CIRRELT
Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport
Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Estimating the Value-Creation Potential of Wood Supply Using a Hybrid Simulation-Optimization Approach

Gregory Paradis
Luc LeBel

May 2018

CIRRELT-2018-24

Bureaux de Montréal :
Université de Montréal
Pavillon André-Aisenstadt
C.P. 6128, succursale Centre-ville
Montréal (Québec)
Canada H3C 3J7
Téléphone : 514 343-7575
Télécopie : 514 343-7121

Bureaux de Québec :
Université Laval
Pavillon Palasis-Prince
2325, de la Terrasse, bureau 2642
Québec (Québec)
Canada G1V 0A6
Téléphone : 418 656-2073
Télécopie : 418 656-2624

www.cirrelt.ca

Estimating the Value-Creation Potential of Wood Supply Using a Hybrid Simulation-Optimization Approach

Gregory Paradis[1,2,*], Luc Lebel[2,3]

1. University of British Columbia, Department of Forest Resources Management, 2424 Main Mall, Forest Sciences Centre 2312, Vancouver, BC V6T 1Z4
2. Département des sciences du bois et de la forêt, Pavillon Abitibi-Price, 2405, rue de la Terrasse, Local 2121, Université Laval, Québec, Canada G1V 0A6
3. Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

Abstract. Current implementation of wood supply optimization models in the province of Quebec, Canada, do not include financial performance indicators. We describe a methodology for compiling a hybrid simulation-optimization model that can be used to estimate the value-creation potential (VCP) of any subset of species-wise annual allowable cut (AAC) volume from the long-term wood supply models used in Quebec. Our model retro-fits financial performance indicators to the optimal solution of the long-term wood supply optimization model, which we link to a network flow optimization model that simulates profit-maximizing fibre consumption behaviour of a network of primary processing facilities. This network flow model can be used to simulate the subset of available fibre supply that a given profit-maximizing industrial network configuration will consume, and provide new insight into the techo-economic factors limiting fibre consumption. If solved repeatedly using different network configurations (e.g., opening or closing one or more facilities, expanding or reducing capacity at exiting facilities, testing different model-exogenous market price assumptions, or testing the impact of different stumpage rate models on network behavior, etc.), the network flow model can be used as a framework to generate scenarios, which can then be compared and analysed. Our methodology uses readily-available wood supply models and input data, and can be applied to any of the 71 management units that make up the managed public forest in the province of Quebec. Thus, we present a methodology that produces state-of-the-art VCP estimates, which could be leveraged to yield new provincial-scale insight into the complex relationships between wood supply modelling and timber licensing policy, stumpage policy, industrial fibre consumption network configurations, and market prices. We run a number of scenarios on management unit UA 064-51, as an example of application of our methodology, and report VCP as a function of the proportion of AAC that is consumed.

Keywords. Forest management, hierarchical planning, value-creation potential.

Acknowledgements. This study was supported by funding from the FORAC Research Consortium.

* Corresponding author: gregory.paradis@ubc.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2018

1 Introduction

Wood supply anlysis plays a central role in strategic forest management planning on public forests. An important output from the wood supply planning process is *annual allowable cut* (AAC), which sets an upper bound on species-wise timber licences that are attributed to industrial fibre consumers for the current planning period. In Quebec, Canada, as in all other Canadian provinces, the provincial government is responsible for natural resource planning, which includes the forest resource. Government analysts in Quebec use a complex software framework to compile and run aspatial *wood supply* models that estimate maximum sustainable species-wise harvest levels (i.e., AAC). The public forest of Quebec (76 million hectares) is divided into 71 *management units* (MU)—one wood supply model is compiled and run for each MU. Wood supply models in Quebec currently do not feature any financial performance indicators.

We use the term *value-creation potential* (VCP) to describe the net financial value of a unit of standing timber (from the perspective of a centrally-managed network of fibre-consuming facilities, i.e., *the network*)—in other words, VCP is an estimation of the marginal profit one might expect a mill owner to make from sale of primary forest products (e.g., lumber, panels, paper, etc.), accounting for all costs and revenues, from standing tree to delivered product. The wood supply models described above estimate maximum sustainable bio-physical fibre output from the forest, regardless of the value-creation potential (VCP) of the fibre in the stands. If a stand has a negative VCP, then harvesting this stand would induce a negative profit (after the timber has been tranformed into primary forest products, these products have been delivered to the market, and all costs and revenues are tallied). A profit-maximizing network will not willingly harvest a stand with negative VCP, unless compelled to do so by some exogenous factor (e.g., contractual or regulatory obligaiton, subsidy, etc.). Thus, we can reasonably assume that only the subset of AAC volume in stands with non-negative VCP will be consumed by the network.

Paradis et al. (2013) show that the classic wood supply model formulation and AAC determination process fail to predict risk of long-term wood supply failures under certain circumstances, even when species-wise harvest volumes are substantially lower than AAC in all planning periods. More specifically, they link risk of unpredicted wood supply failure to incoherence between bio-physical AAC models and industrial fibre consumption behaviour, and conjecture that risk of wood supply failure could be mitigated using an alternative model formulations (that explicitely anticipate industrial fibre consumption behaviour). Paradis et al. (2018) explore this conjecture, and describe a new bilevel wood supply model formulation that mitigates this risk, albeit at a relatively high cost in terms of reduced AAC. To anticipate industrial fibre consumption, they embed a linear network flow optimization model within the existing wood supply model, which simulates the behaviour of a network of profit-maximizing fibre processing facilities, assuming that only fibre with a positive net VCP (for which there is a demand, and matching processing capacity) will in fact be consumed.

In Quebec, the AAC planning cycle length is currently 5 years, and may be increase to 10 years after the end of the current planning cycle (ending in 2018). At this point, it is unlikely that the government will implement a bilevel wood supply modelling workflow by 2018 (modelling for the current planning cycle is already well underway), so we cannot expect the aforementionned risk of wood supply failure to be addressed *a priori* in the wood supply planning process until at least 2028 or later.

Meanwhile, it may nonetheless be beneficial to implement the technical capabilty to estimate VCP of existing wood supply model solutions, and to explore the effect of modifying local industrical processing capacity (e.g., by modelling opening, expansion, or de-comissionning of a one or more facilities) on fibre flow, network profit, and proportion of AAC consumed. We therefore propose a methodology for compiling a hybrid simulation-optimization model that can be used to estimate the value-creation potential of any subset of species-wise annual allowable cut (AAC) volume, using existing wood supply models. Although our modelling framework is conceptually similar to the framework presented in Paradis et al. (2018), the framework presented here is specifically designed to be compatible with the complex wood supply models currently used to determine AAC in Quebec (i.e., our model can import and interpret these models directly, for any of the 71 MUs in the public forest of Quebec). Also, our framework is tightly integrated with the volume-disaggregation and VCP indicator retro-fitting methodology described in Paradis and LeBel (2017b,a, 2018).

Our hybrid simulation-optimization modelling approach makes it possible to perform a *post hoc* analysis of the volume gap between species-wise AAC and anticipated industrial fibre consumption, which can provide government policy makers with heretofore unavailable information to guide the strategic forest management process.

Our overall methology can be described in terms of two *phases*. The first phase links the existing wood supply models to the new network flow model. This involves deriving best-fit statistical distributions of stem diameter to sample plot data (see Paradis and LeBel 2017b), compiling volume disaggregation coefficients (see Paradis and LeBel 2017a), and retro-fitting financial performance indicators to an existing wood supply model solution (see Paradis and LeBel 2018). The second phase, which is the focus of this document, then links the enhanced wood supply model to a network flow optimization model.

