lender's offer is the unconditional expectation of the conditional Good and Bad price-response elasticities:

$$\begin{aligned}\varepsilon^{(T)}(\mathbf{x}, r) &= \varepsilon_G^{(T)}(\mathbf{x}, r)p(G|T, \mathbf{x}, r)\\ &\quad + \varepsilon_B^{(T)}(\mathbf{x}, r)p(B|T, \mathbf{x}, r). \qquad (30)\end{aligned}$$

We emphasize that, in general, (28) and (29) do not hold (eg, as can occur when there is a Capacity effect); when they do, conditional and marginal risk elasticities can be expressed in terms of their response counterparts. From Tables (2) and (3) when $o_A = 10$ we obtain (recall that our response elasticities are the negative of the traditional definitions):

$$\mathbf{E} = \begin{bmatrix} -2.16 & -1.20 & -2.00 \\ 0.49 & 0.70 & 0.50 \end{bmatrix},$$

$$\mathbf{\Delta} = \begin{bmatrix} -0.16 & 0.80 \\ -0.015 & 0.195 \\ 0 & 0 \end{bmatrix}$$

Using the notation in (24), we confirm that $\varepsilon_B^{(T)} - \delta_T^{(B)} = -(1.20 + 0.8) = -2.0 = \varepsilon^{(T)}$. With BAS, the response elasticity of Bads is larger than that of Goods while conditional default risk elasticities are negative for Goods and positive for Bads.

Risk and response scores

Thus far, the definitions of adverse selection and elasticities have not explicitly required either risk or response scores described in the graphs of Figures 1 and 2. While they were originally developed and used to measure relative risk performance of borrowers they are often used as guidelines to help design loan offers that simultaneously recognize borrower risk, preferences and attractiveness of customized loans. Unfortunately, there are a number of difficulties that complicate the design of risk-based pricing policies. The first is the identification of timely and relevant characteristics that influence risk and response scores, a second is the specification of the conditional independencies that influence outcomes and a third is the degree to which adverse selection may complicate preference assessments.

We define the *baseline log odds default score* as the score when no offer rates or terms are included. It is well known that the score includes two additive components: one term depends on the log of population odds, the other on the weight of evidence or log of information odds which measures the relative importance of the Good and Bad profiles:

$$s_p(r) = s_p(\mathbf{x}) \triangleq \ln o(G|\mathbf{x}) = \ln\frac{p(G|\mathbf{x})}{p(B|\mathbf{x})} = \ln\frac{p(G)}{p(B)} + \ln\frac{f(\mathbf{x}|G)}{f(\mathbf{x}|B)}, \quad \mathbf{x} \in \chi. \tag{31}$$

When the default score includes the new information associated with each offer, whose attractiveness to the borrower is uncertain until it has been Taken or Not-taken, we modify the definition in (31) so that the action-based score is defined as

$$s_p(r) = s_p(\mathbf{x}, r) \triangleq \ln o(G|\mathbf{x}, r) = \ln\frac{p(G|\mathbf{x}, r)}{p(B|\mathbf{x}, r)}, \quad \mathbf{x} \in \chi, r \in \mathcal{R} \tag{32}$$

The key role that risk and response scores provide is that (i) over many decades they have exhibited remarkable stability and reliability in providing trustworthy assessments of well-defined risk outcomes and (ii) a well-calibrated scalar score is a sufficient statistic for large vectors of behavioural, demographic and financial data. A well-calibrated default score provides as much information as is available from the original data on which the score was based so that

