Examining Economic Census Reporting Patterns

Eric B. Fink[1], Elizabeth Hoeffel, Joanna Fane Lineback, and Sarah S. Vetting

In this paper, we examine reporting patterns for the 2012 Economic Census to inform future data collections. The economic census is a quinquennial survey of U.S. business establishments in the eight major business trade areas: manufacturing, construction, mining, retail, services, wholesale, finance-insurance-real estate, and utilities-transportation. Information collected includes employment labor costs and output, assets, expenditures, inventory, and other industry-specific items. For the 2012 Economic Census, data were primarily collected using two-self-administered modes: mail-out/mail-back and electronic. Electronic data collection was through one of two custom-built Census Bureau software products, depending on the size of the business. Available data for this analysis include 2007 and 2012 Economic Census data, both of which include establishment responses to the survey and survey process data. We examine establishment-reporting behavior by sector. We also use multinomial modeling to understand business characteristics that predict which establishments are most likely to switch (or not switch) response mode. Here, we present our initial findings, and hope that they might be used to influence discussions on improving data collection efficiencies, including tailoring survey contact strategies.

Key phrases: Economic Census, administrative data, response rates, establishment surveys, logistic modeling.

1. Introduction

In this paper, we examine 2012 Economic Census (EC) reporting patterns. Here, we introduce the reason for this work and lay out the remainder of the paper.

The 2017 EC can expect major methodological changes. One such change is the switch from a multi-mode to single-mode survey, for both regular data collection and nonresponse follow-up. For the 2012 EC, respondents had the option of reporting by paper or electronically through one of two software products, depending on the size of the business. For the 2017 EC, only electronic reporting will be offered. To mitigate the potential for nonresponse bias, there needs to be a strategy for moving paper respondents to electronic reporting and getting nonrespondents and new businesses to report. Additionally, the 2017 EC electronic collection instrument should make the reporting process as seamless as possible. The issue of moving paper respondents to electronic reporting is the focus here.

In a single-mode setting, the options to reduce nonresponse error are limited, because alternative modes are not offered during nonresponse follow-up. However, there is the ability to carefully plan contact strategies based on what is known about business reporting patterns. In this vain, we have begun to develop a profile of respondents using available information: administrative data, response data, and paradata from the 2007 and 2012 ECs. We examine response rates at the business sector level and on key subgroups. We examine time spent and frequency of actions in the electronic instruments. We use multinomial and logistical modeling to understand business characteristics that predict which establishments are most likely to switch (or not switch) response mode.

For this paper, an establishment and unit are defined as a single physical location where business is conducted or where services or industrial operations are performed, and an enterprise is defined as a

---

business organization consisting of one or more establishments that were specified under common ownership or control. The enterprise and the establishment are the same for single-unit business organizations. Each multi-unit business organization forms one enterprise.

The remainder of this paper is organized as follows. Section 2 discusses the basic EC methodology, including sample design, data collection and estimation. Section 3 discusses the analysis methods used for this research. Section 4 presents the results, and section 5 discusses the results and next steps.

2. Background
This section covers basic 2012 EC methodology (for more information, please see http://www.census.gov/econ/census/), including basic sample design, data collection, and estimation.

The EC is conducted quinquennially, in years ending in 2 and 7. It covers the eight major trade areas of the U.S. economy including manufacturing, construction, mining, retail, services, wholesale, finance-insurance-real estate (FIRE), and utilities-transportation. Collected information consists of employment labor costs and output, assets, expenditures, inventory, and other industry-specific items. The information gathered by the EC is used as input to calculate gross domestic product, among other economic performance measures, is used as a benchmark and an update to the sampling frame for other economic survey programs, and is extensively analyzed by the business community.

2.1 Sample Design
Of the roughly 8.8 million U.S. business establishments in the eight major trades areas, approximately 4.1 million were mailed as part of the 2012 EC. Another 2.4 million were in sample, but not mailed a survey form. As this suggests, there was a sampling component to the 2012 EC. In fact, as estimation was at the sector level, the sample design differed among the trade areas. Table 1 summarizes the sample designs in the major trade areas[2].

Wholesale	Manufacturing and Mining	Retail, Services, and Utilities-Transportation	Finance, Insurance, Real Estate	Construction
•Take all with certainty	•Take all units with certainty. Based on payroll cutoffs, unit receives a long form, short form, or no form. Units that receive no forms are imputed using administrative data.	•Take those with the largest payroll with certainty •Fill those with the smallest tax liabilities with administrative data from the BR •Sort remaining cases on 8-digit NAICS code and payroll and take a systematic sample. Non selected units' key items** filled with administrative data	•In select industries, take all with certainty In all other industries: •Take those with largest payroll with certainty •Fill those with the smallest tax liabilities with administrative data from the BR •Sort the remaining cases on 8-digit NAICS and payroll and take a systematic sample. Non selected units' key items filled with administrative data	• Take all multi-unit businesses with certainty • For single-unit businesses with complete 6-digit NAICS, stratify by state and NAICS, small strata taken with certainty, within remaining strata take a probability-proportional-to-size (payroll) sample. • For single-unit businesses with incomplete 6-digit NAICS, take those with largest payroll with certainty, take a simple random sample of remaining cases

Table 1: 2012 Economic Census high level sampling schemes across the eight major trade areas.

[2] MU's are taken with certainty across all trade areas.
**Key items include receipts, payroll, and number of employees

In general, all multi-unit business operations were always in sample. For the wholesale, manufacturing and mining trades, all single-unit businesses were also in sample, with administrative payroll used in manufacturing and mining to determine which form (long or short), if any, an establishment would receive. For the construction trade, there was no "non-mail" strata. In general, single-unit businesses were selected using probability-proportional-to-size sampling with target coefficient-of-variation constraints. For the remaining trade areas, in general, the single-unit businesses with largest administrative payroll were in sample, those with the smallest tax liabilities were not mailed a survey form and imputed using administrative data, and of the remaining units, some were selected using a systematic sampling procedure.

2.2 Data collection

Although the sample designs varied across trade areas, 2012 EC data collection was one massive operation. In general, the EC reporting unit is the establishment. The main exception is multi-unit business operations that have a predetermined arrangement that one reporting unit will report for multiple establishments. For the 2012 EC, survey data were primarily collected using two self-administered modes: mail-out/mail-back paper and electronic. Electronic collection was through one of two custom-built Census Bureau software products, depending on the size of the business. For multi-establishment businesses and large single-establishment businesses, data were collected through the downloadable software Surveyor. For other single-establishment businesses, data were collected through the Web application Centurion. There were other collection modes, such as telephone, for a small portion of cases. For the 2017 EC, data collection will be electronic.

