

This PDF is a selection from an out-of-print volume from the National Bureau of Economic Research

Volume Title: Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation

Volume Author/Editor: F. Thomas Juster and Robert P. Shay

Volume Publisher: NBER

Volume ISBN: 0-87014-402-2

Volume URL: <http://www.nber.org/books/just64-2>

Publication Date: 1964

Chapter Title: Introduction and Summary to "Consumer Sensitivity to Finance Rates: An Empirical and Analytical Investigation"

Chapter Author: F. Thomas Juster, Robert P. Shay

Chapter URL: <http://www.nber.org/chapters/c1283>

Chapter pages in book:

CHAPTER 1

Introduction and Summary

SINCE the end of the Second World War it has been increasingly clear that variations in consumer expenditures often play an active rather than a passive role in the generation of business cycles. Variations in consumer demand appear to be relatively, and perhaps absolutely, stronger in the durable goods component of consumer spending, especially automobiles, than in nondurable goods or services. And it is argued by many that changes in the factor variously called "consumer anticipations," or "attitudes," or "expectations" are importantly related to independent variations in expenditures, and that changes in consumer anticipations are capable of being measured and predicted. This monograph is concerned with the use of survey data on consumer anticipations as an aid to prediction of durable goods purchases.

Concentration on the problem of prediction—as opposed to explanation—bypasses a subject of considerable importance. Even if consumer surveys are able to provide an accurate measure of the current state of consumer anticipations, it does not necessarily follow that they will also be able to explain the way in which anticipations are formed. But for making predictions, it is unnecessary to decide whether anticipatory variables are in some sense basic determinants of consumer behavior, or whether they are themselves wholly predictable from purely objective (i.e., historical) factors like income levels, or past rates of income change, and so on. I am, of course, concerned with whether or not surveys of consumer anticipations can improve predictions of purchase behavior *relative to predictions that use only objective variables obtainable at the same date*. Analysis of the empirical relations between durable goods purchases and objective variables like past changes in income or liquid assets is thus important here because it provides a bench mark against which to assess the predictive value of anticipatory variables.

Cross-Section vs. Time Series Analysis

Although the basic objective of the study is improved prediction of changes over time, the data comprise a cross section of households covering a single time period. Cross sections provide substantially greater flexibility than time series in that they contain an almost complete spectrum of possible differences in household circumstances as well as potentially unlimited numbers of observations. This makes it possible to investigate the simultaneous influence of a large number of factors; more

INTRODUCTION AND SUMMARY

important, perhaps, complex interrelations among explanatory variables can be examined.

However, the use of cross-section data to illuminate the behavior of time series presents its own difficulties. An analysis of factors associated with behavior differences among households at a single point in time does not necessarily yield valid inferences about factors responsible for differences in behavior over time. To illustrate: cross-section data will ordinarily show that households with incomes of \$5,000 per year act differently than households with incomes of \$4,000 per year. Can it be inferred that households with an income of \$4,000 in one period and \$5,000 in a subsequent period will show the same difference in behavior? Probably not. In a cross section, behavior differences among households with different incomes are mainly a reflection of long-run adjustment to a particular income position: for many—perhaps the majority—of households, income has been at the same level for some time in the past and is expected to remain there for some time into the future. But the behavior of households whose current income represents a change from the past is bound to reflect recent arrival at a new income status, the fact of income change rather than income stability, and, perhaps, the influence of concurrent changes in the income of neighbors.

It is true that many of these differences, e.g., the first two just noted, can be handled by an appropriate selection of households from a cross section or by inclusion of enough variables in the analysis. But the third kind of difference, which exists because a given change over time never takes place in a vacuum and frequently takes place in conjunction with similar changes in other households, cannot really be handled by any data that currently exist.¹

Another obvious difference between cross-section and time-series analysis lies in the relative importance assigned to certain variables. Imagine a community with constant population, birth rate, death rate, and marriage rate, but with income varying through time. In a cross-section analysis it would probably be found that the most important factors associated with purchases of durable goods are the demographic structure of the household and its current and past income level. At a given income level,

¹ In principle, even this problem could be managed with cross-section data covering a number of consecutive years for identical households. Then one could talk about the effects of an income change for household A during a period when community income was also changing, and compute the effect on A's behavior of a change in A's income, community income unchanged, or of a change in community income, A's income unchanged. Daniel Suits of the University of Michigan has been working in this area, both in the development of statistical techniques and in fitting empirical relationships. I am not aware of any published results.

