Hazen Technical Memorandum

Bay Area Water Supply & Conservation Agency

June 12, 2026

To: Danielle McPherson, Senior Water Resources Specialist

From: Jack Kiefer, PhD
Luke Wang, PE
Devon Becker, PE

QC: Andrea Zimmer, PhD, PE

cc: Tom Smegal, General Manager
Tom Francis, Water Resources Manager

2025 Regional Demand Study: Methodology to Incorporate Future Droughts

Technical Memorandum

Job no 20095-017

Table of Contents

Executive Summary

This technical memorandum presents a statistical analysis of historical drought restriction episodes in the Bay Area Water Supply and Conservation Agency (BAWSCA) region, with the objective of improving understanding of how water demand responds to drought conditions and how those responses may persist after restrictions are lifted. The analysis supports ongoing demand planning by providing a structured framework for representing drought impacts in future scenarios.

Study Design and Goals

The analysis examined how water use changed during and after three recent droughts (2006-2009, 2014-2017, and 2021-2023), using historical data across the region. The goal of the analysis was not to produce a new water demand forecast, but to better understand:

- How much demand drops during drought restrictions;
- Whether demand reductions persist after drought restrictions are lifted; and
- How seasonal demand patterns (especially summer peaks) are affected.

Key Findings

Several key findings emerged from the analysis including:

- **Drought restrictions can create lasting reductions in demand.** The strongest evidence comes from the 2014-2017 drought, where water use dropped significantly and did not fully return to pre-drought levels after restrictions ended. This suggests that drought restrictions can permanently change customer behavior (e.g., landscaping or water-use habits).
- **Impacts are strongest in the residential sector.** Single-family residential use shows the most consistent and lasting reductions. By contrast, overall system-wide metrics (like total per capita demand) can obscure these changes because they combine many different types of users. This implies that single family residential behavior is a primary driver of long-term demand shifts.
- **Peak (summer) demand is reduced and may stay lower.** Drought restrictions disproportionately reduce outdoor water use, leading to smaller summer peaks. In some cases, these reduced peaks persist after the drought restrictions have been lifted.
- **Not all droughts have the same impact.** The magnitude and persistence of demand reductions vary across droughts. As discussed above, restrictions during the 2014–2017 drought showed clear and lasting impacts. Conversely:
  - Restrictions during the 2006-2009 drought did not significantly reduce water demand.
  - Restrictions during the 2021-2023 drought had a less significant reduction in demand as compared to the 2014-2017 drought and demand appears to be “rebounding” at a faster pace.

Planning Implications

Overall, the analysis shows that drought restrictions can produce lasting changes in both the level and seasonal structure of water demand, particularly in the residential sector. However, the magnitude and persistence of these impacts vary across drought events, and the timing and severity of future restrictions remain uncertain and closely tied to supply conditions. Given this uncertainty, future drought impacts are best evaluated through a scenario-based planning framework rather than embedded in a single deterministic forecast. This memorandum outlines a practical drought-scenario framework and identifies key opportunities to enhance future demand studies.

1. Introduction

The BAWSCA region has endured multiple droughts over the last quarter-century. Each drought period was accompanied by the implementation of water use restrictions, requesting residents and businesses to reduce water consumption, particularly outdoor uses associated with irrigation.

BAWSCA exercised an optional task of the 2025 Demand Study (Task 12) to statistically evaluate the dynamics of these drought episodes, with specific focus on restriction-induced demand reductions and whether and to what extent demand reductions persisted or rebounded once restrictions were lifted. The purpose of the analysis was to enhance understanding of drought impacts to support long-term planning and scenario development.

The recent completion of the 2025 Demand Study presented the opportunity to evaluate regional demand trends through the lens of drought-management. Historical water use, prices, socio-demographic, and conservation data collected and modeled for that study provided the foundation for this analysis.

Unlike the 2025 Demand Study, which derived tailored demand projection models for each BAWSCA member agency, this analysis focuses on a regional characterization to infer the general impacts of implemented drought restrictions. The analysis involves statistical modeling of two regional indicators of unit water use: Average single-family use per account per day and total use per capita per day, including all sources of water and all retail customer classes. The analysis finds evidence that restrictions associated with certain drought episodes—particularly the 2014–2017 event—produced lasting reductions in demand levels and seasonal variability.

2. Conceptual Relationship Between Drought Restrictions and Water Demand

Evaluating the effects of drought on the demand for water involves many aspects and deserves nuanced treatment. For example, from a hydrological perspective, droughts are related to dry spells that decrease soil moisture, diminish streamflows, and affect water supply storage conditions. From a water demand perspective, less precipitation and lower soil moisture (perhaps also accompanied by warmer air temperatures) are expected to increase water use. If water demand is unconstrained during droughts, water suppliers face a risk of not having enough supply available for all users. Consistent with the California Water Code, urban water suppliers are authorized to declare shortage conditions and may respond with rationing and restrictions on water use.

Water use restrictions for urban water suppliers in California are implemented within a statutory framework established by the Water Code through required Water Shortage Contingency Plans (WSCPs). These plans incorporate six standardized shortage levels— ranging from shortages of up to 10 percent to shortages exceeding 50 percent —that correspond to progressively more severe shortage conditions and increasingly stringent response actions. For each shortage level, urban water suppliers must identify locally appropriate shortage response measures, which may include supply augmentation, water use restrictions, enforcement mechanisms, penalty structures, communication protocols, operational changes, and other demand reduction actions.

While the WSCP framework provides a consistent statewide structure, urban water suppliers retain significant discretion in how they design and implement actions within each stage. As a result, specific restrictions, enforcement practices, penalties, and customer outreach measures vary widely across agencies, and in some cases, agencies face legal or policy constraints on their ability to levy monetary penalties on consumers. Although suppliers typically determine when to activate shortage stages based on supply availability assessments, the State has, during periods of severe drought, exercised its authority to mandate specified conservation actions or restrictions statewide, effectively standardizing certain drought response measures across agencies.

At the end-user / customer level, reactions and responses to restrictions depend on a broader context, involving water using technology and behaviors and myriad factors that influence choices and attitudes. In other words, water use can vary based on many factors, not solely from the implementation of drought management policies.

Figure 1 provides the conceptual framework employed herein to estimate the impacts of drought restrictions. Prior to the imposition of restrictions, there is a trajectory of water use driven by demographic and land use trends, pricing and price levels, and on-going water efficiency. The assumption is that once drought restrictions are imposed, demands decrease relative to some preceding level, which can be thought of as a reference baseline for estimating conservation induced by the restrictions. Once restrictions are lifted, water use may or may not recover to the pre-restriction trend, but will settle to some "equilibrium" path, at least until some future shock occurs, including the possibility of future drought restrictions.

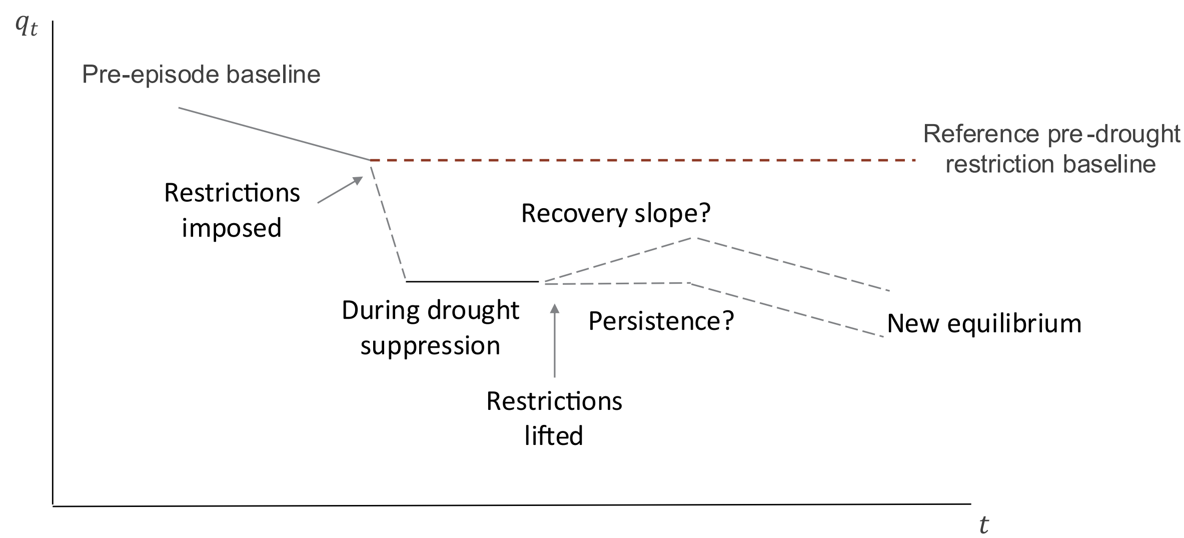

Figure 1 Conceptual Framework for Evaluating Effects of Drought Restrictions

This conceptual framework relates directly to the statistical specification used in this analysis, where:

1) In-episode indicators capture immediate suppression,

2) Post-period indicators capture persistent level shifts, and

3) Interactions with seasonal terms allow for changes in the amplitude and timing of seasonal demand.

