Metal Price Volatility: A Study of Informative Metrics and the Volatility Mitigating Effects of Recycling

by

Nathan Richard Fleming

Bachelor of Science in Mechanical Engineering
Northeastern University, 2004

Submitted to the Department of Mechanical Engineering
and the Engineering Systems Division
in partial fulfillment of the requirements for the degrees of

Master of Science in Mechanical Engineering

and

Master of Science in Technology Policy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2011

Author . . . .
Department of Mechanical Engineering
and the Engineering Systems Division
May 13, 2011

Certified by . . . .
Joel P. Clark
Professor, Department of Materials Science and Engineering
and the Engineering Systems Division
Thesis Supervisor

Certified by . . . .
Timothy G. Gutowski
Professor, Department of Mechanical Engineering
Thesis Reader

Accepted by . . . .
David E. Hardt
Professor, Department of Mechanical Engineering
Graduate Officer, Department of Mechanical Engineering

Accepted by . . . .
Dava J. Newman
Professor, Department of Aeronautics and Astronautics
and the Engineering Systems Division
Director, Technology and Policy Program

Metal Price Volatility: A Study of Informative Metrics and the Volatility Mitigating Effects of Recycling

by

Nathan Richard Fleming

Submitted to the Department of Mechanical Engineering
and the Engineering Systems Division
on May 13, 2011, in partial fulfillment of the
requirements for the degrees of
Master of Science in Mechanical Engineering
and
Master of Science in Technology Policy

Abstract

Metal price volatility is undesirable for firms that use metals as raw materials, because price volatility can translate into volatility of material costs. Volatile material costs and can erode the profitability of the firm, and limit material selection decisions. The undesirability of volatility gives firms an incentive to try to gather advanced information on fluctuations in price, and to manage—or at least control their exposure to—price volatility.

It was hypothesized that since price can be a measure of the scarcity of a metal, that other metrics of scarcity risk might correlate with price. A system dynamics simulation of the aluminum supply chain was run to determine how well some commonly used metrics of scarcity correlated with future changes in price, and to explore some conditions that strengthened or weakened those correlations. Additionally, prior work has suggested that increased recycling rates can lower price volatility. The study of the correlation of scarcity risk metrics with price is accompanied by a study on how the technical substitutability of secondary metal for primary, termed secondary substitutability, affects the price volatility.

The results show that some scarcity risk metrics modeled (alumina price, primary marginal cost, recycling efficiency, and the static depletion index) weakly correlate with future primary metal price, and hence volatility. Other metrics examined (recycling rate, mining industry Herfindahl Index, the acceleration of the mining rate, and the alumina producer's marginal cost) did not correlate with the future primary price. Correlations were stronger when the demand elasticity was high, the secondary substitutability was high, or the delays in adding primary capacity were low. Regarding managing price volatility, greater secondary substitutability lowers price volatility; likely because it increases the elasticity of substitution of secondary for primary metal—this result is explored mathematically.

The model results show that some scarcity risk metrics do weakly correlate with future primary price, but the strength of the correlation depends on certain market conditions. Moreover, firms may have some ability to manage price volatility by increasing the limit for how much secondary metal they can use in their product.

Thesis Supervisor: Joel P. Clark
Title: Professor, Department of Materials Science and Engineering
and the Engineering Systems Division

Acknowledgments

I would like to thank my team of advisors, Elisa Alonso, Frank Field, Randolph Kirchain, and Rich Roth for their guidance, their contributions, their helpful prodding, and their apparently inexhaustible patience. I could not have written this thesis without all of their help. I must also thank my wife, Bib, for her unflagging support, and my very young daughter, Mira, for being such a sport when Daddy was too busy to play.

Contents

List of Figures

List of Tables

Chapter 1

The Detrimental Effects of Price Volatility for Firms

1.1 Introduction

Raw material price volatility—variation of the price over time—can be detrimental to firms; mainly because the prices paid for raw materials are costs to the firm. Volatility in costs can influence firm profitability and their material selection decisions. For these reasons, firms have an interest in understanding what conditions may lead to changes in their raw material costs; moreover, they have an interest in trying to control or mitigate the risks that volatile prices pose. On the first issue, understanding price volatility, it is known that price can be considered a proxy for scarcity. The literature offers a number of methods for measuring scarcity risk, so this raises a question: can other scarcity risk metrics inform firms about the risk of price changes? And, if so, under what market conditions? On the second issue, managing volatility, the literature discusses conditions that are thought affect price volatility. For instance, larger secondary (recycled material) markets for the material and the ease of substituting for the material are thought to increase price stability (reduce volatility.) This seems to suggest that, as a corollary, if firms can increase the amount of secondary material that they are capable of substituting for primary, that they can increase price stability. This does not automatically follow, but the question of whether and when it is true is also addressed by this thesis.

1.2 Problems for the Firm

The price of a raw material is one of the costs of doing business of a firm, so volatility in raw material price translates into volatility in costs. Unless the firm can pass these costs on to its customers and increase their revenue in tandem with these costs, firms' profitability will be volatile when their costs are volatile (since profit is revenue minus costs.) Volatility (also referred to as instability) means motion both up and down: the risk of increases in prices is obvious, since increases in material prices can lead to reduced (or negative) profitability, unless the cost increases can be passed along to customers. Low prices are, of course, good for profitability but can present an enticing trap: an uncharacteristically (temporarily) low price of a material may open it up as as a viable choice for a product; however, the profitability of that choice may evaporate when the price returns to a higher level, but the firm may be 'locked in' to that choice. Switching away from the material may have its own costs.

So, the volatilities of the prices of the raw materials that a firm uses are of interest to the firm because they affect their profitability; however, the price volatilities of materials that the firm does not currently use—but are potential substitutes—are also of interest to the firm. Firms may shy away from using a material with a volatile price, thus volatility can limit material selection decisions. For example, Urbance et al. argue that the price volatility of magnesium (as well as its high price) deters automotive manufacturers from using it, despite its very desirable material properties—namely, its extremely light weight (Urbance, 2002).

Price volatility can also be detrimental to *producers* of materials: high prices can cause over-investment in capacity, low prices can cause underinvestment (Slade, 1988). Volatility can cause customers to substitute away from the material, lowering the demand for the material, to the possible detriment of the producer. This paper will focus on the considerations of firms that use raw materials, not produce them.

1.3 Minerals & Metals

The discussion so far pertains to materials in general. This broad category of substances covers everything from Alpaca wool to Zinc. While there are perhaps some features that all materials have in common, there are others that they don't: Alpaca wool is not mined, nor is it recycled (to any appreciable degree). Because price volatility is associated with a wide range of characteristic features of raw materials, this study will focus on a subset of materials whose commonalities will admit to some more pointed analysis and conclusions.

Materials can be divided into those that come from renewable sources, such as wood, rubber, or natural fibers; and those that are non-renewable, such as petroleum and minerals. Non-renewables are distinguished by the fact that they are drawn from stocks of finite supply that cannot be replenished. Non-renewable resources represent over 90% the total materials consumed yearly in the US since the mid-1950's, as shown in Figure 1-1. Of those non-renewables, industrial minerals and metals typically comprise more than half of the total, as shown in Figure 1-2. This study will focus on non-renewable resources, because, as will be explained, material scarcity contributes to price volatility and scarcity is frequently associated with non-renewable resources, particularly minerals and metals—the later being used as the example case for modeling.

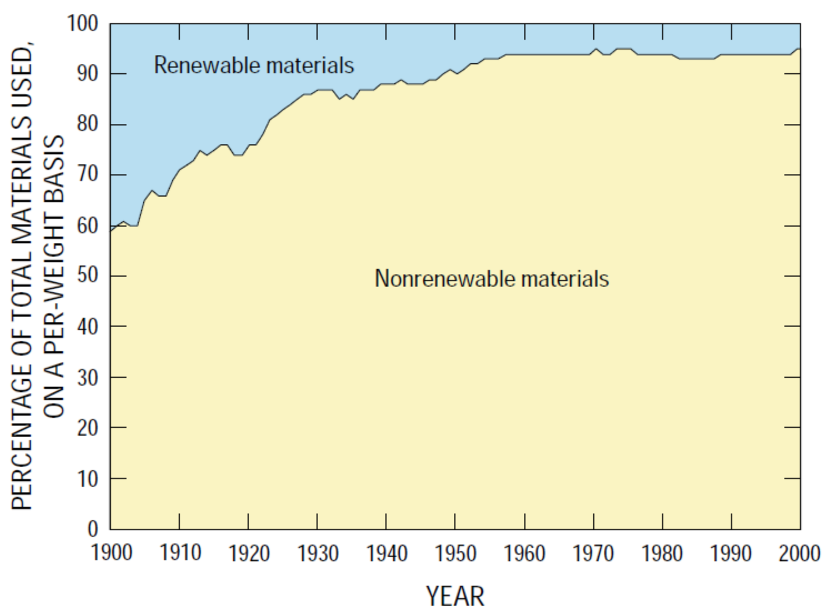

Figure 1-1: Percentage of renewable and nonrenewable materials used in the United States from 1900-2000. Use of nonrenewable resources has increased dramatically in the United States during the 20th century (Wagner, 2002)(modified from Matos and Wagner (1998))

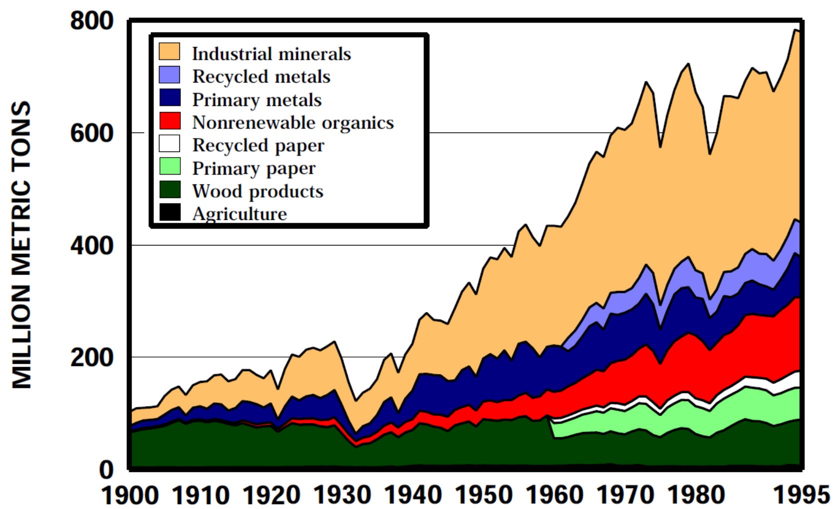

Figure 1-2: Measurement of the amount of raw materials (excluding crushed stone and construction sand and gravel) consumed in the United States (Matos and Wagner, 1998).

1.4 Pricing Systems & Strategies for Stabilization

In the minerals/metals markets there are a number of strategies used by firms to protect against the risks of price volatility: One is the choice of pricing system. There are two systems for the pricing of metals: prices are either set by the major producers (producer pricing) or by trade on an open market in a commodity exchange (commodity-exchange pricing) (Slade, 1988). Producer prices tend to be more stable, because they are often tightly controlled by the producers (Rees, 1985; Slade, 1988). However, producer prices may be higher than market fundamentals would suggest; they are also opaque, such that customers don't know what their competitors are paying (Slade, 1988), possibly causing them to pay more than they would in an open market.

A commodity exchange pricing system can have the advantage of transparency, and opens up the option of futures trading. Firms can provide themselves with a form of insurance against price volatility by hedging, which is "defined as the establishment of an opposite position on a futures market from that held and priced in the physical commodity"(Slade, 1988). Exchange prices are also influenced by speculation, which, as will be discussed, can be both good and bad for price stability. Each pricing system has its costs and benefits; firms choose between the different systems based on their own calculation of the costs and benefits to themselves.

The choice by a raw material consuming firm to buy through the producer system or the open market can be seen as a strategy to limit the risk of exposure to price volatility, by choosing either a more stable, but possibly higher, producer price; or trading on the open market which is less stable, but provides financial instruments that can protect against losses due to volatility.

Another strategy to stabilize price is through acquiring large inventories (Newbery, 1982; Slade, 1988). Large inventories can provide a buffer, in a sense, absorbing fluctuations in supply and demand. However, this is typically only done by larger firms, because the benefits of price stabilization are accrued by all firms in the market, but the costs are born by the firm with the large inventory; only firms with a large market share tend to see enough benefit from large inventories to make holding them worth the cost (Newbery, 1982; Slade, 1988).

1.5 Causes of Price Volatility

Price can be described as a signaling mechanism (Rutherford, 2002) that conveys information about the relative magnitudes of supply and demand. In a perfect competitive market, changes in either the supply of or the demand for a material would result in a quick and smooth adjustment to the price. For instance, if demand increases, prices should rise; the rising price will both lower the demand and increase the supply until a new equilibrium is reached. In a perfect market, price would only be as volatile as changes in supply or demand. Real markets contain imperfections which cause price adjustments to be not so smooth, and cause increased volatility in the price. Rees (1985) identifies four major inefficiencies that are commonly found in markets for mineral resources and contribute to price volatility.

1. **Delay in Adding Capacity:** When demand increases, new supply takes time to bring on-line—at least four years (Rees, 1985). During this time price rises more than it would if supply could be brought online more quickly.

2. **Delay in Shutting Down Capacity:** When demand decreases, producers resist taking capacity off-line "as long as it is making some contribution to overhead" (Rees, 1985). This can cause the prices to fall more than they would if the producers adjusted their capacity to the new demand more promptly.

3. **Growth-Sensitive, but Inelastic Demand:** "The major end-uses for a number of metals are the construction industry and the production of machinery, commercial vehicles and other intermediate goods" (Rees, 1985). The demand for many minerals depends on *growing* demand for final goods. Only when final goods producers are adding capacity are intermediate goods in high demand. Moreover, the fact that for final goods producers, the mineral costs are only a fraction of their total costs of producing goods, and would require a transaction cost to substitute, means that their demand for minerals is relatively unresponsive to price. So, in short, mineral demand is tied to economic growth and is relatively price inelastic.

4. **Small Open Market:** Only a fraction of the minerals produced are traded on the open market: Most are traded either internally within an integrated organization or are sold on long term contracts. This means that the trading price of a small portion of all the production must reflect the supply and demand imbalance for the whole market, making the market price very volatile.

In addition to the imperfections commonly found in minerals markets, there are other factors that affect the stability of price. The following factors are gathered from the literature and are with respect to metals; however, with the possible exception of the secondary market effect, would apply to minerals in general.

1. **Speculation:** Economic theory predicts that

   > "speculation should stabilize the market since purchases would be made when prices were low and selling would occur when prices were high. However, many of those dealing in the exchanges are metal fabricators. During a recession such firms will not be able to afford to hold large metals stocks and will, therefore, tend to reduce their holdings, so driving prices down still further. However, when trade improves, the will have the financial capability of busing more stocks, so fueling the metal price rise."(Rees, 1985)

   Another reason, that speculation does not stabilize market prices is even more general: as Hart and Kreps (1986) point out

   > "It is sometimes asserted that rational speculative activity must result in more stable prices because speculators buy when prices are low and sell

when they are high. This is incorrect. Speculators buy when the chances of price appreciation are high, selling when the chances are low. Speculative activity in an economy in which all agents are rational, have identical priors, and have access to identical information may destabilize prices, under any reasonable definition of destabilization. It takes extremely strong conditions to ensure that speculative activity (of the commodity storage variety) 'stabilizes' price, even in a very weak sense."

2. **Government Intervention:** Because government stockpiles are relatively large compared to the open markets, governments adding material to or releasing material from their stockpiles can have significant effects on prices (Rees, 1985). Theoretically government intervention could also stabilize prices.

3. **Primary Market Structure:**

   - *Horizontal Integration:* Markets in which fewer of the producers control more of the market tend to have more stable prices. "In a competitive market, firms are price takers. As prices fluctuate, producers can alter the level of their output, but they cannot control price directly. In a concentrated industry, in contrast, firms can choose to vary price, output, inventories, or some combination of the three in response to fluctuations in demand. The higher degree of control usually results in more stable conditions." (Slade, 1988)
   - *Vertical Integration:* The more vertically integrated a firm, the less it is exposed to the instabilities of prices; however, the more vertically integrated firms there are, the smaller the open market and the more unstable the price (Slade, 1988).

4. **Secondary Markets:** Larger secondary markets lead to more stable prices (Slade, 1988). When demand increases, secondary production can satisfy some of that demand. The costs of recycling depend on the recycling efficiency (fraction of waste recovered,) as prices increase it can become economical to recycle at a higher rate (Slade, 1988). Furthermore, the secondary producers can sometimes respond more quickly than primary producers to the increased demand (Alonso, 2010). The downside is that if there is a large low-cost secondary producer, prices will fall more drastically when demand declines (Slade, 1988).

5. **Byproduct Production:** If the mineral or metal is a byproduct of another, then production of that material will not be very responsive to price, leading to greater price fluctuations (Plunkert, 1999; Slade, 1988).