INTRODUCTION AND SUMMARY

newly married couples with young children typically buy more durables than older couples whose children have married. But in such a community, demographic structure would clearly be of little or no importance to an explanation of time series changes in purchases because it is invariant over time.²

Another reason for caution in using single-time cross sections to illuminate time series problems is the "personality correlation" problem,³ the difficulty encountered in isolating real behavior relationships from cross-section data. Observed differences among families may reflect mainly differences in the characteristics of the families themselves rather than in objective circumstances like financial status. To illustrate: the existence of personal debt must have a tendency to inhibit purchases, other things equal: a family with large debt would surely be less apt to buy than if the *same* family had a smaller debt but all other relevant circumstances were the same. Yet cross-section data from the Survey of Consumer Finances⁴ indicate that relatively high levels of debt in a given period tend to be associated with relatively large purchases during a subsequent period. Does this mean that debt encourages purchases? Certainly not.

The empirical results presumably reflect the tendency of families with large debts to be spending-oriented and of families with low debts to be saving-oriented. By definition, spenders buy more than savers. But the empirical data do not indicate what the effect would be if all households—spenders and savers alike—had an increased level of debt and other factors remained the same. The answer to the latter question cannot be obtained because the empirical tests are dominated by personality differences that are in turn correlated with debt level. An independent test is needed to distinguish "spenders" from "savers," so that the effect of debt on households with the same personality traits can be investigated. For example, one might test for the effect on purchases of debt differences at two points in time for identical households, thus holding personality constant.

² Even in this extreme case demographic structure might be highly relevant for predicting changes over time. It cannot be assumed that the aggregate pattern of spending or saving is independent of the distribution of income change among households. If income increases were generally concentrated in younger households, the effect on purchases would be quite different than if the same aggregate income increases were concentrated in older households. Here, even when the distribution of powerful cross-section variables remains absolutely invariant over time, it would still be necessary to include these variables in an analysis concerned solely with movements over time.

³ See James Tobin, "Consumer Debt and Spending: Some Evidence from Analysis of a Survey," *Consumer Instalment Credit: Conference on Regulation*, Board of Governors of the Federal Reserve System, 1957, Part II, Vol. 1.

⁴ *Ibid.*

INTRODUCTION AND SUMMARY

As already noted, tests of the simultaneous influence of a large number of variables constitute one of the basic advantages of cross-section over times-series data; even more important, perhaps, complex interrelationships among variables can be tested with much greater precision in a cross section of households. In my view, it is quite likely that the relation between consumer behavior and many of the potential explanatory variables is statistically complex rather than simple; i.e., the appropriate relations are likely to involve some form of nonlinearity. For example, income change may well exert more influence on the purchases of the young than the old; debt may inhibit purchases for the old but not the young; optimism may exert a more than proportional influence if it pervades all facets of household expectations than if it is mixed with some pessimism; fears of depression probably affect the behavior of wage earners more than that of salary earners; and so forth.

These nonlinearities and interactions can generally be isolated only from cross-section data. Moreover, they can be integrated into forecasting models only if survey data are regularly available. If it becomes necessary to know, not average income change for the entire population but income change for the population under forty-five years old, not average consumer debt but debt for households in which the chief wage earner is over forty-five, etc., the customary time series materials clearly become inadequate and need to be supplemented by survey data. It is self-evident that the existence of substantively important relations between consumer behavior and consumer anticipations can only be exploited by survey data.

An investigation of differences among households in a cross section is thus essential, in my view, to the construction of an appropriate model for the prediction of time series changes. Some aspects of the relation between cross-section differences and time series movements, in connection with the problems of predicting consumer purchases from surveys of consumer buying intentions, have been examined by Arthur Okun.⁵ Okun shows that the accuracy of purchase predictions over time, based on surveys of buying intentions, depends on three factors: the size of the cross-section difference in purchase rates between households who reported buying intentions (intenders) and those who did not (nonintenders), the time series correlation between the proportion of intenders in the population and the purchase rates of both intenders and nonintenders, and the amount of variation over time in the proportion of intenders in the population.⁶

⁵ See "The Value of Anticipations Data in Forecasting National Product," *The Quality and Economic Significance of Anticipations Data*, Princeton University Press for National Bureau of Economic Research, 1960.

