Imputing for Extraordinary Sample Events: A Story of Targeted Donor Pools and Administrative Data December 2023

Bradley D. Rhein
U.S. Bureau of Labor Statistics, 2 Massachusetts Ave., NE
Washington, DC 20212

Proceedings of the 2023 Federal Committee on Statistical Methodology (FCSM) Research Conference, October 24-26, 2023

Abstract

The Current Employment Statistics (CES) program produces monthly employment estimates by subnational geography and industry through a survey of about 666,000 establishments. Occasionally, a large-scale event, such as a hurricane, occurs near the survey reference period, significantly affecting data collection for that month, and challenging an implicit missing-at-random assumption in the estimator. Several methods have been activated to catch the ensuing employment drop, though none of these methods account for the nonresponse directly. This paper proposes the potential use of a random cold deck imputation approach, constructing a targeted donor pool from a similar circumstance occurring in past administrative data from the Quarterly Census of Employment and Wages (QCEW) program. The author explores using a donor pool from an earlier hurricane to impute for missing reporters due to hurricanes in 2017 and 2018.

1. Introduction

Large-scale, or extraordinary, events, such as a hurricane, may cause employment movements at the state and area level that are difficult to capture in estimation due to increased nonresponse. While the Current Employment Statistics (CES) estimators account for missing-at-random (MAR) nonresponse implicitly, increased nonresponse due to an extraordinary event challenges the usual assumption of missingness – that is, the increased nonresponse is not missing at random because it is a consequence of an extraordinary event.

This paper explores a possible imputation solution to account for the nonresponse caused by an extraordinary event. The next section gives a brief background of the survey, followed by brief descriptions of the sampling, estimation, and variance methods for the survey. Next, a discussion of the research question, an elaboration of the proposed methods, and results follow. Conclusions and future work wrap up the paper.

2. Background Information on CES

The CES program [1] produces detailed industry estimates of employment, hours, and earnings of workers on nonfarm payrolls. Each month, the Bureau of Labor Statistics (BLS) collects data through the CES survey, which includes about 122,000 businesses and government agencies and represents approximately 666,000 individual worksites. The sample is drawn from a sampling frame of unemployment insurance (UI) tax accounts for roughly 12 million establishments (covering about 97% of all US jobs).

CES publishes estimates for all 50 states, the District of Columbia, Puerto Rico, the U.S. Virgin Islands, and about 450 metropolitan areas and divisions. BLS releases these estimates around the third week of the month, along with the Local Area Unemployment Statistics (LAUS) estimates. Each publication includes a revision of the previous month's estimates as more sample data becomes available between releases.

CES benchmarks the state and area employment data on a yearly basis, primarily using information from the Quarterly Census of Employment and Wages (QCEW) [2]. The process creates benchmarked employment levels for all months up through September, with a publication of the updated levels and benchmark revisions to the estimates occurring in mid-March the following year.

3. Overview of the CES Sampling, Estimation, and Variance Methods for States and areas

The CES Handbook of Methods [3] elaborates further on each of the sections below.

Sample Design

The CES draws a stratified, simple random sample of establishments from the QCEW, an administrative data source collected quarterly through unemployment insurance accounts and covering about 97% of the jobs in the US economy. Stratification is by state, aggregate industry, and employment size class; selection is with optimum allocation.

The North American Industry Classification System (NAICS) [4] defines industries by 6-digit codes. CES may combine one or more NAICS industries into a single CES industry code to ensure complete coverage of all nonfarm businesses and to prevent disclosure of payroll data for individual establishments.

Industry codes and employment size class definitions appear below:

Table 1: Industry codes and employment size classifications [3]

CES Industry Code	Major Sector Name
10-000000	Mining and logging
20-000000	Construction
31-000000	Durable goods manufacturing
32-000000	Nondurable goods manufacturing
41-420000	Wholesale trade
42-000000	Retail trade
43-000000	Transportation and warehousing
44-220000	Utilities
50-000000	Information
55-000000	Financial activities
60-000000	Professional and business services
65-000000	Private education and health services
70-000000	Leisure and hospitality
80-000000	Other services
90-910000	Federal government
90-920000	State government
90-930000	Local government

Size class	Number of employees
1	0–9 employees
2	10–19 employees
3	20–49 employees
4	50–99 employees
5	100–249 employees
6	250–499 employees
7	500–999 employees
8	1000+ employees

Estimation

CES features two types of estimators for private industry employment series: sample-based and a model. As a time series, CES relies on a weighted link relative - the ratio of current month employment to prior month employment - to estimate the over-the-month (OTM) change, building off the prior month's employment level. The employment estimate also includes a series-level forecast of the net business births and deaths [5] and an adjustment for non-covered employment for religious organizations (part of the other services industry). The equation for the sample-based estimator is below:

Equation 1: Sample-based Estimator

$$\widehat{\mathrm{AE}}_c = \left[\widehat{\mathrm{AE}}_p - \sum_j \left(\mathrm{ae}^*_{p,j}\right) - \mathrm{NCE}^{RO}\right] \times \left[\frac{\sum_i \left(\mathrm{ae}_{c,i} \times w_i \times r_i \times d_i\right)}{\sum_i \left(\mathrm{ae}_{p,i} \times w_i \times r_i \times d_i\right)}\right] + \sum_j \left(\mathrm{ae}^*_{c,j}\right) + b_c + \mathrm{NCE}^{RO}$$

For all $i \in I$ and $j \in J$

where,

$i =$ matched sample unit;

$j =$ matched sample unit where the current month is atypical;

$w_i =$ weight associated with the CES report;

$r_i =$ down $-$ weight factor associated with the CES report;

$d_i =$ differential response rate factor associated with the CES report;

$\mathrm{ae}_{c,i} =$ current month reported all employees;

$\mathrm{ae}_{p,i} =$ previous month reported all employees;

$\mathrm{ae}^*_{c,j} =$ current month reported all employees where the current month is atypical;

$\mathrm{ae}^*_{p,j} =$ previous month reported all employees where the current month is atypical;

$\widehat{\mathrm{AE}}_c =$ current month estimated all employees;

$\widehat{\mathrm{AE}}_p =$ previous month estimated all employees;

$b_c =$ current month net birth $-$ death forecast; and

$\mathrm{NCE}^{RO} =$ noncovered employment in religious organizations (NAICS 813110).