3. Statistical Analysis of Historical Drought Restrictions and Water Demand

The historical modeling database developed for the 2025 Demand Study[1] served as the basis for this drought analysis. The existing BAWSCA water demand model dataset, derived from billing data provided by member agencies presents the most complete, consistent, comprehensive historical dataset available in the BAWSCA region. The dataset includes standardized historical customer class water use data, as well as water produced, for each BAWSCA member agency, and, as described in the 2025 Demand Study Final Report, formed the basis of BAWSCA member agency forecast models and regional projections. Consistent with sector models development for the 2025 Demand Study, data spanning the period 2000-2023 was used for this analysis.

Unlike the modeling described in the 2025 Demand Study, this analysis focuses on regional estimates of water use and not water use for individual BAWSCA member agencies. A more simplified approach to evaluating variability in regional demands is useful for testing novel statistical approaches, without introducing additional variability in individual member agency consumption data. Total water use per capita and regional average water use per single-family account were selected as the primary water use metrics for the analysis. The single-family class is the most homogenously defined model sector, where one account generally serves a single household. Total per capita use, while lacking sector detail, is a standard gauge of overall regional trends. The use of both SF use per account and total per capita demand allows for comparison between a behaviorally focused residential metric and a system-wide aggregate metric.

3.1 Derivation of Regional Estimates of Unit Rates of Water Use

For the 2025 Demand Study, single-family (SF) water use per account for each BAWSCA member agency (smoothed to account for differences in billing frequency) used in the estimation of the long-term projection models. This analysis weighted the smoothed data by the annual average number of accounts in each agency to derive a regional estimate of SF use over the 2000-2023 time frame. The weighting was applied to available information for each historical monthly time period, such that the regional average rate of use reflected the same screening undertaken in the development of the single-family models employed in the 2025 Demand Study to address missing data or suspicious values.

Member agency-specific estimate of total per capita use (GPCD) were derived by dividing total water production in each month for each BAWSCA member agency (inclusive of all water using sectors and supply sources) by corresponding population estimates derived from the California Department of Finance and member agency records for the 2000-2023 time frame. An overall regional average of total per capita use was then derived by weighting available member agency per capita use data by member agency population. Unlike the single-family class data, the per capita dataset did not require smoothing to

[1] BAWSCA, 2025 Regional Water Demand and Conservation Projections Study Final Report, prepared by Hazen and Sawyer.

account for customer billing cycles and was complete without any screened or missing values prior to averaging up to the region.

3.2 Specification of Drought Restriction Episodes

The analysis defines three historical drought restriction episodes inferred from water restriction histories used in the 2025 Demand Study, including restrictions imposed by SFPUC and the State.[2] Based on these data, the timing and span of three episodes are defined as:

- Episode 1: Starting April 2006 and ending in June 2009
- Episode 2: Starting January 2014 and ending in April 2017
- Episode 3: Starting August 2021 and ending in March 2023

According to historical data, not all BAWSCA member agencies were subject to the exact same timing or severity of restrictions.[3] Since the regional single-family data were derived from the same screened dataset used in the 2025 Demand Study, for simplicity the single-family analysis defines the episodes as binary variables according to the dates above. The total GPCD analysis employs population weighted binary variables, which accounts for the possibility that not all agencies were subject to restrictions uniformly. In the population weighted binary framework, a variable equal to 1 indicates all agencies were subject to a restriction, while a value lower indicates that fewer than all agencies were subject to restrictions.

The defined span of drought restriction episodes permitted the development of additional variables to aid in examining persistence and/or recovery from restriction-induced changes in water use. Additional variables were constructed to define post-restriction periods for each episode; hereafter abbreviated as POST1, POST2, and POST3. Each post-restriction period begins from the time period when restrictions were lifted until the next restriction episode begins. Furthermore, the number of months since restrictions were lifted is defined for each episode; hereafter abbreviated as Months-Since-Lifting (MSL) 1, MSL2, and MSL3. These variables serve as time counters that measure the cumulative number of months from the beginning of each post-restriction period until the next episode.

3.3 Supplemental Statistical Controls

In addition to the core drought episode and post-period indicators, the models incorporate a set of supplemental variables to account for short-term variability and longer-term structural changes in water use. These variables serve two primary purposes: (1) to isolate drought-related effects from other drivers

[2] The data compiled for the 2025 Demand Study identified the timing and requested water use reductions associated with SFPUC voluntary and SFPUC Tier 2 restrictions, as well as State required percent reductions.

[3] The variation of modeled drought restrictions among BAWSCA member agencies is the due to the formulation of SPFUC Tier 2 restrictions, which included agency-specific variables such as seasonal use of all available water supplies and residential per capita use, as well as the formulation of State required reductions, which varied depending on agency-specific residential per capita use.

of demand, and (2) to provide a stable baseline against which restriction-induced changes can be evaluated.

Short-term variability in demand is captured through weather variables, including departures from long-term normal temperature and precipitation. These variables account for fluctuations in outdoor water use driven by climatic conditions and help distinguish weather-related variation from policy-induced demand reductions. The same weather data for each BAWSCA member agency used in the 2025 Demand Study were weighted by population to create the regional climatic representation.

Systematic seasonal variation is represented using a set of annual harmonic terms, which provide a smooth and continuous characterization of intra-annual demand patterns. Interactions between these seasonal terms and drought-restriction indicators allow for changes in seasonal amplitude and timing associated with restrictions.

To account for broader economic and structural influences, the models include either (a) explicit socioeconomic variables – such as price, conservation indices, and economic indicators – or (b) a reduced-form linear time trend. Historical real prices and estimates of passive conservation used in the 2025 Demand Study for each BAWSCA member agency were either account- or population-weighted to derive regional measures. The trend-based specification, which is the primary focus of this analysis, absorbs gradual changes in water use associated with long-term efficiency improvements, evolving customer behavior, and unobserved structural factors. This approach reduces potential multicollinearity among regional characterizations of price and passive efficiency and provides a more stable framework for isolating and drawing inferences about drought-related effects.

Finally, a COVID indicator and associated seasonal interactions are included to capture structural disruptions to water use patterns during and following the pandemic period.[4] These variables account for shifts in both overall demand levels and seasonal usage patterns that are not directly attributable to drought conditions.

3.4 Model Estimation and Fit

All models were estimated using ordinary least squares (OLS) applied to monthly data over the period 2000 through 2023. This estimation window captures the three distinct drought episodes and corresponding post-restriction periods, providing a basis for identifying both short-term demand suppression and longer-term persistence effects. Given the time-series nature of the data, standard errors are computed using Newey–West heteroskedasticity and autocorrelation consistent (HAC) estimators to account for potential serial correlation and ensure valid statistical inference. The OLS framework provides a transparent and interpretable structure for estimating average effects and interaction terms, which is well-suited for statistical decomposition and policy interpretation. Furthermore, the monthly frequency allows for identification of both seasonal dynamics and the timing of restriction impacts.

[4] The COVID indicator assumes a value of 1 from March 2020 to the end of the historical modeling period, so that any lingering pandemic effects are addressed independently from lingering effects of drought restrictions.

Appendix A provides a description of model variables and tabular estimation output, including model coefficients and diagnostic information. Meanwhile, Figure 2 and Figure 3 compare observed and model-predicted demand for the SF and GPCD series, respectively.[5] The fitted values closely track observed demand over the full study period, capturing both the timing and magnitude of drought-related reductions as well as subsequent recovery dynamics. With achieved R^2 values of more than 95 percent and mean absolute percentage errors of under 4 percent, the models reproduce key features of the data, including structural declines in demand and changes in seasonal variability, providing confidence in their use for interpreting drought-related effects.

While the models provide a strong statistical fit to historical data, they are not intended for direct use as forecasting models. In particular, the inclusion of a reduced-form time trend is designed to absorb long-term structural changes within the estimation period, rather than to represent a causal or policy-invariant driver suitable for extrapolation.

