Peer-reviewed research article

Journal of Forest Business Research 5(1), 61-97, 2026

Introducing pySRTS: The next generation of the Sub-Regional Timber Supply Model

Richard H. Manner, [a,*] Gaurav Dhungel, [b] Robert C. Abt, [a] Justin S. Baker [a]

a: North Carolina State University, Department of Forestry and Environmental Resources, Raleigh, NC, USA;
b: Oak Ridge Institute for Science and Education (ORISE), Oak Ridge, TN, USA;
[*]Corresponding author: E-mail: rhmanner@ncsu.edu

ABSTRACT

The Sub-Regional Timber Supply (SRTS) model has been a staple in the analysis and research of the Southeastern U.S. timber market for over two decades. In this paper, we take the existing model and modernize it through a ground-up rewrite and reimagining. The updated model includes adding new functionality, such as basin-level price premiums when solving the market equilibrium, and dynamic product weights in the harvest goal program. We document model assumptions, input data, and mathematical formulations in addition to identifying opportunities for users to modify and extend the model. By focusing on flexibility and modularity during the initial development, the pySRTS model is more adaptable to a wide range of research applications. We present a set of illustrative demand scenarios in which a new sawmill opens in North Carolina, U.S. In these scenarios, we explore how the new local prices and model parameterization interact to produce projections of future removals, prices, and forest inventories.

Keywords

forest economics model, SRTS, timber markets, U.S. South

Citation

Manner RH, Dhungel G, Abt RC, Baker JS. 2026. Introducing pySRTS: The next generation of the Sub-Regional Timber Supply Model. J.For.Bus.Res. 5(1): 61-97.
https://doi.org/10.62320/jfbr.v5i1.93

Received: 11 June 2026
Accepted: 10 August 2026
Published: 20 August 2026

INTRODUCTION

The Sub-Regional Timber Supply (SRTS) model has a long history in the Southeastern U.S. The core concepts of the model, introduced by Abt et al. (2000), have been used to model the Southeastern forestry sector for decades by forest sector practitioners in academic, government, industry, and non-governmental organizations. The model is a short-run partial equilibrium forest economic model focused on the interaction between market-level demand for stumpage (e.g., pulpwood and sawtimber), resource availability, and age class growth dynamics. SRTS is designed to find the market equilibrium between a monolithic stumpage demand signal and a collection of basin-level supply curves for stumpage, which are governed by inventory dynamics, forest growth and yield, and previous harvest patterns. The flexible market and forest inventory components make it one of the more spatially refined economic models available in the forest economics community.

SRTS has been applied to a range of forest market contexts. The model is used for making regional projections of market change and its impact on forest land use (e.g., Prestemon and Abt 2002, Abt et al. 2009, Manner et al. 2026). With frequent hurricanes in the region, Henderson et al. (2022) extend the model to consider their impact on inventory and future market conditions. Additionally, the model has been applied extensively to analysis of forest bioenergy markets and policy in the Southeastern U.S. (Guo et al. 2011, Abt et al. 2012, Galik et al. 2015) and carbon fertilization (Henderson et al. 2020), with incremental features being added to the existing model to address a new and widening range of questions facing the forestry sector in the Southeast.

SRTS has a regionally-defined hierarchy with basins rolling up to form a market. This makes it a useful tool when downscaling projections from other national or global models like FASOMGHG (Galik et al. 2015) or FOROM (Manner et al. 2026) where sub-regional detail is potentially lacking. In addition, previous research and development efforts have attempted to expand the model beyond the original Southeastern U.S. scope. Sendak et al. (2003) connected the SRTS framework with the ATLAS model to produce a bio-economic model for the Northeastern United States that is similar to SRTS. Creekmore (2015) expanded the scope of the model into Uruguay and Brazil to study timber markets in a South American context. The Hardwood SRTS model, described by Dhungel et al. (2023), expanded the model into the central hardwood region to study the sustainability of high-quality white oak inventory under alternative scenarios of future market demand. A common theme across these regional expansion efforts and related model developments was the need for extensive code changes and, in some cases, other inventory accounting models (i.e., ATLAS) to accomplish these changes in model functionality.

Our goal with the next generation of the SRTS model is to retain the core functionality that is used by the research community and forest sector practitioners while simultaneously modernizing the SRTS platform, streamlining the code base, and improving flexibility, all so the model can be more readily deployed in new spatial contexts. That last objective of this SRTS redesign is multi-pronged; comprised of a modular design more accepting of new functions, subroutines, and external inputs (i.e., other models) and a dynamic statement of problem parameters derived from input data with minimal restriction placed on end use.

This paper is organized in three sections. The next section introduces the new pySRTS framework with a focus on model innovations - including a general model description, discussion of underlying economic and biological components, and selected differences from the legacy SRTS approach. Next, a case study application of pySRTS on a hypothetical new demand source (e.g., a forest product mill) in North Carolina highlights new functionality and economic intuition in an empirical setting. Finally, we wrap up with a summary of key differences, methodological and empirical, and offer a vision for future research and development.

THE (py)SRTS MODEL

The SRTS model is composed of three main phases:

1) A short-term partial equilibrium model that finds the market-clearing conditions: a price such that demand equals the sum of basin-level supply.

2) A basin-level harvest module that uses forest inventories, product densities, and market supply quantities to allocate harvests onto the landscape.

3) A forest “book-keeping" step that manages starting inventory, harvests, growth, and advances the forest through time.

The first phase is driven by user-defined projections of future demand for stumpage. The model assumes that the demand levels specified in these projections are at initial prices. The second phase consumes inventory and Diameter at Breast Height (DBH) information from forest inventories, in this case the USDA Forest Service's FIA data (Forest Inventory and Analysis Database 2026), and converts it into a goal program that seeks to achieve the supply targets from the first phase by harvesting the landscape. This process is done individually for each basin in the projection. The final step is to apply these harvests to the landscape. To accomplish this, the model goes back to the inventory data, using area, volume, and growth estimates to apply removals, clear-cut acres, increment the stocks, and advance the age class structure. This process is repeated year after year, with ending inventory in year t becoming the starting inventory in year $t + 1$.

The new generation of the SRTS model, called pySRTS, is written in Python 3.14.4, leveraging the popular pandas (pandas-dev-team 2020) package for data management and workflow, SciPy (Virtanen et al. 2020) for the goal programming (specifically *scipy.optimize.linprog*), and NumPy (Harris et al. 2020) for vector operations and sorting. We retain the spirit of the original model, though much of the original code is rewritten as part of this ground-up reimagining. The model is composed of four main classes: settings, forest, market, and harvest. The basic overview of the pySRTS process is presented in Figure 1.

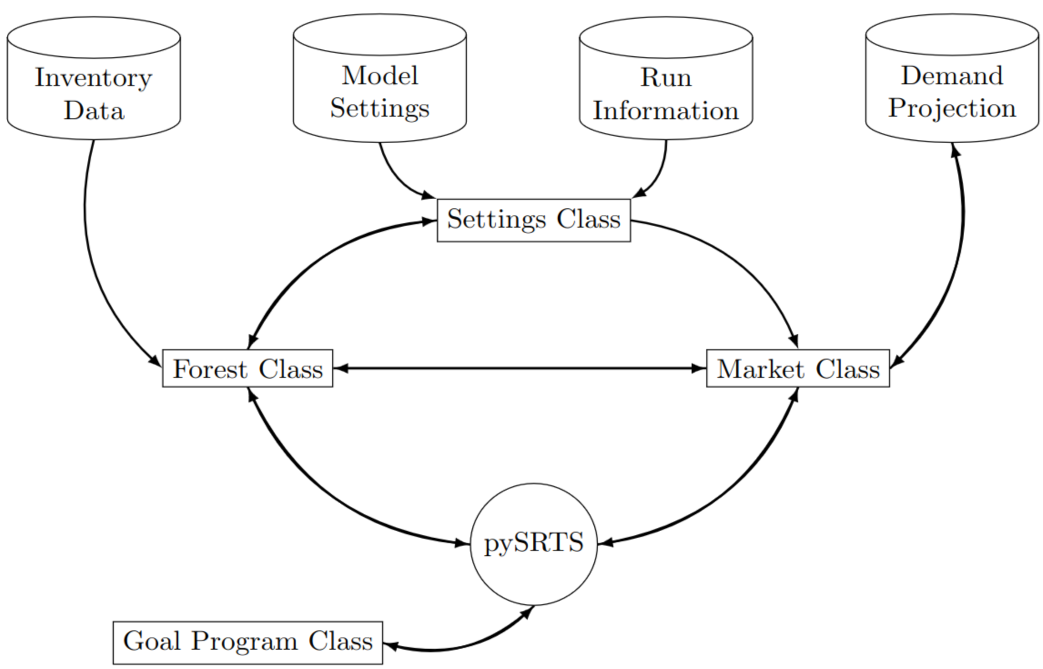

Figure 1. High-level process overview for the pySRTS model code.

The settings class houses most user inputs, like the product definition file and scenario settings. The settings class holds a single set of key model parameters defined universally. The forest class houses and maintains the forest inventory information, which in the U.S. context is derived from the Forest Inventory and Analysis (FIA) database and recent updates to the National-Scale Volume Biomass equations (Westfall et al. 2024, Restrepo et al. 2025). The forest class provides inventory and growth information to the market and harvest classes and houses the default forest aging algorithm that advances the forest through time. The market class is used to find market equilibrium based on available inventory and the demand projection file (PRJ). It also houses the land use change module and passes information to the forest module to age the forest. The last core module is the harvest class, responsible for setting up and solving the harvest goal program. This class uses information and parameters from the settings, market, and forest classes to harvest inventory from the individual timber supply basins and support the market solution.

The model pipeline is defined as a function within the pySRTS module. The base model formulation is set up for the Southeastern U.S. region, and is called by the function *se.pysrts(...)*. This function makes a series of class-level calls to execute the SRTS model and move both economic and physical information through time. In each simulation period, the model reads starting inventory conditions, determines the market solution (prices and quantities), solves the harvest goal program, and advances the forest resource to create the initial inventory conditions for the next period (see Figure 2). This sequential process requires the harvest program to minimize harvest miss vs. maximizing welfare directly, as would be the case in a more tightly coupled optimization. This limits the model's ability to produce truly economically-optimal solutions in cases where the ratio of product demand does not align with landscape realities. However, there are a few benefits to this decoupling. First, this approach improves computational tractability and speed. By limiting the size of the individual optimization problem, the model is able to converge more quickly and, in cases where required, be distributed across processors. This approach also allows a separation of economic and operational logic; the potential to estimate welfare losses due to sub-optimal operational constraints and imperfect information is thus preserved. These steps form the core of the pySRTS model and are made available to users who would like to tailor the model to their specific needs.

