The effects of rounding on the Consumer Price Index

Calculating percent changes in a price index rounded to three decimal places mitigates a problem that can arise when percent changes are based on the same index rounded to a single decimal place

Elliot Williams

The Bureau of Labor Statistics (BLS, the Bureau) rounds the Consumer Price Index (CPI) to a single decimal place before it is publicly released. In 1984, the index was rebased to 100.0, and it stands near 200 today. Because the index value is so large, one might think that the difference between, for instance, a CPI of 189.7 and a CPI of 189.72 would be negligible. However, it is not negligible for the *percent change* between two CPI values, or CPI inflation. Because the actual changes in the CPI have been quite small recently (the rate of inflation has been relatively low), the small differences incurred in rounding up or down can create a misleading picture of monthly price inflation.

And that difference matters in the economy. For example, this Reuters news article is representative of the impact that the release of the February 2005 index value had, signaling a surprisingly large increase in the rate of inflation to financial markets:

> The core CPI, which strips out volatile food and energy costs, rose 0.3 percent. It was the biggest rise in the core rate since September and broke a string of four straight 0.2 percent gains.
>
> Wall Street economists had braced for a milder 0.3 percent rise in overall consumer prices and had expected another 0.2 percent gain outside food and energy.
>
> The report added to financial market inflation jitters and increased speculation [that] the Federal Reserve, which raised credit costs on Tuesday, might step up the pace of its rate rise to keep inflation under wraps.[1]

Both the stock and bond markets moved on the news that the index for all items less food and energy increased from inflating at a steady 0.2-percent rate to 0.3 percent, a relatively large growth in the rate of inflation. But in this case, the apparent increase is an artifact of using already rounded index values to calculate the inflation rate. Calculating the "core" inflation rate (the CPI for all items less food and energy) by using an unrounded CPI index series gives 0.2 percent instead of 0.3 percent and would have constituted essentially no news for inflation projections or bond prices. This article demonstrates how such an artifact can arise and investigates how frequently there is a discrepancy between inflation rates calculated from unrounded indexes and those calculated from rounded indexes under different possible rounding policies.

Although the rounding error in recent months' CPI inflation can cause a passing stir in the financial markets, some effects of rounding are still more marked in the historical CPI series. Plotting the percent changes in the published CPI all-items series with points in addition to the usual lines makes the rounding apparent to the naked eye, as shown in chart 1.

There is nothing fancy about this plot; each month, a percent change is calculated and plotted. The fanning horizontal lines that the eye picks up are evidence of the fact that rounding the original series constrains changes in the level of the CPI to integer multiples of 0.1. The percent changes after rounding are thus integer multiples of 0.1, divided by the level of the CPI in the beginning period. The horizontal lines that appear in the series correspond exactly to $\{\ldots, -0.2, -0.1, 0.0, 0.1, 0.2, \ldots\}/CPI_t$, where the subscript t denotes time.

Elliot Williams is a research economist in the Division of Price and Index Number Research, Office of Prices and Living Conditions, Bureau of Labor Statistics.

Chart 1. CPI inflation

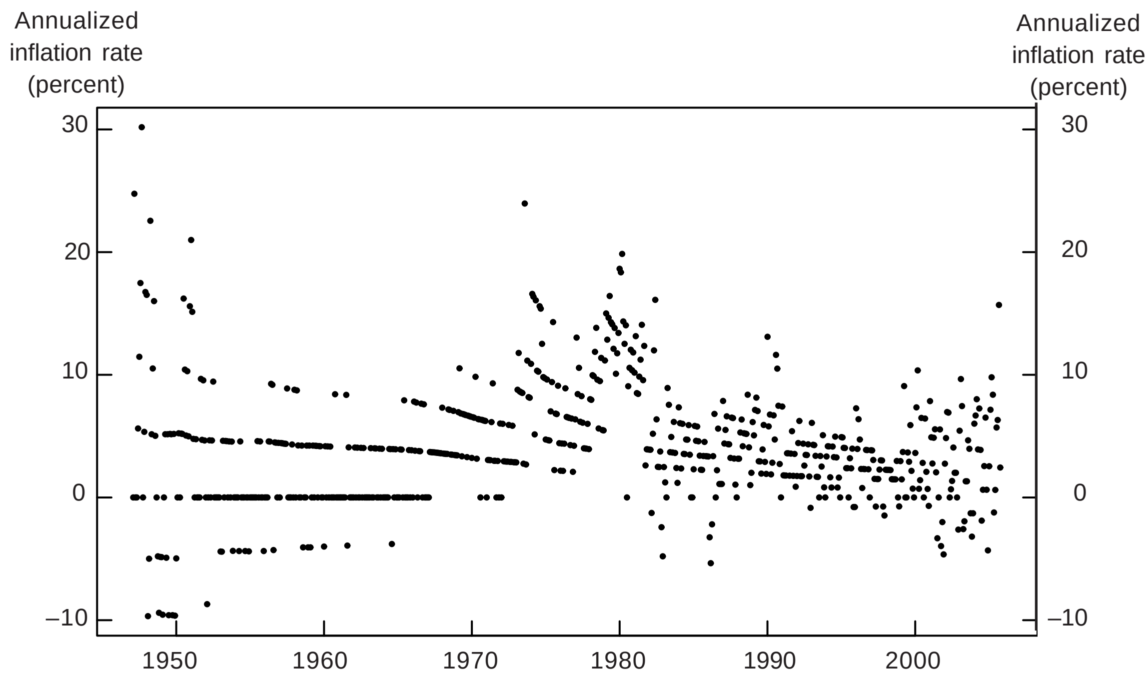

Chart 2 demonstrates this fact by overlaying the plot of points with lines corresponding to the different allowed rounded CPI inflation values in the historical 1984-base-year series. All of the percent changes in the reported series line up nicely, as they must.