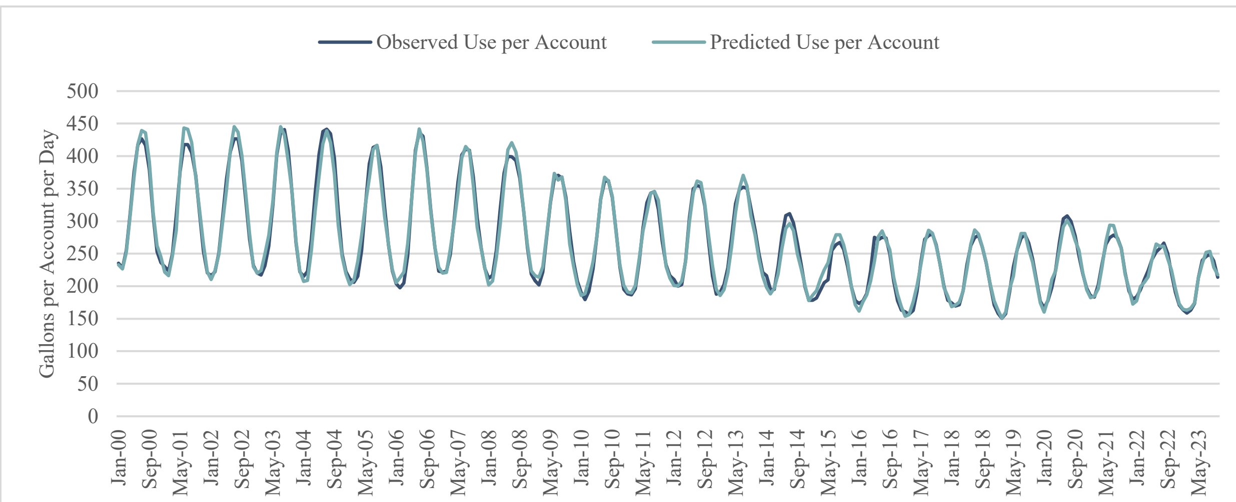

Figure 2 Model Predicted versus Observed Regional Single-Family Use per Account

[5] Specifically, the predictions displayed in Figure 2 and Figure 3 are obtained from model coefficients shown in Table A-3 and Table A-5 of Appendix A.

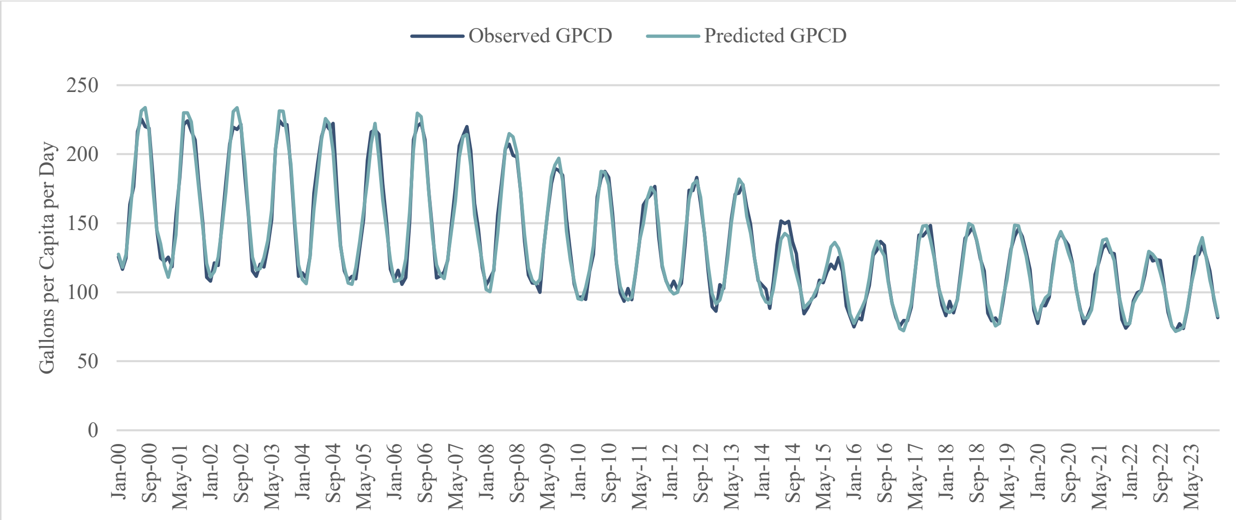

Figure 3 Model Predicted versus Observed Regional Total GPCD

Instead, the primary purpose of these models is inferential—to isolate and quantify the magnitude, persistence, and seasonal characteristics of demand responses to drought restrictions. For forecasting applications, these effects are more appropriately applied as adjustments or overlays to baseline projections derived from sectorally disaggregated demand models.

3.5 Decomposition Framework

While the estimated coefficients in Appendix A provide a complete statistical description of the models, direct interpretation is complicated by the interaction of multiple components operating simultaneously over time. In particular, observed demand reflects the combined influence of long-term structural trends, short-term weather variability, and drought-related effects that alter both the level (average) and seasonal pattern of use.

To facilitate interpretation, the model is reformulated as a decomposition into additive components in logarithmic space (or multiplicative components in level terms). This representation separates baseline conditions, weather-driven variability, and drought-seasonal effects, allowing each to be examined independently. The resulting framework provides a clearer lens through which to interpret how drought restrictions influence both the magnitude and temporal structure of water demand. This decomposition is particularly useful for isolating persistence effects, as it distinguishes between temporary fluctuations and structural changes that remain after restrictions are lifted.

The estimated models can be interpreted as decomposing demand into three components:

$$\widehat{q_t} = B_t * W_t * DS_t \qquad \text{Equation 1}$$

where $\hat{q}_t$ is a prediction of unit water use (single-family water use per account or total gpcd) and t represents time (year-month). The components are defined as:

- B (Trend / Structural factors): long-term baseline trajectory driven by trend, economic conditions, and COVID effects

- W (Weather): short-term variability driven by temperature and precipitation anomalies

- DS (Drought-Seasonal): component capturing both drought-restrictions and seasonality, including base seasonality policy-induced deviations, including:
    - In-episode suppression
    - Post-episode persistence
    - Changes in seasonal amplitude

This decomposition allows drought effects to be interpreted as structural shifts in seasonal amplitude and persistence, separate from weather-driven variability and long-term baseline change.

Figures 4 and 5 illustrate the decomposition of predicted demand into the baseline (B), weather (W), and drought-seasonal (DS) components for the GPCD and SF models, respectively. Multiplying the three components exactly reproduces the predictions shown in Figure 2 and Figure 3.

The figures highlight the distinct roles of each component. The baseline component (B) captures the long-term downward trajectory in demand, reflecting structural factors such as efficiency improvements and broader economic influences. The weather component (W) introduces short-term variability, while the drought-seasonal component (DS) captures both the sharp reductions during drought episodes and the subsequent persistence and reshaping of seasonal demand patterns.

Notably, the DS component shows both level shifts and reductions in seasonal amplitude during drought periods, with partial persistence of these effects in post-episode periods.