$$p(B|\mathbf{x}, r) = p(B|\mathbf{x}, r, s_p(r)) = p(B|s_p(r)). \tag{33}$$

As mentioned earlier it should be understood that the score s_p is shorthand for the baseline $s_p(\mathbf{x})$ and $s_p(r)$ for the action-based score $s_p(\mathbf{x}, r)$. These scores are scalars even though $\mathbf{x}$ is usually a high dimensional vector and r might include financial information other than the loan rate. An example of a baseline risk score could be a Bureau or Agency score where r is not explicitly included; although Bureau records contain important financial information of each borrower, details of financial terms associated with a particular loan are usually missing. A late-payment Bureau score should not substitute for a default score, even though the former is often used as a characteristic in development of the latter. In general, baseline and action-based scores are computed at different times, with different information and with different performance outcomes in the life cycle of the origination process. In the case of a finite number of different offers, there are as many default scores, as there are offers plus one, the latter being a baseline score. In comparing baseline and action-based scores as defined in (31) and (32) we understand that the probability of default derived from a baseline score augmented by a loan rate is not the same as the probability of default calculated from an action-based score using the original data, $\mathbf{x}$, augmented by a loan rate:

$$p(B|\mathbf{x}, r) = p(B|s_p(r)) \neq p(B|s_p, r). \tag{34}$$

In the right-hand side the data available to the lender would consist of a two-element vector: the baseline score and loan rate whereas on the left-hand side the relevant behavioural/demographic/financial data as well as loan rate is available in the construction of the action-based scorecard. The score weights for attributes, other than the loan rate itself, associated with an action-based score can be non-linear functions of the loan rate so that two individuals with different behavioural/demographic/financial data but identical baseline risk scores may have very different action-based risk scores *even when the loan rate offers are identical*. To say it differently, prospective borrowers with identical baseline scores can have larger or smaller action-based scores; thus, probabilities of default that include the effect of the loan rate r may be larger or smaller than those associated with the baseline score. The fact that the coefficient of a continuous variable, r, is negative in a regression of log odds is not a guarantee that the probability of a Bad decreases with increasing loan rate, r! We illustrate the differences between baseline and action-based risk scores in the scatter diagram of Figure 3 comparing some typical action-based and baseline default scores of several hundred distinct borrowers—the diagonal line corresponds to equality in both scores and each dot corresponds to a unique individual borrower.

The concept of risk or default scores extends to response scores; there is seldom any need for a baseline response score as the primary influence on response is the loan rate or price. When T and N denote Take and Non-Take borrower outcomes we have

$$s_q(r) = s_q(\mathbf{x}, r) \triangleq \ln\frac{p(T|\mathbf{x}, r)}{p(N|\mathbf{x}, r)} \tag{35}$$

As with risk scores, we assume that response scores represent sufficient statistics for the prediction of Take and Not-Take outcomes. It is clear that risk and response scores in (32) and (35) are dependent not only because they each depend on r but because there may be several common characteristics in the $\mathbf{x}$ vector that affect both scores. As mentioned earlier, our convention is that a larger risk score implies a less risky borrower and a higher response score corresponds to a more attractive offer and higher probability of a Take.

In the diagram of Figure (2a) that displays a baseline score, the prior probability of a Bad is *indirectly influenced* by the offer rate but, because of the possible presence of BAS or GPS, the differential responses of Good/Bad borrowers to different offers inherit the influence of loan rate in the posterior probability of borrowers

