

IMPLEMENTING FAUSTMANN: STOCHASTIC DYNAMIC PROGRAMMING IN SPACE

Harry J. Paarsch*
Department of Economics
University of Melbourne

John Rust
Department of Economics
University of Maryland

March 2008

Preliminary Draft: Please do not quote without permission of the authors.

Abstract

We construct an intertemporal model of rent-maximizing behaviour on the part of a timber harvester under potentially multi-dimensional risk as well as geographical heterogeneity. Subsequently, we use recursive methods (specifically, the method of dynamic programming) to characterize the optimal policy function, the rent-maximizing timber-harvesting profile. One noteworthy feature of our application to forestry in British Columbia is the unique and detailed information we have organized in the form of a dynamic geographical information system to account for site-specific cost heterogeneity in harvesting and transportation as well as uneven-aged stand dynamics in timber growth and yield across space and time in the presence of stochastic lumber prices. Our model is a powerful tool with which to conduct policy analysis.

JEL Classification Numbers: C61, C81, D92, H32, L73, Q23.

Keywords: stochastic dynamic programming; optimal timber rotation; spacial economics; rent maximization.

* Corresponding author: Fifth Floor, Economics and Commerce Building;
Melbourne, Victoria, 3010, AUSTRALIA. Telephone: 61 3 8344 4655; E-mail:
hpaarsch@unimelb.edu.au

Acknowledgements

We are grateful to several individuals at the British Columbia Ministry of Forests and Range (prior to 16 June 2005, this ministry was known as the Ministry of Forests) for their coöperation and help. Peter Fuglem, then Manager of the Development and Policy Section of the Timber Supply Branch (which is now known as the Forest Analysis Branch), took an early interest in the project after Bill Howard, then Director of the Revenue Branch, organized an inter-ministry meeting at which the basic idea was pitched. Thereafter, Dave Waddell, Tim Bogle, and Albert Nussbaum (all of the Forest Analysis Branch) were instrumental in helping our research assistant, Mark Weldon, construct the geographical information system that was central to organizing the data for the project.

Cam Bartram and Rob Drummond, of the Ministry of Sustained Resource Management, provided much early helpful advice necessary when implementing the programme *VDYP* for our problem; Shelley Grout, of the Research Branch at the Ministry of Forests and Range, provided us a copy of the programme *VDYP*. Ken Polsson, of the Research Branch, gave extensively of his time using the programme *TASS* to create the data necessary for us to estimate the growth-yield equations. Jim Goudie and Ken Mitchell, also of the Research Branch, provided advice that was useful in evaluating our timber growth model.

We cannot overstate the contribution of our research assistant, Mark Weldon. A mechanical engineer by training, a forester in a previous profession, and a Ph.D. student in environmental engineering (at that time), Mark toiled quietly and efficiently for two years to build us an extraordinary data set; we are in his debt. At a later stage in the research, Clinton J. Levitt also provided invaluable research assistance, particularly by using animation to depict the solutions to our dynamic programmes.

The financial support of the National Science Foundation under grant SES-0241509 made the entire project possible; we are also grateful to the NSF for ongoing support and, in particular, to Dan Newlon for his encouragement.

1. Motivation and Introduction

The application of recursive methods and, in particular, the method of dynamic programming to solve economic decision problems involving both risk and time is by now quite standard, especially in the natural-resource economics literature. In this paper, we report research that goes beyond that which builds on Faustmann (1849): *viz.*, that reported by Kaya and Buongiorno (1987); Brazee and Mendelsohn (1988); Morck, Schwartz, and Strangeland (1989); Reed and Clarke (1990); Haight and Holmes (1991); Thomson (1992); Reed (1993); Provencher (1995); or Reed and Haight (1996). First, we take geography seriously, both in the planar sense and in the three-dimensional sense. Second, we take site-specific heterogeneity seriously both on the cost side in terms of harvesting and transportation and on the growth and yield side in terms of heterogeneous stands of timber. Third, we model initial conditions. In particular, we do not take as a starting point a steady-state allocation, or even an optimal allocation. Instead, we take the existing, potentially uneven, age distribution of timber as given and derive the optimal policy function—the optimal timber-harvesting profile. Fourth, we use best-practice biological methods to model the dynamics of uneven-aged forest growth and yield. Fifth, in the past, economists have typically demonstrated their methods by solving simple examples in closed-form or they have imposed conditions sufficient to sign comparative-static predictions. Below, we harness recent developments in computational methods to solve numerically for the optimal policy function on hundreds of thousands of potential harvest sites.

We are able to accomplish these advances because we have had access to information from unique and elaborate databases maintained by the Ministry of Sustainable Resource Management, Terrestrial Information Branch, and the Ministry of Forests and Range, Forest Analysis Branch (formerly, the Timber Supply Branch), in the province of British Columbia, Canada. From these different databases, we have constructed a dynamic geographical information system whose different relations we have then exploited to develop estimates of site-specific cost heterogeneity in harvesting and transportation as well as uneven-aged stand dynamics in timber growth and yield across space and time.

Our framework is quite rich: using it, one can conduct a variety of different policy

experiments that previous researchers could not. For example, below, we compare our predicted optimal harvesting policy with the harvests that have occurred during the last eight years. In another vein, from the perspective of industrial organization, in Paarsch and Rust (2008), we have also investigated how the the province of British Columbia might behave as a “big player” in the lumber market.¹

2. Previous Theoretical Structure

In order to place our research in an historical context, we first develop a notation and then outline a simple theoretical framework which some previous researchers have used to investigate the optimal harvesting of timber. This work allows us to isolate a variety of different and important features of the timber-harvesting problem. Subsequently, we go on to extend the existing research in section 3.

2.1. Biological Environment

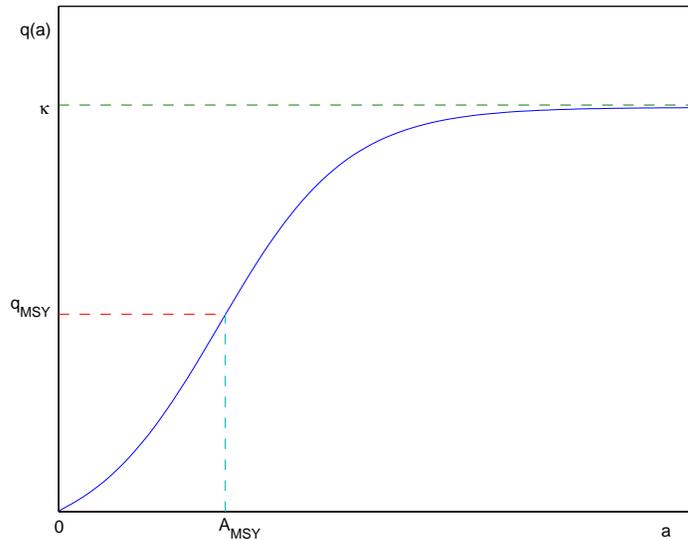
Forests are biological assets whose net returns typically depend on, among other things, the age of the standing timber. In Figure 1, we present a stylized graph of the relationship between the age a of an even-aged stand of timber on a harvest block of a particular area and its average (mean) volume $q(a)$. Note that with aging the mean volume rises, initially quite quickly, but subsequently at a slower rate. In the absence of disaster or disease, the mean volume will approach an asymptote κ . Both the rate-of-change of mean volume $\Delta q(a)$ and the asymptote of mean volume κ are species dependent and can be influenced by a variety of environmental factors, some of which are under the direct control of the forester—*e.g.*, the density of newly-planted stems.

2.2. Maximum Sustainable Yield

One commonly-used criterion for determining the harvest of timber as well as other biological resources, such as fish, involves the concept of *maximum sustainable yield*.

¹ Sawmills in British Columbia produce about twenty-five percent of the softwood lumber supply in North America.

Figure 1
Age Path of Timber Volume



In Figure 2, q_{MSY} denotes that mean volume at which a stand's rate-of-change in mean volume is at its maximum Δq_{MSY} . Under the maximum-sustainable-yield management criterion, an age A_{MSY} exists at which to harvest the stand, yielding q_{MSY} units of timber; see figure 1 as well.

Suppose there are a total of N harvest blocks. Consider dividing all of these blocks into A_{MSY} different types of plantations, each of a different age, but each with the same species and stem density. Under the maximum-sustainable-yield criterion, in a steady state, a uniform age distribution $f(a)$, often referred to as a *normal forest*, will obtain. Such an age distribution is depicted in Figure 3. Given the N harvest blocks, this implies that (N/A_{MSY}) sites will be harvested in each period, yielding an average total volume of timber $[Nq(A_{MSY})/A_{MSY}]$, which we shall denote \bar{t} .

2.3. Lumber Production

A useful approximation to the milling process of timber is the Leontief production function. For this technology, hours of labour input h and cubic metres of timber t are combined according to a fixed-coefficient production function to yield an output,

Figure 2
Biological Rate-of-Change

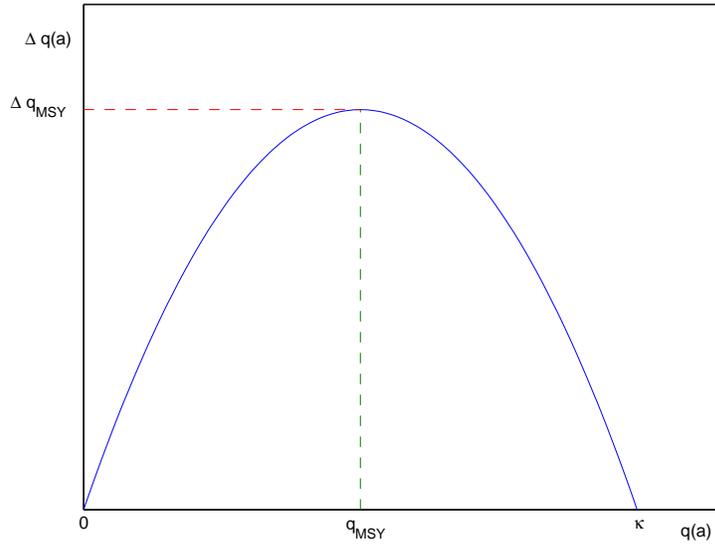


Figure 3
Age Distribution, Normal Forest

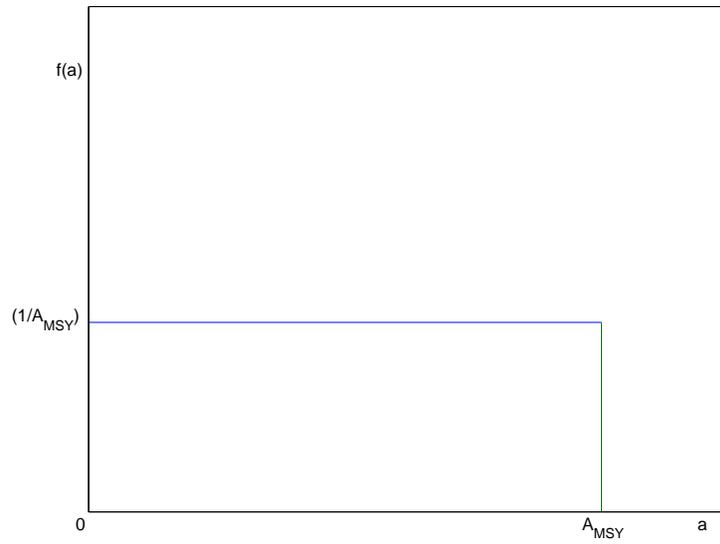
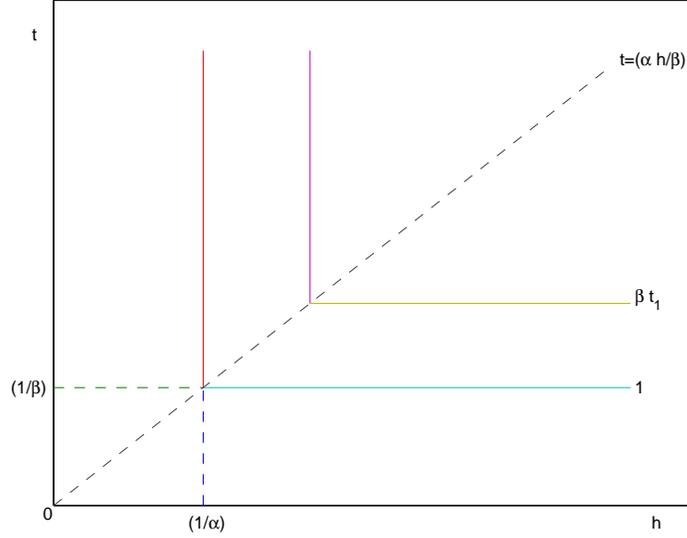