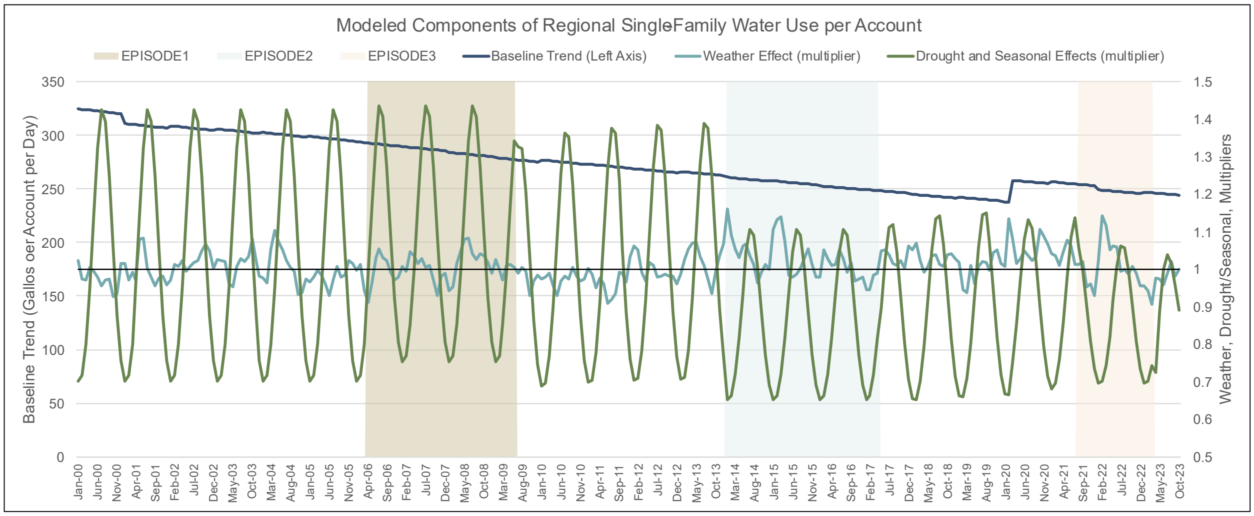

Figure 4 Estimated Components of Regional Single-Family Model

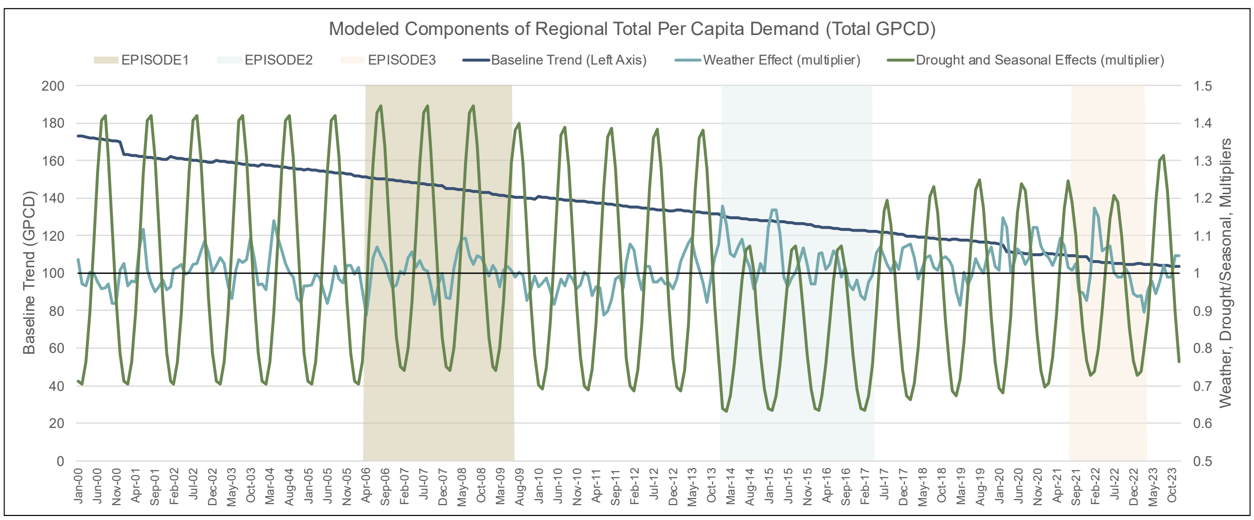

Figure 5 Estimated Components of Regional Total GPCD Model

The decomposition framework separates long-term baseline trends, short-term weather variability, and drought-related structural effects. Having demonstrated that the model reproduces both level shifts and seasonal reshaping observed in the data, the drought-seasonal (DS) component can be used as a multiplicative adjustment factor to isolate and quantify restriction-induced suppression, persistence, and recovery dynamics. The following section interprets these effects across drought episodes.

4. Model Interpretation and Analysis of Drought Restriction Events

This section summarizes the key findings from the regional demand models, with emphasis on specifications that include a linear time trend in place of explicit price and conservation variables (specifically the models of Table A-3 and Table A-5 of Appendix A). These models provide a more stable and interpretable framework for evaluating drought-related dynamics by absorbing long-term secular changes into a single trend term, thereby reducing collinearity and isolating the effects of drought episodes and post-restriction behavior.

The results are organized around three dimensions of drought response:

- During-drought suppression,
- Post-drought persistence and recovery, and
- Changes in seasonal demand patterns.

Across these dimensions, a consistent pattern emerges: the 2014–2017 drought (Episode 2) provides the most robust and system-wide evidence of demand reduction and persistence, while the 2021–2023 drought (Episode 3) shows comparable magnitudes in the single-family sector but with greater uncertainty due to a shorter post-period and overlapping structural changes. Accordingly, Episode 2 serves as the most reliable benchmark due to its consistency across both models and the longer period available for evaluation.

4.1 During Drought Effects

Table 1 summarizes the estimated average during-drought effects across the three episodes for both the single-family (SF) and total per capita (GPCD) models. The reported values translate the estimated coefficients into percentage changes relative to baseline conditions, allowing for direct comparison across episodes and model specifications.

The estimated coefficients for the drought episode indicators reflect the average reduction in water use during periods when restrictions were in place. Across both the single-family (SF) and total per capita (GPCD) models, the magnitude and statistical significance of these effects vary notably by episode.

Episode 2 (2014–2017) exhibits the most consistent and statistically significant reduction in demand across both models.

- In the SF model, the estimated coefficient implies a reduction in use on the order of approximately 15–20 percent relative to baseline conditions.
- The GPCD model shows a similar directional effect, though somewhat attenuated in magnitude due to aggregation across sectors.

Episode 3 (2021–2023) also shows evidence of demand suppression, particularly in the SF model. However, the aggregate (GPCD) response is smaller and less precisely estimated, reinforcing that system-wide effects are less certain for this episode relative to Episode 2. While Episode 3 exhibits similar magnitudes in the SF model, Episode 2 remains the most reliable benchmark due to consistency across both models and a longer post-period for evaluation.

Episode 1 (2006–2009) presents weak or counterintuitive results, with coefficients that are not statistically significant and in some cases positive. This likely reflects a combination of milder restrictions, less consistent implementation, or confounding influences from the Great Recession.

Table 1 Model Estimates of During-Drought Effects

Model	Episode	Estimated average demand change*	Statistical support**	Interpretation
SF	Episode 1	+4%	borderline	Not statistically distinguishable from zero
SF	Episode 2	−15%	strong	Large drought-period suppression
SF	Episode 3	−14%	strong	Large drought-period suppression
Total GPCD	Episode 1	+3–4%	weak	Counterintuitive aggregate result
Total GPCD	Episode 2	−18%	strong	Major system-wide reduction
Total GPCD	Episode 3	−6–7%	weak	Moderate aggregate reduction

*Percent effects are computed as 100[exp(β)−1]. For GPCD, episode effects are scaled by the restriction proportion variable.
** Statistical support refers to conventional significance levels ($p < 0.10, 0.05, 0.01$) based on HAC standard errors.

4.2 Persistence and Recovery

A central objective of the analysis is to determine whether demand reductions observed during drought conditions persist after restrictions are lifted. This is captured through the POST indicators and the associated MSL terms.

4.2.1 Level Shifts Following Restrictions

The POST coefficients indicate whether water use remains below its pre-drought trajectory after restrictions are lifted.

In the SF model, the post-Episode 2 and post-Episode 3 indicators are both negative and statistically significant, indicating that a portion of the drought-induced reduction persists beyond the end of restrictions. The post-Episode 3 effect is particularly large in magnitude; however, consistent with earlier discussion, this estimate should be interpreted with caution given the short post-period and overlapping structural changes.

In contrast, the GPCD model shows weaker and less consistent post-period effects, particularly for Episode 3, where coefficients are not statistically significant. This suggests that persistence at the aggregate level is more muted or obscured by variability in non-residential demand components.

4.2.2 Recovery Dynamics

Post-drought demand recovery is represented by the post-period indicator and its interaction with the natural logarithm of months since lifting (MSL) restrictions. For episode j, the implied post-period effect at month m is given by Equation 2:

$$\Delta_j(m) = 100 * \left[exp\left(\beta_{POST,j} + \beta_{MSL,j} * \ln(1+m)\right) - 1\right] \qquad \text{Equation 2}$$

The estimated number of months for half-recovery and return-to-baseline are calculated using Equation 3 and Equation 4, respectively.

$$m_{\frac{1}{2}} = exp\left(-\frac{\beta_{POST,j}}{2\beta_{MSL,j}}\right) - 1 \qquad \text{Equation 3}$$

$$m_{base} = exp\left(-\frac{\beta_{POST,j}}{\beta_{MSL,j}}\right) - 1 \qquad \text{Equation 4}$$

The implied half-life corresponds to the number of months required for the initial post-period reduction to be reduced by 50 percent, based on the estimated recovery slope.[6]

Persistence and recovery dynamics are summarized through the estimated post-period level shifts and recovery slopes, which together describe both the magnitude and duration of drought-induced reductions (**Error! Reference source not found.**).

The MSL interaction terms provide insight into the rate at which demand recovers over time:

- In the SF model, the post-Episode 3 MSL coefficient is positive and statistically significant, indicating gradual rebound following the initial drop in demand, but not a full return to pre-drought levels within the observed time frame.
- For Episode 2, the recovery effect is weaker and not consistently significant, suggesting either slower recovery or stabilization at a lower level.
- In the GPCD model, recovery dynamics are generally weaker and less precisely estimated, reinforcing the interpretation that persistence effects are more clearly identifiable in the residential sector.

Taken together, these results suggest that drought-induced reductions in demand are only partially reversible, with evidence of both immediate persistence (level shifts) and gradual recovery over time. However, the recovery metrics of

[6] Note that these expressions are meaningful only when the post-period effect is negative and the estimated recovery slope is positive.

Table 2 are derived from the estimated functional form and should be interpreted as indicators of persistence rather than literal forecasts of recovery duration.