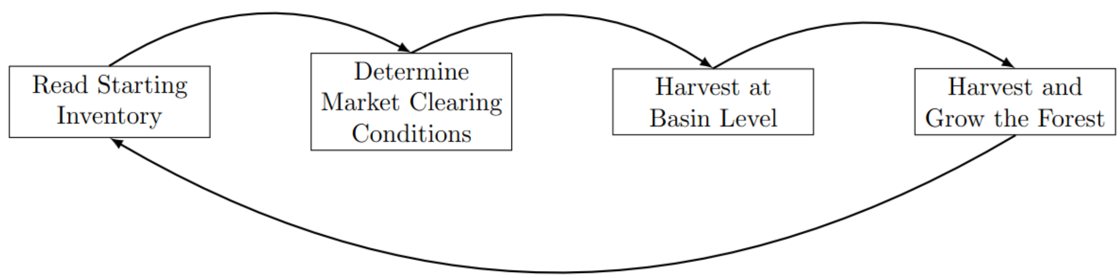

Figure 2. Southeastern pySRTS workflow.

Given the extensive legacy of the model, an effort was made to keep as many of the input files consistent between the original model (SRTS) and the new model (pySRTS) as possible. There are a few notable new files: a separate file for holding default or user-defined elasticity parameters, and a residue flow file. These allow for a more seamless implementation of the problem dimensions and user flexibility in defining market responsiveness and the flow of residual biomass between markets. In addition, initial prices have been added to the product file, offering the ability to calibrate starting price points to spatially varying stumpage prices and to validate price projections against historic trends.

One key feature that makes the model more portable is the dimensioning of the problem dynamically from the input inventory data. The dimensions of the INV file (a breakdown of inventory, growth, area, and removals by species, age class, management type, owner, and basin) are joined with the dimensions of the DBH Pct File (a mapping of age class level volume to DBH) to create the "merchantable" forest, or initial inventory conditions for merchantable stumpage. The dimensions of the harvest goal program, in the harvest class, are established according to this intersection. Similarly, the original dimensionality of the INV file, plus a new setting defining the number of years per age class, are used to set the dimensions of the forest aging procedure, where harvests are applied, growth is calculated, and the forest age class structure advances through each time step.

The incorporation of a new residue market setup allows the flow of off-cuts, sawdust, and other mill by-products to flow into other product pools (e.g., pulpwood or biomass for energy). The solution analyzes these source-to-destination product flow percentages (defined by users) to identify the appropriate solve order and ensure that changes in year-to-year production in source products are propagated to destination products correctly.

The key scientific enhancement of the model is the introduction of local prices that fluctuate based on changes in mill capacity and demand. LURA (Latta et al. 2018) is currently the primary U.S. forestry-focused partial equilibrium model that demonstrates the effect of local changes in mill capacity on local delivered wood costs and output prices. The SRTS/pySRTS model's focus on markets, ability to project into the more distant future, and endogenous land use change is a key differentiator for SRTS/pySRTS, which may be preferable in longer-term projections contexts where we expect market changes to influence both wood allocation and subsequent management outcomes. pySRTS links local prices with more explicit control of residue and cull percents to produce more detailed studies of local and regional dynamics in the face of new emerging industries with potentially unique pulp or biomass demands.

The market procedure

The market class houses the new local prices driven by revised direct change logic as well as the endogenous, price-based, product weights. These new price-related functions enable a more nuanced view of how a regional stumpage market may separate when the law of one price begins to break down due to market friction. In addition, the finer scale control over the residue market allows for more robust modeling of non-forest sources of biomass, particularly important as residues become a focus of bio-energy and sustainable fuel initiatives.

In demand mode, the market class is responsible for finding the equilibrium between demand and the sum of supply across the various regions and owners in a given period (t). The demand function (equation 1) is made up of a calibration term (α_ρ), a demand shifter ($vsBase^D_{\rho,t}$), a price ($P_{\rho,t}$), and a price elasticity of demand (ϵ^D_ρ). The rho (ρ) subscript denotes the product.

$$D_{\rho,t} = \alpha_\rho \times vsBase^D_{\rho,t} \times P_{\rho,t}^{-\epsilon^D_\rho} \tag{1}$$

In the initial period, we assume that the total demand, derived from the INV file and product definitions, is satisfied at the initial price ($P_{\rho,0}$). Thus α_ρ is simply the ratio of the initial demand quantity over the initial price raised to the price elasticity of demand. In the initial period, residues are assumed to be included in the removals signal. In subsequent years, residues are calculated on the year-on-year change. As the model projection proceeds, the demand curve is shifted based on the current PRJ value net of residues relative to the PRJ value of the starting year. This is captured in the $vsBase^D_{\rho,t}$ dynamics in equation 2.

$$vsBase^{D}_{\rho,t} = \frac{PRJ_{\rho,t} - Residues_{\rho,t}}{PRJ_{\rho,0}} \tag{2}$$

The supply side (equation 3) adds subscripts for basin (b) and owner (o). The calibration term is then beta ($\beta_{\rho,b,o}$), a supply shifter ($vsBase^{S}_{\rho,t,b,o}$), available product inventory, a proxy for costs, in a given basin/ownership combination ($I_{\rho,t,b,o}$), an inventory elasticity ($\epsilon^{I}_{\rho,o}$), price, and a price elasticity of supply ($\epsilon^{S}_{\rho,o}$). A similar calibration process is used to set the values of $\beta_{\rho,b,o}$ based on initial conditions.

$$S_{\rho,t,b,o} = \beta_{\rho,b,o} \times vsBase^{S}_{\rho,t,b,o} \times I^{\epsilon^{I}_{\rho,o}}_{\rho,t,b,o} \times P^{\epsilon^{S}_{\rho,o}}_{\rho,t} \tag{3}$$

In SRTS, the equilibrium price is calculated based on a log-log approximation of percent change. Thus, the legacy model calculates the price change year to year that equates the changes in demand, driven by the PRJ, and supply, driven by inventory. In pySRTS, the equilibrium price is found by directly solving the partial equilibrium each year. We implement a vector local search that seeks the market price ($P^{*}_{\rho,t}$) that satisfies equation 4. This is repeated for each product market.

$$D_{\rho,t}\left(P^{*}_{\rho,t}\right) = \sum_{b,o} S_{\rho,t,b,o}\left(P^{*}_{\rho,t}\right) \tag{4}$$

In harvest mode, the price elasticities of demand are assumed to be zero (0) and thus the price represents what is required to induce landowners to supply a given level of product demand (equation 5).

$$\alpha_{\rho} \times vsBase^{d}_{\rho,t} = \sum_{b,o} S_{\rho,t,b,o}(P^{*}_{\rho,t}) \tag{5}$$

The residue market

The user can configure the residue market as a set of percentage flows ($\Phi_{s,d}$) from a source (subscript s) product to destination (subscript d) product(s). For example, assume there are three softwood products: pulp, chip-and-saw, and sawtimber. The user could define 20% of sawtimber residues flow to the pulp market, while another 10% flow to chip-and-saw. pySRTS analyzes this network and determines the solve order for the products that ensure source products' equilibria are known before destination products are considered. In this case, sawtimber would be solved first, and then in arbitrary order the chip-n-saw and pulp markets. If there was another dependency, such as 5% of the chip-n-saw volume flowing to the pulp market as residues, then the solve order would be fully defined as sawtimber → chip-n-saw → pulp.

Unlike mill surveys, the FIA data used to construct the input files for SRTS are based purely on removals from the forest. As a result, any mill demand that is satisfied by residues and non-forest sources of biomass is not included. The SRTS model assumes that the removals present in the INV file already reflect the residue market. pySRTS carries this assumption forward; therefore, only the year-to-year change in removals for a source product is eligible for the residue market. Residues are calculated after the sub-regional market equilibrium is satisfied (equation 4); therefore, the model can consider only the year-to-year variations in projected removals when determining the residue market. Equation 6 shows the dynamic nature of this calculation.

$$Residues_{d,t} = \sum_{s} \left(\phi_{s,d} \times \left(D_{s,t}\left(P^{*}_{s,t}\right) - D_{s,t}\left(P^{*}_{s,t-1}\right) \right) \right) \quad (6)$$

Direct Change file

The Direct Change (DC) file was previously used in SRTS to update a particular year's market solution for a given basin and product combination (or more than one). This functionality has been overhauled in the pySRTS model. An intuitive use case may be the addition of a new softwood sawmill in a specific timber supply basin. One might expect total sawtimber demand to increase in this basin based on the increase in local mill capacity. Further, depending on the size of the procurement zone for the new mill, one might not expect wood procurement to be fluid across the broader region. In the original SRTS formulation, the increase in demand would be satisfied ex-post the market equilibrium by adjusting the solution directly. These two methods are compared in Table 1.

Table 1. A comparison of the direct change file implementation between SRTS and pySRTS.

SRTS	pySRTS
• Changes to basin-level supply are directly applied after market equilibrium • The equilibrium displacement method means changes are automatically accounted for in the market for the next year • Price effects are revealed only at the regional level	• Changes to basin-level supply are determined using the basin-level supply and demand curves • The direct estimation of market equilibrium each year allows direct-change entries to migrate from local to regional demand • Price effects are revealed at the basin level and in subsequent year's regional prices • Local price premiums are available at for each basin

In pySRTS the DC file information defines a "mini-market", where the timber supply basin demand and supply curves are in equilibrium at the market price plus a premium that reflects a near-term market shift associated with the new mill capacity. To solve this mini-market, the broader region's market price ($P^*_{\rho,t}$) is established without considering any DC file information (see section 2.1). For the basin(s) with a DC file entry, the supplied quantity based on the market solution ($S_{\rho,t,b,o}(P^*_{\rho,t})$) is used to calibrate a basin-level demand curve at the market price. This basin demand curve then shifts according to the DC file to position the mini-market demand curve (see equation 7). It is important to note that a timber supply basin can have multiple forest ownership classes, in keeping with the market definition, this mini-market demand curve represents the total demand for stumpage across ownerships, allowing each ownership group to supply a different quantity determined by inventory and price elasticities of supply.

$$D_{\rho,t,b} = \frac{\sum_o S_{\rho,t,b,o}(P^*_{\rho,t})}{(P^*_{\rho,t})^{-\epsilon^d_\rho}} \times \frac{\sum_o S_{\rho,t,b,o}(P^*_{\rho,t}) + DC_{\rho,t,b}}{\sum_o S_{\rho,t,b,o}(P^*_{\rho,t})} \times P_{\rho,t}^{-\epsilon^d_\rho} \tag{7}$$

The supply curves are carried over from the region-wide market analysis for any of the ownership categories that are present in the target basin. The mini-market equilibrium is then solved across the ownerships in a given supply basin (see equation 8). $PP_{\rho,t,b}$ is then the price premium that is paid above (or below) the current regional market price.

$$D_{\rho,t,b}\left(P^*_{\rho,t} + PP_{\rho,t,b}\right) = \sum_{o} S_{\rho,t,b,o}\left(P^*_{\rho,t} + PP_{\rho,t,b}\right) \tag{8}$$

This equilibrium condition assumes that the market-level price elasticity of demand is appropriate for the mini-market demand function, which is consistent with other forest sector models that disaggregate national-scale demand to the regional or mill levels (Visser et al. 2022). If the model is run in harvest mode, the price elasticity of demand is set to zero (0), just as it is in the market equilibrium. After an equilibrium is found, the market solution demand and supply totals are then updated to reflect the new mini-market realities. The model tracks the market price, basin level price premiums, and calculates a regional average price, weighting removals by the market price (P^A_ρ) plus any appropriate basin level price premiums.