We can then use the linked models to simulate fibre-consumption behaviour of a network of primary processing facilities. We solve the network flow optimization model repeatedly, constraining the model to consume an evenly-spaced range of proportions of total AAC. Naturally, our profit-maximizing network flow model consumes the subset of available fibre supply with the highest marginal value-creation potential at each iteration. Plotting output from such a simulation (i.e., total network profit as a function of proportion of AAC consumed) provides insight into the interaction between proposed wood supply and available industrial fibre processing capacity, for a given management unit. If configured to emulate the *status quo* situation in a given region (i.e., existing processing facilities and capacities, fibre supply, costs, values, stumpage policy, etc.), our model can be used to estimate a baseline value-creation potential.

This baseline scenario can also be modified (e.g., add or remove processing facilities, modify lower and upper facility-wise capacity bounds, modify stumpage policy, simulate various forms of commodity-wise subsidies) to derive alternative scenarios for further strategic level analysis.

We deliberately designed our methodology around the latest generation of government wood supply models, using only input data that is readily available for the entire province and accessible to government analysts. Thus, our methodology could hypothetically be integrated into the government workflow. Furthermore, our methodology can be applied to model fibre supply from several several management units in a single model, making regional (or provincial scale) analyses possible. The software framework that is currently in use to model wood supply in Quebec cannot easily be adapted to allow this sort of multiple-scale wood supply analysis.

Finally, we implemented our methodology using only freely-available or open-source software components. This includes development of the `ws3` software library[1], which implements the core wood supply and network flow modelling functions, as well as a host of other functions to help us link the many pieces of this puzzle together.

As an example, we apply our methodology to management unit UA 064-51, and present simulation results for a number of scenarios.

The remainder of this paper is organized as follows. We describe our methodology in §2. Results are presented in §3, followed by discussion in §4 and concluding remarks in §5.

2 Methods

We present the methods in two subsections. In the first subsection we describe the formulation of the network flow optimization model, and provide some information on data sources and methodology for compiling input data for the model. In the second subsection we describe a test case and a number of test scenarios, illustrating the application of our methodology to management unit UA 064-51 in Quebec.

2.1 Network flow optimisation model formulation and compilation

The MERIS database[2] contains data we can use to estimate unit value-creation-potential of any standing tree stem in Quebec. The MERIS database is de-

[1] See `http://ws3.readthedocs.io` for documentation of the `ws3` software library, which is freely downloadable from `http://github.com/gparadis/ws3`. The use-case described here was implemented using Jupyter Notebooks—the notebooks are available from the corresponding authour upon request. Please note that running the notebooks requires specific datasets—although these datasets are readily available upon request, terms of use of these datasets do not allow us to distribute the data directly.

[2] See `https://bmmb.gouv.qc.ca/analyses-economiques/outils-d-analyse/` for more information on MERIS and to download the software and database.

veloped for and maintained by the emphBureau de mise en Marché des bois (BMMB), a branch of the Quebec provincial government responsible for marketing the portion of the wood supply that is sold through a public auction process.

Another branch of the Quebec government, the *Bureau du forestier en chef* (BFEC), compiles and solve a wood supply models to determing species-wise AAC, for each management unit in Quebec. We want to map value-creation–potential data from the MERIS database to each component (i.e., applied action) in the wood supply model solution, for a given MU. Data in the wood supply model is more hightly aggregated than data in the MERIS database. For example, a harvesting action in the wood supply model will specify the development type, the action type, and the treated area. Applying this action will generate an array of outputs, including net merchantable harvest volume (expressed in m^3, by species group). There are 11 standard species groups in wood supply models (see Paradis and LeBel 2018, Table A1).

The MERIS database stores value data using several different levels of aggregation. The finest level of data aggregation used in the MERIS database (for the data we need) expresses unit VCP in terms of a combination of one of 2 processor profiles (hardwood, softwood), 45 species codes, 26 stem size classes, and 6 product classes. In step 1 of phase 1 of this project (see Paradis and LeBel 2017a for details), we compiled *disaggregation coefficients* that can be used to disaggregate (i.e., explode) each unit of harvested volume output from a BFEC wood supply model solution into smaller volume sub-units. Each volume sub-unit exactly matches the finest aggregation level used in the MERIS database, thus allowing us to retrofit MERIS value-creation-potential data to BFEC wood supply model solutions.

In this context, the *value* of the stem of a standing tree is defined from the perspective of a network of primary processing facilities seeking to maximize profit from sale of primary products (and primary co-products) to external markets. Given this definition of value, unit VCP is equivalent to the sum of unit costs and revenues along a given trajectory from standing tree to delivered primary product.

We designed a linear programming (LP) optimization model formulation that can emulate profit-maximizing fibre consumption behaviour of any hypothetical network of primary fibre processing facilities. We use a network flow model design pattern (i.e., a directed graph of processing nodes connected by transportation arcs). Fibre flows through the network from *source* (`src`) nodes to emphsink (`snk`) nodes. We model three additional layers of nodes between source and sink nodes,*viz.*, *dispatch* (`dsp`), *commodity* (`cdt`), and *processor* (`prc`) nodes, for a total of five layers of nodes. Figure 1 gives a schematic representation of the topology of arcs and nodes in our network model.

We model a 1:1 mapping between source nodes and decisions in the harvest schedule we import from the BFEC wood supply model. In step 2 of phase 1, we use the `ws3` software library to help us retrofit data from the MERIS database to the BFEC wood supply model, which allows us now to programatically interface the wood supply model with the network flow model. The integration of

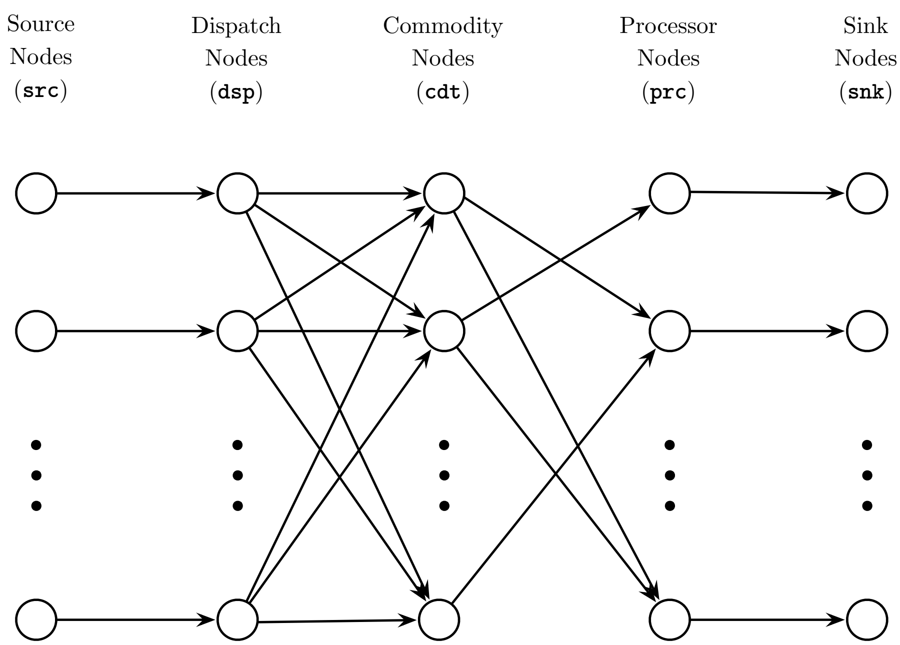

Figure 1: Schematic representation of network flow model.

this software interface of these two model components within the common `ws3` modelling framework is key to keeping this complex techical task manageable.