⁶ These relations are discussed at greater length in Chapter 3.

INTRODUCTION AND SUMMARY

Cross-section data must be used to examine the first of these three factors, that is, differences in purchase rates between intenders and nonintenders.⁷

Nature of Sample and Data

Most of the data presented in this paper are drawn from a reinterview of some 24,000 households, which were (or still are) member-subscribers to Consumers Union of U.S., Inc. (CU), the product-testing and rating organization. The background of these surveys is discussed in Appendix C, and is summarized below.

In April 1958 questionnaires were mailed to some 33,000 CU households that had volunteered to participate in a survey designed to investigate consumer spending and saving behavior. Since I was especially interested in analysis of buying intentions, I split the sample into five randomly selected subgroups; each subgroup was sent a question about intentions to buy that varied with respect to planning period and certainty specification. Over 26,000 returns were received; more than 24,000 of these respondents agreed to further interviews with the understanding that reference numbers would be attached to future questionnaires to permit a matching of responses from identical households.

A follow-up survey was mailed to these 24,000 households in October 1958, six months after the date of the original survey; about 20,000 returns were obtained. Additional reinterview surveys were taken in April 1959 and in April 1960. Some 16,000 responses were received in April 1959,

⁷ A finding that no difference exists would be conclusive evidence that buying intentions have no predictive value over time; on the other hand, the opposite finding would not in itself demonstrate that intentions are necessarily useful for predicting purchases. The existence of intender-nonintender differences in purchase rates constitutes a necessary but not a sufficient condition for demonstrating the proposition that intentions have predictive value over time.

This analysis of the relation between cross-section differences and time series prediction may be applied generally. For example, one could examine the effects of optimism or age on purchases by classifying the sample into optimist-pessimist or young-old groups. If a larger fraction of "optimists" buy than "pessimists," other things being equal, it can be shown that a time series measuring changes in population optimism will be positively correlated with purchases, provided that the fraction of the sample expressing optimism is independent of the fraction of optimists or pessimists that buy, and provided that the fraction of optimists varies over time (see Chapter 3). Similarly with age, income, or any other variable.

There would be no disagreement with the formal statement of this conclusion. There is disagreement, however, about the adequacy of the empirical (cross-section) tests that have been used to determine whether or not particular variables have net predictive value in time series. (See the comment by Eva Mueller following Okun's article in *Anticipations Data*; see also George Katona, *The Powerful Consumer*, New York, 1960, pp. 254-256.)

INTRODUCTION AND SUMMARY

and about 8,500 in April 1960.⁸ Responses from the first three surveys—April 1958, October 1958, and April 1959—form the empirical basis for this monograph.⁹

The Consumers Union Panel, as it might now be called in view of the number of reinterviews available, is unique in many respects. In the first place, the sample size is large enough to permit detailed stratification by relevant household characteristics while still retaining relatively large cell sizes. Secondly, the sample comprises families whose educational level and cooperativeness are such that they apparently can give accurate responses to mail questionnaires; it is evident that many respondents consulted records for data they could not reproduce from memory.

These families are obviously not a random selection from the population as a whole, nor even from the population of CU subscribers.¹⁰ But because of their social and economic status their behavior may well throw considerable light on the causes of fluctuations in durable goods purchases for the population as a whole. Broadly speaking, these fairly young, middle- or upper-middle-income families with college backgrounds own homes and comparatively large stocks of durables. At present, households like these account for a relatively small but steadily increasing fraction of total United States families; and such families are responsible for a relatively heavy share of the fluctuations in durable goods buying that are the main focus of research on consumer discretionary expenditures.¹¹

Further, the factors that are associated with fluctuations in the purchases of this group are possibly easier to disentangle, since these families have a greater-than-average ability to articulate their preferences, attitudes, and expectations. In sum, my view is that a study of behavior patterns in

⁸ These two surveys were conducted under the auspices of the Columbia University Expectational Economics Center, directed by Albert G. Hart. The last questionnaire contained substantially more quantitative detail with respect to income, assets, debt, and durable goods stocks than previous ones; the comparatively low response rate was apparently due to the time-consuming nature of the questionnaire rather than to the sensitivity of respondents to providing information, judging from the (unsolicited) comments that accompanied many of the returned questionnaires.