Figure 4
Leontief Production Function



lumber ℓ , in thousands of board according to the following:

$$\ell = \min(\alpha h, \beta t) \quad 0 < \alpha, \quad 0 < \beta.$$

The parameter β is often referred to by lumbermen as the *lumber recovery factor* (LRF). The unit isoquant for this production function is depicted in Figure 4 by the right-angled curve with 1 beside it; in this case $(1/\alpha)$ units of h are combined with $(1/\beta)$ units of t to produce one unit of ℓ . Another representative isoquant, which exceeds that for 1, is depicted to the northeast of 1 and denoted βt_1 .

Given the Leontief technology and assuming that producers take input prices as given, one can derive the following total- (C) and marginal- (c) cost functions:

$$C(\ell; w, s) = \left(\frac{w}{\alpha} + \frac{s}{\beta} \right) \ell$$

and

$$c(\ell; w, s) = \left(\frac{w}{\alpha} + \frac{s}{\beta} \right)$$

where w is the wage of labour and s is the price of timber, often referred to as the *stumpage rate*. Let c^0 denote $c(\ell; w, 0)$, the marginal cost when s is zero. With a

Leontief technology, c^0 is the marginal cost of all factor inputs, excluding the natural-resource input timber.

2.4. Rent Extraction

In many jurisdictions, government agencies often dispose of publicly-owned timber using administratively-set prices. In such situations, questions of what to charge for such publicly-owned assets arise naturally. Conditional on a maximum-sustainable-yield volume \bar{t} , for example, the principle behind determining the optimal stumpage rate \bar{s}^* dates back to Rothery (1945). Basically, absent formal markets for timber, the price that a “central planner” should charge for each cubic metre of timber is that resource’s residual value, its rent—*viz.*,

$$\bar{s}^* = \beta(p - c^0) = \beta \left(p - \frac{w}{\alpha} \right).$$

This can be deduced from the graph in Figure 5. Here, the parameter β in front of $(p - c^0)$ simply ensures that the units match; p and c^0 are in dollars per thousands of board feet of lumber while \bar{s}^* is in dollars per cubic metre of timber. Thus, the units of β are thousands of board feet per cubic metre, those of the LRFs.

2.5. Site Heterogeneity

Different stands of timber often have different LRFs. In the absence of other information, one convenient way to model stand-specific differences in LRFs is as random draws from a probability density function $g(\beta)$. An example of $g(\beta)$ is depicted in Figure 6. Random differences in LRFs then mean that the rent-maximizing stumpage rates depend on stand-specific factors, so the stumpage rate $\bar{s}^*(\beta)$ will be a function of the stand-specific LRF β . Thus, stumpage rates will reflect differential factor rents in the sense of David Ricardo (1817).

Different stands of timber are often located at different distances d (in kilometres) from timber-processing facilities. Typically, transportation costs per unit volume γd are significant, where γ is the cost of transporting the timber equivalent of one

thousand board feet of lumber one kilometre. In this case, the stumpage rate is determined according to the following:

$$\bar{s}^*(\beta, d) = \beta(p - c^0 - \gamma d) = \beta\left(p - \frac{w}{\alpha} - \gamma d\right).$$

This stumpage rate has a location-specific rental component in the sense of Johann von Thünen (1826) as well as a Ricardian component.

2.6. Faustmann's Solution

Most of the analysis considered above has been couched in terms of a steady-state, maximum-sustainable-yield harvest \bar{t} . An unusual feature of this solution, at least from the perspective of an economist, is that \bar{t} , the volume of timber brought to market in each period j , is independent of economic variables and determined solely by biological parameters.

The German forester Martin Faustmann (1849) introduced the rent-maximizing way in which to rotate a forest, which is often referred to as the *Faustmann solution*. To begin, we shall introduce the Faustmann solution in terms of the management of an even-aged stand of timber on a single harvest block.

Assume that k , the costs of planting a harvest block at a particular stem density, are incurred in period 0, while in period A a net revenue of $\beta(p - c^0 - \gamma d)q(A)$ is realized from the sale of the harvested timber as lumber. When the discount rate is δ , the present-discounted profit π from the sale of a single rotation of the timber is

$$\pi(A) = -k + \beta(p - c^0 - \gamma d)q(A) \exp(-\delta A).$$

In a multi-rotation setting, after the harvest of the first stand, another stand will be planted and then harvested and, after that, another, and so forth. Thus, the value of the scarce land deriving from an infinite number of rotations of the same species of

Figure 5
Demand and Supply of Lumber

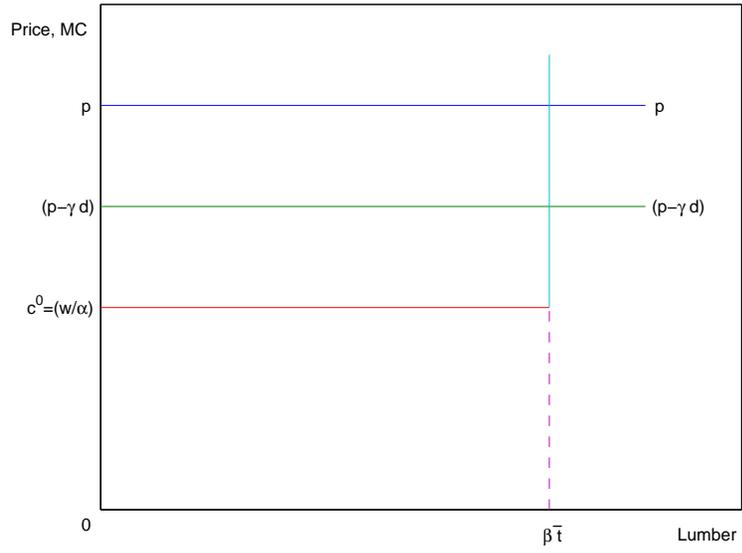
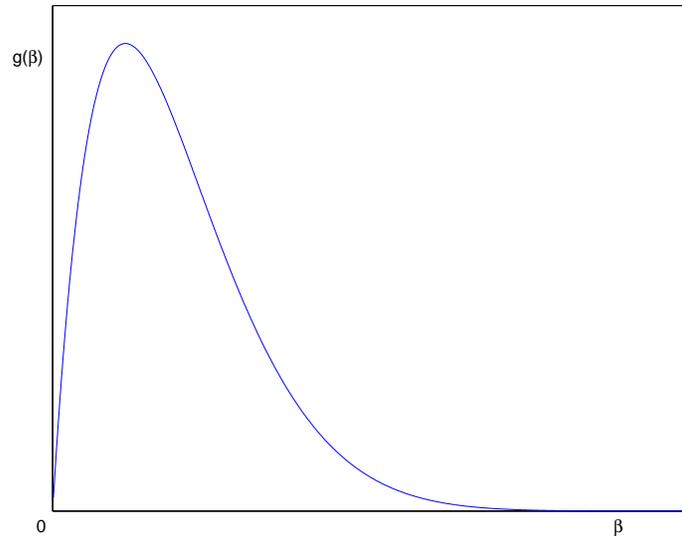


Figure 6
Site-Specific Heterogeneity in Lumber Recover Factors



stand is

$$\begin{aligned}
V(A) &= \pi(A) + \pi(A) \exp(-\delta A) + \pi(A) \exp(-2\delta A) + \dots \\
&= \pi(A) \sum_{i=0}^{\infty} \exp(-i\delta A) \\
&= \frac{\pi(A)}{[1 - \exp(-\delta A)]} \\
&= \frac{[\beta(p - c^0 - \gamma d)q(A) \exp(-\delta A) - k]}{[1 - \exp(-\delta A)]}.
\end{aligned}$$

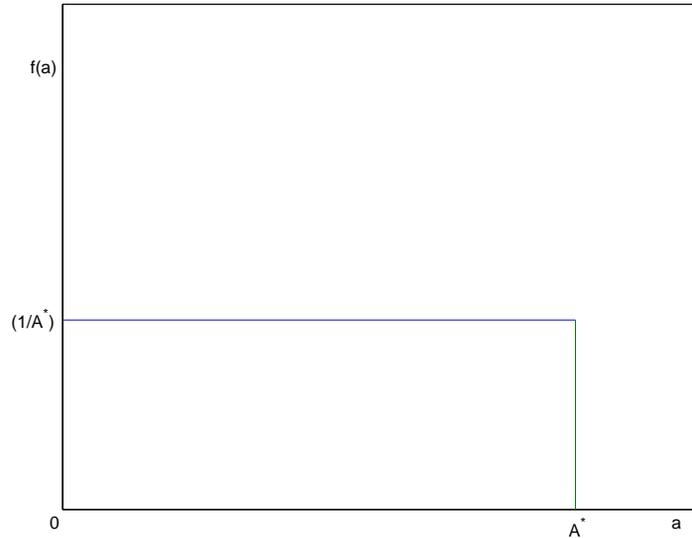
For this block of land, the value-maximizing harvest date A^* is characterized by the following first-order condition:

$$\beta(p - c^0 - \gamma d)q'(A^*) = \delta\beta(p - c^0 - \gamma d)q(A^*) + \delta V(A^*).$$

The term on the left-hand side of the equal sign represents the marginal benefit from holding the stems an extra “period,” while the two terms in the sum on the right-hand side represent the marginal cost. The marginal benefit is the rent-maximizing value of the timber multiplied by the change in volume, while the first term of marginal cost is the opportunity cost of interest on the net revenue and the second term is the rent on the land, often referred to as the *Faustmann land rent*. Under an optimal rotation A^* , the value of the land is $V(A^*)$, so the rent for one period is $\delta V(A^*)$.

Again, consider dividing the N harvest blocks into a number of different types of plantations, each of a different age, but the same species and stem density. This time, however, let there be A^* different ages instead of A_{MSY} . In a steady-state, a normal forest will obtain which is depicted in Figure 7. One can replace \bar{t} with t^* , which equals $[Nq(A^*)/A^*]$, and much of the economic analysis of the maximum-sustainable-yield case presented above carries through without any major modifications. Now, however, A^* depends on the biological parameters (κ for example) as well as p and w (through c_0), and also α , β , γ , δ , and d . Thus, the optimal volume of timber harvested t^* (the supply curve) as well as the optimal stumpage rate s^* depends on these too—*viz.*, $t^*(p; w, \alpha, \beta, \gamma, \delta, d, k)$ and $s^*(p; w, \alpha, \beta, \gamma, \delta, d, k)$. Also, s^* is potentially different from \bar{s}^* .