There is a 1:1 mapping between source nodes and dispatch nodes. The number of source and dispatch nodes will vary, depending on the number of harvesting actions in the wood supply model solution. Each dispatch node has one inbound arc—flow along source-dispatch arcs represents *harvested area* (measured in ha). By setting appropriate upper and lower flow capacity bounds on these arcs, we can simulate harvesting and consumption of any subset of the Woodstock solution. In step 2 of phase 1, we compiled a net VCP coefficient for each harvesting decision in the Woodstock model—these coefficients represent the sum of all revenues and costs, from standing tree to delivered primary product, expressed on the basis of harvested area (i.e., $\$ \cdot \text{ha}^{-1}$). The objective function value of our network flow model is defined as the scalar product of source-dispatch flow and net VCP coefficient vectors.

Each dispatch node has one outbound arc connecting it to each of 12 commodity nodes. The notion of *commodity* is used in the MERIS database as an intermediate aggregation scheme combining species group and product class. Outbound flow from dispatch nodes represents *harvested volume* (measured in m^3). Thus, dispatch nodes convert harvested area to commodity volumes. Unit stumpage cost in the MERIS database is defined as a function of stumpage zone and commodity, so we can map coefficients to dispatch-commodity flows representing unit stumpage cost.

Commodity nodes have one inbound arc for each dispatch node, and one outbound arc for each processing node. Outbound flows represent volume of merchantable fibre (in m^3). The number of processing nodes will vary from one instance to another, depending on the number of fibre processing facilities modelled. For our test dataset, we model fibre flows from management unit UA 064-51. We used publicly-available government documentation of timber licence (TL) agreements to compile the list of processing facilities that procure fibre directly from the target management unit. Since 2013, TL volumes in Quebec are attributed on a regional basis (i.e., for a set of management units). Our test dataset models fibre flows from a single management unit, so we need to disaggregate regional TL volumes to define appropriate upper bounds on flows between commodity and processing nodes. We use the most recent pre-2013 (management-unit-wise) TL volume data, for the set of management units in the target region, to estimate commodity-wise proportion of regional TL volumes corresponding to our target management unit, which we use as upper bounds on commodity-processor arc flows in our network model.

Note that sawmill processing nodes generate a substantial volume of chips as a co-product. For each sawmill processor node, we add an outbound arc to each pulpmill processor node so that sawmill chip co-products can flow to pulpmills.

We model one sink node for each processor node, with matching arcs. Facility capacity constraints can be defined as flow bounds on these arcs.

Note that we chose this formulation for the model because it maps directly onto the data aggregations used in the wood supply model and the MERIS database. This simplified software implementation and use of the model, but may limit potential extensions of the model. For example, convergence of volume flow at commodity nodes implies loss of *traceability* of volume between specific stand-mill combinations. We do not need this level of traceability for our case study, but this might limit future use of this network model implentation. Other, more general, network model formulations could easilty be implemented within `ws3`-based software framework, and still leverage the rest of the functionality of the framework (including the methodology and software functions we developed to retro-fit VCP indicators to existing wood supply models).

The mathematical formulation of the network flow model is given by

$$
\begin{aligned}
\max \quad & \sum_{s\in S} c_s^{\mathrm{SD}} f_s^{\mathrm{SD}} \\
+ \quad & \sum_{d\in D}\sum_{c\in C} c_{dc}^{\mathrm{DC}} f_{dc}^{\mathrm{DC}} \\
+ \quad & \sum_{c\in C}\sum_{p\in P} c_{cp}^{CP} f_{cp}^{\mathrm{CP}} \\
+ \quad & \sum_{p\in P} c_p^{PK} f_p^{\mathrm{PK}}
\end{aligned} \tag{1}
$$

$$
\begin{aligned}
\text{s.t.} \quad & b_{dd}^{\mathrm{SD\_L}} \le f_{dd}^{\mathrm{SD}} \le b_{dd}^{\mathrm{SD\_U}}, & \forall d \in D \quad & (2)\\
& b_{dc}^{\mathrm{DC\_L}} \le f_{dc}^{\mathrm{DC}} \le b_{dc}^{\mathrm{DC\_U}}, & \forall d \in D, c \in C \quad & (3)\\
& b_{cp}^{\mathrm{CP\_L}} \le f_{cp}^{\mathrm{CP}} \le b_{cp}^{\mathrm{CP\_U}}, & \forall c \in C, p \in P \quad & (4)\\
& b_{pp}^{\mathrm{PK\_L}} \le f_{pp}^{\mathrm{PK}} \le b_{pp}^{\mathrm{PK\_U}}, & \forall p \in P \quad & (5)\\
& b_{pp'}^{\mathrm{PP\_L}} \le f_{pp'}^{\mathrm{PP}} \le b_{pp'}^{\mathrm{PP\_U}}, & \forall p \in P \quad & (6)\\
& v_{dc}^{SC} f_{dd}^{\mathrm{SD}} = f_{dc}^{\mathrm{DC}}, & \forall d \in D, c \in C \quad & (7)\\
& \sum_{d\in D} f_{dc}^{\mathrm{DC}} = \sum_{p\in P} f_{cp}^{\mathrm{CP}}, & \forall c \in C \quad & (8)\\
& \sum_{c\in C} f_{cp}^{\mathrm{CP}} = f_{pp}^{\mathrm{PK}} + \sum_{p'\in P} f_{pp'}^{PP}, & \forall p \in P \quad & (9)\\
& \sum_{p'\in P} f_{pp'}^{PP} = \alpha_p \sum_{c\in C} f_{cp}^{\mathrm{CP}}, & \forall p \in P \quad & (10)\\
& A^{\mathrm{L}} \le \sum_{s\in S} f^{SD} \le A^{\mathrm{U}} & \quad & (11)
\end{aligned}
$$

where

$$
\begin{aligned}
S &:= \text{set of source nodes} \\
D &:= \text{set of dispatch nodes} \\
C &:= \text{set of commodity nodes} \\
P &:= \text{set of processor nodes} \\
K &:= \text{set of sink nodes} \\
f_{sd}^{\mathrm{SD}} &:= \text{[variable] flow between nodes } s \in S \text{ and } d \in D \text{ (area [hectares])} \\
f_{dc}^{\mathrm{DC}} &:= \text{[variable] flow between nodes } d \in D \text{ and } c \in C \text{ (volume [cubic meters])} \\
f_{cp}^{\mathrm{CP}} &:= \text{[variable] flow between nodes } c \in C \text{ and } p \in P \text{ (volume [cubic meters])} \\
f_{pk}^{\mathrm{PK}} &:= \text{[variable] flow between nodes } p \in P \text{ and } k \in K \text{ (volume [cubic meters])} \\
f_{pp'}^{\mathrm{PP}} &:= \text{[variable] secondary hardwood chip flow between nodes } p \in P \text{ and } p' \in P \\
&\quad\ \text{(volume [cubic meters])} \\
c^{\mathrm{SD}} &:= \text{objective function coefficient for flow variable } f_{sd}^{\mathrm{SD}} \\
&\quad\ \text{(default value: total net value-creation potential)} \\
c^{\mathrm{DC}} &:= \text{objective function coefficient for flow variable } f_{dc}^{\mathrm{DC}} \text{ (default value: 0)} \\
c^{\mathrm{CP}} &:= \text{objective function coefficient for flow variable } f_{cp}^{\mathrm{CP}} \text{ (default value: 0)} \\
c^{\mathrm{PK}} &:= \text{objective function coefficient for flow variable } f_{pk}^{\mathrm{PK}} \text{ (default value: 0)} \\
b_{sd}^{\mathrm{SD\_L}} &:= \text{lower bound for flow variable } f_{sd}^{\mathrm{SD}} \text{ (default value: 0)} \\
b_{dc}^{\mathrm{DC\_L}} &:= \text{lower bound for flow variable } f_{dc}^{\mathrm{DC}} \text{ (default value: 0)} \\
b_{cp}^{\mathrm{CP\_L}} &:= \text{lower bound for flow variable } f_{cp}^{\mathrm{CP}} \text{ (default value: 0)} \\
b_{pk}^{\mathrm{PK\_L}} &:= \text{lower bound for flow variable } f_{pk}^{\mathrm{PK}} \text{ (default value: 0)} \\
b_{pp'}^{\mathrm{PP\_L}} &:= \text{lower bound for flow variable } f_{pp'}^{\mathrm{PP}} \text{ (default value: 0)} \\
b_{sd}^{\mathrm{SD\_U}} &:= \text{upper bound for flow variable } f_{sd}^{\mathrm{SD}} \\
&\quad\ \text{(default value: harvest area from Woodstock model)} \\
b_{dc}^{\mathrm{DC\_U}} &:= \text{upper bound for flow variable } f_{dc}^{\mathrm{DC}} \text{ (default value: inf)} \\
b_{cp}^{\mathrm{CP\_U}} &:= \text{upper bound for flow variable } f_{cp}^{\mathrm{CP}} \text{ (default value: inf)} \\
b_{pk}^{\mathrm{PK\_U}} &:= \text{upper bound for flow variable } f_{pk}^{\mathrm{PK}} \text{ (default value: inf)} \\
v_{sc}^{SC} &:= \text{volume of commodity } c \in C \text{ generated by harvesting one unit of area} \\
&\quad\ \text{from source node } s \in S \\
\alpha_p &:= \text{proportion of volume processed at node } p \in P \text{ that is redirected to} \\
&\quad\ \text{outbound hardwood chip arcs} \\
A^{\mathrm{L}} &:= \text{lower bound on total harvested area} \\
A^{\mathrm{U}} &:= \text{upper bound on total harvested area}
\end{aligned}
$$