Even without any mathematical analysis, some features of the percent difference series are immediately apparent. First, inflation takes on discrete values that widen and become more separated as the level of the CPI decreases as one traces the series backwards in time. In the period between 1955 and 1970, for example, inflation took on one of four values, and only two of them with any regularity. The most that can be said about this period from the rounded data is that monthly inflation was at an annualized rate somewhere between 0 percent and 5 percent.

Second, gradual inflation, especially in the earlier part of the series, is replaced with months of zero inflation followed by months with too-large inflation. The effect is visible on inspection: too many of the earlier months in the series register zero inflation. Over all the postwar data, 19.5 percent of the monthly changes are exactly zero, and each month that is rounded down to zero is offset by other months that are rounded upwards by the same amount. Thus, rounding tends to inflate the time-series variance of inflation, making it appear that monthly inflation was swinging wildly during the period, when, in fact, it was relatively calmer.

In recent times, rounding error has increased the variability of CPI inflation significantly. Because the Bureau now collects a very large number of prices for goods and services every month, the sampling variation of CPI inflation is very small—on the order of 0.0036 percent monthly. As will be shown in the analysis that follows, rounding error adds a further 0.0026 percent error variance to the reported figures—a further 72 percent of the variation itself. In the case of monthly CPI all-items inflation, the current BLS rounding criterion obscures a reliably estimated figure.

The picture of rounding in CPI inflation is not entirely bleak, however. In the long run, the rounding errors *do* average out, so rounding is not a source of long-term bias in the index. Rounding also is a less important source of error in the annual inflation series than in the monthly series, because the average of 12 rounding errors is closer to zero than a single error is and the magnitude of a year's inflation (for most years) is larger than the magnitude of a single rounding error. Moreover, the Bureau still makes available a 1967-base-year series that is less subject to rounding error due simply to the fact that the index values are larger; rounding to the tenths place produces a smaller relative error.[2]

Finally, one could, in principle, calculate a more accurate monthly inflation series by going back to the original publications

Chart 2. CPI Inflation, overplotted with allowed levels

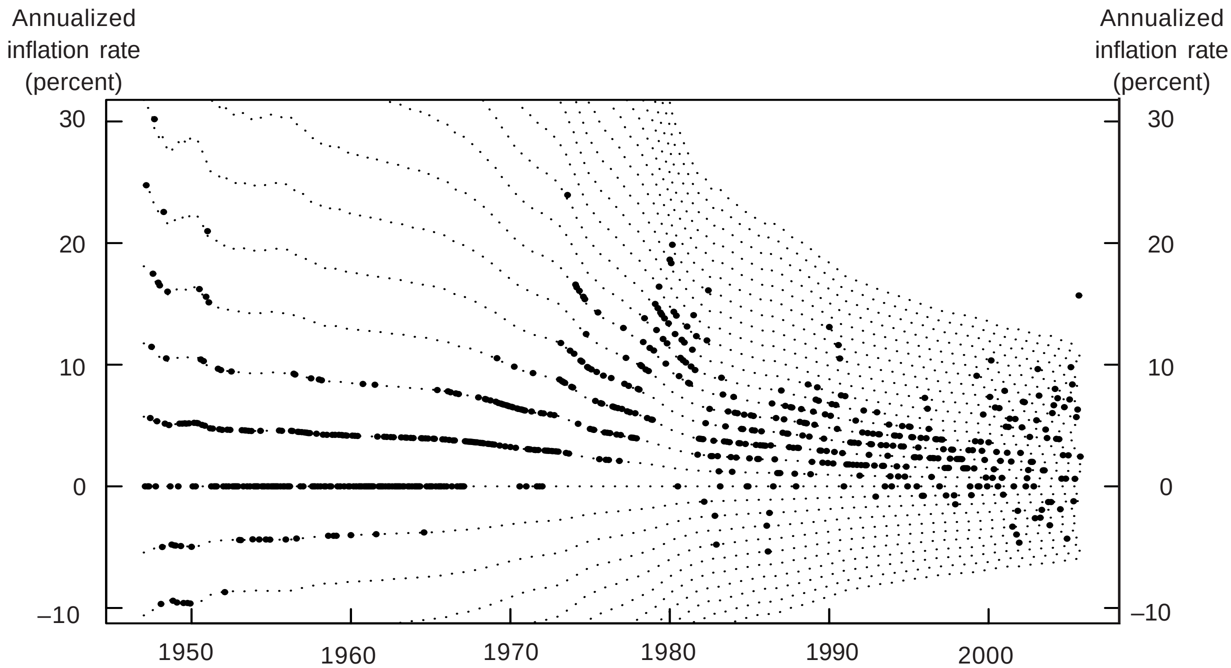

(when the values were higher because the series had not yet been rebased) and converting them to current values, retaining the extra precision.[3]

Nevertheless, for short-run inflation based on the current CPI series, rounding to the first decimal place affects the accuracy of contemporaneous and historical data. The next section details how the CPI series are rounded and demonstrates by example how discrepancies can arise. The section after that examines the effects of rounding error on recent inflation data for which an unrounded counterpart is available. The final section mathematically analyzes the effect of rounding error to extend these results to the entire historical CPI data series.

