Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique et le transport

Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation

Extending the Classic Wood Supply Model to Anticipate Industrial Fibre Consumption

Gregory Paradis
Mathieu Bouchard
Luc LeBel
Sophie D'Amours

February 2015

CIRRELT-2015-06

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

Extending the Classic Wood Supply Model to Anticipate Industrial Fibre Consumption

Gregory Paradis[1,2,*], Mathieu Bouchard[1,2], Luc LeBel[1,2], Sophie D'Amours[1,3]

[1] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

[2] Department of Wood and Forest Sciences, 2405, rue de la Terrasse, Université Laval, Québec, Canada G1V 0A6

[3] Department of Mechanical Engineering, 1065, avenue de la Médecine, Université Laval, Québec (Québec), G1V 0A6

Abstract. The classic wood supply optimisation model maximises even-flow harvest levels, and implicitly assumes infinite fibre demand. In many jurisdictions, this modelling assumption is a poor fit for actual fibre consumption, which is often a species-unbalanced subset of total fibre allocation. Failure to anticipate this bias in volume and species mix of industrial wood fibre consumption has been linked to increased risk of wood supply failure. In particular, we examine the distributed wood supply planning problem, which is a variant of the general wood supply planning problem where the roles of forest owner and fibre consumer are played by independent agents (e.g. wood supply planning on public forest land in Canada, where government stewards control wood supply and forest products industry firms consume the fibre). We use agency theory to describe the source of antagonism between public forest land owners (the principal) and industrial fibre consumers (the agent). We show that the distributed wood supply planning problem can be modelled more accurately using a bilevel formulation, and present an extension of the classic wood supply optimisation model which explicitly anticipates industrial fibre consumption behaviour. The general case of the bilevel wood supply optimisation problem is *NP*-hard, non-linear, and non-convex - it is difficult to solve to global optimality. By imposing certain restrictions on agent network topology, we show that the general case can be decomposed into convex sub-problems. We present a solution methodology that can solve this special case to global optimality, and compare output and solution times of classic and bilevel model formulations using a computational experiment on a realistic dataset. Experimental results show that solution time for the bilevel problem is comparable to solution time for the classic single-level problem, and that the bilevel formulation can mitigate risk of wood supply failure.

Keywords: Forest management, wood supply, distributed planning, value creation networks, principal-agent problem, bilevel programming.

Acknowledgements. This study was supported by funding from the FORAC Research Consortium, the Fonds de recherche et de développement en foresterie (FRDF), the Fonds de recherche du Québec-Nature et technologies (FQRNT), the Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique, et le transport (CIRRELT), and the Natural Sciences and Engineering Research Council of Canada (NSERC) Strategic Research Network on Value Chain Optimization (VCO).

* Corresponding author: Gregory.Paradis@cirrelt.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2015

1 Introduction

In Canada, provincial government authorities allocate timber licences (TL) to industrial fibre consumers. These TLs, which grant rights to harvest fibre on public forest land, set species-wise upper bounds on periodic harvest volume, however there is typically no policy requirement to set matching lower bounds. Timber licences are typically valid for a pre-determined period (e.g. 5 years), after which point licences may be renewed, subject to re-evaluation of available wood supply. Maximal sustainable TL volume is commonly referred to as *annual allowable cut*, or AAC.

The term *wood supply planning* describes the process by which AAC is determined. In practice, this often amounts to solving a linear programming (LP) optimisation model, which finds the maximum even-flow harvest levels (Gunn, 2007). Periodic fluctuation of harvest levels in the LP solution is controlled using *even-flow constraints*, which specify an upper bound on the difference between highest and lowest periodic harvest volumes. Even-flow constraints are conceptually associated with the *sustainability* of the forest management process, although the scientific basis for this association is questionable.

The concept of sustainable forest management (SFM) has long been a central theme of public forest policy in Canada (Canadian Council of Forest Ministers, 2008). In concrete terms, sustainable forest management policy is implemented via silviculture treatments, notably harvesting treatments. According to the *Criteria and Indicators of Sustainable Forest Management in Canada* (CCFM C&I) (Canadian Council of Forest Ministers, 2003), harvesting is generally deemed to be sustainable if it is below AAC[1]. This notion, that any harvest level below AAC is sustainable, is also implicitly used at the forest management unit (FMU) scale when government uses species-wise even-flow AAC as contractual upper-bounds in TLs.

The classic model simulates a finite alternating sequence of harvesting and growth, and implicitly assumes that all available fibre will be consumed in every planning period (regardless of quantity, quality, cost, or value creation potential). In practice this assumption is rarely respected. Figure 1 shows the species-wise proportion of Canadian AAC consumed from 1990 to 2012. On average, 80% of softwood AAC and 45% of hardwood AAC were consumed, indicating a clear industrial preference for softwood during this period. The data show a species-skewed negative consumption bias, relative to AAC. This bias is related to the infinite-demand assumption implicitly embedded in the classic wood supply optimisation model. The bias can, to a certain extent, be attributed to a poor alignment between industrial fibre demand and wood supply planning. Local industrial processing capacity may be insufficient to consume some parts of AAC. Other parts of AAC may be economically unattractive (i.e. have a negative net value) or be operationally inaccessible (due to fragmented forest landscape, prohibitive access cost, or constrained by regulations that govern harvesting). Thus, the optimal solutions will likely never be executed, and the

[1] In particular, see criterion *5.3.1* of the CCFM C&I, which compares aggregated national AAC and harvest volumes as an indication of sustainability of forest management practices.

long-term state of the forest will be systematically different from that predicted by the wood supply model.

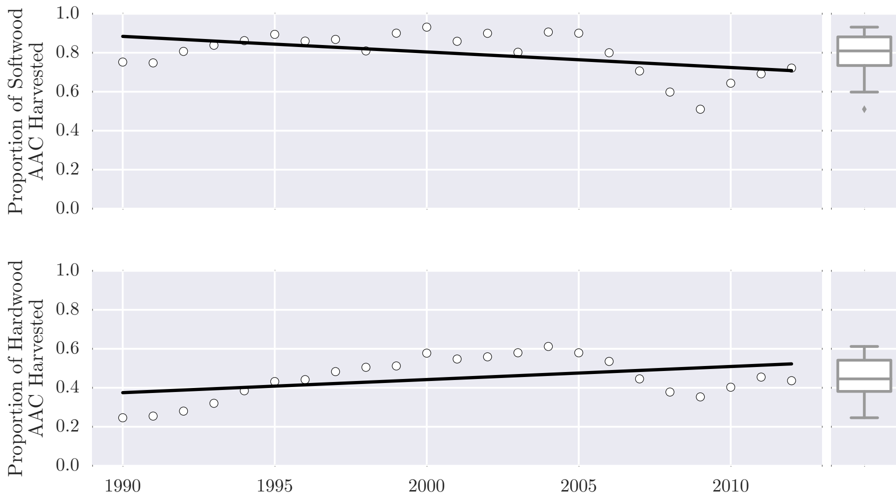

Figure 1: Proportion of species-wise AAC consumed in Canada for period 1990 to 2012 [source: National Forestry Database (2014)]

Mathey et al. (2009) use a spatially-explicit harvest scheduling model to estimate financial outputs for various timber supply levels, using a case study dataset from northern Ontario. Their simulation results confirm that a subset of AAC may be uneconomic.

Paradis et al. (2013) simulated repeated wood supply planning cycles, using a principal-agent paradigm to model the interaction between government wood supply planners and industrial fibre consumers. They show that the classic wood supply model may fail due to the aforementioned species-skewed negative fibre consumption bias, and conclude that the wood supply planning process currently in place on public land in Canada may not provide credible assurance of the long-term sustainability of the wood supply. Given the pervasiveness of this bias in practice, the classic wood supply optimisation model formulation does not constitute a rational basis for the implementation of sustainable forest management.

For management and planning purposes, the forest landscape is subdivided into stands. The stand is the basic silviculture decision unit, and can be described as a contiguous forested area with uniform vegetation and growth characteristics. Projection of future forest condition and fibre availability is based on aggregated result of stand-level growth and yield simulation. This requires four distinct types of information: *(1)* detailed starting inventory of forest condition (i.e. stand ages and types), *(2)* hypothetical projection of stand condition over time for each type[2], *(3)* hypothetical state transitions induced by stand events (i.e. shift in stand age and stand type), and *(4)* hypothetical intensity, location and timing of planned future stand events (e.g. clear-cut harvesting of stand i in period j).

Wood supply optimisation models typically use the first three information types (i.e. starting inventory, yield curves, and state transitions) as input, leaving the fourth information type (i.e. location and timing of future stand events) as variables in the objective function. Assembling this information into a coherent model, and subsequent analysis of model output, is referred to as the *wood supply planning problem*. We focus on a particular problem variant, which we call the *distributed wood supply planning problem* (DWSPP), where the roles of forest land owner and industrial fibre consumer are played by independent agents. The DWSPP is common in Canada and other jurisdictions, where public forest land is managed by government stewards on behalf of the general population.

We can describe the DWSPP, from a game-theoretic perspective, as an instance of the principal-agent problem (Laffont and Martimort, 2002; Schneeweiß, 2003). The role of principal is played by the forest owner (or government steward), and the role of agent is played by the industrial fibre consumer. The principal has the long-term responsibility to ensure a sustainable wood supply (hence the even-flow constraints in the wood supply model), but aims to maximise economic activity by exploiting the forest resource (hence the wood-supply-maximisation objective function). The agent aims to maximise short-term profit by transforming a subset of wood supply into forest products. Antagonism between the principal and the agent stems from either *(a)* binding agent capacity constraints[3] or *(b)* the presence of negatively-valued subsets of the wood supply [4]. Either of these factors may deter the agent from consuming the entire wood supply, which in turn induces the problematic negative fibre consumption bias.

There remains a gap in the literature with respect to the incidence of the principal-agent problem on the DWSPP, although a number of recent papers use a bilevel approach to model forest-sector decision problems. Bogle (2012) recently modelled optimal government policy response to a mountain pine beetle epidemic in British Columbia, Canada, using a principal-agent framework. Emphasis is placed on determining optimal government policy to incite fast liquidation of rapidly deteriorating beetle wood[5]. Their approximate solution methodology,

[2]Foresters typically refer to these as *yield curves*.

[3]We model agent behaviour using a network flow model, which we describe in more detail in §2. Each business unit in the agent network encapsulates one or more processes. Product flows between business units are defined by directional links. Both processes and links have capacity constraints, which can become saturated. When saturated capacity constraints in the agent network limits further improvement in the (profit-maximising) agent objective function, we can describe them as *binding agent capacity constraints*.

