

Simultaneous versus Sequential Information Processing*

Ross M. Hoffman

John H. Kagel

Dan Levin

Department of Economics

Ohio State University

1/7/2011~~12/30/2010~~

Abstract

We study how experimental subjects update exogenously induced prior information when receiving two information signals simultaneously versus receiving the same signals sequentially. Under the simultaneous treatment, subjects' updates of likelihoods vary a lot, yet remarkably the average is within the second decimal place of the correct Bayesian estimate. In contrast, under sequential information subjects' mean estimates of the two treatments (good news preceding bad news or vice versa) are significantly different from each other and on opposite sides of the mean estimate from the simultaneous treatment. The results suggest that both sequencing and the order in which information arrives matters.

*We thank Hal Arkes for helpful discussions and the comments of an anonymous referee. This research was partially supported by grants from the National Science Foundation and private contributions to the Behavioral Decision Making fund at the Ohio State University.

In many real life situations important updating information arrives over time and/or simultaneously. In theory, in individual choice tasks starting from a given set of prior beliefs, a decision maker's final set of updated beliefs should not be influenced by whether they receive the additional information simultaneously or sequentially. We know that without special training decision makers are not good Bayesians on a number of dimensions (see Camerer, 1995 for a review of the literature). However, there has been little exploration of systematic biases in information updating as a consequence of receiving information simultaneously or sequentially, or whether good news precedes bad news or vice versa. This short note reports on a study in which we begin to explore this issue experimentally.

Our results show that for the simple case studied, sequencing matters. When subjects receive information simultaneously, although their individual updated estimates vary significantly from the correct Bayesian estimate, remarkably enough the average of these estimates is the same as the correct Bayesian posterior to the third decimal place.¹ However, mean estimates of the two sequential treatments (good or bad news first) are significantly different from each other and on opposite sides of the mean of the simultaneous treatment, even when subjects received the same information signals as in the simultaneous treatment. These results suggest that both sequencing and the order in which information is received matters.

There have been few previous experiments that can be directly compared to ours. Beach and Wise (1969) studied differences between sequential and final estimates regarding which of two different decks of cards was being used. For each deck, a letter (A-F) was written on each of 100 cards, with the relative frequencies of each letter varying between decks. One of the decks was then randomly selected and sequences of cards were drawn to provide information about which randomly selected deck of cards was being used. The estimates were more conservative for extreme probabilities relative to the correct Bayesian probability, but there were no major differences between the final estimates of subjects making sequential estimates (as the cards were drawn) versus those making a single, final estimate. Peterson and DuCharme (1967), using dice and colored chips found that sequential estimates tended to be slow to follow the correct Bayesian values, with estimates less resistant to moving up (towards 100%) than down. Finally,

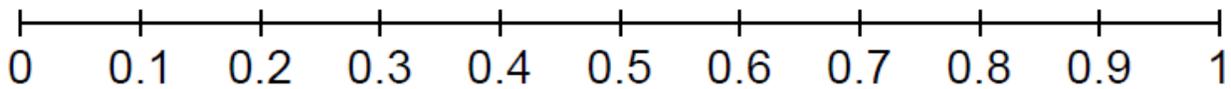
¹ This result is consistent with the notion of the “wisdom of the crowd” (Surowiecki, 2004).

in an experiment not involving estimating probabilities, Bruner and Potter (1964) studied subjects' ability to recognize pictures of common objects that gradually came into focus. When the initial image was less focused, subjects' recognition of the object was much worse than at clearer starting points when subjects were provided with the same final image. To the extent that the brain updates visual images and probabilities through the same, or an analogous, mechanism these results are directly related to ours.²

Experimental Design:

Subjects' baseline condition was purposely framed in terms of a potential genetic disorder that requires a laboratory test to diagnose. This was done in order to examine the effect of good and bad news on probability estimates, rather than abstract signals such as dice or cards. Subjects were told that the prior probabilities for having the disorder were "...a 30% chance you have the condition, and a 70% chance that you do not." Further, the laboratory test for having the disorder, "...like most medical tests", is not completely accurate, so that if they had the disorder, the "...test comes out positive 80% of the time, and comes out negative (a "false negative") 20% of the time." Further, if they did not have the condition, "...the test comes out negative (as it ought to) 90% of the time and comes out positive (a "false positive") 10% of the time."