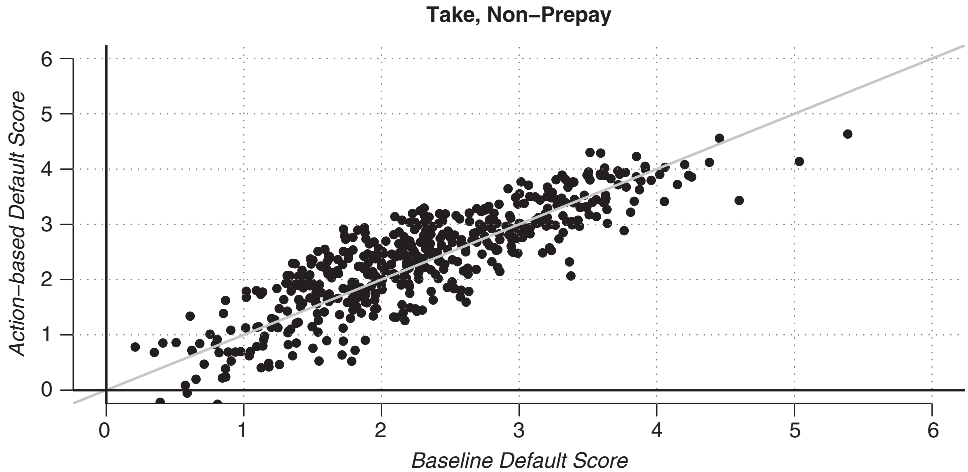

Figure 3 Scatter diagram for baseline and action-based scores.

who Take. We point out that even in the case of an action-based default score $s_p(r)$, where the loan rate directly affects the priors, there may be no adverse selection if, from (8), we have:

$$p(B|T,\mathbf{x},r) = p(B|\mathbf{x},r) = p(B|s_p(r)), \quad B,G\perp T,N|\mathbf{x},r. \qquad (36)$$

Again, a Take may provide a second-order effect or indirect influence by r that represents adverse selection. Obviously, our view of direct and indirect influences of loan rate on adverse selection differs from P&R (2010) who use a measure for Total Adverse Selection that is the sum of expected default rate derivatives when the profile of risk scores is held constant (Indirect Adverse Selection) and expected default rates of booked accounts due to changes in the profile of risk scores (Direct Adverse Selection).

Default rates in booked portfolios

We thought it would be useful to see how probability statements of individual borrowers can be incorporated into the prediction of average default rates in booked portfolios. Such calculations obviously depend on the rate/score composition of a booked portfolio which is obtained from the rate offers and the score profiles of the accepted borrowers. Consider, for example, a portfolio of borrowers whose default score cutoff by the lender is s_C (subscript C). For purposes of simplicity in this paragraph we use the notation s without a subscript p to denote baseline risk score. It is well known that the categorical forecast for default of a randomly selected member of this portfolio can be expressed in terms of the tail distribution (small superscript (c) as distinct from capital subscript C for cutoff):

$$p(B|s\geqslant s_C) = \frac{\Pr(s\geqslant s_C|B)}{\Pr(s\geqslant s_C)}p(B) = \frac{F^{(c)}(s_C|B)}{F^{(c)}(s_C)}p(B)$$

with $F^{(c)}(s|B,T)$ being the fractional tail or complimentary distribution of Bad loans that have scores greater than or equal to s_C. With a similar notation for fraction of booked Bad borrowers who have loan rates greater than or equal a cutoff r_C and scores greater than or equal s_C, one can derive the expected default rate of a randomly selected borrower in the booked sub-portfolio. An obvious generalization to examine the default rate of the booked quadrant of high score/ high loan rate customers yields:

$$p(B|T,r\geqslant r_C,s\geqslant s_C) = \frac{F^{(c)}(r_C,s_C|B,T)}{F^{(c)}(r_C,s_C|T)}p(B|T). \qquad (37)$$

There is no difficulty in this formulation when the prices or loan rates, r, are randomly assigned or where a default score value implies a unique value of r. The probabilistic interpretation is slightly more difficult when r is a decision or control variable. It has been suggested that the Bayes' factor on the rhs of (37) can be decomposed into independent loan rate and score factors. We have found, both theoretically and experimentally, that there is strong dependence between risk and response outcomes because r and several $\mathbf{x}$ characteristics are usually common to both scores. As Simpson's Paradox is at work one must be careful to include conditioning arguments required by the Chain Rule or use evidence-based assumptions to justify the separability of risk score and loan rate factors.