Many series do not qualify for a sample-based estimator due to small sample sizes, low average response rates, or modest employment levels. For these series, a model estimator determines the current month employment level using information from a variety of sources, including the sample, the relevant group's state/area/series/month trends, state-related effects, industry effects, and historical state/area/series information. The model [6] employs a regression tree to determine state/area/series groupings, an unsupervised clustering mechanism to add industry (aka domain) effects, and variances to establish the influence of each component. The small area model equation is below:

Equation 2: Small Area Model Estimator

$$\hat{Y} = m_k + X\beta + u_{st} + u_d$$

where,

$\hat{Y}$ = model link;
m_k = intercepts derived by an unsupervised clustering process;
X = matrix of predictor variables;
β = vector of coefficients for intercepts and slopes;
u_{st} = state random intercepts; and
u_d = domain random intercepts.

All estimators feature outlier detection, non-sample events, seasonal adjustment, and key non-respondent (KNR) imputation.

A routine identifies outliers during the estimation process. Depending on the degree of outlierness, a sample unit either becomes self-representative or significantly downweighted, dampening its influence on the estimate. Self-representative outliers are not included in the weighted link relative.

Unique business closures, strikes, or layoffs not caught by the sample can be added as non-sample event adjustments. CES verifies the occurrence of the event, the number of employees affected, and the timing by collaborating with state partners. Non-sample events must occur during the survey reference period.

CES uses X-13ARIMA-SEATS software [7], developed and maintained by the U.S. Census Bureau, to seasonally adjust estimates. Seasonal adjustment eliminates the piece of the series that is attributable to normal seasonal variation, allowing for the observation of the underlying employment trend [8]. Seasonally adjusted series are published for selected nonfarm payroll employment series.

On occasion, large, seasonal, in-sample reporters fail to submit data for a month that, historically, has a seasonal movement. Such reporters, or key non-respondents (KNRs), have important seasonal movements that would be noticeable if missing from estimation. To account for these movements, should any KNRs fail to submit data, CES-imputes the over-the-month employment change using administrative values from the previous year.

Variance Estimation

Reliability measures for CES employment estimates are published on the BLS website for statewide aggregate industries and MSA total nonfarm employment [9]. Many estimates have small domains, causing variance estimates to be quite unsettled. To stabilize the variances, CES uses a model-based Generalized Variance Function (GVF) to calculate a measure of error for employment estimates [10].

4. Overview of the Research Question

Occurrences of extraordinary events near the reference period tend to increase nonresponse rates and challenge the missing at random assumption inherent to statistical surveys – that is, the extraordinary event directly causes some of the nonresponse. While this paper focuses on hurricanes, there are other types of extraordinary events as well. Future work could investigate solutions to floods, wildfires, tornados, snowstorms, and other extreme weather.

There are plenty of circumstances that could challenge an establishment's ability to submit data in a month where an extraordinary event occurs. Some storefronts could be boarded up and temporarily closed. Retail, leisure, and hospitality businesses tend to fall into these scenarios. Areas may have been evacuated. Construction projects could be delayed or postponed. However, some industries may add jobs, such as relief workers or cleanup crews belonging to the administrative, support, and waste services industry. For CES, if employees were paid during the reference period, even if the business closes temporarily, they are considered employed.

Currently, CES takes the following actions in the case of an extraordinary event:

- Monitoring response rates
- Watching for last-minute microdata submissions
- Altering the state/area/industry groupings used in model estimation (perhaps)
- Relying on the sample more (or entirely) in modeled series

Response rates tend to decrease in a month with an extraordinary event. Economists that monitor data collection may detect changes in response rates at the county level. Matching low response rates at a county level with areas designated as disaster zones may help identify geographic areas to monitor. Meanwhile, the monthly production cycle allows for a brief window to include additional microdata submittals before the Wednesday after calculating estimates (shown below). Estimates must be finalized by Wednesday (week 2) to ensure that seasonal adjustment and publication activities finish in time for the release.

Table 2: Example Production Cycle Calendar

Sunday	Monday	Tuesday	Wednesday	Thurs	Friday	Saturday
Week 1				Calculate estimates		
Week 2			Finalize estimates			

While the model estimator usually determines state, area, and industry groupings by using a regression tree, allowing small domains to borrow information from similarly trending domains, a model feature permits special groupings of state, area, and industry data if desired. In the case of an extraordinary event, affected states and areas may benefit from an isolated grouping, away from state, area, and industry combinations that are not experiencing such an event. Isolated groups could depend on county-level information, designated disaster areas, or simply states and areas near the site of the extraordinary event.

The existing sample-based estimator routinely captures current labor market trends due to a combination of sample design, cooperation of businesses in submitting data, coordination of field offices, and rigorous analysis and cleaning of microdata by economists working both in the region and at the BLS national office in Washington, DC.

Models, however, tend to perform well during "normal" times when there are no shocks to the system. To counteract this tendency during an extraordinary event, CES may alter the model to increase the reliance on the sample. While the model estimator may decrease the influence of the sample in a month containing an extraordinary event, a CES analyst may review the microdata, consider the sample to be representative, and choose to place more emphasis on the sample data (or rely on the sample entirely). As another fallback, the analyst may revisit that choice for the revised release of estimates, after several more weeks of data collection have taken place.

There can be cases that are longer term or have especially significant shocks to the labor market, such as the COVID-19 pandemic or Hurricane Katrina. In these cases, there are more severe solutions available for use, including the use of reported zeros, excess zeros, manual adjustment factors, UI claims, and modifying the business birth-death methodology temporarily.

Establishments that report zero employment in the current month or prior month do not contribute to the calculation of estimates since the birth-death forecast captures these movements. However, when Hurricane Katrina made landfall, reported zeros were used in estimation. Post-Katrina analysis revealed the method overstated the employment drop [11].

Since using reported zeros in estimation risks a too heavy-handed intervention, using “excess zeros,” which account for a baseline frequency of reported zeros, became available as an upgrade in methodology. The use of Excess Zeros (EZ) involves including establishments that report zero in the current or prior month, but then downweighting the effect of such a movement by applying the following ratio for each state, area, series, and month combination. Calculations occur separately for a reported zero in the current month and prior month. The ratio accounts for the number of zeros in excess of historical averages for a given state, area, series, and month combination.