Rounding the CPI

The Bureau long ago standardized on one decimal place as the level of precision for reporting all of its CPI series. Both the level of, and the percent change in, the CPI are rounded to the tenths place before being released to the public as official statistics. However, because the Bureau wishes to have the released inflation series match the released index series, CPI inflation is calculated from the *rounded* CPI index values. Chart 3 illustrates the way the CPI is rounded.

Notice that the final inflation figure has been rounded twice, once before a percent difference is taken to calculate inflation and once afterward. The two stages of rounding the CPI inflation series have qualitatively different effects. The final stage of rounding merely shortens the figure and provides a signal of how much confidence the Bureau has in the estimate of inflation. Indeed, it is possible to motivate the choice of rounding the inflation rate to the nearest 0.1 percent by appealing to the Bureau's estimates of the sampling variation in the CPI. Approximate 95-percent confidence intervals can be constructed around the reported inflation figure by adding or subtracting 0.12 percent. Thus, when the Bureau reports a change of 0.2 percent in the inflation rate, one can be 95 percent confident that the true value lies between approximately 0.1 percent and 0.3 percent. Releasing the final inflation figure with less precision would obscure detail the Bureau measures well, while releasing more precision would give the appearance of more confidence in the estimate of inflation than is warranted.

The first-stage rounding—of the CPI index level—is the cause of the problems just documented. A numerical example demonstrates how rounding the CPI index before calculating the inflation rate can result in discrepancies between rounded

Chart 3. CPI rounding procedures

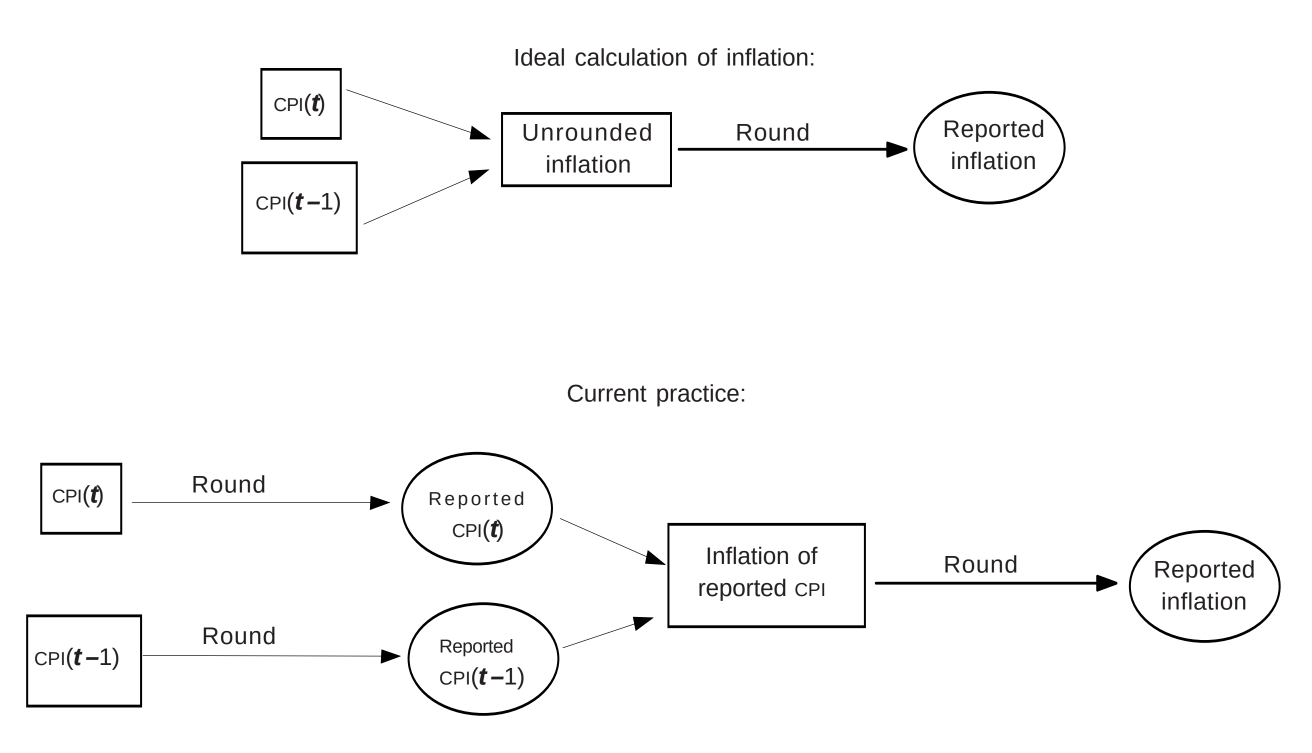

and unrounded figures:

Calculation	*Unrounded*	*Rounded*
Raw CPI$_t$	192.345	192.3
Raw CPI$_{t-1}$	192.770	192.8
Change in CPI$_t$	.425	.500
Percent change in CPI$_t$	.221	.260
	(.425/192.770)	(.5/192.8)
Rounded percent change in CPI$_t$	.2	.3

The first column of the tabulation corresponds to the ideal inflation calculation method presented in chart 3. The change in the CPI is calculated by subtracting the previous period's unrounded CPI from the current period's unrounded value, dividing this difference by the previous period's CPI, and, finally, rounding the result. The second column corresponds to the current BLS practice: the same procedure is carried out, but starting instead with the rounded CPI values.

