

The background is a solid teal color. There are four light blue spheres scattered across the page: one in the upper left, one in the middle left, one in the lower left, and one in the lower right. Each sphere has a soft shadow cast to its right.

Close Counts

Expanding the use of estimates for
better, faster financial reporting

BY SHAYNE KAVANAGH
AND MICHELE MARK LEVINE

It might seem strange—or even subversive—to suggest that finance professionals rely more on estimates in financial reporting. But financial reporting already depends heavily on estimates for some of its most consequential numbers. Pension liabilities, Other Post-Employment Benefit (OPEB) obligations, and other long-term measures rely on assumptions about uncertain future events.

Estimates are accepted because they are accurate enough to support decisions. Yet in other areas of reporting, preparers spend a great deal of time trying to achieve near-exact precision—even when that extra precision adds little value for users. This work is costly. It consumes staff time, delays reports, and diverts professional attention from areas where judgment and analysis matter more.

This article argues that governments should use estimates more broadly when they provide sufficient accuracy for decisions at lower cost. The question is not whether estimates are appropriate—financial reporting already depends on them—but when additional precision meaningfully improves a user's ability to make decisions and when it does not.

The report provides a framework for using estimates responsibly and defensibly in financial reporting. We'll focus on accruals because they are common, labor-intensive, and often well suited to estimation. The report shows how accurate estimates can play a valuable role in financial reporting and how estimates can be governed through controls and documentation, and how to determine when detailed accounting work is necessary—and when it is not.

In short, we make the case that governments can maintain accuracy without overspending on unnecessary precision.

Precision vs. Accuracy

The hidden tension that unnecessarily limits the use of estimates

"Precision" and "accuracy" are often treated as synonyms. They are not.

Understanding the distinction helps us see when additional accounting effort improves the usefulness of financial information—and when it does not.

Precision refers to how closely a reported amount reflects the underlying value. Highly precise figures result from detailed measurement, careful verification, and little tolerance for approximation. Accuracy, by contrast, refers to whether information is fit for its purpose—whether it supports the judgments and decisions users make based on financial reports. Reporting unrestricted fund balance as \$34,538,423.21 is precise, but reporting it at \$34.5 million may be just as accurate if it is equally helpful to the decision-maker.¹

Financial reporting relies on accuracy, not precision. Its purpose is to reduce uncertainty about a government's financial condition, performance, and risks. Once uncertainty is reduced enough to support user judgments, additional precision may add little value. To illustrate, if sufficient accuracy can be achieved in two hours but achieving greater precision would take two more weeks, then that precision is probably not worthwhile. This is the diminishing returns of additional precision.


In some areas of financial reporting, though, preparers pursue high levels of precision even when it provides no added benefit to accuracy. For example, staff may devote substantial effort to identifying, classifying, and measuring many small accruals with great exactness—even though modest differences in those amounts would not change a reader's understanding of the government's financial position or its results of operations.


We must therefore consider when more precision improves accuracy. When it does not, the added effort creates an opportunity cost: Staff time and attention are taken up without increasing the usefulness of the information.

For example, preparers give up precision—but not accuracy—when they round figures to the nearest thousand or million dollars. This loss of precision does not lower standards.

Given that less precision does not necessarily mean less accurate or useful reports, our goal is to align the cost of measurement with the informational value of the result. The remainder of this article explains how to do that, using accruals as our primary example.

Reporting unrestricted fund balance as **\$34,538,423.21** is precise, but reporting it at **\$34.5 million** may be just as accurate if it is equally helpful to the decision-maker.





Accruals as the Archetype

Accruals are ubiquitous across governments, labor-intensive, and often have limited influence on user judgments. At the same time, efforts to measure them precisely can be substantial. This combination makes accruals a natural candidate for estimation.

One reason for this is the concentrated nature of their impact. A small number of accruals often account for a large share of the total dollar amount, while many others are individually small. Errors in the larger items are far more likely to affect reported results in a meaningful way than errors spread across many smaller items. This means the effects of estimation error are not evenly distributed.