Equation 3: Excess zero downweight formula

$$EZ\ Downweight = \frac{(number\ of\ reported\ zeros - historical\ average\ number\ of\ reported\ zeros)}{number\ of\ reported\ zeros}$$

The use of manual adjustment factors and alterations to the business birth-death forecast were explored in a *Monthly Labor Review* article written by Steven Mance [12].

Note that current methods do not address the scenario of not missing at random, which the author assumes is causing at least some of the nonresponse when an extraordinary event occurs. To clarify, the assumption is that the extraordinary event affects the act of submitting data for some non-respondents. So, the nonresponse may be due to numerous factors (too busy, evacuated, closed, or temporarily closed) that may or may not be related to the number of employees on the payroll during the reference period in a month with an extraordinary event. The proposed method, described in the next section, targets both non-respondents and donor pools to help ensure that the imputation solution applies to nonresponse that is not missing at random.

5. Exploring an Imputation Solution

Due to the nature of extraordinary events, selection of both the recipients of imputation and the donors providing the imputed values would ideally reflect the unique circumstance.

Defining the recipients of imputation

The last section notes that during an extraordinary event some nonresponse is not missing at random. One strategy for determining recipients that are not missing at random depends on whether the non-respondent “usually” submits data in time for preliminary estimates. Define these Usuals as establishments that have provided data for at least five out of the last six months in time for preliminary estimates (the earliest release of CES data). Such a definition strengthens the assumption that The Usuals’ failure to submit data was due to the occurrence of an extraordinary event.

Before imputing values for recipients at the establishment level, a couple modifications must occur. CES samples by UI account (stratified by state, aggregate industry, and size), assigning sample weights and employment size classes. Many establishments could fall under one UI account. Since imputation is set to occur at the establishment level, there is a need to modify the sample weights and employment size classes for all Usuals. As a result, The Usuals receive a substitution of appropriate sample weights and a reclassification of employment size class based on the individual establishment’s employment level.

Determining the donor pool

For the donor pool, consider historical data from the source of the CES sampling frame: the Quarterly Census of Employment and Wages (QCEW). This administrative data source contains monthly microdata, reported quarterly, for establishment employment and wages (based on total payroll). Establishments in the QCEW have NAICS designations, state, county, and various other business characteristic information. The QCEW also covers about 97% of the jobs in the US economy. These features make the QCEW an excellent data source for donor pools.

Since The Usuals were selected with the assumption that their act of submitting data was affected by the extraordinary event, the donor pool must be targeted in a similar way – by determining which historical events are similar to the extraordinary event affecting The Usuals. Upon finding an appropriate match in circumstances, the administrative data from the past forms a targeted donor pool from which to determine values for imputation.

Hurricanes could be measured and compared in various ways. The National Hurricane Center and the Central Pacific Hurricane Center [13], divisions of the National Weather Service that track and predict tropical weather systems, classify hurricane intensity along the path of the storm using the Saffir-Simpson Scale [14]. News reporting on hurricanes mention areas along the path and the location of where the storm makes landfall. Affected areas could be mined for population density (number of workers by county), industry composition (number of workers for each aggregate industry), geographic details (flood plains, mountains, forest areas, gulf, or ocean coast), and disaster designations. Using this information, for example, the effects of a past hurricane (e.g., Hurricane Irma on Naples, FL) could serve as a donor pool for a set of Usuals from a more recent hurricane (e.g., Hurricane Ian on Punta Gorda, FL).

The following table lists viable donor pools from recent hurricanes. Note that Hurricane Ian, from October of 2022, would be a good future addition to the list.

Table 3: Donor Pool Examples

Donor Pool	Hurricane	Category
Naples	Irma	4
FL	Irma	4
Corpus Christi	Harvey	5
Houston	Harvey	5
Victoria, TX	Harvey	5
Beaumont, TX	Harvey	5
Lafayette	Harvey	5
Panama City	Michael	5
Pensacola	Michael	5
Tallahassee	Michael	5
Jacksonville	Florence	1

Since the donor pool source is administrative, the number of available donors is vast and includes all employment sizes, NAICS designations, and county information. Actual behavior, on an establishment level, from a past extraordinary event is available for use as a direct donor to a non-respondent of interest. With this setup, the author suggests that there is no need to synthesize or model imputed values.

Narrowing the imputation method

There still needs to be a method to match a donor to a recipient. Consider two imputation methods used frequently in practice: random cold deck (RCD) and nearest neighbor (NN) from a cold deck. Since the donors arrive from an outside data source, the QCEW, the "deck," or list of potential donors, is a cold deck. Both methods utilize donor-to-recipient matching, transferring actual values from a donor establishment to an establishment-level recipient. The use of a specific imputation cell establishes the assumption that the donor and recipient are similar and behave similarly under similar circumstances. In this case, both donors and recipients have similar establishment characteristics and are subject to similar extraordinary events.

Effective cold deck imputation methods rely on a strict definition of the imputation cell, a sufficiently large donor pool, a modest need for imputed values, and the use of multiple imputation to avoid bias [15]. The setup for an extraordinary event arguably meets these qualifications due to a vast number of donors available (administrative data) for relatively few recipients (The Usuals), one missing value of interest (establishment employment), and a specific imputation cell definition consisting of the affected area, full 6-digit NAICS designations, and eight employment size classes. If zero donor-to-recipient matches exist under the initial imputation cell definition, a potential fallback could use the 4-digit NAICS designation and potentially a more aggregate grouping of employment size. There is some debate in the literature over how many multiple imputations to perform; the author briefly investigates a couple scenarios.

Using nearest neighbor imputation involves specifying an additional measure to compare the “distance” between a recipient and potential donor establishment. Note that the imputation cell ensures that the matching process already narrows the available donors by area, NAICS designation, and employment size class. However, size class is a range of “distances” – finding the closest match in employment level is one measure that could determine the nearest neighbor. For example, both the donor and recipient could be establishments with exactly 18 employees in the month prior to the month of interest. Again, the assumption is that similarly sized establishments behave the same under similar circumstances.