A comparison of the rounded and unrounded CPI levels reveals that they differ only in precision. However, the change in the CPI calculated from the rounded data differs from the change calculated from the unrounded data. This discrepancy then carries over into the final percent change.

Note that the difference between the two changes is small relative to the size of the index; thus, the difference between the rounded and unrounded percent changes is relatively small (0.1 percent), though not insignificant. For the historical series, the relative size of the errors increases because, although the difference between the rounded and unrounded numbers remains constant as one traces the series back in time, the level of the index drops. This makes the error in the percent change series become larger and larger as the level of the index falls.

For instance, if this example were based on 1960s data, when the index level was around 30, the unrounded CPI values could have been 30.345 and 30.770 for the previous and current periods, respectively. The error in the change in levels would remain the same at 0.075, while the error in the percent change would be 0.075/30 = 0.25 percent, twice as large as the average monthly inflation rate in 1960. A rounding error of this size would obscure the actual monthly changes in the inflation rate for a large part of the historical series.

Given that the Bureau desires to use prerounded indexes to calculate monthly CPI inflation and desires to report an accurate statistic to 0.1-percent precision, how many digits should it retain in the CPI levels series? How often do the reported BLS inflation numbers differ from an accurate measure of inflation calculated with unrounded CPI figures, and when the figures do differ, by how much do they differ? Finally, what effect does rounding the CPI levels have on the error variance of inflation?

To answer these questions, this article takes two paths. Where rounded CPI data are available, the published CPI inflation values are compared with those calculated before rounding and the question is asked, How often do they match at the reported level of precision? To address the frequency and magnitude of these differences in pre-1986 data, the article relies on some simple statistical analysis.

Rounding policies and real data

To construct a measure of inflation that is free from rounding error, this section uses the CPI's Research Database index data files. The data employed here include all of the major indexes from January 1986 to July 2005 at the full level of precision used internally at the Bureau. For this article, the CPI all-items index and its top-level components are considered. In addition, the information technology and personal-computer indexes are included because they have seen rapid declines in price and are probably the worst-case scenario for rounding error in the post-1986 period.

A monthly benchmark inflation series is calculated from the unrounded data and then is rounded to the one-tenth-of-a-percent level to match the published inflation series, as in the ideal method presented in the previous section. To copy current and possible BLS procedures, the data from the Research Database also are rounded to one, two, and three decimal places initially, and inflation rates are calculated. The resulting inflation series is then rounded to the tenths place in percentage terms. The only difference between the benchmark and rounded series is the precision in the first stage of rounding.

Table 1 reports the percentage of the sample for which the inflation rates in the rounded data differ from the benchmark series at a 0.1-percent level of precision. Results are presented for both the non-seasonally-adjusted series and the seasonally adjusted series. Because the two series are similar and the rounding errors should be independent between them, the differences in the percentages shown give an indication of the variability of the estimated percentage.

The table shows that following the current practice of rounding the CPI index to the tenths place results in a derived monthly inflation that is materially different from the benchmark inflation rate roughly 25 percent of the time. This finding is basically consistent across the various series, with a few exceptions. The relatively low percentage of differences in the medical index and the index for other goods and services is due to the fact that those sectors saw high inflation over the 1986–2005 period and consequently have large index values for most of the period. In contrast, information technology and personal computers decreased in price dramatically over the same period and so have very small index values, making the first-stage rounding error large enough to change the monthly inflation rate as often as 75 percent of the time.

A look at the columns corresponding to retaining two and three decimals in the CPI indicates that the frequency of discrepancies between the inflation series can be reduced to nearly zero for most series (though not the problematic personal-computer series) by reporting the index rounded to three, rather than two, decimal places.

If the inflation series created from CPI data rounded to the tenths place differs from the benchmark series roughly 25 percent of the time, by how much is it off? Fortunately, the rounded data are precise enough that the difference is always limited to ±0.1

Table 1. Percent of cases in which rounded monthly inflation rates differ from unrounded rates, by digits of precision

[In percent]

Index	Not seasonally adjusted			Seasonally adjusted		
	One digit	Two digits	Three digits	One digit	Two digits	Three digits
All items	26	1.7	0.4	24	0.9	0.0
Food	23	1.7	.0	17	.9	.0
Energy	32	4.3	.4	30	2.1	.4
All items less food and energy	25	1.3	.4	16	3.4	.9
Apparel	27	1.7	.9	26	1.7	.0
Education and communication	35	3.3	.0	37	2.0	.7
Food and beverages	21	1.7	.0	16	3.0	.0
Other goods and services	13	1.3	.0	13	2.1	.4
Housing	26	3.0	.4	19	.9	.0
Medical	13	.9	.0	12	2.1	.0
Recreation	23	3.3	.7	24	3.3	1.3
Transportation	26	2.1	.4	24	2.1	.0
Information technology	57	10.3	.0	([1])	([1])	([1])
Personal computers	75	18.7	2.2	([1])	([1])	([1])

[1] Index not seasonally adjusted.

percent from 1986 to the present. In recent times, however, monthly inflation rates have been around 0.2 percent, which makes the rounding error as a percentage of the actual monthly change quite large indeed. Table 2 summarizes the distribution of the magnitude of the rounding errors relative to the unrounded inflation rate for the all-items index.