This has implications for how accounting effort is used. Precision matters most for accruals that are large or volatile, or that strongly influence reported amounts. For these items, additional effort to identify, measure, and verify amounts may meaningfully improve accuracy. For the remaining accruals, however, the same level of effort often yields diminishing returns.

Modest differences in their reported amounts are unlikely to change a reader's understanding of the government's financial condition or results of operations.

Accruals are good candidates for estimation because the underlying economic activity has often already occurred, but invoices, usage data, or final documentation are not yet available at year-end. This makes accruals well suited to estimation techniques that rely on historical patterns, classification, or other simplifying assumptions.

The cost of pursuing near-exact precision for all accruals can be high. Finance staff may spend significant time tracking down invoices, reviewing small transactions one by one, or reclassifying amounts late in the close process.

Accruals, as a category, offer many opportunities to align measurement effort with the informational value of the result. Doing so requires distinguishing between accruals when precision meaningfully improves accuracy and when it does not. This pattern can be explained with a simple but powerful mental model.

The 80/20 Rule

A model for distinguishing when precision improves accuracy—and when it doesn't

In many areas of public finance—including accruals—a small number of items account for a large share of the total dollar amount. The remaining items are individually small and have little influence on reported results. This pattern is described by the well-known "Pareto Principle," also called the 80/20 rule. In many situations, about 20 percent of the causes produce about 80 percent of the effects.² The math behind this pattern provides an important insight. The formula is: $E \sim 1/[1-A]$, where E is effort, and A is accuracy. As accuracy gets closer to 100 percent, the effort required increases rapidly, approaching infinity.

This concentration of impact is common in accruals. Large contracts, significant service arrangements, or major year-end obligations often dominate the total accrual amount. At the same time, many smaller accruals—routine services, minor purchases, or recurring costs—usually account for a much smaller share of the total dollar value.

The implication is that the total contribution of smaller items to decision-relevant information is limited. An imprecise estimate of many small accruals is unlikely to change a user's judgment. By contrast, even small changes in a few large accruals could affect those judgments.

We illustrate the Pareto Principle with two charts. First, Exhibit 1 shows—using data sets that follow this pattern—how much of the total effect [e.g., total reported value] comes from a given share of the items [e.g., percent of items in the data set]. We see that a relatively small share of items—the largest ones—account for most of the total impact. For example, about 40 percent of items account for about 90 percent of the impact. About 70 percent of items account for about 95 percent of the impact.

Next, let's examine Exhibit 2. First, imagine that the reported figure must be at least 99 percent accurate to avoid a material impact. Exhibit 2 shows how much error is allowed in the estimated portion while still reaching 99 percent accuracy overall. We can read the exhibit in three steps:

Step 1: Select a percentage on the horizontal axis to represent the portion of the total the preparer chooses to measure precisely. It makes sense to use the same percentage used for Exhibit 1.

Step 2: That percentage represents how much of the total is known precisely.

Step 3: Move to the vertical axis to see how much error is allowed while still reaching 99 percent accuracy.

For example, if you know 85 percent of the true value with certainty [the horizontal axis], you can make an error of up to 7 percent when estimating the remaining 15 percent and still reach 99 percent accuracy.

Exhibit 2 shows that the allowable error increases sharply as we move to the right. If you select 95 percent on the horizontal axis, the allowable error grows much larger: even a 20 percent error would still allow for 99 percent accuracy. Exhibit 1 showed that measuring 95 percent of the total precisely requires focusing on the largest 70 percent of items—so the remaining 30 percent can be estimated.

To see how this works in practice, we examined a real-life data set from a smaller city to see whether the pattern holds. Using recreation department spending, we assumed a year-end scenario in which payroll, benefits, utilities, and other predictable expenditures were already known or readily accrued through established systems and contracts. For these items, achieving precision is relatively inexpensive. The remaining accounts—those affected by invoice timing and

operational variability—were most likely to require judgmental accruals because obtaining precision for them is more costly.