6. **End Use Stability:** If the end-use product of the mineral is subject to unstable prices, this market instability may affect the stability of the mineral's price (Slade, 1988). For instance, if automotive demand fluctuates, the prices of automotive materials may fluctuate as well.

7. **Substitutability:** The easier it is to substitute for a mineral, the more stable its price is expected to be (Slade, 1988). As prices increase, customers substitute for the alternative, thus diminishing demand and lowering the price. Likewise, when prices decline, customers substitute for the mineral in favor of its alternative.

8. **Exchange Rates:** For metals that have production costs measured in one currency, but are traded in another, the prices will change with exchange rates between the two currencies (Slade, 1988).

9. **Inflation Rates:** Changes in inflation rates translate into changes in costs for firms; these can translate into price changes (Slade, 1988).

10. **Disruption:** Disruptions to production, such as strikes, trade embargoes, wars, etc. can affect the supply and therefore the price of metals (Slade, 1988).

11. **Cost Changes:** Changes in the costs of producing firms, such as energy costs, affect the prices of the metals (Slade, 1988).

With so many factors that can influence price stability (and its opposite, volatility,) it is useful to try to organize them into a simpler arrangement. Brunetti and Gilbert (1995) group the sources of metal price volatility into three classes: those arising from informational considerations, hedging/speculative pressure, and the physical availability of material. All of the items listed above would fit into that system, with the possible exception of inflation and exchange rates, which might be shoehorned into 'informational considerations.' Many of the factors that affect price stability from the previous lists are characteristics of either supply or demand. Of the three classes of sources, Brunetti and Gilbert (1995) find that most of the medium-term (approximately monthly) variation in price is driven by physical availability, and the volatility was highest when the ratio of stocks to consumption was low. By physical availability Brunetti and Gilbert (1995) mean the fundamentals of supply and demand, consumption and stocks. Since supply and demand fundamentals are so important to the volatility of the price it will be worthwhile to explore them in more detail.

1.5.1 Demand

The term demand has two meanings in an economic sense: Demand can either be the 'schedule' of quantities demanded for a range of prices, as expressed by a demand curve; or demand can mean the quantity demanded at a given price (Friedman, 1976). Assuming there is no shortage of material and that which is demanded is received, the quantity demanded and consumption will be identical quantities. For clarity, when referring to the demand schedule, it will be called as such or will be referred to as the 'demand curve;' when speaking of quantity demanded, the term 'demand' will be used. Likewise with supply and supply schedule (or curve.)

In the US, one of the primary drivers of demand for minerals (and metals) is population growth; with that population growth comes construction of new homes and businesses and transportation infrastructure and hence demand for minerals and metals (Sznopek, 2006). In fact, the consumption of minerals has increased faster than the population; from 1900-2000 the average annual growth of the US population was 1.3 percent, while mineral consumption increased at 3.1 percent (Sznopek, 2006). This means that per capita minerals consumption has been increasing over that period. Moreover, the patterns of material consumption change with consumer preferences. For instance, in the transportation sector, a preference for larger vehicles increases the amount of material in an automobile. The preference for hybrid vehicles increases the demand for rare earths and lithium.

Legislation also affects demand: Regulations on toxic metals can reduce their demand, but some regulations can increase the demand for other metals (Plunkert, 1999). For instance, vehicle emissions standards increased the demand for platinum group metals, which are used in catalytic converters to lower emissions (Sznopek, 2006). Likewise, CAFE standards make lightweight vehicles more desirable, likely leading to an increased use of aluminum and a lower reliance on iron and steel, which has been observed (Kelly, 2005).

The drivers of demand for the US are likely to apply to the rest of the world, especially if the rest of the world industrializes, and becomes more like the US in terms of income.

The world population is expected to increase until about 2075 (United Nations, 2004), and world GDP has been growing at a much faster rate than population (World Bank, 2006); if consumption of minerals can be assumed to grow with GDP, then it is likely that the global consumption of minerals per capita has been growing faster than the population on the global scale, as has been seen in the US.

So the overall story of minerals demand is that it has been increasing and can be expected to continue to do so well into the future.

1.5.2 Supply

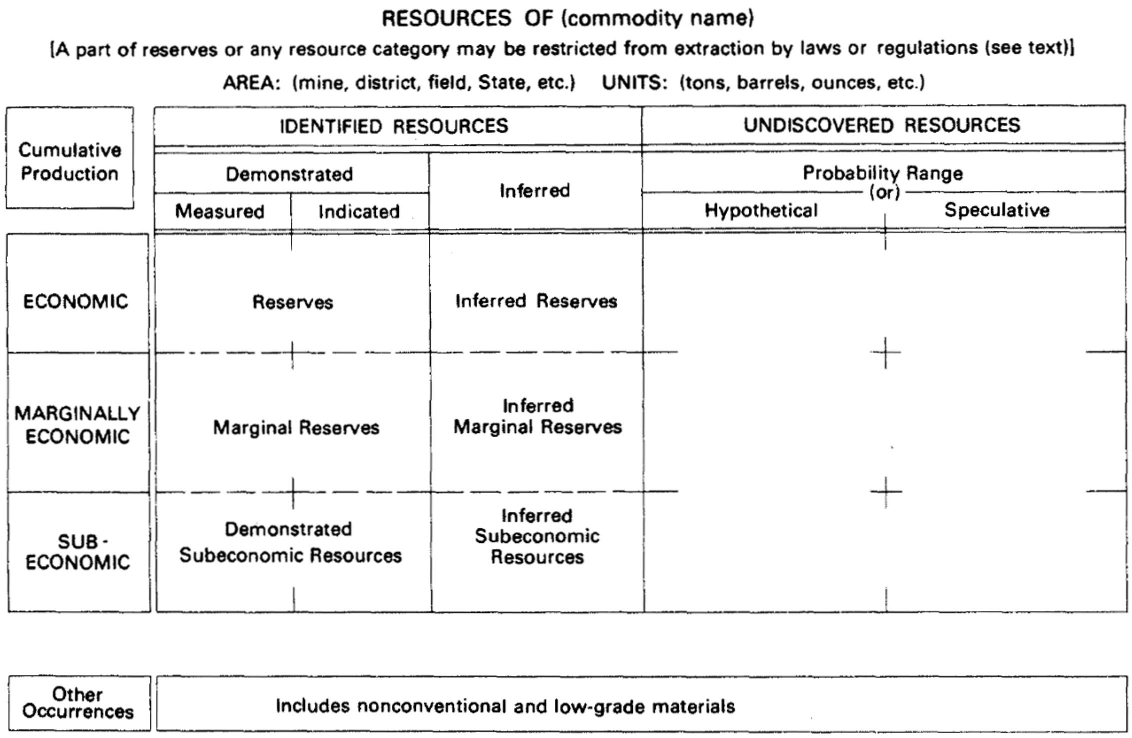

Figure 1-3: McKelvey Diagram: Classification of Resources (USGS, 1980)

The supply of minerals depends on the distribution of minerals in the earth's crust and the associated costs of extracting them. The costs to a mining firm are that of finding the deposits, setting up a new mine, processing the ore, and getting the mineral to market. The costs of extracting from larger deposits are lower than those of extracting from smaller ones because they require less discovery costs per unit ore extracted, and have lower fixed costs per unit to set up a mining operation. Higher ore grade is typically less costly because less rock must be processed per unit of mineral; and, finally, it is less costly to be close to market, or in a location with good infrastructure for getting product to market.

Resources in the ground can be divided into those that have been *identified* and those that are as yet *undiscovered* but are assumed to exist. Of those that have been identified, there are those that have been *demonstrated* to exist, because they have been *measured* directly or are strongly *indicated* to exist, and those that are *inferred* to be there by "assuming continuity beyond measured resources."(USGS, 1980). Across all of these categories, there are resources, that because of their ore grade and location, are expected to be *economic* to extract at current market prices; others are considered *marginally economic* if their cost of extraction is close to being economic, while the remainder are termed *subeconomic*. The quantity of minerals that is demonstrated and economic is known as the *reserve*. This naming scheme is presented in what is known as a McKelvey diagram, as shown in Figure 1-3 in terms of reserves.

On the one hand, the amount of minerals in the earth's crust is absolutely immense. Wagner (2003) cites Brooks (1973) as claiming "A single cubic mile of average crustal rock contains a billion tons of aluminum, over 500 million tons of iron, 1 million tons of zinc, and 600 thousand tons of copper " However, what is more important is what portion of that material is of a grade that would be economic to extract, or, at least, that will be in the near future. Questions of how much material is in the earth's crust, how or when will it become economic, and will we run out—questions of resource scarcity—deserve their own treatment and are a focus of mineral economics.

1.6 Mineral Scarcity

The discipline of economics, itself, is in essence the study of the interaction of nearly boundless human desires and the scarce resources available to meet them (Wetzstein, 2005). The study of mineral (including material) scarcity is a subdiscipline of economics (Gordon and Tilton, 2008). Minerals have the peculiar feature, that unlike other economic resources such as labor or capital, minerals are absolutely finite in quantity and they are only depleted, not expanded, through use or other activity that destroys them (Barnett and Morse, 1963). This startling fact has led to a long-standing fear that we will run out of mineral resources. This point of view has been labeled 'Malthusian' in that it echoes Thomas Malthus's concerns in the 18th century about the implications of an exponentially growing human population using a finite land area to grow its food (Barnett and Morse, 1963; Malthus, 1963). Meadows et al. (1972) *Limits to Growth* study in 1972 predicted that present economic growth was unsustainable, echoing Malthus's concern.

The problem with the Malthusian point of view is that it does not take into consideration the varied quality of mineral resources. Mineral ores are typically span a range of quality—translating into a range of extraction costs. As higher quality minerals become depleted, prices will rise and lower quality ores will become economical, thus expanding the quantity of available resources. Also, as prices rise, demand will fall, and alternatives will become feasible. Views of scarcity that take into consideration this theoretical march down the quality scale are termed 'Ricardian,' after the economist David Ricardo who first proposed them (Barnett and Morse, 1963; Ricardo, 2006). The concern about scarcity in the 'Ricardian' view is not so much that minerals will be entirely depleted, just that they will become increasingly costly to extract. These terms 'Malthusian' and 'Ricardian' will be useful in classifying metrics for scarcity, since the metrics that will be considered make implicit assumptions that fall into either of those two categories.

1.6.1 Mineral Scarcity & Price

In 1931, Harold Hotelling wrote a paper that was key to the discipline of mineral economics, predicting that competitive producers of exhaustible resources would do so at a socially optimal rate, and that as a result the price of the resource would increase with the rate of interest (Hotelling, 1931; Devarajan and Fisher, 1981; Solow, 1974). Hotelling's model did not include an extraction costs. With the inclusion of extraction costs the result changes slightly such that it is the component of the price that is beyond the extraction costs that increases at the rate of interest. This value, the difference between the price and the extraction costs, is termed 'scarcity rent' or 'Marginal User Cost' (Tietenberg, 2009). From Hotelling's theory, we expect that scarcity rent will increase over time. However, the scarcity rent is difficult to observe, because it is not something that can be observed alone; it can only be observed indirectly through price (because it is a component of price) where it is combined with the extraction costs. In their seminal study, Barnett and Morse found that natural resource prices had actually been in decline at the time of the writing, contrary to expectations (Barnett and Morse, 1963). Slade (1982) suggested that the price increase due to scarcity had simply been masked by decreasing extraction costs, and that price paths were beginning to take on a U-shape as increasing scarcity rent began to overtake diminishing extraction costs. However, time has not born out Slade's predictions. In the subsequent decades since her 1983 study, the prices have continued to decline in line with Barnett and Morse's study (Krautkraemer, 1998; Tilton, 1999). Figure 1-4 shows a composite price index for 5 metals and 7 industrial minerals compiled by the USGS (Sullivan, 2000) that demonstrates the decline in prices seen over the course of the 20th century. Note the increases in the index that occurred in the 1970's and early 1980's. That was the era that saw Meadow's dire predictions Meadows et al. (1972), and led Slade (1982) to offer the theory of the U-shaped price path.

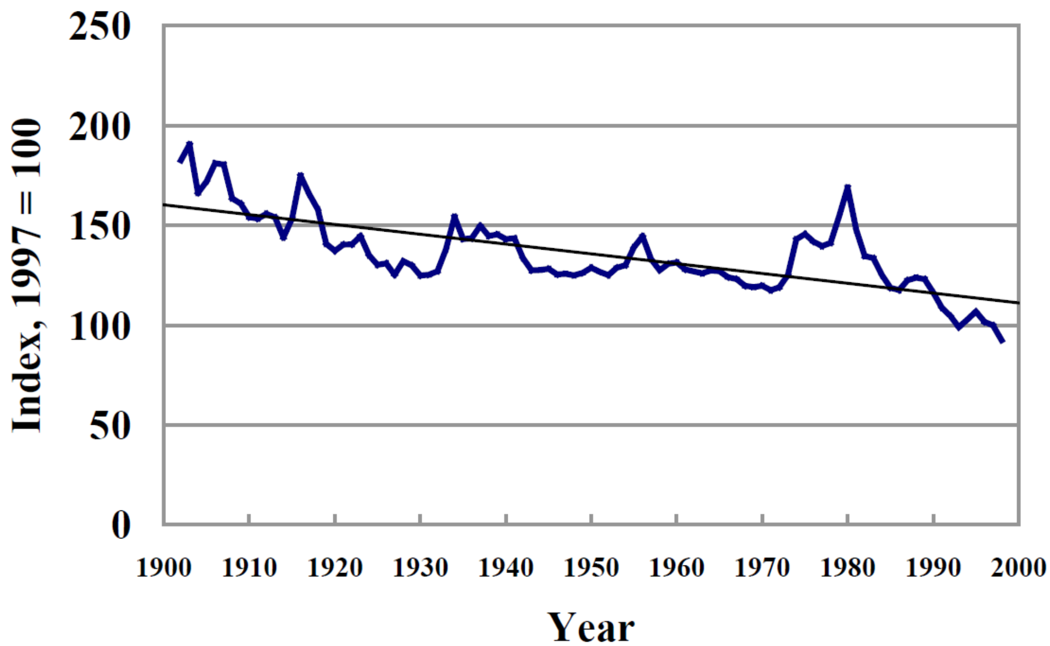

Figure 1-4: USGS composite price index for five metal commodities (copper, gold, iron ore, lead, zinc) and seven industrial mineral commodities (cement, clay, crushed stone, lime, phosphate rock, salt, sand and gravel) in 1997 dollars (Sullivan, 2000)

So, why have prices not increased despite a growing human population that is increasingly industrialized? There are three factors that tend to mitigate scarcity: exploration and discovery, technological progress, and substitution (Tietenberg, 2009). Exploration expands the stock of identified and economic resources, thus diminishing scarcity. Technological progress refers to improvements to the machinery or processes mineral extraction that lower costs by requiring less expenditure of labor, capital, and energy per unit of output. This is likely a major reason the declining resource prices that have been observed.

Substitution could occur as a resource becomes more scarce and prices rise. Eventually, as prices become high enough, a substitute 'backstop' technology will become more economical, and consumers will substitute the backstop for the scarce resource (Solow, 1974; Nordhaus et al., 1973).

Over the long term, mineral scarcity has been decreasing, as evidenced by the declining prices of minerals. Slade (1992) found that there is evidence that scarcity has indeed been decreasing as the declining prices suggest, and that the markets are not underpricing resources—with one exception: The decreasing prices of minerals are not reflecting the costs of environmental damage of mineral extraction.

1.6.2 Short-Term Scarcity

Although scarcity has been declining in the long-term, short-term scarcity can still present an issue. As was discussed earlier, shorter-term relative changes in supply and demand cause fluctuations in prices that are exacerbated by market imperfections and other factors that amplify price instability. A dramatic example of a short-term scarcity event is the so-called 'cobalt crisis' of the late 1970's. In a period of increasing demand, a series of events diminished the supply of cobalt on the world market: The US government halted sales from its stockpiles; much of the world's cobalt producing capacity was in Zaire, which already had low inventories, when a political upheaval further reduced production (Plunkert, 1999). This sudden scarcity lead to a dramatic price spike, as shown in Figure 1-5.