The first column of the table indicates that, of the 234 total observations of the rounded CPI all-items inflation index, 19 (8.1 percent) are in error by between 25 percent and 50 percent of the magnitude of the unrounded monthly change. Summing down the columns reveals that 62 observations (26.5 percent) differ by more than 5 percent of the benchmark inflation rate. Slightly more than 21 percent of the time, the reported CPI inflation rate differs from the benchmark inflation rate by 25 percent or more. More than 6 percent of the time, the inflation rate derived from the CPI rounded to one decimal place is off by 100 percent or more.

Reading across the table makes it clear that raising the initial level of rounding to the hundredths place eliminates all of the very large relative errors. Reporting the index rounded to the thousandths place would reduce the frequency of discrepancies to under 1 percent, and the magnitude of the error would be greatly diminished.

An alternative measure of the importance of rounding error for CPI inflation is a comparison of rounding error variance with the intrinsic sampling error variance. Sampling error arises because the Bureau is unable to collect all prices on all goods in the market and instead takes a sample of these prices. To assess the reliability of the sample of prices collected, the Bureau reports an estimate of error variance due to its sampling procedure. Currently, the monthly sampling error variance of all-items CPI inflation is about 0.0036 percent.

The estimates of the sampling error variance were created from unrounded figures, so adding rounding error to the CPI increases the variance of the reported inflation series relative to an unrounded series. The following tabulation compares the contributions to total error variance made by sampling error and by rounding error:

Type of error variance	Number of decimal digits reported: One	Two	Three
Total	0.0062	0.0038	0.0036
Sampling	.0036	.0036	.0036
Rounding	.0026	.0002	.0000

One can see that rounding error variance is approximately 72 percent as large as sampling error variance. Reducing the rounding error variance would reduce the total error variance by 42 percent.

Mathematical analysis

To get a better feel for how rounding error affects the historical inflation record further back into the past, it is worthwhile to undertake the same experiments as those just presented and compare the percent changes in the unrounded figures with their rounded counterparts. The Bureau, however, does not produce a full-precision historical series. Instead, a mathematical analysis, despite being in some sense approximate, provides additional insight that the data alone could not supply. The analysis that follows parallels the steps the Bureau takes in producing the CPI inflation figures, as summarized in the previous section.

Given a rounded CPI index value, the true (unrounded) value must lie within a known range, but which particular digits have been rounded away remains unknown. For instance, if the reported CPI level is 145.2, the true value can lie anywhere between 145.15 and 145.25 with equal likelihood. That is, the true and rounded levels can differ by one-half of the precision in either direction, or

$$\mathrm{CPI}_t = \mathrm{CPI}_t^* + \varepsilon_t,$$

where CPI_t is the rounded CPI level in month *t*, CPI_t^* is the

Table 2. Density (percentage error count) of relative errors, by digits of precision, in CPI all-items inflation

[In percent]

Relative errors	One digit: Number	One digit: Percent	Two digits: Number	Two digits: Percent	Three digits: Number	Three digits: Percent
0–5	172	73.5	230	98.3	233	99.6
5–15	1	.4	0	.0	0	.0
15–25	11	4.7	3	1.3	1	.4
25–50	19	8.1	1	.4	0	.0
50–100	16	6.8	0	.0	0	.0
100–200	7	3.0	0	.0	0	.0
> 200	8	3.4	0	.0	0	.0

NOTE: Sample size = 234 observations.

unrounded value, and the ε_t's are independent, uniformly distributed random variables that take values between plus and minus one-half of the first-stage rounding precision, δ. In the one-digit rounding case, $\delta = 0.1$.

After the CPI levels are rounded, the difference between two adjacent month's values is calculated as

$$\begin{aligned}\Delta \mathrm{CPI}_t &\equiv \mathrm{CPI}_t - \mathrm{CPI}_{t-1}\\ &= \mathrm{CPI}_t^* + \varepsilon_t - \mathrm{CPI}_{t-1}^* - \varepsilon_{t-1}\\ &\equiv \Delta \mathrm{CPI}_t^* + \Delta\varepsilon_t,\end{aligned}$$

where $\Delta \mathrm{CPI}_t^*$ is defined as the difference between the two unrounded values and $\Delta\varepsilon_t$ is defined as the difference between the two errors.