Each single-establishment business received one of the following three types of survey forms: a classification form, a "long" form, or a "short" form. The classification form collects industry classification information. Establishments for which we have insufficient identifying information receive a classification form. The long and short forms both collect information on key economic items such as employment, payroll, and employment. The long form, however, collects this information at a more detailed level. For example, where the long form asks for annual payroll by job, the short form asks for total annual payroll. Single-establishment businesses that received the long form also receive a supplemental form, which collects information about the ownership and control of the establishment. Some long form recipients also receive a second supplemental form, which collects information on items such as foreign ownership, research and development activity, royalties, and manufacturing activities. Most reporting units received a long form. Only small reporting units in the manufacturing and mining sectors may have received a short form.

Although response to the EC is required by law, historically, the Census Bureau has employed an intensive contact strategy, especially targeted to large, multi-unit establishment businesses, in order to maintain high return rates. For the 2007 and 2012 ECs, contact initially began in the autumn of the reference year, when large multi-establishment companies were first provided with EC forms and notified about electronic reporting procedures. The purpose of this early contact was to facilitate businesses' planning for EC response as the reference year ends. Mail-out to the rest of the units eligible for the EC occurs mid-December of the reference year, with a due date of February 12 the following year. There were as many as four follow-up mailings to non-respondents, the first of which began approximately one week after the due date. Intervals between subsequent mail-outs varied, but were roughly 30 days. The final mailing for the 2012 EC occurred in July 2013. Follow-up plans have not been solidified for the 2017 EC.

2.3 Estimation

In general, the 2012 EC produced estimates of totals at the sector level and on key subgroups. Key estimates across all trade areas include total revenue, payroll, and employment. Estimates are being

released on a flow basis, with final estimates by industry and zip code. For the 2012 EC, no measures of sampling or nonsampling error are being published, except for the construction trade area.

To prepare the data for estimation, they went through complex editing, imputation and weighting steps. These steps differed depending on the trade area. There were both manual and automated editing processes. Manual corrections made up a large percentage of edited data. Automated editing occured in conjunction with imputation. Imputation was the nonresponse adjustment method used for key items. For some trade areas, a factor adjustment was used to adjust for missing product lines data, where the sum of product lines revenue makes up total revenue. For instance, an animal food manufacturer establishment's total revenue might be the sum of cat food, dog food, and specialty feed revenue. In this case, cat food, dog food, and specialty feed are the product lines. It is beyond the scope of this paper to go into detail about these methods, but in general each data item must pass all edits, and, where possible, missing data were imputed using administrative data.

3. Analysis
3.1 Analysis Questions
Much of the research conducted to this point was exploratory in nature, and many of the initial research questions necessitate only descriptive statistics to answer. Our overarching research question is what can existing data tell us about business reporting patterns? To further refine our question into several more manageable parts about business-respondent behavior, we developed the following initial questions:
1. What do paradata reveal about business reporting patterns?
2. What is the cumulative unit response rate?
3. What is the cumulative total quantity response rate?
4. What changes in reporting trends do we notice since the previous survey cycle?
5. What are the characteristics of early versus late responders?
6. Are there strong predictors for switching from paper to electronic reporting? From electronic to paper?

The analysis that follows also examines respondent behavior overall and at the sector level. As the EC aims to produce estimates at the sector level, and sampling methodology and questionnaires differ between sectors, understanding if/how respondent behavior varies between sectors is important. Here, we have limited out analysis to the retail and manufacturing sectors.

3.2 Limitations of the Analysis
There are some limitations to our analysis. Because all establishments in sample report using their parent company ID, from the Surveyor paradata it is not possible to distinguish between multiple establishments downloading or uploading the Surveyor software under a company, and a single individual establishment downloading or uploading the software multiple times. Thus, our Surveyor analysis is restricted to only the initial download/upload event. Additionally, the response rates given below are only approximate, because 2012 EC data are still going through the editing process as of May 31, 2014 and EC systems will only be updated for the 2017 EC to allow us to calculate standard response rates (http://www.census.gov/quality/standards/Quality_Standards.pdf).

4. Results
4.1 Surveyor and Centurion Results
The results presented below stem from a similar analysis conducted on the Annual Survey of Manufactures (ASM). The ASM is a subset of the manufactures sector questionnaire in the EC, and thus the ASM is often used to pretest new processing or data collection procedures. Like the findings in an analysis (Fink and Lineback, 2013) conducted on the ASM, 2012 EC Surveyor results show that most multi-unit respondents upload their data the same day they download the Surveyor software package (see Figure 1). In fact, 87.1% of Surveyor respondents uploaded their data the same day they downloaded the

software. Figure 2 indicates that for those respondents downloading and uploading on the same day, 80.87% of respondents took under three hours, with a mode of three-quarters of an hour.
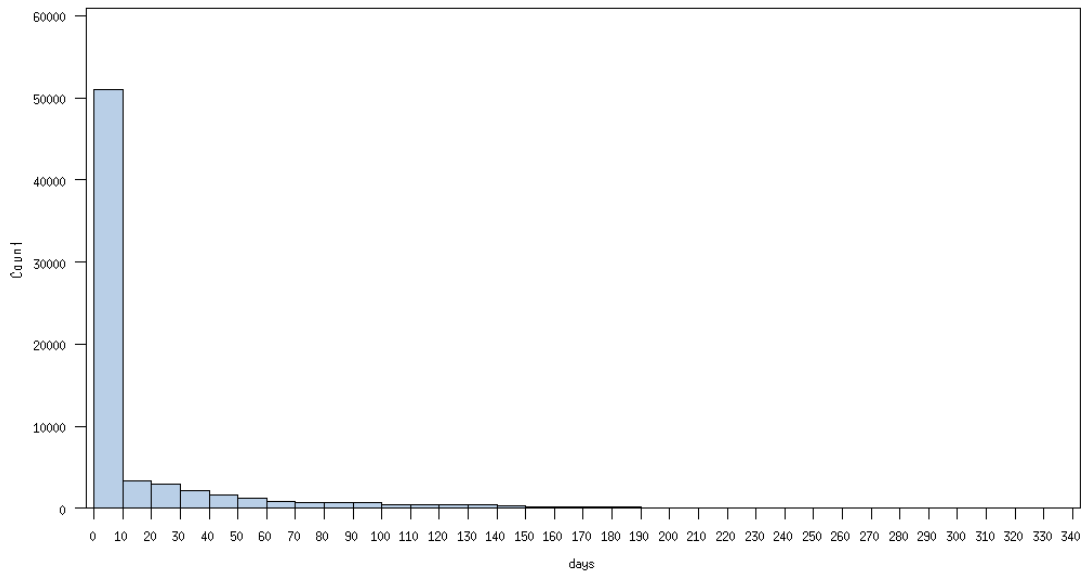