The harvest procedure

In pySRTS, the harvest class has been more tightly coupled with the market solution through the inclusion of dynamic, price-based product weights. These weights begin to mimic economic intuition during the harvest allocation by weighting harvest decisions towards the highest value use. In addition, the user-settable exogenous upper and lower bounds on harvest as well as internally scaling path dependency weights enable the study of operational constraints' influence on the forest resources base.

Both SRTS and pySRTS use a goal program to downscale the market solutions for each basin (b) and ownership (o) in a particular year (t) and represent the interconnectedness of forest products (ρ) harvests. For simplicity, we drop the b and o subscripts for much of this section. The market solution serves as product level harvest targets ($S_{\rho,t}\left(P^*_{\rho,t} + PP_{\rho,t}\right)$), though the forest is managed and harvested, not by product but by management type (m, traditionally Planted Pine, Natural Pine, Mixed Pine, Upland Hardwood, and Bottomland Hardwood in SRTS) and Age Class (a, 5 year groupings of ages in SRTS). pySRTS is not limited to these management type and age class definitions, though it expects these definitions in the default *se.pysrts(...)* formulation. The goal program is run individually for each basin and ownership (and year given the sequential nature of the algorithm); therefore, we drop t from the notation for the remainder of this section.

The decision to harvest a given volume from a particular management type/age class is driven by three main criteria; 1) the amount of a product that is created, 2) the amount of other products that are created, and 3) how the forest has been historically harvested (a path dependency). The first two are determined by product densities ($\Psi_{\rho,m,a}$), which represent the percentage of the harvested volume that ends up satisfying the product level demand. Therefore, the sum of the product densities for a given management type and age class (i.e., across products) is equal to one (see equation 9).

$$\sum_{\rho} \Psi_{\rho,m,a} = 1 \tag{9}$$

Path dependency is represented as distance away from, above ($\alpha_{m,t}$) or below ($\beta_{m,t}$), the historical removal patterns. pySRTS calculates ideal basin-level harvest targets based on the initial (i.e. t_0) product-level pattern of removals across management types and age classes ($Removals_{\rho,t_0,m,a}$). This pattern represents the percentage of product removals that occurs in each of the buckets on the landscape within a given basin. The pattern is fixed for the run but is scaled each year based on the volume demanded of each product coming from the market solution (S_ρ). The resulting targets, $\Omega_{m,a}$, are described by equation 10 including the final aggregation to management type and age class.

$$\Omega_{m,a} = \sum_{\rho} \left(\frac{Removals_{\rho,t_0,m,a}}{\sum_{\rho} Removals_{\rho,t_0,m,a}} \times S_\rho \right) \tag{10}$$

A set of dynamic weights are applied in pySRTS, $W_{m,a}$, to any path miss to calibrate their importance in the overall model performance; this is expressed as a percentage of the lowest average product price.

The objective of the goal program in both SRTS and pySRTS is to minimize "harvest miss", or the distance above (A_ρ) or below (B_ρ) the market solution target and the aforementioned path dependency (see equation 11). As noted in equation 9, a unit volume harvested from a particular management type and age class can supply various products. It therefore may be necessary to over or under harvest some products to achieve the optimal solution. In SRTS this harvest miss is weighted by the product weights, which are typically 1. pySRTS however uses the average price (P_ρ^A) as an endogenous weight. Therefore more scarce and higher value products are given a higher weight when solving for optimality (e.g., a higher sawtimber to pulpwood price ratio would assign more weight to sawtimber removals).

$$
\begin{aligned}
&\text{minimize}_{H_{m,a}} \left\{ \sum_{\rho} P_\rho^A \times \left(A_\rho + B_\rho\right) + \sum_{m,a} W_{m,a} \times \left(\alpha_{m,a} + \beta_{m,a}\right) \right\} \\
&\text{Subject to:} \\
&\qquad \sum_{m,a} \left(\Psi_{\rho,m,a} \times H_{m,a}\right) - A_\rho + B_\rho = S_\rho(P_\rho^* + PP_\rho) \\
&\qquad H_{m,a} - \alpha_{m,a} + \beta_{m,a} = \Omega_{m,a} \times AI_{m,a} \\
&\qquad H_{m,a} \geq l_{m,a} \times \Omega_{m,a} \\
&\qquad H_{m,a} \leq u_{m,a} \times AI_{m,a}
\end{aligned}
\tag{11}
$$

The last two constraints in equation 11 represent bounds on the harvest quantity. The first is a lower bound on harvests, a parameter ($0.0 \leq l_{m,a} \leq 1.0$) that sets the percentage of the ideal removal target ($\Omega_{m,a}$) that must be harvested. The second serves as an upper bound, stating we cannot harvest more than we have available. Thus, the parameter $u_{ma,a}$ is a percentage applied to the available inventory ($AI_{m,a}$) in the domain $[0,1]$. These parameters, l and u, are configurable in the model at the management type and age class level. Situations arise where the ideal harvest is greater than the upper bound on harvests (and potentially the amount available). In these cases, the lower bounds are adjusted to equal the upper bound. Bounds can be directly adjusted to achieve various results, in fact the lower bounds on harvests for planted pine age class 8 and 9 are elevated in the Southeastern context to prevent the accumulation of older age class planted pine. In this specific case, the lower bounds are set to the minimum of the ideal harvest level or the upper bound on harvests. This adjustment is introduced in the Southeastern-specific function (*se.pysrts*). This design decision was made to allow for the development of custom and regionally specific workflows that do not interfere with the core model classes. The model is allowed to remain more general in nature while supporting the idiosyncrasies of specific markets. A user could develop their own pipeline with different lower or upper bounds as required based on their region, market, management objectives, or research interests.

Aging the forest

The final component of the model is propagating the harvests and growth forward through time. Much of the logic in pySRTS mimics that in SRTS. However, the separation of the forest "book-keeping" process lays the groundwork for future study, where growth and mortality can be independently modeled. This approach offers new opportunities for research and application, including potentially layering in stochastic disturbance events and more detailed damage assessments that replace existing hurricane and disturbance logic. SRTS and pySRTS assume that the inventory in the INV file represents a starting position. The harvest procedure operates on management types and age classes, and the output is removals at this level. However, most forests are not monolithic in nature; even planted pine in the Southeast has some amount of hardwood mixed in. To account for this, we maintain the forest inventory with species-level (sp) information. As part of this forest aging process, we allocate removals to the species level based on the product densities ($\Psi_{\rho,m,a}$).

Next, the procedure to age the forest classifies removals into thinnings and clear cuts. The classification is done using a volume per acre (VPA) target derived from the INV file. If a particular management type and age class is above the target, removals are treated as thinning. If the VPA is reduced down to target, and there is additional harvest volume that needs to come off the landscape, it is done as clear-cutting.

It is important to note that SRTS and pySRTS both maintain an annual view of the forest. This view breaks down the age classes into annual time steps (ag). For example, if age class 1 represents one- to five-year-old trees, the annual inventory file would have five separate records, one for each year. Initially, inventory is broken out uniformly across ages, but the age data is persisted throughout the model run, so it begins to represent the reconfiguration of the forest's age structure as soon as the first set of harvests is made.

After all of the removals are accounted for, the clear-cut acres are reset to an age of zero (0). At this point, the growth per acre (GPA) values come into play. Based on the number of acres in each age, a growth volume is calculated. The model calculates ending inventory according to equation

12. It is at this time that the land-use change module is called (Henderson et al. 2024). The new Python implementation follows the original closely.

$$\begin{aligned} EndingInv_{b,o,t,sp,m,ag} = StartingInv_{b,o,t,sp,m,ag} - Removals_{b,o,t,sp,m,ag} \\ +GPA_{b,o,t,sp,m,ag} \times Acres_{b,o,t,m,ag} \end{aligned} \quad (12)$$

Finally, we advance the ending inventory's time parameter (t) one step into the future such that the ending inventory in the previous time step then becomes the starting inventory in the following period (see equation 13).

$$StartingInv_{b,o,t+1,sp,m,ag} = EndingInv_{b,o,t,sp,m,ag} \quad (13)$$

CASE STUDY: MODELING A NEW FOREST PRODUCT MILL DEMAND

When a new mill enters the market, there are two potential effects. In a market with low friction (i.e., transportation costs), new demand is absorbed across the region as a whole with all basins settling into a new equilibrium. In pySRTS, this would be modeled by introducing new demand into the projection (PRJ) file. When there is more friction in the market, demand could be confined to a specific locale or basin within the broader region. The Direct Change module (see section 2.1.2) allows the model to bound the influences of the new mill's demand and introduces basin-level price premiums. We also might consider a hybrid option, where an initial shock is locally defined, and transitions to a region-wide demand shift. In this case, the PRJ and DC files can be used in tandem to initially locate new demand in the source basin and then transition to higher market demand through time.

In this section, three demand shocks are explored. The state of North Carolina serves as the market. There are four basins defined by the USDA Forest Service Survey Units: the Southern Coastal Plain (SCP), Northern Coastal Plain (NCP), Piedmont, and Mountains. We project from 2022 for a 30-year horizon through 2052. A new softwood sawmill is introduced in 2027 that is expected to consume 10,000 MCF[1] of softwood sawtimber per year. We ran 5 scenarios in total, summarized in Table 2.

Table 2. A summary of new mill scenarios considered and their high-level implementation.