Table 2 Model Estimated Recovery Dynamics

Model	Episode	Immediate post effect	Implied half-recovery	Implied return to baseline	Interpretation
SF	Episode 1	Moderate rebound	~9 months	~8 years	Gradual recovery
SF	Episode 2	Large persistent effect	>30 years	effectively none	Persistent drought legacy
SF	Episode 3	Rapid rebound	~2 months	~7 months	Short-lived effect
Total GPCD	Episode 1	No clear recovery path	—	—	Weak persistence
Total GPCD	Episode 2	Slow rebound	~28 months	~70 years	Long-term legacy
Total GPCD	Episode 3	Discrete event	—	—	No clear post path
Note: Recovery metrics reflect the implied behavior of the estimated functional form and should be interpreted as indicators of persistence rather than literal forecasts of recovery duration.					

4.3 Changes in Seasonal Demand Patterns

In addition to affecting overall demand levels, drought restrictions appear to influence the seasonal structure of water use, as captured by interactions between drought indicators and the seasonal harmonic terms (S1 and C1).

The annual seasonal cycle is represented by the first-order harmonic terms S1 and C1. The baseline magnitude of seasonal variation is summarized by the harmonic amplitude (A), which refers to the coefficient estimates for the annual harmonics (β_{S1} and β_{C1}):

$$A_{base} = \sqrt{\beta_{S1}^2 + \beta_{C1}^2} \qquad \text{Equation 5}$$

For drought or post-drought regimes, the adjusted seasonal amplitude is computed by adding the relevant interaction terms:

$$A_j = \sqrt{\left(\beta_{S1}^2 + \beta_{S1,j}^2\right) + \left(\beta_{C1}^2 + \beta_{C1,j}^2\right)} \qquad \text{Equation 6}$$

where j denotes the episode or post-episode period and ($\beta_{S1,j}$ and $\beta_{C1,j}$) represent the coefficient estimates for the periods defined by the POST variables. In percentage terms, change in amplitude relative to the baseline is calculated as:

$$\%\Delta A_j = 100 * \left(\frac{A_j}{A_{base}}\right) \qquad \text{Equation 7}$$

4.3.1 Seasonal Compression During Drought

During Episode 2, both models show statistically significant interactions with seasonal terms, indicating that peak-season (summer) demand is reduced disproportionately relative to off-season use, resulting in a compression of seasonal amplitude. This pattern is consistent with the targeting of outdoor irrigation during drought restrictions.

Table 3 summarizes the estimated seasonal amplitudes across baseline, in-episode, and post-episode periods for both the SF and GPCD models. As shown in Table 3, both models exhibit clear evidence of seasonal compression during drought periods, particularly for Episode 2 and Episode 3. While both episodes show strong reductions in seasonal amplitude, the Episode 2 results are supported by a longer post-period and more stable parameter estimates, and therefore provide the most reliable evidence of systematic seasonal restructuring.

The reduction in amplitude reflects a disproportionate decline in peak-season demand, consistent with restrictions targeting outdoor water use. These results indicate that drought persistence operates not only through sustained reductions in average demand, but through lasting changes in the seasonal structure of use, particularly reductions in peak-season outdoor demand.

Table 3 Estimated Seasonal Effects

Regime	SF amplitude	SF % change	GPCD amplitude	GPCD % change	Interpretation
Baseline	0.358	NA	0.359	NA	Normal seasonal cycle
Episode 1	0.325	-9%	0.34	-5%	Mild compression
Episode 2	0.266	-26%	0.266	-26%	Strong seasonal compression
Episode 3	0.249	-30%	0.268	-25%	Strong seasonal compression
Post 1	0.345	-4%	0.356	-1%	Little persistent compression
Post 2	0.284	-21%	0.308	-14%	Meaningful persistent compression
Post 3	0.213	-41%	0.352	-2%	Persistent compression in SF; limited evidence in GPCD

4.3.2 Persistence of Seasonal Effects

Post-episode results indicate that these seasonal effects are not fully reversed. The persistence of seasonal effects is most clearly observed in the SF model, where post-episode amplitudes remain meaningfully below baseline levels. In particular, the post-Episode 3 amplitude is estimated to be substantially reduced relative to baseline. However, consistent with earlier discussion, this result should be interpreted as suggestive rather than conclusive given the limited number of post-period observations for Episode 3.

In contrast, the GPCD model shows weaker and less consistent persistence, indicating that these effects are more concentrated within the residential sector. Aggregate per capita measures therefore tend to obscure seasonal restructuring that is more clearly observed within the single-family sector.

5. Key Takeaways and Limitations

These findings provide empirical support for the hypothesis that drought restrictions can induce lasting changes in water demand, though the magnitude and durability of these effects vary by episode and sector. A consistent theme across the analysis is the difference in persistence signals between the SF and GPCD models. Specifically:

- The SF model exhibits stronger and more statistically robust evidence of both level persistence and seasonal restructuring, reflecting the direct impact of restrictions on residential outdoor water use.

- The GPCD model shows weaker and less consistent persistence, likely due to
  - Inclusion of commercial, industrial, and institutional uses,
  - Differing behavioral responses across sectors, and
  - Greater aggregation and measurement variability.

This comparison suggests that long-term drought impacts in the BAWSCA region are concentrated in the residential sector. Aggregate demand metrics may therefore mask important structural changes occurring within the residential sector. This divergence highlights the importance of sectoral resolution when evaluating drought impacts, as aggregate measures blend changes occurring within individual sectors.

While Episode 3 exhibits some of the largest estimated post-drought shifts in the single-family model, these results should be interpreted with caution. The Episode 3 period is relatively short and coincides with broader structural changes associated with COVID, making it difficult to fully disentangle persistence from concurrent effects.

In contrast, the 2014–2017 drought (Episode 2) provides the most robust and consistent evidence of drought-induced demand reduction. This episode shows statistically significant suppression across both SF and GPCD models, clear seasonal compression, and more stable parameter estimates over a longer post-period.

Accordingly, Episode 2 is interpreted as providing the strongest evidence of system-wide drought impacts, while Episode 3 is more suggestive of potential persistence dynamics that warrant further observation as new data in the post-restriction period become available.

The three drought episodes analyzed exhibit distinct patterns in suppression, persistence, and seasonal response, despite broadly similar stated levels of restriction severity. This is despite the fact that severity, at least as defined by requested reductions, did not vary dramatically. This highlights the critical role of defining the baseline demand trends, which suggest that the context of water use is quite different today than 20 years ago. The same will likely hold true 20 years from now. Thus, it may be unreasonable to assume that all future drought events will result in restrictions that have the same range of effects as estimated in this analysis.

Finally, it is important to note that the models developed in this analysis are designed for inference rather than direct forecasting. In particular, the estimated time trend captures historical structural changes within the sample period and is not intended to be extrapolated as a forecast driver. Accordingly, the results are best used to inform the magnitude and persistence of drought-response effects, which can then be applied to existing baseline projections in a scenario-based framework reflecting different drought response characteristics.

6. Summary and Next Steps

The findings of this study suggest that drought restrictions can produce lasting reductions in demand, particularly through sustained changes in seasonal outdoor water use. The principal inference from the study is that the 2014–2017 drought restrictions produced the largest and most persistent demand reductions of the three restriction episodes occurring since 2000. Observed patterns suggest a partial rebound following restrictions; however, persistence is driven in part by a sustained compression in seasonal variability. Estimated persistence effects are stronger in single-family demand than in aggregate demand. Other drought episodes show less consistent or statistically weaker effects and may be less reliable. In particular, there are relatively few post-drought observations for most recent drought (Episode 3), since the modeling dataset extends only through 2023.

6.1 Application to Baseline Forecasts

The results of this analysis are not intended to serve as stand-alone forecasting models, but rather to inform the magnitude, persistence, and seasonal structure of drought-related demand responses. In particular, the inclusion of a reduced-form time trend reflects historical structural change within the estimation period and is not intended to represent a causal or policy-invariant driver for extrapolation.

Instead, the estimated drought-response relationships are most appropriately applied as overlays to existing baseline projections, which are developed using sectorally and spatially disaggregated models. These baseline forecasts already incorporate assumptions regarding population, economic activity, weather, and long-term efficiency trends under non-drought conditions.

Within this framework, drought scenarios can be represented by applying a multiplicative adjustment factor to the baseline forecast, where the adjustment reflects three components derived from the estimated models:

- During-drought suppression, representing the immediate reduction in demand associated with active restrictions;
- Post-drought persistence and recovery, capturing both the initial level shift following the lifting of restrictions and the gradual return toward baseline conditions over time;

- Seasonal reshaping, reflecting changes in the amplitude and timing of demand, particularly reductions in peak-season outdoor use.

This approach allows the drought-response effects estimated in this study to be incorporated into planning scenarios without altering the underlying structure of the baseline forecasting models. It also provides flexibility to simulate alternative drought conditions by varying the assumed severity, duration, and persistence characteristics of restrictions, while maintaining consistency with the utility's existing forecasting framework.