Obviously, Non-Take Inference plays a critical role in assessing the number of Bads among Non-Takes and Accepts so, difficult as it may be, greater understanding and experimentation with this important topic should be encouraged. It is instructive and valuable to revisit the closely related topic of Reject Inference in Classification and Credit Scoring in Hand and Henley (1997) (italics for our substitutions):

> … Typically it (*the risk score*) is the set of people who were classified as good risks by an earlier score-card. …. If the new score-card is based on a superset of the characteristics used in the original score-card then the true classes in the reject (*Not Take*) region are missing, but those in the accept (*Take*) region are not. In this case, the available data can be used to construct an accurate model, without taking into account the rejected (*Not Take*) cases, but only over the 'accept' (*Take*) regions of the space, as defined by the original classifier. Extrapolation over the reject (*Not Take*) region is then needed, (*provided we believe that preferences by borrowers for loan rates from the new scorecard are 'similar' to those with the old*) Improved classification could be produced if information was available in the reject (*Not Take*) region—if some applicants who would normally be rejected were accepted (*Taken*).

As suggested by Hand and Henley what is urgently needed is access to new information (characteristics) among the Take and Non-Take populations. For example, the knowledge as to whether a borrower did or did not have loan offers, other than the one booked, is a valuable first step as would post-mortem 'after-the-fact' surveys of borrowers who turned down one lender at a given rate but booked with a different lender at the same or a possibly different rate. Unlike the classification schemes for Reject Inference where one can run small side-experiments to accept borrowers, ordinarily classified as Rejects, in order to gain behavioural and performance information, this opportunity is not available in the analysis of adverse selections of Non-Takes because a loan rate enticement that encourages more borrowers to respond to offers is, by itself, influencing the 'price' of the loan along with the adverse selection composition within the new preferences.

It is highly unlikely that data substitution or surrogates (see Finlay, 2010, p 233; Thomas, 2009, p 67) from existing credit bureaus or agencies can be useful for the estimation of the required Non-Take inference even though bureau data contain loan and credit line repayments for a single borrower from different lenders. Bureaus receive performance data for Takes but, to our knowledge, do not archive the terms of loans for Non-Takes. Even though there are some cases where the offer rate taken by a borrower can be inferred from the loan type and contractual payments required by the lender, terms and size of loan and other relevant information are largely unrecorded. Because credit bureaus do not record the linkages between repayment histories along with offers turned down by each borrower, relevant data are unavailable for making Non-Take Inference and calculating action-based risk scores such as $s_p(r)$.

An important but much more expensive data acquisition strategy is to audit the terms of offers, responses, behaviour and performance of borrowers who Take as well as do Not Take from each lender. This requires extensive tracking capabilities and possible agreements among competitive lenders. Apparently such data were available to Agarwal *et al* (2010) in their study of Federal Reserve data. Unfortunately, they made no effort to quantify and estimate the expected number of adverse selections, other than providing comparisons of the cumulative distribution functions for risk scores of Takes and Non-Take populations.

Default score revisions for Take and Non-Take subpopulations

We can compare the relative odds of Goods among Takes as suggested by Good (1961) which, from (7), expresses the posterior odds of a Good among Takes as the prior odds times a ratio which depends on differential Take rates:

$$\begin{aligned}\frac{p(G|T,\mathbf{x},r)}{p(B|T,\mathbf{x},r)} &= \frac{p(G|\mathbf{x},r)}{p(B|\mathbf{x},r)}\times\frac{p(T|G,\mathbf{x},r)}{p(T|B,\mathbf{x},r)}\\ &= \frac{p(G|\mathbf{x},r)}{p(B|\mathbf{x},r)}\times\frac{q_G(\mathrm{x},r)}{q_B(\mathrm{x},r)}.\end{aligned}\tag{38}$$

The ratio of interest in (38) is the relative Take rates of Goods and Bads. What the Bayes' factors tell us is that the Good/Bad odds of the Take group is the Good/Bad odds for the entire population multiplied by the price-dependent ratio of Take rates for Goods and Bads. The adjustment is the weight of evidence in favour of Take rates for Goods against Bads at each offer price and default score. By including this weight of evidence in favour of Goods who take the offer, the posterior default score inherits the influence of differential response rates:

$$\begin{aligned}s_p(r|T) &\triangleq \ln\frac{p(G|T,\mathbf{x},r)}{p(B|T,\mathbf{x},r)}\\ &= \ln\left(\frac{p(G|\mathbf{x},r)}{p(B|\mathbf{x},r)}\times\frac{q_G(\mathrm{x},r)}{q_B(\mathrm{x},r)}\right)\\ &= s_p(r)+\ln\frac{q_G(\mathrm{x},r)}{q_B(\mathrm{x},r)}.\end{aligned}\tag{39}$$