#	Scenario Name	New Mill Location	Implementation Notes
1	Baseline Constant Demand	No New Mill	Demand Derived from FIA Removals
2	Regional Demand	North Carolina	Demand introduced in PRJ
3	SCP Only	SCP	Demand introduced via DC file
4	NCP Only	NCP	Demand introduced via DC file
5	Transitional	SCP → North Carolina	DC file and PRJ

To parameterize the model, we use the elasticities (Table 3), which are largely pulled from the literature (Robinson 1974, Newman 1987, Carter 1992, Brown and Zhang 2005, Polyakov et al. 2005; 2005, Liao and Zhang 2008, Polyakov et al. 2010) and as used in Manner et al. (2026). Generally, mills are less price sensitive in the hardwood market compared to softwood. Suppliers are fairly similar, except for hardwood sawtimber, which is less price sensitive. Inventory tends to be unit elastic, except in the case of softwood sawtimber, where a 1% increase in softwood sawtimber inventory would only result in a 0.55% increase in supply.

Table 3. Elasticity assumptions for pySRTS new Sawmill demand scenarios.

Elasticity	SW Pulpwood	SW Sawtimber	HW Pulpwood	HW Sawtimber
Demand	0.6	0.5	0.35	0.3
Supply	0.45	0.4	0.4	0.25
Inventory	1.0	0.55	1.0	1.0

Baseline

A constant demand scenario is used as the baseline, where the initial removals from 2022 by product (softwood pulpwood, softwood sawtimber, hardwood pulpwood, and hardwood sawtimber) are held constant in the PRJ file (see Figure 3). North Carolina is predominantly a softwood-dominated timber market, with roughly two-thirds of removals coming from softwood pulpwood and sawtimber. Hardwood pulpwood and sawtimber then make up the final third and are roughly evenly split. Both of these numbers are influenced by our 30% product cull factor. If

[1] In forestry and timber measurements, MCF stands for 1,000 cubic feet (M = 1,000, CF = cubic feet) of solid wood volume.

cull factors were lower, we would expect more harvests in the grade products compared to pulpwood.

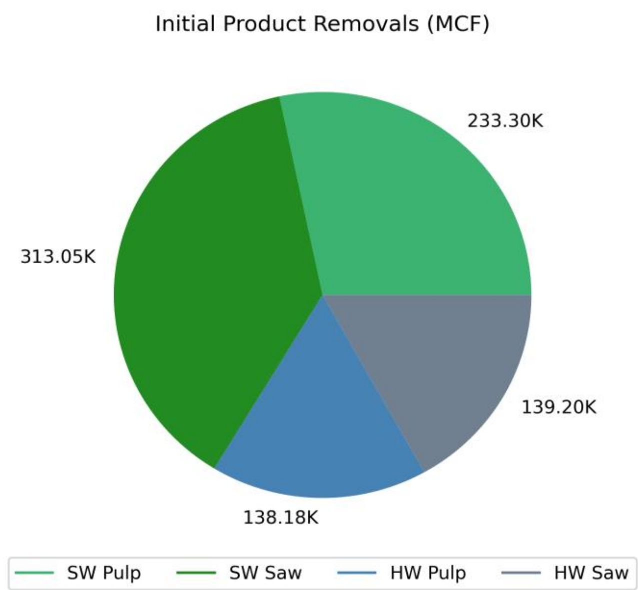

Figure 3. Initial removal levels used as in the baseline projection.

The price paths (see Figure 4) for all products are generally downward sloping as growth outpaces removals and inventory grows. Given the constant demand in the PRJ file, this follows as the third term in equation 3 drives prices lower as inventory rises through time. The prices in the pulp markets fall at similar rates, given the similarity in elasticities. However, we see a divergence in the sawtimber products, with hardwood sawtimber prices falling faster and more linearly than their softwood counterparts. Initial prices come from the 2022 Q4 quarterly price report published by NC State University's Extension Forestry (Bardon 2022). Prices are converted from U.S. dollars ($) per green ton to U.S. dollars per MCF using weight factors derived from the 2024 North Carolina Timber Product Output (TPO) state-level production survey results (U. S. Department of Agriculture 2024) available in Table A1.

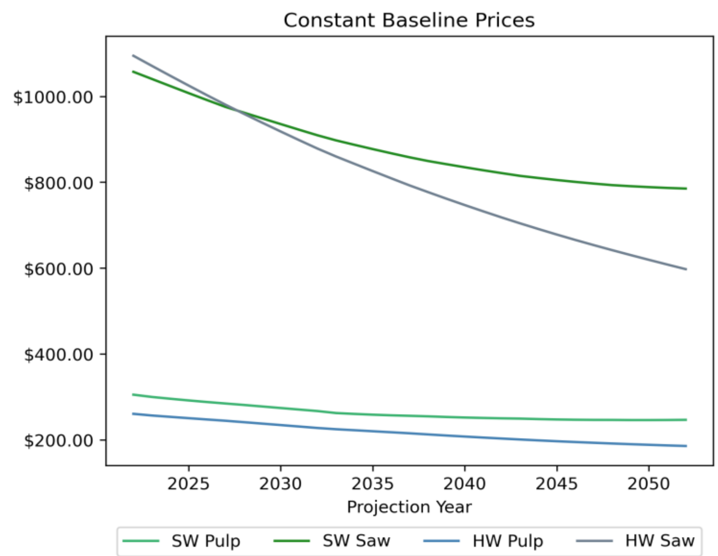

Figure 4. Baseline price trends.

Given our interest in softwood sawtimber dynamics in the state, we present trends in total softwood removals and inventory in Figure 5. There are two distinct patterns present: removals and inventory are rising over time. This indicates a growth-to-drain of greater than one, and the lack of scarcity pictured here leads to the falling prices observed through time. We present these trends for three of the four basins in NC. The fourth, Mountains, does not have a very robust pine inventory and is excluded for that reason (see Figure 1A). The bulk of removals comes from the Southern Coastal Plain (SCP). The Piedmont starts out in second place, but relatively low inventory growth leads to stagnant removals through time. By the midpoint of our projection, the Northern Coastal Plain (NCP) overtakes the Piedmont as the second largest supplier of softwood. It also ends up with the largest inventory, though not in sawtimber.

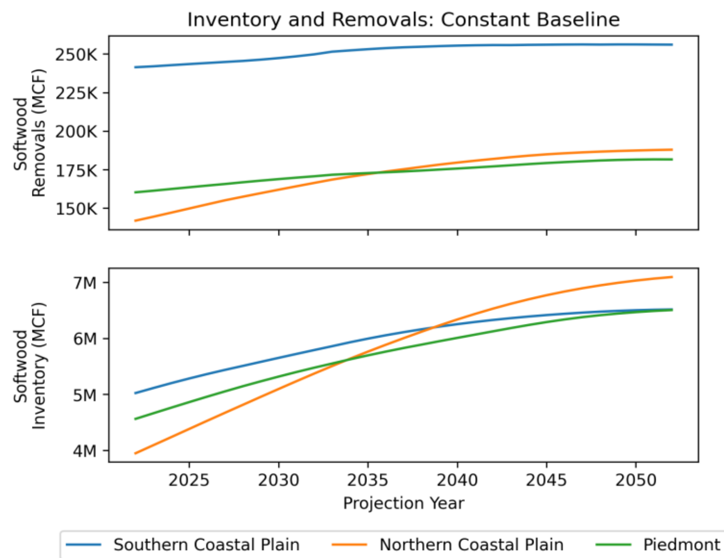

Figure 5. Trends in removals and inventory through time under baseline conditions.

New mill scenarios

Initially, we construct three scenarios for the new mill to enter the NC region. A low-friction regional scenario, where 10,000 MCF of new softwood sawtimber demand is added to the PRJ file. A scenario where the new mill is located in the SCP and demand is confined there. The addition of the new mill to the NCP serves as the third scenario. In the second and third scenarios, we expect the new mill to source from either the SCP or NCP and drive a local price premium via the new DC file logic. A fourth scenario is added to explore the transitional path from a local shock to a market reorganization.

Average prices rise when new mill demand is added to the market. Figure 6 outlines these changes across the three scenarios. Prices increase by the largest percentage over the baseline scenario when the new demand is added to the market directly. These price increases start off at about 3.55% in 2027 and rise through time to a peak near 3.8% between 2040 and 2045 before starting to fall again. The NCP scenario results in the second highest prices initially, with the average regional price also starting at roughly 3.55% over the baseline. However, the price tends to fall through time. The SCP scenario shows the lowest initial average price increase at just below 3.5%. This price difference stays relatively constant through time. Notably, the range of price responses are all very similar across the three scenarios, differing by less than 50 basis points at their peak.

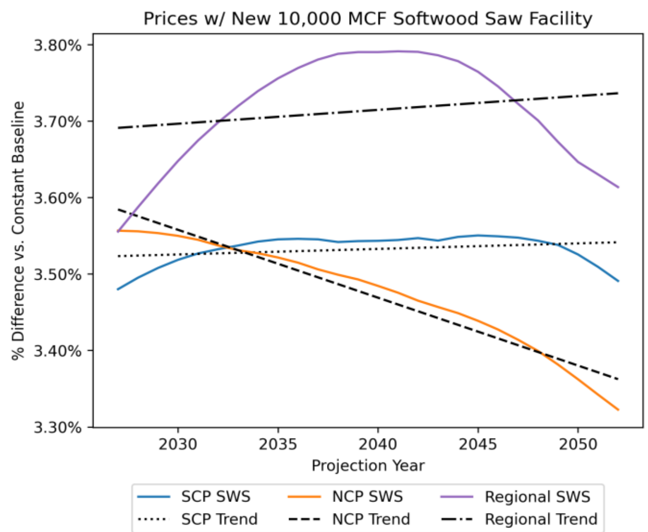

Figure 6. Scenario price results relative to the baseline. Note: SWS stands for Softwood Sawmill.

There is a notable feature in these results, namely that prices are highest under the most flexible scenario, where demand from the new mill can be spread across the regions. Inventory changes are driving these dynamics. From the perspective of the regional scenario, falling inventories in all three basins lead to a positive price response. When considering either of the DC file-driven scenarios, we see high harvests and large inventory decreases for the target basin, while the other two differ slightly from baseline. In both scenarios, the two regions that are not the site of the mill act to moderate prices as their inventory remains high through time. In this way, more wood comes off of the landscape initially at a lower average market price when using the DC file, a result that is potentially counter-intuitive. This behavior is evident when comparing the three panels in Figure 7.