⁹ The bulk of the empirical analysis in this monograph is based on the 20,000 October 1958 responses and the matching April 1958 data, although in Chapters 2 and 3 the 16,000 April 1959 responses (and the matching October 1958 and April 1958 data) are used extensively.

¹⁰ For a thorough discussion of the characteristics of the CU sample, see Appendix C, below, and F. Thomas Juster, "The Predictive Value of Consumers Union Spending-Intentions Data," *The Quality and Economic Significance of Anticipations Data*, Princeton for NBER, 1960.

¹¹ The cyclical pattern of expenditures on durables is somewhat more pronounced in the CU sample than in the population as a whole. This can be seen in my *Consumer Expectations, Plans, and Purchases: A Progress Report*, Occasional Paper 70, New York, NBER, 1959, Tables 1 and 2.

INTRODUCTION AND SUMMARY

this sample tells a great deal about the factors that can be used to predict or explain differences in the spending behavior of consumers generally. I would judge that these data have provided definitive answers to some of the unsettled questions concerning the relation between consumer anticipations and purchases and have added materially to the forecasting value of anticipations data. At the very least the CU data have provided a fertile field for the exploration, development, and refinement of hypotheses about these relations.

Summary of Findings

A basic objective of this study is analysis and evaluation of consumer intentions to buy durable goods. The empirical results are presented in Chapter 2. For each of thirteen consumer durable products (ranging from new automobiles to garbage disposal units) the subsequent purchase rates of households who reported buying intentions (intenders) and of those who did not (nonintenders) are summarized. Purchase rates are calculated for both six and twelve months subsequent to the intentions survey, and for seven alternative intender-nonintender classifications based on differences in the type of intentions question asked. The data show that:

1. Intenders always have higher purchase rates than nonintenders: this generalization holds for each of the thirteen products, for both time periods, and for all seven intender-nonintender classifications.

2. For a given product and time period, the purchase rates of both intenders and nonintenders are systematically related to the proportion of intenders in the alternative classifications: the larger the proportion of intenders, the smaller the purchase rates of both intenders and nonintenders.

3. For a given product and time period, the *difference* between intender and nonintender purchase rates is systematically related to the proportion of intenders: the larger the proportion of intenders, the smaller the difference between intender and nonintender purchase rates.

4. Of the alternative intender-nonintender classifications tested, one is consistently superior to the other six in predicting differences in purchases among a cross section of households. This superior classification is experimental, and has never been used by any of the surveys that currently report consumer buying intentions.

The data presented in Chapter 2 suggest the hypothesis that buying intentions are essentially statements about subjective purchase probability: those reporting that they "intend to buy A" are simply saying that their probability of purchasing A is at least as high as the minimum probability

INTRODUCTION AND SUMMARY

implied by the intentions question; those reporting that they "do not intend to buy A," that their purchase probability is less than the minimum implied by the question. Further, the data suggest that the distribution of purchase probabilities in the population is a continuous, rather than a discrete, function: households are located at every probability level between zero and unity, rather than at two or three distinct points.

Given this interpretation of buying intentions, a number of generalizations are developed in Chapter 3. These involve the relations among mean purchase probability in the sample, mean probability among intenders, mean probability among nonintenders, the minimum probability associated with the respective intentions questions, and the proportions of the sample reporting intentions to buy when asked the respective intentions questions. Although none of these variables except the last are directly observable from a survey of buying intentions, data from a follow-up or reinterview survey yield information about the purchase rate in the sample as a whole, among intenders, and among nonintenders; and these constitute estimates of mean purchase probability in the respective groups.