Since the donor pool is an administrative source, there could be many donors qualifying as nearest neighbors, perhaps with the same exact employment as the recipient. For these cases, a random number becomes the deciding factor, leaving the recipient with just one nearest neighbor. As a precaution against biasing the variance by using a single donor many times, each donor establishment may only contribute imputed values to only one recipient regardless of how many times the process chooses a given donor as the nearest neighbor for various recipients.

Examples of imputed values for some Usuals appear below, using Corpus Christi, TX from Hurricane Harvey as the donor pool. Ten and twenty-five multiple imputations tended to produce very similar results in terms of imputed value distribution and variance for all test cases. Multiple imputation results with high volatility were examined further for issues with donor values and were removed from consideration when necessary.

Graph 1: Multiple Imputation Results for Usuals in Punta Gorda, FL

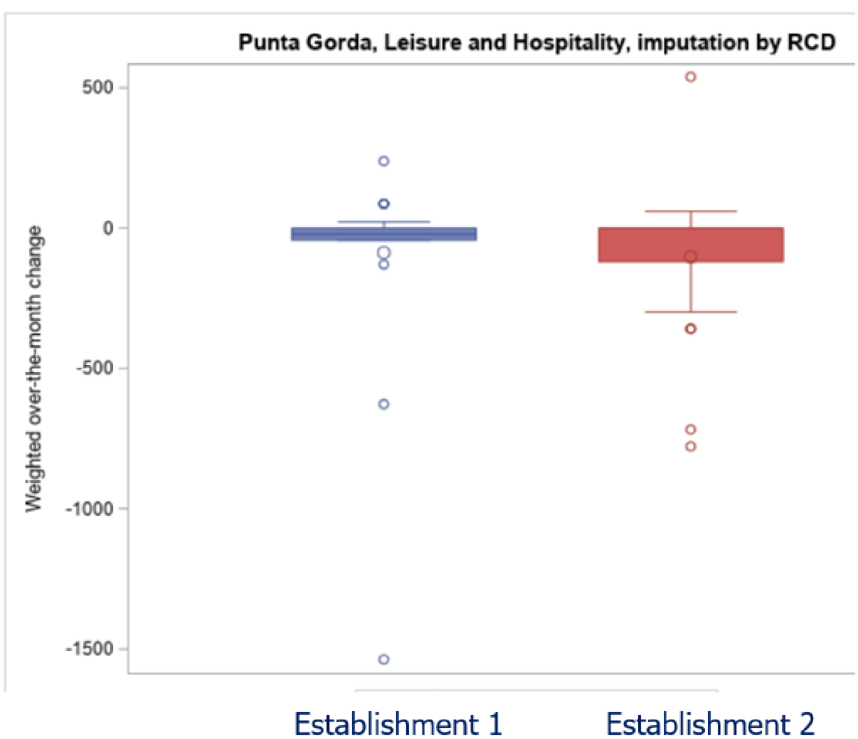

Establishment 1 falls into size class 2 (10 to 19 employees) and had a tight dispersion of imputed values as illustrated by the narrow quartile range. There was at least one extreme outlier found in 25 multiple imputations, having a weighted over-the-month (WOTM) change of -1,500. Note that while this outlier has little effect on the mean of 25 imputations, such an outlier could lead to larger-than-desired volatility if chosen by the nearest neighbor approach. Establishment 2 falls into size class 4 (50 to 99 employees) and had a slightly larger quartile range. Both companies have narrow distributions of imputed values that support the use of a mean as the final imputed value.

Graph 2: Multiple Imputation Results for Usuals in Cape Coral, FL

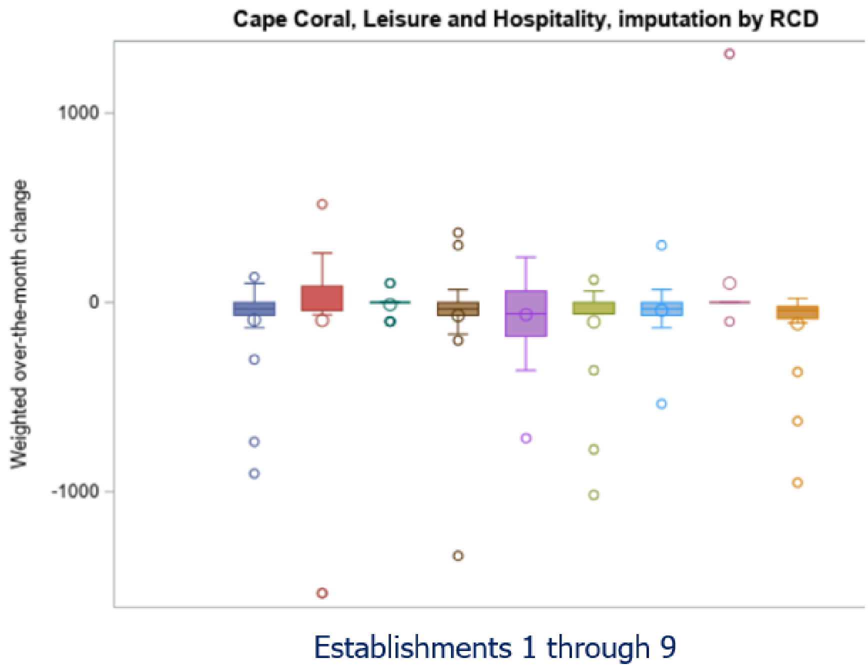

Nine Usuals in the leisure and hospitality industry were evaluated for Cape Coral, FL in the month of Hurricane Ian, October 2022. While there were some outliers, and most outliers were decidedly negative, the twenty-five multiple imputations seem to find means with small overall deviations for each Usual.

Once imputation is complete for all recipients in a series, the WOTM values become the input for a calculation of an estimate adjustment. Estimate adjustments are applied directly and account for the estimator, sample circumstances, and model inputs (if appropriate).

Implementation of the methodology

To illustrate the implementation of this methodology, consider the real-time runup to Hurricane Idalia, an extraordinary event that made landfall in Florida at the writing of this paper (October 2023). The category 3 hurricane made landfall near Perry, FL on August 30, 2023, causing significant damage to northern Florida. Along the way to Perry, FL, there was notable damage due to a storm surge near Tampa, FL where the hurricane reached its peak strength as a category 4 hurricane [16].