The results were striking. First, expenditures were highly concentrated at the account level. Across three years of data, an average of 27 percent of accounts accounted for 80 percent of annual costs. These accounts included payroll, benefits, and other categories where precision is relatively inexpensive. Those same accounts made up about 70 percent of costs in period 12. The lower concentration at year-end reflects increased activity among smaller accounts—especially office and operating supplies. (The city also places delayed invoices from earlier in the year in period 12.) Period 12 spending becomes somewhat more spread out, but it does so in predictable accounts rather than randomly across the ledger.

Second, within the remaining accounts that require accrual estimation, period 12 transactions also showed a Pareto pattern. Across all three years, the top 20 percent of invoices accounted for 87 percent of the dollar value. Notably, all of these large invoices were concentrated in a single account: "office and operating supplies," which was just one of 23 accounts analyzed. In other words, the largest accrued payables were located in a clearly identifiable place.

EXHIBIT 1 | Illustration of the Pareto Principle






EXHIBIT 2 | Maximum Allowable Error for Achieving 99% Accuracy in Estimates





In short, the Pareto pattern appeared twice—both across accounts and within individual transactions. When accrual judgments are applied mainly to many small items, rather than to a few large obligations, the practical risk of estimation error can be materially reduced.

To illustrate, assume actual expenditures are known precisely for the first 11 periods. Also assume that payroll, benefits, utilities, and other predictable items in period 12 can be measured with high precision. In this data set, those items account for 97.8 percent of the true annual total. Now suppose the finance team identifies the largest period 12 invoices in the office and operating supplies account—the top 10 percent by dollar value. Even if this effort is only 80 percent effective, the annual total would still be 99 percent accurate—before any attempt to estimate and include the remaining 90 percent of smaller invoices in the total.

These figures are illustrative, and different assumptions would yield somewhat different results. The broader point, though, is that when the Pareto structure is present, there is a lot of room to achieve accurate totals.

Before we move on, we should recognize that not all data sets in financial reporting follow the Pareto Principle. There will be opportunities for estimation in other kinds of data sets, but the most obvious opportunities are likely to be found in Pareto-type data sets, which are the best place to begin expanding the use of estimates in financial reporting.

So far, we have focused on why and how estimates should be used in financial reporting. The next question is how they should be governed.

Accounting Policies and Estimates

There are many opportunities to use estimates more widely in financial

reporting. But how can they be applied in a disciplined, transparent, and consistent way that preserves the integrity of financial reporting?

The answer is good governance. Governance includes the rules (policies) and decision rights about when and how estimates are used.

Governance begins with clarity of purpose. Estimates are used to achieve sufficient accuracy for decisions—not to speed up reporting at any cost or to avoid detailed accounting work when precision is warranted. Estimation is not cutting corners; it is a deliberate choice about how to allocate professional effort.

Materiality plays a central role in this governance framework. An estimate is material if it impacts the decision usefulness of information.

An estimate could have a material impact if the difference between the “true” value and the estimated value could affect a user’s judgment. Let’s walk through an example using accruals. Suppose the largest 40 percent of accrued items account for about 90 percent of the total value. You would not want to estimate those 40 percent. In fact, most data sets may not follow the Pareto Principle—but enough do that the opportunity to apply it to estimation in financial reporting is meaningful. You would measure them precisely. That would give you confidence in 90 percent of the total. Because accruals often follow predictable patterns in where they come from and their size, it may be reasonable to estimate the remaining 60 percent, since any difference between the true value and the estimate is unlikely to affect user judgments. The general policy rule is to be more cautious when the possible difference between an estimate and the true value could change a user’s judgment.

Consistency across periods is another key part of governance. Estimates applied inconsistently can undermine credibility, even if they are reasonable on their own. By contrast, estimates that follow the same general approach from year to year allow users to focus on trends and changes in underlying

What Does Estimation Actually Look Like?