Figure 1: Time elapsed, in 10-day intervals, between the initial download to upload of the 2012 EC Surveyor software.
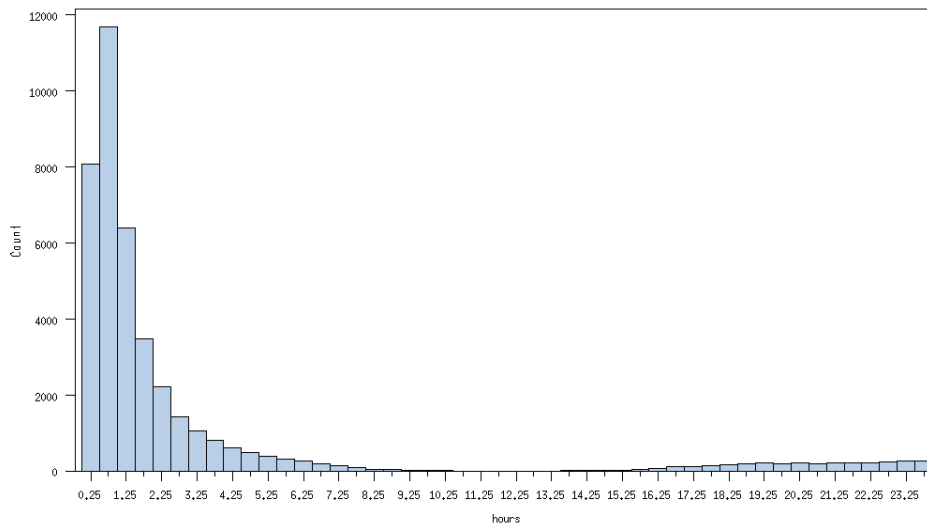

Figure 2: Same day 2012 EC Surveyor responders in 15 minute intervals.

Of particular interest is the patterning among those respondents taking more than 10 days to upload their data (See Figure 3). A bulk of responses do not occur until about 148 days after mail-out. This suggests respondents either: (1) tended to other matters of higher priority in their respective businesses, waiting to complete the survey until it best suited their schedules, (2) did not have the requisite information at-hand to be able to respond within a short time frame after downloading the software, or (3) a combination of (1) and (2).
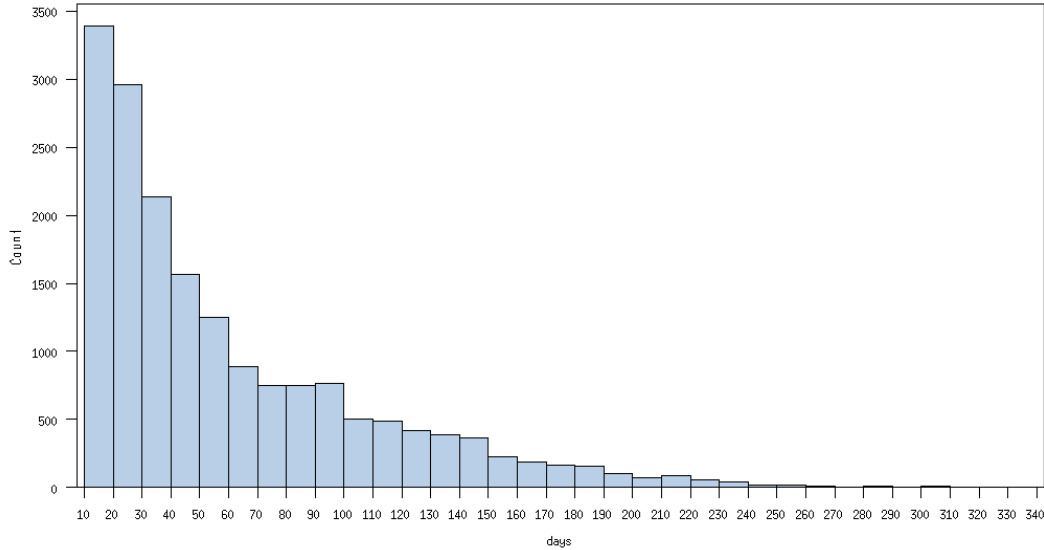
Figure 3: Time elapsed, in 10-day intervals, for respondent taking longer than 10 days between the initial download to upload of the 2012 EC Surveyor software.

Figure 4 shows the single-unit business mode and median response times in Centurion based on responses to the two of the most frequently used EC forms. This is based on units that did not have any failed log-in attempts, logged in only once, and submitted. Combined, the forms are 13 pages long. The mode response time was 10 minutes, and the median response time was 16 minutes. Given that the amount of time taken to complete these forms is relatively short compared to the length of the forms, this suggests that respondents used the paper version of the form to first gather the requested information.
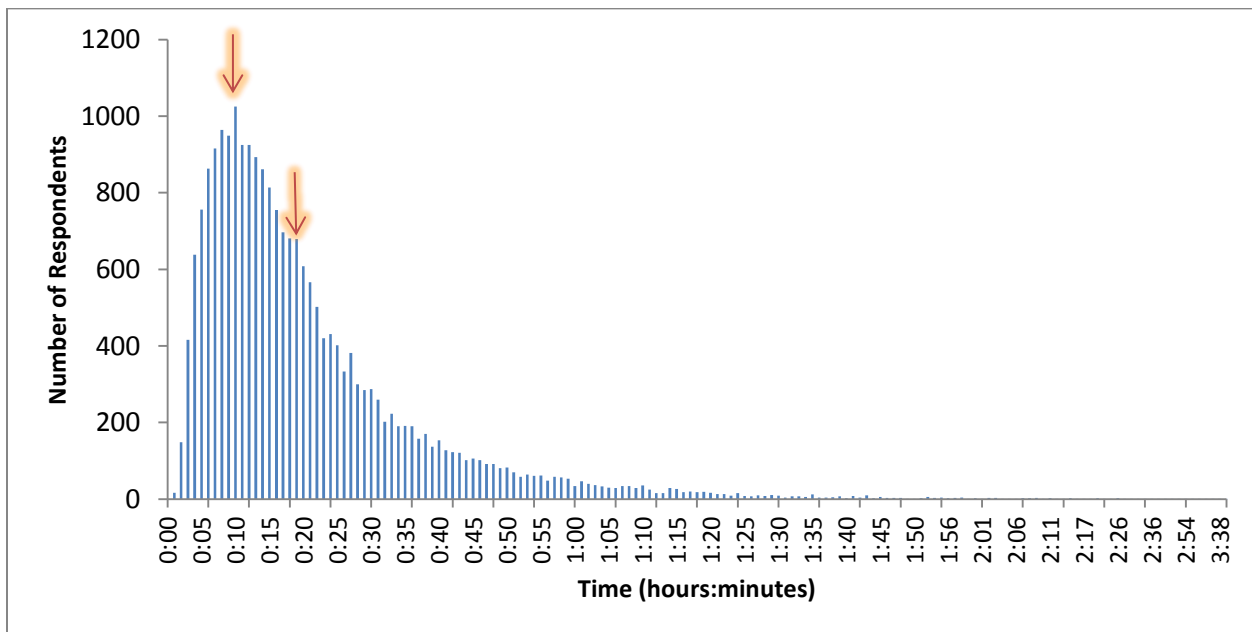