While August 30 is well past the reference period for August, the effects of the hurricane could possibly show up in the September job numbers for parts of Florida and perhaps Georgia. As a result, CES began monitoring microdata reports for the September reference period. Preliminary estimates for September were set to be run on October 4, so by the last week of September the following actions were taken:

- Monitored response rates

- Read news reports
- Checked for FEMA disaster designations
- Scanned for any Usuals

Although National Guard members were deployed and a federal emergency declaration was issued for the state of Florida, the response rates and microdata monitoring did not turn up any collection abnormalities compared to recent months. A search for Usuals came up empty as well.

If significant Usuals had been found, a search of donor pools would be the next step. After determining a donor pool, the imputation method and estimate adjustment calculations would produce results for consideration.

6. Analysis of the Empirical Results

The following test cases were explored, using the imputation methodology described above. Performance of the imputation methodology was evaluated by measuring the improvement (or worsening) in the percent revision – the percent difference between the CES published estimate and the benchmarked employment. CES published estimates can be found in archived news releases [17] while CES benchmarked employment can be found using the BLS data tool [18]. The equation for calculating the percent revision appears below.

Equation 4: Calculating the percent revision

$$Percent\ Revision = \frac{(CES\ Published\ Estimate - Benchmarked\ Value)}{Benchmarked\ Value}$$

Circumstances with both large and small revisions were of interest, as the imputation method ideally must not only improve the capture of employment drops when such drops occur (i.e., Irma), but also avoid exaggerating employment drops that did not happen (i.e., Elsa). Note that revisions for Hurricane Ian are preliminary (based on published data by QCEW), as CES has not yet benchmarked employment for October 2022. All other revisions could be found using archived press releases and the BLS data tool.

Table 4: Hurricane Test Cases

Hurricane	Category	Month	Year	Area Affected	Revision
Irma	4	September	2017	Naples	-6.70%
Irma	4	September	2017	Florida	-1.20%
Ida	4	September	2021	New Orleans	-1.04%
Ida	4	September	2021	Lafayette	-0.73%
Ian	5	October	2022	Cape Coral**	-2.80%
Ian	5	October	2022	Punta Gorda**	-3.60%
Elsa	1	July	2021	Tallahassee	0.10%
Idalia	3	September	2023		

In September 2017, Hurricane Irma made landfall near the Naples, FL metropolitan area, causing considerable damage and a temporary decline of about 10,000 workers in various industries [19]. Job losses at the total private industry level appear below, illustrating the noticeable difference between the CES estimate – a drop of 2,000 jobs [20] – and the benchmark value for September of 2017. A downward revision of about 6.7% was found after completing the benchmark. Note that the confidence intervals are crudely calculated – prior month CES published values are assumed correct - and were added to give the reader a sense of the variance.

Graph 3: Effect of Hurricane Irma on Naples, FL

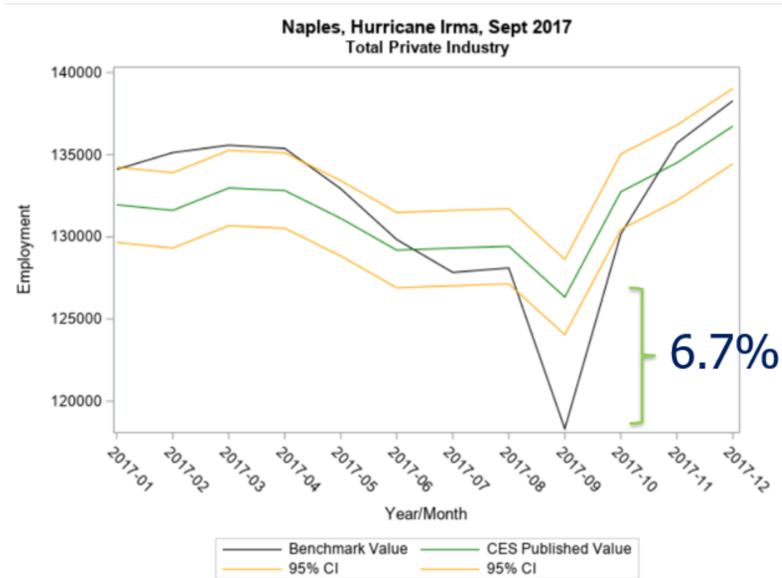

At the industry level, consider the following examples.

Graph 4: Effect of Hurricane Irma on Retail Trade in Naples, FL

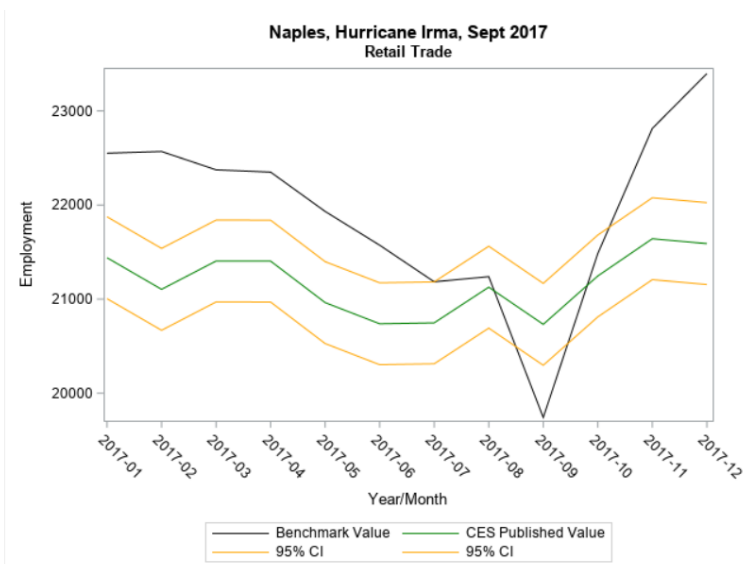

Retail trade had a noticeable downward revision of about 5% for September 2017. The sample-based estimator operated below the benchmark level for most of the year, up until August 2017. Some employment drop was captured in September. Could the imputation method provide a more accurate result?