1.7 Efficient Market Hypothesis

It has been proposed in the past that market prices immediately reflect all available information; this is known as the Efficient Market Hypothesis (EMH) (Durlauf, 2008). It was first developed by Samuelson and Fama, independently in the same year (Fama, 1965; Samuelson, 1965). The implication of this theory for material scarcity, if it held, would be that the market price would contain all of the available information about material scarcity, making it the best indicator of scarcity, in that no other measure of scarcity would contain information not reflected in the price. While the EMH has its proponents, there have been a number of strong arguments against it, the first and most notable by Grossman and Stiglitz (Grossman, 1980) who argued that it would be impossible for markets to be completely

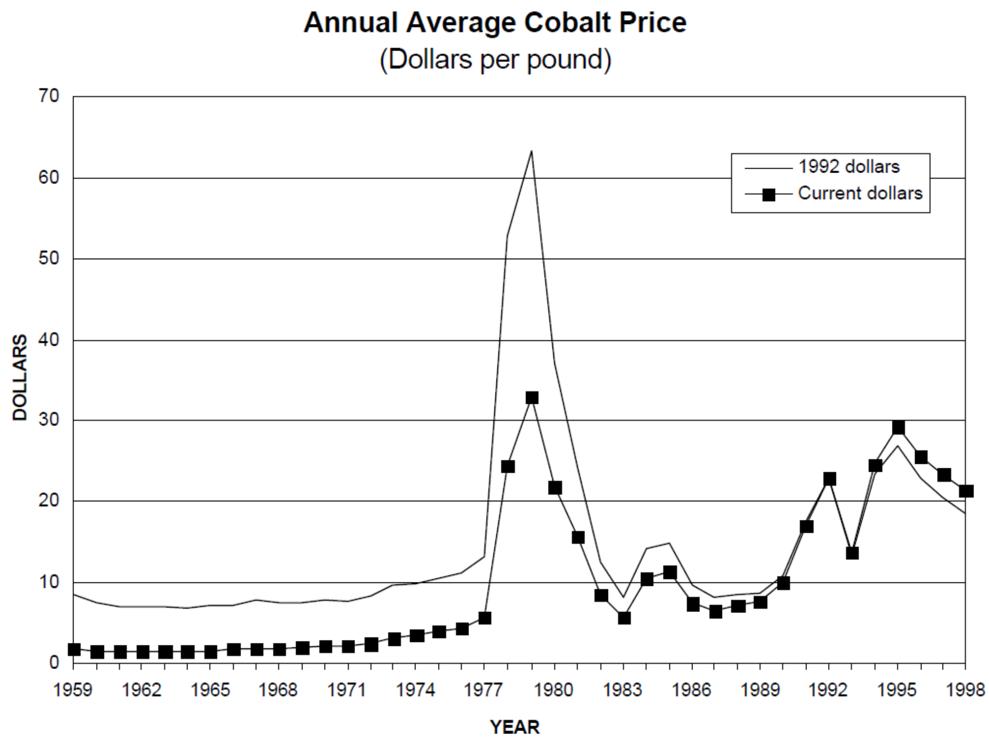

Figure 1-5: Cobalt Prices, 1959 to 1998 (Plunkert, 1999). Note the dramatic price spike.

informationally efficient (which is what the EMH is proposing.) Another theoretical argument against it is that made almost a decade before by Simon (Simon, 1956) that it is never possible in economic decision making to have perfect information or be completely rational in the face of that information. These are theoretical arguments; there is also empirical evidence against the Efficient Market Hypothesis. For metals, Davutyan has shown that for several commodity materials examined, there is some cyclicity (periodic behavior) to the prices and they are not a random movement (Davutyan and Roberts, 1994) as the Efficient Market Hypothesis would predict. This study takes as a fundamental assumption, that markets are not perfectly informationally efficient, and that other scarcity risk metrics can provide information on the scarcity risk of a material, not already captured in the price.

1.8 Scarcity Risk Metrics

So can scarcity be measured? A number of metrics have been devised to provide measures of scarcity risk. Alonso et al. (2007a) have identified a range of metrics in the literature that have been used to measure scarcity risk, and have developed a taxonomy for organizing

Type	Assumption	Metric	Units
Institutional		Geographic Structure based on Supply	
Efficiency		Geographic Structure based on Production	
		Institutional Structure based on Production	%
		Institutional Structure based on Consumption	%
		Recycling Rate	%
		Recycling Efficiency or Recovery Rate	%
		Market price	$
Physical	*Malthusian*	Static Index of Depletion	years
Constraint		Exponential Index of Depletion	years
		Relative Rate of Discovery and Extraction	
		Time to Peak Production	years
	Ricardian	Average Ore Grade	%
		Costs	$
		Market Price	$

Table 1.1: Taxonomy of Scarcity Metrics (Alonso et al., 2007a)

them. First they have distinguished between metrics of risk that are measures of the efficiency of markets, firms and governments (institutional efficiency metrics); and between metrics related to physical quantities in the material system (physical constraint metrics) (Alonso et al., 2007a). They further subdivided the physical constraint metrics into those that are Malthusian in that the due not capture any notion of resource quality in them, and those that are Ricardian (in that they do capture such a distinction.) The metrics she has identified are listed in Table 1.1 and can be summarized as follows.

- *Geographic Structure based on Supply:* The distribution of reserves by geographic location (Chapman and Roberts, 1983). It is assumed that greater diversity leads to greater efficiency (Alonso et al., 2007b).

- *Geographic Structure based on Production:* The distribution of production by geographic location. It is assumed that greater diversity leads to greater efficiency (Alonso et al., 2007b; Plunkert, 1999; Chapman and Roberts, 1983).

- *Institutional Structure based on Production:* Distribution of market shares of producers (McClements and Cranswick, 2001). It is assumed that greater diversity leads to greater efficiency (Alonso et al., 2007b).

- *Institutional Structure based on Consumption:* Distribution of market shares of consumers (McClements and Cranswick, 2001). It is assumed that greater diversity leads to greater efficiency (Alonso et al., 2007b).

- *Recycling Rate:* The ratio of recycled material production to total material production (U.S.G.S, 2009; Ruhrberg, 2006). It is assumed that markets that depend more on recycled metal are more efficient (Alonso et al., 2007b).

- *Recycling Efficiency or Recovery Rate:* The fraction of material waste recovered for recycling (Ayres et al., 2003; Ruhrberg, 2006). As with the recycling rate, it is assumed that markets that depend more on recycled metal are more efficient (Alonso et al., 2007b).

- *Market price:* Assuming the market is efficient, price should reflect the scarcity of the material. (Chapman and Roberts, 1983; Barnett and Morse, 1963; Cleveland, 1993).

- *Static Index of Depletion years:* The ratio between the reserves of material and the use rate(Meadows et al., 1972; MMSD, 2006). This metric does not include the effects of recycling and discovery, and measures how soon the current reserve will run out at the current extraction rate.

- *Exponential Index of Depletion years:* The depletion index based on an exponential rate of growth $\frac{1}{r}ln(r\frac{supply}{c_0}+1)$, where r is the growth rate, and c_0 is the rate of consumption of the present year (Meadows et al., 1972; MMSD, 2006).

- *Relative Rate of Discovery and Extraction:* The ratio of the rate of discovery to the rate of extraction—a measure of how the reserve is depleted or expanded. (Malthus, 2006; Gordon et al., 2006).

- *Time to Peak Production years:* Hubbert's theory of peak production predicts that the extraction of resources will follow a bell-shaped curve, with extraction eventually peaking and going into decline. The time to peak production measures the distance between the present moment and the point of declining extraction rates (Hubbert, 1962).

- *Average Ore Grade:* In the Ricardian model of resource scarcity, the ore grade is a measure of scarcity in that higher-grade and presumably cheaper ores will be extracted first, lower ore grades later. Declining ore grade is a sign of increasing scarcity (Ayres et al., 2003; Chapman and Roberts, 1983).

- *Costs:* Related to average ore grade, costs increase as higher quality and more easily accessible resources are used up and the most accessible discoveries are made. (Barnett and Morse, 1963; Cleveland, 1993; Chapman and Roberts, 1983)

1.9 Summary

Price volatility of raw materials is detrimental to firms, because it can affect their profitability and constrain their material selection decisions. Firms try to protect themselves from price volatility by their choice of price system or possibly by holding large inventories. Price volatility is due to changes in supply and demand that are amplified by market imperfections and by other characteristics of the market. Among those characteristics are the size of the secondary market and the degree to which the material can be substituted for. There are other sources of volatility, such as speculation, but much of the volatility in price is due to supply/demand tensions. The supply of minerals has the special characteristic that the supply is ultimately exhaustible. This has led to the specialized study of mineral scarcity. A conclusion from the scarcity literature is that long-term scarcity risks have not yet manifested, but short term scarcity can cause price volatility. As a result of this extensive interest in scarcity, a number of metrics have been devised that provide insight into the risk of a material becoming scarce.

1.10 Addressing Gaps in the Literature

Because price volatility of raw materials is detrimental to firms, they have an interest in understanding their risk due to price instability; moreover, they have an interest in trying to control or manage price volatility.

The existing literature has established that a major source of price volatility is the scarcity of the material. Price, itself, can be considered a measure of scarcity. The literature has identified a number of other metrics that identify risks for increased scarcity. That raises a question: If price volatility is related to scarcity, then can other scarcity risk metrics provide information on price? No one has tried to determine and compare in a structured manner to what degree different scarcity metrics correlate with price. One of the objectives of this thesis is to explore the correlation between scarcity risk metrics and future price. The literature has also established that there are many market characteristics that affect price stability. As a corollary to the question about how metrics relate to price, this thesis will try to determine how some market conditions, that are known to affect price volatility, affect the correlations between metrics and price.

In regards to managing price volatility, the literature suggests that larger secondary markets and greater availability of substitute materials increases the price stability of a given material. Substitutability has been framed as substitution for *other* materials. No one has explicitly treated how the degree to which secondary metal can be substituted for primary metal affects price stability. It is expected that increasing the secondary substitutability would provide a stabilizing effect on price, because it increases the availability of a substitute and is likely to increase the size of the secondary market. However, substituting secondary for primary metal is not the same as substituting for a different metal: The supply of secondary metal is not independent of primary metal, the prices and supplies are linked. And, the effect of changing the technical substitutability of primary metal is not just a change in the relative size of the secondary market. In short, the effect of secondary substitutability on primary price is not unambiguous, and will be explored.

The overarching goal of the research questions will be to provide strategies for firms to gain insight into their scarcity related price volatility risk, and to provide them with strategies for diminishing the price volatility of their raw materials, or at least their exposure to that volatility. It will be shown that these two sets of strategies are linked.

Chapter 2

Model

2.1 Introduction

Both of the research questions were explored using a system dynamics simulation of a commodity metals market. The aluminum market was chosen because it can be considered representative of major metals markets and, since primary and secondary aluminum are not always perfectly substitutable, it lends itself to a study of secondary substitutability. The model uses standard 'textbook' supply chain modeling components; including structures for setting price, setting costs, and changing capacity. Particular attention is paid to the setting of price, since it is so important for addressing the question on the relationship between scarcity metrics and primary price. The data for the model were obtained from USGS sources and relevant literature. The model was roughly calibrated to the historical aluminum prices and primary production for 1985 though 2008. This chapter contains, first, a general description of the model, and, second, the specifics of the implementation in aluminum.

2.2 Model Intent

The goal of the model is to provide a coherent framework for evaluating the interactions of the major components of a commodity metals market. It needs to be fine-grained enough that the scarcity metrics outlined in the first chapter can be explicitly modeled as endogenous variables.

2.3 System Dynamics Commodity Market Models

System dynamics is a field that was developed at MIT by Jay Forrester in the 1960's (Forrester, 1961). System dynamics simulations consist of a range of variables and constants. The variables are related to each other through functions—including first-order differential equations. The simulation starts with the variables set at some initial conditions; based on the relationships between variables (functions) it solves for the value of the variables at the next time step, and then repeats until the final timestep is reached. The output of the simulation is the value of the variables at each time step. It is very useful for determining how variables behave over time (dynamically,) hence the name. It is particularly useful for modeling complex systems that behave in a non-linear fashion, such as commodity markets. Most of the methodology in this study comes from a textbook by Sterman (2000) and work by Meadows (1970) into modeling commodity markets.

2.4 General Model

2.4.1 Model Structure

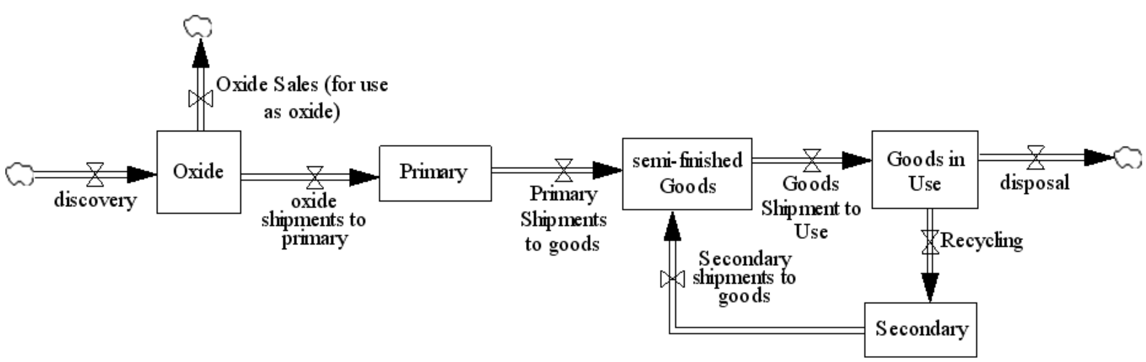

Figure 2-1: Model Major Stocks and Flows

The model captures the major stocks and flows of metal starting from discovery in the ground all the way to the use in products and the eventual disposal or recycling of the metal. Each of the major stocks is accompanied by prices, inventories, production capacities, orders, and the other variables. There is significant commonality to the overall structures of how those variables are functionally related to each other across each of the major stocks in the model. So, it is convenient in terms of explaining the model to discuss how, for instance, price is set, as opposed to having to separately explain how the oxide price is set, how the primary price is set, and so on. Stocks on either 'end' of the model, mining and goods demand, have their own peculiarities, as does secondary metal, so those portions of the model will be explained separately.

Market Actors

The market simulation has four major market actors that are endogenously modeled: Mining firms (regions), primary metal producers, secondary metal producers, and goods producers. Consumers of goods and consumers of oxide (not for primary production) are modeled exogenously. Descriptions of the market actors are as follows:

- **Mining regions:** This model component represents the aggregation of firms that explore for, mine, and process metallic ores, segregated by region. Aggregating all firms and functions simplifies the model. Data for ore body size and quality is often collected on a country by country basis; dividing mining firms into regions allows each region to be matched to ore bodies.

- **Primary producers:** This category represents those firms that produce primary metal from oxide. The primary producers were subdivided into three identical firms to allow one of those firms to be shut down as part of one of the thesis experiments.

- **Secondary producers:** The aggregate of firms that collect old scrap, sort it, and provide it to market in a form that can be used by remelters or alloy producers. New scrap, scrap that is a byproduct of manufacturing processes, is not modeled since it does not add to the overall material supply—unlike old scrap (goods at the end of their lives) which does (Radetzki and Duyne, 1985). There is actually only one secondary producer in the model, it has been divided into three sections, like the primary producers, to allow the flexibility to model different market segments explicitly; however, in this model all secondary producers were treated as one.

- **Goods producers:** This category models firms that consume primary and/or secondary metal to produce goods. Those goods could be finished products, or semi-finished goods that will then be used in other products. This category includes both firms that use only primary, only secondary, or some mix of the two—it is the aggregation of those firms. As with secondary producers, the goods producers are actually divided into three sections, but all are identical, so there is, in essence, only one goods producer that represents the aggregate of all.

Demand

There are two exogenous demands in the model: That of metal-containing goods (termed just 'goods') and that of oxide to be used for purposes other than making primary metal.

- **Goods demand:** demand for goods is modeled through a constant elasticity demand function of the form $Q = A \cdot P^{-\epsilon}$, where Q is demand, P is price, and ϵ is the price elasticity of demand. An estimation of 0.3 for ϵ provided the best match to historical data for the aluminum implementation.

- **Oxide demand** (not for primary production): Oxide has other industrial uses besides making primary metal.

For the study on the correlation between scarcity metrics and price, a stochastic element was introduced to the goods demand function. This was accomplished by adding noise to the goods demand function. The deterministic demand at any time was multiplied by a randomly chosen number from a specified range: Meaning that the demand for any given price would not always be exactly the same. The number by which the demand was multiplied was actually not completely random, but in fact depended on past values of that multiplier; a purely random number would be termed 'white noise,' but what was used is termed 'pink noise.' Sterman (2000) recommends using 'pink noise' over white noise when modeling noise, due to the former being more realistic. The stochastic element allows the model to be run multiple times with the same characteristics, but with slightly different 'paths' for the variables due to the noise. This was useful for the metrics study because the question in that study is how do the paths of different variables (metrics) predict the path of price. To only look at one path in this case would beg the question to what extent did the specific path determine the result. The variation introduced by the noise was set at the maximum value that did not cause variation in production to deviate beyond historical values.

Because of the stochastically determined goods demand, no two model runs will be exactly the same—even with the same initial conditions. For these reason when looking at the results of a particular model run, the term 'representative' run will be used. The term means that the run represents the typical results for a model run with the same initial variables.