Under the regional scenario, removals are spread across the three basins according to inventory availability. In this case, the SCP initially supplied most of the new demand, while the NCP and Piedmont supply roughly half as much. Through inventory accumulation, the Piedmont leads the SCP to overtake the SCP in later years and ends the projections with roughly equal levels of removal. The largest inventory differences are in the Piedmont, likely driven by the unique age class structure in that basin. The NCP and SCP scenarios are similar in behavior, though the basins in which removals occur are, by design, different. Both regions show a reduction in softwood inventory of roughly 40K MCF. This happens earlier in the simulation horizon in the NCP than in the SCP, and both regions show some rebound as forest dynamics play out. Of interest is the slight increases in removals in the other basins relative to the baseline, evidence of a marginal harvest reallocation across basins driven by lower initial inventories in the basins where the mill is sited.

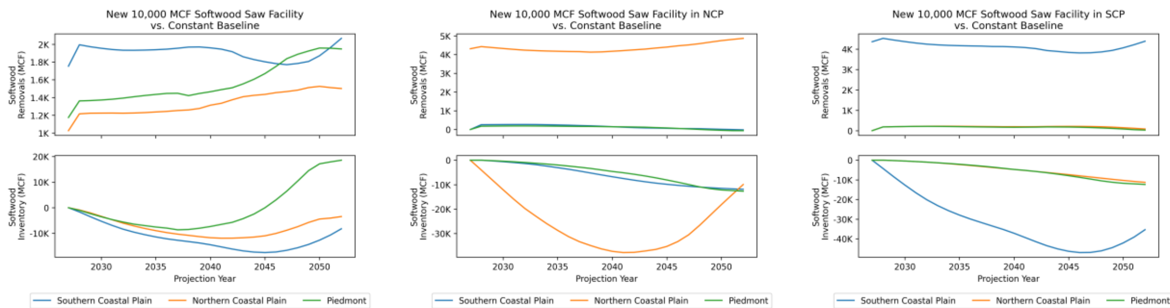

Figure 7. Removals and Inventory vs the Baseline. Left: Regional Scenario, Center: Northern Coastal Plain, Right: Southern Coastal Plain

Age class dynamics

Age class structure and dynamics explain some of the behavior we see in the scenario outcomes. Namely, why the price in the SCP scenario holds fairly constant while the price in the NCP scenario falls through time. In Figure 8, we compare the age class structure in the NCP for the baseline scenario and the NCP located mill. We see that a little more than 60% of the inventory is in age class 7 and younger (35 years old and younger). As the projection moves through time, that number shrinks to less than 10% with inventory accumulating in the older age classes. This drives the price of sawtimber down as those upper age classes are predominantly grade.

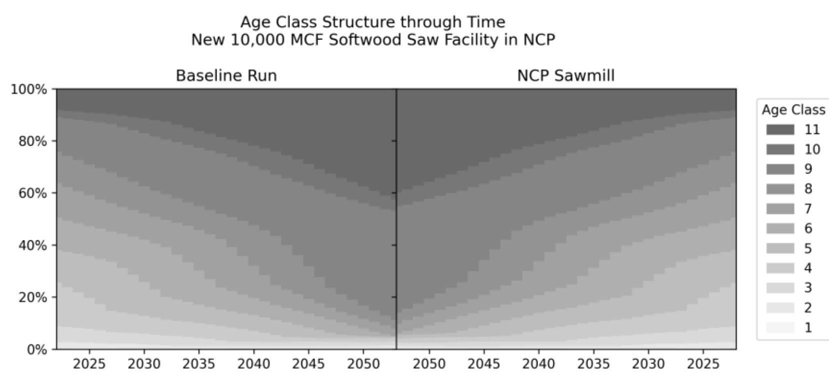

Figure 8. The share of softwood volume (percent of total) in the Northern Coastal Plain by age class through time. LEFT: Baseline scenario. RIGHT: NCP Scenario. Note that the time axis is reversed on the right, providing a comparison of the end of the projection.

A different story emerges for the SCP (Figure 9). A larger proportion of the inventory was already in age class 11 (natural pine 50+ years) when the projection started. This inventory of older age stands grows through time, but the other relatively old age classes (8, 9, and 10) do not continue to expand. The result is less of the softwood inventory on the landscape ending up in the older age classes, other than age class 11. The other grade age classes do not grow in share and in some cases shrink. This implies more of a balance between removals and growth. As a result, prices stay slightly flatter through time.

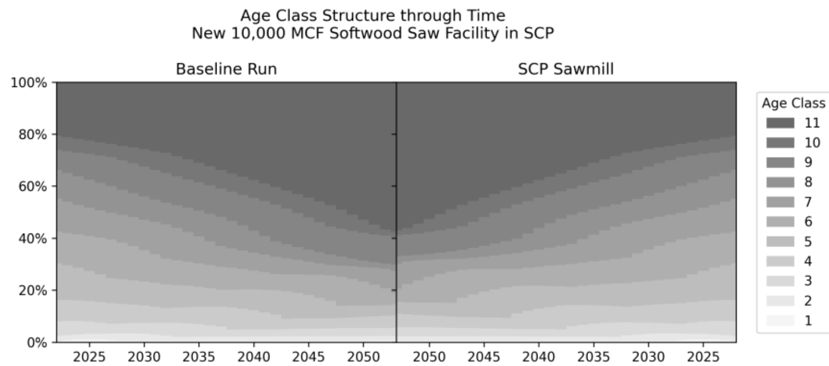

Figure 9. The share of softwood volume (percent of total) in the Southern Coastal Plain by age class through time. LEFT: Baseline scenario. RIGHT: SCP Scenario. Note that the time axis is reversed on the right, providing a comparison of the end of the projection.

In both scenarios, we see a mismatch between the baseline (left) and DC file (right) where the graphs meet. We generally see that the age class skews younger when compared against the baseline. This, coupled with the inventory elasticities, explains the increase in sawtimber prices throughout the projection. Sawtimber becomes scarcer with the added demand pressure from the captive new mill.

Transitional demand scenario

In this case, demand is initially put in the SCP basin via the Direct Change file and then “transitioned" to the broader market. This mimics a simple case where a local shock occurs, a new mill opening, causing a local increase in demand. Eventually, this shock is integrated into the broader market as other actors adapt to new market conditions. In this case, a new softwood sawmill enters the SCP through a 10,000 MCF increase in demand via the DC file. That demand then transitions to the regional market at 1,000 MCF per year for 10 years. The DC file represents the change in a year, so the first year is set up with a 10,000 MCF increase, and then each year from 2028 through 2037 has a 1,000 MCF decrease (see Table 4).

Table 4. DC file for the SCP demand transition.

Year	Basin	SW Pulp	SW Saw	HW Pulp	HW Saw
2027	1	0	10000	0	0
2028	1	0	-1000	0	0
2029	1	0	-1000	0	0
...					
2037	1	0	-1000	0	0

The price path for this new scenario is a hybrid between the regional and the SCP scenarios. Prices start out in the SCP scenario neighborhood and move up towards the regional levels as the model state transitions from a local disturbance to a new market paradigm. Figure 10 shows a rapid increase in price between 2028 and 2037, with prices slightly overshooting the regional market scenario. The gap between the Regional and Transition scenarios is the cost imposed by overharvesting the SCP initially.

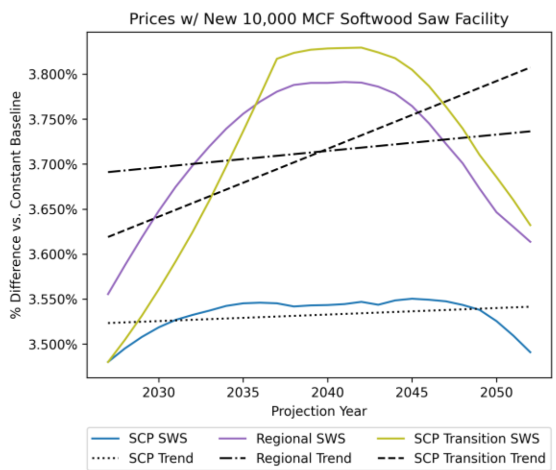

Figure 10. Price trends for the Regional, SCP, and SCP Transition scenarios.

There are initially high removals in the SCP associated with the new mill (see Figure 11). From 2028 to 2032, removals in the SCP fall while removals in the other two basins rise, settling into a removal pattern very similar to the left-most graphic in Figure 7. The inventory versus baseline in the SCP is at or below the inventory level in the Piedmont and the NCP for the entire projection horizon (see Figure 11). Similar to the regional demand scenario, the SCP's inventory vs. baseline is below the NCP for the projection. The key difference being the larger gap that forms in the transitional scenario, indicating that the local shock in the SCP depletes inventory at a faster rate in the early years of the projection. Then, as demand is allowed to flow into other basins, there is some recovery, but not complete recovery. This result is in line with expectations as a large shock to inventory has lingering ramifications for years to come, as suggested in previous literature (Henderson et al. 2022). This result also reflects the efficiency gains as a local shock settles into a new regional equilibrium. As demand transitions into the region, we see removals come into parity, and SCP inventory starts to stabilize and rebound relative to the baseline run.

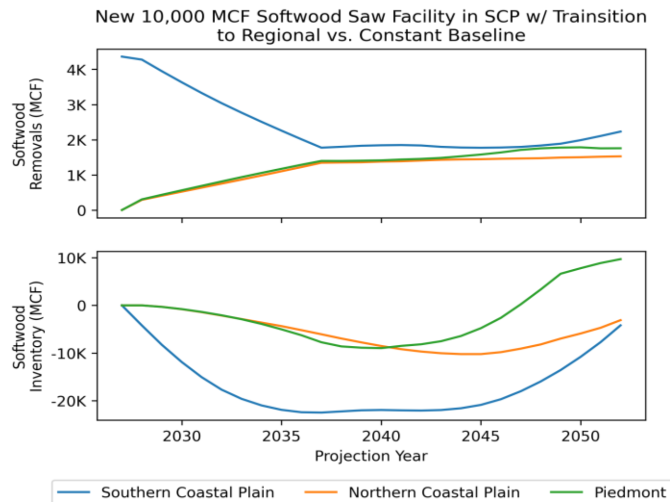

Figure 11. The transitional scenario of a new mill opening in the SCP.

Scenario comparison

Figure 12 captures the trends in total softwood removals and inventory across the four scenarios presented here. The regional and transition scenarios result in the highest observed prices, but they also sustain the highest rates of removal and inventory levels over the long term. The ability to efficiently allocate removals across regions, coupled with a market response to added demand in all basins, drives this result. The SCP scenario results in the lowest rate of removals and inventory through time.