Using Panama City, FL, an MSA near the landfall of Hurricane Michael (October 2018, category 5), as a donor pool, three Usuals were imputed for, and an estimate adjustment was calculated and added, as illustrated below. Panama City during Hurricane Michael was selected as the donor pool due to its employment size, similar location traits, and hurricane intensity at the time of landfall.

Graph 5: Imputation Results for Retail Trade, Naples, FL

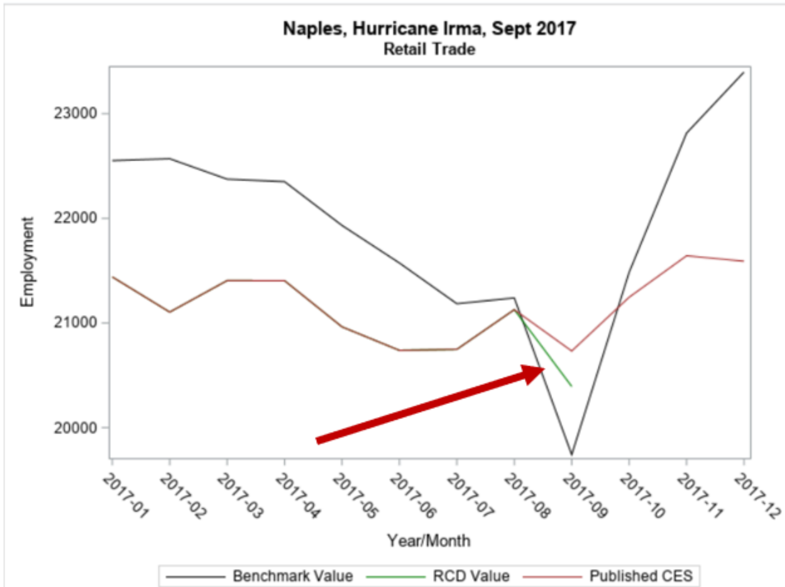

With 25 multiple imputations for each of the three Usuals, the random cold deck (RCD) method captured a further drop in employment for September 2017, decreasing the downward revision from 5% to 3.3%. The nearest neighbor (NN) method did not provide any results for retail trade.

As another example, consider the education and health services industry in Naples, FL. By August there was already a slight downward revision, which expanded notably to about 5.4% in the month affected by Hurricane Irma. In Naples, FL, this industry typically has seasonal movements, and the CES model estimate captured the customary seasonal increase in September 2017, as shown by the upward movement in the green line below.

Graph 6: Effect of Hurricane Irma on Education and Health Services in Naples, FL

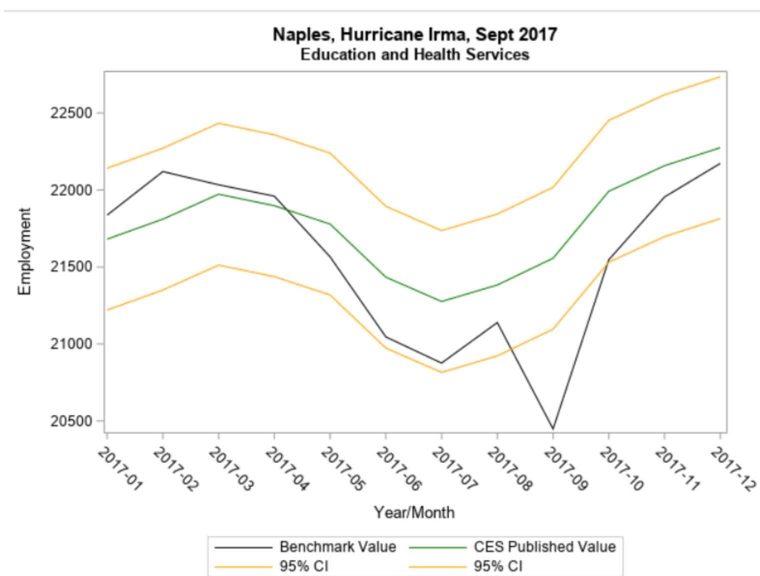

Both imputation methods were able to catch more of the employment drop in September, imputing for four Usuals and continuing the use of Panama City during Hurricane Michael as the donor pool. Revisions improved to about 3.5% and 3.8% utilizing the random cold deck and nearest neighbor imputation methods, respectively.

Graph 7: Imputation Results for Education and Health Services, Naples, FL

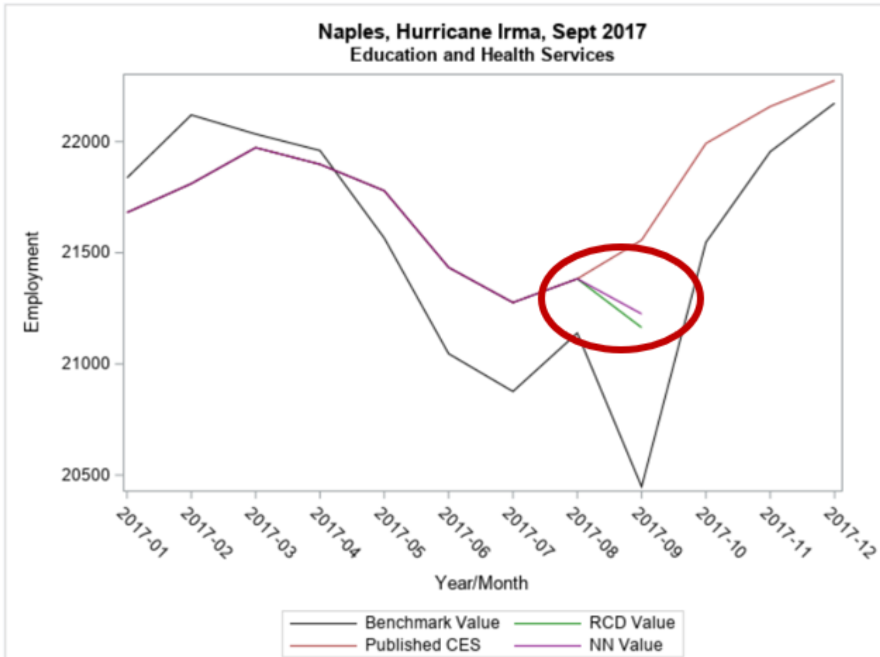

It is worth noting, for the purpose of the next example, that most hurricanes tend to strike between August and October, a time of the year that is furthest from the annual benchmark for CES at the state and area level. While the estimates and benchmark values line up perfectly in the previous September, by the time the hurricane season approaches, the estimates have been potentially deviating from the benchmark level for ten to twelve months. Such deviations are realized and measured as percent revisions once the employment has been benchmarked.