Figure 7
Faustmann Normal Forest

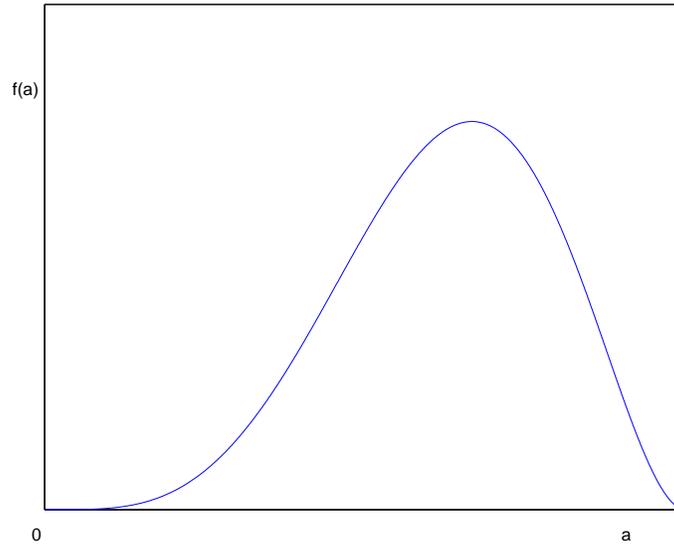


2.7. Common Criticisms of the Faustmann Framework

Practitioners attempting to implement the Faustmann solution often complain that the framework provides little guidance concerning what to do when a forest has previously been unmanaged, so the current age distribution is uneven. For, in many jurisdictions, some stands of timber have never been harvested previously, while others have been harvested, but allowed to regenerate naturally, so the age distributions in such stands are not those of a normal forest in either the maximum-sustainable-yield or the Faustmann sense. This is perhaps the most common criticism of Faustmann's work. Moreover, this criticism does not go away once an optimal, steady-state obtains. For even if the initial forest were in a steady-state, even-aged Faustmann distribution, a change in the economic environment (for example, because of an increase in the price of lumber p), would induce an uneven-age distribution. After several shocks, a representative age distribution might look something like that sketched in Figure 8.

Another criticism of the Faustmann framework is that selective harvesting is implicitly assumed. Often, however, the economic and engineering reality of harvesting timber implies that selective harvesting a particular strata of the age distribution is

Figure 8
Age Distribution of Old-Growth Forest



impossible to do effectively or even responsibly: under certain circumstances, selectively harvesting stems from a forest can weaken the remaining stems—they may be unable to withstand winds. Thus, unlike in even-aged plantation tree farming, where selectively cutting a particular part of the age distribution is both feasible and desirable, in old-growth forests, clear-cutting the entire heterogeneous age distribution in a stand of timber is a fact of life.

A third criticism of the Faustmann framework is that forests exist in space: they are in different planar locations as well as at different elevations. This heterogeneity implies that their growth and yield functions as well as their harvesting and transportation costs are heterogeneous.

A final criticism of the Faustmann framework is that economic variables, such as lumber prices, and biological variables, such as volumes of merchantable timber, are typically subject to stochastic variation over time. These features change markedly the decision problem as the research of Kaya and Buongiorno (1987); Brazee and Mendelsohn (1988); Morck, Schwartz, and Strangeland (1989); Reed and Clarke (1990); Haight and Holmes (1991); Thomson (1992); Reed (1993); Provencher (1995);

as well as Reed and Haight (1996) has shown.

As a practical matter, however, the stochastic variation in lumber prices is much more important than the stochastic variation in timber volumes per hectare. For example, consider data at the monthly level. In discussions with professional foresters, we have learned that, in any month, the standard deviation, at the hectare level, of timber volumes is probably less than 0.1 percent—*i.e.*, less than ten basis points. On the other hand, using market data, we have estimated the standard deviation in the monthly first-differences in the logarithm of representative lumber prices to be around 8.6 percent—around hundred times larger than the variation in timber volumes. This difference persists at other time intervals of measurement. Thus, in our research below, we have focussed on stochastic price variation, treating (mean) growth as deterministic.

3. A Geographical, Intertemporal, and Stochastic Model

Below, we develop a theoretical model which admits the four features of either natural forests or the economic environment discussed in the previous section: first, initial conditions; second, engineering and physical constraints; third, geography; and fourth, stochastic variation in the price of output.

3.1. Recursive Solution via the Method of Dynamic Programming

To provide a solution to the problem having the above features, we adopt a recursive modelling strategy and, in particular, the method of dynamic programming. What we want to do is take the infinite horizon faced by the decision-maker, and break it into a decision to be made this period, and then a continuation into the future. All of our assumptions are made to ensure that such a recursive decomposition can be constructed in a computationally-tractable way.

3.1.1. Assumptions concerning the Economic and Physical Environment

We begin by assuming that the central planner (the government as represented by the Forest Service, but also referred to as the Crown below) has timber-bearing land

which is divided into individual harvest blocks. We shall formulate a dynamic-programming problem whose solution will determine the optimal time at which to clear-cut the timber on a particular block. The objective is to maximize the expected discounted value of rents earned from managing a portfolio of blocks over an infinite horizon. Below, we shall refer to the decision to clear-cut any particular block as the decision “to harvest” the timber on that block. We assume that clear-cutting is optimal because the costs of selectively harvesting individual stems on a block are prohibitively high.

We assume that the timber growing on each block is relatively homogeneous in terms of the age, biomass, and species of trees. In fact, later, to the extent possible, when we implement our framework, the size of an harvest block is chosen to ensure this homogeneity.

We assume that no capacity constraints exist on the resources needed to harvest different blocks; *i.e.*, any particular harvester in a Timber Supply Area (TSA) or any particular Tree Farm Licensee (TFL) can procure additional harvesting capacity at a constant marginal cost.² This assumption allows us to treat different blocks separately, greatly simplifying the dynamic-programming problem. Otherwise, we would have to keep track of the available harvesting capacity and consider how to ration this available capacity during periods when the number of blocks scheduled for harvesting exceeds the harvesting capacity.

A related assumption is that the cost of harvesting timber is independent of the number of blocks to be harvested or the volume of timber harvested on a given block, although the cost of harvesting could depend on the particular characteristics of each block.³ This is equivalent to the assumption that harvesters face a perfectly elastic

² In British Columbia, nearly 90 percent of all timber is on government-owned (Crown) land. Basically, the Crown, through the Minister of Forests and Range, sells the right to harvest the timber on this land in two different ways. During our sample period, the most common way was charging administratively-set prices to a small number of firms who held Tree Farm Licenses or other similar agreements. The terms of these agreements were negotiated over the last three-quarter century, and require that the licensee adopt specific harvesting as well as reforestation plans. About ninety percent of all Crown timber is harvested by firms holding Tree Farm Licenses or similar agreements. The second, and less common way, to sell timber is at public auction—formerly through the Small Business Forest Enterprise Program.

³ In fact, we shall make use of elaborate and unique site-specific data gleaned from a GIS to

supply of firms willing to fell trees on any particular block, and that a decision to cut a larger number of blocks will not significantly bid up the prices these firms charge.

We also assume that British Columbia is a small player in the international market for lumber, so that at any particular time the Crown faces a perfectly elastic demand for timber at the current market-determined spot price.⁴

Finally, we assume that any given block will always be used for growing trees and harvesting timber and that the block has no alternative use. Later, we show how the problem could be modified to allow for a decision to convert permanently the block to a best alternative use, such as conversion to parkland or the sale of a block for housing or industrial use, and so forth.

3.1.2. Solution Heuristic

In our model, harvesting decisions are made at discrete points in time, such as the beginning of each month, quarter, or year. The periodicity of the model can be changed easily, assuming sufficient data exist to estimate transition probabilities of the state variables at sufficiently fine time intervals. The basic decision in the dynamic-programming model is binary: to harvest a particular block or to delay harvesting to a future period.

The dynamic-programming problem we analyze is a considerable generalization of the simple Faustman timber-harvesting problem outlined above because we specify much richer and more realistic models of both timber volume growth and future lumber prices. Our model is often referred to as a *regenerative, optimal-stopping problem*. Similar problems have been analyzed and solved previously by Rust (1987) and Provencher (1995). The *stopping* decision is equivalent to the harvest decision. The problem is referred to as *regenerative* because, once a harvest occurs, we assume

control for this sort of heterogeneity.

⁴ Our framework is sufficiently general that it can be also used elsewhere—*viz.*, Paarsch and Rust (2008b)—to investigate the case where the Crown is a “big player” in the timber market, so its harvesting decisions have an impact on the spot price of lumber. As in the case of capacity constraints on harvesting decisions, this creates an inter-dependency in the harvesting decisions since, on the margin, increasing the volume of timber harvested can depress current spot prices and this may make it optimal to delay harvesting timber on certain plots to avoid unduly depressing the current spot price of lumber.

that the new seedlings to be planted are identical genetically to those stems that came before them. As in the Faustman problem, an infinite sequence of harvests over an infinite horizon occurs. However, unlike in the Faustman problem, the optimal harvests in the stochastic version of the problem will occur at random intervals of time depending on the expected rate of change in the spot price of lumber and the condition of timber on the block.

Thus, the harvest decision at time j will depend on a vector \mathbf{x}_j of *state variables* that describe the state of a block, the price of lumber, and other macroeconomic variables useful in forecasting future lumber prices and harvesting costs. In our analysis, the vector \mathbf{x}_j will consist of just two variables (q_j, p_j) where q_j denotes the current volume of merchantable timber on the block, measured in cubic metres, and p_j denotes the current spot price of lumber.

Randomness in the evolution of the spot price of lumber is reflected in the transition probability $\tau_2(p_{j+1}|p_j)$. Next period's spot price p_{j+1} is random, but its probability distribution can depend on this period's spot price p_j .

3.1.3. Bellman's Equation

Let $V(\mathbf{x}_j)$ denote the expected present discounted value of the profits earned from optimally harvesting and selling timber on the tract. $V(\mathbf{x}_j)$ is the solution to the following *Bellman equation* of dynamic programming, where for notational simplicity we drop the j and $(j + 1)$ subscripts and let p denote a current period variable and p' denote its (random) value next period.

$$V(q, p) = \max \left[p\beta q - C + \delta \int V(0, p')\tau_2(p'|p), \delta \int V(q', p')\tau_1(q'|q)\tau_2(p'|p) \right]. \quad (3.1)$$

The transition $\tau_1(q'|q)$ denotes the mean growth dynamics on the harvest site. In the above Bellman equation, the current value of the block $V(q, p)$ is the maximum of two options: 1) to harvest or 2) to delay. If the block is harvested, then we assume that the volume harvested q is sold at the current spot price for lumber p resulting in revenue $p\beta q$ where β is the LRF. However, this revenue is reduced by C —which includes harvesting, transportation, milling, and replanting costs.

The first term inside the “max” operator in the Bellman equation is the current net profits from harvesting a block plus the expected discounted profits from future harvests. Since the harvest is assumed to reduce the effective volume of timber on the block to 0, the expected value function has q' equal 0, reflecting the fact that the harvest has occurred. If the harvest is delayed, then we assume that there are no revenues or costs associated with allowing the block of land to remain untouched another period, so the value of this option is simply the expected discounted value of profits from some future harvest (and sequence of subsequent harvests). This option has a q' not equal to 0, reflecting the fact that a harvest has not occurred.