6.2 Planning Implications

The results of this analysis have several implications for water supply planning and demand forecasting under drought conditions.

First, the findings indicate that drought restrictions can produce persistent reductions in demand, particularly within the single-family residential sector. This suggests that demand may not fully return to pre-drought trajectories following major restriction events, and that baseline projections should be interpreted with caution in the periods immediately following droughts.

Second, the analysis highlights the importance of distinguishing between average demand reductions and structural changes in seasonal demand patterns. The observed compression of seasonal amplitude, especially reductions in peak-season use, implies that drought impacts may disproportionately affect maximum-day and peak-period demands. This has direct implications for supply reliability assessments and evaluation of capacity-related investments.

Third, the results suggest that aggregate demand metrics may mask underlying behavioral changes occurring within specific customer classes. In particular, persistence effects are more clearly observed in the single-family sector than in total per capita demand, indicating that sectorally disaggregated analysis remains important for understanding and projecting demand under changing conditions.

Fourth, the variation in estimated effects across drought episodes indicates that drought response is not uniform, and depends on factors such as restriction severity, duration, and external conditions. Accordingly, planning analyses should consider a range of drought-response scenarios rather than relying on a single deterministic assumption.

Finally, the framework developed in this study provides a practical means of incorporating these insights into planning analyses. By applying empirically derived drought-response overlays to existing baseline forecasts, utilities can evaluate the potential impacts of alternative drought scenarios on both average and peak demands, while maintaining consistency with established forecasting methodologies.

6.3 Potential Simulation Prototype

Since the sectoral demand projections developed for the 2025 Demand Study were anchored to recent history by means of calibration, they effectively embed the persistence effects of recent droughts into the forecasts. As demonstrated herein, additional forecast scenarios can be designed to account for uncertainty in both level persistence and reduced seasonal amplitude of potential future droughts, using models like those constructed for this analysis. However, the addition of such scenarios should not be weighed any more or less than scenarios surrounding other demographic, pricing, efficiency, and economic factors employed in BAWSCA's demand forecasting model, which also drive expectations of long-term secular trends. The drought-response effects estimated in this study are therefore most appropriately applied as overlays to the existing baseline projections, rather than as stand-alone forecasting models.

The modeling framework developed for this study can be used as a tool to simulate the effects of future drought-restriction episodes. The framework is generalizable so that one may evaluate the characteristics of restriction periods that may or may not share the exact same outcomes as the episodes addressed in this study. Appendix B provides a structured prototype for implementing this framework, including the definition of restriction paths and corresponding behavioral response parameter sets. The appendix illustrates how empirically derived coefficients from this analysis can be translated into scenario-based simulations that vary in severity, duration, and persistence characteristics.

6.4 Recommendations for Future Demand Studies

The results of this regional assessment of drought-restriction impacts lead to multiple recommendations for future water demand studies that could enrich understanding of drought impacts and planning implications. It is recommended that BAWSCA:

Evaluate persistence and recovery patterns for BAWSCA member agencies. The analysis documented in this technical memorandum concentrated on regional metrics and as a result provides a composite view of drought-restriction impacts. The techniques defined for this analysis could be applied to water use histories of either individual agencies or defined subgroups of multiple agencies to examine the range of impacts that lie behind the average regional impacts estimated for this study. Specific agencies or agency groups could be selected based on similarities and differences among member service areas, including the mix of customers, density of development, socio-demographic characteristics, and other attributes. Such an analysis might reveal heterogenous responses to drought-restrictions and point to the general factors leading to such differences.

Conduct detailed comparisons of historical drought episodes. The estimated impacts of the drought-restriction episodes defined in this study differ in multiple aspects. A retrospective analysis of both how the restrictions were implemented, publicized, and enforced by BAWSCA member agencies may reveal factors influencing longer-term behavior associated with recovery and persistence. Furthermore, a closer assessment of active conservation programs implemented during different episodes (for example Episode 2 and Episode 3) may reveal the extent to which turf replacement, and other technological and behavioral programs offered to customers, may be tied to estimated difference in post-drought consumption patterns. A regional survey could also be designed to obtain information on conservation actions taken by consumers independently or in conjunction with active programs.

Perform on-going refinements to sectoral demand models. The demand forecasting models estimated for the 2025 Demand Study specified multiple variables to capture the effects of drought restrictions. However, the underlying specifications focus exclusively on during-restriction effects and do not employ the specifications developed in this study to address additional dynamics related to persistence and seasonal suppression. Future modeling efforts could attempt to enhance the specification of the drought-restriction components of the sector models, which may further improve model predictions and the coefficients of other non-drought variables.

Appendix A: Summary of Model Estimates

Table A-1 Variable Key for Estimated Models

Variable Name	Description	Source/notes
C	Model intercept	Estimated
SF_USE_PER_ACCT	Account weighted smoothed SF sector use per account per day - represents regional estimate of SF class use	Derived from non-missing observations in single-family model data series
BAWSCA_GPCD	Regional total per capita use per day - defined as sum of BAWSCA agency production from all sources divided by estimated population	Production data provided by BAWSCA as part of 2025 Demand Study; population estimated from CA Dept of Finance
SILICON_INDEX	Annual Change in Inflation-Adjusted GDP (Santa Clara & San Mateo Counties, San Francisco, California, and the United States)	Downloaded from: https://siliconvalleyindicators.org/data/economy/innovation-entrepreneurship/productivity/economic-growth/
SF_CONSERVATION_INDEX	Account-weighted estimate of cumulative SF sector passive savings index estimated from A4WE tracking tool (where year 2000=1)	Derived from non-missing observations in single-family model data series
TOTAL_CONS_INDEX	Population-weighted estimate of cumulative SF, MF, and CII passive savings estimated from A4WE tracking tool (where year 2000=1)	Derived
SF_PRICE	Account-weighted inflation adjusted single-family price for the 10th CCF (2022$)	Derived from non-missing observations in single-family model data series
SF_P10CCF	Population-weighted inflation adjusted single-family price for the 10th CCF (2022$)	Derived from single-family price history and population estimates (serves as price instrument for total per capita models)
@TREND	Linear time counter (2000M01=0…2023M12=288)	Derived
COVID2	Covid-19 pandemic indicator (2020M03 - 2023M12=1; 0 otherwise)	Derived
S1	Annual sine harmonic	Derived as sin(2*pi*1*month/12)
C1	Annual cosine harmonic	Derived as cos(2*pi*1*month/12)
DLNTMAX	Departure from 30-year normal avg maximum daily temperature (natural log space)	Derived from population weighted BAWSCA agency observed and long-term normal avg maximum daily temperatures (original source: PRISM)
L1DLNTMAX	1-month lag of temperature departure	Derived
DLNPRCP01	Departure from 30-year normal monthly (precipitation +0.01) (natural log space)	Derived from population weighted BAWSCA agency observed and long-term normal monthly precipitation (original source: PRISM)

Variable Name	Description	Source/notes
L1DLNPRCP01	1-month lag of precipitation departure	Derived
L2DLNPRCP01	2-month lag of precipitation departure	Derived
L3DLNPRCP01	3-month lag of precipitation departure	Derived
EPISODE1	Binary (0/1) drought restriction indicator (2006M04 - 2009M06=1; 0 otherwise)	Derived as presence of any restrictions associated with Episode from SFPUC or State implemented restrictions
EPISODE2	Binary (0/1) drought restriction indicator (2014M01 - 2017M04=1; 0 otherwise)	Derived as presence of any restrictions associated with Episode from SFPUC or State implemented restrictions
EPISODE3	Binary (0/1) drought restriction indicator (2021M08 - 2022M03=1; 0 otherwise)	Derived as presence of any restrictions associated with Episode from SFPUC or State implemented restrictions
RESTRICT_PROP	Proportion of regional population under drought restrictions associated with Episodes 1-3	Derived; used as a weight in total per capita model to account for not all BAWSCA agencies under implemented restrictions
POST1	Binary (0/1) variable defining post restriction period of Episode 1 prior to Episode 2 (2009M07 - 2013M12=1; 0 otherwise)	Derived
POST2	Binary (0/1) variable defining post restriction period of Episode 2 prior to Episode 3 (2017M05 - 2021M07=1; 0 otherwise)	Derived
POST3	Binary (0/1) variable defining post restriction period of Episode 3 (2023M04 - 2023M12=1; 0 otherwise)	Derived
MSL1	Time counter for number of months since lifting of restrictions for Episode 1; resets to value of 0 with introduction of Episode 2	Derived
MSL2	Time counter for number of months since lifting of restrictions for Episode 2; resets to value of 0 with introduction of Episode 3	Derived
MSL3	Time counter for number of months since lifting of restrictions for Episode 3; extends to end of historical data	Derived