[4]If total cost of pushing a unit of fibre through the agent network exceeds potential revenu from sale of products to end-clients, then this unit of fibre has a negative net unit value, and will not be willingly consumed by a profit-maximising agent.

[5]The mountain pine beetle acts as a vector for a fungus, which spreads through the sapwood and typically kills the affected trees (Byrne et al., 2006). Although these trees can still be harvested and transformed into valuable forest products, the blue fungus that kills the trees stains the wood, making it less attractive than a similar volume of healthy clear wood.

which is tested only on a very small synthetic dataset, does not guarantee convergence on optimal solutions and is may be intractable for typical (i.e. large) datasets. Zhai et al. (2014) use a bilevel approach to model hierarchical planning in the case of fast-growing plantation management. Their explicit treatment of multiple lower level decision makers is interesting, however they do not explicitly address the case where only a subset of the harvest quota is economically attractive. Yue and You (2014) present a bilevel optimisation model to analyse the impact of adding new biorefinery capacity to an existing timber supply chain.

Paradis et al. (2013) also use a principal-agent framework to model failure of the classic wood supply model, and conjecture that extending the model to explicitly anticipate the principal-agent relationship should improve the coherence of the distributed wood supply planning process. Beaudoin et al. (2010) successfully used an agent-based modelling approach to anticipate interaction between participants in a distributed fibre procurement process at the tactical-operational planning level. They leverage this anticipation to better coordinate the planning process between horizontally-linked agents (i.e. firms sourcing fibre from the same procurement area in the public forest), resulting in local and global profitability gains relative to the *ad hoc* procurement planning process currently in place in many jurisdictions. An agent-based anticipation approach could help coordinate the vertical integration between government (principal) and industry (agent) planners, at a strategic-tactical planning level.

We endevour to close this gap by extending the classic wood supply optimisation model to explicitly anticipate industrial fibre consumption behaviour. Our anticipation function is based on a network flow model called *LogiLab*, developed by the FORAC Research Consortium (Jerbi et al., 2012). Thus, we model agent behaviour as material and financial flows through a forest product value-creation network. The agent seeks to maximise profit, subject to wood supply, capacity, and demand constraints. The model assumes centralised network planning, and exogenous end-product prices.

From an operations research perspective, bilevel programming subsumes the principal-agent problem (Colson et al., 2007). Hence, any principal-agent problem can be formulated as a bilevel optimisation problem. Although a bilevel modelling approach could potentially address the principal-agent aspect of the DWSPP, solving a bilevel problem to global optimality is typically not trivial. Even the simplest bilevel optimisation problems are known to be $\mathcal{NP}$-hard, and non-convex, non-linear solution spaces are common (Dempe, 2003; Colson et al., 2007). This paper addresses the shortcomings of the classic formulation identified in Paradis et al. (2013), namely the species-skewed fibre consumption bias that is implicitly embedded in the classic wood supply model. We developed both a bilevel model formulation for the DWSPP, and a novel methodology to solve the bilevel problem to global optimality.

We hypothesised that the bilevel model formulation would improve stability

Furthermore, the physical properties of the dead wood can negatively affect the quality of lumber products manufactured from this material, with degradation worsening overtime such that beetle wood is typically considered to have a 5 to 10 year shelf life (Trent et al., 2006).

of the wood supply. To verify this hypothesis, we designed a computational experiment comparing output from classic and bilevel model formulations after 30 sequential rolling-horizon replanning cycles.

The remainder of this paper is organised as follows. The mathematical formulation of the bilevel problem, solution methodology, and experimental methods are presented in §2. Results from the computational experiment are presented in §3, followed by discussion in §4. Concluding remarks are presented in §5.

2 Methods

We present a mathematical formulation of the bilevel wood supply optimisation model in §2.1, followed by a solution methodology we developed to solve the bilevel problem in §2.2. We also present the methodology for a computational experiment, in which we compare the performance of classic and bilevel model formulations, in §2.3.

2.1 Bilevel Wood Supply Problem Formulation

The bilevel wood supply problem extends the DWSPP to anticipate industrial fibre consumption. We present mathematical formulations for both upper level (principal) and lower level (agent) problems before defining the bilevel optimisation problem solution space.

2.1.1 Upper-Level Problem Formulation

The following formulation of the upper level model is a simplified representation of the *de-facto* standard (i.e. *classic*) wood supply planning model implemented in many jurisdictions, including public forest land in Canada[6]. This is equivalent to a *Model I* formulation, using the modelling terminology in Davis et al. (2001).

Maximise

$$\sum_{i \in Z} \sum_{k \in P_i} c_{ik} x_{ik} \tag{1}$$

subject to

$$\sum_{k \in P_i} x_{ik} = 1, \qquad \forall i \in Z \tag{2}$$

$$y_o \leq \sum_{i \in Z} \sum_{k \in P_i} \alpha_{ikot} x_{ik} \leq (1 - \varepsilon_o) y_o, \qquad \forall o \in O', t \in T \tag{3}$$

$$l_{ot} \leq \sum_{i \in Z} \sum_{k \in P_i} \alpha_{ikot} x_{ik} \leq u_{ot}, \qquad \forall o \in O, t \in T \tag{4}$$

[6] Adapted from Paradis et al. (2013).

where

$$
\begin{aligned}
Z &:= \text{set of spatial zones} \\
P_i &:= \text{set of available prescriptions for zone } i \in Z \\
O &:= \text{set of forest outputs} \\
O' \subset O &:= \text{set of sustainable forest outputs} \\
T &:= \text{set of time periods in the planning horizon} \\
\varepsilon_o &:= \text{admissible level of variation on yield of output } o \in O \\
\alpha_{ikot} &:= \text{quantity of output } o \in O \text{ produced in period } t \in T \text{ by prescription} \\
&\quad\; k \in P_i \text{ in zone } i \in Z \\
l_{ot} &:= \text{lower bound on yield of output } o \in O \text{ in period } t \in T \\
u_{ot} &:= \text{upper bound on yield of output } o \in O \text{ in period } t \in T \\
c_{ik} &:= \text{global value of including cost and benefits of prescription } k \in P_i \text{ in} \\
&\quad\; \text{zone } i \in Z \\
x_{ik} &:= \text{fraction of zone } i \in Z \text{ on which prescription } k \in P_i \text{ is applied} \\
y_o &:= \sum_{i \in Z} \sum_{k \in P_i} \alpha_{iko1} x_{ik}, \text{ which corresponds to first-period harvest volume} \\
&\quad\; \text{for output } o \in O'
\end{aligned}
$$

The objective function (1) maximises harvest volume over the planning horizon. Constraint (2) is an accounting constraint, and simply ensures that the entire forest area is assigned to a prescription (including the null prescription, i.e. do nothing). Constraint (3) models an *even-flow* policy, which constrains simulated harvest level to be more or less level throughout the planning horizon (i.e. stabilise periodic flow of timber from the forest). We have expressed the even-flow constraint in terms of y_o, which represents first-period harvest volume for output $o \in O'$ (i.e. species-wise periodic allowable cut[7]). Constraint (4) bounds minimal and maximal periodic yield on the general outputs (e.g. area converted to plantation in a given period).

Thus, y_o corresponds to the wood supply, for a given species group, that the principal might offer the agent in a given planning cycle—y_o constitues the primary linkage interface between upper and lower level models, as we will see in the next section. This is analogous to the actual wood supply planning process, where AAC constitutes the primary policy interface between the principal and the agent.

2.1.2 Lower-Level Problem Formulation

The following formulation of the lower-level (agent) problem is adapted from Jerbi et al. (2012).

[7] Annual allowable cut, or AAC, is essentially the same as periodic allowable cut, but expressed on an annal basis.

Maximise

$$\sum_{u\in U}\left(\sum_{p\in P|d_{up}>0}\rho_{up}D_{up}-\sum_{w\in W_u}c_wY_{uw}\right)-\sum_{e\in E}\sum_{p\in P}c^f_{ep}F_{ep} \tag{5}$$

subject to

$$\beta_{uo}+\sum_{w\in W_u}(\gamma_{pw}-\alpha_{pw})Y_{uw}+\sum_{e\in\delta^+_u}F_{ep}-\sum_{e\in\delta^-_u}F_{ep}-D_{up}=0, \qquad u\in U, p\in P \tag{6}$$

$$\sum_{w\in W_u}\gamma_{kuw}Y_{uw}\leq q_{ku}, \qquad u\in U, k\in K \tag{7}$$

$$D_{up}\leq d_{up}, \quad \forall u\in U, p\in P \tag{8}$$

$$\sum_{p\in P}F_{ep}\leq f^u_e, \qquad \forall e\in E \tag{9}$$

$$f^l_{ep}\leq F_{ep}\leq f^u_{ep}, \quad \forall e\in E, p\in P \tag{10}$$

$$\sum_{u\in U}\beta_{uo}\leq y_o, \qquad \forall o\in O' \tag{11}$$

where

$U :=$ set of business units
$K :=$ set of resource capacity types (machine capacities, stock limits)
$W :=$ set of processes (machines, inventories)
$W_u \subset W :=$ set of processes available at business unit $u\in U$
$P :=$ set of products
$O' \subset P :=$ set of sustainable forest outputs, from the upper-level model
$E :=$ set of links between business units
$\delta^+_u \subset E :=$ set of inbound links for business unit $u\in U$
$\delta^-_u \subset E :=$ set of outbound links for business unit $u\in U$
$q_{ku} :=$ capacity of type $k\in K$ at business unit $u\in U$
$f^l_{ep} :=$ lower bound on flow of product $p\in P$ through link $e\in E$
$f^u_{ep} :=$ upper bound on flow of product $p\in P$ through link $e\in E$
$f^l_e :=$ upper bound on flow of all products through link $e\in E$
$c_w :=$ unit cost of process $w\in W$
$c^f_{ep} :=$ unit cost of transporting product $p\in P$ on link $e\in E$
$\alpha_{pw} :=$ quantity of product $p\in P$ required for one unit of process $w\in W$
$\gamma_{pw} :=$ quantity of product $p\in P$ produced for one unit of process $w\in W$
$\lambda_{kuw} :=$ capacity of type $k\in K$ utilised by process $w\in W$ at business unit $u\in U$
$d_{up} :=$ demand for product $p\in P$ at business unit $u\in U$
$\rho_{up} :=$ price of product $p\in P$ at business unit $u\in U$
$\beta_{uo} :=$ external supply of sustainable forest output $o\in O'$ at business unit $u\in U$
$y_o :=$ maximum external supply of sustainable forest output $o\in O'$ (from upper-level model)

and

$$Y_{uw} := \text{quantity of process } w \in W \text{ performed at business unit } u \in U$$
$$D_{up} := \text{quantity of product } p \in P \text{ sold by business unit } u \in U$$
$$F_{ep} := \text{flow of product } p \in P \text{ on link } e \in E$$

The objective function (5) maximises network profit (i.e. revenue from sale of products, net of production and transportation cost). Constraint (6) ensures flow conservation in the network. Constraint (7) limits process utilisation to production capacity upper bounds. Constraint (8) limits product sales to demand upper bounds. Constraint (9) limits total flow to link capacity upper bounds. Constraint (10) ensures that product-wise flow upper and lower bounds are respected for each link. Constraint (11) limits consumption of fibre by agent network to the species-wise wood supply determined in the upper-level model.

