ARTICLES

The Use of Response Propensity Modeling (RPM) for Allocating Differential Survey Recruitment Strategies: Purpose, Rationale, and Implementation

Paul J Lavrakas¹, Michael Jackson², Cameron McPhee²

¹ Independent Consultant, ² AIR

Keywords: survey costs, response rates, nonresponse

<https://doi.org/10.29115/SP-2018-0023>

Survey Practice

Vol. 11, Issue 2, 2018

Response Propensity Modeling (RPM) is an empirical process that identifies a multivariate statistical model to predict the likelihood (propensity) that a given element in an initial sample will cooperate with a forthcoming survey request. This predicted probability ranges from 0 to 1 and reflects an element's unique combination of characteristics that are expected to affect the relative likelihood of obtaining a response from the element. RPM is consistent with Leverage-Salience Theory in that it is used to tailor different recruitment strategies to different sampled elements, rather than using a one-size-fits-all (OSFA) approach whereby all sampled elements receive the same recruitment strategies. The appeal of using an RPM approach to allocating differential recruitment strategies is that, in theory, it should perform better than an OSFA approach, in terms of gaining a higher overall response rate, gaining a more representative unweighted final sample, and being more cost-effective. As described here, RPM fits into a Tailored Design Method of trying to most efficaciously utilize total survey costs to reduce total survey error. We explain what we mean by RPM, compare it to past work on allocating differential recruitment strategies within an initial sample, and explain the three-stage process used to identify, implement/test, and refine the response propensity model.

1. An Overview of Response Propensity Modeling for Survey Recruitment

It is well known that the years have taken their toll on survey response rates and that traditional means of gaining cooperation from sampled respondents has become much harder and costlier since the 1980s. Thus, new means for recruiting respondents must be devised and tested. In this article, we present one of those new approaches that we believe merits extensive attention, *Response Propensity Modeling* (RPM). This paper describes what RPM is and how to implement it when the context is an ongoing cross-sectional survey. The major goals for the approach we describe are to increase response rates, yield more representative unweighted final samples, and/or reduce total survey costs.

In using the phrase, Response Propensity Modeling, we mean an empirical process that identifies a multivariate statistical model to predict the likelihood (propensity) that a given element in an initial sample will cooperate with a forthcoming survey request. This predicted probability (an "RP score") ranges from 0 to 1 and reflects an element's unique combination of characteristics that are expected to affect the relative likelihood of obtaining a response from

the element.¹ The statistical model used for prediction is developed in Stage 1 of the RPM process and then is used to allocate tailored recruitment strategies (e.g., advance contacts, types and amounts of incentives, number of contacts) to sampled elements in Stage 2 of the process. Stage 2 could be (1) an experiment to test the efficacy of the model and its operationalization, or it could be (2) the deployment of differential strategies tailored for different elements in the part of or the entire initial sample of a forthcoming survey. An optional Stage 3 involves continued refinement of the model.

2. Rational for Response Propensity Modeling and the Tailored Allocation of Survey Recruitment Strategies

What we call RPM can be viewed as fitting within the broader realm of “Responsive Design” or “Adaptive Design” (cf. Groves and Heeringa 2006; Tourangeau et al. 2017; Wagner 2008). In our view, RPM provides a convenient and easily justifiable means of helping researchers determine *how* differential recruitment strategies should be allocated. That is, once it is determined that a design should be “adaptive” in the sense that different cases should get different treatments, RPM provides an objective way to determine (on the basis of patterns observed in at least one prior similar survey) which elements should get which recruitment treatments.

What RPM is meant to accomplish also can be viewed as fitting under the umbrella of the “Tailored Design Method” (TDM; cf. Dillman, Smyth, and Christian 2014), but it is not something that is commonly associated with what heretofore has been used in TDM approaches. That is, we are unaware of anyone previously articulating the approach described here — other than some work at Nielsen (e.g., Lavrakas, Burks, and Bennett 2004)— about why the RPM approach is an ideal method for tailoring recruitment methods to known characteristics of sampled elements.

Traditionally, initial recruitment strategies (e.g., survey incentives) have been distributed with a “one size fits all” (Luiten and Schouten 2013) approach, whereby all sampled elements are treated the same. An exception to this has been the approach used during the past several decades by Nielsen in its various TV audience surveys/panels, whereby certain demographic subgroups with traditionally low rates of cooperation have been given higher valued incentives than those groups with traditionally high rates of cooperation (cf. Trussell et al. 2006).