The objective function (1) expresses the sum of products of each type of flow variable and corresponding coefficients. By default, only f_{sd}^{SD} flows are assigned non-zero coefficient values, which correspond to total net VCP (measured in $\$ \cdot \mathrm{ha}^{-1}$). Constraints (2) through (6) express lower and upper bounds on flow variables. Constraints (7) through (9) enforce flow conservation in the network. Flow between source and dispatch nodes (i.e., f_{sd}^{SD} variables) represent harvested area. All other flows represent commodity volume. Conversion from area to volume occurs at the dispatch nodes, by way of flow conservation constraints (7). Hardwood sawmills may produce seconary hardwood chip flows. Each hardwood sawmill processor node has outbound arcs to all hardwood chip processing nodes. Total outbound secondary hardwood chip flow from a given hardwood sawmill node is defined as a proportion of total inbound flow at that node, given by constraint (10). Note that secondary hardwood chip flow volume is accounted for in processor node flow conservation constraints (9). Total harvested area bounds can (optionally) be set using constraint (11). Most of the scenarios in our case study output value-creation potential as a function of the proportion of area harvested. We use constraint (11) to implement this (in combination with a loop that iteratively solves the model).

Figure 2 provides a schematic overview of the various steps in our modelling methodology.

We divide the methodology into two phases. Phase 1 isolates the process of compiling input data for the network flow model. Phase 2 isolates the process of compiling and running the network flow model.

We divide phase 1 into two steps. The most complex task in phase 1 is retro-fitting value-creation-potential performance indicators to the wood supply models that we imported into `ws3`. Much of the financial data used to compile the VCP indicators comes from the MERIS database described earlier. Step 1 involves compiling *volume disaggregation coefficients*, which we can use to link the (highly aggregated) wood supply model output to the (hightly detailed) data in the MERIS database. Each volume disaggregation coefficient vector is compiled from three components: a stem diameter distribution vector, a diameter-class-wise harvest probability vector, and a diameter-class-wise stem form factor vector (see Paradis and LeBel 2017a for a detailed description of the methodology used in step 1 of phase 1). Step 2 involves retro-fittting VCP performance indicators to the wood supply model harvest schedule—we import one (or more) wood supply models into `ws3`, disaggregate harvest volume, link financial data from the MERIS database (and other sources), re-aggregate harvest volume, and inject the new VCP indicators into the wood supply model (see Paradis and LeBel 2018 for a detailed description of the methodology used in step 2).

We divide phase 2 into two steps. Step 1 involves importing the augmented wood supply model and other input data (timber licence volumes, facility-wise capacity bounds, etc.), compiling network flow model, and defining scenarios. Step 2 involves running scenarios, compiling output (figures, tables, etc.), and analyzing output.

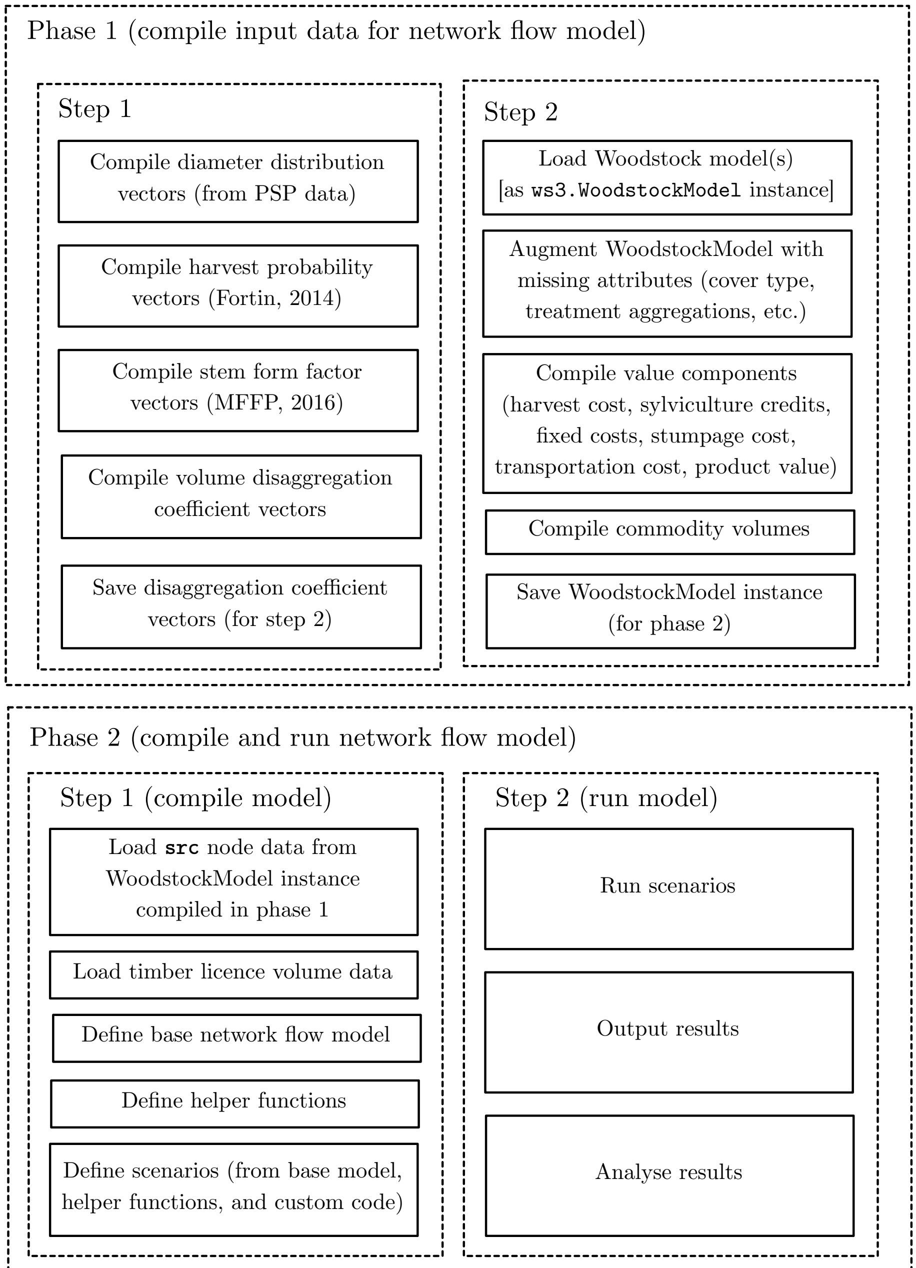

Figure 2: Schematic representation of modelling methodology.