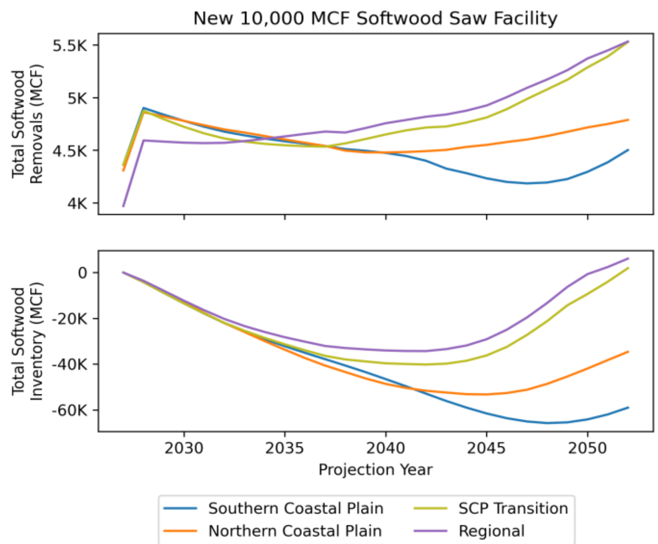

Figure 12. Total softwood removals and inventory across the four scenarios.

Relative to the regional scenario, the lower removals and inventory in the transition scenario, coupled with higher local prices, highlight a loss of efficiency. When the new mill is confined to the NCP, prices in that basin start at over 9% above the baseline price (orange line in Figure 13). Prices fall through time as inventories decline dramatically over the baseline scenario and the basin's age class structure changes dramatically over time. When the mill is in the SCP, that basin experiences prices that are roughly 4% higher than baseline. This price premium is consistent as the age class structure does not swing as dramatically towards older stands in this scenario. In the transitional scenario, where removals move from the SCP to the broader market, we see the same initial local price jump (≈4%), but local prices fall rapidly as demand transitions into the broader region, eventually falling to near 0% difference, in line with the price in the regional scenario (see Figure 10).

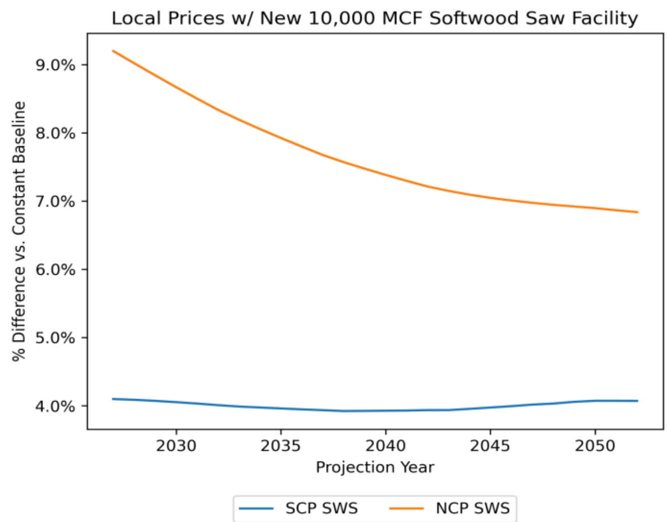

Figure 13. Local prices for the NCP and SCP scenarios.

The results from the new mill analysis highlight a couple of important features of this model and the timber stumpage market. The first is the efficiency of the market. While prices are higher in the regional scenario, so too are removals and inventory levels relative to the other new mill scenarios through time, as higher prices induce management and age class changes that enable higher levels of removal and inventory over time. This leads to more value in the market and on the landscape in the long run. The introduction of local shocks that inhibit the ability of the market to function in a frictionless way imposes a cost on the broader system. Relative to a pure demand-driven scenario analysis, price acts as a mitigating factor in this context, limiting harvests in a period based on inventory/cost sensitivity. If run in harvest mode, the model would remove what is necessary to meet demand, regardless of the price or short-run inventory implications. pySRTS (and SRTS) is not an intertemporal model, and we see the temporal dependency play out. Decisions made in earlier periods have lasting and potentially profound influences through time on the amount harvested and the stock of inventory in the future. The second builds on this, specifically that the transition from local shocks to regional demand can mitigate these costs. Forests grow, and removal patterns will chase volume (when allowed to). A local shock will introduce friction in the market, but economic and ecological forces work together to balance this out through time if demand becomes more mobile. This insight is driven by new functionality that was not possible in the legacy SRTS model.

Comparison with SRTS

To round out our analysis of the new pySRTS model, we compared results with the SRTS model. We compare the results for the constant baseline, regional demand change, and Southern Coastal Plain (SCP) scenarios. We present the SCP version in detail, while the other two are available in Appendix A. Before diving into the price comparison (see Figure 14), there are a couple of expected differences here. First, pySRTS calculates the equilibrium each year, while SRTS approximates the shift in the equilibrium using a log-log transformation. In addition, the prices serve as endogenous weights in the pySRTS model while product weights are fixed in SRTS. The immediately obvious difference in Figure 14 is in Softwood Pulpwood, where the SRTS price is more than $50 per MCF lower by the end of the projection.

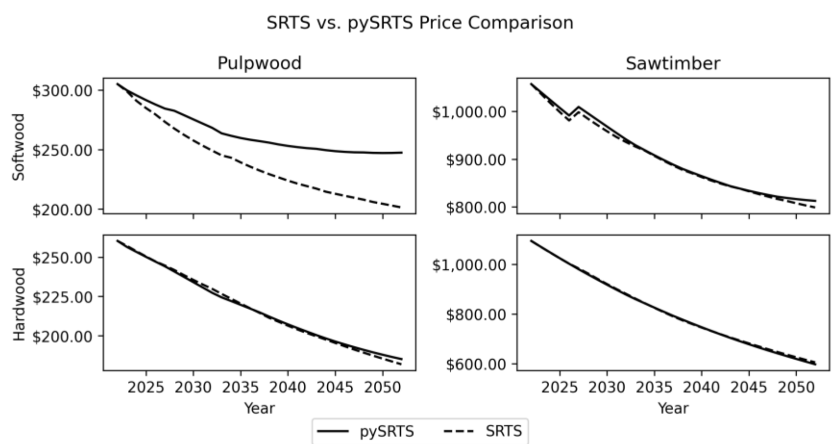

Figure 14. Comparison of SRTS vs. pySRTS projected prices for the Southern Coastal Plain scenario.

To investigate the differences, we examine inventory, demand, and harvests (taking landscape friction into account). To start, the inventory, plotted in Figure 15, is similar across all products except softwood pulpwood, where the pySRTS projections are substantially lower than those from the SRTS model. They begin to diverge after 2035, with the pySRTS volume beginning to level off by the end of the projection, while the SRTS projection continues to grow. This lends intuition to the prices observed in Figure 14, with higher prices in pySRTS where inventory is less available, resulting in a smaller rightward shift in the supply curve.

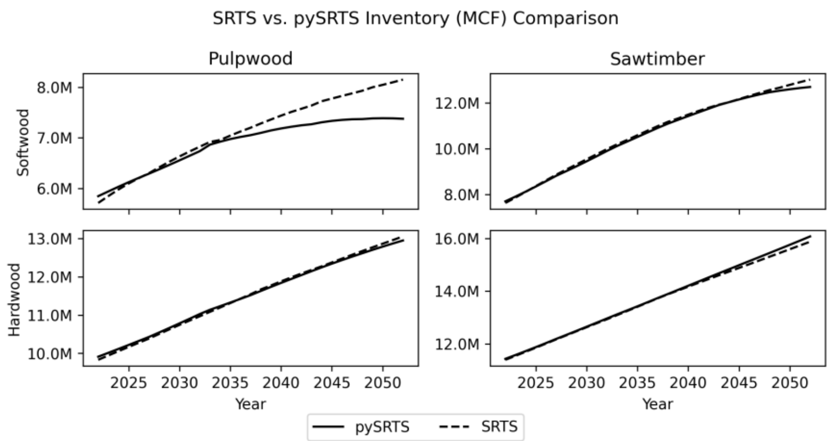

Figure 15. Comparison of SRTS vs. pySRTS inventory by product for the Southern Coastal Plain scenario.

Demand for harvests should then be lower in the pySRTS scenario compared to the SRTS results. In Figure 16, we see that up through 2035 (and a bit beyond), demand between the two scenarios is very similar. But beyond this point, demand for softwood pulpwood harvest in pySRTS starts to lag behind that of SRTS. Again, this follows from prices and inventory presented previously in this section. The remaining question is, what is driving the difference in the first place? The next level down is to investigate actual harvests, the optimization result spreading basin-level supply to the management type and age class level.

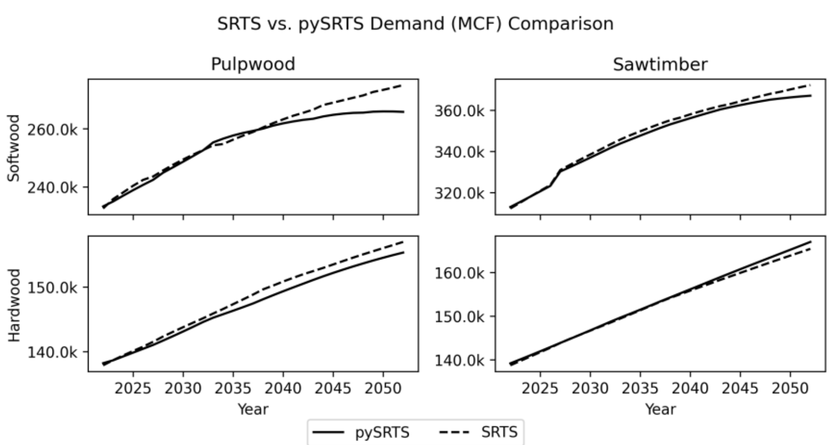

Figure 16. Comparison of SRTS vs. pySRTS harvest demand by product for the Southern Coastal Plain scenario.

At this point, it is important to note that we see similar patterns across the Constant Demand and Regional Demand scenarios. To understand what is driving these differences in model behavior, we compare harvest miss. Harvest miss is the friction that results from mismatches between demand for a product at the basin level and the configuration of availability and constraints on the landscape in that basin. In this case, we see no harvest miss in the pySRTS model, while significant harvest miss occurs in the softwood sawtimber. Figure 17 outlines this trend across the three scenarios.