Figure 4: Mode and Median Response Times in 2012 EC Centurion (Forms NC-99510 and AF-72201)

4.2 Overall Response Metrics

The unit response rate (URR) is the ratio of the number of "forms" returned with sufficient information to be deemed a response to those mailed to eligible units, as well as units of unknown eligibility. The URR serves as a measure of data collection performance. The curve in Figure 5 shows a gradual, constant increase in the URR, achieving an overall rate just over 62%. We see an initial spike in the URR between 10 and 60 days after mail-out, with a second increase between 90 and 110 days.
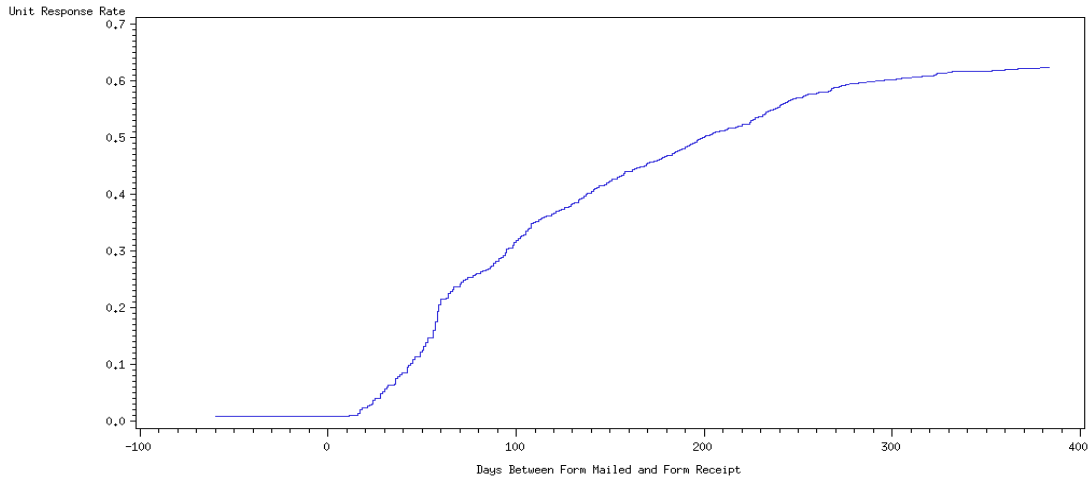

Figure 5: The unit response rate for the 2012 EC from when forms were initially mailed to respondents.

More telling is Figure 6, showing the URR for single and multi-unit cases. Multi-unit cases have a relatively monotonic increase in the URR after 65 days, achieving a rate just over 80%. Single-units experience significant diminishing returns to any nonresponse follow-up efforts, both failing to significantly increase the rate during the same period and achieving a final URR of over 50%.
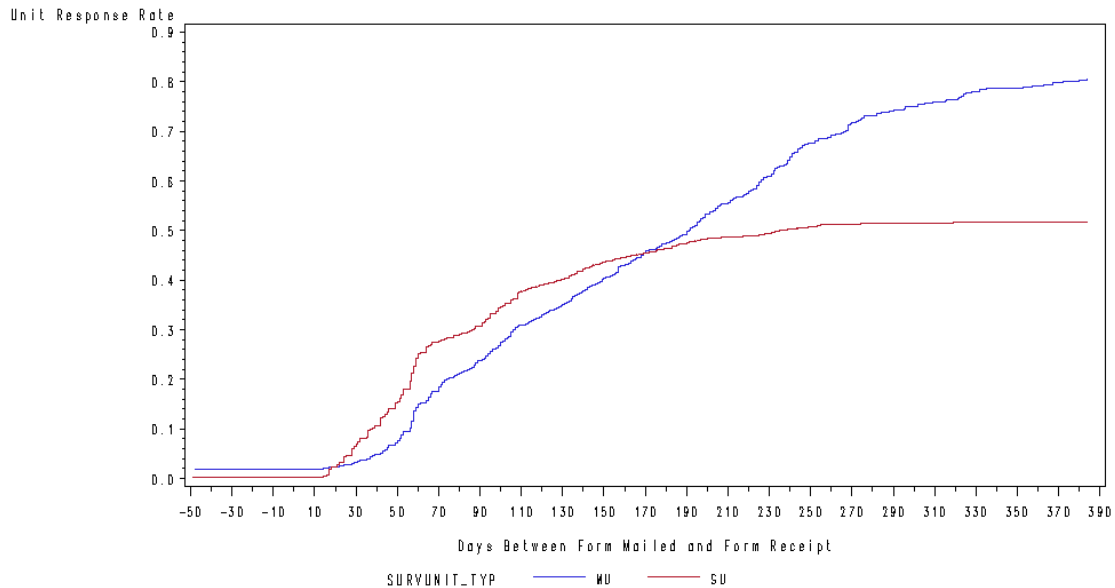

Figure 6: The unit response rate for single and multi-unit firms for the 2012 EC from when forms were initially mailed to respondents.