Table A-2 Single-Family Model 1

Model Specifications
Dependent Variable: LN(SF_USE_PER_ACCT)
Method: Least Squares
Sample: 2000M01 2023M10
Included observations: 286
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 6.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	5.9034	0.0636	92.7655	0.0000
COVID2	0.0545	0.0136	3.9960	0.0001
SILICON_INDEX	0.0009	0.0008	1.0688	0.2862
SF_CONSERVATION_INDEX	-0.1134	0.1128	-1.0054	0.3157
LN(SF_PRICE)	-0.0432	0.1656	-0.2608	0.7944
DLNTMAX	0.4970	0.0907	5.4774	0.0000
L1DLNTMAX	0.2698	0.0840	3.2105	0.0015
DLNPRCP01	-0.0156	0.0022	-7.1546	0.0000
L1DLNPRCP01	-0.0136	0.0022	-6.0974	0.0000
L2DLNPRCP01	-0.0060	0.0019	-3.1676	0.0017
L3DLNPRCP01	-0.0024	0.0018	-1.3441	0.1801
S1	-0.2193	0.0127	-17.2912	0.0000
C1	-0.2815	0.0063	-44.9139	0.0000
COVID2*S1	0.0602	0.0084	7.1407	0.0000
COVID2*C1	-0.0207	0.0213	-0.9684	0.3338
EPISODE1	0.0258	0.0191	1.3492	0.1785
S1*EPISODE1	0.0196	0.0183	1.0751	0.2834
C1*EPISODE1	0.0265	0.0106	2.5115	0.0127
POST1=1	-0.0819	0.0321	-2.5495	0.0114
S1*(POST1=1)	-0.0095	0.0181	-0.5279	0.5980
C1*(POST1=1)	0.0248	0.0091	2.7281	0.0068
(POST1=1)*LOG(1+MSL1)	0.0115	0.0111	1.0398	0.2995
EPISODE2	-0.1841	0.0493	-3.7358	0.0002
S1*EPISODE2	0.0546	0.0220	2.4851	0.0136
C1*EPISODE2	0.0742	0.0180	4.1358	0.0000
POST2=1	-0.2108	0.0557	-3.7828	0.0002
S1*(POST2=1)	0.0170	0.0138	1.2370	0.2172
C1*(POST2=1)	0.0854	0.0083	10.3305	0.0000
(POST2=1)*LOG(1+MSL2)	0.0137	0.0122	1.1230	0.2625
EPISODE3	-0.1648	0.0742	-2.2213	0.0272
S1*EPISODE3	0.0131	0.0174	0.7522	0.4527
C1*EPISODE3	0.1451	0.0257	5.6449	0.0000

Variable	Coefficient	Std. Error	t-Statistic	Prob.
POST3=1	-0.6472	0.1277	-5.0684	0.0000
S1*(POST3=1)	0.1552	0.0539	2.8802	0.0043
C1*(POST3=1)	0.0843	0.0248	3.4046	0.0008
(POST3=1)*LOG(1+MSL3)	0.2950	0.0775	3.8055	0.0002

Model Diagnostics				
R-squared	0.9802			
Adjusted R-squared	0.9774			
S.E. of regression	0.0409			
Mean dependent var	5.5531			
St. Dev. dependent var	0.2720			
F-statistic	352.7581			
Prob(F-statistic)	0.0000			
Durbin-Watson stat	0.9331			

Table A-3 Single-Family Model 2

Model Specifications
Dependent Variable: LOG(SF_USE_PER_ACCT)
Method: Least Squares
Sample: 2000M01 2023M10
Included observations: 286
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 6.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	5.7697	0.0099	582.8325	0.0000
COVID2	0.0839	0.0136	6.1890	0.0000
SILICON_INDEX	0.0012	0.0008	1.4257	0.1552
@TREND	-0.0013	0.0003	-4.5471	0.0000
DLNTMAX	0.4889	0.0862	5.6709	0.0000
L1DLNTMAX	0.2707	0.0820	3.3015	0.0011
DLNPRCP01	-0.0156	0.0023	-6.9165	0.0000
L1DLNPRCP01	-0.0133	0.0022	-6.0172	0.0000
L2DLNPRCP01	-0.0057	0.0019	-2.9943	0.0030
L3DLNPRCP01	-0.0021	0.0017	-1.1950	0.2332
S1	-0.2222	0.0130	-17.0563	0.0000
C1	-0.2811	0.0062	-45.1029	0.0000
COVID2*S1	0.0560	0.0137	4.0966	0.0001
COVID2*C1	-0.0128	0.0184	-0.6963	0.4869
EPISODE1	0.0395	0.0211	1.8737	0.0621
S1*EPISODE1	0.0210	0.0178	1.1798	0.2392
C1*EPISODE1	0.0262	0.0103	2.5453	0.0115
POST1=1	-0.0648	0.0335	-1.9321	0.0545
S1*(POST1=1)	-0.0093	0.0176	-0.5302	0.5964
C1*(POST1=1)	0.0252	0.0094	2.6730	0.0080
(POST1=1)*LOG(1+MSL1)	0.0142	0.0105	1.3616	0.1745
EPISODE2	-0.1627	0.0511	-3.1822	0.0016
S1*EPISODE2	0.0548	0.0224	2.4493	0.0150
C1*EPISODE2	0.0739	0.0184	4.0067	0.0001
POST2=1	-0.1875	0.0496	-3.7797	0.0002
S1*(POST2=1)	0.0160	0.0137	1.1704	0.2430
C1*(POST2=1)	0.0863	0.0085	10.1164	0.0000
(POST2=1)*LOG(1+MSL2)	0.0156	0.0063	2.4832	0.0137
EPISODE3	-0.1503	0.0654	-2.2995	0.0223
S1*EPISODE3	0.0193	0.0207	0.9367	0.3498
C1*EPISODE3	0.1360	0.0227	5.9877	0.0000
POST3=1	-0.6344	0.1247	-5.0864	0.0000

Variable	Coefficient	Std. Error	t-Statistic	Prob.
S1*(POST3=1)	0.1621	0.0549	2.9511	0.0035
C1*(POST3=1)	0.0767	0.0227	3.3755	0.0009
(POST3=1)*LOG(1+MSL3)	0.3000	0.0768	3.9062	0.0001

Model Diagnostics				
R-squared	0.9806			
Adjusted R-squared	0.9780			
S.E. of regression	0.0403			
Mean dependent var	5.5531			
St. Dev. dependent var	0.2720			
F-statistic	373.5070			
Prob(F-statistic)	0.0000			
Durbin-Watson stat	0.9149			

Table A-4 Total Per Capita (GPCD) Model 1

Model Specifications
Dependent Variable: LOG(BAWSCA_GPCD)
Method: Least Squares
Sample: 2000M01 2023M12
Included observations: 288
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 6.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	5.4502	0.0572	95.2316	0.0000
TOTAL_CONS_INDEX	0.0232	0.0707	0.3284	0.7429
LOG(SF_P10CCF)	-0.3725	0.1094	-3.4034	0.0008
SILICON_INDEX	0.0005	0.0008	0.6625	0.5083
COVID2	-0.0783	0.0173	-4.5273	0.0000
DLNTMAX	0.4988	0.0951	5.2431	0.0000
L1DLNTMAX	0.5112	0.0980	5.2147	0.0000
DLNPRCP01	-0.0168	0.0023	-7.3565	0.0000
L1DLNPRCP01	-0.0165	0.0025	-6.6104	0.0000
L2DLNPRCP01	-0.0050	0.0020	-2.4724	0.0141
L3DLNPRCP01	-0.0018	0.0019	-0.9133	0.3620
S1	-0.2715	0.0097	-27.9268	0.0000
S1*EPISODE1*RESTRICT_PROP	0.0092	0.0172	0.5350	0.5931
S1*EPISODE2*RESTRICT_PROP	0.0685	0.0229	2.9921	0.0030
S1*EPISODE3*RESTRICT_PROP	0.0327	0.0202	1.6188	0.1067
S1*(POST1=1)	-0.0052	0.0124	-0.4172	0.6769
S1*(POST2=1)	0.0243	0.0116	2.0979	0.0369
S1*(POST3=1)	-0.0477	0.0373	-1.2806	0.2015
C1	-0.2368	0.0064	-36.8330	0.0000
C1*EPISODE1*RESTRICT_PROP	0.0172	0.0118	1.4638	0.1445
C1*EPISODE2*RESTRICT_PROP	0.0680	0.0189	3.6067	0.0004
C1*EPISODE3*RESTRICT_PROP	0.0909	0.0244	3.7246	0.0002
C1*(POST1=1)	0.0106	0.0101	1.0519	0.2938
C1*(POST2=1)	0.0501	0.0094	5.3048	0.0000
C1*(POST3=1)	0.0774	0.0336	2.3032	0.0221
EPISODE1*RESTRICT_PROP	0.0025	0.0196	0.1290	0.8975
POST1=1	-0.0621	0.0235	-2.6406	0.0088
(POST1=1)*LOG(1+MSL1)	0.0069	0.0092	0.7534	0.4519
EPISODE2*RESTRICT_PROP	-0.2056	0.0527	-3.9029	0.0001
POST2=1	-0.1341	0.0466	-2.8777	0.0043
(POST2=1)*LOG(1+MSL2)	0.0056	0.0107	0.5247	0.6003
EPISODE3*RESTRICT_PROP	-0.1420	0.0813	-1.7460	0.0820
POST3=1	-0.0658	0.1299	-0.5069	0.6127