The solution to the dynamic-programming problem partitions the state space into two regions: 1) a *continuation region* in which harvests do not occur, and 2) a *stopping region* in which it is optimal to harvest. The stopping region will be a subset of the space (q, p) in which the value of harvesting now exceeds the value of waiting to harvest later. We can describe general properties of the stopping region, but its precise characterization will depend on the numerical solution of the dynamic-programming problem (3.1). For example, the stopping region will generally have the property that if it is optimal to harvest at volume q it will also be optimal to harvest at all volumes \hat{q} greater than q . Similarly, when the stochastic process for lumber prices is sufficiently mean-reverting then, if it is optimal to harvest for a particular spot price p , it will also be optimal to harvest for all higher spot prices \hat{p} which exceed p .⁵

The only way to calculate detailed predictions of the optimal harvesting strategy is to solve the dynamic-programming problem numerically. In previous work, Rust (1996,1997) has developed efficient computational algorithms that make it feasible to solve dynamic-programming problems of the type outlined above. The key to an accurate solution to the problem is access to good data on the volume of timber on particular blocks as well as the spot price of lumber, so that we can estimate τ_1 and τ_2 , as well as good data concerning site-specific harvesting, transportation, manufacturing, and reforestation data to estimate C .

⁵ With mean reversion, the higher the current spot price, the greater the likelihood that the spot price in the next period will fall.

3.1.4. Modifying the Framework

Once the dynamic-programming problem has been solved, it can be used to determine the *sale value* of a particular block. The predicted value of a block will, in general, be $V(q, p)$, which equals the expected present discounted value of the stream of harvesting profits. We could generalize the dynamic-programming problem to consider the case where the block of land might have alternative uses that might have a potentially higher value, such as using the land for parks or for commercial, housing, or resort developments. In this case, two basic choices exist: 1) to sell outright the block to the highest bidder or to devote it to its highest-value alternative use, such as conversion to park land; or 2) to continue to use the land for harvesting timber. Of course, the land might be used for park/recreational activities in periods when timber is not being harvested, but in this analysis we have assumed there is no rental value for these uses. It would be easy to adapt the analysis above to include the case where there is a *shadow rental value* representing alternative use of the block in periods where harvesting does not occur.

We can introduce an additional state variable v_j representing the value of the best alternative use of the block of land. Accounting for the possibility of optimally converting or selling the block for its best alternative use, the Bellman equation would be modified to

$$V(v, q, p) = \max \left[v, R(q, p) + \delta \int V(v', 0, p') \tau_0(v'|v) \tau_2(p'|p), \right. \\ \left. r(q, p) + \delta \int V(v', q', p') \tau_0(v'|v) \tau_1(q'|q) \tau_2(p'|p) \right].$$

In this version of the Bellman equation, we let $R(q, p)$ denote the expected revenue from auctioning the right to harvest the block in the current period, $r(q, p)$ represents the shadow rental value of alternative uses of the block when harvesting does not occur, and v denotes the expected present discounted value of the best alternative use of the block. The transition probability $\tau_0(v'|v)$ represents the stochastic evolution of this best alternative value. Ordinarily, $V(v, q, p)$ will exceed v , in which case it is better to use the block for harvesting timber since the expected present value of the profits from continued use of the block for timber harvesting exceed its next best

use. However, if $V(v, q, p)$ equals v , then it will be optimal to sell or to convert the block to its next best use. We have assumed this sale or conversion is irreversible for simplicity, although it would be possible to consider cases where we allow the Crown to lease the block or to buy back the block after a previous sale.

Our framework is also quite rich and allows us to conduct a variety of different policy experiments that previous researchers could not. For example, as alluded to above, we can compare our estimates of the optimal harvesting policy with the harvests that have occurred during the past eight years, or so.

4. Geographic Information System

How will the rent map and, thus, the harvesting profile, of a particular region be calculated? The analytic device we chose to organize our data is a geographic information system (GIS). A GIS is a computer system capable of assembling, storing, manipulating, and displaying geographically-referenced information; *i.e.*, data identified according to their locations.

Our GIS contains several types of information. For example, first we have political maps which define the boundaries of the TSA in terms of planar coördinates. Second, we have maps in which those areas of the TSA that are on Crown land and which are not protected against harvest are listed; protected areas include federal and provincial parks as well as ecologically sensitive areas. Third, we have maps of elevations as well as maps of natural creeks, lakes, and rivers as well as man-made roads. Fourth, we have maps showing soil characteristics and vegetation types along with age distributions and stem densities. We exploit the different relations in this GIS to construct measures of different economic concepts.

To illustrate how we use the GIS, consider the following simple example: in Figure 8, we depict a stylized map of the political boundaries which determine an area available for harvest—in our case a TSA; we abstract from protected areas for presentational parsimony. Clearly, these can be introduced easily. In constructing a final map of costs, we shall impose regulations, such as the prohibition of harvesting near streams and watersheds.

In Figure 9, we present a digital elevation model (DEM) of the contours of

Figure 8
Political Map of Timber Supply Area



Figure 9
Digital Elevation Model

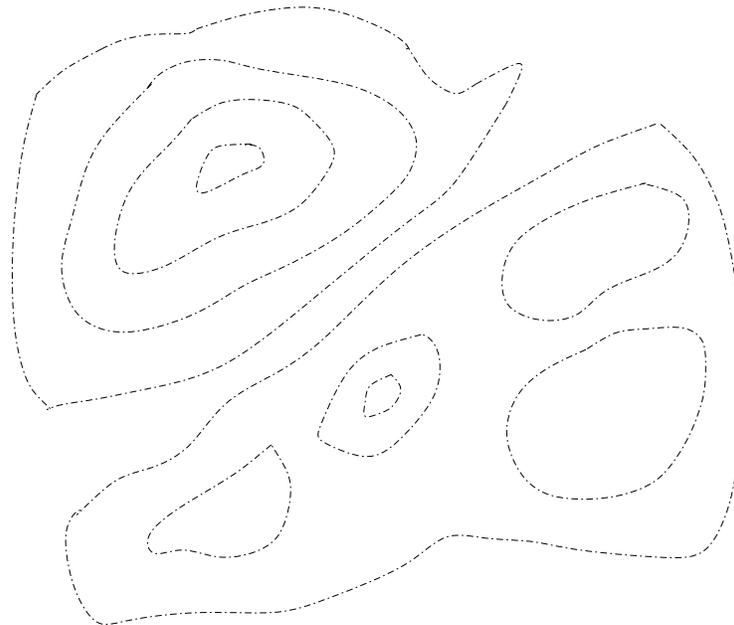


Figure 10
Map of Creeks, Lakes, and Rivers

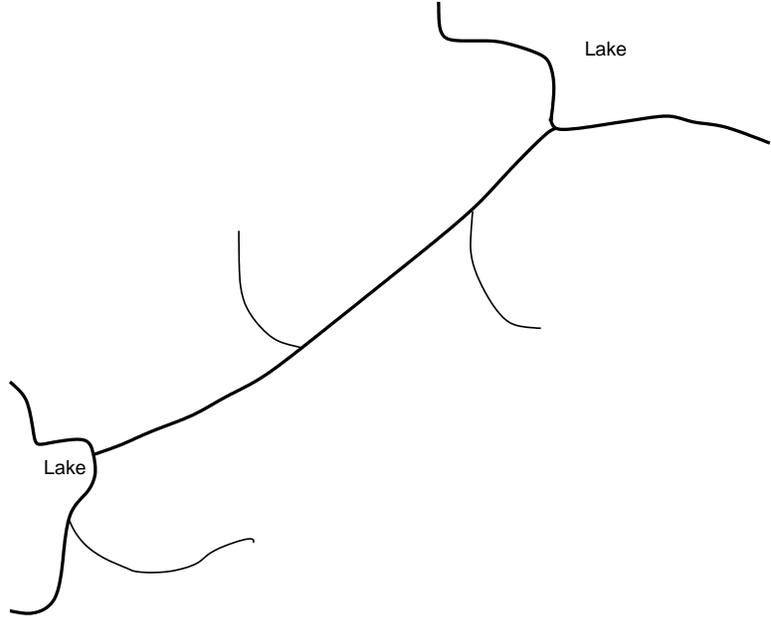
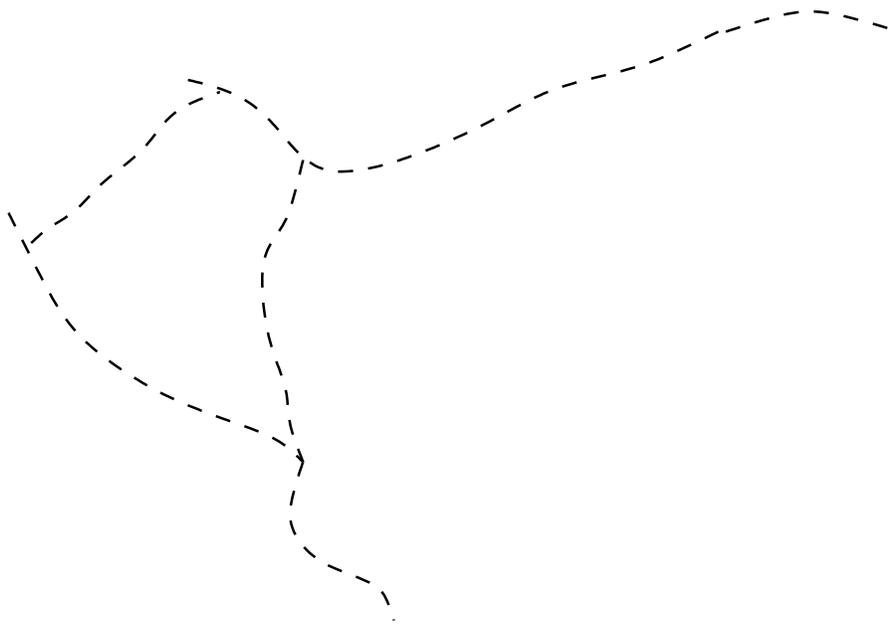


Figure 11
Map of Roads and Highways



elevation. This information will be important in estimating harvesting costs as these vary considerably by elevation and slope as well as transportation cost since steep grades are difficult to drive.

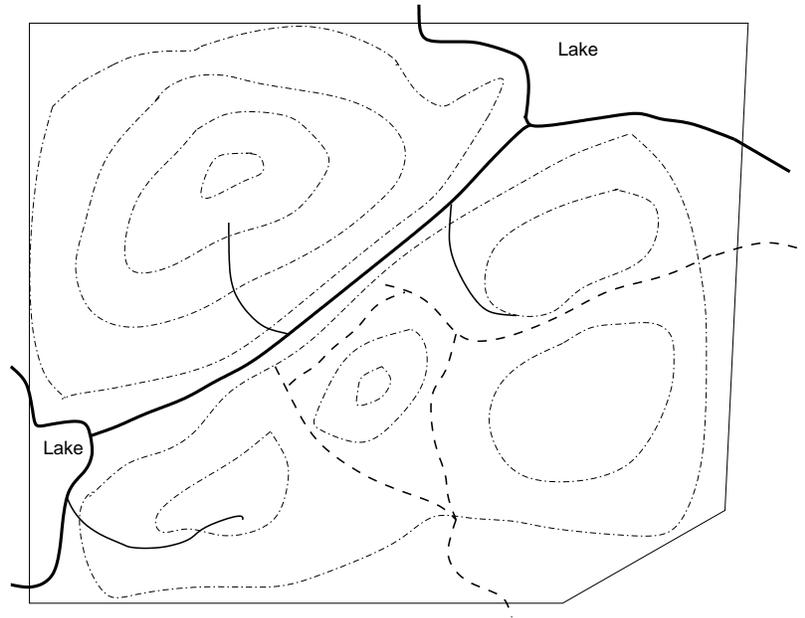
In Figure 10, a map of the creeks, lakes, and rivers is presented, while in Figure 11 a map of the roads and highways is presented. This information will be important in determining whether timber harvesting will affect watersheds and, thus, the ecology of the region as well as transportation costs.