Price

In the model, the price for any marketable stock in the model—oxide, primary metal, secondary metal, goods—is set as a function of inventory coverage; where inventory coverage is the ratio of the current inventory over the shipments (sales) of the product. Increases in inventory coverages lower price; decreases raise price. This relationship mimics the result of the classical economic approach to price setting, where there is a 'schedule' of what quantity of goods will be supplied at a given price—a supply curve—and there is, likewise, a schedule of what quantity of goods are demanded at a given price—a demand curve. The equilibrium price is the price at which the quantity supplied equals the quantity demanded—the intersection of the two curves. In the classical approach, when there is a change in either supply or demand, represented by a shift in either the supply or demand curve, the equilibrium price changes: Increases in supply lower price, increases in demand raise price. In the inventory coverage model the behavior is similar but supplied by a different mechanism. An increase in supply will lead to an increase in inventory, and, hence, an increase in inventory coverage, and a lower price. Increases in demand will increase customer orders and therefore shipments, which will lower inventory coverage and raise price. As Sterman observes, inventory at any given time is the integral of all production (supply) up to that point minus the integral of all shipments (demand) up to that point (Sterman, 2000). So, in a very real sense, inventory is the difference between supply and demand—a further justification for modeling price as a function of inventory.

There is more to the price-as-a-function-of-inventory story: Sterman (2000) cites Meadows's work in commodity production cycles (Meadows, 1970) as a reference for that relationship. Meadows in turn cites a study of the dynamics of the world cocoa market by Weymar (1968). Weymar introduced the idea of price being related to inventory coverage in econometric studies of the cocoa-market, but he himself was building of the 'supply of storage' theory developed by Brennan (1958) years earlier.

The supply of storage theory can be summarized as follows: There are costs to holding inventory; namely, it costs money to warehouse goods and there is a theoretical cost associated with the risk of holding inventory (holding inventory is risky in the sense that the price

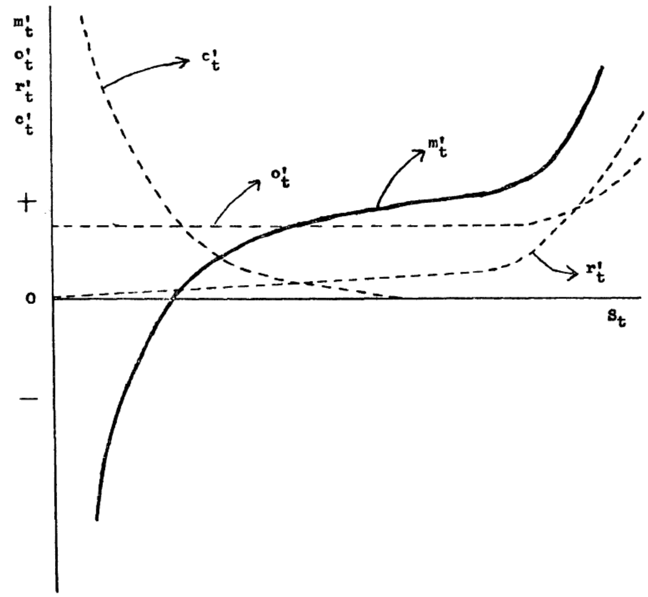

Figure 2-2: Supply of Storage: The marginal net cost of storage (m_t') is the sum of the marginal physical storage costs(o_t'), the marginal risk costs(r_t'), minus the marginal convenience cost(c_t')(Brennan, 1958). The inventory is given by S_t. Note that subtracting the convenience cost is equivalent to adding the convenience benefit.(Brennan, 1958)

could suddenly drop and the inventory is then worth less.) There are also benefits to holding inventory, that Brennan terms as (negative) convenience costs (Brennan, 1958). The idea is that not having inventory on hand is inconvenient and can be costly. For instance if a firm does not have the inventory to meet customer orders, it may lose that business or have to provide the inventory late and at a discount. Brennan postulated that convenience benefits would be very high for low inventories and rapidly diminish to zero with increasing inventories. He suggested that physical storage costs would be constant per unit of storage, until the warehousing space was depleted; at that point the costs would increase rapidly. The risk costs per unit would be gradually increasing until they reached a level at which a change in price would affect the firm's position, at which point they would increase rapidly. These relationships can be seen in Figure 2-2. The marginal net cost of storage is the sum total of all the marginal costs and benefits.

The demand for storage, in Brennan's conception, depends on the expected change in price. That is best explained through example. Assume the price of the good in question was expected to increase over the next period; there would be a 'demand' to hold on to that good to reap the benefits of selling the good at the later period at the higher price. The equilibrium level of storage then is where the supply of storage equals demand. In other words, where the expected change in price equals the costs of storing one more unit of good (marginal net storage costs.) The intuition is as follows, from the equilibrium level, to store one more unit of goods would cost more to store than the gains expected from holding on to it; to hold less would forgo the gains from the difference between the price change and the costs of storing that unit of goods.

For goods with a futures market, the expected price in the next period would be the futures price; otherwise the price is just the expectation of how price will move. Weymar further developed Brennan's theory and added the observation that it is more convenient to talk of inventory divided by consumption (inventory coverage), because that provides meaning and context to the inventory quantity (Weymar, 1968). Brennan found support for his theory in the econometric modeling of butter, wheat, oats, eggs, and cheese (Brennan, 1958); while Weymar did likewise for the world cocoa market (Weymar, 1968).

The price as a function of inventory coverage paradigm used in system dynamics is simply an extrapolation of the supply of storage theory. If, at equilibrium, the expected change in price between now and some future time period correlates with the inventory coverage, then, knowing the current price and the inventory coverage, we can postulate what the price should be in the next period. Sterman, however, recommends that when modeling the price, the price be treated as being influenced by both the inventory coverage and the costs of the producer (Sterman, 2000), since producers tend to want to pass along some of their costs to the purchasers. Furthermore, Sterman (2000) suggests that over time, firms adapt to any given price level. This concept is captured in the model variable 'trader expected price,' which forms the baseline price from which changes are made based on costs and inventory. A simplification of the model structure for setting price, is reproduced in Figure 2-3. Each item of text represents a variable, the arrows show which variable affects the others. In the model, price is a function of both costs and inventory coverage, the relative sensitivities being set in the model calibration phase.

Costs

The variable costs per unit of output to each of the market actors is modeled as the sum of material costs, energy costs, labor and overhead, and production costs. The energy, labor and overhead components were subjected to a learning effect to model the declining

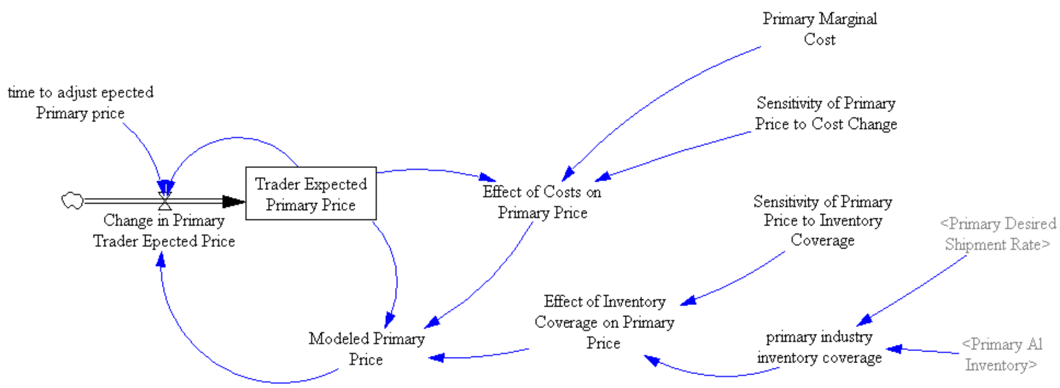

Figure 2-3: Model Structure for Price

production costs for materials that has been reported by Slade (1982) and Barnett and Morse (1963). The learning effect were achieved by multiplying those costs by a fraction of one that decreases at a decreasing rate with time—a learning curve. The cost components are described in more detail below:

1. *Material Costs:* For primary and goods producers the material costs are assumed to be simply the price of their raw materials. For primary metal that is the price of oxide. For goods producers, it is the primary and secondary metal prices times the ratio at which the materials are used. The material costs for mining regions is the total of the costs per unit of rock mined, divided by ore grade, which is the mining cost for unit of ore. For secondary metal, the material cost was an increasing function of the recycling efficiency; it was assumed that, as has been reported (Tilton, 1999; Slade, 1988), the costs of collecting scrap increase as the easier to find scrap is collected leaving more difficult to obtain sources.

2. *Energy Costs:* Energy costs are modeled as the price of fuel or electricity for processing the materials. The literature was consulted for estimates of the energy requirements per unit product.

3. *Capacity Costs:* It is assumed that there is a desired capacity utilization that corresponds with a minimum per unit cost. Exceeding the desired capacity utilization corresponds to higher costs per unit—despite there being more units produced; falling short means that per unit costs go up, because there are less units of production over which to divide the capacity costs.

4. *Labor and Overhead:* Labor and Overhead Costs have been used as a 'catch-all' term to cover the remainder of the per unit costs necessary to model historical prices.

A simplified version of the model structure for cost setting is shown in Figure 2-4

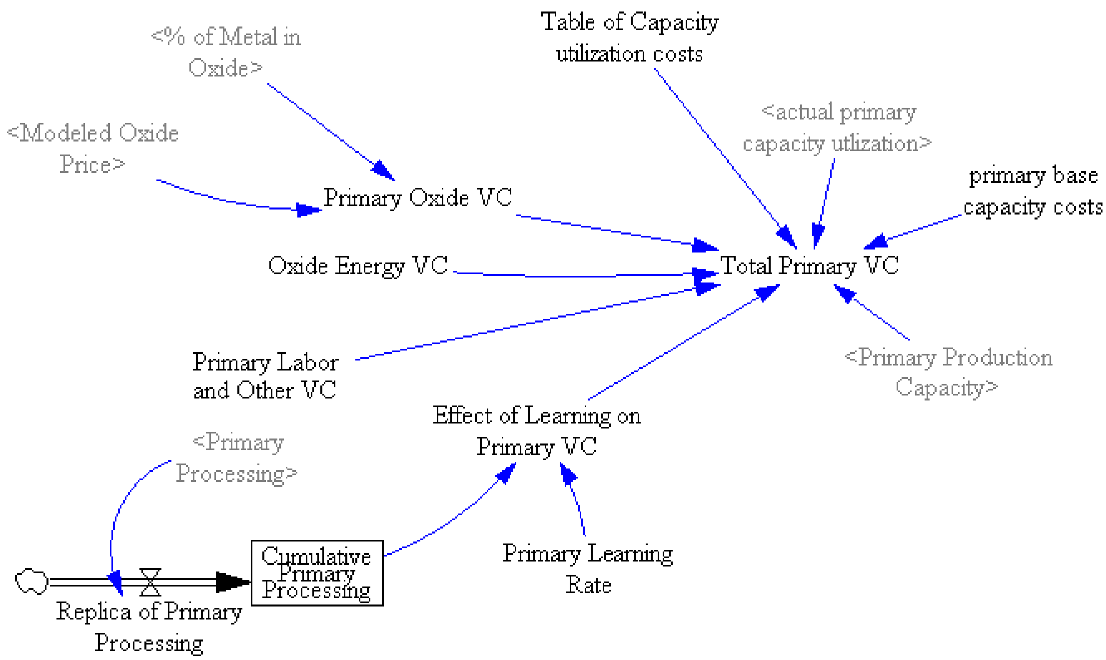

Figure 2-4: Model Structure for Costs

Recycling & Secondary Price

When goods reach the end of their life in the model, some fraction is dissipated, but the remainder is available for recycling. All that is not recycled is disposed of and can not be recovered. The cost of collection to secondary producers is modeled as an increasing function of the recycling efficiency.

Like primary price, the price of secondary (recycled) metal is a modeled as a function of inventory coverage and costs; however, it is also influenced by the price of primary. Xiarchos (2009, 2006) has shown that for several commoditity metals, including aluminum, while primary and secondary prices may not show a consistent relationship over the short term, they do over the long term. That effect was modeled by having the ratio between primary and secondary prices change with inventory coverage. That is, secondary price was set as a fraction of primary price; fluctuations in secondary demand would change that fraction; but, over the long term, the fraction would return to its long-term value. Figure 2-5 presents a simplified graphic of the model structure used to calculate secondary price.

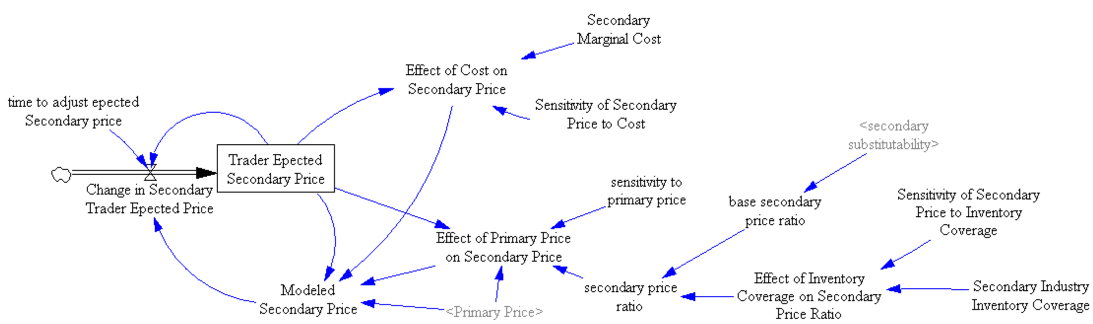

Figure 2-5: Model Structure for Secondary Price

Demand for secondary metal is modeled as a function of both the technical substitutability of secondary for primary metal and the price difference between the two. The particulars of this relationship will be discussed at length in Chapter 3, but suffice it to say that there is a trade-off between the cost savings of using the cheaper secondary metal and the risks of exceeding batch limits for any tramp or alloying elements that may be contained in the secondary metal.

The model's supply of secondary metal is constrained by the quantity of metal containing goods that are being disposed of. The model is a simplification of real life where there may be stocks of scrap metal available to be 'mined' when the secondary price is high enough to make it economical to do so. The costs of recycling are treated as a function of the recycling efficiency—percentage of the metal in the waste stream that is captured for recycling. Recyclers tend to capture the most inexpensive and easily accessible waste metal first; after that is depleted metal becomes progressively more costly to obtain (Tilton, 1999; Slade, 1988). The costs of recycling were modeled as increasing with recycling efficiency.

Capacity & Capacity Utilization

Production capacity, or just capacity for short, is the model term for the upper limit on production for a market actor. To reach the absolute maximum capacity, the firms may have to pay employees overtime, rent additional equipment, or perform other activities that are not cost effective. The cost effective, and therefore desirable capacity utilization is somewhat less than the theoretical maximum. The actual capacity utilization is modeled as depending on two factors: The first being the profit consideration, which was framed as the product 'markup'—the difference between the market price of the product and its variable cost (Sterman, 2000). A higher markup provides greater incentive to produce thus putting pressure on the firm to increase the capacity utilization. The other pressure on capacity utilization is the desire of the firm to maintain its desired inventory level. Based on the costs and benefits of maintaining inventory, the firm will likely have a desired level of inventory coverage that they like to maintain. When this becomes low or high, they will seek to change their production to correct the inventory level. This methodology was adopted from Sterman (2000).

The total capacity is also subject to change. It modeled as being influenced by two factors: One, the capacity utilization; two, the profitability of additional capacity. When capacity utilization is consistently high (or low) the firms have an incentive to increase (or reduce) capacity. Consistently operating away from the optimal capacity utilization suggests that the total capacity is not optimal, and should be adjusted accordingly. Moreover, the firms are assumed to consider the expected future price of their products (based on current trends) and compare those expected returns with the costs of adding capacity and add or reduce capacity accordingly. This methodology was adopted from Sterman (2000) and Alonso (2010).

Inventory

As discussed in section 2.4.1, price is modeled as a function of inventory coverage. As such, price functions as a method of controlling inventory by modifying demand: Inventories get to high, prices go down, demand goes up, and inventories diminish to a desired level. These values were set in the process of calibrating the model. Inventory also serves as method for firms to smooth their production (Pindyck, 1994); in other words, inventory provides a buffer between production and customer orders so that firms do not have to change their production levels with every perturbation in customer orders. Firms are assumed to maintain raw material inventories to provide a buffer on the other end of production—so that perturbations in raw materials deliveries do not interrupt production. While goods orders and some oxide orders are modeled exogenously as discussed in section 2.4.1, all other product orders are modeled exogenously: firms order raw materials at a rate that maintains their desired raw material inventory coverage.

Mining

The mining sector is model like the other market actors, with one key difference: the ore resources must be tracked. At the start of the model simulation the the ore resources identified in Bray (2009) are divided up into bins by ore grade—in order to track the grade of the ore being used. The grade being the fraction of the ore rock that is the desired mineral or metal. The identified resources can be divided into those that are economical to extract and those that are not—termed economic and subeconomic, respectively. What makes the subeconomic reserves not economical is that they would cost more to extract than the current price; in the real world, this could be do to their location, or, as is assumed in the model, that the ore grade is too low. Mining expenses are typically in terms of the quantity of rock extracted, not the quantity of ore Alonso (2010); so, lower quality ores typically have a higher expense per unit of ore, making them less economical. The model used the assumption that mining expenses were on a per-unit-rock basis.