Note that storage of unsold products and end-product consumption by external clients are modelled as special cases of processes. Fibre procurement from the forest is also modelled as a process, although we present it here using a separate parametre β_{uo}, as this facilitates description of the linkage with the upper-level model, via parametre y_o.

We have chosen to model industrial fibre consumption in the lower level for the first planning period only. We could have chosen to extend anticipation of agent behaviour to an arbitrary number of periods without loss of generality.

2.1.3 Bilevel Solution Space

The bilevel feasible region is the subset of the upper-level feasible region for which the lower-level model consumes the entire wood supply. In other words, a bilevel-feasible wood supply solution must be upper level feasible *and* be entirely consumed by the lower-level model (i.e. $\sum_{u \in U} \beta_{uo} = y_o, \forall o \in O'$).

The classic definition of AAC fails to account for the species-skewed negative fibre consumption bias (i.e. fails to account for realistic industrial fibre consumption behaviour). We propose an extended definition of AAC, which is *the maximum periodic harvest level that respects species-wise even-flow constraints, and will be entirely consumed by profit-maximising industrial consumers*. Given this extended definition of AAC, even-flow wood supply solutions that will not be entirely consumed by the agent are no longer feasible. If deterministic assumptions hold true[8], this extended definition will completely eliminate the problematic species-skewed negative fibre consumption bias.

Although defining bilevel AAC is relatively straightforward, the resulting bilevel solution space is non-convex, which makes it difficult to solve the bilevel wood supply problem to optimality.

[8]The assumption of determinism is implicit to linear programming, which we use to model both upper and lower levels of the problem.

2.1.4 Bilevel Solution Space Convexity Issues

The objective of this research is to formulate and solve a bilevel optimisation model for the DWSPP. Early model development efforts focused on what we will hereafter refer to as the *general bilevel case*, which we hoped to solve to global optimality using an iterative solution methodology, based on the pioneering work of Fortuny-Amat and McCarl (1981) and Bard and Falk (1982). After developing the mathematical formulation for the general bilevel case, we recognised the potential for non-convexity of the bilevel solution space. This limited the worst-case performance of our initial solution methodology approach to local optimal solutions. Given the strategic importance of wood supply planning, we considered local optimal solutions to be of limited interest, particularly in the absence of reliable bounds on the optimality gap. Thus, we decided to focus subsequent efforts on identifying special cases of the general bilevel problem that we could solve to global optimality.

Non-convex cases can occur when the agent value-creation network features one or more convergent processes with saturated capacity constraints[9]. For example, this could occur if all chips in a network must flow through a single pulpmill (i.e. the convergent process), with a single digester that runs non-stop (i.e. shared resource with saturated capacity constraints). Appendix B presents a counter-example proving non-convexity of the general bilevel solution space for the DWSPP.

By restricting the problem domain to non-problematic (i.e. convex) instances, we can define a special case of the bilevel problem that can be solved more easily. The special case, which is defined by excluding problematic instances from the general problem domain, occurs in two types of instances. Special case 1 corresponds to strictly divergent networks (i.e. no convergence of product flows at any process), and can be easily identified by analysing network topology of the agent model. Special case 2 corresponds to partially convergent networks with sufficient capacity at convergent processes so these will not be saturated at the bilevel optimal solution. Special case 2 may only be observable through empirical analysis of the agent model, and may not be detectable *a priori*, as checking for potential capacity constraint saturation requires us to first solve the lower level model.

Both types of species cases are separable in O' (i.e. the sustainable forest outputs used to segment AAC in the upper level model). Thus, we can develop a solution methodology for the special case that decomposes the lower level problem into convex species-wise subproblems, which can be solved to determine optimal species-wise upper bounds on harvest levels in the upper level problem. These species-wise upper bounds essentially represent valid cuts for the upper level solution space. These cuts form the basis for our solution methodology to solve the special cases. We cannot generate similar valid cuts for the general case, using this methodology, due to non-convexity of the general bilevel solution space.

[9] A convergent process accepts input from two or more sustainable forest outputs $o \in O'$ (e.g. a mix of softwood and hardwood).

The species-wise upper-bound volume for output $o \in O'$ (i.e. solution to lower level subproblem for output o) is equal to $\arg\max_{v_o} p(v_o)$, where $p(v_o)$ is a concave piece-wise linear function describing profit. The piecewise profit function is concave for each separable output o, for any non-negative consumption volume v_o (see Figure 2). At the leftmost end of each curve (i.e. the lower extremity of the piecewise linear function domain), the agent utilises his limited resources to produce his most profitable product mix until one of his resource constraints is saturated. The agent then switches production to the next most profitable product mix. The slope of each line segment in the piecewise function represents unit profit from a given product mix. By definition, the slopes of these line segments will be monotonically decreasing, hence the concave shape of the piece-wise linear profit function.

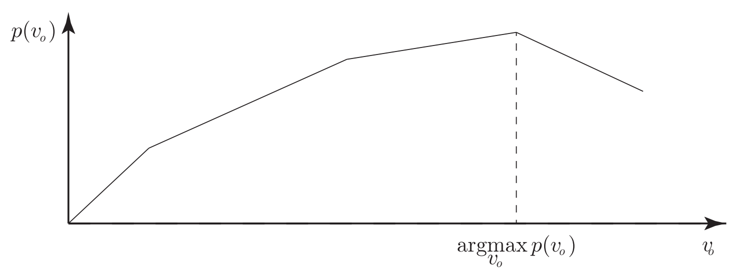

Figure 2: Illustration of concave profit function of a hypothetical product o

2.2 Solution Methodology

We developed a solution methodology to solve the bilevel problem presented in §2.1. Our methodology is implemented on top of an iterative simulation framework, originally described in Paradis et al. (2013). This framework, which was developed to simulate principal-agent interaction after several sequential two-stage rolling-horizon distributed wood supply planning cycles, links SilviLab and LogiLab software platforms[10] to model principal and agent behaviour. See Appendix C for more detailed framework implementation notes.

We originally set out to design an iterative solution methodology for the general case, however convexity issues[11] led us to focus algorithmic development efforts on the (convex) special cases, which we can solve to global optimality. By limiting the problem domain to the special cases, we eliminate the possibility of interaction between outputs $o \in O'$ (i.e. the bilevel problem becomes separable in O'). This property allows us to compute output-wise upper bounds on agent consumption behaviour, which can be used to optimally constrain the upper-level

[10]SilviLab and LogiLab software platforms are developed by the FORAC Research Consortium to model wood supply problems and forest products industry value creation network problems, respectively.

[11]See §2.1.4 and Appendix B.

problem to produce the global optimal bilevel wood supply solution. This is the basis of our solution methodology for the special case. Algorithm 1 describes this solution methodology using pseudo-code.

Algorithm 1: Bilevel model solution algorithm (special case)

Output : Global optimal wood supply solution x^*

1 **foreach** *output* $o \in O'$ **do**

2   Determine upper bound μ_o on agent consumption of output o (i.e. solve lower-level sub-problem).

3   Set upper bound constraint $y_o \leq \mu_o$ on quantity of output o produced by the upper level problem.

4 **end**

5 Solve constrained upper-level problem.

Algorithm 1 essentially augments the upper-level model formulation presented in §2.1.1 with output-wise upper bounds (i.e. $y_o \leq \mu_o, \forall o \in O'$) on harvesting. Values for μ_o are derived from the optimal solutions to *lower-level subproblems* (referenced in line 2 of Algorithm 1).

Each subproblem is similar to the lower-level model formulation presented in §2.1.2, but with constraint (11) relaxed[12] and non-targeted outputs[13] disabled. Disabling non-targeted outputs ensures that only the targeted output o affects the subproblem solution (see Appendix C for more detailed notes on lower-level model implementation, and how we disable the non-targeted outputs). The optimal solution to each lower-level subproblem yields the maximum volume μ_o of output o that would be voluntarily consumed by the agent if an infinite supply of output o was available. Conceptually, μ_o corresponds to $\arg\max_{v_o} p(v_o)$ (see Figure 2).

For the special case, each μ_o constitutes a valid cut for the bilevel solution space. Applying similar cuts to the general case would possibly exclude the global optimal solution, due to non-convexity of the solution space. Algorithm 2 describes a methodology to solve the general case to global optimality by simple enumeration (i.e. iterate over the set of all possible convex subproblems, and return the best solution). We use the term *super-saturated process* to describe a process $p \in P_S$ (where $P_S \subset P$), whose capacity constraints are exceeded when subproblem solutions are summed in line 5 of Algorithm 2. It is possible that some process capacity constraints are violated when subproblem solutions are summed, due to non-convexity of general case solution space (i.e. the general case is not strictly separable in o, hence the need for enumeration to find feasible solutions). Subproblem definition is the same as for line 5 of Algorithm 1.

Algorithm 2 is of limited practical interest, as computational effort required to solve realistically-sized instances by enumeration would be prohibitive. However, efficiency of this algorithm could potentially be improved by replacing subproblem

[12] Equivalent to infinite external supply of output o.

[13] Non-targeted outputs correspond to $O' \setminus o$, that is the set O' excluding targeted output o.

enumeration with a custom branching algorithm, thereby taking advantage of problem structure. Although we have not tested this approach, further development of computationally tractable methodologies to solve the general case represents an interesting direction for further research.

Algorithm 2: Bilevel model solution algorithm (general case)

Output : Global optimal wood supply solution x^*

foreach *output* $o \in O'$ **do**
    Generate subproblem (isolate output o in lower-level problem).
    Solve subproblem.
end
Sum subproblem solutions.
Build super-saturated process set P_s (i.e. find violated capacity constraints).
foreach *combination of super-saturated process* $p \in P_s$ *and output* $o \in O'$ **do**
    Generate subproblem (restrict certain combinations of output o and process p).
    Solve subproblem.
    if *subproblem feasible (i.e. no violated capacity constraints)* **then**
        Add solution x to feasible solution set X.
    end
end
return *Global optimal solution* x^* *(i.e. best solution in feasible set* X*).*

In practice, most lower-level datasets will correspond to one or the other of the special cases, which can be solved with relative ease using Algorithm 1. The test dataset used in our computational experiment, which we describe in the next section, is an example of special case 2.