In all treatments subjects got two test results, from different laboratories, with the need for a second set of results framed in terms of getting a second opinion. For the simultaneous treatment, subjects were told that they received both results at the same time, with the tests coming back positive from one lab ("...indicating that you have the condition") and negative from the other ("...indicating that you do NOT have the condition"). Subjects then were asked to determine "...considering the results of **both** tests" the likelihood of actually having the condition. Subjects had a number line like the one shown below to place an X indicating their answer, along with a box below it to fill in a numeric value:



Numeric Answer:

² None of these experiments appear to have used financially motivated subjects.

Subjects in the sequential treatment had exactly the same information set as in the simultaneous treatment, only the lab results were separated by “a two week period.” In one treatment, which we will refer to as the “Good News First” (GNF) treatment, subjects were first told that the lab results came “...back negative, indicating that you do NOT have the condition”, while the second set of results came “...back positive, indicating that you do have the condition.” Further, subjects were required to fill out a number line and to provide a numeric answer like the one shown above following their first lab results, followed by a second number line and numeric answer following the second lab results, for which they were requested to consider “...the results of **both** tests”. The other sequential treatment, which we will refer to as the “Bad News First” (BNF) treatment, was exactly the same as the first, but with the lab results provided in the opposite order.

To incentivize respondents to consider their answers carefully they received a cash payment of \$10 for the correct answer (relative to that of “a medical expert”), with \$1 subtracted for every 5% their answer deviated from the correct value. For the sequential treatments these incentives were in effect for both answers so that subjects could have earned a total of \$20 in these treatments. Subjects were recruited from undergraduate economics classes at Ohio State University. These were largely introductory or lower level, classes so that most students would have had little, if any, familiarity with Bayes’ rule. With instructors’ permission the last 15-20 minutes of class time was used to distribute the questionnaires, grade them and pay subjects. Everyone in the class was invited to stay and participate with the understanding that we would randomly select 10% of them to pay. We had a total of 167 subjects from 5 classes roughly divided equally between the three treatments.

Experimental Results:

Table 1 provides descriptive statistics for the likelihood of having the condition based on the results of both lab tests for all three treatments. In the Simultaneous treatment, the mean estimate of the likelihood of having the disease is 0.433, almost the same as the correct value of 0.432. While the number of subjects above and below the correct value is almost evenly distributed, the estimates are quite variable ranging from 0.08 to 0.90. However, 66% of the subjects had estimates greater than the prior value of 0.30 indicating that they correctly identified that the combined impact of the test results was to increase the likelihood of having the disease.

Table 1
Likelihood of Having the Condition after Both Lab Reports

Treatment (number of subjects)	Mean Estimate (standard error of the mean)	Number of Subjects Above/Below Mean	Number of Subjects Above/Below Correct Value
Simultaneous (53)	0.433 (0.028)	25/28	25/28
Bad News First (56)	0.367 (0.034)	24/32	19/37
Good News First (58)	0.475 (0.034)	28/30	32/26

Correct value: 0.432.

In the BNF treatment the mean estimate of the likelihood of having the disease after both lab tests is 0.367 which is below the estimate for the simultaneous treatment (as well as the correct Bayesian value). In contrast, in the GNF case the mean estimate after both lab results is 0.475, which is above the estimate for the simultaneous treatment. Neither the BNF nor the GNF estimate is significantly different from the Simultaneous estimate at conventional significance levels ($|t| = 0.95$ and 1.49 , $p > 0.10$, two-tailed test in both cases). But the BNF estimate is significantly lower than the GNF estimate ($|t| = 2.23$, $p < 0.05$, two-tailed test).³ The results indicate that sequential arrival of information yields biased results compared to simultaneous arrival of information. Subjects were more responsive to the latest piece of information they got than to the combined information.