Reserves are depleted through extraction, but they are also supplemented through discovery. Discovery was modeled on the assumption that, as producers have desired inventory, mining firms have a desired static depletion index that they prefer to maintain. This value was set at 30 years. As firm's static depletion indices fall below this level, they increase the effort that they put into discovery. The actual amount and grades of ore discovered are treated as a stochastic function of the discovery effort; in other words, discovery increases with effort, but there actual degree of success is partly random, since firms don't know for certain where stocks of ore reside.

2.5 Implementation in Aluminum

The model was designed to have components common to any commodity metal, so that the results would not be specific to only one peculiar metal. As has been explained in this chapter so far, the model coheres with some theoretical relationships between variables. However, it is difficult to have confidence in the results of a model, unless 'reasonable' inputs produce 'reasonable' outputs—unless the model produces results that correspond to some degree with the reality that the model represents. A difficulty here, is that perfect correspondence with observed results does not make a model perfect; in fact, the process of tuning the model to produce an exact set of circumstances may decrease its ability to behave as the real system does; it may reduce the robustness of the model—its ability to respond to a wide range of conditions as the real system might. For these reasons, the model was implemented with data for a specific commodity metal, aluminum. The model model was then calibrated to historical data, but the exactness of the calibration was balanced by the need to keep the model robust enough to run the range of tests performed in this study.

2.5.1 Choice of Aluminum

Aluminum was chosen as the metal to be modeled for four reasons: First, aluminum is a major commodity metal and can be thought of as representative for that class of metals. Second, the aluminum market has the characteristic that primary and secondary metal are not always perfectly substitutable, which is a necessary precondition for the second experiment—testing the effect of secondary substitutability on primary price stability. Third, due to its wide use, a study of aluminum is interesting in its own right. Fourth, and finally, there is an advantage to modeling aluminum in that there is a lot of data available to provide the initial conditions for the model and to compare the results to.

2.5.2 Simulation Length and Granularity

The simulation begins in 1985 and continues until 2030 for the study on secondary substitutability and until 2035 for the metrics study. The metrics study ran slightly longer so as to provide more data points for calculating the correlations between metrics and primary price. The time step is 1/32 years.

2.5.3 Model Variables

Ore Deposits & Mining Regions

Mining operations are modeled by region. The world's largest bauxite producing regions (as of 2009) (Bray, 2009) were chosen: Australia, Brazil, China, Guinea, and India. Data for the ore grade and size of their deposits was obtained from a USGS report for 1985, which was chosen as the year to begin the simulation (Patterson, 1986). Mining firms were aggregated with alumina producers. So the output product from the mining regions is alumina (aluminum oxide).

Energy

Energy costs for primary production are based on a 15.6 kWh/kg-Al (Choate and Green, 2003) energy requirement for electrolyzing aluminum, and an industrial energy cost of $0.051/kWh, which is the average industrial electricity cost in the US from 1985 to 2009 (EIA, 2010). For processing scrap, (Choate and Green, 2003) estimates the energy requirement for casting secondary ingots as 2.5 kWh/kg, but the EcoInvent LCA database quotes 6.6 kWh/kg (Swiss Centre for Life-Cycle Inventories, 2010). In this model an intermediate number of 5.5 kWh/kg was used, this was based on an initial approximation. The energy requirements for processing primary and secondary for producing goods were approximated at 0.75 kWh/kg, which is slightly lower than the energy costs for casting primary (1.01 kWh/kg), but well above the theoretical minimum for casting 0.33 kWh/kg (Choate and Green, 2003). Again these numbers were approximations that satisfied the dual requirements both being within a reasonable range themselves, and producing model behavior that approximated historical behavior.

Delays & Desired Inventory

The delays and desired inventory levels for the market actors are summarized below.

Variable	Oxide	Primary	Secondary	Goods
Desired Inventory Coverage	3	1/2	1/2	1
Desired Raw Material Inventory Coverage	N/A	1	1/12	1
Capacity Acquisition Delay	1	1	1	1
Time to Adjust Capacity	4	3	1	5
Time to Adjust Cap. Utilization	1	1/4	1/4	1/4

Table 2.1: Delays and Desired Inventory Coverages (all values are in years)

The values in Table 2.1, were selected in the model calibration as giving levels of primary price and production in the range of historical values. There selection was constrained by limits of what are considered 'reasonable' values in the system dynamics literature (Sterman, 2000), and by what characteristics are known about the aluminum industry. As an examples of the later, it is known that adding mining capacity takes at least 4 years to bring capacity online (Rees, 1985); this is captured in the mining capacity acquisition delay, which is the delay between receiving a signal an acting on it, and the time to adjust capacity, which is the time required for a change to be fully realized. Also, it is expected that secondary capacity is added more quickly than primary capacity (Alonso, 2010).

2.5.4 Calibration

The model was roughly calibrated to historical aluminum price and primary production, as obtained from the USGS (Buckingham, 2010). The modeling goal was not to completely reproduce the real world events, but simply to provide a coherent framework for testing some hypotheses about the relationship between different variables in a stylized aluminum market. The model calibration goal was simply to show that the major model variables, primary price and production, were in the same ballpark as historical values. This can be seen in Figure 2-6. Note that the largest discrepancies occur in the first ten years of running the model. This is due to the difficultly of setting the initial model conditions. As the model is dynamic, with hundreds of endogenous variables, the initial relationships set may may be far from equilibrium. As this is an almost unavoidable consequence of this type of modeling, the chosen solution was to disregard the first 10 years of data. The model take time to 'settle.' For that reason in both of the two experiments the first ten years of the model run were not used for any experiments.

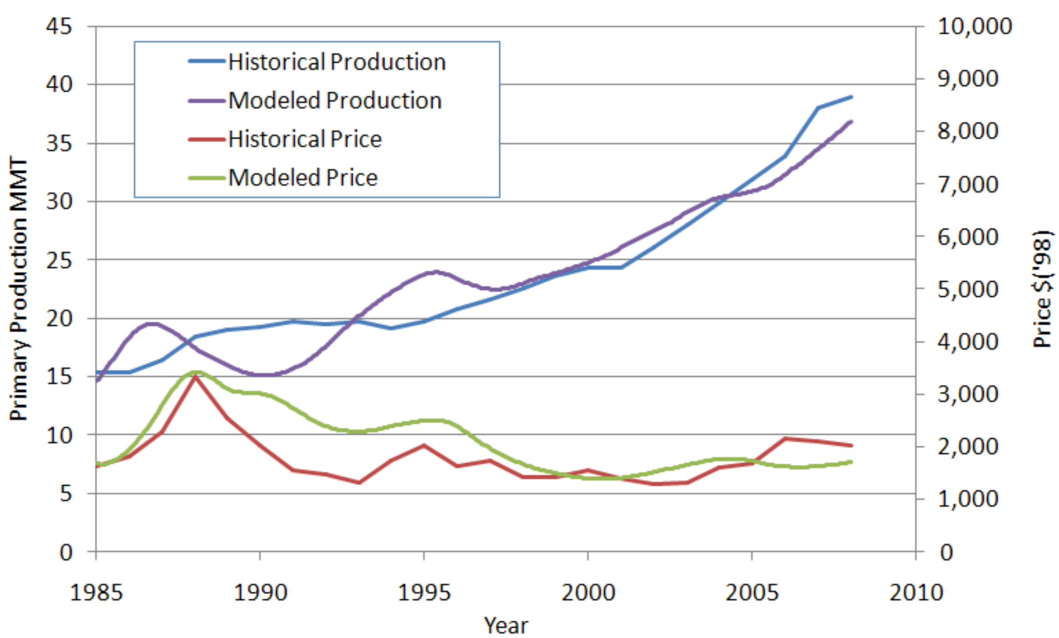

Figure 2-6: Model Validation for primary price and primary production. Note that the model took time to equilibrate: the first ten years of data have the worst fit and were not used in the experiments performed in this thesis.

Chapter 3

Secondary Substitutability & Price Volatility

3.1 Introduction

This chapter explores how the degree to which primary metal can be substituted for secondary metal—secondary substitutability—affects the stability of the primary price.

3.2 Problem Statement

Rees (1985) and Alonso (2010) show that larger secondary markets can stabilize price. One would expect that higher secondary substitutabilities would lead to larger secondary markets by increasing the demand for secondary material. Moreover, since a greater ability to substitute for another material also stabilizes price (Slade, 1988), it would be expected that being able to substitute secondary metal for primary would stabilize price. But, this does not necessarily follow: since secondary metal is generally cheaper than primary, one would expect firms to be using as much as was available–right up to their technical limit—the technical limit being the maximum percentage of secondary that they could use. A change in the primary price would not likely lead to a change in the amount of secondary used, because it would still be desirable to use the maximum amount of secondary. Only if primary became cheaper than secondary (unlikely) would the ratio used change. In other words, irrespective of higher secondary substitutability, there may be no substitution for secondary due to price changes, but there would probably be a larger secondary market.

There is evidence, that even for a given technical limit, the amount of secondary used (relative to primary) is actually sensitive to the price difference between the two metals (Fowler, 1937). The following analysis will explore how the secondary substitutability affects the degree of substitution between primary and secondary and, more broadly how secondary substitutability affects the stability of primary price.

The question of how secondary substitutability affects primary price stability will be addressed through an experiment with the system dynamics model, where a perturbation is introduced that produces an increase in price. That perturbation is a suddenly imposed shortage of primary metal. The size of the price increase for different levels of secondary substitutability will be tested. So, formally stated, the question explored in this chapter is:

> *How does secondary substitutability affect the stability of the primary price in the presence of an exogenously introduced shortage of primary metal?*

3.3 Prior Work

Prior work by Alonso (Alonso, 2010) showed that, in a platinum market model, higher recycling rates led to smaller price spikes in the presence of an imposed shortage of primary metal: in other words, when there was a primary metal shortage, primary prices increased, but less so when the recycling rates were higher. This can behavior is shown in Figure 3-1.

3.4 Secondary Substitutability

Secondary substitutability is defined, for the purposes of this study, as the degree to which secondary metal can be substituted for primary metal in the final aluminum goods. For example setting the secondary substitutability to 50% would indicate that 50% of the total aluminum in goods could be secondary metal. For alloys, secondary substitutability is a function of the limits of the alloy for other elements contained in the secondary metal. Specifically, it is driven by the most limiting element. These other elements may be alloying elements purposely added to the scrap when it was a new product, or they may be 'tramp' elements that were picked up in the recycling process. In this analysis, the amount of tramp elements in the secondary metal is not modeled explicitly. Rather the secondary substitutability is given as the ratios that follow from the levels of tramp elements in the

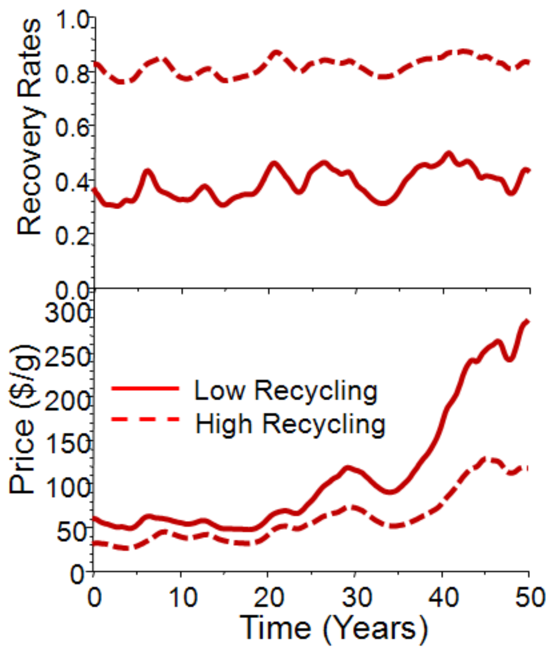

Figure 3-1: Effect of recycling on price stability in a platinum market model. Recycling efficiency (fraction of metal recovered when goods are disposed of, labeled as recovery rate in the plot) is displayed at top; metal price is displayed at the bottom. Prices are lower with higher recycling efficiencies.

scrap and the limits in the alloys/goods that use them.

3.5 Cost-Error Trade Off

In choosing the amount of secondary metal to use, firms must trade-off between savings from using more of the cheaper scrap, and costs incurred from missing batch targets. There is evidence in literature that for iron, as the primary price increases, a greater percentage of scrap is used; and that this increase in percentage per change in price diminishes as primary prices get higher (Fowler, 1937). This suggests that the cheaper secondary metal is relative to primary, the more secondary metal is used, up to some limit.

A way of visualizing the trade off between the risk of missing batch targets and having a high secondary metal composition in an alloy can be seen in Figure 3-2. In any supply of scrap, there will be a distribution of tramp/alloying element compositions from any one batch to the next. The mean of that distribution may or may not higher than the limit for the alloy, but some portion of that distribution is likely to be, so that the scrap will have to be diluted with primary metal or other scraps that are not limited by the same element

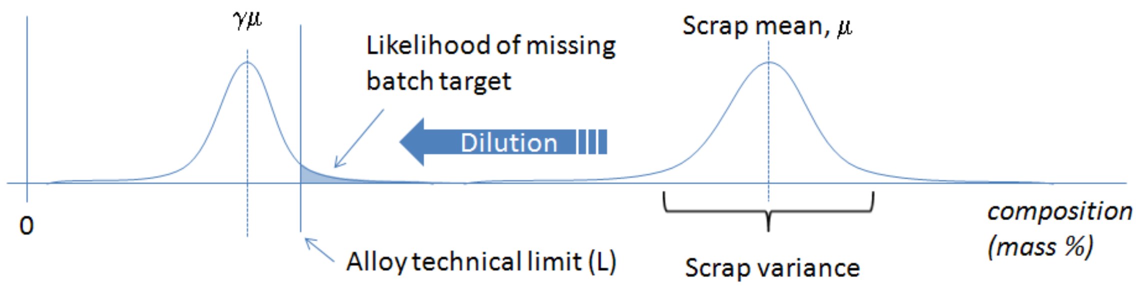

Figure 3-2: Cost-Error Trade-off: Dilution of scrap metal with primary erodes savings (less of the cheaper scrap is used per unit of final product), but it brings more of the distribution of the alloy/tramp element below alloy limit—lowering the expected value of the error rate.

such that only an acceptable 'tail' of the distribution exceeds the batch limits (Gaustad, 2009). So the trade-off is more dilution, less errors; less dilution, more savings.

3.5.1 Optimal Rate of Recycling

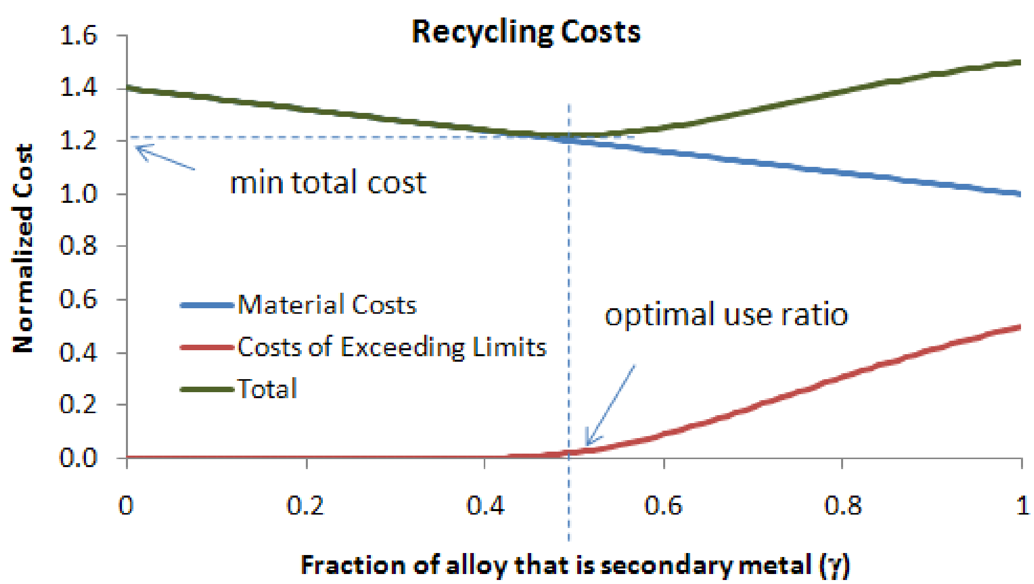

Figure 3-3: Cost-Error Trade-off: Cost is normalized to secondary price. Primary price is set at 1.4 times secondary price, which would not be atypical for aluminum.