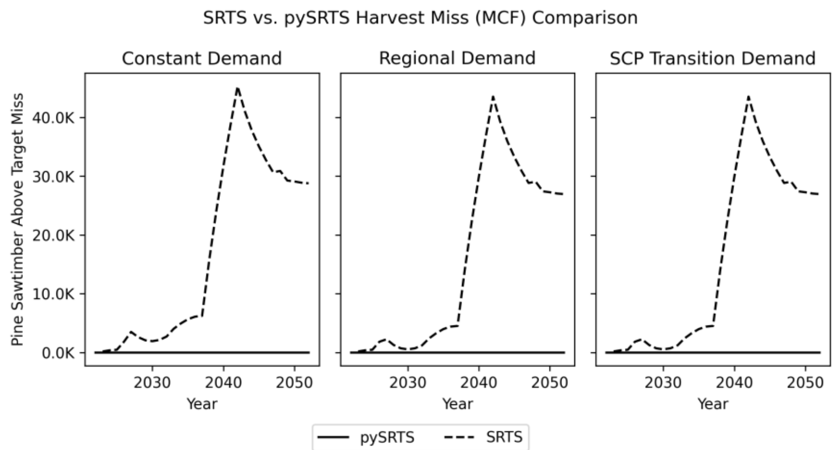

Figure 17. Softwood sawtimber harvest miss comparison between pySRTS and SRTs across the Constant Demand, Regional, and SCP scenarios.

We see a few key points here. First, harvest miss is above the demand for softwood sawtimber removals, indicating more sawtimber is being cut, and more acres are being reset to age zero. This mechanism increases the area and volume in younger stands that are dominated by pulpwood inventory. Second, this over-harvest starts to take off right around the same time that softwood pulpwood inventories between pySRTS and SRTS start to diverge. Lastly, harvest miss is marginally higher in the constant-demand scenario than in the other two. This follows the fact that there is an increase in demand for softwood sawtimber in the latter two cases, and this demand consumes some of the overharvest.

What is driving this overharvest? Hard-coded lower bound constraints in SRTS that apply to age classes 8 (40-45 years old) and 9 (45+ years old) in planted pine. These constraints are designed to prevent inventory from piling up in old-age classes when there is no demand to harvest them. In pySRTS, these lower bounds are configurable and set to allow more of that inventory to accumulate and ultimately drive price down, hence the lower price trend for pine pulpwood present in the legacy SRTS results.

The pySRTS model provides the user finer control over these lower bounds and the ability to set them endogenously based on evolving prices, inventory, or a scientist's curiosity. These lower bounds are hard coded in the legacy model and cannot be modified without changing the codebase itself. Further, the transitional scenario was excluded because there is no straightforward way to implement this use case in the legacy SRTS model. The transition from a local shock to a global one is handled by the year-to-year change in inventory and price. Therefore, the explicit leakage of local shocks into the broader market cannot be modeled in the SRTS frame as it is in pySRTS.

CONCLUSIONS

We set out to update the SRTS model to a more modular, flexible version while maintaining the core competencies of the existing framework. In addition, we explored the influence of new basin-level price premiums on the model's solutions. We find that the addition of local price premiums allows us to study assumptions about market friction and demand mobility through a new lens.

While the pySRTS model is a ground-up rewrite, it does share a lot in common with the original SRTS model. First and foremost, they are both partial equilibrium economic models coupled with a harvest goal program and a forest growth module that advance ecological and economic processes over time. They both allow for activities in the softwood sawtimber market to influence the pulp market. The endogenous land-use change module, brought over from the legacy model, is similar in function. pySRTS extends the original model formulation, as in the residue market, where the user now has control and can introduce additional market linkages.

There are a few key new features in the pySRTS model. The pySRTS model solves the partial equilibrium directly, rather than the change year-to-year. Prices are used as product weights in the goal program, making the weights dynamic, evolving as management and growth influence inventory availability. The introduction of basin-level price premiums as a way to account for local shocks is not present in the legacy model, but this functionality is available in pySRTS.

There are also technical innovations worth noting. The model formulation itself has been set up with flexibility in mind, making it easier to specify pySRTS in new regions without code changes (excluding new functional development). More detailed inventory accounting is available, allowing the end user to audit demand, harvests, growth, and inventory with ease. The modular nature of the code allows users to develop their own functions and modules to suit their specific purposes. The transition to Python also allows for the parallel processing of models, enabling more efficient batch processing when multiple runs are required (e.g., Monte Carlo simulations for uncertainty analysis). There is also cross-platform support for Python, which VB and VBA did not have (i.e., Linux, MacOS).

We show that even under conditions that are outside the ideal economic assumptions, the model performs in an intuitive way. Further, allowing the model to transition to a frictionless market yields insight into the lasting impacts and ramifications of near-term demand shocks on long-run inventory and removals. The comparison with SRTS demonstrates the similarity in results after accounting for differences in bounds and model structure. Effort has been made in the pySRTS model to endogenize parameters and bounds. However, lower bounds on harvest are still largely exogenous, considering only the amount of inventory available and user input by default. This is somewhat by design, as constraints of this nature allow the analyst an opportunity to impart sub-optimal harvest behavior into the model for the purposes of modeling a timbershed. One area of future investigation is embedding declining saw timber efficiency into product densities. If the productivity of an age classes falls in response to the harvest activity, a more natural spreading of harvests could occur, further limiting the influence of bounds.

There is functionality that has not yet been ported over to the pySRTS model, for example, carbon fertilization growth rate multipliers and hurricane damage. These functions are under consideration for future development. The modular nature of the pySRTS model makes their inclusion less critical, as one can adjust the parameterization of the model by customizing the new pipeline architecture rather than relying on bespoke functions and logic that muddies the model's broader context. In addition, the tight integration between SRTS and Excel in recent years has improved usability and adoption of SRTS by a wider user group of forest sector practitioners in industry, government, and non-governmental organizations. This integration with Excel does not currently exist for pySRTS, but is slated for development.

There are several avenues for potential future model development with pySRTS. One logical next step is to codify harvest miss as an economic measure. To this end, the minimization of dead weight loss could be beneficial, but would require some fundamental changes to the model structure and may require branching out a new model altogether. It may also, at that time, make sense to introduce some amount of future foresight into the model, or to provide an intertemporal version of the model. In addition, the use of elasticities from the broader market to parameterize local market contexts could be enhanced with subsequent releases of the model. A set of basin-level elasticities could be built in to mimic the influence of spatial factors and heterogeneous market factors on the aggregate regional elasticity parameters.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the insightful micro-economic discussions with Dr. Daniel Tregeagle.

FUNDING

This research was supported by Southern Forest Resource Assessment Consortium (SOFAC) at NC State University. Results and perspectives expressed in this article are solely the authors and do not necessarily represent those of SOFAC member organizations.

CONFLICTS OF INTEREST

The authors confirm there are no conflicts of interest.

REFERENCES CITED

Abt KL, Abt RC, Galik C. 2012. Effect of bioenergy demands and supply response on markets, carbon, and land use. For Sci. 58(5):523-539. https://doi.org/10.5849/forsci.11-055.

Abt RC, Cubbage FW, Abt KL. 2009. Projecting southern timber supply for multiple products by subregion. For Prod J. 59(7-8):7-16.

Abt RC, Cubbage FW, Pacheco G. 2000. Southern forest resource assessment using the subregional timber supply (SRTS) model. For Prod J. 50(4):25-33.

Bardon R. 2022. North Carolina standing timber price report - Q4 2022. NC State Extension Forestry. Available from: https://forestry.ces.ncsu.edu/documents/887/NCPR4q2022.pdf.

Brown R, Zhang D. 2005. The Sustainable Forestry Initiative's impact on stumpage markets in the US South. Can J For Res. 35:2056-2064. https://doi.org/10.1139/x05-135

Carter DR. 1992. Effects of supply and demand determinants on pulpwood stumpage quantity and price in Texas. For Sci. 38(3):652-660. https://doi.org/10.1093/forestscience/38.3.652

Creekmore LT. 2015. Modeling timber supply from planted forests in selected South American markets: applications of the SRTS modeling framework [master's thesis]. Raleigh (NC): North Carolina State University.

Dhungel G, Rossi D, Henderson JD, Abt RC, Sheffield R, Baker J. 2023. Critical market tipping points for high-grade white oak inventory decline in the Central Hardwood Region of the United States. J For. 121(3):224-234. https://doi.org/10.1093/jofore/fvad005

Forest Inventory and Analysis Database. 2026. St. Paul (MN): U.S. Department of Agriculture, Forest Service, Northern Research Station. Available from: https://apps.fs.usda.gov/fia/datamart/datamart.html.

Galik CS, Abt RC, Latta G, Vegh T. 2015. The environmental and economic effects of regional bioenergy policy in the Southeastern U.S. Energy Policy. 85:335-346. https://doi.org/10.1016/j.enpol.2015.05.018.

Guo Z, Hodges DG, Abt RC. 2011. Forest biomass supply for bioenergy production and its impacts on roundwood markets in Tennessee. South J Appl For. 35(2):80-86. https://doi.org/10.1093/sjaf/35.2.80

Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, Wieser E, et al. 2020. Array programming with NumPy. Nature. 585(7825):357-362. https://doi.org/10.1038/s41586-020-2649-2

Henderson JD, Abt RC, Abt KL. 2024. Forest carbon under increasing product demand and land use change in the US Southeast. For Policy Econ. 167:103296. https://doi.org/10.1016/j.forpol.2024.103296

Henderson JD, Abt RC, Abt KL, Baker J, Sheffield R. 2022. Impacts of hurricanes on forest markets and economic welfare: the case of Hurricane Michael. For Policy Econ. 140:102735. https://doi.org/10.1016/j.forpol.2022.102735

Henderson JD, Parajuli R, Abt RC. 2020. Biological and market responses of pine forests in the US Southeast to carbon fertilization. Ecol Econ. 169:106491. https://doi.org/10.1016/j.ecolecon.2019.106491

Latta GS, Baker JS, Ohrel S. 2018. A land use and resource allocation (LURA) modeling system for projecting localized forest CO2 effects of alternative macroeconomic futures. For Policy Econ. 87:35-48. https://doi.org/10.1016/j.forpol.2017.10.003

Liao X, Zhang Y. 2008. An econometric analysis of softwood production in the US South: a comparison of industrial and nonindustrial forest ownerships. For Prod J. 58(11):69-74.

Manner RH, Guo J, Clutter M, Sodiya OE, Abt RC, Sheffield RM, Baker JS. 2026. Globally informed and locally refined: market, management, and carbon projections for the Southeastern United States forest sector. For Policy Econ. 186:103755. https://doi.org/10.1016/j.forpol.2026.103755

Newman DH. 1987. An econometric analysis of the southern softwood stumpage market: 1950-1980. For Sci. 33(4):932-945. https://doi.org/10.1093/forestscience/33.4.932

pandas-dev-team. 2020. Pandas-dev/pandas: pandas. Zenodo. https://doi.org/10.5281/zenodo.3509134.

Polyakov M, Teeter LD, Jackson JD. 2005. Econometric analysis of Alabama's pulpwood market. For Prod J. 55(1):41-44.