One can think of a number of possible mechanisms generating these biases, with some sort of recency effect (Murdock, 1962) coming most immediately to mind.⁴ However, a closer look at the data suggests that the source of the bias is the incomplete adjustment to the initial test result received. For the BNF treatment, the mean updated estimate of the likelihood of having

³ Mann-Whitney non-parametric tests for mean differences yield similar results except that BNF is significantly different from simultaneous at the 5% level ($Z=2.13$, two-tailed test)

⁴ Also see Hogarth and Einhorn (1992) who note that recency effects are characteristic of sequential arrival of information compared to primacy effects for simultaneous evaluation of the same information, when relatively simple bits of information are involved, as would be the case here. In contrast, our subjects get it right, on average, with simultaneous arrival of information, compared to sequence effects when the information arrives sequentially. There are important differences between our task, and our evaluating outcomes relative to Bayes' rule, compared to the experiments summarized in Hogarth and Einhorn which, more often than not, involve qualitative evaluations of a variety of stimuli such as trait adjectives and behavior statements.

the disease following the first, negative (no disease) lab result is 0.605 (0.034) (with the standard error of the mean in parentheses). This move is in the right direction relative to the prior probability, but not nearly enough, so that it is well below the correct value of 0.774. For the GNF treatment the mean estimated likelihood of having the disease changes very little to 0.296 (0.030) which is not significantly different than the prior probability, and well above the correct Bayesian value of 0.087.⁵

For the sequential subjects we also compared their final estimate to an updated Bayesian estimate, but one that starts from their individual estimates following the first set of lab results. Taking the average of these individual estimates, for the BNF treatment this updated Bayesian estimate is 0.330 (0.029) as opposed to the final predicted estimate of 0.367, indicating under adjustment once again. Repeating this exercise for the GNF treatment, this updated Bayesian estimate is 0.680 (0.026) compared to the final predicted estimate of 0.475, once again indicating under adjustment. Thus, the bias in the sequential data is driven by systematic under adjustment of subjects' estimates compared to the correct Bayesian estimate, albeit this underestimation is systematically stronger for the GNF group than the BNF group.⁶ This suggests that the systematic bias in the sequential estimates rests on the fact that subjects are implicitly required to do two sets of their own, imperfect, Bayesian updating. This conservatism in updating probabilities is well established in the psychology literature (see, for example, Edwards, 1966). In contrast to the sequential case, the simultaneous provision of the additional information helps to limit whatever bias there is in updating.

Summary and Conclusions:

We have reported systematic bias in updating beliefs when information arrives sequentially compared to having the same information arrive simultaneously. We trace the source of the bias to incomplete updating of initial beliefs in response to the first piece of information provided relative to the correct Bayesian updated probabilities, as well as to the subsequent information provided. More experiments are needed to explore this phenomenon. But to the extent that our observations generalize, we have identified another adverse effect of faulty Bayesian updating that needs to be taken into account in decision making, which has

⁵ Using the Wilcoxon rank sum sign test the medians in both cases are significantly different from the correct Bayesian values at better than the 1% level.

⁶ The latter more than likely is a sampling effect rather than anything to do with the order in which the lab information was provided.

policy implications regarding the release of information (sequentially or simultaneously) . While it remains an empirical question to what extent our results generalize, it is already known that the “conservatism” in updating underlying our results is well established in the psychology literature. This suggests that similar results are likely to be found in a variety of economic settings; e.g., business hiring decisions which are conditional on sequential arrival of information about the state of the economy or patients evaluating alternative medical medical procedures after a second or third opinion.

References

- Beach, Lee Roy, and J. Wise. Subjective probability revision and subsequent decisions. *Journal of Experimental Psychology* 81, no. 3 (1969): 561-565.
- Bruner, Jerome S., and M. Potter. Interference in visual recognition. *Science* 144, no. 3617 (1964): 424-425.
- Camerer, Colin. Individual Decision Making in Kagel, John H. and Roth, Alvin e. (eds) *The Handbook of Experimental Economics*. Princeton University Press (1995).
- Edwards, W. Conservatism in a simple probability inference task. *Journal of Experimental Psychology* 72 (1966): 346-354.
- Hogarth, Robin M. and Hillel J. Einhorn. Order effects in belief updating: the belief adjustment Model. *Cognitive Psychology* 24 (1992): 1-55.
- Murdock, B. The serial position effect of free recall. *Journal of Experimental Psychology* 64 (1962): 482-488.
- Peterson, Cameron R., and W. DuCharme. A primacy effect in subjective probability revision. *Journal of Experimental Psychology* 73, no. 1 (1967): 61-65.
- Surowiecki, J. *The Wisdom of the Crowd*. Doubleday (2004).