Figure 3-3 shows the cost-error trade-off in cost space. At high dilutions (low fractions of secondary metal), material costs are relatively high, because most of the metal used in primary, which is assumed to be more expensive than secondary. As more secondary metal is used, the material costs go down. The material cost M can be expressed as a function of the fraction of the alloy that is secondary metal γ and the primary and secondary metal prices, P and S, respectively. The material cost M decreases linearly with γ.

$$M = (1 - \gamma)P + \gamma S \tag{3.1}$$

If an alloy producer misses a batch target for composition, it may be able to correct the error by adding more aluminum, but the limited capacity of the furnace limits how much dilution is possible. It also may be possible to remove some of the offending element, but that is very time and energy intensive and may be more expensive than the final option: if a batch is out of specification and not fixable by one of the other methods, the whole batch will typically be cast as large castings known as 'sows' and re-used later as if it was scrap (Schmitz, 2006). There are a number of costs associated with this. There is the wasted labor, energy, and overhead of producing the batch—which will eventually have to be re-melted and incur those same costs. There are also costs associated with the missed production target, maybe a customer order is late or business is lost. And, finally, there is a storage cost to holding the sows in inventory. The costs associated with missing production targets are likely non-linear, since missing one or two batches here and there would likely be less disruptive than the extreme of missing every target; however, the other costs of missing the target, lost labor, energy, overhead, and inventory costs are likely to be constant to first order. For the purposes of this study, the costs associated with missing a batch target are assumed to be constant. If that is so, then the per unit cost of error B is the expected value of batch error costs, which, for C being the cost of a batch error is simply

$$B = E[C] \tag{3.2}$$

where $E[\cdot]$ is the expected value operator. If we assume that the mass composition of tramp elements is normally distributed, which is not unreasonable (Gaustad, 2009), then Equation 3.1 becomes C times the complement of the normal cumulative distribution, Q.

$$B = C \cdot Q(x) = C \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{\frac{-t^2}{2}} \, \mathrm{d}t. \tag{3.3}$$

$$x = \frac{(L - \gamma\mu)}{\sigma} \tag{3.4}$$

where μ is the mean element composition of the scrap, and σ is the standard deviation.

The expected error costs B increase with γ while the material costs decrease with γ. The sum total of the two, the total costs T, will at first decrease with γ and then, depending on the size of B, increase. This means that so long as error costs begin to increase with γ faster than total costs decrease with material costs, there exists a γ at which costs are at a minimum. It is important to note that if the costs of errors are always small, the minimum-cost gamma will be $\gamma = 1$, or full recycling. There are of course situations and industries where this is always the case and there are no restrictions on recycled metal use. However this is not expected to be true industry-wide for any metal. It was probably not true in the case of iron cited before, where the percentage of scrap used changed as a function of price (Fowler, 1937). If the smelters were using all the scrap metal that they could, their use ratio would not have changed with price.

In order to put scenarios with different technical limits on a common footing, to make B a function of γ, and to avoid having to track mean and standard deviation, it is useful to make the following substitutions. The secondary substitutability will be given by λ.

$$\lambda = \frac{L}{\mu} \tag{3.5}$$

$$c_v = \frac{\sigma}{\mu} \tag{3.6}$$

$$x = \frac{\lambda/\gamma - 1}{c_v} \tag{3.7}$$

Putting Equations 3.1, 3.3, and 3.7 together gives the total cost TC function.

$$TC = M + B = (1 - \gamma)P + \gamma S + C\frac{1}{\sqrt{2\pi}}\int_{\frac{\lambda/\gamma - 1}{c_v}}^{-\infty} e^{\frac{-t^2}{2}}\,\mathrm{d}t. \tag{3.8}$$

By setting the first derivative of this function with respect to λ to zero the local minima and/or maxima can be obtained.

$$\frac{\partial TC}{\partial \gamma} = S - P + C\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}(\frac{\lambda/\gamma-1}{c_v})^2}\frac{\lambda}{c_v\gamma^2} = 0 \tag{3.9}$$

The second derivative of Equation 3.8 with respect to γ is not always positive. The rightmost term in the first derivative (Equation 3.9) is the normal probability density function (PDF). Because of the bell-shape of the PDF function, the slope of the PDF function will be equal to $(P - S)/C$ at two different points; the first being a local minimum the second a local maximum. This is demonstrated in terms of the total cost function and its derivative in Figure 3-4. In example A, the derivative of the cost function equals zero only once over the range of γ equals zero to one. However, in example B, due to a lower cost of batch errors, the derivative of TC with respect to γ equals zero twice on the range γ equals zero to one: the first being a minimum the second a maximum. Caution should be exercised when using Equation 3.9 to be sure it is finding a local minimum. However, it is expected that batch errors are typically very low. Anecdotal evidence suggests they occur anywhere from every 1/10 to 1/200 batches—that would put firm's operating range in the tails of the PDF function, where the solution to Equation 3.9 is always a minimum.

This logic for determining desired secondary use was included in the system dynamics model, where Equation 3.9 was solved for current prices. As the solution to that equation deviated far from zero, increasingly aggressive corrections are made to the desired secondary use rate, with a delay. It was not possible to solve directly for the optimal use ratio in the system dynamics software, but the approach used is effectively equivalent; it actually has an advantage in that it may more accurately capture the human decision making process—small deviations from optimality are overlooked as not worth the trouble of making a change, larger deviations drive more aggressive changes, with the potential to overshoot the mark due to delays between the signal triggering action and the response.

3.6 Experimental Set-Up

To answer the question how does secondary substitutability effect price volatility, perturbations were introduced to the model that forced volatility, which took the form of a large price spike; and the height of this price spike was compared for a range of secondary substitutabilities in order to determine the relationship between the secondary substitutability and the height of the price spike. Formally, the height of the price spike was the maximum

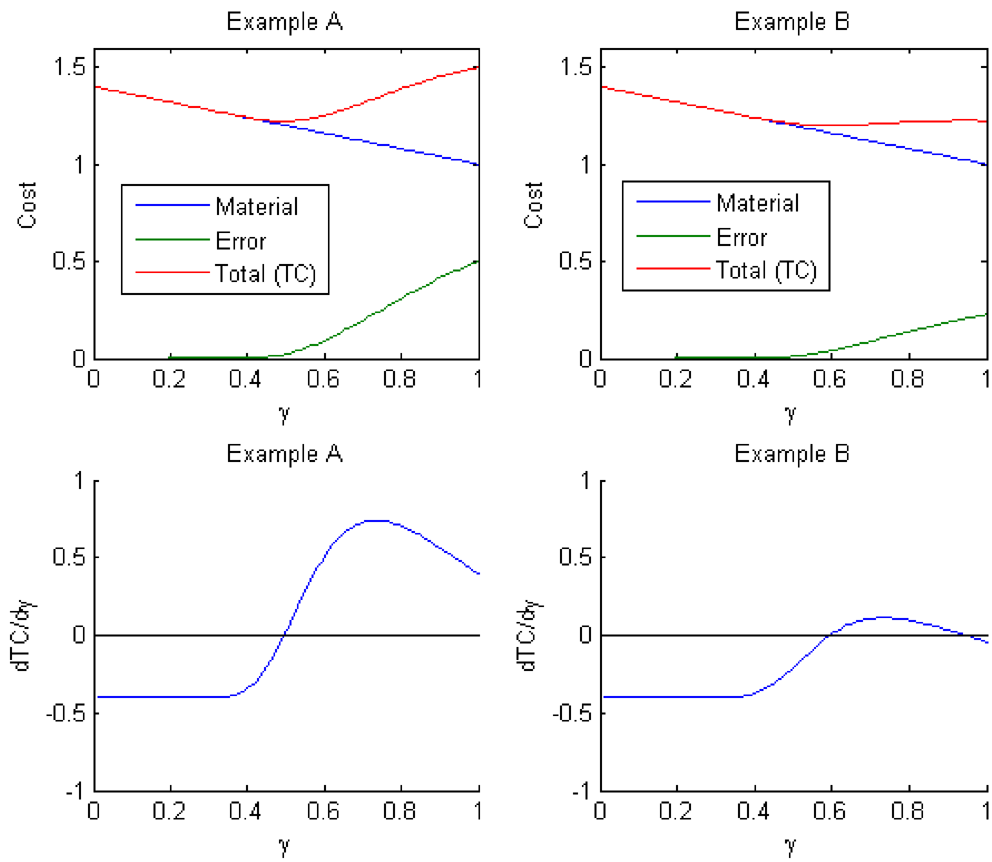

Figure 3-4: Cost Function: Local Minima and Maxima. In example A, which is considered more likely, the derivative of the total cost function equals zero at the minimum value of TC. In example B, both a minimum and a maximum are found, with the minimum occurring first (lower γ).

price within 10 years of the perturbation minus the price at the time of the perturbation. The term 'relative price spike height,' which is used in the results in Figure 3-6, is the price spike height divided by the price at the time of the perturbation.

The price spike was introduced by simulating a permanent shutdown of 1/3 of all primary production, thus introducing a shortage of primary. Because price is set by inventory coverage, shutting down a fixed fraction of primary production will produce the same effect regardless of the size of the primary market.

To explain this experimental design, it is important to remember that markets with higher secondary substitutabilities will exhibit relatively larger secondary markets and relatively smaller primary markets. If the price spike were to be initialized imposing the shutdown of a fixed amount of primary capacity, the effects would be strongly dependent on the size of the primary market. Because the modeled price is set by inventory coverage, the imposition of a shutdown of a fixed proportion of the primary market should have comparable effects on primary price, independent of primary market size. An example of a price spike caused by such a perturbation is shown in Figure 3-5.

While this experimental design removes the effect of market size, the dynamic nature of the model presents another difficulty: the time at which the price spike occurs will affect its height. By design the model is never in equilibrium. Inventory coverage is usually slightly higher or lower than the desired inventory coverage. If the perturbation occurs when the coverage is already low, then the price spike will be higher than if it happened to be high (because firms would have additional inventory to fall back on before having to enter the market for new material.) To limit the influence of time, price spikes were introduced across the entire range of model times with the exception of the very beginning and end. That way price spikes occur at at times of low and high inventory with equal likelihood, and the results can be averaged. As mentioned in Chapter 2, the first 10 years of the model run are thrown out; the last years were not used because they would not leave enough time to allow the price spike to develop.

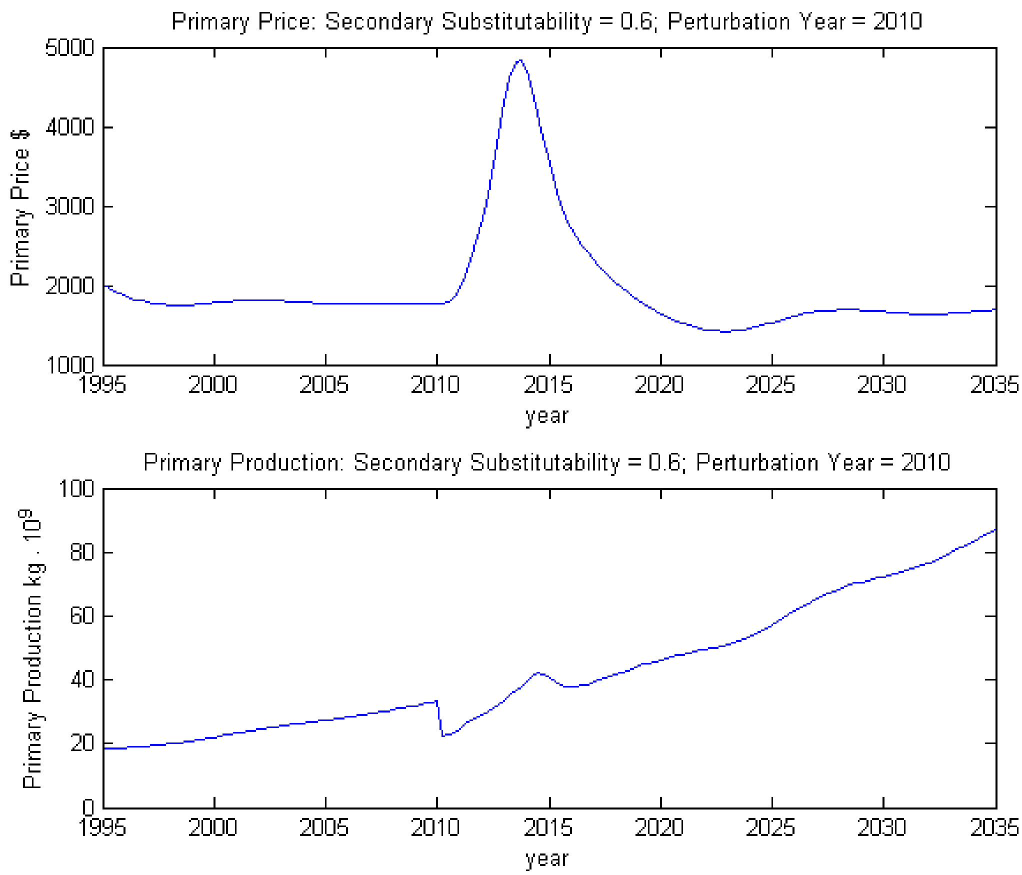

Figure 3-5: Sample Result: Primary Price and Primary Production, with a perturbation at year 2010.

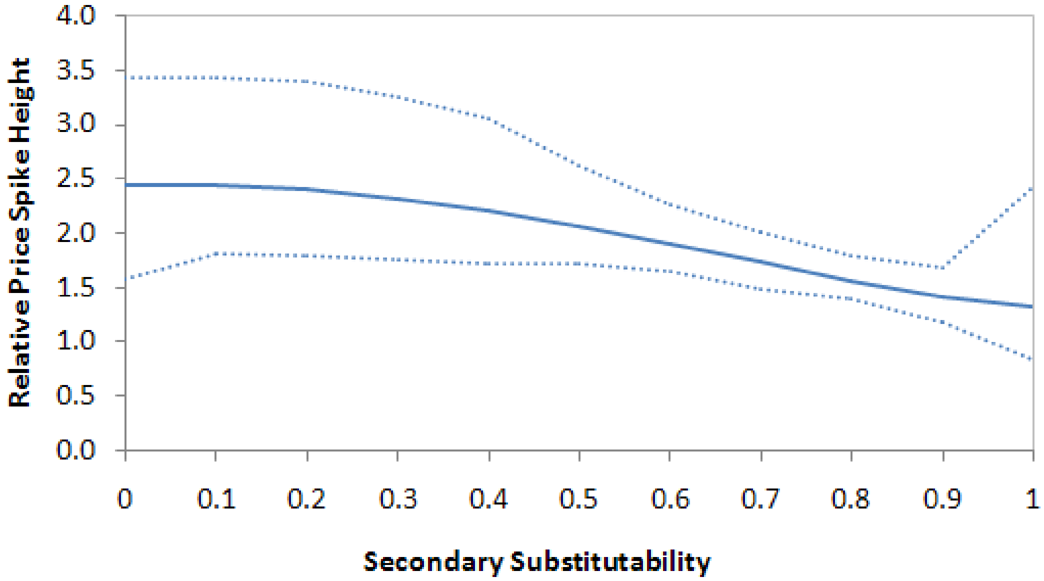

Figure 3-6: Main Result: Effect of Secondary Substitutability on Primary Price. This figure shows the average height of the price spike relative to the price at the time of perturbation for 200 different perturbation times, evaluated for secondary substitutabilities from 0 to 1 at increments of 0.1. The solid line is the average. Hashed lines represent minimum and maximum values.

3.7 Results

3.7.1 Main Result

The main result of the experiments is shown in Figure 3-6, which shows the average height of the price spike relative to the price at the time of perturbation for 200 different perturbation times. With greater secondary substitutability, the price spikes are smaller. At the rightmost end of the graph, where secondary substitutability equals 100% we see what is essentially Alonso's result (Alonso, 2010). The range of price spike heights also decreases with increased secondary substitutability, with one exception: when secondary substitutability is equal to one, there is an increase in the maximum values of the price spike. This result was caused by the price spike heights from perturbations early in the model. Perturbations at later times were smaller, hence the observation that the average and minimum values of the price spike are decreasing. The higher price spikes earlier in the model were very likely due to the way the model 'settled' from its initial conditions, as discussed in section 2.5.4. It seems that the model settled a little differently at that higher secondary substitutability.