By layering the maps in Figure 8 to 11 one atop the other, one can put together a physical map of the TSA which we depict in Figure 12. Using engineering information as well as harvesting regulations (such as prohibitions against harvesting near streams), one can calculate for each raster of the map in Figure 12 an estimate of harvesting costs as well as one of transportation costs. The estimates are derived from the topography of the land as well as site-specific information contained in the GIS relations. In Figure 13, the number in each square (raster) represents some measure concerning how much each cubic metre of timber will cost to harvest and to transport to a sawmill. When we implement this framework, we assume that the cost of milling lumber is the same for all timber and add this to the site-specific harvesting and transportation costs. Reforestation costs depend on a stem density and we add those, too. Higher numbers represent higher costs, while the “**x**”s represent rasters which cannot be harvested under any circumstance—*e.g.*, because of harvest regulations.

5. Modelling Growth and Yield in Stands of Timber

From a capital-theoretic perspective, perhaps the most important feature of this controlled stochastic, dynamic decision problem is the growth and yield of merchantable timber from the forest. Basically, two types of forests exist: old-growth and newly-planted. Typically, old-growth forests, and even some second-growth forests which have regenerated naturally, are heterogeneous in species and age as well as stem density. Such heterogeneity is difficult to model. Foresters have used a variety of different methods to estimate the growth and yield of uneven-aged forests; these have been summarized by Peng (2000). On the other hand, newly-planted forests are typically

Figure 12
Physical Map



Legend

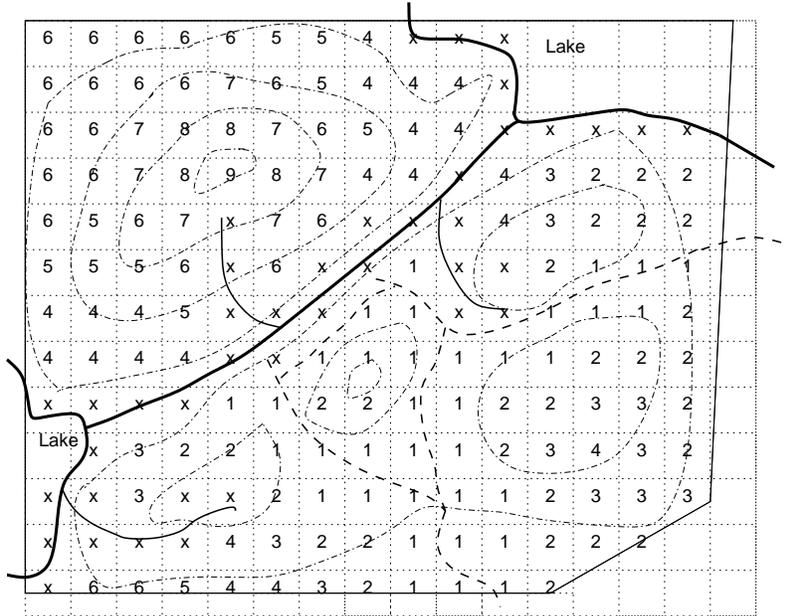
- | | | | |
|-----------|----------------------|-----------|--------------------|
| ————— | Creek | ————— | Political Boundary |
| ————— | River, Lake Boundary | - - - - - | Elevation Contour |
| - - - - - | Road, Highway | | |

homogeneous with respect to species as well as stem age and density. Modelling these forests is straightforward, relatively speaking, of course.

5.1. Predicting Growth and Yield in Planted Forests: *TASS*

The Ministry of Forests and Range in British Columbia has devoted considerable time and resources to investigating the growth of seedlings of the same species planted according to a particular stem density on sites having different productivity. Foresters in the Research Branch have a computer programme, *TASS* (Tree and Stand Simulator), that can be used to simulate the growth and yield of a particular species, planted according to a pre-specified stem density, on a block of a particular site index, produc-

Figure 13
Rasters of Costs



tivity.⁶ To run *TASS*, one must provide a species (or species) as well as a site index and an initial stem density. Based on experiments done by foresters, *TASS* will simulate a future forest, and then estimate the volume of merchantable timber at any age. We used *TASS* to simulate a variety of different forest yields for different combinations of single species as well as site indices and stem densities. In particular, we simulated 100 forests over a 150-year life-span for the following 64 combination of species, site index, and stem density: (Fir,Spruce) \times (10, 15, 20, . . . , 40, 45) \times (1200, 1400, 1600, 1800).⁷ We then used the output to estimate $\tau_1(q_{j+1}|q_j)$ for different species as well as different site indices and stem densities.

⁶ A site index is a summary statistic concerning the productivity of a particular block. It is derived by foresters during surveys. Our site index is the diameter at breast height of a twenty-five year-old stem. Historically, this measure has correlated quite well with the height and, consequently, the volume of a stem. For a particular stem density, one can then estimate relatively accurately the volume of merchantable timber on an hectare of land.

⁷ We are grateful to Ken Polsson of the Ministry of Forests and Range, Research Branch, for running the simulations with *TASS* and providing us the output.

5.2. Predicting Growth and Yield in Old-Growth Forests: *VDYP*

VDYP (Variable Density Yield Prediction) is a computer programme designed to implement a prediction system to estimate average yields and to provide forest inventory updates over large areas. It is intended for use in unmanaged natural stands of pure or mixed species composition. Basically, field observations on the yields from forests of different species compositions and ages under different site indices were used to develop a regression model, the output of which is an estimated merchantable volume of timber at some point in the future. Because natural forests can be quite heterogeneous, *VDYP* is not as “accurate” as *TIPSY* (Table Interpolation Program for Stand Yields), which uses interpolation methods and thousands of *TASS* runs to estimate timber yields. Of course, the prediction problem under the conditions assumed for *VDYP* is much more difficult than the one for *TASS* and *TIPSY*, so the comparison is somewhat unfair. We used the output of *VDYP* to estimate $\tau_1(q_{j+1}|q_j)$ for different species compositions and ages as well as different site indices and stem densities.

6. Implementation

We applied the framework discussed in the previous sections to the Fraser Timber Supply Area (FTSA) in British Columbia. In Figure 14, we depict the location of the FTSA in the province, the southwest corner of the mainland. The FTSA is sometimes referred to as the Chilliwack Forest District because its boundaries are nearly coincident with those of that Forest District whose district office is located in Chilliwack, a town about an hour by automobile from Vancouver.

6.1. Some Relevant Features of the Data Set

Below, we describe briefly the main features of our data set, while in an appendix to the paper we describe in detail the mechanics of how we built the data set.

The Chilliwack Forest District and, thus, the FTSA are about 1.4 million hectares in area, around 5,400 square miles.⁸ Not all of this land, however, is under the

⁸ An hectare is 100 metres square or 10,000 square metres or, approximately, 2.4711 acres.

Figure 14
British Columbia

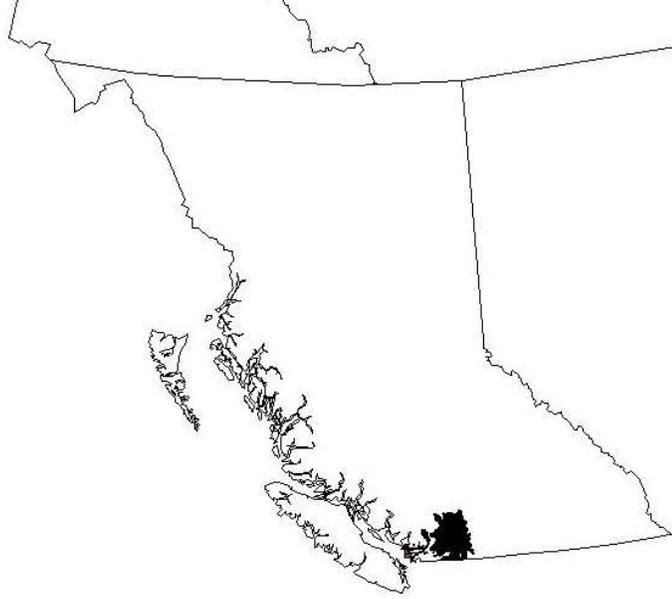
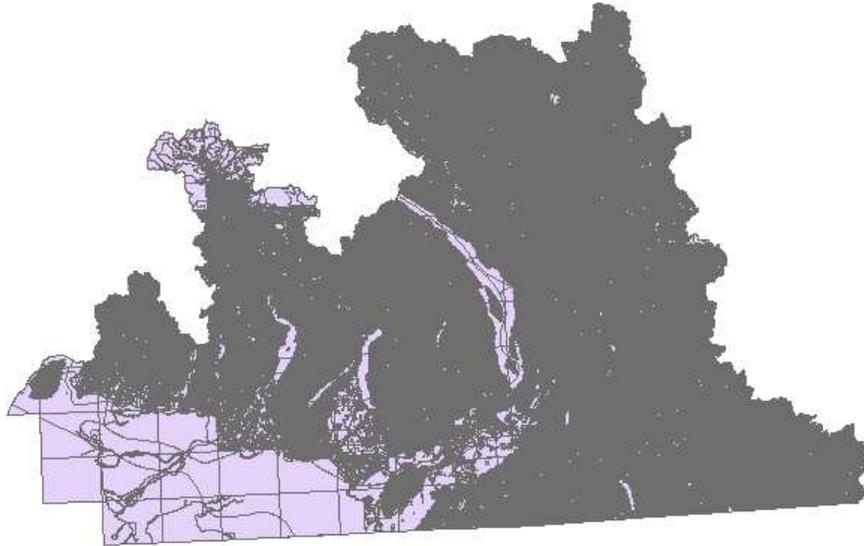


Figure 15
Fraser Timber Supply Area



jurisdiction of the Ministry of Forest and Range. In Figure 15, we depict all Crown land, the darkly-shaded areas. The lightly-shaded areas are bodies of water, inhabited areas, or private land. Having imposed this screen left us an area of about 706,603 hectares, which we represented as grid squares, hectares of land.

On these sites, we estimated harvesting costs based on which harvesting technology could be used. For relatively flat sites, which we defined as a slope less than 75 percent, standard yarding technology can be used, while on very steep sites at high elevations, only helicopter logging can be pursued.⁹ We deduced the topography of the land from a DEM of the FTSA.