This report explains why we need more estimation in financial reporting—but how do we actually make estimates? Future GFOA content will provide detailed guidance on estimation methods, but preparers can already rely on simple, well-understood approaches:

- **Use recent history as a baseline.** Prior-year accrual amounts can serve as a starting point, especially for routine services or recurring costs. These amounts can be adjusted for known changes such as inflation, contract terms, or workload.
- **Apply known drivers.** For example, utility accruals may be estimated using year-to-date spending patterns, budgeted amounts, or seasonal trends.
- **Validate estimates over time.** Comparing estimated amounts with subsequent actuals helps confirm that the approach remains reasonable.

conditions rather than on shifts in measurement methods. When estimation methods need to change, those changes should be deliberate, documented, and clearly explained. For example, if departments are required to stop using available funds to purchase supplies earlier than in past years, fewer supply invoices are likely to be outstanding at fiscal year-end, which may require adjusting the accrual estimate.

Good documentation explains why and where estimation is appropriate, identifies the assumptions used, and shows how those assumptions relate to the decisions and judgments users make based on the reported amounts.

There are many opportunities to apply estimation to both revenues and expenditures. Some may be universal, or nearly so, across local governments (such as certain accruals). Others may vary. Larger governments may have more opportunities because they process more transactions. As a result, individual estimation errors are more likely to balance out, producing reliable totals even when each item is not calculated with perfect precision. (This is known as the “law of large numbers” in statistics.) That said, many opportunities exist for governments of all sizes. Start with the most promising candidates. As your experience and skill with estimation grow, you can expand to other areas. Data sets that follow the Pareto Principle are especially strong places to begin.

Consultation with auditors is an important part of making materiality determinations, but it should not be the starting point. Auditors assess whether judgments are reasonable. Clear communication about where estimates are used, why they are reasonable, and the government’s policies helps align expectations and reduce the risk of disputes. Responsibility for materiality judgments rests with preparers. Auditors evaluate those judgments. A clear governance framework gives preparers a strong foundation for these conversations.

Finally, governance helps determine when reevaluation is needed. Estimates that work under stable conditions may become less appropriate when circumstances change. Significant growth, greater volatility, policy changes, or shifts in operations may require more precision or different estimation approaches. Governance sets clear expectations for when estimates should be revisited.

Taken together, these elements—purpose, materiality, consistency, documentation, auditor coordination, and reevaluation—create a governance framework that makes the use of estimates both defensible and sustainable.

The next section turns to a concern raised about estimation: that many individually immaterial estimates could combine to create a materially misleading result.

Death by 1,000 Immaterial Cuts?

The risk of a cumulative material misstatement

A common concern about the use of estimates is that many individually small estimates could combine to produce a materially misleading result. This concern is legitimate, but that doesn’t mean we must return to the detailed accounting work that estimation is meant to avoid. Instead, we can manage the risk of the “1,000 cuts” by applying well-established principles about how estimation error behaves in the aggregate.

- Be consistent in the process and document it to show people how the estimates are arrived at.
- Observe rules for prudent estimates within any given category of items; if your categories are estimated well, you will be fine at the fund level.
- Open communication with auditors about what you are doing and where.
- Understand that small and large governments may have different opportunities.

The first principle is that estimates should strive to be objective, not systematically conservative. While conservatism may feel prudent for an individual estimate, applying conservative assumptions across many items creates a bias in the estimate that stacks up. Every estimate will be wrong to some degree—it is simply a question of how much and in what direction. If each estimate is wrong in the same direction, the cumulative effect can be materially misleading, even when each estimate appears reasonable on its own. Objective estimation—aimed at the best representation of actual value—reduces this risk and improves the reliability of aggregate results.