3.7.2 Causal Linkage

The next question is why are price spikes smaller with greater secondary substitutability. This can be explained using the graphs in Figure 3-7. When secondary substitutability was higher the price spike height was smaller because, as the price began to rise, primary orders dropped off more quickly. This means demand for primary is lower and the inventory coverage will increase. The reason that primary orders dropped off more quickly in the higher secondary substitutability runs is that the firms were able to switch to using more secondary metal. This last detail is especially important. Without it, one might assume that the reason that the price spike was smaller for the high secondary substitutability cases was because the secondary market was relatively larger and the primary market was relatively smaller. In other words, with a smaller primary market there is a smaller hole to fill. However, what the bottom graph in Figure 3-7 is showing is that, even though the hole to fill might be relatively smaller, the response of the firms in the higher secondary substitutability runs is *more* drastic not less. Despite their secondary use ratio being higher to begin with they change it more, percentage-wise, than do the firms in the lower secondary substitutability runs. So, the smaller price spikes seen with higher secondary substitutability are caused by something other than the larger secondary

3.8 Discussion

The derivative of the total cost function (3.9) can be used to explain why a greater relative shift towards secondary usage occurs with higher secondary substitutability. Equation 3.9 can be rearranged such that all of the price and cost terms are on one side and the substitutability related characteristics are on the other.

$$\frac{P-S}{C} = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(\frac{\lambda/\gamma - 1}{c_v})^2} \frac{\lambda}{c_v \gamma^2} \tag{3.10}$$

In this arrangement, there is no simple closed-form solution for the variables on the right hand side but, for any given coefficient of variation, a plot like Figure 3-8 can be made. The vertical axis is the optimal (minimum cost) secondary use rate γ. The horizontal axis is the difference between primary and secondary price, divided by the costs of a batch error. This plot presents the optimal secondary use rate γ as a function of the prices of primary

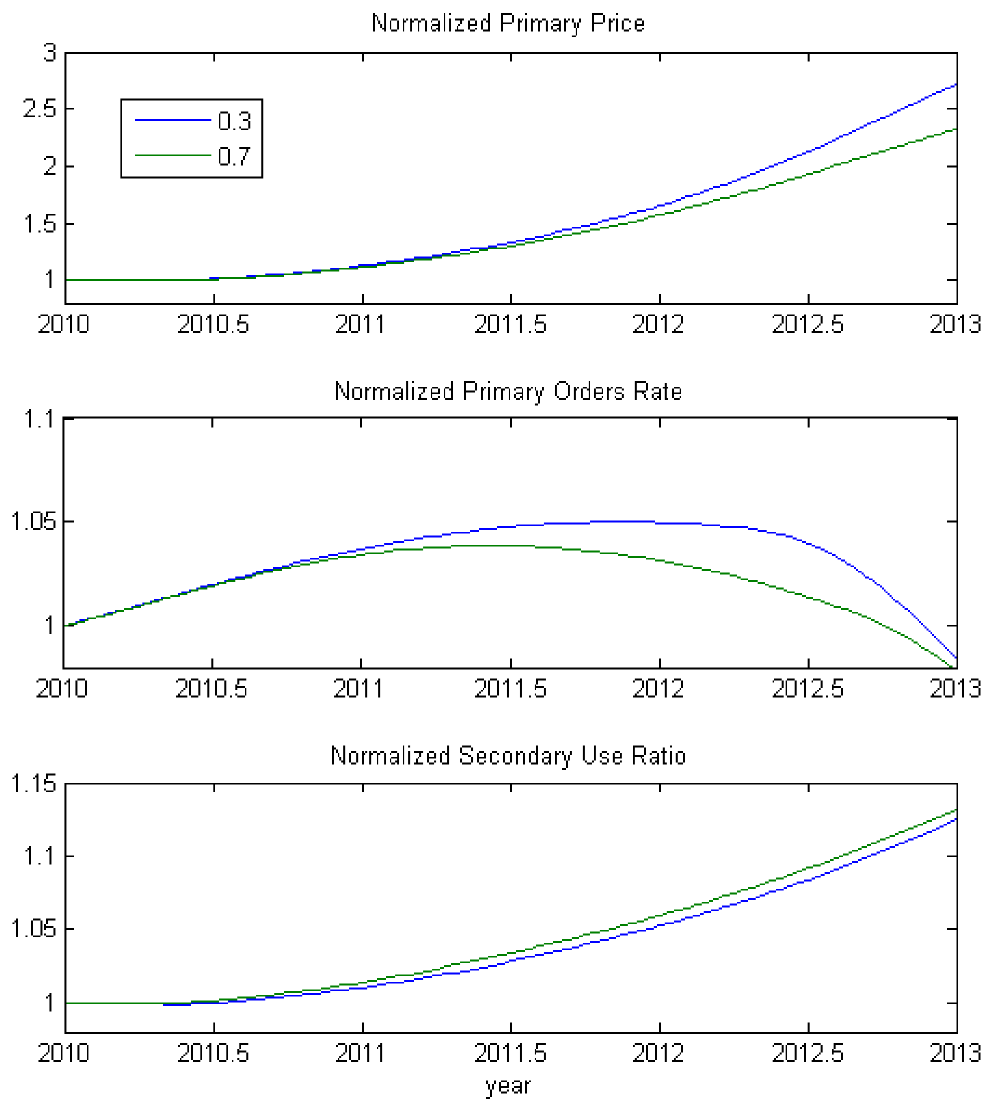

Figure 3-7: Causes of result: For two levels of secondary substitutability, 0.3 and 0.7, these three charts show why the price spike is smaller. The price (top graph) is lower with higher secondary substitutability (0.7) because primary orders (middle graph) drop off more quickly in response to the price increase. That is because the in the higher secondary substitutability run, the firms switched to using more secondary metal (bottom graph), relative to what they had been using before.

metal, secondary metal, and the costs of batch errors. Figure 3-8 shows three curves, one for a secondary substitutability (λ) of 0.3, one of 0.7, the other 1.0. At any point on the x-axis, higher secondary substitutability is always associated with higher optimal recycling rates. Moreover, with higher secondary substitutability, the slope of the optimal curve is steeper. All else being equal, an increase in the primary price, leads to a rightward shift along the x-axis; for any change to primary price, the change in the amount of secondary used will be larger for higher secondary substitutabilities. For instance, going from 0.5 to 1 on the x-axis changes the optimal recycling rate from about 0.5 to 0.6 for a secondary substitutability of 1. For secondary substitutability of 0.3, the same change on the x-axis leads to a change from a recycling rate of about 0.125 to 0.14. The change at the higher secondary substitutability was larger both in absolute terms and as a percentage change from the original value. This could be described as the higher secondary substitutability firms exhibiting a greater price elasticity of substitution.

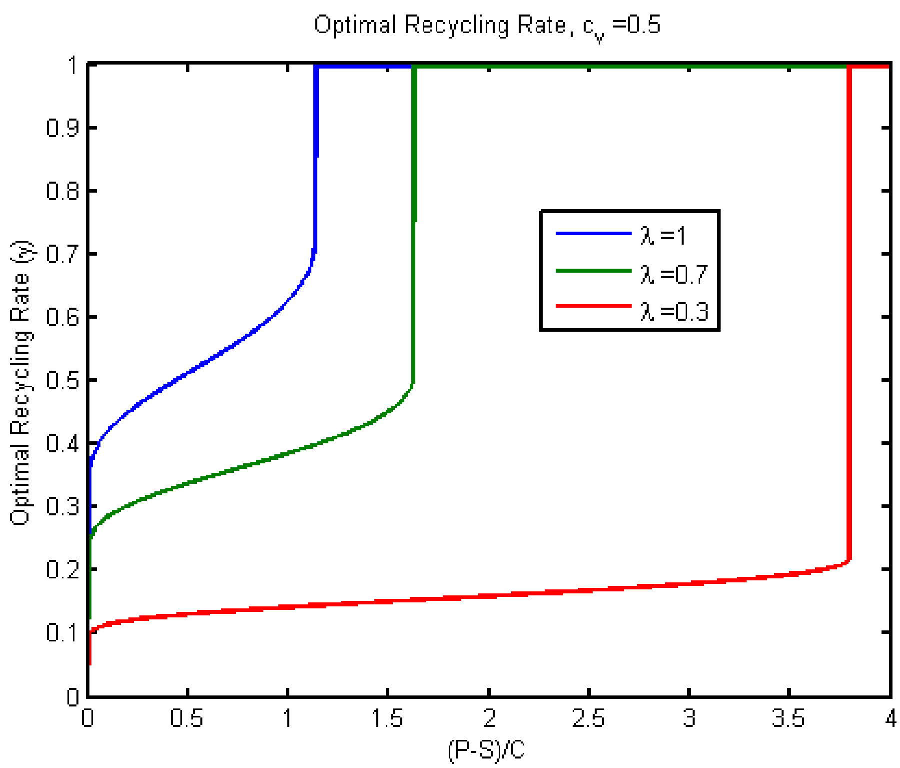

Figure 3-8: Optimal secondary use rate as a function of prices and costs. The y-axis variable γ is the optimal secondary use rate at a given primary price (P), secondary price (S), and cost of batch error (C).

If the optimal secondary use rate curves shown in Figure 3-8 are divided by their respective technical limits, an additional interesting result can be seen. At higher secondary substitutabilities (higher technical limits—the terms are interchangeable in this discussion) the optimal secondary use rate represents a larger portion of that technical limit. So, if a firm were to increase its technical limit, it would increase the amount of secondary it used by more than the percentage by which it increased the technical limit: it would receive gains both from the higher limit and the higher portion of that limit that was optimal for it to use.

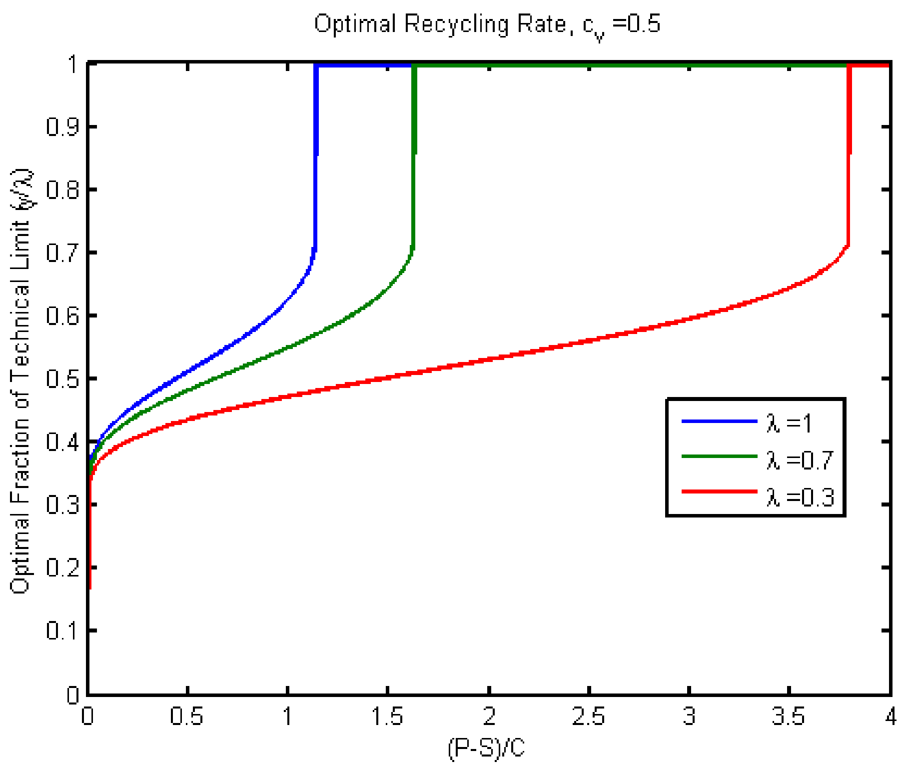

Figure 3-9: Optimal fraction of the technical limit as a function of prices and costs. The y-axis variable γ/λ is the optimal fraction of the technical limit at a given primary price (P), secondary price (S), and cost of batch error (C).

To tie Figures 3-8 and 3-9 back to the modeling results: the simulations suggest that firms with higher secondary substitutability made proportionally larger adjustments to their secondary use ratios for any given change in price. This larger change in secondary use meant a larger drop in primary orders, which helped slow the price increase of the primary metal—shortening the price spike.

The benefits of the higher secondary substitutability are on two levels. As the results show, there is an industry-wide benefit to higher secondary substitutability, in that the price is less volatile in the event of a primary supply disruption; but there is also a benefit to individual firms. Even if only one firm had a higher secondary substitutability, it may not have much effect on the market price (depending on its market share,) but its greater elasticity of substitution would mean that it would allow it to use less of the more expensive primary metal in the event of a primary price increase.

Some notes on the form of Figure 3-8: The curves are monotonically increasing until some x-axis value, where the curve becomes vertical. For all x-axis points on and to the right of this vertical portion of the curve, the optimal recycling rate is 100%. Intuition for this discontinuity can be provided by an examination of Figure 3-4. Lower costs of batch errors 'pull' the location of the minimum cost to the right, but eventually, at a low enough cost (or high enough P-S) the minimum cost occurs at the rightmost extent of the total cost curve, at a recycling rate of 100%—a corner solution to the optimization problem, as it is known. However, as mentioned earlier in the chapter, due to the fact the batch errors occur with relatively low frequency, it is expected that the costs of batch errors are relatively high relative to the difference between primary and secondary price—meaning that most firms would operate on the monotonically increasing portion of the curve, not past the discontinuity; otherwise, they would be recycling at 100%. Now, of course, some firms do use only recycled metal, but this discussion pertains to firms that are bound by technical limits of the amount of secondary that can be used.

The secondary substitutability is driven by desired material properties of the metal product. Based on those properties, there are limits to the various other elements that can be in the alloy or product. However a degree of conservatism is likely built in, and there is some flexibility in how those results can be obtained (Peterson, 1999). This suggests that firms may have the option to increase the secondary substitutability of their products. For firms constrained by a technical limit on the maximum secondary metal they can use, relaxing those limits may increase price stability, or at very least, decrease their exposure to increased primary prices by giving them more flexibility to substitute away from it. Increasing secondary substitutability may also have other benefits. As the next chapter will discuss, increased secondary substitutability has an effect on the ability of other scarcity risk metrics to provide information on movements of the primary price.

Chapter 4

Informative Metrics

4.1 Introduction

This chapter presents an exploration into how well, and under what market conditions, certain scarcity metrics can provide information on the movement of primary metal price in the aluminum simulation. The ability to 'inform' is expressed in terms of the correlation coefficient between each of the metrics (at a range of time lags) and the primary price.

4.2 Problem Statement

As discussed in Chapter 1, price volatility can be detrimental to firms. This creates an incentive for firms to know when price may be changing, so that they can plan accordingly. Price can be a measure of scarcity, so perhaps other measures of scarcity may act as a leading indicator for price change, and provide some information on the motion of price. To be clear, it is probably impossible to predict the motion of a price with any degree of fidelity, but it may be possible to say that price may be moving in in a general direction over a broad time frame. For example, if the mining rate were to increase, and the reserve size remained unchanged, the static depletion index would drop. The increase in the amount of metal mined might provide the conditions for an increase in primary production capacity; leading to a decline in price. So, a change in the static depletion index preceded a change in price. However, the degree to which this would be true would depend on other market conditions. For instance, if the primary producers were slow to add capacity, this increase in supply in mined ore, would translate into an increased supply of primary more slowly, and price may not decrease by much. To what degree to scarcity risk metrics correlate with future changes in price? Under what conditions? These are the questions that this chapter addresses.

Formally stated, the problem addressed in the model experiment presented in this chapter is.

> *To what degree to scarcity risk metrics correlate with future primary price?*

With a related corollary:

> *How do market conditions affect the correlation (if any) between scarcity risk metrics and future primary price?*

4.2.1 Selection of Metrics

Of the metrics listed in Table 1.1, not all were explored in this analysis. Geographic structure was excluded because it is as much a qualitative variable, with many dimensions (politics, transportation, infrastructure,) as a quantitative one. The institutional structure of supply was explored through the Herfindahl Index, but not the institutional structure of primary supply or demand. This was done to simplify the model. Ore grade was a casualty of the choice of aluminum as the metal to be modeled. Current reserves of aluminum consist of large bodies of relatively homogeneous quality ore; ore grade would then be expected to change little as reserves are drawn down. Likewise, the relative rates of discovery and extraction were excluded, because of the massive size of the aluminum reserves, it was assumed that not much effort would be put into discovery, making that less meaningful. Each metric tracked, increases the size of the results data set. To simplify the analysis, exponential index of depletion was not modeled, because it is similar to the static depletion index.

Modeling the remaining metrics for aluminum required some degree of interpretation. For instance, time to peak production could not be modeled explicitly and was modeled by a proxy, as will be explained. The implementation of the scarcity risk metrics in the context of the aluminum model is explained below.

- **Alumina Marginal Cost:** This is variable cost of the highest cost alumina producing region. It is not a value that is likely to be available to the general public, but may be accessible to industry insiders.

- **Herfindahl Index:** This is a measure of market concentration, given by $\sum_{i=0}^{n} s_i^2$, where s_i is the market share of firm i (Hirschman, 1964). Low values are indicative of low concentration–many smaller firms. High values are indicative of higher concentrations, with a value of 1 corresponding to a monopoly. Despite its name it was originally invented (in a slightly different form) by (Hirschman, 1964).

- **Alumina Price:** The market price of aluminum oxide paid by exogenous purchasers and primary aluminum smelters.

- **Normalized Mining Acceleration:** This value is the the second derivative of the mining rate, normalized to the mining rate (so as to be expressed as a percentage.) This is the closest proxy to time to peak production that could be approximated in Vensim®. Hubbert's model for peak oil production assumes that production should slow down until it reaches a peak and then decline. A declining mining acceleration would suggest the imminent approach of peak production (Cavallo, 2004).

- **Primary Marginal Cost:** This is variable cost of the highest cost Primary producer. Like Alumina Marginal Cost, it is not a value that is likely to be available to the general public, but may be accessible to industry insiders.

- **Recycling Efficiency:** The fraction of total waste stream that is recycled.

- **Recycling Rate:** The fraction of total aluminum production that is recycled metal.

- **Static Depletion Index:** The known reserves divided by the extraction (mining) rate. It is expressed in units of years, and represents the number of years-worth of metal remaining at the current extraction rate.