However, during the past decade, Lavrakas (e.g. 2009, 2011) has opined that there is no theoretical rationale to support the notion that giving all respondents (or all sampled persons within a certain demographic subgroup)

¹ In this paper, we focus on the application of RPM to predict the future behavior of a sampled element in an upcoming survey by using data that are known *about* that element prior to the recruitment of that element. RPM can be applied in other instances that need not rely solely on information known prior to data collection or used to help with initial recruitment (cf. Tourangeau et al. 2017).

the same initial “recruitment package” will achieve the greatest overall gain in response rates that is possible within a given survey budget. Furthermore, the OSFA approach when applied to a survey’s entire initially designated sample does nothing to address the issues of differential nonresponse and nonresponse bias. In some cases, the OSFA approach may even exacerbate differential nonresponse, e.g., when those most likely to respond are differentially stimulated to do so at even higher rates when they receive the OSFA recruitment protocol(s) than are those least likely to respond.

In addition, the OSFA approach to deploying survey recruitment protocols is contrary to Leverage-Salience Theory (Groves, Singer, and Corning 2000), which posits that for each sampled person in a given survey there is a unique mixture of factors that motivate, and other factors that inhibit, a person to/from participating in that survey. That theory directly implies that different persons will need different recruitment protocols. For example, for some sampled persons, incentives (no matter their value) will play little or no role in their decision to participate in a given survey, whereas for others the value of an incentive will be an important factor, if not the key determining factor, in whether they participate. Furthermore, not all members of a given demographic cohort (e.g. African Americans or Spanish-dominant Hispanics) will be impacted equally by a given recruitment protocol.

In addition to the implications of Leverage-Salience Theory, total survey costs should be more cost-effectively spent if the pool of money available for recruitment is differentially allocated so that more costly strategies are given to those inherently least likely to cooperate, and less costly strategies are given to those inherently most likely to cooperate.²

Thus, a logical approach for deciding how to more cost-effectively allocate a finite total amount of funding for recruitment across a given sample is through the use of RPM, rather than an OSFA approach.

3. Previous Research on Differential Tailored Incentives and Other Recruitment Protocols

This section provides a review of some past research that is relevant to RPM, although it is not meant to be a comprehensive literature review; for additional relevant literature, see Tourangeau et al. (2017). In the section, we focus mostly on the allocation of incentives, because that is the type of recruitment strategy that has been most often reported about in the literature. But RPM can, and likely should, be considered for use in differentially allocating any type of recruitment strategy, including advance contacts, recruitment in more than one language, the number of follow-up contacts, the use of informational

² This assumes that a primary goal of recruitment is to gain an unweighted final sample that is the most representative one that is possible of the survey’s target population.

brochures, varying the mode of recruitment, leaving voice mail messages in telephone recruitment or leave-behind packets in in-person recruitment, use of refusal conversion protocols, etc.

As mentioned, as we define it, RPM fits under the rubric of using a “tailored design.” However, what Dillman, Smyth, and Christian (2014) describe as the TDM includes no mention of using an a priori statistical approach to tailoring the allocation of differential strategies in ways that lead to different sampled elements receiving certain recruitment strategies but not others. Rather, our RPM tailoring of differential recruitment strategies builds upon what now appear to be simplistic approaches that Nielsen and other U.S. media audience measurement companies (e.g. Arbitron and Simmons) traditionally have used to allocate survey incentives.

Starting in the 1980s, Nielsen decided to use different values of incentives for recruiting different demographic cohorts. It did this because Nielsen had found that there was considerable variation among cohorts (e.g., Blacks vs. Whites; young adults versus older adults) in the rates at which they returned a completed Nielsen TV diary. Nielsen used a dual-mode process for their diary surveys whereby household demographic information was first gathered in an RDD (random-digit dialing) survey and then used to decide the amount of incentive to enclose with a subsequent mailed diary. In doing this, Nielsen was using a differential incentive approach that was “tailored” at the level of demographic cohorts; i.e., everyone in a particular cohort received the same incentive. However, this essentially was an OSFA approach within a given cohort.