Next, the percent change is calculated by dividing through by the previous period's CPI value:

$$\frac{\Delta \mathrm{CPI}_t}{\mathrm{CPI}_{t-1}} = \frac{\Delta \mathrm{CPI}_t^*}{\mathrm{CPI}_{t-1}} + \frac{\Delta\varepsilon_t}{\mathrm{CPI}_{t-1}}.$$

Then the resulting percent change is rounded again, this time at the final precision level, α (which is 0.1 percent), yielding

$$\frac{\Delta \mathrm{CPI}_t}{\mathrm{CPI}_{t-1}} = \frac{\Delta \mathrm{CPI}_t^*}{\mathrm{CPI}_{t-1}} + \frac{\Delta\varepsilon_t}{\mathrm{CPI}_{t-1}} + v_t, \qquad (1)$$

where

$$v_t \sim U(-\alpha/2, \alpha/2).$$

Equation (1) shows that the reported CPI inflation figure is the sum of three terms: the true CPI inflation figure,[4] plus the first-stage rounding error scaled by the CPI, plus the second-stage rounding error. The two error terms are qualitatively different. As the level of the CPI increases, the first-stage rounding error matters less and less and $\Delta\varepsilon_t/\mathrm{CPI}_{t-1}$ gets smaller and smaller, while v_t stays the same magnitude. Conversely, as one can see in chart 1, the first-stage rounding term increases in size as the value of the CPI decreases upon tracing it backwards in time. Alternatively, if the Bureau increased rounding precision in the reported CPI, $\Delta\varepsilon_t$ would become smaller and smaller, leaving only the final difference in rounding error between the true and rounded inflation values.

Now the question previously posed—How frequently does the total difference between the reported percent change and the true percent change cause the reported inflation rate to differ?—can be answered analytically. The answer is given by

$$\mathrm{Prob}\left(\left|\frac{\Delta\varepsilon_t}{\mathrm{CPI}_{t-1}}\right| + v_t > \frac{\alpha}{2}\right). \qquad (2)$$

For a given first-stage precision δ and a given desired final precision α, both at a given CPI level, equation (2) can be evaluated by computer simulation. First, two uniform random numbers are drawn from the interval $(-\delta/2, \delta/2)$. Then, the one number is subtracted from the other, and the resulting difference is divided by the CPI value. Next, a third uniform number is drawn from the interval $(-\alpha/2, \alpha/2)$ and is added to the preceding difference, and the resulting sum is compared with $\alpha/2$. If this procedure is now repeated many times, the average number of times that the absolute value of this sum exceeds $\alpha/2$ will be equal to the probability of a first-stage rounding error resulting in an erroneous inflation report. The results of repeating the entire simulation with 1 million repetitions per CPI level are presented in chart 4.

From the chart, it is clear that extending the number of digits of precision at which the CPI is reported will go a long way toward reducing the probability of rounding errors affecting the final result. For example, for an index value of 100, there is a 33-percent chance of a different figure when the series is rounded to one decimal place. The probability drops to 3.3 percent for two digits and 0.36 percent for three digits. For an index value of 50, the chances are 58 percent, 6.7 percent, and 0.64 percent, respectively. For reference, January CPI values for selected years are plotted on the graph. The reported series differs from the benchmark series more than 60 percent of the time prior to 1970.

Note from equation (2), however, that the only way to guarantee a CPI inflation series that is free from rounding error is not to round the CPI levels at all before calculating the percent change. Not rounding corresponds to $\Delta\varepsilon = 0$. Because the second-stage rounding error v_t is between $-\alpha/2$ and $\alpha/2$, the probability that its absolute value exceeds $\alpha/2$ is exactly zero. By contrast, *any* first-stage rounding makes this probability greater than zero. Intuitively, when the second-stage rounding error is very close to being as large as it can be, namely, $\alpha/2$, even a tiny first-stage error can push it over the edge, and the inflation figure will round the wrong way.

As with the real-data experiment presented, the question can be asked, How big are the rounding errors relative to monthly inflation rates for the historical series? The answer is given by an estimate of the distribution of

$$r_t^* = \frac{\left|\Delta\varepsilon_t / \mathrm{CPI}_{t-1} + v_t\right|}{\Delta \mathrm{CPI}_t^* / \mathrm{CPI}_{t-1}},$$

the magnitude of the rounding error relative to the benchmark inflation rate. The numerator can be calculated as in the preceding simulation, while the denominator is the benchmark inflation rate for the month in question. With this equation, the relative percent error can be calculated as a fraction of the