If the Take probability for Goods is less than for Bads their ratio is less than one and the second log term is negative which is equivalent to BAS, adverse selection of Bads. When this occurs with a higher price-offer the default score for the higher priced offer is less than the default score for the lower priced offer. When the ratio is greater than one, we have GPS. Thus, the default score conditional on a Take is larger or smaller than the action-based default score for the borrower depending on whether the log of the ratio of differential Take is greater than or less than zero. Similar comments apply to Non-Takes. The importance of the conditional Take rates is apparent when one realizes that the rate of change of the posterior Good/Bad odds with respect to r equals zero iff

$$\frac{\varepsilon_G^{(T)}}{\varepsilon_B^{(T)}} = \frac{q_G(\mathrm{x}, r)}{q_B(\mathrm{x}, r)}. \tag{40}$$

There are some applications where the Take/Not-Take correction factors in (39) are almost entirely due to price effects so that a useful model for the ratio of Good/Bad response rates can be obtained without a requirement for a response score. One can also calculate the posterior score in (39) conditioned on Non-Takes rather than Takes; on subtracting the former from the latter, we find that the prior score cancels and we are left with

$$s_p(r|T) - s_p(r|N) = \ln\frac{q_G(\mathbf{x}, r)}{q_B(\mathbf{x}, r)} - \ln\frac{1 - q_G(\mathbf{x}, r)}{1 - q_B(\mathbf{x}, r)} \tag{41}$$

which may help to explain the gaps in the cumulative score distributions in Figure 1. For the numerical results reported in the top portion of Table 6, the update to the default score is

$$\ln\frac{q_G(\mathbf{x}, r)}{q_B(\mathbf{x}, r)} = \ln\frac{547/3597}{36/272} = 0.14,$$

that is a positive (not a negative!) increase in the risk score which leads to a lower (not a higher) default rate. The difference in the posterior scores is therefore

$$\begin{aligned} &s_p(r|T) - s_p(r|N) \\ &= \ln\frac{q_G(\mathbf{x}, r)}{q_B(\mathbf{x}, r)} - \ln\frac{1 - q_G(\mathbf{x}, r)}{1 - q_B(\mathbf{x}, r)} \\ &= 0.14 + 0.02 = 0.16. \end{aligned}$$

For the bottom half of Table 6 with BAS, the reduction in the posterior scores is $s_p(r|T) - s_p(r|N) = -0.69 - 0.09 = -0.78$ which represents an example of a case where the cumulative distribution of Take scores dominates the Non-Takes, as might be expected.

Estimation of scores and experimental results from validation samples

The results that are reported below include an analysis of a proprietary database of mortgage applications that included prime and sub-prime paper. The analysis and results are associated with fixed-rate 1st mortgages although preliminary results suggest that similar conclusions can be drawn from other loan types. All risk and response scores used in this study were based on data available at the time the offer decision was made; the action-based default/non-default and response/no response scores depend on offer rates as well as the premia of quoted rates over LIBOR-3. The first step is to infer the G/B outcomes on Non-Takes to build action-based risk score within the lender's own database. This should use whatever relevant data are available up to and including the time at which offers are made. The second step is to estimate expected counts under the NAS condition and the final step is to use the deviations between NAS counts and observations to estimate the number of adverse selections.