4.3 Manufacturing Sector Response Metrics

We now move to examine respondent behavior at the sector level. Figure 7 has the same general pattern for the URR as the overall rate. However, Figure 8 which breaks down the URR by establishment type indicates a much better response outcome from single-unit firms in the manufacturing sector than is seen in the overall rates. Multi-unit firms still have a better URR than their single-unit counterparts, but as the EC measures totals, and the larger multi-unit firms account for more of that estimated total, nonresponse follow-up tends to target these large multi-unit firms.
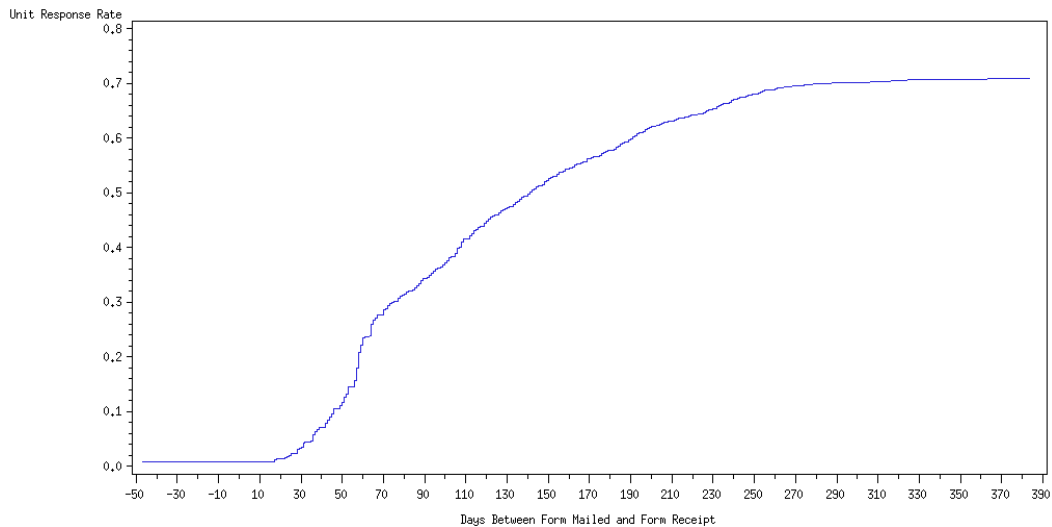

Figure 7: The URR for the manufacturing sector in the 2012 EC from when forms were initially mailed to respondents.
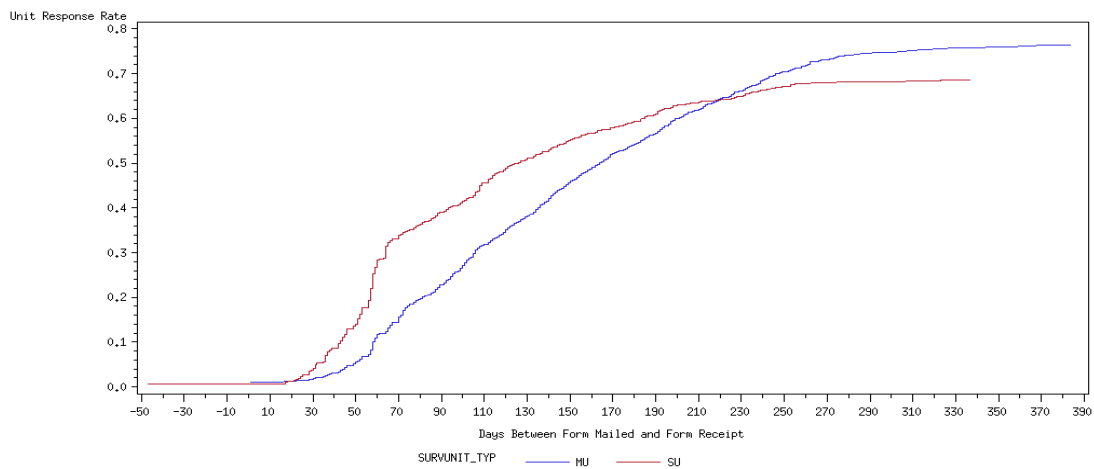

Figure 8: The URR for single and multi-unit firms in the manufacturing sector in the 2012 EC from when forms were initially mailed to respondents.

The quantity response rate (QRR) is an alternative rate used by economic survey programs at the U.S. Census Bureau (U.S. Census Bureau, 2012) to gauge collection performance. The QRR is calculated as follows:

$$QRR = \left[\frac{\sum_{i=1}^{N} w_i r_i t_i}{\sum_{i=1}^{N} w_i t_i}\right] * 100$$

where:

w_i is the design weight of tabulation unit i,
r_i is the indicator variable for reported data for tabulation unit i,
t_i is the data value for tabulation unit i,
N is the total number of eligible tabulation units.

The QRR is therefore the proportion of the estimated total obtained from directly reported data. One serious limitation with the QRR is that the denominator is an estimate, which is not ideal.

Figure 9 shows the QRR calculated for receipts in the manufacturing sector. Between 50 and 270 days after mail-out, we see a relatively monotonic increase in the QRR before leveling off, with a spike around the 135-day mark. Referring to Figure 10, although we receive more reported data from multi-unit than single-unit firms as a proportion of the total receipts, we can attribute at least some of this discrepancy to administrative data substitution in lieu of response for smaller single-unit firms.
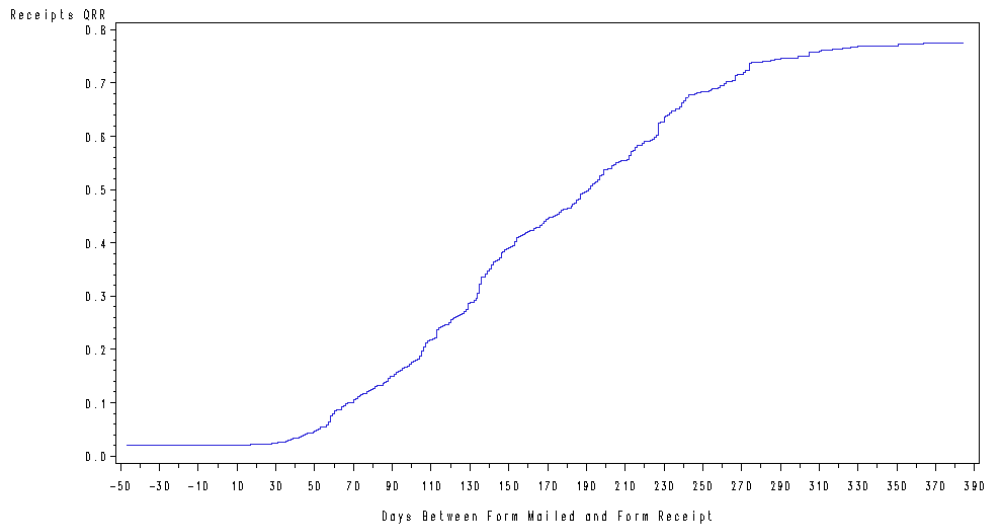