Chart 4. Simulated probability of an inaccurate percent change due to rounding

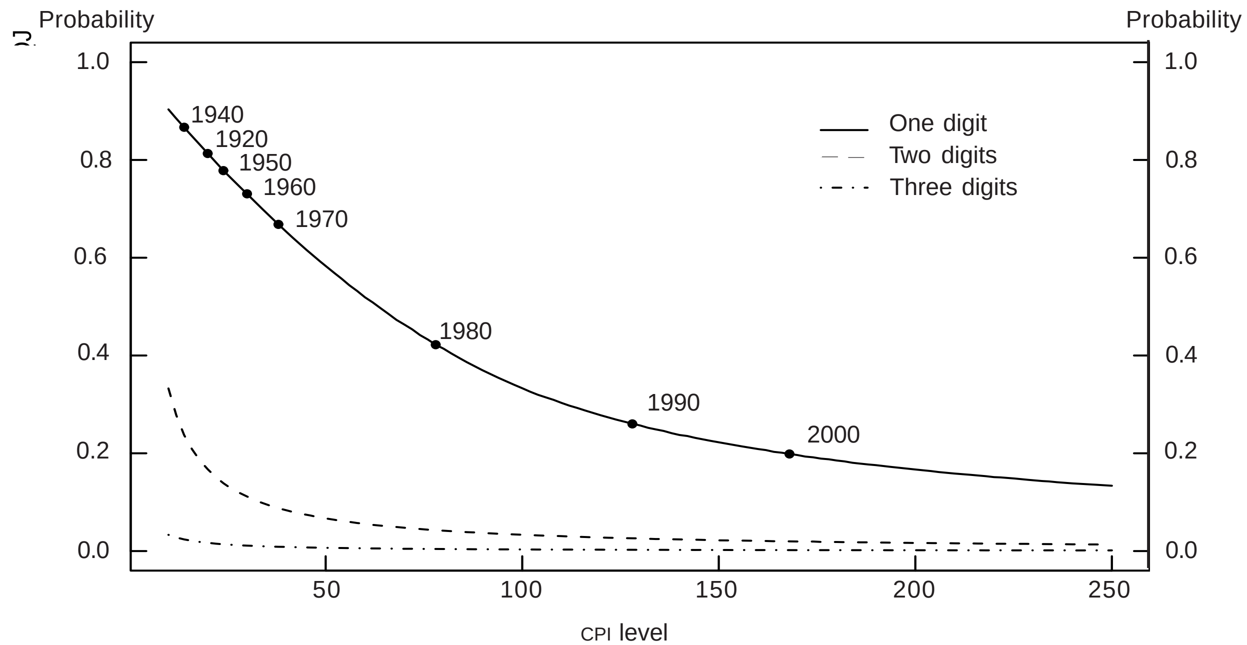

benchmark, as can the number of errors exceeding a given size relative to the true inflation rate.

However, the unrounded inflation rate for the denominator of r_t^* is unknown, and the reported inflation rate is unsuitable, which is indeed part of the motivation for this article in the first place. Because the rounded series contains a (misleadingly) large fraction of months with no change, using the reported inflation series in the denominator results in dividing by zero in many months, exaggerating the relative size of the errors for those months.

For the purpose of showing overall trends in rounding error, a smoothed version of the inflation series was used in place of the "true" inflation rate in the denominator of r_t^*.[5] For each month, 1 million samples are taken from the rounding error distribution corresponding to that month's CPI level, as in the simulation presented earlier. Each of these million simulated errors is divided by that month's smoothed inflation value. The resulting simulated sample of a million values of r_t should approximate the true distribution of the relative errors.

Chart 5 shows the frequencies of different values of the relative error on one plot. The topmost curve represents the probability that $r_t > 0$, or the probability of any discrepancy at all. Consequently, the zero-percent curve corresponds to the simulated results presented in chart 4. For comparison, at the 1990 CPI value of 130, chart 4 shows that inflation rates calculated with the rounded CPI differ from those calculated with an unrounded CPI 25 percent of the time. Correspondingly, in chart 5, the chance of a relative error greater than zero percent in 1990 is approximately 25 percent.

The other curves show the probabilities of errors of various sizes relative to the underlying smoothed inflation rate. For instance, for 1990, errors larger than 5 percent of the true inflation rate occur roughly 15 percent of the time, and errors larger than 15 percent of the true inflation rate occur very seldom—just about 2 percent of the time. In the past—particularly, prior to 1975—not only are the errors much more frequent, but their magnitudes relative to the inflation rate are much larger. For instance, through the early 1960s, the reported inflation rate differs from the benchmark roughly 70 percent of the time, with errors as large as 50 percent of the inflation rate occurring around 45 percent of the time. Perhaps surprisingly, errors larger than the actual inflation rate occur around 20 percent of the time in the period from 1950 to 1968.

Two opposing tensions underlie the probability distribution of the relative size of the discrepancies due to rounding. On the one hand, in modern times the value of the CPI is relatively large, so rounding error should be small as a fraction of the CPI level, and consequently differences should be infrequent. On the other hand, the rate of inflation decreased through the 1990s, so when rounding errors do occur, they would be expected to be larger