The sample time frame used for scorecard models built in this study was calendar year 2004. The Take/No Take performance was obtained from application records. The Good/Bad outcomes among Takes were available for a period of two to three years. Booked loans over 120 days past due, foreclosed, or bankrupt were tagged as defaults, that is Bads, while the others were tagged as Goods. The resulting sample contained over 50 000 records, of which 40% were held out for validation. Scorecards predicting Good/Bad and Take/No Take performances were built. A typical default scorecard included:

rate premium	employment status
loan to value ratio	term of loan requested
FICO score	loan type
income	home price appreciation
assets (type and size)	back-end ratio

where front-end ratio is the ratio of existing debt payments to income of the prospective borrower excluding terms of the new loan being considered and back-end ratio is the ratio that includes terms of the new loan, that is borrower's total new debt payments if the loan were to be booked. Predictors for the typical response scorecard included:

rate premium	down payment amount
loan amount	property type
loan type	back-end ratio
% chg in HPI (one qtr)	borrower years @ residence
broker fee amount	borrower years @ job
FICO score	front-end ratio

The 'rate premium' was adjusted for the most recently available LIBOR3 rate prior to a loan's application date; this was an attempt to immunize the models from rate changes taking place over a 12-month window. There was not a material difference in the observation date of the 'rate premium' variable and the observation date of the other predictors. HPI variables attempted to capture possible speculative motivations of the borrowers. HPI values used are the quarterly, state-level HPI values published by the US Federal Housing Finance Agency. As expected, the 'rate premium' variable was the strongest variable in the default scorecard models. Not surprisingly, 'FICO score' and 'loan to value' variables were also strong predictors of risk. The single most important variable in the response scorecards was the loan rate premium, the same characteristic that appeared in the default scorecards. Although score weights were not optimized over these performance measures, the K-S, AUC and Divergence values (see Thomas, 2009) in our development samples were (Table 5).

A significant number of the Takes were Pre-pays, that is early termination of the loan contract with full payment of all outstanding interest and unpaid balances. Thus, there are four rather than three rows (states) that might have been considered in Tables 1 and 2: Non-prepays, Takes who Prepay, Non-Takes and Accepts. Although it might make sense to extend our framework to include a Prepay scorecard, only default and response scorecards were developed. The need for inference of Goods and Bads extends to all four states but in our development of scorecards we replaced the Take group by the Non-Prepay state and combined the Prepays with the Non-Takes to keep the format consistent with the theoretical models in this paper. The availability of loan rates, Takes, Non-Takes and Prepay data has made it possible to obtain Non-Take inferences for Goods and Bads in the combined groups.

Many prepays in the analyses exhibited various degrees of Delinquency with Default being extremely rare. In our initial analysis, we treated prepays as Takes and allowed their re-payment history to drive their status as Good or Default. However, given the relatively low payment exposure when the prepay event occurred (clustered around the first six months) and the longer performance window over which the Good/Default outcome was observed (over two years exposure), we determined that the best way to address prepays was to treat them as a censored performance variable and infer Good/Bad (Default) performance. Prepays remain an interesting and important subgroup of outcomes and worthy of separate study.

Table 5 Performance measures for development sample scorecards

Score	*K-S (%)*	*AUC*	*DIV*
Baseline 'good'/'bad' score	33	0.72	0.81
Action-based 'good'/'bad' score	36	0.74	0.92
Baseline 'take'/'no take' score	18	0.62	0.45
Action-based 'take'/'no take' score	23	0.66	0.57

Table 6 examines records in a validation sample within a risk segment with baseline default scores in the intervals (2.0–3.0) and two different price tiers with loan rate premiums (3–5%) and (5–7%). This risk segment corresponds to an average default rate of approximately 7%. The numbers in parentheses in the Non-Prepay/Bad cells are our estimate of the number of adverse selections for BAS (positive) or the change in the Bad count for GPS (negative numbers).

Table 6 indicates that with the higher priced loan rate premium the number of Bad Adverse selections is slightly less than half of the Bad counts among Takes; there are possibly a small number of Good adverse selections in the lower priced tier but further analysis would have to be made to decide whether this number is large enough to be significant. It should be noted that different analyses using either a superior scoring technology and/or more informative data should be able to obtain improved estimates of adverse selections and risk/response elasticities.