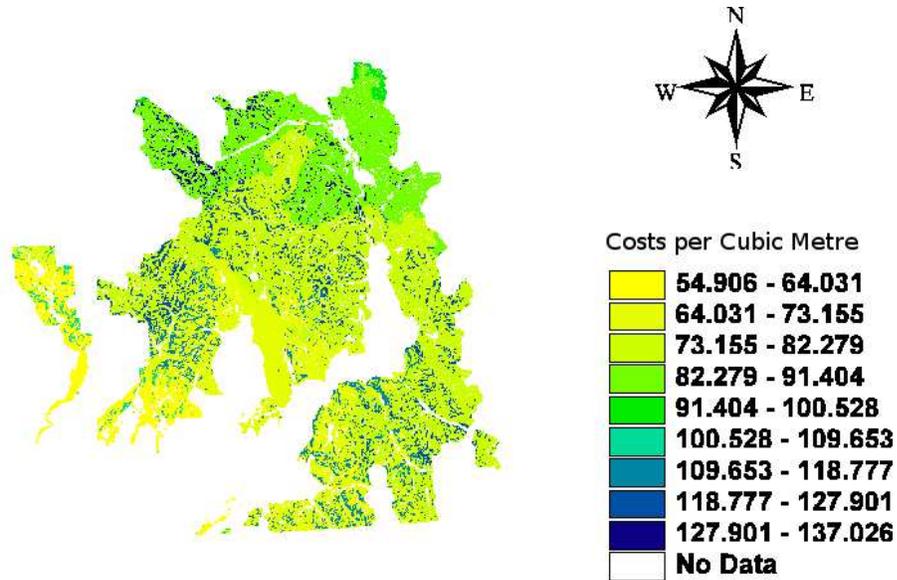
We then used the *Coast Appraisal Manual*, which is published by the Revenue Branch of the Ministry of Forests and Range and which was used previously to determine stumpage rates throughout the coastal region of British Columbia, to estimate site-specific harvesting costs as well as manufacturing costs. We also used the extant road network in the FTSA, and contained in our GIS, to estimate the distance to the nearest sawmill. In the FTSA, several sawmills exist, but the bulk of these are located near Chilliwack. We chose the centre of gravity of these mills, weighting each mill by the volume of lumber it could produce, as the destination of all timber.¹⁰ From these distance estimates, we then formed estimates of transportation costs. Thus, for each site, we have both harvesting-, transportation-, and manufacturing-cost estimates. In Figure 16, we depict our map of costs using different colours or, sometimes in printed copy, different shades of gray.

Our next task was to determine growth and volume for each block. On each block, we were able to obtain some 134 pieces of biological information—*e.g.*, the species composition, a site index, and so forth. However, not all of the grids were suitable for harvesting. In fact, some 113,667 had no species codes for trees, presumably because these were comprised mostly of rock. For another 4,937 hectares, the site indices were incredibly low, so that no merchantable timber was predicted to grow, while another 4,750 hectares had no volumes for other reasons, and 2,173 hectares had missing

⁹ A slope of 100 percent is defined to be a 45° pitch, so 75 percent would be a pitch of about 33.3°.

¹⁰ We are grateful to Steven Fletcher, of the Revenue Branch, for providing the GIS coördinates as well as the volumes of each sawmill in the FTSA.

Figure 16
Estimated Cost Map



parameters for our growth and volume programmes, *VDYP* and *TASS*. In the end, we were left with 581,076 viable hectare blocks; these made up our analysis unit. The reader should note that this area is substantially larger than the 206,910 hectares reported by Larry Pedersen (at the time the chief forester of British Columbia) in the December 2003 Fraser Timber Analysis.¹¹ Presumably, our larger area obtained because we did not constrain ourselves by the rules contained in the *Forest Practices Code* of British Columbia.¹²

6.2. Computational Issues

Solving nearly 600,000 stochastic dynamic programmes is extremely time consuming, even under the best of circumstances. We divided our computations into two parts.

¹¹ *Fraser Timber Supply Area Analysis Report*. Victoria, Canada: British Columbia Ministry of Forests, Forest Analysis Branch, 2003.

¹² *Forest Practices Code of British Columbia Act* of 2002. Victoria, Canada: Queen's Printer, 2002.

In the first part, we solved for the optimal value functions assuming that a particular block had just been harvested. We partitioned the sum of harvesting, transportation, and manufacturing costs into intervals of \$5.00 CAD per cubic metre from a minimum of \$55.00 to a maximum of \$140.00—*i.e.*, eighteen different cost regions. Reforestation cost were included as a fixed cost of \$600 per hectare: we assumed a stem density of 1,200 per hectare with seedlings costing \$0.25 a piece to grow and another \$0.25 a piece to plant. For each of these different regions, we entertained 64 different combinations of replanted sites. Thus, in total, we solved (18×64) , or 1,152 continuation stochastic dynamic programmes. Given that our software can solve a stochastic dynamic programme in under three seconds by discrete policy iteration, the 1,152 optimal continuation value functions took just under five hours to solve on a desktop computer.

Solving for the optimal policy functions for the unmanaged natural stands is much more time-consuming than in the newly-planted case. Basically, in the FTSA, for our 581,076 available harvest blocks, there are 46,360 different covariate combinations—*i.e.*, combinations of site indices, species compositions, age compositions, and so forth. Given our computational technology, this could involve a thousand hours of computer time to solve them all. We adopted an alternative strategy. For each block i , we approximated the volume profiles generated by *VDYP* by the following three-parameter function:

$$q_{j+1} = a_0^i + a_1^i q_j + a_2^i q_j^2 + U_i.$$

We estimated a_0^i , a_1^i and a_2^i by the method of least squares for all 706,603 blocks. The majority of the R^2 for these models was above 0.99. Those that were not were flagged as anomalies. The anomalies were then examined and dealt with separately. Some 113,667 had no species codes for trees, presumably because these were comprised mostly of rock. For another 4,937 hectares, the site indices were incredibly low, so that no merchantable timber was predicted to grow, while another 4,750 hectares had no volumes for other reasons, and 2,173 hectares had missing parameters for our growth and volume programmes, *VDYP* and *TASS*. In the end, 581,076 blocks could be dimension-reduced in the above way. We then took the triplets of estimated parameters $(\hat{a}_0^i, \hat{a}_1^i, \hat{a}_2^i)$ and used cluster analysis to assign them to particular sets.

We chose 64 sets. Again, we partitioned the sum of harvesting, transportation, and manufacturing costs into eighteen different cost regions. For each of these different regions, we entertained the 64 different triplets, and solved another 1,152 stochastic dynamic programmes, where we conditioned on the appropriate continuation optimal value function discussed above.

6.3. Some Illustrative Results

In Figure 17, we present the optimal harvesting rule for a representative block, along with a sample path for the lumber price and the timber volume in (q, p) -space. As one might expect, it is very difficult to describe the outcome on each and every one of the 581,076 harvest blocks, in each month for, say, the next century. Even presenting a snapshot at a point in time, as we did for the static costs depicted in Figure 16, may not be particularly illuminating to some. Because our results are dynamic, we have chosen to present the aggregate output from solving thousands of different stochastic dynamic programmes against the Annual Allowable Cut (AAC) for the FTSA as well as the actual volume harvested for the last five years. This is depicted in Figure 18. Our solution is quite different from what has transpired during the past eight years, or so.

What are the sources of these differences? An autopsy of our micro-level solutions reveals that the harvesting decisions are clumped together: to wit, when the price of lumber crosses a certain threshold, then a majority of the sites are harvested at once—timber cycles obtain.

7. Summary and Conclusions

In this paper, we have constructed an intertemporal model of rent-maximizing behaviour on the part of a timber harvester under potentially multi-dimensional risk as well as geographical heterogeneity. Subsequently, we have used the method of dynamic programming to characterize the optimal policy function, the rent-maximizing timber-harvesting profile. We then applied our theoretical framework to analyze unique and detailed information from the databases of the British Columbia Ministry of Forest and Range concerning the Fraser TSA. We have organized these data in the

Figure 17
Optimal Decision Rule

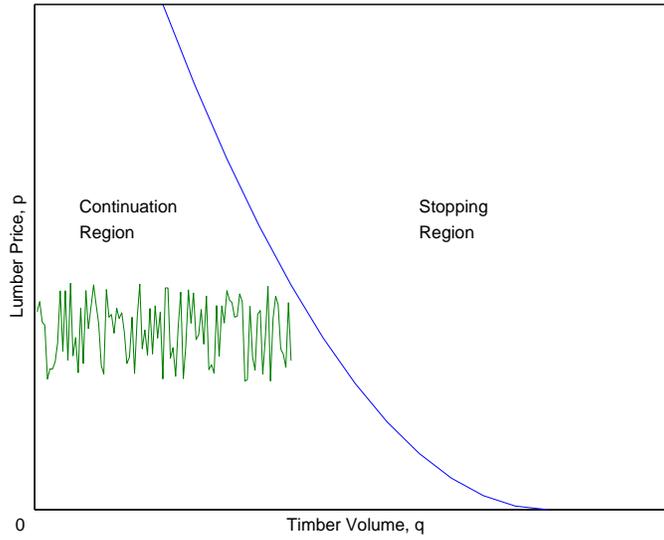
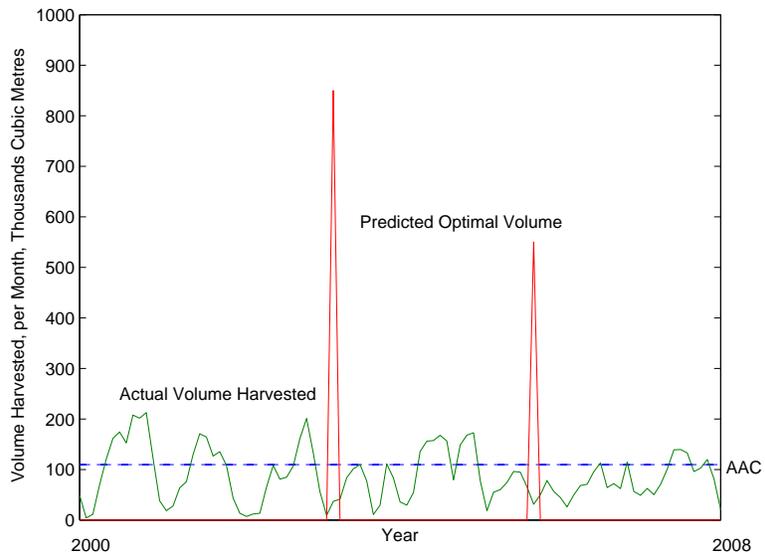


Figure 18
Optimal and Actual Volumes Harvested



form of a dynamic geographical information system to account for site-specific cost heterogeneity in harvesting and transportation as well as uneven-aged stand dynamics in timber growth and yield across space and time in the presence of stochastic lumber prices.

Our model is a powerful tool with which to conduct policy analysis for a number of reasons. First, we take geography seriously, both in the planar sense and in the three-dimensional sense. Second, we take site-specific heterogeneity seriously both on the cost side in terms of harvesting and transportation and on the growth and yield side in terms of heterogeneous stands of timber. Third, we model initial conditions. In particular, we do not take as the starting point a steady-state allocation, or even an optimal allocation. Instead, we take the existing uneven-aged timber stand as given and derive the optimal policy function—the optimal timber-harvesting profile. Fourth, we use best-practice biological methods to model the dynamics of uneven-aged forest growth and yield. Fifth, in the past economists have typically demonstrated their methods by solving simple examples in closed-form or they have imposed conditions sufficient to sign comparative static predictions. We have harnessed recent developments in computational methods to solve numerically for the optimal policy function.