As suggested by Lavrakas, Burks, and Bennett (2004) and Burks, Lavrakas, and Bennett (2005), the use of an RPM approach to more finely tailor recruitment strategies, such as incentives, at the level of the individual household/person theoretically should achieve greater gains in response rates for a fixed amount of recruitment funding in a given survey. This approach should also have a more favorable impact on reducing differential nonresponse, which may in turn reduce nonresponse bias. The research reported by Lavrakas, Burks, and Bennett (2004) and Burks, Lavrakas, and Bennett (2005) focused on whether a useful multivariate statistical model could be identified to predict the response propensity of households in a prior Nielsen diary survey. Their RPM analyses used auxiliary data (e.g., local census characteristics appended to the prior survey’s initial sample), as well as paradata and data gathered about each household during the telephone mode of the prior survey to build a dataset to test and develop their RP model. Their logistic regression model was found to accurately predict ($p < .001$) whether a household that agreed to complete the diary in the telephone interview actually returned a completed diary in the mail stage of the Nielsen survey.

More recently, Link and Burks (2013) reported on a process for allocating recruitment protocols differentially that shared somewhat similar goals to what we describe as RPM, but the process by which they tested differential incentives protocols varied considerably from our RPM approach. These Nielsen researchers used a small number of local area (block group) census variables to identify addresses that were likely to reach households populated by relatively low-responding demographic cohorts, such as Blacks and Hispanics. In their study, Link and Burks used “sample-frame indicators” to develop two different ways to identify these sampled elements. They then targeted higher incentives to these sampled elements.

In contrast to the “tailored” approach used by Link and Burks, an advantage to the RP approach that we have articulated is that an RP score assigned using a multivariate model is a more efficient and easily interpretable way to “aggregate” the impact of several different characteristics that are known about each element in the initially designed sample. For example, while Black households may be less likely, as a group, to respond to a survey request than are Whites as a group, a Black household whose other characteristics are predictive of higher response propensity (e.g. higher income, highly educated, etc.) may have a relatively high overall probability of responding. Therefore, an RP score would better reflect the full combination of factors that determine how likely the household is to respond and therefore which recruitment strategies are most prudent for that type of household.

Apart from the work at Nielsen, Luiten and Schouten (2013) reported a successful approach that appears in principle to be similar to the RP approach that we have developed, but they did not provide details about how they devised their model nor about how their model was applied to allocate the differential recruitment strategies (e.g., recruitment modes, number of contacts, timing of contacts) that they tested in their experiment. Furthermore, these scholars deployed their approach as part of an ongoing panel study in a European country where there are population registries that provide frame data at the level of individual residents. Thus, Luiten and Schouten had government registry data that were measured at the individual level — data that are not readily available in the United States. In addition, in using a panel for their research, Luiten and Schouten (2013) had myriad prior data gathered from each panelist upon which to model a response propensity for their experiment.

Of note, Tourangeau et al. (2017) report on the use of response propensity approaches within the context of surveys using adaptive and response designs but also do not provide details about how the modeling was carried out and applied. The details of how we propose that RPM should be operationalized follow in this manuscript.

4. A Two-stage Approach to Testing the Allocation of Recruitment Protocols with RPM

In Stage 1 of RPM research, statistical investigations are carried out to determine if a parsimonious set of predictor variables³ can be identified that will reliably predict ($p < .05$) a person's/household's likelihood to cooperate in a *forthcoming* survey.

If such a set of variables can be identified, then in Stage 2 of the process, the RP model from Stage 1 is applied, so as to tailor differential recruitment strategies (Luiten and Schouten 2013) to the various sampled respondents/households.

4.1. RPM Stage 1. Data from a previously completed survey ----- where researchers know whether the sampled units ended as a completion or not⁴ — are used to identify the most parsimonious and effective RP model for predicting the binary outcome, response/nonresponse, in that survey.

Stage 1 begins by assembling the initial sample in that prior survey, and any auxiliary data that were appended to each sampled element. If this is an address-based sample, there may be many variables that were matched to each sampled address. If this is a telephone sample, there also will be variables that may be matched to all the sampled numbers, but the number of variables likely will be fewer than with an address-frame, and the precision of the data matched to the telephone numbers often will be lower than when matching such data to addresses (cf. Harter et al. 2016). If it is a sample from a panel, there likely will be myriad other data gathered from panelists in the panel's prior surveys that can be appended to the sampled panelists.