The observed data are generally consistent with the probability hypothesis: intenders always have higher purchase rates than nonintenders; for any given commodity, the difference between intender and nonintender purchase rates is larger the smaller the fraction of intenders; and the variance of intender-nonintender purchase rate differences among the several intentions questions is greater for commodities purchased by relatively few households in the sample. But these are comparatively weak empirical tests, since the hypothesis was developed from certain of these same observed relationships. Other empirical tests, based on indirect implications of the model, are constructed; the observed data are shown to be generally consistent with these implications. In neither case, however, can it be argued that the tests do more than suggest that the model is empirically plausible. I see no way, given the available data, to devise a convincing test of the basic hypothesis; such a test would require a survey specifically designed for the purpose.

Sampling variation aside, the probability hypothesis implies that current surveys of buying intentions are likely to yield poor predictions of purchases on some occasions, although it does not follow that they will necessarily do so. The objective of a survey is to predict the population purchase rate during some future period. The best estimate of this purchase rate is presumably mean (*ex ante*) purchase probability in the population. But existing surveys yield an estimate of the proportion of intenders, that is, the proportion of the population with purchase probabilities

INTRODUCTION AND SUMMARY

above some unknown cut-off point. It cannot be assumed that, over time, the proportion of intenders is perfectly correlated with mean purchase probability; even if it were, a linear relation cannot be assumed. And forecasts of the population purchase rate based on the proportion of intenders in a survey make precisely these assumptions. One test of the accuracy with which surveys of intentions can predict purchases shows that only about one-third of the time series variance in the latter can be explained by the former. In addition, a series that constitutes a reasonable but necessarily crude proxy for purchase probability among non-intenders is able to explain a substantial part of the residual variance—the variance in purchases not explained by intentions.

Finally, it can be shown empirically that the accuracy of predictions based on dichotomous (yes-no) survey questions can be substantially improved. Aside from sampling errors, such surveys might yield poor predictions for either of two reasons: (1) the survey does not obtain enough information about *ex ante* probability because respondents are asked a dichotomous question; (2) the *ex post* purchases of respondents differ from their *ex ante* purchase probability because of events that were unforeseen or imperfectly foreseen at the survey date. Improvements in survey design can reduce errors in the first category; those in the second category cannot be avoided unless the incidence of unforeseen events can itself be predicted. Using the most extreme set of assumptions, the data suggest that even complete elimination of the influence of unforeseen events still leaves a good deal of residual error in predictions based on dichotomies. It follows that the potential improvement from better survey design is quantitatively important.

Chapter 4 explores some characteristics of the probability distributions for a number of different commodities and for different groups of households. The observed differences in purchase rates (hence, estimated mean probability) among those reporting intentions to buy automobiles, refrigerators, clothes dryers, etc., suggest that the minimum probability associated with any given question about buying intentions is not likely to be the same for all commodities. In contrast, the minimum probability for a given intentions question and commodity appears to be independent of household characteristics such as income, age, etc.; ownership of an item may be a possible exception.

It is also shown that, given the commodity and the intentions question, the distribution of probabilities below the cut-off point (i.e., among non-intenders) is systematically related to household characteristics such as income, age, etc. Generally speaking, the higher the mean probability

INTRODUCTION AND SUMMARY

for any class of households, the higher the mean probability among non-intenders in that class. There is no evidence that this generalization holds for intenders. By implication, therefore, the variation in mean probability among classes of households is a function of two factors: (1) differences among the classes in the proportion of intenders, and (2) differences in mean probability among nonintenders. I regard this as an additional a priori indication that the change over time in the proportion of intenders is unlikely to be a consistently good proxy for change in mean *ex ante* probability. Further, the data suggest that differences in the level of buying intentions among groups of households are seriously inadequate as a predictor of differences in purchase rates among these same groups.

Chapter 5 is concerned with the relation between alternative measures of aggregate buying intentions, aggregate purchases, and a limited number of explanatory variables other than intentions. It is shown that the strong association between aggregate intentions to buy durable goods and aggregate purchases of durables, again based on analysis of differences among households for a given time period, cannot be attributed to the common influence of other variables, such as income, life-cycle status, or intentions to purchase houses. The variance in durable goods purchases explained by intentions is far in excess of the maximum that can be explained by income or life-cycle status, and the intentions-purchases relation is about as close among households that did not report house-buying intentions as among those that did. In short, buying intentions for durable goods are not a proxy for other variables that we have been able to measure.