2.3 Computational Experiment Methodology

This section describes the computational experiment we conducted to compare the performance of classic and bilevel wood supply models.

2.3.1 Test Dataset

We tested our bilevel solution methodology on a realistic synthetic dataset from Quebec, Canada. The study area is a forest management unit (FMU 031–53) located in the boreal forest region. It covers an area of approximately 102 thousand hectares. Approximately 88% of initial growing stock is from softwood species, with the remaining 12% of initial growing stock in hardwood species. Although some pure softwood stands are present, forest cover is primarily composed of softwood-rich mixed-wood stands.

Output from the upper-level (forest) model is aggregated into two outputs: *softwood* and *hardwood*. The lower-level (industrial) model has limited capacity for transforming hardwood (approximately 1/3 of potential sustainable wood supply). The classic wood supply model therefore systematically over-estimates short-term hardwood fibre consumption.

We use the same test dataset as in Paradis et al. (2013), which is an instance of special case 2. Although chip flows from both hardwood and softwood sawmills converge at the pulp mill, we have determined empirically that its capacity is sufficient to avoid saturation problems. We can therefore use Algorithm 1 to solve the bilevel problem to optimality for our test instance.

2.3.2 Iterative Simulation Framework

We use the same two-stage rolling-horizon simulation framework described in Paradis et al. (2013) as a testbed in which to compare the performance of classic and bilevel wood supply model formulations. At each simulated planning cycle, the principal and the agent make their moves sequentially, in a two-stage game. For our computational experiments, we chose to simulate 30 (5-year) replanning cycles, as this corresponds to the length of our wood supply model planning horizon. The framework simulates forest growth between each 5-year rolling-horizon replanning cycle. The principal has the advantage of the first move, which means he can set AAC to any level of his choosing.

The simulation algorithm can be summarised as follows:

1. *First stage: the principal determines his wood supply offer.* We simulate the wood supply planning process by solving a wood supply optimisation model. Which model we solve at this stage—either the classic (single-level) model or the extended (bilevel) model—depends on the scenario. The wood supply offer is communicated to the agent in terms of species-wise upper bounds on volume that can be harvested in the second stage.

2. *Second stage: the agent consumes a subset of the wood supply..* We simulate fibre consumption by solving a network flow model, to determine the profit-maximising subset of wood supply that the agent would willingly consume. Species-wise upper bounds from the first stage (i.e. AAC) are applied at this stage. We also implement *line-wise profitability constraints* (see §C), which simulate existence of multiple profit centres in the agent network.

3. Simulate rolling horizon forward one period (simulate evolution of forest state using growth and yield curves from long-term wood supply model).

2.3.3 Experimental Methodology

We present five scenarios. Within each scenario, simulation parametres[14] for the industrial fibre consumption network are held constant for all 30 planning cycles. Table 1 summarises simulation parametres for each scenario.

[14] Mill capacities, costs, prices, client demand, etc.

Scenario 1 simulates *status quo* behaviour for both principal and agent, and acts as a control scenario. At each planning cycle, the principal maximises even-flow AAC (30-period horizon) using the classic wood supply model, then the agent maximises first-period profits (1-period horizon) by consuming an optimal subset of the wood supply offered by the principal. The principal does not consider the agent's fibre consumption capacity when determining AAC.

Scenario 2 presents a perfect-implementation bilevel scenario—rather than being allowed to replan harvesting on a one-period horizon (as is the case for the other scenarios), the agent is forced to exactly implement the first period of the principal's wood supply solution. This scenario shows the best-case performance of the bilevel model solution.

Scenario 3 is the basic bilevel scenario. The principal uses the bilevel model to determine AAC, and the agent is allowed to replan harvesting on a one-period horizon, choosing the profit-maximising subset of available wood supply. Due to the optimal formulation of the bilevel model and perfect anticipation of volume consumption, the agent always chooses to harvest the entire wood supply. However, the agent may select to harvest this volume from a different combination of forest types than that which was prescribed in the first period of the principal's optimal solution. This reflects the distributed nature of forest management planning on public forest land in many jurisdictions.

Scenarios 4 and 5 are variants of scenario 3, simulating reduction of softwood supply allocated to the agent to 80% and 60% of bilevel AAC. Adjusting AAC allocation indirectly creates a buffer stock to protect against the effects of agent harvest replanning (i.e. compensation for the principal's incomplete control of the fibre procurement process).

Table 1: Summary of scenario parametres

Scenario	Principal Model	Agent Model
1	Classic	Basic
2	Bilevel	Slave
3	Bilevel	Basic
4	Bilevel (80% attribution†)	Basic
5	Bilevel (60% attribution†)	Basic

†Of bilevel softwood AAC.

3 Results

We present experimental results in two stages. First, we present detailed results comparing output from the first planning cycle of scenarios 1 and 3. Next, we show results of simulating 30 sequential rolling-horizon planning cycles for scenarios 1 through 5.

For scenario 1, the control scenario, potential hardwood fibre supply is

$64\,583\,\mathrm{m}^3$, whereas actual consumption is only $20\,800\,\mathrm{m}^3$. The difference between planned and executed hardwood fibre consumption volumes is due to the limited processing capacity at the (single) hardwood sawmill in the agent processor network. The entire softwood fibre supply is consumed by the agent, as we would expect, as both end-product demand and processing capacity for the softwood line are high enough to accommodate more softwood fibre than the forest can supply. This phenomenon (of consuming certain components of the wood supply entirely while other components are only partially consumed) can be observed to varying extents in practice, as processing capacity and market demand are often misaligned with the proposed wood supply. To consume the full softwood supply, while only harvesting a third of the hardwood supply, the agent must favour harvesting stands that have a higher proportion of softwood and lower proportion of hardwood.

For scenario 3 (i.e. the basic bilevel scenario), results show that our bilevel anticipation mechanism completely eliminates the over-estimation of hardwood fibre consumption volume. Fibre consumption by the agent is exactly equal to wood supply volumes offered by the principal ($20\,800\,\mathrm{m}^3$ for hardwood, and $323\,759\,\mathrm{m}^3$ for softwood). This is the desired outcome from the bilevel model, and represents a global optimal solution for this instance. Note that the agent plans his own harvesting in the second stage of the simulation (using his single-period profit-maximising model), so harvest areas will not typically match the first period of the principal's plan, although harvest volume is exactly equal in scenario 3.

Table 2 presents intermediate results from each step of the bilevel solution method[15], for scenario 3. We present this data as an example of how we derive the optimal upper bounds to the wood supply problem from solutions to output-wise sub-problems. Hardwood consumption capacity ($20\,800\,\mathrm{m}^3$) is the binding constraint in this case. Our anticipation mechanism shows that the softwood line would have willingly consumed up to $584\,861\,\mathrm{m}^3$ of softwood in the first planning period, however the species-wise even-flow constraints on the upper-level wood supply model limit long-term softwood harvest level to $323\,759\,\mathrm{m}^3$. As expected, fibre volume consumed by the agent in the second phase of the simulation is exactly equal to the wood supply. The bilevel model eliminates the gap between planned and executed fibre consumption levels, thereby fulfilling its intended purpose. Figure 3 presents detailed results from the first planning cycle of scenarios 1 (classic model) and 3 (basic bilevel model).

We solve the bilevel model in less than twice the time required to solve the classic model. The classic model can be solved in a single step, which corresponds to approximately 13 seconds of CPU time for our test setup. The bilevel model requires $|O'| + 1$ steps to solve, which corresponds to approximately $(4 + 6) + 10$ seconds of CPU time using our test setup. We ran our tests on an Intel® Xeon® E5–2670 processor (20 MB cache, 2.60 GHz).

Figures 5 and 6 show sequential replanning simulation results for five scenarios described in §2.3.3. For each scenario, Figure 5 plots species-wise AAC and fibre

[15]See Algorithm 1.

consumption for each of the 30 rolling-horizon replanning cycles. The same data are shown in Figure 6 using box-plots to illustrate the variability of periodic AAC and fibre consumption data across scenarios. The boxes encompass the inter-quartile range (IQR) with the median marked. The whiskers extend to 1.5 IQR past the nearest quartile. Observations outside this range are marked as outliers using a dot symbol.

Figure 4 shows planned and executed softwood harvest volumes for bilevel scenarios 3, 4, and 5. We include this figure to show that witholding a portion of bilevel AAC tends increase mean AAC, decrease mean harvested volume, and improve wood supply stability throughout the horizon (as show by increased tightness of the boxplots).

Table 2: Bilevel solution method (intermediate results)

Stage	Description	Volume (m^3) Hardwood	Softwood
1 (principal)	Upper bound on *hardwood* consumption	20 800	–
1 (principal)	Upper bound on *softwood* consumption	–	584 861
1 (principal)	Maximum even-flow wood supply levels	20 800	323 759
2 (agent)	Agent fibre consumption	20 800	323 759

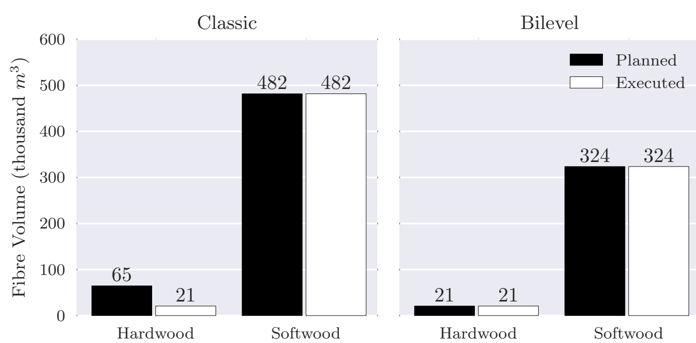

Figure 3: Comparison of planned and executed volumes

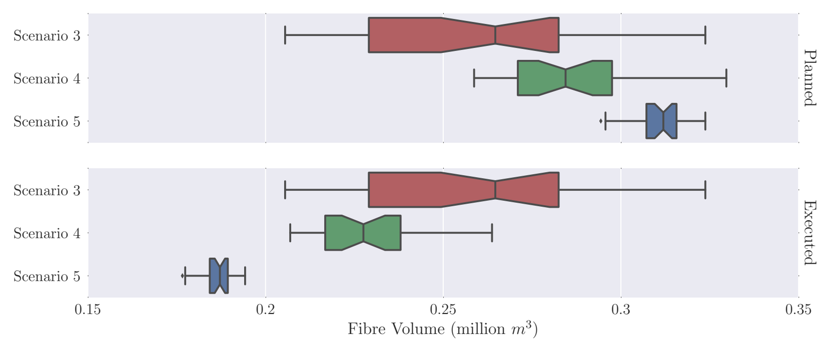

Figure 4: Planned and executed softwood harvest volume for bilevel scenarios 3, 4 and 5 (correspond to 100%, 80%, and 60% softwood AAC attribution policies)