Variable	Coefficient	Std. Error	t-Statistic	Prob.
(POST3=1)*LOG(1+MSL3)	-0.0379	0.0601	-0.6305	0.5290
COVID2*S1	0.0676	0.0132	5.1049	0.0000
COVID2*C1	-0.0466	0.0144	-3.2331	0.0014

Model Diagnostics				
R-squared	0.9754			
Adjusted R-squared	0.9720			
S.E. of regression	0.0500			
Mean dependent var	4.8610			
St. Dev. dependent var	0.2986			
F-statistic	285.6223			
Prob(F-statistic)	0.0000			
Durbin-Watson stat	1.6457			

Table A-5 Total Per Capita (GPCD) Model 2

Model Specifications
Dependent Variable: LOG(BAWSCA_GPCD)
Method: Least Squares
Sample: 2000M01 2023M12
Included observations: 288
HAC standard errors & covariance (Bartlett kernel, Newey-West fixed bandwidth = 6.0000)

Variable	Coefficient	Std. Error	t-Statistic	Prob.
C	5.3318	0.0375	142.3112	0.0000
@TREND	-0.0016	0.0002	-6.5931	0.0000
SILICON_INDEX	0.0016	0.0007	2.4146	0.0165
COVID2	-0.0330	0.0091	-3.6085	0.0004
DLNTMAX	0.4704	0.0968	4.8609	0.0000
L1DLNTMAX	0.4958	0.1043	4.7511	0.0000
DLNPRCP01	-0.0176	0.0024	-7.4723	0.0000
L1DLNPRCP01	-0.0170	0.0025	-6.8182	0.0000
L2DLNPRCP01	-0.0052	0.0020	-2.6068	0.0097
L3DLNPRCP01	-0.0018	0.0019	-0.9260	0.3553
S1	-0.2681	0.0088	-30.6104	0.0000
S1*EPISODE1*RESTRICT_PROP	0.0110	0.0166	0.6638	0.5074
S1*EPISODE2*RESTRICT_PROP	0.0675	0.0250	2.6972	0.0075
S1*EPISODE3*RESTRICT_PROP	0.0537	0.0230	2.3343	0.0204
S1*(POST1=1)	-0.0042	0.0118	-0.3568	0.7215
S1*(POST2=1)	0.0248	0.0111	2.2484	0.0254
S1*(POST3=1)	-0.0403	0.0388	-1.0403	0.2992
C1	-0.2387	0.0061	-39.4241	0.0000
C1*EPISODE1*RESTRICT_PROP	0.0169	0.0114	1.4781	0.1406
C1*EPISODE2*RESTRICT_PROP	0.0644	0.0194	3.3260	0.0010
C1*EPISODE3*RESTRICT_PROP	0.0779	0.0170	4.5728	0.0000
C1*(POST1=1)	0.0094	0.0101	0.9387	0.3488
C1*(POST2=1)	0.0495	0.0093	5.3286	0.0000
C1*(POST3=1)	0.0693	0.0317	2.1827	0.0300
EPISODE1*RESTRICT_PROP	0.0355	0.0206	1.7231	0.0861
POST1=1	-0.0079	0.0228	-0.3458	0.7298
(POST1=1)*LOG(1+MSL1)	-0.0051	0.0079	-0.6449	0.5196
EPISODE2*RESTRICT_PROP	-0.2004	0.0510	-3.9292	0.0001
POST2=1	-0.1674	0.0421	-3.9772	0.0001
(POST2=1)*LOG(1+MSL2)	0.0249	0.0053	4.6842	0.0000
EPISODE3*RESTRICT_PROP	-0.0686	0.0602	-1.1406	0.2551
POST3=1	0.0077	0.1145	0.0669	0.9467
(POST3=1)*LOG(1+MSL3)	-0.0307	0.0598	-0.5136	0.6080

Variable	Coefficient	Std. Error	t-Statistic	Prob.
COVID2*S1	0.0573	0.0178	3.2242	0.0014
COVID2*C1	-0.0375	0.0094	-3.9932	0.0001

Model Diagnostics				
R-squared	0.9749			
Adjusted R-squared	0.9715			
S.E. of regression	0.0504			
Mean dependent var	4.8610			
St. Dev. dependent var	0.2986			
F-statistic	289.0400			
Prob(F-statistic)	0.0000			
Durbin-Watson stat	1.6487			

Appendix B: Prototypical Drought-Response Simulation Approach

This appendix specifies a prototypical approach for simulating future demand responses to drought restrictions. It is designed generically to use the modeling method, and specifically to use the decomposition technique described in the main body of the report.

Each simulated drought scenario would consist of two parts, a "restriction path" and a "response parameter set."

The restriction path defines:

- Start date
- End date
- Severity (qualitative ordinal)
- Post-period start date (begins period after restrictions end)

The response parameter set defines the characteristics of a single restriction Episode as a bundle or vector of behavioral responses tied to severity:

$$Severity \rightarrow \phi = (\beta_E, \beta_P, \beta_R, \beta_{S1E}, \beta_{C1E}, \beta_{S1P}, \beta_{C1P})$$

Where terms in parentheses represent model coefficients:

β_E = During drought suppression (model variable EPISODE)
β_P = Post drought suppression (model variable POST)
β_R = Post-drought recovery (model variable ln(1+MSL))
β_{S1E} = During drought seasonal effect (model variable interaction of EPISODE*S1)
β_{C1E} = During drought seasonal effect (model variable interaction of EPISODE*C1)
β_{S1P} = Post drought seasonal effect (model variable interaction of POST*S1)
β_{C1P} = Post drought seasonal effect (model variable interaction of POST*C1)

The response parameter set associated with each restriction set reflects historically informed behavioral responses and is not assumed to be determined solely by severity, but may also reflect policy design, duration, and customer adaptation.

A simulation involves a selection of a restriction set (*r*) and a corresponding response parameter set of coefficients. The coefficients of the response parameter set (ϕ) feed the generalized structural form of the drought restriction response or multiplier *D*:

$$D_t^{(r)} = \exp\left[E_t^{(r)}(\beta_E + \beta_{S1E}S1_t + \beta_{C1E}C1_t) + P_t^{(r)}\left(\beta_P + \beta_R \ln\left(1 + MSL_t^{(r)}\right) + \beta_{S1P}S1_t + \beta_{C1P}C1_t\right)\right]$$

The drought multiplier can be interpreted as the combination of (a) a level component governing average demand suppression and recovery, and (b) a seasonal component governing changes in the amplitude and timing of demand.

A simulated forecast is given as:

$$q_t^{scenario} = q_t^{baseline} \times D_t^{(r)}$$

where, for example, $q_t^{baseline}$ represents the BAWSCA baseline scenario of SF sector per unit use or total per capita use (obtained by dividing projected baseline projections of volume by population).

A volumetric scenario would thus become:

$$Q_t^{scenario} = N_t \text{ x } q_t^{scenario}$$

where N is the appropriate driver unit.

Table B-4 Example structure of restriction set

Restriction Set	Severity	Duration	During suppression	Post period level shift	Recovery speed	During seasonal compression/ reshaping	Post seasonal persistence
RS-1	Mild	12 mo	small	none	fast	light	none
RS-2	Moderate	18 mo	moderate	moderate	moderate	moderate	light
RS-3	Severe transient	12 mo	large	large	fast	strong	moderate
RS-4	Severe persistent	24 mo	large	moderate	slow	strong	strong

Table B-5 Example Response parameter sets (each row defines ϕ)

Restriction Set	Response parameter set	During beta	Post beta	Recovery beta	During S1	During C1	Post S1	Post C1
RS-1	ϕ-1	-0.05	-0.01	0.05	0.01	0.01	0.00	0.00
RS-2	ϕ-2	-0.12	-0.05	0.03	0.03	0.04	0.01	0.02
RS-3	ϕ-3	-0.16	-0.20	0.20	0.05	0.1	0.03	0.04
RS-4	ϕ-4	-0.2	-0.15	0.02	0.07	0.06	0.08	0.05