In this chapter also, I examine alternative statistical procedures for combining responses to "standard" and "contingent" buying intentions questions, i.e., questions with, respectively, a relatively high and a relatively low probability cut-off. For the most part, the questions are such that standard intenders are likely to be more homogeneous with respect to actual purchase probability than are standard nonintenders. As a consequence, purchase rates among standard nonintenders are correlated with variables such as income, age, or education, while purchase rates of standard intenders do not appear to be related in this way. By the same token, it appears that purchase rates are more closely associated with contingent buying intentions among standard nonintenders than among standard intenders.

Chapter 6 is given over to a detailed investigation of consumer attitudes (or expectations) and their interrelations with buying intentions and purchases of durables. The first part of this chapter examines some of the problems involved in the construction of an index of consumer attitudes.

INTRODUCTION AND SUMMARY

The data suggest that the relation between purchases and attitudes may be nonlinear; households with very optimistic (very pessimistic) responses to a set of questions appear to purchase relatively more (less) than would be consistent with the assumption of linearity. Further, households with a moderately optimistic set of attitudes do not show a significantly different purchase rate than those with a moderately pessimistic set. In short, optimistic households apparently purchase more than pessimistic ones, other things being equal, entirely because of the behavior of the relatively small number of households with extreme views rather than because of a consistent relation between the degree of optimism and the purchase rate.

The last part of Chapter 6 is concerned with the relation among unexpected developments, buying intentions, and purchases. The investigation fails to turn up any convincing evidence that the difference between actual and expected change in family income has a systematic influence on the relation between intentions and purchases. There may, however, be a nonlinear relation among these variables. Households were classified into three groups: those that experienced unexpectedly favorable developments with respect to changes in both family income and in general business conditions, designated (*F*); those that experienced unexpectedly unfavorable developments with respect to both, designated (*U*); and those whose experience was mixed, designated (*N*). Within the *F* group, the association between contingent (probable-possible) buying intentions and purchases tends to be *stronger* than in the other groups; within the *U* group, in contrast, the association between standard (definite) buying intentions and purchases tends to be *weaker* than in the other groups. All the results in Chapter 6 are best regarded as suggestive of possible relations, since some of the available evidence does not support the findings summarized above.

The last chapter summarizes the results of a multivariate regression analysis of durable goods purchases on fifteen independent variables. Seven of these—designated as initial-data variables—are responses to questions about income, assets, expectations, etc., that were included on the buying intentions survey. Five are responses to questions included on the follow-up survey; these are designated as intervening variables, since they measure the degree to which households are affected by events that were wholly or partly unforeseen. The remaining three variables are responses to questions about buying intentions. In addition, results are presented from a regression of so-called standard buying intentions (one of the three intentions variables) on initial-data and intervening variables. Regressions are computed for nine separate groups of house-

INTRODUCTION AND SUMMARY

holds, classified by life-cycle status and characteristics of the buying-intentions variable.

The hypotheses tested are mainly implications of the proposition that buying intentions represent probability judgments on the part of respondents. In general terms, the hypotheses—and their analytical bases—are as follows:

1. Among alternative buying-intentions variables, that constructed from an intentions question with a relatively high cut-off probability should have a relatively large net regression coefficient. According to the probability model, the regression coefficient of intentions measures the difference in mean purchase probability between intenders and nonintenders. In the Consumers Union survey, the higher the cut-off probability associated with a given intentions question, the larger the difference in mean probability between intenders and nonintenders; hence, the higher the probability cut-off, the larger should be the regression coefficient of intentions.

2. Since the association between purchase probability and actual purchases is weakened by the impact of unforeseen events, initial-data variables ought to be more closely related to buying intentions—which are a proxy for purchase probability—than to purchases. Conversely, intervening variables ought to be more closely related to purchases than to intentions, since they represent events that were not perfectly foreseen at the survey date.