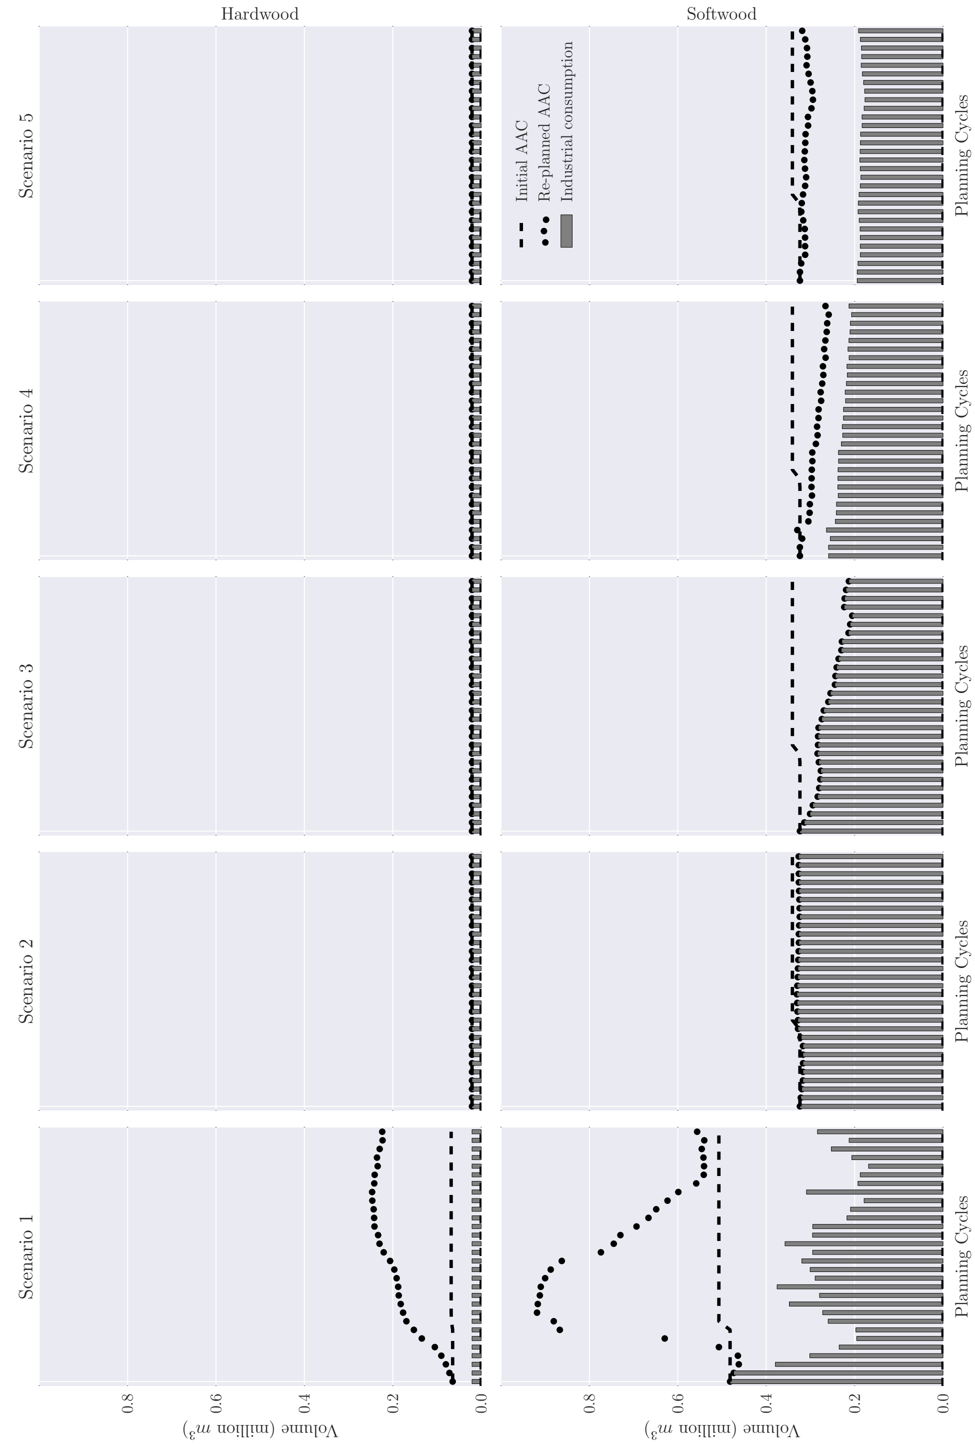

Figure 5: Species-wise AAC and fibre consumption for scenarios 1 to 5 (time series)

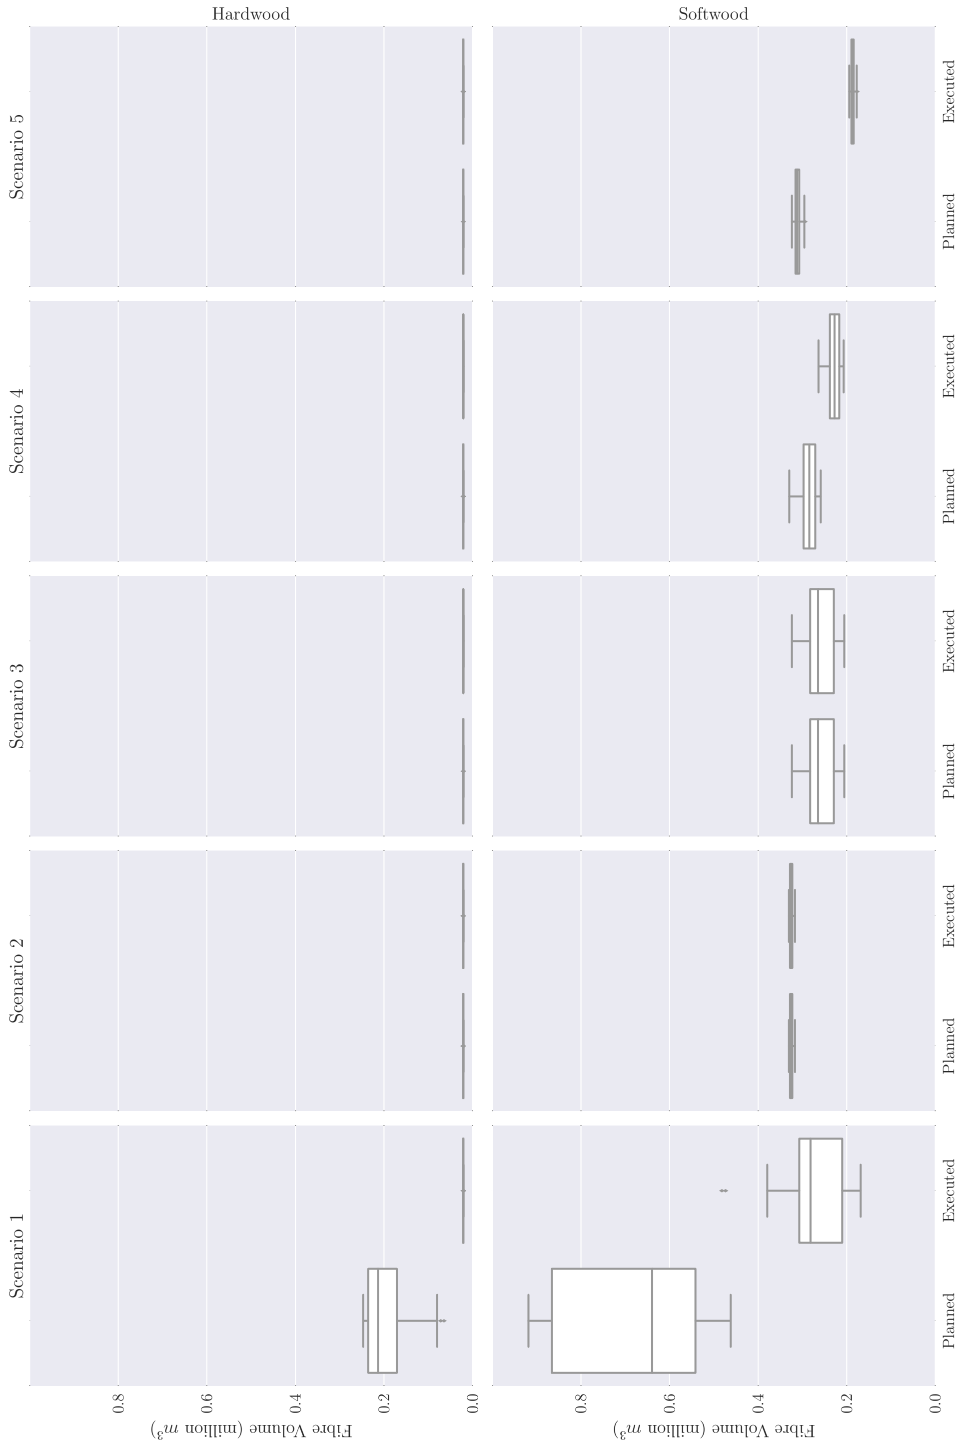

Figure 6: Species-wise AAC and fibre consumption for scenarios 1 to 5 (boxplots)

4 Discussion

Scenario 1 shows the relative instability of the classic wood supply model. This is attributable to the species-skewed gap between AAC and fibre consumption volumes. Despite the harvest levels being systematically lower than AAC, the agent's preference for harvesting high-softwood-content stands gradually shifts the composition of the residual forest cover towards a higher hardwood content. This undesirable shift in forest composition is not predicted by the classic wood supply model. Paradis et al. (2013) use the term *systematic drift effect* to describe this phenomenon.

The principal uses the bilevel model to determine AAC for scenarios 2 through 5. By virtue of its formulation, the bilevel model completely eliminates the volume gap between AAC and fibre consumption. Note that we simulate perfect anticipation of agent fibre consumption *volume*. For scenarios 3 through 5, we allow the agent to plan his own harvesting in the second phase of each planning cycle simulation—this explains the residual instability in long-term wood supply.