2.2 Description of illustrative test case and scenarios

To demonstrate the application of our framework, we compiled a test model based on management unit UA 064-51. First provide some background information describing UA 064-51, summarizing the harvest schedule we import from the wood supply optimization model, and comparing these volumes to GA allocations. Next, we describe a number of scenarios that we ran using our hybrid simulation-optimization model.

Note that we present these scenarios for illustrative purposes—to show the type of analysis that our methodology and companion software framework can enable—rather than as definitive simulations leading to policy recommendations.

2.2.1 Area description

Our illustrative test case is based on management unit UA 064-51, which is located in the Laurentides region in Quebec. We use a Woodstock wood supply model compiled by BFEC as the starting poitn for our analysis. This model includes a harvest schedule, which is the optimal solution to an even-flow AAC-maximization problem that was solved by BFEC analysts using the Woodstock modelling software. The harvest schedule for the first 5 planning periods is composed of approximately 30000 lines, with each line prescribing simulation a specified *action* to an area (in ha) of a given *development type* (i.e., combination of theme values and age class). In a previous phase of this project, we retro-fitted financial performance indicators (VCP cost and revenue components) to every line of this harvest schedule that corresponds to a harvesting action (see Paradis and LeBel 2018 for a detailed description of the methodology we developed to do this).

Figures 3 and 5 summarize volume from this harvest schedule in terms of cover type, species group, treatment type, and stem size class. This representation includes some disaggregation and reaggregation (relative to the original harvest schedule, which is expressed on an area basis), which we compiled using the `ws3` modelling platform. Harvest volume data on the ordinate axis of all subfigures is shown in the same range (i.e., 0 to 20000).

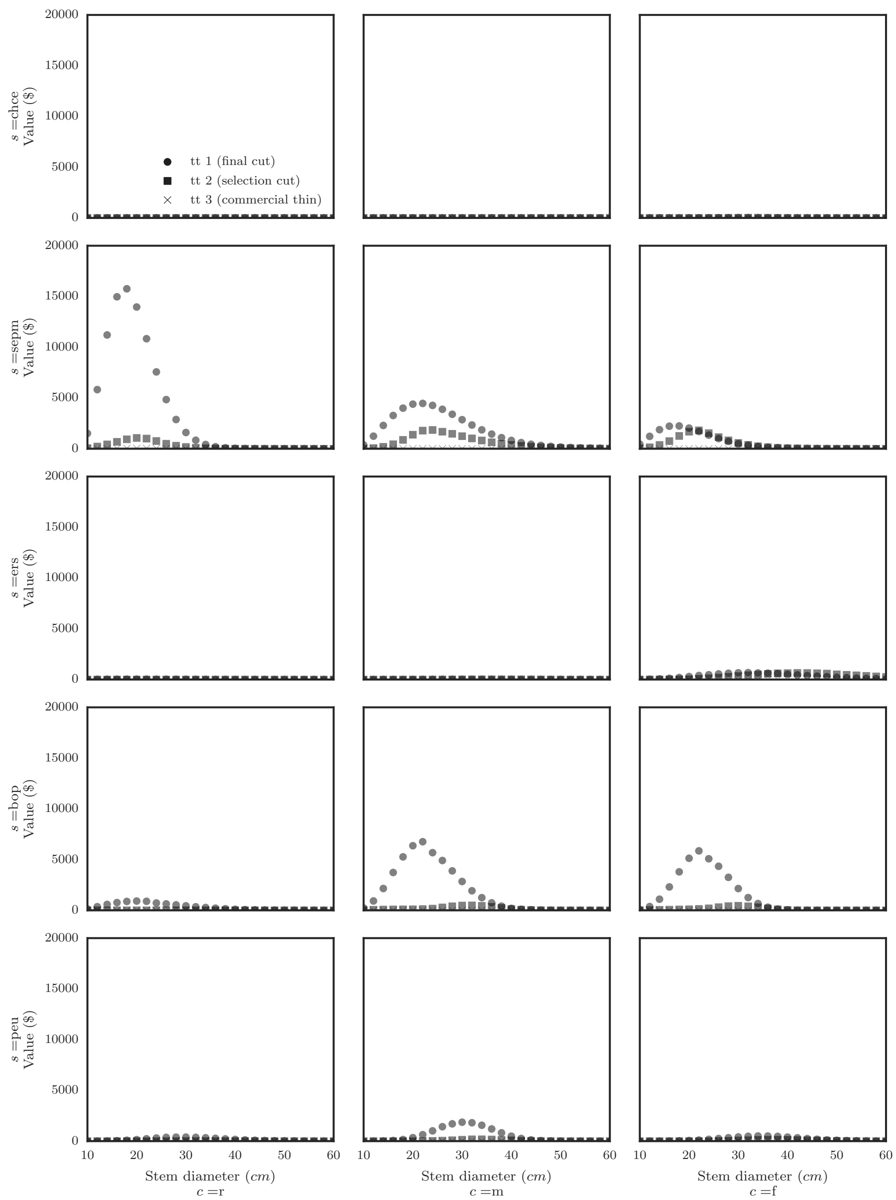

Figure 3: Harvest schedule volume compiled for management unit UA 064-51 in Quebec, Canada. Species group is fixed for a given row of subfigures, and cover type is fixed forn a given column of subfigures. Treatment type 1 (circles) corresponds to clearcut harvesting, treatment type 2 (squares) corresponds to selection cut, and treatment type 3 (crosses) corresponds to commercial thinning.