Polyakov M, Wear DN, Huggett RN. 2010. Harvest choice and timber supply models for forest forecasting. For Sci. 56(4):344-355. https://doi.org/10.1093/forestscience/56.4.344

Prestemon JP, Abt RC. 2002. Southern forest resource assessment highlights: the southern timber market to 2040. J For. 100(7):16-22. https://doi.org/10.1093/jof/100.7.16

Restrepo HI, Munro HL, Prisley SP, Schilling E, Radtke PJ, Bush R, Coulston J, Woodall CW. 2025. Population estimate differences for forest biomass, merchantable volume, and carbon pools as a result of modeling system updates in the United States. For Ecol Manag. 595:122965. https://doi.org/10.1016/j.foreco.2025.122965

Robinson VL. 1974. An econometric model of softwood lumber and stumpage markets, 1947-1967. For Sci. 20(2):171-179. https://doi.org/10.1093/forestscience/20.2.171

Sendak PE, Abt RC, Turner RJ. 2003. Timber supply projections for northern New England and New York: integrating a market perspective. North J Appl For. 20(4):175-185. https://doi.org/10.1093/njaf/20.4.175

U.S. Department of Agriculture, Forest Service. 2024. North Carolina state production timber product output survey. National Resource Use Monitoring. Available from: https://research.fs.usda.gov/products/dataandtools/national-resource-use-monitoring-data-downloads.

Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, et al. 2020. SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods. 17:261-272. https://doi.org/10.1038/s41592-019-0686-2

Visser L, Latta G, Pokharel R, Hoefnagels R, Junginger M. 2022. Exploring the limits to sustainable pellet production for international markets: the impact of increasing pellet production in the US Southeast on feedstock use, production cost and carbon sequestration in forest areas. GCB Bioenergy. 14(8):896-917. https://doi.org/10.1111/gcbb.12946

Westfall JA, Coulston JW, Gray AN, Shaw JD, Radtke PJ, Walker DM, Weiskittel AR, et al. 2024. A national-scale tree volume, biomass, and carbon modeling system for the United States. Gen Tech Rep WO-104. Washington (DC): U.S. Department of Agriculture, Forest Service. 104 p. https://doi.org/10.2737/WO-GTR-104

APPENDIX A

Initial prices and conversion

Table A1. Prices and conversion factors used to initialize the model's market

Product	$/Green Ton	Green Ton/MCF	$/MCF
Softwood Pulp	8.79	34.7	304.90
Softwood Saw	30.48		1,057.26
Hardwood Pulp	6.94	37.5	260.25
Hardwood Saw	29.19		1,094.61

Starting inventory

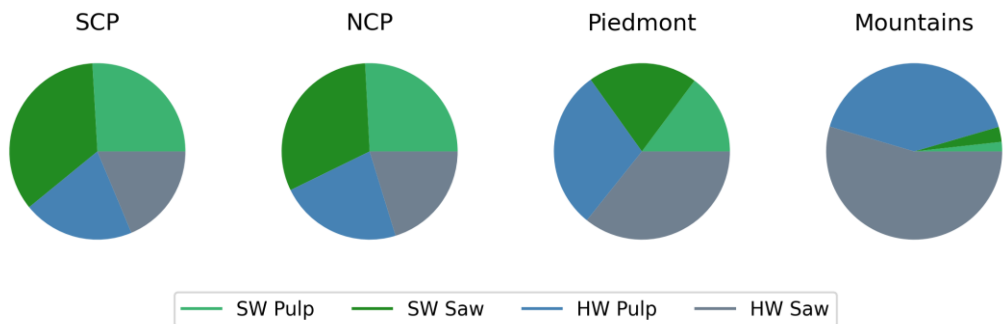

Figure 1A. Starting inventory by basin and product.

SRTS vs. pySRTS Comparison

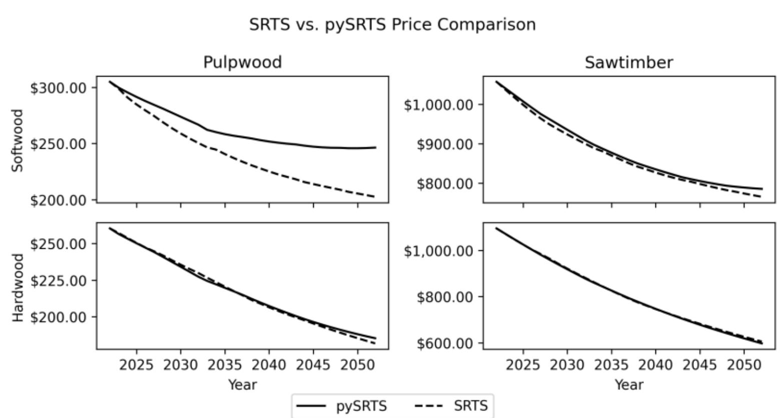

Figure 2A. Comparison of SRTS vs. pySRTS projected prices for the Constant Demand scenario.

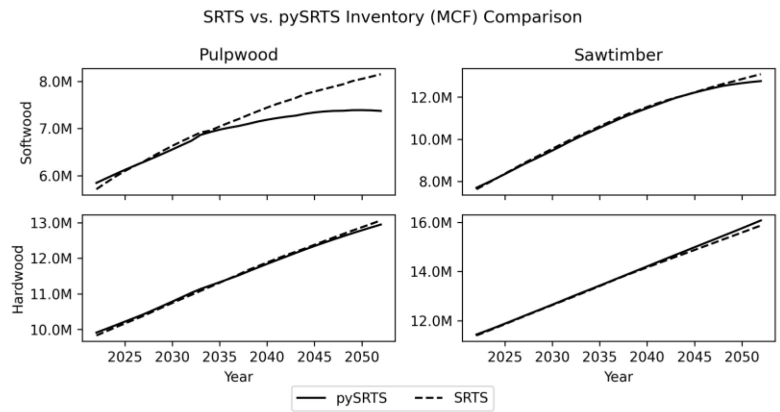

Figure 3A. Comparison of SRTS vs. pySRTS projected inventories for the Constant Demand scenario.

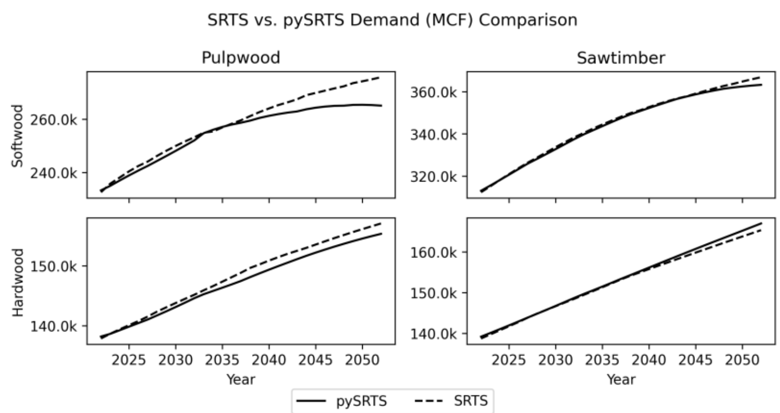

Figure 4A. Comparison of SRTS vs. pySRTS projected demand for the Constant Demand scenario.

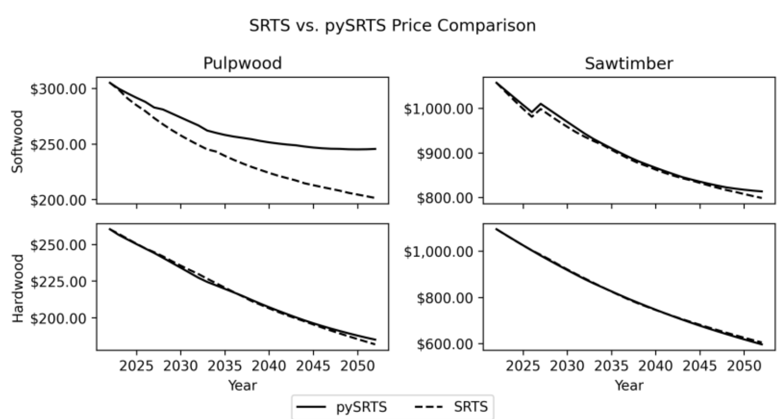

Figure 5A. Comparison of SRTS vs. pySRTS projected prices for the Regional Demand scenario.

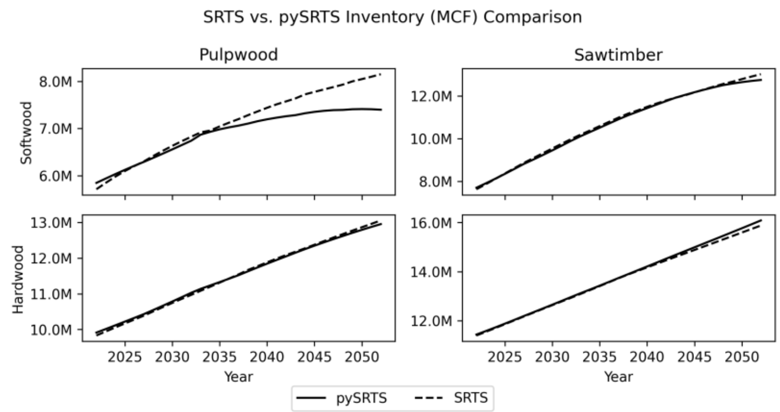

Figure 6A. Comparison of SRTS vs. pySRTS projected inventories for the Regional Demand scenario.

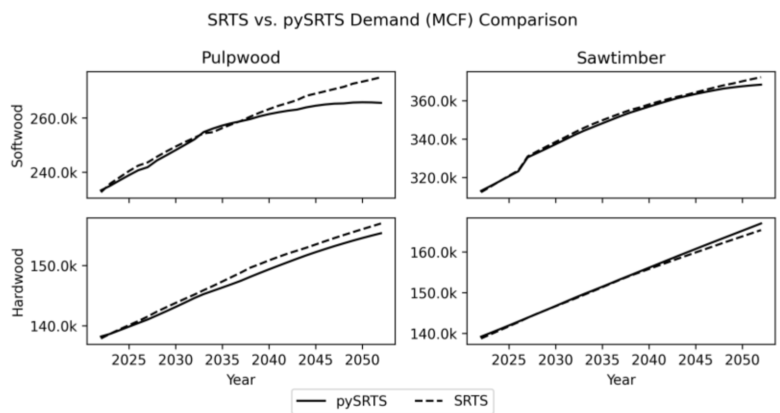

Figure 7A. Comparison of SRTS vs. pySRTS projected demand for the Regional Demand scenario.