Summary

The authors have defined (i) conditions for no adverse selection and have provided (ii) a simple theoretical model to compare counts in the presence of adverse selections with counts expected when there is no adverse selection. We also define (iii) price-risk elasticity and derive conservation equations that reveal the exchanges between risk and response preferences as loan rates are changed. By comparing theoretical predictions with observed response and risk outcomes, we offer (iv) limited experimental results for different price-risk segments where default risk and response scores are used to quantify borrower preferences and the magnitude of adverse selections. We identify

Table 6 Observed/Inferred cell counts for Goods/Bads, Takes/Not-Takes for one risk, two rate segments

	Observed/Inferred #Goods	*Observed/Inferred #Bads (Adverse selects)*	*Totals*
3–5% Rate Premium			
Non-Prepays	547	36 (−5.3)	583
Non-Takes and Prepays	3050	235.5	3286
Accepts	3597	271.5	3869
5–7% Rate Premium			
Non-Prepays	493	76 (31.8)	569
Non-Takes and Prepays	5382	450.5	5832
Accepts	5875	526.5	6401

(v) the critical role of Non-Take inference with the hope that these results provide (vi) further incentives for statistical testing and experimentation and a better understanding of the role and magnitude of adverse selection in marketing and risk assessment within the credit loan and mortgage industry.

Acknowledgements—The authors are indebted to the referees, who made us think carefully about the sources and content of data that influence adverse selection. We are extremely grateful for the careful reading, comments and suggestions by Brian Bloechle, particularly his emphasis on probabilities and Bayes' factors conditioned on data rather than scores.

References

Agarwal S, Chomsisengphet S and Liu C (2010). The importance of adverse selection in the credit card market: evidence from randomized trials of credit card solicitations. *Journal of Money, Credit and Banking* **42**(4): 743–754.

Akerlof GA (1970). The market for 'lemons': Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* **34**(3): 488–500.

Altman D, Cutler DM and Zeckhauser RJ (1998). Adverse selection and adverse retention. *American Economic Review* **88**(2): 122–126.

Ausubel LM (1999). *Adverse selection in the credit card market*. Working Paper, University of Maryland, Department of Economics.

Cressy R and Toivanen O (2001). Is there adverse selection in the credit market? *Venture Capital* **3**(3): 215–238.

Edelberg W (2004). *Testing for adverse selection and moral hazard in consumer loan markets*. Finance and Economics Discussion Series 2004:1-09, Division of Research and Statistics, Federal Reserve Board: Washington, D.C.

Fahner G (2012). Estimating causal effects of credit decisions. *International Journal of Forecasting* **28**(1): 248–260.

Finlay S (2010). Credit Scoring. In C Cherry (ed.) *Response Modeling and Insurance Rating. A Practical Guide to Forecasting Consumer Behaviour*. Palgrave Macmillan: Basingstoke, pp 232–235.

Good IJ (1961). Weight of evidence, causality and false-alarm probabilities. In: Cherry C (ed). *Information Theory: Fourth London Symposium*. Butterworth: London, pp. 125–136.

Hand DJ and Henley WE (1997). Statistical classification methods in consumer credit scoring. *JRSS. Series A* **160**(3): 523–541.

Lewis EM (1992). *An Introduction to Credit Scoring*. 2nd edn, Athena Press: San Rafael, CA, republished by Fair Isaac & Co.

Phillips R and Raffard R (2009). Theory and empirical evidence for price-driven adverse selection in consumer lending. Columbia University Business School and Nomis Solutions, Edinburgh Credit Scoring and Credit Control Conference XI. Revised as Working Paper on 20 July 2010.

Rasmussen E (1994). *Games and Information*. 2nd edn, Blackwell Publishers: Oxford, pp. 223–245.

Thomas LC (2009). *Consumer Credit Models*. Oxford University Press: New York.

Received March 2011;
accepted January 2012 after one revision