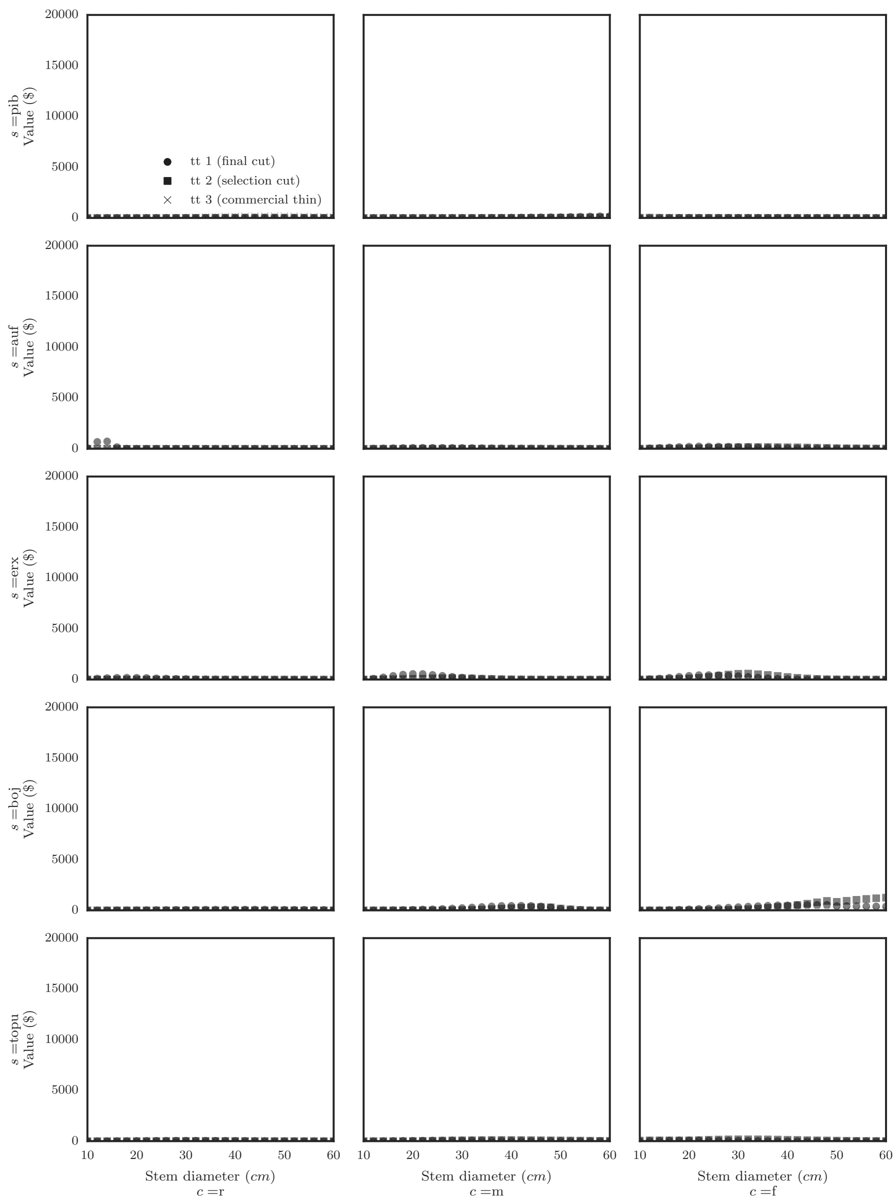

Figure 4: [Continued from Figure 3] Harvest schedule volume compiled for management unit UA 064-51 in Quebec, Canada. Species group is fixed for a given row of subfigures, and cover type is fixed forn a given column of subfigures. Treatment type 1 (circles) corresponds to clearcut harvesting, treatment type 2 (squares) corresponds to selection cut, and treatment type 3 (crosses) corresponds to commercial thinning.

Note that most of the harvest volume concentrated in a relatively small number of bins (i.e., combinations of cover type, species group, treatment type, and stem size class)—specifically, the *sepm* species group (all cover types) *bop* species group (mixedwood and hardwood cover types) contain the bulk of harvested volume.

Figure 5 shows aggregated harvest schedule in terms of three commodity groups—SPFL sawlogs and pulpwood (*SPFL*), hardwood pulpwood (*HW pulp*), and other fibre (*Other*)—split into GA volume (i.e., fibre allocated in a timber license contract) and unallocated volume. Note that a smaller proportion of hardwood pulpwood (*HW pulp*) is allocated than the other commodity aggregates. This situation is not unique to management UA 064-51—supply of low-value hardwood fibre in Quebec (and several other other Canadian provinces) systematically exceeds market demand and industrial processing capacity. The effect of this imbalance between low-value hardwood supply and demand on long-term wood supply stability is discussed in Paradis et al. (2013) and Paradis et al. (2018).

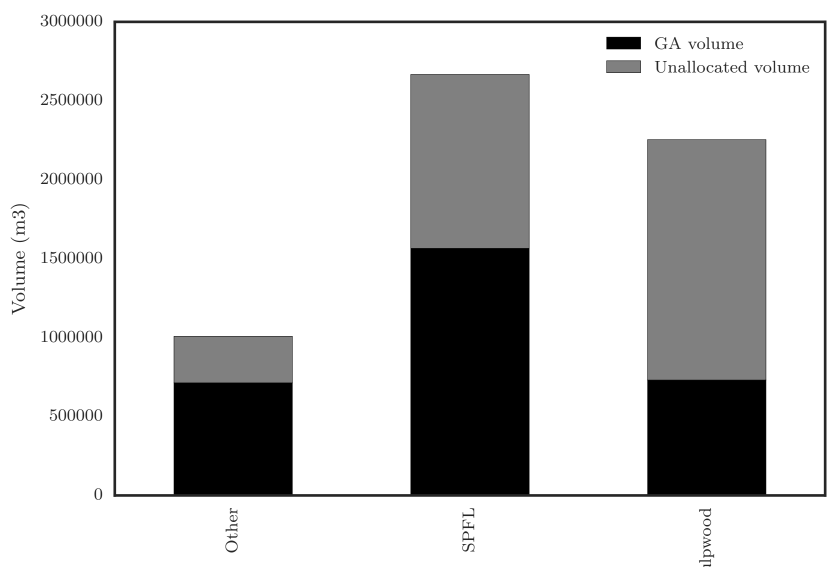

Figure 5: Harvest schedule volume by aggregated commodity, split into GA and unallocated volume.

2.2.2 Scenarios

We ran 4 scenarios using our test case (one base scenario, and three alternate scenarios), to illustrate the capabilities of our hybrid simulation-optimization modelling platform for exploring how strategic forest management policy can affect industrial fibre consumption behaviour. Specifically, we experiment with different hypothetical model-exogenous commodity-wise unit *subsidies*.

Note that this is only a simple example of potential functionality of our new modelling platform. This list of scenarios is by no means definitive—our objective here is simply to show an application of our methodology on a real dataset, to illustrate potential usefulness of our `ws3`-based platform for building complex decision-support modelling interfaces to help strategic forest-sector policy-makers tackle previously inaccessible problems.

Scenario 1 repeatedly solves the base model, constraining harvest area to evenly-spaced values between 0 and 100% of total available area (from the long-term wood supply model solution).

Note scenarios 2.1, 2.2, and 2.3 (presented below) all build on scenario 1.

Scenario 2.1 adds subsidies on certain hardwood pulpwood commodity-processor combinations—subsidy amounts are based on Belzile and Riopel (2015). Scenario 2.2 adds a 10$/m^3 subsidy to the SPFL commodity, as lobbied by the *Conseil de l'industrie forestière du Québec* (CIFQ) in 2015 following the release of a report documenting a recent increase in unit fibre procurement cost in Quebec (see Tremblay 2015). Scenario 2.3 combines the hardwood pulpwood and SPFL subsides from scenarios 2.1 and 2.2.

3 Results

This section presents aggregated simulation results from the scenarios described in the previous section. We first present the base scenario (scenario 1), followed three groups of alternate sub-scenario series (2, 3, and 4).

3.1 Scenario 1

Figure 6 shows VCP by aggregated commodity, as a function of proportion of total available area harvested, for scenario 1. This is the base scenario, so subsidies have been applied to the system to affect simulated fibre consumption behaviour. Initial VCP (at x=0) is negative—this is due to our simulation of the *rente* policy, which requires industrial fibre consumers to pre-pay of 20% of total stumpage fees associated with their GA allocation before they can begin harvesting. VCP for *SPFL* and *HW pulp* commodity aggregates is negative for all proportions of area harvested, and that the *Other* aggregate is only marginally positive. Note that the absolute value of the marginal loss on SPFL volume in the base scenario is less than the average unit stumpage fee—thus, it may be possible to at least break even on SPFL volume by reducing stumpage fees (which, in principle, are supposed to reflect the market value of the fibre rather than induce a net loss). Total VCP is negative for all harvest area values, and that rate of decline of total VCP increases as a function of area harvested—this is expected behaviour, as our profit-maximizing industrial agent preferrentially harvests stands with the most positive (or least negative) marginal VCP, leaving the least profitable stands for last. Assuming that payment of the *rente* has already been made, the loss-minimizing optimal solution (from the perspective of the network) is to harvest 18% of total available area.