Once the already matched data are identified, the researchers then should contact companies that specialize in address- or telephone-based databases to determine other variables that can be matched to the prior sample. These data cover a wide range of demographic and behavioral characteristics aggregated at the local area (e.g., age, ethnicity, race, residence type, employment, income, education, etc.) and some data at the individual level. Of great value here is a relatively new variable provided by the Census Bureau named the “Low Response Score,” which is inversely proportional to the rate at which local area residents completed the questionnaire for the previous Census (see Erdman and Bates 2017). These companies also can append additional information about an address or telephone number from other sources, although there often are many sampled elements with missing values on these variables.

3 We believe that a relatively parsimonious set of variables is preferred because (1) an unnecessarily complex model may be more likely to “overfit” the dataset (e.g., to uncover apparent “relationships” in the dataset used to identify the RP model that are really just statistical noise and therefore may not exist when the model is applied to a new dataset) and (2) for surveys in which the model coefficients used to assign recruitment strategies like incentives need to be documented (e.g. federal surveys requiring very detailed documentation), a simpler model is more interpretable and thus more transparent as to the characteristics that made a person/household more or less likely to receive a particular recruitment protocol such as a higher incentive.

4 It is possible, and it may be preferable, to predict different types of nonrespondents (e.g., noncontacts vs. refusals) rather than simply predict one catch-all cohort of nonrespondents. But in this explanation of RPM, we explain it in its simplest form, using a binary dependent variable.

However, and fortuitously for RPM researchers, it often is the case that “missingness” on these variables is positively correlated with survey nonresponse.

Once the “best” database is assembled for the prior survey, then the Stage 1 analyses generally use some form of multivariate procedures to predict the binary dependent variable, response/nonresponse (Tourangeau et al. 2017). One viable method, which is likely familiar to survey practitioners, is logistic regression. However, nonparametric methods that automatically identify the independent variables most strongly predictive of response, such as CART (classification and regression tree analyses), CHAID (chi-square automatic interaction detector), and random forests, also can be used, either to directly assign RP scores or to select predictors for inclusion in a logistic regression analysis. These methods also have the advantage of identifying 2nd and 3rd order interactions.

In predicting this binary outcome, the model that is identified can be used to generate an RP score for each case, ranging from 0.00 to 1.00. This score represents the predicted likelihood that an individual element complied with the survey request. Because RPM is best done with large initial-sample datasets, statistically significant results are easily achieved. Therefore, it is the size of the effects, both individually and collectively, that matter more. Since there is not much empirical data on RPM, it is premature to suggest with confidence how “accurate” the RPM must be before there may be value in applying it. But a crude rule of thumb from our somewhat limited past experience is that the model should be at least 60% accurate in differentiating respondents from nonrespondents. However, the greater the accuracy that can be achieved in Stage 1, the greater will be the value of the model in conducting Stage 2. The most important consideration is whether the model generates a reasonably accurate *ordering* of cases by response propensity scores, such that if these scores are used to divide the completed prior survey sample into cohorts, the low-RP cohorts show substantially lower response rates than the high-RP cohorts.

Predictive power at Stage 1 will depend on many factors, but mainly on (a) the robustness of the predictor variables available, (b) the precision of the local area data, and (c) the extent to which there are strong predictor variables. Furthermore, interactions and polynomial terms, in addition to main effects, should be investigated to learn if they improve the model’s predictive accuracy.

An approach to try to further refine the RPM during Stage 1 is to build it using only a random subset of the cases (the “training” dataset) from the previously completed survey. Then the model can be applied to the other random subset (the “test” dataset) on which the model was *not* built. This kind of cross-validation may provide further insights about the likely generalizability of the RPM.

4.2. RPM Stage 2. The approach we are describing assumes that the design and topic of the forthcoming survey, to which the RP model from Stage 1 will be applied, is very similar, if not essentially identical, to the prior survey on which Stage 1 work was based.

The RP scores (ranging from 0 to 1) generated in Stage 2 for each element in a new survey sample are used to make decisions about how to tailor recruitment strategies to different sampled cases. For example, if there is \$100K available for noncontingent incentives in the upcoming survey, and were it decided that there are four RP cohorts that will receive differential treatment, and it were decided that only incentives would be varied, the RP scores could be used to determine which cases will get \$10, which will get \$5, which will get \$2 and which will get no (\$0) incentive, if the researchers decided that these four levels of incentives are the ones to use for the new survey.