The next example illustrates the complications of attempting to capture an employment drop while operating under a notable pre-existing revision. Leisure and hospitality had about a 7% revision (see the red bracket below, the difference between the black and green line) in August of 2017, and this pre-existing revision “helped” the CES estimator nearly perfectly capture the benchmarked employment level after the drop caused by Hurricane Irma.

Graph 8: Effect of Hurricane Irma on Leisure and Hospitality in Naples, FL

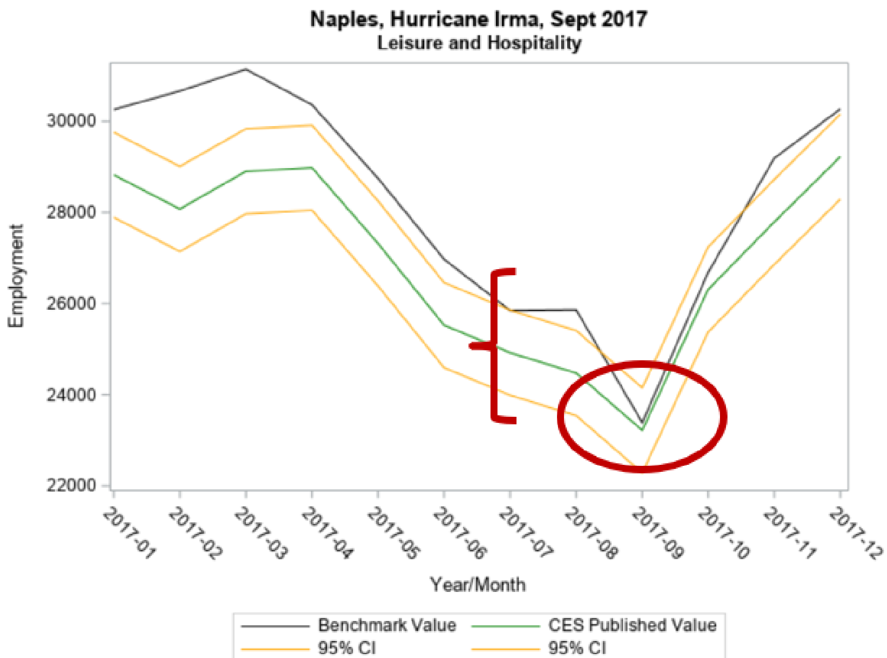

There were four Usuals in leisure and hospitality, and the estimate adjustment calculated from the imputed values drove the revision to increase significantly. However, considering the capture of the over-the-month change, using the imputation method was quite an improvement. Notice how the slope of the benchmark line (arrow) compares to the slope of the circled, imputed over-the-month change. CES reports over-the-month changes, and so while the post-imputation employment level is noticeably lower than the benchmark value for September, the over-the-month change has been captured quite accurately.

Graph 9: Imputation Results for Leisure and Hospitality, Naples, FL

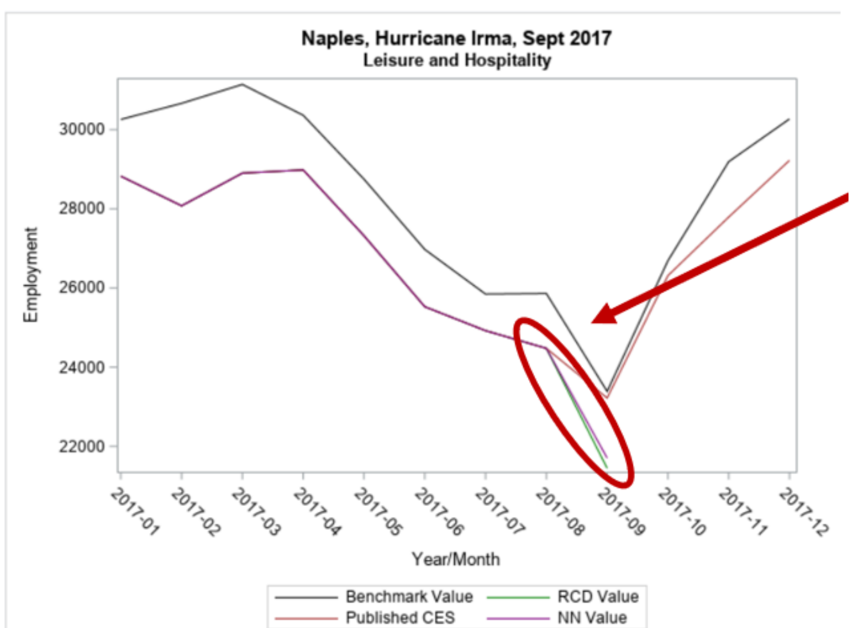

All results for Naples, FL appear below. Only four industries had non-negligible imputation effects, three of which notably improved revisions and one that did not (though, as discussed above, the over-the-month change was captured more accurately). The state of Florida had 1,265 Usuals but zero imputation effects due to estimate adjustments offsetting each other. A statewide revision of 1.2% (yellow highlight) is notable - about 116,000 workers - and well outside the 95% confidence interval for September 2017. An improvement in estimates at the MSA level could filter up to aid the statewide estimate values.

Table 5: Imputation Results for Industries affected by Hurricane Irma in Naples, FL

Event	Area Affected	Donor Pool	Industry	# of Usuals	Employment Level	Orig Revision	RCD Revision	NN Revision
Irma	Naples	Panama City (Michael)	Retail	3	20,000	5.0%	3.3%	.
Irma	Naples	Panama City (Michael)	Professional	2	15,000	9.6%	6.8%	2.0%
Irma	Naples	Panama City (Michael)	Health/Ed	4	20,500	5.4%	3.5%	3.8%
Irma	Naples	Panama City (Michael)	Leisure	4	24,000	0.7%	-8.3%	-7.2%
Irma	Florida	Houston (Harvey)	All Industries	1,265	9,688,000	1.2%		

Other results appear below. Cells highlighted in yellow identify cases with somewhat small revisions where there may be no need for further imputation effects (and none were found). Hurricane Ian was recent enough that CES benchmark data was not yet available – revisions shown here are preliminary, based on fourth quarter data published by the QCEW. Blue highlights identify series with model estimators that were plausibly switched to sample-based estimators as a result of sufficient sample representation, response rates, and conflicting historical trends that the model happened to prioritize (much like the education and health services example in Naples, FL). In a couple of cases, multiple donor pools were used, each finding similar results.