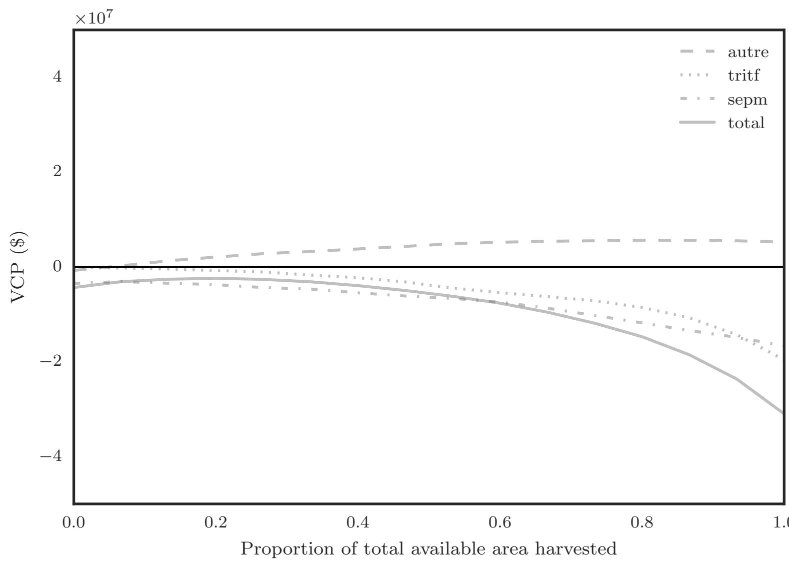

Figure 6: *Scenario 1.* VCP by aggregated commodity, as a function of proportion of AAC harvested. Dash-dot line represents SPFM sawlogs and pulpwood, dotted line represents hardwood pulpwood, and dash-dot line represents all other commodities.

3.2 Scenario 2

Scenarios 2.1, 2.2, and 2.3 model various combinations of subsidies, as an example of how industrial timber consumption behaviour could potentially be manipulated through deliberate policy targetting key VCP parameters.

Figure 7 shows VCP by aggregated commodity, as a function of proportion of total available area harvested, for scenario 2.1. This scenario builds on scenario 1, adding a 10 $\$ \cdot \mathrm{m}^{-3}$ subsidy to hardwood pulpwood flowing from the forest to the Fortress mill, a 25 $\$ \cdot \mathrm{m}^{-3}$ subsidy to white and yellow birch pulpwood flowing from the forest to the Domtar mill. Note that these subsidies do not apply to chip and residue co-products flowing from sawmills to pulpmills. VCP for the *HW pulp* commodity aggregate is now positive for all harvest area proportions, although the slightly negative slope in the later part of the curve indicates that the hardwood pulp component of some stands remains negative despite the subsidy, and that our model is harvesting these stands last (as expected). VCP for *SPFL* and *Other* commodity aggregates are not affected by this policy, and behave similarly to scenario 1. Total VCP is now positive for all harvest area proportions, although it peaks at 64% of total available area—a profit-maximizing industrial network would not willingly harvest the last 36% of available area.

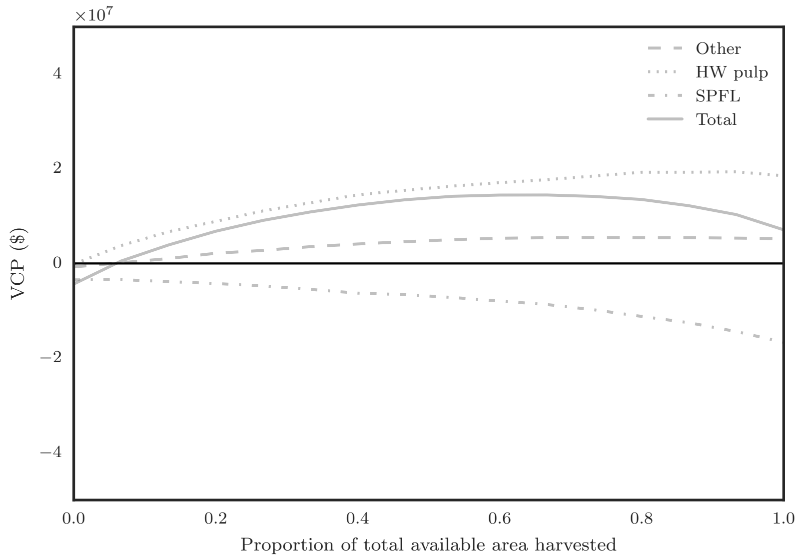

Figure 7: *Scenario 2.1.* VCP by aggregated commodity, as a function of proportion of AAC harvested. Dash-dot line represents SPFM sawlogs and pulpwood, dotted line represents hardwood pulpwood, and dash-dot line represents all other commodities.

Figure 8 shows VCP by aggregated commodity, as a function of proportion of total available area harvested, for scenario 2.2. This scenario builds on scenario 1, adding a 10 $\$ \cdot \mathrm{m}^{-3}$ subsidy to all SPFL volume. VCP for the *SPFL* commodity aggregate is now positive for all harvest area proportions, although the slightly negative slope in the later part of the curve indicates that the SPFL component of some stands remains negative despite the subsidy, and that our model is harvesting these stands last (as expected). VCP for *HW pulp* and *Other* commodity aggregates behave similarly to scenario 1. Total VCP is now positive for all harvest area proportions, although it peaks at 52% of total available area—a profit-maximizing industrial network would not willingly harvest the last 48% of available area.

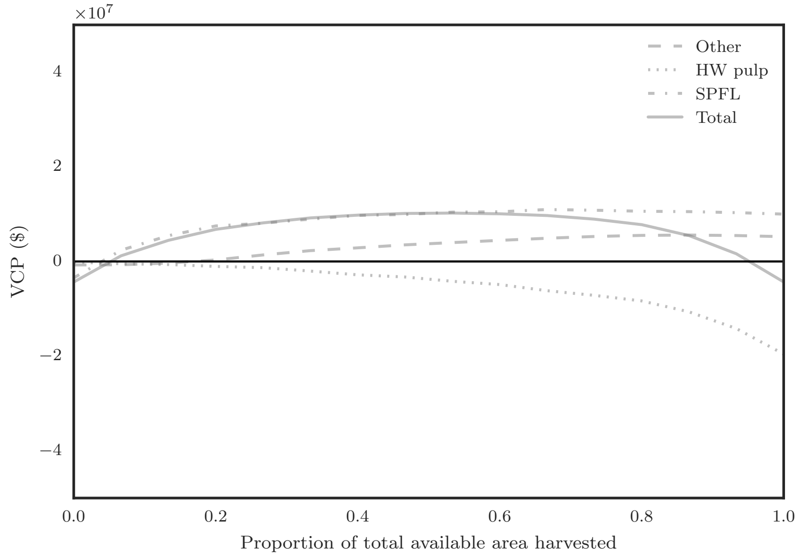

Figure 8: *Scenario 2.2.* VCP by aggregated commodity, as a function of proportion of AAC harvested. Dash-dot line represents SPFM sawlogs and pulpwood, dotted line represents hardwood pulpwood, and dash-dot line represents all other commodities.

Figure 9 shows VCP by aggregated commodity, as a function of proportion of total available area harvested, for scenario 2.3. This scenario combines subsidies from scenarios 2.1 and 2.2. Total VCP is now positive and monotonically increasing until 90% of available area is harvested—a profit-maximizing industrial network would not willingly harvest the last 10% of available area.