Scenario 2 forces the agent to harvest the exact forest units that form the basis of the first period of the principal's optimal bilevel solution. The purpose of this scenario is to show that wood supply tracks almost perfectly along the initial bilevel AAC solution (even after 30 rolling-horizon replanning cycles) under best-case conditions (i.e. when the principal controls wood procurement planning and execution all the way to the mill gate). In practice, the decoupling point between the principal and the agent is not typically located this far downstream. Scenarios 3 through 6 simulate a more conventional decoupling point.

Scenario 3 shows vastly improved stability softwood supply levels, relative to the control scenario. Note the gradual downward trend of the softwood fibre supply for scenario 3, which contrasts with the even supply profile simulated in scenario 2. The contrast between scenarios 2 and 3 shows that the bilevel model is sensitive to deviations from the optimal wood supply model solution. Sensitivity to deviations from the optimal solution is typical of deterministic optimisation models, as optimal solutions are invariably located along the boundary of the feasible region—even the slightest deviations from the optimal solution (or error in constraint right-hand-side values) can induce problem infeasibility.

Scenarios 4 and 5 show the effect of reducing the proportion of AAC that is allocated to the agent, in an attempt to compensate for the residual drift seen in scenario 3. Reducing allocation is an indirect way for the principal to induce a buffer stock in the standing timber inventory. This tends to move the executed (second-stage) solution away from the feasible boundary of the planned (first-stage) solution space, thereby improving the robustness of the distributed wood supply planning process. Scenario 4 shows a marked reduction in drift compared with scenario 3. Residual drift is virtually eliminated in scenario 5. Intuitively, withholding 40% of AAC seems like a high penalty to eliminate the residual drift in the bilevel model. We conjecture that, using a more direct management approach to maintaining a buffer stock in standing timber inventory, as described in Raulier et al. (2014), it may be possible to achieve higher stable bilevel AAC levels. This represents a promising direction for further research.

By stabilising the long-term wood supply, scenario 5 succeeds in restoring credibility to the wood supply planning process, albeit at a relatively high cost in terms of withheld AAC. Furthermore, scenario 5 makes less optimistic assumptions regarding agent behaviour than scenarios 2 (which simulates a perfectly compliant agent). Scenario 5 respects the even-flow pattern prescribed by the wood supply model constraints. Assuming that the even-flow constraints are valid and necessary (although not sufficient) conditions for sustainability of the forest management plan[16], and that the principal's responsibility to ensure sustainability must absolutely supersede any desire to maximise short-term TL volume allocations, scenario 5 represents the only example of a principal-feasible policy in this study.

We simulated the distributed wood supply planning process as a two-stage sequential game, where the principal proposes his wood supply in the first phase and the agent consumes a profit-maximising subset of the wood supply in the second phase. At this point, we can conjecture that stable increases in AAC may be achievable if the principal and the agent were allowed to iteratively adjust their respective supply and demand offers within a given planning cycle. This represents a promising direction for further wood supply policy research. From a game-theoretic perspective, extending the two-stage game simulated in this study to include an iterative negotiation dimension corresponds to a *repeated game* or *supergame* in game theory. Under certain conditions *supergames* are known to converge on *socially optimum* equilibrium solutions (i.e. collaborative solutions) that are globally superior to the (optimal) selfish behaviour in the context of non-repeated (i.e. one-shot) games (Fudenberg and Tirole, 1991). The concept of supergames could also be used on a larger scale, to model principal and agent anticipation of upcoming planning cycles (and, potentially, memory of past planning cycles). Ultimately, both scales could be nested (i.e. iterative negotiation within each planning cycle, combined with anticipation of upcoming planning cycles). Although technically challenging, these hypothetical nested-supergame models might be harnessed for practical application using a *metagaming* approach (Howard, 1971), potentially providing a wealth of valuable insight to guide high-level government policy-makers.

[16]There has been considerable debate over the validity and necessity of including even-flow or non-declining yield constraints in wood supply optimisation models (Gunn, 2009). Nonetheless, one or the other of these constraint formulations has traditionally been included in almost all wood supply models in Canada since the advent of the use of linear programming to optimise wood supply planning (with the notable exception of the province of Ontario). We have included even-flow constraints in both the classic and bilevel optimisation model formulations used in this study, as this allows us to measure the impact of extending the *status quo* wood supply model formulation to include explicit anticipation of industrial fibre consumption behaviour. For more information on the effects of even-flow constraints and alternative model formulations, we invite the reader to consult Luckert and Williamson (2005).

5 Conclusion

Paradis et al. (2013) conjecture that extending the classic wood supply model formulation to anticipate industrial fibre consumption would improve wood supply stability (i.e. mitigate risk of wood supply failure). We test this conjecture.

We framed this problem using agency theory, and proposed mathematical formulations to describe the optimisation problems of the principal and the agent. We then combined principal and agent problems into a bilevel optimisation model.

Using a counter-example, we showed that the general case of the bilevel problem is non-convex. We presented a solution algorithm to solve the general case to global optimality, through enumeration of feasible solutions. However, an enumeration-based strategy is computationally intractable for realistically-sized instances. By imposing a restrictive condition on the topology of the agent's problem, we isolated a special case of the bilevel problem, which can be decomposed into output-wise convex subproblems. We presented an algorithm that solves the special case to global optimality.

We tested our solution methodology on a synthetic dataset of realistic size and complexity, and compared results to output from the classic (single-level) wood supply optimisation model. Using a series of five scenarios, we showed that the bilevel model improves long-term wood supply stability, although instability is not completely eliminated by the bilevel model. We showed that the principal can compensate for this residual instability by withholding (i.e. not attributing) a large fraction of bilevel AAC. We conjecture that a similar stabilising effect could be achieved more efficiently (i.e. at a lower cost in terms of withheld bilevel AAC) using a more direct buffer stock modelling approach.

The bilevel solution algorithm for the special case converges on a global optimal solution in less than twice the time required to solve the classic (single-level) model formulation. Considering that these wood supply models are solved infrequently (i.e. once per planning cycle), this increase in solution time is not obviously problematic. The bilevel model has the same output data format as the classic model and can be solved using comparable computational effort. As such, the bilevel model formulation constitutes a technically compatible and conceptually superior alternative to the classic model.

The current study examines the performance of a bilevel model formulation in the context of a two-stage principal-agent game. We recommend that research effort on bilevel wood supply model formulations be extended to *supergame* contexts, both in terms of intra-cycle principal-agent negotiations, and inter-cycle anticipation of future wood supply planning games. To cope with the complexity of using these hypothetical nested-supergame models in a practical government-policy-setting environment, we suggest the adoption of a metagaming approach to wood supply planning as an appropriate starting point for further study.

6 Acknowledgements

This study was supported by funding from the *FORAC Research Consortium*, the *Fonds de recherche et de développement en foresterie (FRDF)*, the *Fonds de recherche du Québec—Nature et technologies (FQRNT)*, the *Centre interuniversitaire de recherche sur les réseaux d'entreprise, la logistique, et le transport (CIRRELT)*, and the *NSERC Strategic Research Network on Value Chain Optimization (VCO).*

Appendices

A Sources of Principal-Agent Antagonism

The principal has the long-term responsibility to ensure a sustained wood supply (hence the even-flow constraints in the wood supply model), but aims to maximise economic activity from the exploitation of the forest resource (hence the wood-supply-maximisation objective function). The agent aims to maximise short-term profit by transforming the wood supply into forest products (e.g. lumber, paper, etc.). The antagonism between the principal and agent is linked to either *(a)* binding agent capacity constraints or *(b)* the presence of negatively-valued subsets of the wood supply. Either of these factors will prevent the profit-maximising agent from consuming the entire wood supply, which in turn induces the problematic negative consumption bias described in Paradis et al. (2013).

The test dataset used in the computational experiments features binding agent capacity constraints. Our test dataset has two lines (which we will refer to as *hardwood* and *softwood*, based on an aggregation of tree species that grow in our test forest). All the hardwood harvested from the forest must pass through a single hardwood sawmill. The hardwood sawmill capacity is approximately one third of the maximum sustainable hardwood supply level determined by the principal using the classic wood supply optimisation model. The softwood line is profitable and has sufficient capacity to process the entire softwood supply offered by the principal. The agent therefore has an incentive to utilise his entire softwood allocation, but limit his hardwood consumption to the capacity of the hardwood sawmill. We have only permitted clear-cut harvesting in our test model, which means the agent only has take-all and leave-all options for each harvestable forest unit. In order to achieve the correct hardwood/softwood mix in his harvesting plan, the agent may select harvest units that have a lower proportion of hardwood than the harvest units appearing in the first period of the principal's optimal wood supply solution. This species-biased deviation from the principal's wood supply plan increases risk of future wood supply shortages.

In the case of wood supply offers with negatively-valued subsets, the only way the principal has to motivate the agent to act is by allowing him to harvest part of the forest (i.e. the agent can choose to consume any subset of wood supply offered by the principal). In practice, it is difficult (impossible) for the principal to force the agent to consume timber at a net loss, so this is a real problem. The principal is incited to propose plans where part of the output is not interesting for the agent, as this allows him to increase his wood supply offer (this is desirable, given his objective function). Including this negatively-valued part in the short-term wood supply allows the principal to increase simulated long-term wood supply offer[17]. However, the agent only plans his consumption on a short-term basis, and has no incentive to consume the negatively-valued subset of wood supply. It may be impossible for the principal to offer the globally optimal plan *and* force the agent to use all of the wood offered. By failing to consume the uninteresting part of the supply, the agent may compromise feasibility of the principal's wood supply plan.

We illustrate this second source of antagonism with an example. Suppose the principal P can offer H_1 to the agent A, which has a value $v(H_1) = 10$. To this offer, the principal can add H_2, which has a value $v(H_2) = 2$. However, H_1 and H_2 can only be consumed sustainably if they are bundled with H_3, which has a value $v(H_3) = -1$. The best long-term solution for both parties is for the agent to consume H_1, H_2 and H_3 for $v(H_1 \cup H_2 \cup H_3) = 11$. However, the principal knows that if he offers all three lots, the profit-maximising agent will only take H_1 and H_2 for $v(H_1 \cup H_2) = 12$. The principal knows that this is unsustainable, so he only offers H_1 which is sustainable but has a lower value of 10.

Note that if the principal could bundle the uninteresting part with a more interesting surplus, then the antagonism would disappear. However, this bundling would require a more highly-constrained contract binding the agent to the principal. This bundling option is not typically available to the principal in practice, leaving him with no rational choice but to lower the wood supply offer until the agent willingly consumes it all. Determining the maximum species-wise even-flow wood supply offer that will be totally consumed by the agent is not a trivial problem. The antagonism between the two levels can induce non-convexity and non-linearity in the solution space, when we constrain the principal's problem such that a wood supply contract is principal-feasible only if the agent consumes it entirely.