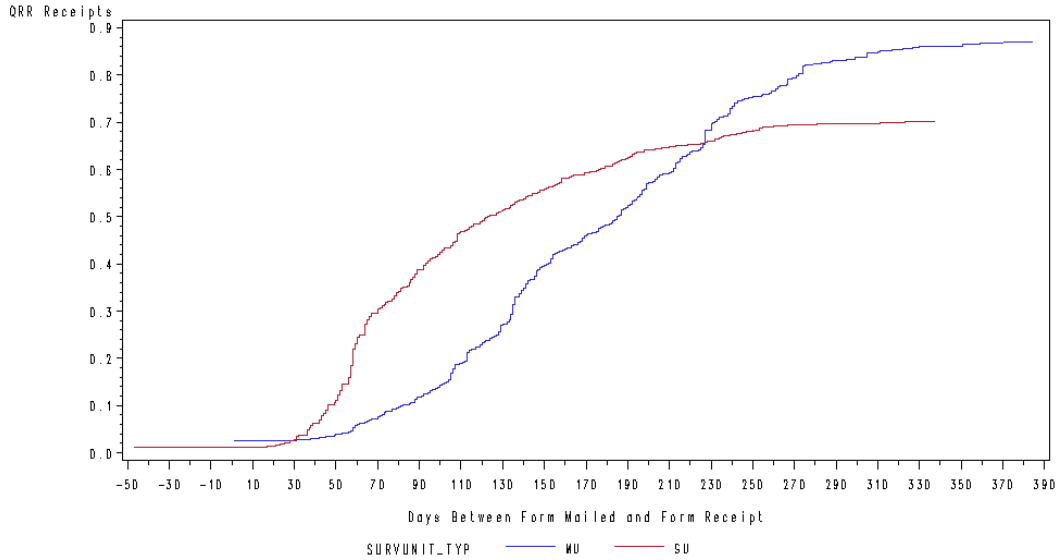**Figure 9:** Overall 20212 EC QRR for the manufacturing sector.

Figure 10: Manufacturing QRR for single and multi-units for 2012 EC.

4.4 Retail Sector Response Metrics

Moving to the retail sector, we see similar patterns in the URR in Figure 11 as we did for the manufacturing sector and the overall EC. However, we again see a much larger disparity in response among single and multi-unit firms (See Figure 12). Similar to the overall EC results, single-units reached a URR of just over 46%.
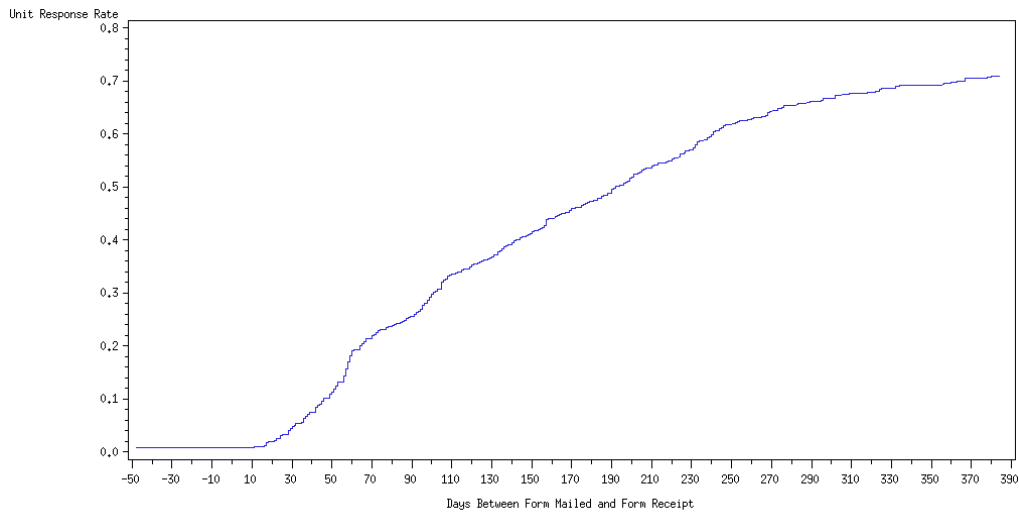

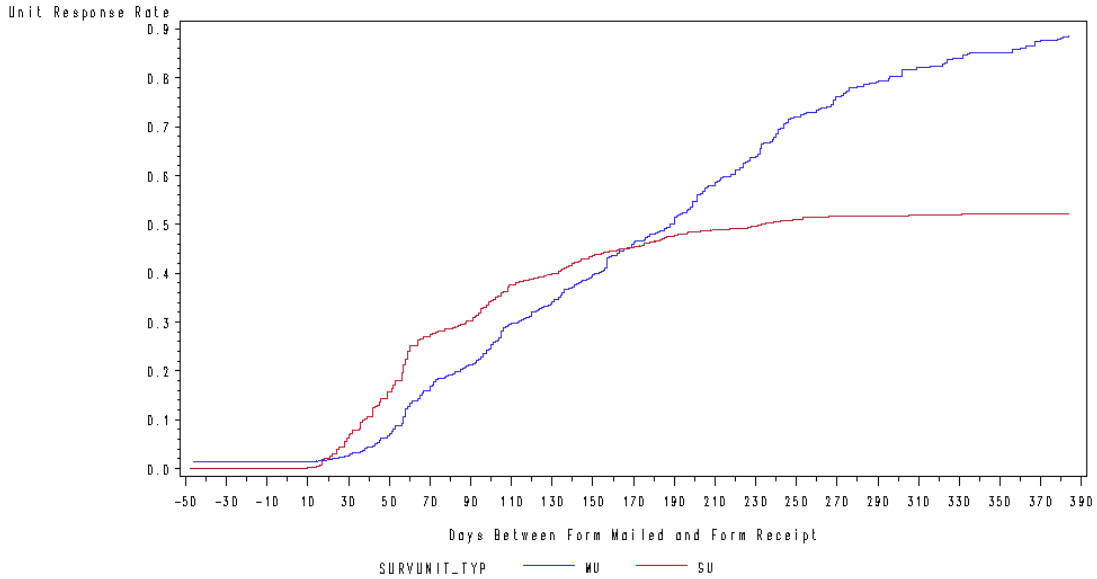**Figure 11:** The URR for the retail sector in the 2012 EC.

Figure 12: The URR for single and multi-unit firms in the retail sector in the 2012 EC.