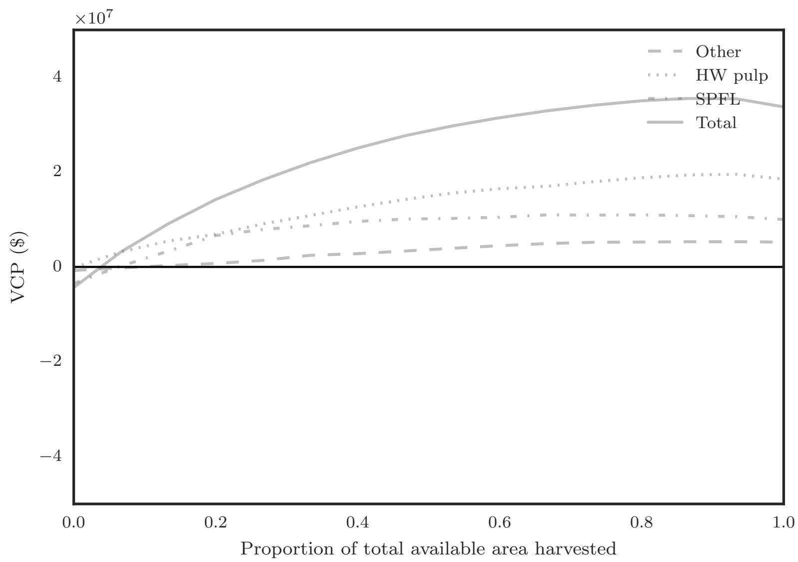

Figure 9: *Scenario 2.3.* VCP by aggregated commodity, as a function of proportion of AAC harvested. Dash-dot line represents SPFM sawlogs and pulpwood, dotted line represents hardwood pulpwood, and dash-dot line represents all other commodities.

4 Discussion

In the previous section, we presented results for four scenarios from an illustrative test case. For each scenario, we run our profit-maximizing network flow optimization model, using different commodity-wise subsidy parameters.

We ran a base scenario (scenario 1), using the financial performance indicators compiled from data we extracted from the MERIS database. Total VCP for scenario 1 is systematical negative, for all proportions of available area harvested. The true optimal solution in this case is to do nothing—harvesting even 1 unit of fibre requires automatic pre-payment of the 20% *rente*, which induces a net loss. Assuming that the rente has already been paid and cannot be refunded, the optimal solution for the network in this case is to harvest 18% of available area, which minimized total loss. The alternative scenarios (2.1, 2.2, 2.3) model different combinations of commodity-wise subsidies.

Scenario 2.1 models a subsidy on hardwood pulpwood, using subsidy amounts proposed for this management unit in a report prepared by a local forester in collaboration with a local consulting firm (see Belzile and Riopel 2015). These subsidies seem to be effective, to the extent that the hardwood pulpwood VCP curve is generally positive (although it has a negative slope at the end). However, the SPFL VCP curve is negative and monotonically decreasing. The combined effect yields a total VCP curve that peaks at 64% of total available area.

Scenario 2.2 models a subsidy on SPFL, using a subsidy amount proposed by the CIFQ (see Tremblay 2015). These subsidies seem to be effective, to the extent that the SPFL VCP curve is generally positive (although it has a negative slope at the end). However, the hardwood pulpwood VCP curve is negative and monotonically decreasing. The combined effect yields a total VCP curve that peaks at 52% of total available area.

Scenario 2.3 combines subsides from sceanrios 2.1 and 2.2. The net result is a total VCP curve that is generally positive, and montonically increasing until 90% of available area is harvested.

Note that total VCP curves are relatively flat near their peaks for all scenarios, so harvest volume could actually vary substantially left or right from peak profit volume with only small reductions in profit. This property can be interpreted in different ways. One might say that the optimal solution is relatively unstable, because harvested volume is likely to shift substantially in one or the other direction with relatively small changes in VCP input parameters. Another interpretation of these results is that government potentially has a powerful lever with which to influence industrial fibre consumption volume—indeed, as small be well-designed manipulation of the financial balance of different components of the wood supply could potentially shift industrial consumption to different subsets of the wood supply—this represents a rich potential area for further research and policy exploration.

Together, these scenarios show the total VCP curve can be substantially influenced by modelling commodity-wise subsidies, and that the shift in this curve can induce an increase in harvested area (and volume) by the profit-maximizing industrial network. Economic interpretation of these subsidies is outside of the scope of this study, as is analysis of political and financial feasibilty of issues—we are not making serious policy recommendations here, but simply illustrating the power of our modelling framework.

We have not validated VCP values compiled using our methodology against other sources or field data. Such validation should be performed, and the results of this exercise used to calibrate the model. Furthermore, we recommend that the calibrated prototype be deployed to the 71 other management units in Quebec, which would allow compilation of a province-wide VCP portrait. This new information could then be used to guide strategic forest-sector policy development at a provincial level.

Next steps in this research trajectory include testing the impact of exogenous market price fluctuations on model output, impact of manipulating industrial network capacity (i.e., adding or removing facilities), and possible game-theoretic applications to explore the concept of optimal benefit sharing.

5 Conclusion

In this paper, we present a hybrid simulation-optimization modelling framework, which can be used to estimate the proportion of available fibre supply that would be consumed by a network of profit-maximizing primary-breakdown forest sector mills. We designed our framework to be compatible with existing wood supply models used for long-term forest management planning in Quebec, and available data on mills and financial indicators. Thus, our modelling framework and methodology can be deployed as-is to the entire forest of Quebec, although we have not tested this.

Compiling the input data for this model is a complex and technically challenging task. We developed a companion model-input-compilation methodology (and software implementation) to assist in this, which we describe in previous publications.

We present an illustrative case study for a complex management unit in Quebec, which shows that only a relatively small subset of total fibre supply will be consumed in base case conditions. We also show how our framework can be used to simulate the hypothetical application of relatively small subsidies (i.e., manipulations of the financial flows of certain parts of the system) can induce substantial changes in industrial fibre consumption behaviour. We hope that this framework will be used by researchers, forest-sector investors, and policy-makers to explore new policy options and industrial capacity investment options to better align long-term fibre supply planning with current and anticipated industrial demand.

6 Acknowledgements

This study was supported by funding from the *FORAC Research Consortium*.

References

Belzile, J.-F. and Riopel, M. (2015). Mise en valeur d'un volume supplémentaire de bois aménagé et récolté dans les UAF 061-52 et 064-51. Technical report, Coopérative Forestière des Hautes-Laurentides.

Paradis, G., Bouchard, M., LeBel, L., and D'Amours, S. (2018). A bi-level model formulation for the distributed wood supply planning problem. *Canadian Journal of Forest Research*, 48(2):160–171.

Paradis, G. and LeBel, L. (2017a). Compiling disaggregation coefficients to link long- and short-term planning models. Technical Report CIRRELT-2017-43, CIRRELT.

Paradis, G. and LeBel, L. (2017b). Diameter Distribution Models for Quebec, Canada. Technical Report CIRRELT-2017-34, CIRRELT.

Paradis, G. and LeBel, L. (2018). Retro-fitting value-creation potential indicators to long-term wood supply models. Technical Report CIRRELT-2018-00, CIRRELT.

Paradis, G., LeBel, L., D'Amours, S., and Bouchard, M. (2013). On the risk of systematic drift under incoherent hierarchical forest management planning. *Canadian Journal of Forest Research*, 43(5):480–492.

Tremblay, A. (2015). Industrie forestière: Sommes-nous compétitifs? Technical report, Conseil de l'industrie forestière du Québec.