B Proof of Non-Convexity

This appendix contains a proof of non-convexity of the solution space for the general bilevel wood supply problem. It may be helpful to recall that, for a convex set of feasible solutions, any linear combination (i.e. convex combination) of two solutions (e.g. $0.5x_1 + 0.5x_2$) will yield a third feasible solution. The basis for our proof of non-convexity is to show, using a simple counter-example, that this property does not hold for all general bilevel problem instances.

[17] For example, simulating harvesting of relatively unproductive or over-mature parts of the forest and regenerating them into higher-productivity stands in a wood supply model may increase the simulated availability of fibre in a future time period. This *allowable cut effect* is a well documented, but potentially problematic, forest policy instrument. For more information, see Luckert and Haley (1995).

Specifically, we describe three solutions to our hypothetical counter-example problem. These three solutions lie along a line segment in solution space. The endpoints of this line segment are bilevel-feasible, however the mid-point is infeasible. By definition, this solution space cannot be convex. Because the counter-example problem is an instance of the general bilevel problem, we can conclude that the general bilevel problem can be non-convex. Any optimisation algorithm for the general bilevel problem would therefore have to assume non-convexity of the solution space, or risk terminating prematurely at a local optimal solution.

Also, it may be helpful to recall the definition of the general bilevel problem solution space. For a wood supply solution to be bilevel-feasible, it must of course be both upper- and lower-level feasible. Furthermore the wood supply must be entirely, and willingly, consumed by the profit-maximising agent in the lower-level model. The second solution of our solution triplet (i.e. the midpoint of our line segment in bilevel solution space) is bilevel-infeasible because it does not respect the second condition for feasibility, *viz.* the agent will not willingly consume the entire wood supply, as it is more profitable for him to leave one unit of hardwood unconsumed.

We now describe the counter-example problem instance as follows. The context for our counter-example is a simple setup where the principal offers a supply of both softwood and hardwood to the agent. Similarly to the dataset we use in our case study (see §2.3.1), the agent is actually composed of two independent sub-agents (i.e. softwood and hardwood lines) which must be independently profitable. Each sub-agent maximises his own profit (i.e. is not willing to reduce his profit for the benefit of the other). There are three types of transformation processes: boards, paper, and cogeneration. The *boards* process is equally profitable for both lines ($+50\,\$/\mathrm{m}^3$ for both softwood and hardwood), and has a transformation capacity of two input units for each line. The *paper* process is profitable for both lines, but at different rates ($+50\,\$/\mathrm{m}^3$ for the softwood line, $+10\,\$/\mathrm{m}^3$ for the hardwood line); it has a transformation capacity of three input units for each line. The *cogen* process is marginally profitable for the softwood line ($+1\,\$/\mathrm{m}^3$), and marginally unprofitable for the hardwood line ($-1\,\$/\mathrm{m}^3$); it has a transformation capacity of one input unit for each line.

Transforming inputs using the paper process requires utilisation of a common resource for both lines (for example, this could correspond to pulp digester capacity). The common resource has a limited capacity of 6 units, which quickly becomes saturated. There is a difference in line-wise efficiency for the utilisation of the common resource. The softwood line uses two units of the common resource for each unit of input transformed, whereas the hardwood line uses one unit of the common resource for each unit of hardwood consumed (for example, this could correspond to softwood chips requiring twice as much time to digest as hardwood chips). Problems with non-convexity of the solution space may arise when this common resource becomes saturated.

We illustrate the counter-example in Figure B.1. We use the symbols S and H to illustrate shared resource capacity utilisation by softwood and hardwood lines, respectively. We show three optimal agent resource allocations, corresponding to three different wood supply offers from the principal. These three wood supply offers correspond to three distinct solutions, which form a line segment in bilevel solution space. The midpoint of this line segment is bilevel-infeasible.

The starting point of the line segment in solution space is a wood supply offer of 4 units of softwood and 4 units of hardwood (see Figure B.1, Solution 1). The optimal allocation of the softwood supply is 2 units to boards and 2 units to paper. The optimal allocation of the hardwood supply is 2 units to boards and 2 units to paper. Shared paper resource capacity is saturated, with 4 resource capacity units utilised by the softwood line and 2 units utilised by the hardwood line. The agent consumes the entire offer, for a profit of \$320, therefore this point is bilevel-feasible.

The endpoint of the line segment in solution space is a wood supply offer of 6 units of softwood and 2 units of hardwood (see Figure B.1, Solution 3). The optimal allocation of the softwood supply is 2 units to boards, 3 units to paper and 1 unit to cogen. The optimal allocation of the hardwood supply is 2 units to boards. Shared resource capacity is saturated, with all 6 resource capacity units utilised by the softwood line. Once again, the agent consumes the entire offer, for a profit of \$351, therefore this point is also bilevel-feasible. This also corresponds to the global optimal solution for this problem.

The midpoint of the line segment in solution space is a wood supply offer of 5 units of softwood and 3 units of hardwood. The optimal allocation of the softwood supply is 2 units to boards and 3 units to paper. The optimal allocation of the hardwood supply is 2 units to boards. Shared resource capacity is saturated, with all 6 resource capacity units utilised by the softwood line. The agent does *not* voluntarily consume the entire offer—his maximum profit of \$350 for this wood supply offer is achieved by leaving one unit of hardwood supply unconsumed (the agent avoids allocating the remaining unit of hardwood to the marginally unprofitable hardwood cogen process). Thus, the midpoint solution is bilevel-infeasible.

Given that all three solutions are located along a line in solution space, infeasibility of the intermediate point proves non-convexity of this solution space[18].

[18] By definition, given a line segment whose endpoints lie inside a convex space, it is not possible for any point along this line segment to lie outside the convex space.

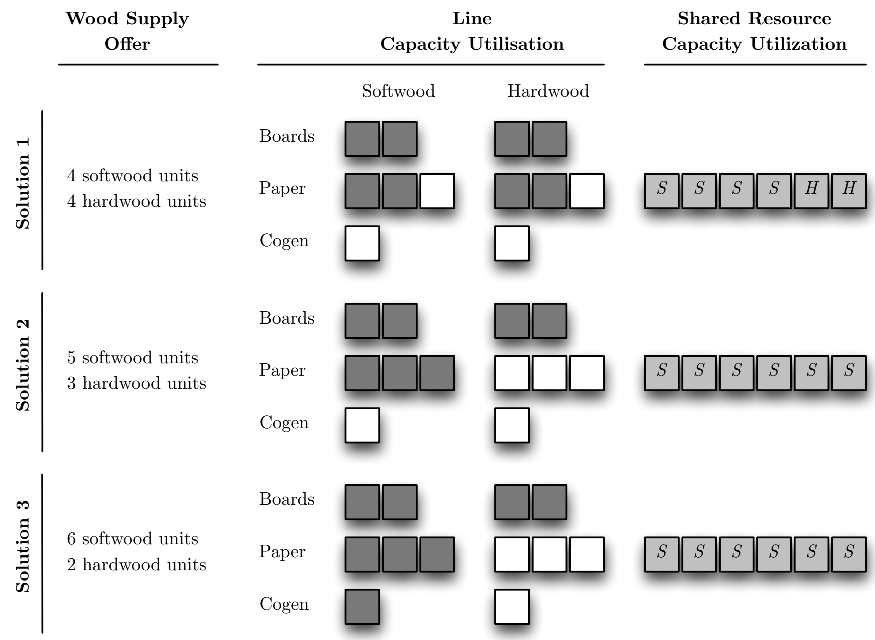

Figure B.1: Simple counter-example illustrating non-convexity of bilevel problem solution space

C Lower-Level Model Implementation Notes

We present a more detailed description of our simulation framework's underlying data model. The LogiLab platform is used to model both the agent-anticipation mecanism in the bilevel model and to simulate actual agent fibre consumption in the second stage of the iterative rolling-horizon replanning simulation process. The LogiLab data model can be described as a network of abstract processors[19] connected by product flows. Each processor node represents a business unit in the value creation network (e.g. sawmill, pulpmill, end-product-client, etc.).

At the upstream (source) end of this network is the interface between upper and lower level models. Outputs from the first planning period in the upper level[20] model (i.e. outputs $o \in O$) represent raw material that can be transformed by the lower-level network. At the downstream (sink) end of the lower level network, external clients are willing to pay exogenously determined unit prices to satisfy a bounded demand for each end product. Profits are induced by pulling a subset of potential fibre supply through the network to satisfy a subset of end client demand.

[19] An abstract processor consumes inputs and resources, and produces outputs. Abstract processors can be topologically connected to form a network.

[20] In the context of our bilevel model, the terms *upper level* and *lower level* refer to principal and agent decision variables, respectively.

When raw upper-level model outputs are first pulled into the lower network, they must go through one of several front-line processor nodes that convert the raw wood supply into species-wise assortments of logs. These front-line processors simulate the interface between the forest and the mills (i.e. the process of harvesting and delivering logs to mills, including transportation cost, which can vary depending on the forest zone from which the raw volume inventory was procured).

Raw upper-level volume is classified by species, and this species-wise distinction may (optionally) be maintained as the outputs are pulled into the lower-level network, depending on configuration of front-line processors. For example, in the case of our test dataset, the front-line processors are configured to convert raw upper-level volume into assortments of either hardwood or softwood logs of various sizes[21].

Due to the abstract nature of the lower-level processor implementation, it is possible to simulate any combination of divergent and convergent product flows. Strictly divergent networks correspond to special case 1, and can be solved to global optimality using Algorithm 1. Special case 2 occurs when the network includes convergent product flows, but no joint capacity constraints are saturated. Although somewhat more difficult to detect, special case 2 is not problematic and can also be solved to global optimality using Algorithm 1. For special cases 1 and 2, each product line $o \in O$ can be treated as an independent subproblem.

The product-wise subproblems can be represented using the lower-level model by disabling all non-targeted outputs[22]. We can then easily solve each subproblem to obtain the optimal subproblem solutions without having to explicitly locate intermediate inflection points of profit function $p(v_o)$ (see Figure 2). This corresponds to the maximum volume that the agent can be expected to consume for a given output o, which we use as an upper bound on harvest level for output o in the final step of Algorithm 1. In other words, optimal solutions of the output-wise subproblems can be used to define valid (and sufficient) cuts for the solution space of the bilevel optimisation problem (for special case instances).