The response profile for the retail trade sector has a much more linear trend to the QRR than the overall or manufacturing sector, as can be seen in Figure 13. We do notice a couple small departures from this trend, most notably in the 15 to 60-day period, as well as at 269 days, where the QRR jumps 6 percentage points. We again notice large disparities in the QRR among single and multi-unit firms (See Figure 14), which we can attribute at least in part to administrative data substitution in lieu of response for smaller single-unit firms.
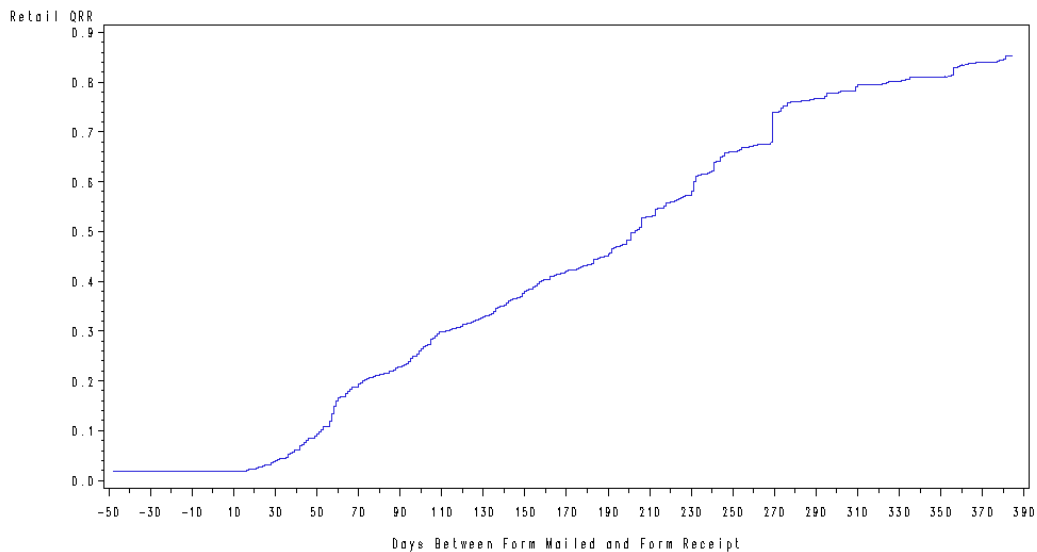

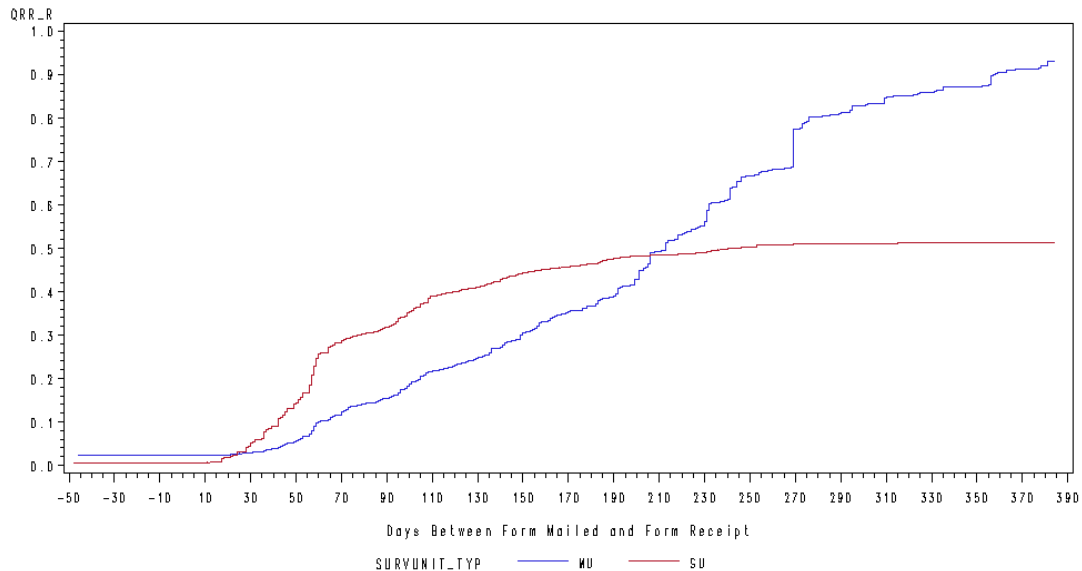**Figure 13:** Overall QRR for the retail sector in the 2012 EC.

Figure 14: Overall QRR for the retail sector for single and multi-units for 2012 EC.

4.5 Mode Change

For ease, we will define "response mode change" as any establishment that changes the mode in which they report data to the Census Bureau from one survey cycle to the next. This change can either be from mailing in a paper form to submitting electronically, or vice versa. We will also define "unit status" as a general term indicating if a unit was a single-unit or multi-unit establishment.

Among the 2,761,938 businesses that reported in both the 2007 and 2012 EC, 1,818,499 (65.84%) businesses responded using a paper form, while 943,446 (34.16%) businesses responded electronically. Additionally, 2,165,382 (78.4%) responded in the same manner during both statistical periods, 86,047 (3.12%) changed from electronic to paper reporting and 510,509 (18.48%) changed from paper to electronic reporting. Previous research (Fink et al., 2013) has been conducted on this topic using the 2011 ASM. The ASM is a subset of the manufactures sector questionnaire in the EC, and thus the ASM is often used to pretest new processing or data collection procedures. From 2010 to 2011, 18.96% of establishments changed the mode in which they reported to the ASM. Among those establishments who changed how they reported the ASM from 2010 to 2011, 69.57% changed from paper submission to electronic, while 30.43% changed from electronic to paper.

A weighted nominal multinomial logistic regression model was constructed for those reporting units that changed response modes, taking into account survey weights by using proc surveylogistic in SAS. The model was run with no mode change as the reference category, with results in Table 2. The variable Diff is the time between form mail-out and return, receipts is the total sales, receipts, and shipments for a company in the 2012 calendar year, and employment is the number of employees a company had in the 2012 calendar year. Other variables were considered, including if the establishment was a single or multi-unit establishment, annual payroll, change in receipts, employment, and annual payroll, and other interaction effects. The model below is the final model after nonsignificant predictors were removed.

Parameter	Change	Estimate	Standard Error	Wald Chi-Square
Intercept	P	-4.2027	0.0113	138501.59
Intercept	E	-0.6281	0.00422	22223.94
Diff	P	0.00341	0.00061	3144.99

Diff	E	-0.00206	0.000027	5867.46
Receipts	P	6.994E-7	1.348E-7	27.05
Receipts	E	7.113E-7	1.337E-7	27.88
Employment	P	0.000810	0.0001	65.73
Employment	E	0.000910	0.000092	97.37
Diff*employment	P	-2.43E-6	4.657E-7	27.10
Diff*employment	E	-3.42E-6	4.013E-7	72.13
Diff*Receipts	P	-1.78E-9	3.93E-10	20.21
Diff*Receipts	E	-1.8E-9	4.09E-10	19.19

Table 2: Results of the nominal multinomial regression model using 2012 EC data. All parameters significant at $p < 0.0001$. The variable change indicates if the respondent changed modes from electronic to paper (P), or from paper to electronic (E), with those respondents who did not change modes as the reference category.