A second principle is that estimation works best when applied to many relatively small items. By limiting estimates to large numbers of small items, we take advantage of what is known as the “law of large numbers.” This principle holds that when many items are involved, an

When estimates are objective, diversified, and applied to a large number of small items, the range of plausible variation in the aggregate can often be shown to be well within thresholds that would affect those judgments.

error in one is more likely to be offset by an error in the other direction in another. As a result, the total tends to move closer to the true amount. In practical terms, this means that while any single estimate may be rough, the combined effect of many independent estimates is more stable than intuition might suggest.

This aggregation effect is strengthened by diversification. When estimates are based on different transactions, vendors, departments, or timing patterns, their errors are unlikely to move in the same direction. Positive and negative differences tend to offset each other, creating a limited range of likely outcomes. It is this diversification effect—not excessive conservatism or exhaustive verification—that helps limit cumulative risk in well-designed estimation approaches.

Finally, cumulative materiality does not mean that differences from one year automatically add to differences from prior years. In most cases, such differences will self-correct over time as periods accumulate. For example, if too much asset value is expensed each year instead of being capitalized, the resulting understatement of depreciation expense in later years would offset the difference. The difference would exist only for the useful life of the uncapitalized assets. Moreover, expensing newly acquired (uncapitalized) assets each year would likely produce a relatively consistent difference from year to year unless the government's small asset acquisition pattern varies significantly.

These principles explain why cumulative risk should be evaluated at the level of user judgments, not individual transactions. Financial statement users care about whether totals, trends, and key indicators are reasonably stated,

not whether every underlying amount is exact. When estimates are objective, diversified, and applied to a large number of small items, the range of plausible variation in the aggregate can often be shown to be well within thresholds that would affect those judgments. This applies to both fund-level and government-wide figures. In either case, remember that if prudent estimation methods and governance are applied to estimating any category of financial information, you will likely arrive at a figure that is materially accurate for that category. If all the figures that make up a total (fund level or government-wide) are materially accurate individually, the total will also be materially accurate.

None of this removes the need for safeguards. Concentration of impact still matters. Large or volatile items require greater precision, and changes in scale, volatility, or operating conditions may call for reevaluation. But managing cumulative risk doesn't require repeating full accounting work year after year. Instead, it requires aligning estimation practices with the numerical realities that govern aggregate behavior.

Conclusion

Financial reporting is not an exercise in exact precision; its purpose is to support informed decisions about a government's financial condition, performance, and risks. Estimates play a central role in fulfilling that purpose. The question addressed in this report is not whether estimates belong in financial reporting, but how to use them deliberately and responsibly.

The framework presented here rests on a simple premise: Accuracy for decision

making does not require maximal precision everywhere. Beyond a certain point, additional measurement effort often yields diminishing returns. When precision is applied in areas where it does not meaningfully reduce uncertainty for users, it consumes professional time and attention that could be better used elsewhere.

None of this diminishes the importance of rigor, controls, or professional judgment. On the contrary, effective use of estimates requires clarity of purpose, thoughtful application of materiality, consistency across periods, and a willingness to reevaluate when conditions change. Estimation is not cutting corners; it is a design choice about how to allocate limited resources in service of decision usefulness.

Ultimately, strong financial reporting provides information that supports understanding and accountability. It also delivers that information as quickly as practical. By aligning measurement effort with decision relevance, governments can strengthen the value of financial reports while reducing unnecessary cost and delay. The framework outlined in this report provides a foundation for doing so. ■

Shayne Kavanagh is GFOA's senior manager of research. **Michele Mark Levine** is director of GFOA's Technical Services Center.

¹ A good argument can be made that, in many cases, \$34.5 million may actually be more helpful to decision-makers than a more precise figure. Humans have limited working memory, and rounding reduces unnecessary cognitive load. This can make it easier for decision-makers to compare the fund balance with other figures, such as revenues or expenditures.

² The 80/20 rule comes from a universally accepted statistical concept called the Pareto distribution, or power law distribution. Like the more famous bell curve, or normal distribution, the Pareto distribution applies to many types of phenomena across different disciplines.