Our framework is quite rich and allows us to conduct a variety of different policy experiments that previous researchers could not. For example, we have compared our optimal harvesting policy with the harvests that have occurred during the past eight years, or so. Finally, in other related work—*viz.*, Paarsch and Rust (2008)—we have investigated how the the province of British Columbia might behave as a “big player” in the lumber market.

A. Appendix

In this appendix, we summarize the geographic work we undertook to estimate harvesting, transportation, and manufacturing costs as well as timber growth and yields for the Fraser TSA of British Columbia, Canada. Our calculations required detailed geographic data because, for example, mountainous terrain led to significant spatial variation in timber growth as well as harvesting costs, while the clustering of sawmills in specific locations and the relative sparseness of suitable roadways influenced transportation costs. Such factors suggest that any economic analyses should be for sites of small area (*e.g.*, by the hectare which is an area 100 by 100 metres or 10,000 square metres). Also, a geographic analysis is essential to provide the site-specific data needed. Consequently, a geographic information system (GIS) becomes very useful when analyzing, processing, storing, and presenting information at the scale necessary to support the economic analyses.

A.1. GIS Software

The GIS software used was the ESRI suite of programmes. This included the following:

- 1) *ArcInfo Workstation* – analysis using command line interface;
- 2) *ArcView GIS 3.3* – analysis using *Windows*-based interface;
- 3) *ArcMap 8.0* – analysis using *Windows* based interface;
- 4) *ArcCatalog* – file management functions;
- 5) *ArcToolbox* – collection of subroutines.

ArcInfo is the oldest of the listed programmes and, generally, the most powerful; *i.e.*, it has the most functionality and allows geographic analyses that cannot be performed using the other programmes. It, however, requires an understanding of the function and format of numerous commands.

ArcView and *ArcMap* perform geographic analyses in a *Windows* environment. *ArcView* is an earlier implementation, but both programmes are still used and supported. *ArcMap* can be more complex, so at times *ArcView* is a quicker and simpler

way to complete some tasks. Both programmes are menu driven and more intuitive than *ArcInfo*.

ArcCatalog is a specialized interface used to manage files. Numerous files can be created during a geographic analysis, often without the user's knowing that a specific file was created. Thus, copying, moving, and deleting files can be tricky without an "intelligent" file manager that recognizes the whole range of geographic data files and their association with other files.

ArcToolbox is a collection of subroutines from within *ArcInfo*. It uses a simpler and *Windows*-based interface to perform specific tasks. Some of this functionality exists within *ArcView* and *ArcMap*.

In general, GIS software has a specific vocabulary used to describe geographic data. Some of the terms in this vocabulary are discussed here; a glossary is presented in Table A.1. Geographic data are usually represented by *areas*, *lines*, or *points*. Areas can be represented by *polygons* or *grid cells*. A polygon is simply a set of lines that encompasses an area and can be of virtually any shape or size as long as the area is completely bounded. On the other hand, grid cells are typically squares of a common dimension and reference point. Within the ESRI system, these types of geographic data can be represented in *coverage*, *shapefile*, or *raster* formats. Coverages and shapefiles apply to polygons, lines, and points; rasters are another name for grid-cell formats. Coverages and shapefiles differ in the amount of information available concerning neighbouring features (topology). Coverages are more complex and include more information concerning neighbouring features than do shapefiles.

A.2. Data Sources

Our data were supplied by the Ministry of Forests and Range of the province of British Columbia, Canada. Data files were provided in ESRI exchange format (E00) and ported to a local workstation via CD.

A.2.1. Method

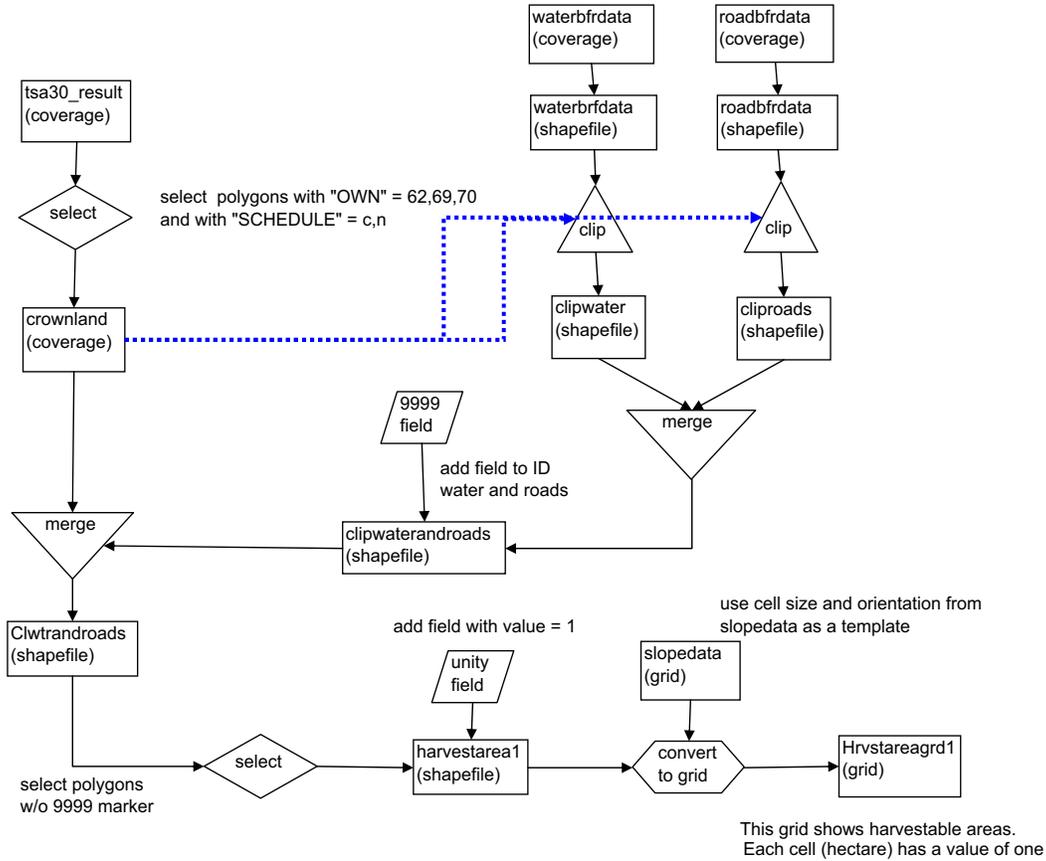
For this research, the geographical analysis can be broadly described by the following:

Table A.1
Glossary

Term	Definition
<i>arc</i>	a line or curve within a GIS coverage or shapefile.
<i>coverage</i>	a GIS format that contains information about topology; <i>i.e.</i> , the explicit spatial relationship among data.
<i>GIS</i>	geographic information system, a computer-based mapping technology.
<i>grid</i>	a GIS format based on square cells, with one unique value for each cell.
<i>hectare</i>	an area represented by a square having sides of 100 metres.
<i>join</i>	a specific GIS technique to associate data from different files or data bases based on a common variable.
<i>point</i>	a position in space, applies to coverages and shapefiles.
<i>polygon</i>	an area represented by a continuous closed boundary, applies to coverages and shapefiles.
<i>raster</i>	another term for grid.
<i>shapefile</i>	a GIS format that contains does not contain specific information about topology, spatial relationships are visible but not explicitly coded within the data.
<i>slope break point</i>	the value (percent) that distinguishes between conventional logging and helicopter logging methods.
<i>timber mark</i>	an identifying number for timber harvested at a particular time and place.
TASS	Timber and Stand Simulator. It simulates the yields of highly-managed stands under pest free conditions. It is based on observed growth trends or research plots and represents the potential of a specific site, species, and management routine. For more information, go to http://www.for.gov.bc.ca/hre/gymodels/tass/
TIPSY	Table Interpolation Program for Stand Yields. It predicts the yields of highly managed stands under pest free conditions. It is based on observed growth trends or research plots and represents the potential of a specific site, species, and management routine. For more information, go to http://www.for.gov.bc.ca/hre/gymodels/tipsy/
VDYP	Variable Density Yield Prediction programme. It predicts the average yield of naturally regenerated forests. For more information see http://www.for.gov.bc.ca/hre/gymodels/vdyp/

- 1) define study area (Figures 14 and 15 in paper);
- 2) estimate harvesting, transportation, and manufacturing costs (Figure 16 in paper);
- 3) import vegetation data;
- 4) create point coverage of vegetation and cost data;

Figure A.1
Flow Chart of GIS Definition



5) export data in ASCII format.

The study area was defined preliminarily as the Chilliwack Forest District, the Fraser TSA, with geographic data (`tsa30_result`) supplied by the Ministry of Forests and Range. The Fraser TSA is located in the far southwestern portion of British Columbia and borders the United States on the south and the Strait of Georgia on the west, which connects with the Pacific Ocean. The Fraser TSA also includes the city of Vancouver. In Figure A.1, we present a flow chart in which refinements of `tsa30_result` are described.

The study area was refined by identifying lands within the Fraser TSA that are owned by the provincial government; *i.e.*, Crown lands. This was done via a selection process where the data file was queried for records that matched the criteria of:

- OWN equals 62 or 69 or 70 and i SCHEDULE = c or n.

Thus, from an ownership perspective, Crown lands became areas where harvesting is feasible. From within this set, it was necessary to identify areas where harvesting is infeasible for technical, cultural, or aesthetic reasons. For our research, we excluded areas near roads and streams. Starting with line coverages of roads and water bodies, we applied a 100 metre buffer around each feature, added an identifying code (9999) and merged these files with the file representing Crown lands. The merged file was then queried for all areas without the 9999 marker, with these areas representing land where harvesting is permitted. This information was saved in a raster format to produce a grid showing permissible harvesting areas. Each grid cell measures 100 metres by 100 metres, one hectare.

In the 2002 *Coast Appraisal Manual*, which is published by the Revenue Branch of Ministry of Forests (at the time), two primary factors affect harvesting and transportation costs: the slope of the harvest block and the distance from the harvest block to the sawmill. The Ministry of Forests and Range used the following expression to estimate cost per cubic metre:

$$\text{Slope Factor} + 0.000153 \times \text{Distance} + 32.05.$$

Here, the “Slope Factor” was either 21.55(\$/m³) or 66.60(\$/m³), depending on whether conventional or helicopter logging was used. The smaller number was applied when a slope was below 75 percent, while the larger number was applied to slopes above 75 percent. Note that a slope of 100 percent represents a 45° angle from horizontal. A break point of 75 percent was chosen based on conversations with Ministry of Forests and Range. It was also compared with historical cost data by examining 2,172 records from the Fraser TSA and computing the error between the historical and predicted costs at the same location. Increasing the break-point slope resulted a slight decrease in the standard deviation of the error between the two data sets (*i.e.*, 24 rather than 22), but it was not considered significant enough to change the base break-point value. Historical costs were adjusted by the Canadian Consumer Price Index (CPI) to be comparable to predicted costs. The land’s slope was derived from a digital elevation model (DEM) and grid files of both slope and elevation were provided by Ministry of Forests and Range. The slope grid file was trimmed to the

extents of the permissible harvesting areas to minimize data processing and storage requirements.

Computing the distance between each hectare of harvestable land and a sawmill involved three steps. First, we had to identify an appropriate point to represent a sawmill receiving wood for each hectare of forest land. Next, we had to compute the road distance from the sawmill location to all points on the road network. Finally, we had to compute the off-road distance from the hectare of forest to the nearest road plus the on-road distance from that point to the sawmill location.