We aggregate all fibre flows into the agent model into a number of input lines, corresponding to the species groups the principal uses to express AAC (e.g. our test case has *hardwood* and *softwood* lines), and require fibre flows from each of these lines to be independently profitable using *line-wise profitability constraints*. The purpose of the line-wise profitability constraints is to model a common situation in many real-world value creation networks.

Subsets of the agent network may be independently owned and managed, and tend to specialise in processing certain species (e.g. hardwood or softwood). Demand for a given species group may be limited by local processing capacity,

[21] Our upper-level dataset uses a more fine-grained classification of tree species, which is aggregated into *hardwood* and *softwood* log types by the front-line processors.

[22] Non-targeted outputs can be disabled by manipulating input parameters of the upper level model, setting conversion efficiency and conversion cost of front-line processors to null values. Instead of converting wood supply units to log assortments, the modified front-line processors now consume all non-targeted outputs at zero cost, which blocks non-targeted outputs from further flow through the lower-level model network.

exogenous end-client demand, or exogenous market prices (i.e. production will cease before capacity is saturated if cumulative unit procurement and processing costs exceed unit revenues, for a given product). The line-wise profitability constraints ensure that all mills that process a certain species group (i.e. each line) ceases production before it drops below maximum profit. The basic LogiLab model formulation assumes centralised network planning (i.e. maximises profit for the entire network)—the model would, were it not for the line-wise profitability constraints, induce one or more lines to continue production beyond the profitability threshold if this was beneficial to the network as a whole.

The line-wise profitability constraints are implemented by updating the species-wise upper bounds set by the principal in the first stage, in the event that AAC exceeds the maximum volume that a given line can profitably consume. To determine these maximum line-wise profitable volumes, we simply solve the output-wise submodels (with non-targeted outputs disabled) described in bilevel solution methodology (see §2.2).

D Computational Experiment Dataset

This section provides more information on the synthetic dataset we used as input to the lower level model for the computational experiments. The case study dataset does not represent an actual agent network, however these parametres were synthesised from realistic data that was compiled for previous research projects realised by the FORAC Research Consortium. See §**??** for a schematic representation of the value creation network used in the computational experiments.

The lower level data model allows for an arbitrary number of external fibre suppliers with each supplier potentially having distinct input parametres. Our test dataset features a single fibre procurement source, thus all input units of fibre in the lower-level model have the same cost. In reality, procurement cost may vary due to a number of factors (e.g. distance from mill, accessibility, terrain, choice of harvesting system and sylviculture prescription, species mix, etc.). However, modelling variable procurement cost was not necessary given the objectives of our computational experiments, and would only have served to obfuscate the results. We model the forest (i.e. external supply of fibre for the network) as a business unit, which encapsulates processes that convert raw fibre to assortments of logs. The forest business units produce three types of softwood logs (small, medium, large) and one type of hardwood log.

The three types of softwood logs can flow to any of the three softwood sawmills in the network. The softwood sawmills produce four types of softwood lumber (2×3, 2×4, 2×6, 2×8) and softwood chips. Softwood lumber can be sold to a single softwood lumber external customer.

The hardwood logs are processed by the single hardwood sawmill, which produces a single type of hardwood lumber and hardwood chips. Hardwood lumber can be sold to a single hardwood lumber external customer.

Both hardwood and softwood chips may flow from the sawmills to the paper

Figure D.1: Schematic representation of test value creation network dataset

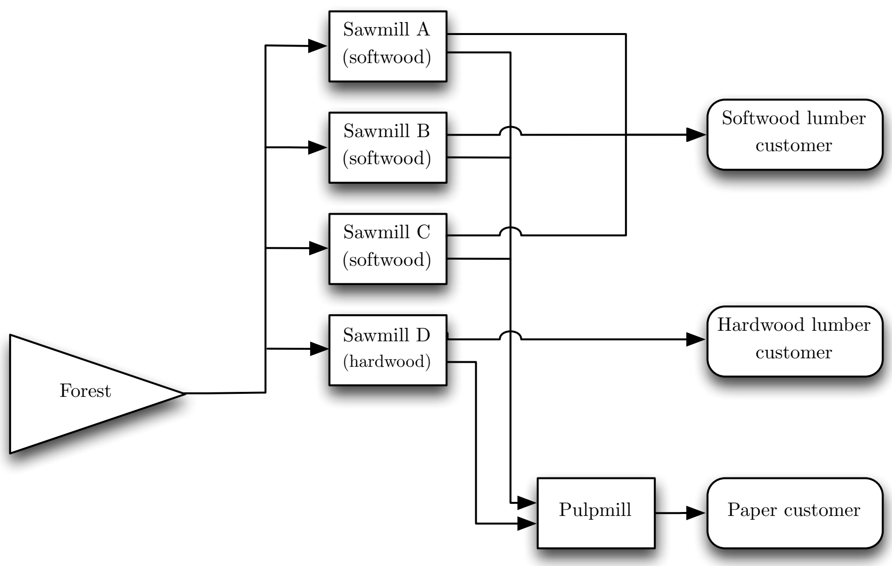

mill, where they can be converted to paper. The paper can be sold to one of two paper external customers. There is no external customer for chips.

The lower level model includes *storage* process at each business unit. Raw volume can be stored at roadside in the forest. Sawmills have storage processes for logs, lumber and chips. The pulpmill has storage processes for chips and paper. Finally, each external client has a storage process for the products they accept. Periodic storage costs are modelled as a proportion of cumulative product cost, which depends on the amount of processing that a unit of product has undergone. Thus, storage costs increase as products move through the network, with the least expensive storage cost being at source nodes and most expensive storage costs being at sink nodes.

References

Jonathan F Bard and James E Falk. An explicit solution to the multi-level programming problem. *Computers & Operations Research*, 9(1):77–100, 1982.

D. Beaudoin, J.M. Frayret, and L. LeBel. Negotiation-based distributed wood procurement planning within a multi-firm environment. *Forest Policy and Economics*, 12(2):79–93, 2010. doi: 10.1016/j.forpol.2009.07.007.

Timothy Norman Bogle. *Timber supply on public land in response to catastrophic natural disturbance: a principal-agent problem.* PhD thesis, University of Victoria, 2012.

Tony Byrne, Cameron Stonestreet, Brian Peter, et al. Characteristics and utilization of post-mountain pine beetle wood in solid wood products. *The mountain pine beetle: A synthesis of biology, management, and impacts on lodgepole pine. Natural Resources Canada, Canadian Forest Service, Pacific Forestry Centre*, pages 233–253, 2006.

Canadian Council of Forest Ministers. Defining Sustainable Forest Management in Canada: Criteria and Indicators. Technical Report Fo75-3/5-2003E, Natural Resources Canada, Canadian Forest Service, 2003.

Canadian Council of Forest Ministers. A Vision for Canada's Forests: 2008 and beyond. Technical Report Fo4-23/2008E-PDF, Natural Resources Canada, Canadian Forest Service, 2008.

Benoît Colson, Patrice Marcotte, and Gilles Savard. An overview of bilevel optimization. *Annals of operations research*, 153(1):235–256, 2007. doi: 10.1007/s10479-007-0176-2.

L.S. Davis, K.N. Johnson, P.S. Bettinger, and T.E. Howard. *Forest Management: To Sustain Ecological, Economic, and Social Values*, chapter 13, page 754. McGraw Hill, 2001.

Stephan Dempe. Annotated bibliography on bilevel programming and mathematical programs with equilibrium constraints. *Optimization*, 52(3):333–359, 2003.

José Fortuny-Amat and Bruce McCarl. A representation and economic interpretation of a two-level programming problem. *The Journal of the Operational Research Society*, 32(9):783–792, 1981.

D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991. ISBN 9780262061414.

E. Gunn. Some Perspectives on Strategic Forest Management Models and the Forest Products Supply Chain. *INFOR: Information Systems and Operational Research*, 47(3):261–272, 2009. ISSN 0315-5986.

Eldon Gunn. Models for strategic forest management. In Andres Weintraub, Carlos Romero, Trond Bjørndal, and Rafael Epstein, editors, *Handbook of Operations Research in Natural Resources*, chapter 16, pages 317–341. Springer Science+Business Media, 2007.

Nigel Howard. *Paradoxes of rationality: theory of metagames and political behavior*. MIT Press, 1971.

W. Jerbi, J. Gaudreault, S. D'Amours, M. Nourelfath, S. Lemieux, P. Marier, and M. Bouchard. Optimization/simulation-based framework for the evaluation of supply chain management policies in the forest product industry. In *2012 IEEE International Conference on Systems, Man, and Cybernetics (IEEE SMC 2012)*. IEEE (accepted), 2012.

Jean-Jaques Laffont and David Martimort. *The theory of incentives: the principal-agent model*. Princeton University Press, 2002. ISBN 9781400829453.

M. K. Luckert and T. Williamson. Should sustained yield be part of sustainable forest management? *Canadian Journal of Forest Research*, 35(2):356–364, 2005.

Martin K. Luckert and David Haley. The allowable cut effect as a policy instrument in Canadian forestry. *Canadian Journal of Forest Research*, 25 (11):1821–1829, 1995.

A.H. Mathey, H. Nelson, and C. Gaston. The economics of timber supply: Does it pay to reduce harvest levels? *Forest Policy and Economics*, 11(7):491–497, 2009.

National Forestry Database, 2014. URL `http://nfdp.ccfm.org`. Last accessed on September 1, 2014.

Gregory Paradis, Luc LeBel, Sophie D'Amours, and Mathieu Bouchard. On the risk of systematic drift under incoherent hierarchical forest management planning. *Canadian Journal of Forest Research*, 43(5):480–492, 2013. doi: 10.1139/cjfr-2012-0334.

Frédéric Raulier, Narayan Dhital, Pierre Racine, Rebecca Tittler, and Andrew Fall. Increasing resilience of timber supply: How a variable buffer stock of timber can efficiently reduce exposure to shortfalls caused by wildfires. *Forest Policy and Economics*, 2014.

C. Schneeweiß. *Distributed Decision Making*. Springer Verlag, 2003.

Tennessee Trent, Val Lawrence, Kathy Woo, et al. A wood and fibre quality-deterioration model for mountain pine beetle-killed trees by biogeoclimatic subzone. 2006.

Dajun Yue and Fengqi You. Game-theoretic modeling and optimization of multi-echelon supply chain design and operation under stackelberg game and market equilibrium. *Computers & Chemical Engineering*, 71:347–361, 2014.

Wenyuan Zhai, Yingzi Zhao, Xiangru Lian, Mingming Yang, and Fadian Lu. Management planning of fast-growing plantations based on a bi-level programming model. *Forest Policy and Economics*, 38:173–177, 2014.