Stage 2 work begins by assembling data for all the variables in the Stage 1 RP model for each element in the initial sample for the forthcoming/new survey. Thus, values for all the variables that are in the RP model from Stage 1 must be added to each sampled element in the new survey. The RP model then is applied to the new sample, and each unit in that sample receives a 0.00 to 1.00 RP score.

Next, the researchers need to identify the “cut points” that will be used to determine a case’s a priori “recruitment package.” To this end, researchers should examine the ordered distribution of RP scores across the new survey sample. These investigations can be based on various statistical procedures, but also should include visual inspection of plots of these scores. The goal of examining the distribution of RP scores is to identify meaningful groupings of cases that may be suggestive of an “optimal” allocation of recruitment strategies.

For example, the RP scores could be plotted on a line graph with the x-axis being the range of scores from 0.0 to 1.0 and the y-axis the frequency of each score within the dataset. This plot may or may not be of value to the researchers, depending on the homogeneity or heterogeneity of the sampled cases. If there are multiple cases that share the same RP score, then this plot likely will prove more valuable in planning recruitment strategies than if the distribution of RPM scores is so heterogeneous that very few cases share the same score. In the latter case, if the variability of the number of cases that are assigned the same score is essentially chaotic, then this plot is not likely to be helpful.

Another approach is to form a histogram plot based on percentile groupings of the RP scores. For example, 50 such bars could be plotted with each bar containing 2% of the cases in the sample. The height of the bar in this histogram would be the average RP score for each of the 50 clusters; or, alternatively or in addition, the response rate within that cluster for the

completed survey could be used. Visual inspection of the histogram distribution would be carried out to help make decisions as to (1) how many meaningful groupings of bars are present in the plot and (2) where are the “cut-points” in RP scores that differentiate each grouping from the adjacent grouping(s).

Based on these analyses and discussions, the researchers would choose the final set of cut points for determining each cluster’s recruitment strategies (e.g., sending an advance letter, incentive types and amounts, number of contact attempts, and/or use of refusal conversion). Once these analyses and discussions are completed, the researchers would use the conclusions they have reached to plan the differential allocation of the recruitment strategies in the new survey or plan an experiment in the new survey testing the RPM set of strategies against the control strategies.

5. RPM Stage 3: Evaluation and Optional Refinement of the RPM and Its Application

Once the new survey or experiment has been conducted, the researcher will have data to evaluate how well the RP model and its application worked. Such data also should be used to investigate ways to improve the efficacy of the RP model and its application. These analyses would seek to identify characteristics that reliably differentiate the respondents and nonrespondents in the new survey. Doing so may help the researchers better understand shortcomings in their application of the RP model in devising the specific recruitment treatment(s) used to motivate cooperation.

There is no guarantee that the RP model and/or its application will work as well as the researchers had hoped. It is possible that the model simply was not precise enough in differentiating the sampled cases in the new survey to allow the researchers to tailor their recruitment strategies to achieve their goals. In such a case, the researchers would need to go “back to the drawing board” and try to build a more effective RP model, or possibly conclude that an effective model simply cannot be built with the auxiliary data available to them. In turn, there may be situations in which the researchers have identified an effective RP model that differentiates cases well according to their actual response propensity, but they failed to apply the model in ways that achieve their goals. That is, the researcher had a good model (one that did not fail them) but their application of the model was what fell short. For example, the researcher used what proved to be ineffective recruitment protocols for some or all of the RP cohorts to raise the overall response rate of the combined cases receiving the RPM tailored recruitment strategies. Or the average cost of recruiting cases under the RPM approaches proved to be too high compared to the average cost of recruitment of the control cases. This type of failure is not unlike the important distinction that Weiss (1972) made in her seminal work on evaluation research in noting that an evaluation may fail either due to a theory failure or an implementation failure.

6. Summary

The appeal of using an RPM approach to allocating differential recruitment strategies is that, in theory, it should perform better than an OSFA approach, in terms of gaining a higher overall response rate, gaining a more representative unweighted final sample, and being more cost-effective. The RPM process as described here assumes that it is applied in an ongoing cross-sectional or panel survey. There are many such surveys, and thus, this RPM approach potentially has wide application. As described here, RPM fits into a TDM of trying to most efficaciously utilize total survey costs to reduce total survey error.