Table 6: Imputation Results for Other Hurricane Test Cases

Event	Area Affected	Donor Pool	Industry	# of Usuals	Employment Level	Orig Revision	RCD Revision	NN Revision
Ida	New Orleans	Corpus Christi (Harvey)	0	69	448,000	1.2%		
Ida	Lafayette	Lafayette or CC (Harvey)	0	26	195,000	0.7%		
Ian	Cape Coral	Corpus Christi (Harvey)	Manufacturing	4	7,700	4%**	0.0%	.
Ian	Cape Coral	Corpus Christi (Harvey)	Wholesale Trade	4	7,700	13.2%**	.	10.8%
Ian	Cape Coral	Corpus Christi (Harvey)	Financial	7	14,000	9.4%**	8.4%	.
Ian	Cape Coral	Corpus Christi (Harvey)	Professional	6	40,000	12%**	.	4.9%
Ian	Cape Coral	Corpus Christi (Harvey)	Leisure	9	37,000	19.8%**	13.1%	11.0%
Ian	Punta Gorda	Corpus Christi (Harvey)	Leisure	2	6,900	11.3%**	5.7%	.
Ian	Punta Gorda	Naples (Irma)	Leisure	2	6,900	11.3%**	-3.0%	3.6%
Ian	Punta Gorda	Naples (Irma)	Retail	2	9,500	10.1%**	7.1%	.
Elsa	Tallahassee	Corpus Christi (Harvey)	0	9	182,000	0.1%		
Idalia	FL, GA, NC							

Note that all imputation effects improved revisions, though effects were only found for some industries. Overall, the random cold deck found more results due to the stability of using an average among multiple imputations. Nearest neighbor results tended to be much more volatile and often found that the nearest neighbor donor selected had zero over-the-month change.

6. Summary of the Results

There appears to be some reduction in revisions with the use of the proposed imputation methods. The random cold deck provides more imputation effects compared to the nearest neighbor method, mostly due to the use of multiple imputations that have a habit of centering nicely around a mean imputed value. More time, practice, and additional donor pools will likely improve results.

Future work could focus on finding more statistically sound methods for choosing a donor pool, further analysis of the variance effects, both for imputed values and post-adjusted estimate values, addressing the concern of capturing the recovery in the months after an extraordinary event, and investigating the use of imputation for other events, such as wildfires, tornados, flooding, snowstorms, and other extreme weather events.

References

[1] CES State and Area, www.bls.gov/sae

[2] Quarterly Census of Employment and Wages (QCEW), www.bls.gov/cew

[3] CES-SA Handbook of Methods, www.bls.gov/opub/hom/sae/

[4] North American Industry Classification System, https://www.census.gov/naics/

[5] Mueller, Kirk, "Impact of business births and deaths in the payroll survey," *Monthly Labor Review*, May 2006, https://www.bls.gov/opub/mlr/2006/05/art4full.pdf.

[6] Gershunskaya, Julie, Savitsky, Terrance D, "Bayesian nonparametric joint model for point estimates and variances under biased domain variances," November 2019, https://www.bls.gov/osmr/research-papers/2019/st190020.htm.

[7] The X-13ARIMA-SEATS software is available on the U.S. Census Bureau website at https://www.census.gov/data/software/x13as.html

[8] More information on the seasonal adjustment process for CES National can be found at this link - https://www.bls.gov/web/empsit/cesseasadj.htm. CES-SA uses a 2-step process, Stuart, Scott, Stamos, George, Sullivan, Thomas, Chester, Paul, "Seasonal Adjustment of Hybrid Economic Time Series," https://www.bls.gov/osmr/research-papers/1994/pdf/st940350.pdf

[9] CES reliability information, https://www.bls.gov/sae/additional-resources/reliability-of-state-and-area-estimates.htm

[10] Gershunskaya, Julie, Dorfman, Alan H., "Calibration and Evaluation of Generalized Variance Functions," October 2013, https://www.bls.gov/osmr/research-papers/2013/pdf/st130050.pdf.

[11] Garber, Molly, Unger, Linda, White, James, Wohlford, Linda, "Hurricane Katrina's effects on industry employment and wages," *Monthly Labor Review*, August 2006, https://www.bls.gov/opub/mlr/2006/08/art3full.pdf

[12] Mance, Steven, "Estimating state and local employment in recent disasters – from Hurricane Harvey to the Covid-19 pandemic," *Monthly Labor Review*, April 2021. https://www.bls.gov/opub/mlr/2021/article/pdf/estimating-state-and-local-employment-in-recent-disasters-from-hurricane-harvey-to-the-covid-19-pandemic.pdf

[13] National Hurricane Center and Central Pacific Hurricane Center, https://www.nhc.noaa.gov/

[14] Saffir-Simpson Hurricane Scale, https://www.nhc.noaa.gov/aboutsshws.php

[15] Discussion in section 7.3, conclusion in section 9; Andridge, Little. 2010. A Review of Hot Deck Imputation for Survey Non-response. International Statistical Review, April 2010, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3130338/

[16] NOAA report on Hurricane Idalia nearing landfall, https://www.nhc.noaa.gov/archive/2023/al10/al102023.public.015.shtml

[17] CES News Release archive, https://www.bls.gov/bls/news-release/metro.htm

[18] CES data tool, https://www.bls.gov/sae/data/

[19] This drop could be found using the CES data tool (https://www.bls.gov/sae/data/) or QCEW data tool (https://www.bls.gov/cew/data.htm)

[20] CES News Release for September 2017, https://www.bls.gov/news.release/archives/metro_11012017.htm

Note: Any opinions expressed in this paper are those of the author(s) and do not constitute policy of the Bureau of Labor Statistics.