Unfortunately, a confusion matrix that was generated to assess model fit demonstrated the model does a poor job of correctly classifying results. The confusion matrix was created by first generating predicted probabilities in proc surveylogistic using the predprobs=individual option for mailed units in the 2007 and 2012 ECs. Proc freq was then used to create a cross-classification table of the actual and predicted response variables for the data used to fit the model (SAS Institute, 2013). The overall classification rate was at 39%. The model did a wonderful job of correctly classifying observations experiencing no mode change, which was expected as the vast majority of cases were in this class. The model was unable to classify any cases that changed from electronic to paper reporting, and only able to correctly classify 0.18% of cases that changed from paper to electronic reporting.

However, developing a weighted binary logit model yielded a more informative result, reported in Table 3. The variables are defined the same as above, with status indicating if the establishment was a single or multi-unit, and payroll defined as total payroll for the 2012 calendar year. Unlike the multinomial model, the logit model is restricted to only those respondents changing modes from 2007 to 2012, with paper-to-electronic mode change as the reference category. With a concordance rate of 69.3%, the binary logit model was able to correctly classify results more successfully than the nominal multinomial model. The model indicates that as the length of time for response, the number of employees, and the amount spent on payroll all increase, and if in the case of a multi-unit establishment, the log odds of a reporter switching from paper to electronic decrease ever so slightly (very close to zero). The table, however, does indicate it is the presence of the status term that is driving the negative relationship in the four-way interaction, as the three-way interaction without the status variable is positively related to the log odds of switching from paper to electronic. Running the model without the status variable, however, does not lead to a significant result for the three-way interaction. Using the same variables and running the model for the two sectors of interest, the models had a concordance rate of 73.4% for the manufacturing sector, and 68.1% for the retail sector.

Parameter	Estimate	Standard Error	Wald Chi-Square	p-value
Intercept	-3.4126	0.0139	60366.11	0.0001
Status (MU)	0.5456	0.0141	1502.96	0.0001
Diff	0.0028	0.000116	578.61	0.0001
Diff*Status	0.00148	0.000117	158.41	0.0001
Employment	0.00032	0.000099	10.38	0.0013
Employment*Status	-0.00068	0.000129	27.58	0.0001

Diff*Employment*Status	3.4E-6	8.684E-7	15.33	0.0001
Diff*Payroll	-3.86E-8	1.423E-8	7.38	0.0066
Employment*Payroll	-3.15E-9	9.93E-10	10.03	0.0015
Employment*Payroll*Status	3.111E-9	1.109E-9	7.87	0.0050
Diff*Employment*Payroll	1.83E-11	5.58E-12	10.71	0.0011
Diff*Employment*Payroll*Status	-175E-13	5.82E-12	9.02	0.0027

Table 3: Results for the binary logit model using 2012 EC data.

5. Discussion/Future Research

The Surveyor and Centurion analyses both present suggestive evidence that we need to compensate for how people have used paper forms as the Census Bureau is planning for the 2017 Economic Census to be an "all-electronic" mode of data collection. This could be as simple as a link to a pdf worksheet. A more elaborate version could be a system whereby multiple users are able to update fields electronically, and have the capability to make comments and notes to each other in those fields.

The Centurion analysis suggested that respondents used the paper form as a tool to first gather the requested information. This was supported by the Surveyor analysis, which showed that for a bulk of those companies taking longer than 10 days in between downloading and uploading the software took up to 148 days. This alluded to the hypothesis that establishments took the time to record requested information on a paper form, downloaded Surveyor, entered the necessary information, and then uploaded the software. Qualitative research supports the idea that within business infrastructure, the data required to respond to surveys is housed in multiple areas within a business necessitating interaction between multiple individuals to answer survey questions (Willimack and Snijkers, 2013). Finally, as over 65% of respondents submitted using paper (among those units reporting in the 2007 and 2012 EC), using paper forms was the mode of choice. The above analysis is not meant to suggest that respondents will not be able to be "pushed" to responding electronically, but does suggest more research needs to be conducted to figure out how to push respondents in a way that will not sacrifice data quality.

As the response rate analysis shows, however, much research needs to be conducted in order to understand the data quality impacts this decision will have. With overall response rates hovering just over 60%, and the smaller, single-unit establishments evidencing relatively poor contributions to overall response, it is important to make this transition in a way that will not further reduce small-business contribution to our estimates, if not find a way to improve their response profiles. Furthermore, because many other economic surveys benchmark their estimates to the EC, we may be introducing error not only into the EC estimates, but also to the benchmarked estimates of other economic surveys given the disparities in response between single and multi-unit establishments.

While we were unable to develop an effective nominal multinomial logistic regression model, the binary logit model did show promise. The logit model showed that we may indeed be able to develop profiles of businesses that will increase the log odds of them switching from paper to an electronic mode of submitting their responses. It must be stressed that we are only at the beginning of understanding this analysis. It may be that a model may have to be built for each sector independently. As business characteristics differ between sectors, so too may the characteristics that predict mode change. Furthermore, as the nature of businesses change over time, we may need to think about building new variables into our models.

While we have addressed important preliminary questions as to respondent behavior in the 2012 Economic Census, there is still much more work we have to do. In particular, while we have a better understanding of response metrics for the EC, we have no idea how representative the respondents are to the initial "sample" drawn for the EC. Thus, we plan to evaluate potential for nonresponse bias in survey

estimates. Furthermore, we hope to incorporate cost data into our analysis to help determine what cost efficiencies might be achieved.

References:

Fink, E. B., & Lineback, J. F. (2013) Using paradata to understand business survey reporting patterns. *2013 Federal Committee on Statistical Methodology*.

SAS Institute (2013). *22603 - Producing an actual-by-predicted table (confusion matrix) for a multinomial response*. Retrieved July 8, 2014, from http://support.sas.com/kb/22/603.html

Willimack, D. K., & Snijkers G. (2013). The Business context and its implications for the survey response process. In G. Snijkers, G. Haraldsen, J. Jones, D. K. Willimack (Eds.), *Designing and conducting business surveys* (pp. 39 – 82). Hoboken, New Jeresey: John Wiley & Sons, Inc.