REFERENCES

- Burks, A.T., P.J. Lavrakas, and M. Bennett. 2005. "Predicting Sampled Respondents' Likelihood to Cooperate: Stage III Research." In *The 60th Annual Conference of the American Association for Public Opinion Research*. Miami Beach, FL.
- Dillman, D.A., J.D. Smyth, and L.M. Christian. 2014. *Internet, Mail, and Mixed-Mode Surveys: The Tailored Design Method*. 4th ed. Hoboken, NJ: Wiley.
- Erdman, C., and N. Bates. 2017. "The Low Response Score (LRS): A Metric to Locate, Predict, and Manage Hard-to-Survey Populations." *Public Opinion Quarterly* 81 (1): 144–56.
- Groves, R.M., and S. Heeringa. 2006. "Responsive Design for Household Surveys: Tools for Actively Controlling Survey Errors and Costs." *Journal of the Royal Statistical Society: Series A* 169 (3): 439–57.
- Groves, R.M., E. Singer, and A. Corning. 2000. "Leverage Salience Theory of Survey Participation." *Public Opinion Quarterly* 64 (3): 299–308.
- Harter, R., M. Battaglia, T. Buskirk, D. Dillman, N. English, et al. 2016. "Addressed Based Sampling: AAPOR Task Force Report." https://www.aapor.org/AAPOR_Main/media/MainSiteFiles/AAPOR_Report_1_7_16_CLEAN-COPY-FINAL.pdf.
- Jackson, M., P.J. Lavrakas, G. Dirks, C. Brazle, L. Lusskin, and B. Ponce. 2018. "A Response Propensity Modeling (RPM) Experiment on Differential Allocation of Recruitment Strategies in the Simmons National Consumer Survey." In *The 73rd Annual Conference of the American Association for Public Opinion Research*. Denver, CO.
- Lavrakas, P.J. 2009. "Using Incentives in Survey Research." Workshop Presented for District of Columbia Chapter of the American Association for Public Opinion Research, Washington D.C.
- . 2011. "The Use of Incentives in Survey Research." In *The 66th Annual Conference of the American Association for Public Opinion Research*. Phoenix, AZ.
- Lavrakas, P.J., A.T. Burks, and M. Bennett. 2004. "Predicting Respondents' Likelihood to Cooperate: Stage I Research." In *The 59th Annual Conference of the American Association for Public Opinion Research*. Phoenix, AZ.
- Link, M., and A.T. Burks. 2013. "Leveraging Auxiliary Data, Differential Incentives, and Survey Mode to Target Hard-to-Reach Groups in an Address-Based Sample Design." *Public Opinion Quarterly* 77 (3): 696–713.
- Luiten, A., and B. Schouten. 2013. "Tailored Fieldwork Design to Increase Representative Household Survey Response: An Experiment in the Survey of Consumer Satisfaction." *Journal of the Royal Statistical Society, Series A* 176 (Part 1): 169–89.
- Tourangeau, R., J.M. Brick, S. Lohr, and J. Li. 2017. "Adaptive and Responsive Survey Designs: A Review and Assessment." *Journal of the Royal Statistical Society A* 180 (Part 1): 203–23.
- Trussell, N., and P.J. Lavrakas. 2004. "The Influence of Incremental Increases in Token Cash Incentives on Mail Survey Response: Is There an Optimal Amount?" *Public Opinion Quarterly* 68 (3): 349–67.
- Trussell, N., P.J. Lavrakas, G. Daily, T. Yancey, M.A. Bennett, J. Bailey, and J. Lai. 2006. "Improving Response Rates among Targeted Demographic Subgroups Using Large Cash Incentives." In *The 61th Annual Conference of the American Association for Public Opinion Research*. Montréal, Canada.
- Wagner, J. 2008. "Adaptive Design to Reduce Nonresponse Bias." Doctoral dissertation., Ann Arbor, MI: University of Michigan. <https://deepblue.lib.umich.edu/bitstream/handle/2027.42/60831/jameswag\1.pdf?sequence=1>.

Weiss, C.H. 1972. *Program Evaluation: Methods for Assessing Program Effectiveness*. Englewood Cliffs, NJ: Prentice-Hall.