Chart 5. Probability of relative errors, by size

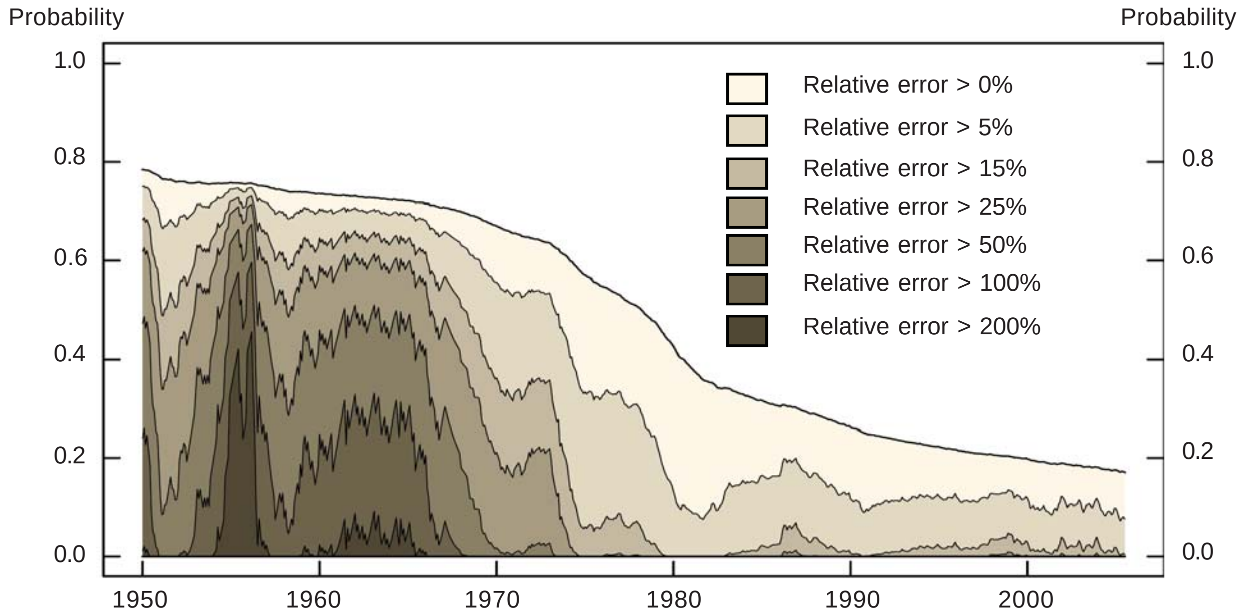

relative to an underlying low rate of inflation. Chart 5 shows how the two forces interact. In 1983, the probability of any error was around 35 percent, while the probability of errors larger than 5 percent of the inflation rate was around 10 percent. Because the CPI level increased over the next two decades, the chance of any error decreased to around 20 percent, but the increase was slow, so the chance of errors greater than 5 percent of the inflation rate remained virtually constant.

Tracing backwards in time, one can see that the two forces act in concert. The smaller CPI values before 1970 lead to a larger (75-percent) probability of a discrepancy, and the periods of relatively low inflation from the 1950s through the 1960s lead to errors that are large relative to the actual inflation rate. Most errors are larger than 5 percent of the inflation rate for that period, with errors larger than the inflation rate itself occurring roughly 20 percent of the time.

The remaining period, between 1970 and the early 1980s, was characterized by two subperiods of high inflation. Consequently, we see the level of the CPI rising and the probability of a rounding error decreasing. At the same time, high monthly inflation rates make the size of errors smaller as a fraction of the inflation rate. One can plainly see the mid-1970s inflation driving the probability of a relative error of 25 percent or greater down from around 20 percent to nearly zero over the course of 2 years. Although the period was hard on the value of a dollar, it was excellent for the accuracy of reported inflation statistics.

THIS ARTICLE DEMONSTRATES HOW ROUNDING A SERIES before calculating its percent changes introduces an additional source of statistical error. Using an unrounded CPI data set as a benchmark reveals that the published values differ from an unrounded benchmark approximately a quarter of the time, and the errors can be large relative to the true underlying inflation rate.

Mathematical analysis and some simulation results demonstrate in more detail how the rounding-induced errors behave with respect to both the level of the CPI and the inflation rate over time. Three regimes emerge: (1) before 1970, both the frequency and magnitude of the errors were large; (2) the inflation of the mid-1970s and early 1980s cut the probability of a rounding error in half and led to moderation in the relative errors; and (3) during the present period, a high CPI value makes the reported inflation rate match the unrounded rate around 75 percent of the time, but the low underlying inflation rate has kept the probability of errors of a given relative size roughly constant and comparatively moderate. These findings certainly have implications for inflation research over the earlier periods.

Finally, the “take-home” message from the real-data experiments presented herein show that increasing the precision of reported CPI levels would go a long way toward making the errors that arise from rounding negligible. Publishing the CPI to three decimal places, for instance, will decrease the error variance of the CPI inflation series by 42 percent and will reduce the likelihood of disagreement between an unrounded index and the reported index from its current 25 percent to under 0.5 percent. □

Notes

[1] "Consumer Prices Jump, Spur Inflation Woes," Reuters, Mar. 24, 2005.

[2] Despite the fact that the 1967-base-year series is more precise, the published inflation series cited in the news and used for official purposes is the rounded 1984-base-year series. Note that, due to rounding discrepancies, the inflation rates calculated from these two published series do not always agree.

[3] Indeed, a large part of the historical series has been converted in just such a manner. However, recreation has been a little more difficult for some parts of the series.

[4] This term is not really the true inflation rate $\Delta CPI^*_t / \Delta CPI^*_{t-1}$, but the difference between it and the true inflation rate is negligible compared with the two rounding errors. For simplicity's sake, the analysis presented here takes $\Delta CPI^*_t / CPI_{t-1} = \Delta CPI^*_t / (CPI_{t-1} + \varepsilon_t)$ as the true inflation rate.

[5] Specifically, the rate used is an exponentially weighted moving average of the reported CPI inflation rate with $\lambda = 0.05$. This rate behaves similarly to a 3-year moving average of inflation. Using the annual inflation rate in the denominator still shows significant rounding effects in the early periods, but otherwise gives results similar to those obtained from the exponentially weighted average.