We chose to represent all sawmills at one central location. This simplified calculating transportation costs because we did not have to assign timber to specific sawmills. We computed the geographic centre of gravity for all sawmills within the Fraser TSA and used this as a reference point for calculating distances. We accounted for production volumes in this procedure, so that larger mills were weighted more heavily than smaller ones and, thus, had more influence on the location of the reference point. Most of the sawmills in the Fraser TSA are clustered in the Chilliwack area, while areas to be harvested are primarily located to the north and northeast of Chilliwack. Thus, errors associated with representing many sawmills by one central or average location is less severe than it might be otherwise.

Computing distances required a data file of the road network; this was provided by Ministry of Forests and Range. This network was converted to a grid and then the grid cells were counted from the sawmill point to each other point on the network. This provided an on-road distance to each part of the forest within the study area. As part of this process, it was necessary to adjust the sawmill reference point so that it was actually on the road network. This was a minor adjustment of only a couple of hundred metres, at most, and was accomplished using a shift function.

The road network grid was also used as the source grid for calculating off-road distances between each forest hectare and its nearest road. The *ArcInfo* function `EUCALLOCATION` was used to compute total distance from each forest hectare to the central sawmill location. `EUCALLOCATION` produces two output grids, `OUT_ALLOC` and `OUT_DIST`. The first of these contained the value of the nearest road cell (on-road distance), and the second output grid contained the distance from each forest hectare

Table A.2
VDYP Input Variables

Variable	Label	Column	Type	Row
Forest Inventory Zone	FIZ	65	character(1)	28
Volume Adjustment Factors	adj_volume_factor	75	real*8	38
Stand Crown Closure	crown_closure	76	integer	39
Stocking Class	class_cd	78	real*8	41
Adjusted Site Index	site_index	79	real*8	42
Species 1	sp1	87	character(4)	50
Species 2	sp2	88	character(4)	51
Species 3	sp3	89	character(4)	52
Species 4	sp4	90	character(4)	53
Species 5	sp5	91	character(4)	54
Species 6	sp6	92	character(4)	55
Percentage Species 1	pct1	93	integer	56
Percentage Species 2	pct2	94	integer	57
Percentage Species 3	pct3	95	integer	58
Percentage Species 4	pct4	96	integer	59
Percentage Species 5	pct5	97	integer	60
Percentage Species 6	pct6	98	integer	61
Projected Adjusted Age	proj_adj_age	102	integer	65
Projected Adjusted Height	proj_adj_height	106	integer	69

to the nearest road. Summing these two output grids produced a new grid containing the total distance from each forest hectare to the central sawmill location. This new grid was converted to a point format for later joining with timber-growth parameters.

The next part of the analysis was to import data concerning vegetation. These data served as inputs to the timber-growth programmes *VDYP* and *TASS*. Vegetation data were also supplied by Ministry of Forests and Range. This file consisted of over 168,000 polygons in a coverage format, with each polygon having of 134 parameters. Of these, 20 parameters were required for the timber-growth programmes. These parameters are listed in Table A.2, along with their column number in the original data file, their variable type, and their position in a data equivalency table.

The above data were joined to a raster representation of harvestable areas within the study area using the original polygon ID numbers as references. The resulting file was converted to a point coverage in preparation for combining vegetation data with cost data.

Once cost and vegetation data were available in point coverages, they were combined into a single file in which they were geo-referenced to each other. Joining these files required an exact match, a one-to-one mapping, between the point coverages. However, the coverages contained some mismatches at the edges of the study

area where data existed for vegetation data (`vdypnt1`), but not for harvest costs (`costrnpnt1`). Edge effects are common in these sorts of analyses because the data are from two different sources and a number of data transformations have been undertaken. These edge errors were identified visually and removed from the data. The function `JOINITEM` in *ArcInfo* was then used to combine cost and vegetation data into a single file for later export.

The combined data set (`polyidpnt1`) was exported in ASCII format, having file-name `siteind4b.asc`. In addition to the 20 timber-growth variables discussed above, four other variables were included in the exported file. They included an identification code, an x -coördinate, and a y -coördinate as well as the total of harvesting-, transportation-, and manufacturing-costs.

A.2.2. Output

The file `siteind4b.asc` contains 706,603 records. Each record represents one hectare of potentially harvestable forest area within our study area and has the 24 parameters listed above associated with it.

Harvesting and transportation costs are calculated from a geographic analysis of the study area at a scale of one hectare. The results of the geographic analysis are discussed below. Timber-growth information was not created in this work, but rather transferred into a form appropriate for analysis.

Harvesting and transportation costs for the study area are depicted in Figure 16. Calculated values range from \$55 to \$137 CAD per cubic metre. In general, costs increase from the southwest to the northeast sections of the study area. This is expected and in keeping with the clustering of sawmill locations in the Chilliwack area which is in the southwest part of the study area. Superimposed over this general trend are local variations in costs due to changes in the slope of the land and due to the remoteness of some sections from the local road network.

From the Revenue Branch of the Ministry of Forests (at the time), we also obtained historical data concerning over 2,172 timber marks that were recently harvested in the Fraser TSA.¹³ Among other things, these historical data contained

¹³ A timber mark is like a tax identification number. It maps the timber harvested from a

information concerning harvesting, transportation, and manufacturing costs as well as fixed costs associated with such things as road construction. By and large, our estimated harvesting, transportation, and manufacturing costs correlate quite highly with the historical data. Fixed costs, however, were impossible to predict in our model. In addition, these fixed costs were sometimes high. We chose to deal with this by subtracting the average of the fixed costs of our sample from the total costs for each site. In essences, we assume that a harvester gets a draw from an urn of fixed costs for each block.

A.3. Stochastic Process of Lumber Prices

Our data set concerning lumber prices contains 250 monthly time-series observations from January 1979 to October 1999. We have the real, average-monthly price per thousand board feet (1MBF) for one box car of Western, Kiln-Dried (KD), Spruce-Pine-Fir (SPF), 2x4s, Number 2 and Better (No.2&Btr), Random Lengths (R/L). The series was constructed from weekly reports in the trade publication *Madison's Canadian Lumber Reporter*, weekly issues, January 1979 to October 1999. This price series is listed as “less 5&2 percent” discounts, and is free-on-board (FOB) mill. Moreover, it is quoted in nominal U.S. dollars. To convert the series into Canadian dollars we used the Canadian/U.S. spot exchange rate from the CANSIM database, matrix 933, series B40001.¹⁴ We converted nominal price data into real terms by dividing by the Canadian Consumer Price Index (CPI) setting January 2002 to one.

We specified $\tau_2(p_{j+1}|p_j)$, the the stochastic process of p_j , according to the following continuous-state, Markov process:

$$\log p_{j+1} = \rho_0 + \rho_1 \log p_j + \sigma \varepsilon_{j+1}$$

particular geographical area into a stumpage rate that must be paid to the Crown for timber harvested. Historically, a timber mark was made with a hammer on the end of a log. Today, it is simply a name for tax record.

¹⁴ Statistics Canada information is used with the permission of Statistics Canada. Users are forbidden to copy the data and disseminate them, in an original or modified form, for commercial purposes, without the expressed permission of Statistics Canada. Information on the availability of a wide range of data from Statistics Canada can be obtained from Statistics Canada's regional offices, from its website at <http://www.statcan.ca> or from its toll-free number 1 (800) 263-1136.

where ε_j is an independent and identically-distributed Gaussian error term having mean zero and variance one. For numerical stability, when solving the dynamic-programming model, we re-scaled the units of prices from dollars per thousand board feet to dollars per board feet. We also imposed reflecting lower and upper bounds on the level of the price process, denoted \underline{p} and \bar{p} , respectively. These were 0.100 and 0.750, respectively. To wit, the lowest price imaginable is \$100 CAD per thousand board feet, in 2002 dollars, and the highest is \$750 CAD per thousand board feet. Our estimates of ρ_0 , ρ_1 , and σ are -0.052412 , 0.945648 , and 0.085376 , respectively.

B. Bibliography

- Brazee, R. and R. Mendelsohn, "Timber Harvesting with Fluctuating Prices." *Forest Science*, 34 (1988), 359–372.
- Coast Appraisal Manual*. Victoria, Canada: British Columbia Ministry of Forests, Revenue Branch, 2002.
- Faustmann, M., "Berechnung des Werthes, welchen Waldboden sowie nach nicht haubare Holzbestände für die Weldwirtschaft besitzen." *Allgemeine Forst und Jagd Zeitung*, 25 (1849), 441.
- Forest Practices Code of British Columbia of 2002*. Victoria, Canada: Queen's Printer, 2002.
- Fraser Timber Supply Area Analysis Report*. Victoria, Canada: British Columbia Ministry of Forests, Forest Analysis Branch, 2003.
- Haight, R. and T. Holmes, "Stochastic Price Models and Optimal Tree Cutting: Results for Loblolly Pine." *Natural Resource Modelling*, 5 (1991), 423–443.
- Kaya, I. and J. Buongiorno, "Economic Harvesting of Uneven-Aged Northern Hardwood Stands under Risk: A Markovian Decision Model." *Forest Science*, 33 (1987), 889–907.
- Madison's Canadian Lumber Reporter*. 1979–1999. Madison's Publications: Vancouver, Canada.
- Morck, R., E. Schwartz, and D. Strangeland, "The Valuation of Forestry Resources under Stochastic Prices and Inventories." *Journal of Financial and Quantitative Analysis*, 24 (1989), 473–487.
- Paarsch, H. and J. Rust, "Timber Cycles." Typescript. Iowa City, Iowa: Department of Economics, University of Iowa, 2008.
- Peng, Changhui, "Growth and Yield Models for Uneven-Aged Stands: Past, Present and Future." *Forest Ecology and Management*, 132 (2000), 259–279.
- Provencher, B., "An Investigation of the Harvest Decision of Timber Firms in the South-East United States." *Journal of Applied Econometrics*, 10 (1995), S57–S74.
- Reed, W., "The Decision to Conserve or Harvest Old Growth Forest." *Ecological Economics*, 8 (1993), 45–69.

- Reed, W. and H. Clarke, "Harvest Decisions and Asset Valuation for Biological Resources Exhibiting Size-Dependent Stochastic Growth." *International Economic Review*, 31 (1990), 147–169.
- Reed, W. and R. Haight, "Predicting the Present Value Distribution of a Forest Plantation Investment." *Forest Science*, 42 (1996), 378–388.
- Ricardo, D., *On the Principles of Political Economy and Taxation*. London: St. Martin's Press, 1817.
- Rothery, J., "Some Aspects of Appraising Standing Timber." *Journal of Forestry*, 34 (1945), 490–498.
- Rust, J., "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica*, 55 (1987), 999–1033.
- Rust, J., "Structural Estimation of Markov Decision Processes." In *Handbook of Econometrics, Volume IV*, edited by R. Engle and D. McFadden. New York: North-Holland, 1994.
- Rust, J., "Numerical Dynamic Programming in Economics." In *Handbook of Computational Economic, Volume I*, edited by H. Amman, D. Kendrick, and J. Rust. New York: North-Holland, 1996.
- Thomson, T., "Optimal Forest Rotation when Stumpage Prices follow a Diffusion Process." *Land Economics*, 68 (1992), 329–342.
- von Thünen, J., *Der Isolierte Staat in Beziehung auf Landwirtschaft und Nationaloeconomie*. Berlin: Wiegandt, Hempel & Parey, 1826.