3. In a stepwise multiple regression of purchases on initial-data variables, intervening variables, and buying intentions, variables in the first two categories ought to be associated with purchases when buying intentions are ignored. Holding intentions constant, however, the influence of initial-data variables on purchases should be reduced or eliminated because these variables are redundant to purchase probability or to its proxy, buying intentions. On the other hand, intentions should have little or no influence on the net association between purchases and intervening variables, since the latter have no systematic relation to intentions except insofar as they are foreseen.

4. Among comparable variables that partly reflect initial-data considerations and partly reflect intervening events, the variable with the larger initial-data component should have the stronger net relation to buying intentions; and that with the larger intervening-event component, the stronger net relation to purchases.

5. Anticipatory variables should be more closely associated with both purchases and buying intentions for families with relatively young heads;

INTRODUCTION AND SUMMARY

for variables that reflect current financial position, in contrast, the association with purchases or intentions should be stronger for families with relatively older heads. This constitutes a test of the hypothesis that consumption decisions are more closely associated with wealth, defined to include the value of discounted future income, than with current income.

The data are generally consistent with all of these propositions, especially the first three: the regression coefficients for the buying-intentions variables are almost completely in accord with the predicted rankings based on probability cut-offs; the partial correlations of both objective and anticipatory initial-data variables are consistently higher with buying intentions than with purchases, while for intervening variables, the partial correlations are consistently higher with purchases; and initial-data variables explain (net of other variables) about twice as much of the variation in purchases before intentions are held constant as after, while intervening variables explain approximately the same amount before as after.

Conclusions

The findings of this study suggest that surveys of consumer anticipations should be able to make a significant contribution to the prediction of durable goods purchases. However, there is a good deal of evidence that existing surveys do not fully exploit these possibilities and that the potential improvement is of considerable magnitude. The most important finding is that consumer buying intentions essentially reflect judgments by respondents about their probability of purchasing a particular commodity. It follows as a matter of course that surveys should attempt to estimate mean purchase probability in the population, not the proportion with sufficiently high probabilities to report that they "intend to buy." It may be, of course, that such a survey is impractical because of cost considerations or because sufficiently precise responses cannot be obtained. Since there is no a priori way to find out, I conclude that considerable experimentation with survey design is called for.

Secondly, the evidence strongly suggests that certain kinds of anticipatory variables—those that implicitly (or explicitly) require respondents to make judgments about the combined influence of the basic causal factors—are much more likely to be useful predictors of household behavior than are the specific anticipations or facts that form the basis for such judgments. Buying intentions are clearly in this category, since they presumably reflect the combined influence of current and prospective financial and other factors on purchase probability. The other survey variables most closely related to purchases—"Is this a good or bad time for you to

INTRODUCTION AND SUMMARY

buy durables?" and "Which items in your current durables inventory are in need of replacement?"—also reflect judgments about combined or joint influence. None of the straightforward variables that measure current financial position or expectations about change proved to be of much use in predicting differences among households. I conclude that households are, in many cases, better equipped to make complicated estimates of joint influence than are statisticians working with computers: each household is free to assign the appropriate weight to each of the relevant variables, or to assign a weight of zero if the variable is irrelevant to a particular decision; in contrast, computers (implicitly) use the same set of weights for all households in the sample.

Finally, the results seem to me highly relevant to anticipations surveys other than those dealing with the behavior of consumers. Broadly speaking, I have argued that consumer surveys of "plans" or "intentions" require respondents to make rather complicated judgments about purchase probabilities; the data suggest that respondents can and do make such judgments. If so, the same logic should apply with even greater force to surveys of business anticipations regarding sales, inventories, and so on, since business anticipations must surely have a more rational structure, on the average, than the anticipations of consumer units.

Surveys of business firms, especially those concerned with intentions to invest in plant and equipment, have yielded relatively accurate predictions for the most part. However, business surveys have been comparatively unsuccessful in some important respects, e.g., in producing an accurate forecast of how business investment plans are revised in the light of unforeseen changes in sales or other operating variables. It is not unreasonable to suppose that the extent of these revisions is partly a function of the probability attached to the original sales forecast, or that the probability attached to "investment intentions" is neither unity nor invariant over time. These and related possibilities are worthy of systematic exploration.