4.2.2 Market Characteristics

It was hypothesized that market characteristics that affect the volatility of price will also affect the ability of scarcity risk metrics to provide information on price. For that reason, the effect of the following four characteristics on the ability of scarcity risk metrics to provide information on primary price were explored: the primary metal and alumina capacity acquisition delays, the price elasticity of goods demand, and the secondary substitutability. All of those characteristics have been shown to affect the stability/volatility of price.

Metric	units	Low	*Level* Medium	High
Alumina Cap. Acquisition Delay	years	1	3	7
Primary Cap. Acquisition Delay	years	1	3	7
Goods Demand Elasticity		0.3	0.8	1.3
Secondary Substitutability		0.2	0.5	0.8

Table 4.1: Factorial Analysis Set Up: Market characteristics to be varied along with the three levels at which they are modeled.

There number of characteristics was limited to four in order to keep the results set to a manageable size. In order to establish the effects of the characteristics, a 3-level full factorial analysis was run. Running a factorial analysis provides a better estimation of the effects of each characteristic than holding the other variables constant and varying each characteristic one at a time (Fisher, 1971). Two levels would be the minimum for an analysis of the effects of changing the characteristics, but three levels were chosen because, running three levels provides information as to the linearity and/or monotonicity of the effect. For instance, if the effect of the 'medium' level is not between the 'low' and the 'high' levels, one cannot say that there is a consistent effect from low to high—a two level test might have obscured this result. To add a 5th characteristic would have increased the size of the factorial analysis by 3 times. To keep the results set manageable, the number of characteristics explored were limited to four.

- **Alumina Capacity Acquisition Delay:** The time elapsed between alumina producers decision to add capacity and the adding of capacity.

- **Primary Capacity Acquisition Delay:** The time elapsed between Primary producers decision to add capacity and the adding of capacity.

- **Goods Demand Elasticity:** Elasticity of demand of aluminum products. The simulation models demand through a constant elasticity of demand function of the form $Q = A \cdot P^{\epsilon}$, where Q is demand, A is a constant, P is the Price of the Aluminum product, and ϵ is the elasticity of demand.

- **Secondary Substitutability:** This is defined as the degree to which secondary metal can substitute for primary metal: it is the mass fraction of total aluminum produced that is secondary metal.

Each of those characteristics was modeled at low, medium, or high levels as shown in Table 4.1. A complete description of the experimental set up will be provided in section ss:setup.

4.2.3 Cross-Correlation

To determine how changes in metrics can provide information on the primary price, a measure of how changes in metrics translate into future changes price is required. In the model, the scarcity risk metrics and the primary price are expressed as time series. A statistical measure of how the changes in one time series are related to the changes in another is the sample cross-covariance function $\hat{\gamma}$, which, for two time series, x and y, is defined as

$$\hat{\gamma}_{xy}(h) = n^{-1} \sum_{t=1}^{n-h} (x_{t+h} - \bar{x})(y_t - \bar{y}) \tag{4.1}$$

Where n is the number of data points in the time series, t is time, and h is the lag time expressed as a negative number (Shumway, 2006). In words, sample **cross-covariance** is the average product of the differences between the lagged x realizations and the mean x, and the y realizations and the mean y. If lagged x and y tend to go up together then the cross-covariance is a large positive number; if they are independent of each other, it is near zero; and if they move in opposite directions is it strongly negative.

The difficulty in using this metric is in its interpretation. What does a cross-covariance of 10 mean relative to one of 200? In order to make the correlation easier to understand, the cross-covariance can be normalized by the by the auto-covariances (cross-covariance of the variable and a lagged version of itself) of x and y at a lag of 0, which can be shown to be equivalent to the standard deviation of the samples x and y. This new value, $\hat{\rho}_{xy}$, is the sample cross-correlation (Shumway, 2006; Chatfield, 2004).

$$\hat{\rho}_{xy}(h) = \frac{n^{-1} \sum_{t=1}^{n-h} (x_{t+h} - \bar{x})(y_t - \bar{y})}{\sqrt{\hat{\gamma}_y(0) \cdot \hat{\gamma}_y(0)}} = \frac{\hat{\gamma}_{xy}(h)}{\sigma_x \cdot \sigma_x} \tag{4.2}$$

The advantage of this normalization is that all values of $\hat{\rho}_{xy}$ will fall between -1 and 1, making interpretation easier. Values close to 1 in magnitude are strong correlations—in fact they are perfect correspondence. Those close to zero are weaker.

4.2.4 Stationarity & Autocorrelation

A time series is said to be stationary if the mean and variance of the time series do not change over time; or in statistical terms if the joint distribution of $x(t_1)$ and $x(t_k)$ is the same as that of $x(t_{1+\tau})$ and $x(t_{k+\tau})$, for all t (Chatfield, 2004). It is said to be weakly stationary if its mean and covariance function $\rho_{xy}(k)$ does not change over time (Chatfield, 2004; Shumway, 2006). Data with trends or that follow a 'random walk' can not be described as stationary(Chatfield, 2004; Shumway, 2006). The cross-correlation function only has meaning to the degree to which the time series in question can be modeled as weakly stationary. The time series from the model result do not appear to be stationary in that some variables, particularly the primary price, exhibit strong trends. To make strong inferences with the cross-correlation function the data will have to be filtered to make them more stationary.

A further difficulty in using cross-correlation functions is that when the time series are strongly auto-correlated, the values of cross correlation can be inflated (Chatfield, 2004). Series are considered strongly auto-correlated when their auto-correlation function shows strong correlation at distant lags (Shumway, 2006). This can occur if the series has a strong trend or is a random walk—in other words if the data is not stationary. The intuition behind the inflated cross-covariances is as follows: suppose y were trending upwards, independent of x, and x was trending upwards independent of y—this would give the illusion that lagged x correlated with increasing y. For example, most men's incomes increase over the course of their life, as their hairline recedes. It does not follow that income, per se, leads to hair loss—Or rather, this would be true only in the weakest sense; however, the two variables would be strongly cross-correlated. What would be really interesting, would be if a sudden increase in income lead to a sudden increase of hair loss–deviations from the trend. This can be seen if the trend is removed from the data, and we convert our time series into being a measure of deviations from the trend.

4.2.5 Data Transformation

A common and often effective method of removing trends and the effects of random walks is first-differencing, whereby the time series is replaced by the differences between each value and the previous one. Instead of x_t, each value in the series is replaced by $\nabla x = x_t - x_{t-1}$

(Chatfield, 2004; Shumway, 2006). However, this is not always enough to make the series stationary, so the process can be repeated again (Chatfield, 2004). The end is result is called the second-order difference ∇^2 of the time series.

$$\nabla^2 = \nabla x_t + \nabla x_{t-1} = x_t - 2x_{t-1} + x_{t-2} \tag{4.3}$$

To make the assumption of stationarity reasonable and also decrease the chance of spurious regression, each time series was transformed by taking the second-order difference. Some of the time series seemed stationary to begin with, but the primary price usually required second-differencing to appear stationary and, since that is the variable that all the metrics are being correlated with, it is logically consistent to apply the same filter to all the metrics. Second-differencing is effective at removing the effects of trends, but maintains the cross-correlation between variables, if it exists (Chatfield, 2004). An example of first and second-order differencing can be seen in Figure 4-1.

4.3 Set-up of Metrics Analysis

The experiment to determine the degree to which the scarcity risk metrics correlated with future price consists of 3^4 runs representing a full-factorial analysis of three levels of four market characteristics. Each combination from the full factorial analysis was replicated 25 times with a different noise seed for the stochastic demand component. Introducing the different noise seeds, means that the variables in each run will follow a slightly different path; introducing a range of price paths ensures that the results are not just a product of one specific model path.

So there are $25 \cdot 3^4 = 2025$ model runs. From each run there are time series for each of the scarcity risk metrics and primary price; the time series span the model years 1995 to 2035, with data recorded quarterly. The cross-correlations between the metrics and primary price for a range of lag-times are calculated. Then the results of the 2025 runs are divided up into groups by the three levels of each of the four market characteristics, so as to determine how the cross-correlations differ depending on the market characteristic.

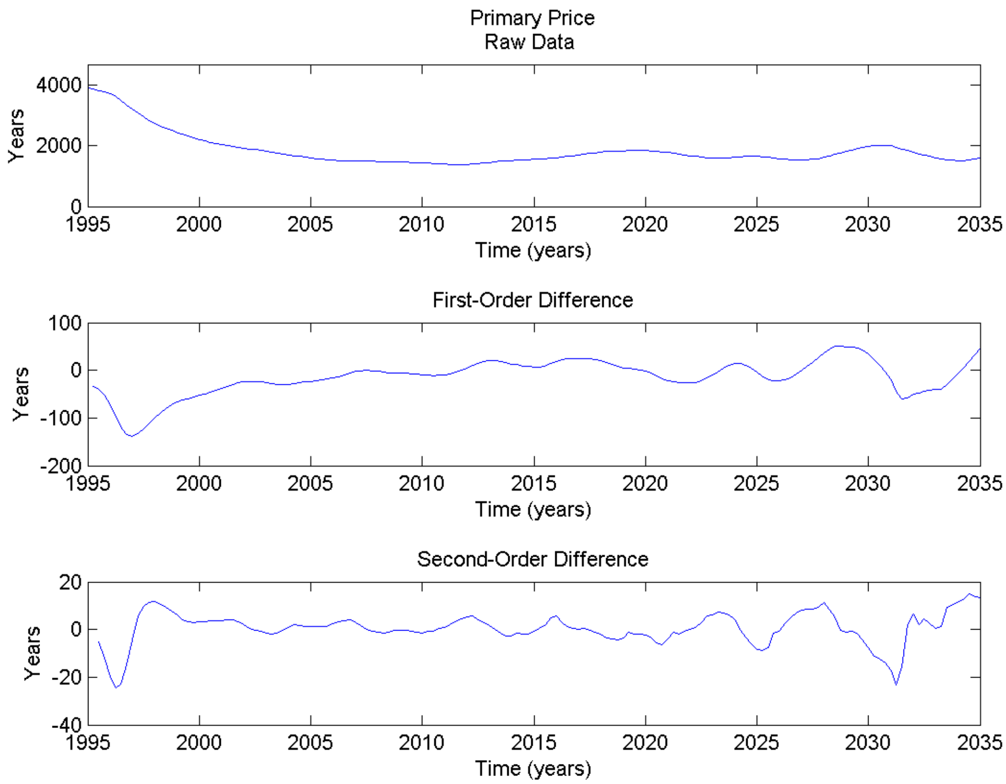

Figure 4-1: Primary Price at different levels of filtering: Unfiltered at the top, first-order differenced in the middle, and second-order differenced at the bottom. Note that there is no visible linear trend in the bottom plot.

4.4 Simulation Results

The results are displayed in the form of correlograms, which are plots of the cross-correlation function, Equation 4.2, for lag times of up to 20 years. The correlograms show the strength of the cross-correlation between series and the direction of the correlation. The results are presented in two groups: first, correlograms showing the cross-correlation of each metric with primary price (including primary auto-correlation) for all 2025 runs (3^4 factorial combinations times 25 replicates;) this is followed by a group of results broken down by the 3 levels of each of the 4 characteristics. The results are summarized in Table 4.2.

4.4.1 Cross-Correlation Plots - All Runs

The cross-correlation plots for all runs are displayed in Figure 4-2. The plots display the average cross-correlation and the extreme values—minimums and maximums—for all 2025 runs. Each correlogram displays the mean correlation between each of the eight metrics and primary price at lags from 0 to 20 years. The primary price correlogram is termed an 'autocorrelogram' in that it is Primary Price's correlation with its own lagged values. So, the value of the mean curve at any lag is the average correlation coefficient for that lag for all runs. For example, in the bottom rightmost plot—for static depletion index—the mean correlation is at a lag of ten years is about -0.1. That means that on average, the static depletion index at any time has a weak negative correlation with primary price ten years from that time. The plot also displays minimum and maximum correlation value for all runs—represented in the lighter dashed lines.

There are two striking features of the results: first, their wide range as demonstrated by the broad spread of extreme values. Correlation coefficients at any lag will range from moderately positive to moderately negative. On any one particular run the lag at any given time could be positive, negative, or zero. This suggests that any two particular runs (out of the 2025) could tell very different stories as to the relationship between the variables. The next most striking feature is that, on average, most metrics show very little correlation with primary price. This is somewhat surprising since the variables are essentially deterministically linked in the system dynamics model, but in such a complex fashion that many variables do not show a strong link with primary price on average.

There are four metrics that seem to cross-correlate with primary price to an appreciable degree: two more strongly than the others. The strongest correlations are between Recycling Efficiency and the Static Depletion Index, followed by alumina Price and Primary Marginal Cost. The following paragraphs contain detailed descriptions of the correlograms created from the simulation results.

Recycling Efficiency:

Recycling Efficiency is positively correlated with primary price at zero lag. Since the correlations are for twice-differenced time series, this suggests that as the second derivative of recycling efficiency is positive, the second derivative of primary price is positive, with a moderate correlation. This correlation declines after about a year and then becomes negative. This cross-correlation mimics the autocorrelation of Primary Price, and this is not surprising: Recycling Efficiency is closely tied to primary price in that changes in Primary Price translate into changes in Recycling Efficiency. In the model, as might be expected in real life, increased primary price increases the demand for secondary metal, which increases the Recycling Efficiency. It also changes over a relatively short timespan, since the secondary industry was modeled to respond relatively rapidly to changes in market conditions. So, it is probably not so much that the Recycling Efficiency correlates with future price as it is that recycling efficiency is closely tied to current price; and, the current primary price correlates with with future primary price (because primary price autocorrelates.)

Static Depletion Index:

The Static Depletion Index has a moderate positive correlation with Primary Price at low lags, which then crosses zero after about three years and alternatives between weak positive and negative correlation. This result is probably the most interesting because a close correlation is not expected. Static Depletion Index is mechanistically disconnected from primary price; moreover, there are a series of lags between the primary price signal and the static depletion index.

Alumina Price:

Alumina Price exhibits an alternating positive and negative weak correlation with Primary Price, but the correlation begins as negative at lag zero–peaking at about one year. This suggests that when alumina prices are turning up, Primary Prices are turning down and will do so to a greater extent a year away.

Primary Marginal Cost:

Primary Marginal Cost is weakly positively correlated with Primary Price, but rapidly becomes negative as the lag increases–peaking at about one year. It mimics the alumina Price correlogram, which is not surprising since Alumina Price is a large component of the Primary Marginal Cost.

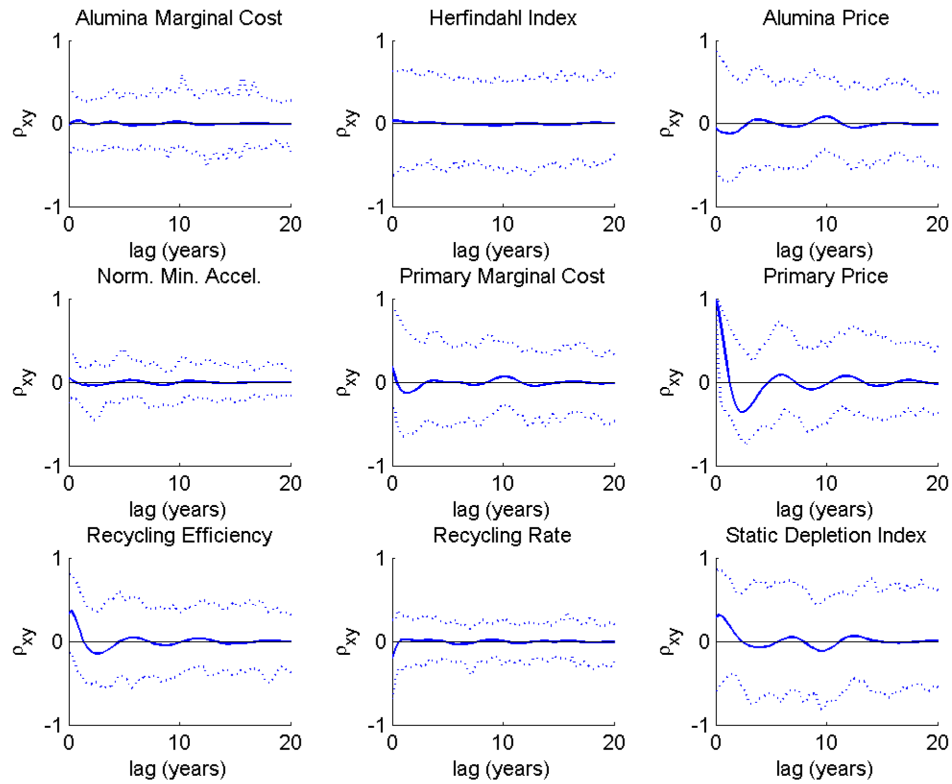

Figure 4-2: Correlograms showing cross-correlation of the nine metrics and primary Price. The solid blue line represents the average over 2025 model runs. The dashed lines represent the extreme minimum and maximum values of all runs.

4.4.2 Break-Downs By Market Characteristic

The more interesting results are the metric correlograms organized by market characteristic: They are displayed in Figures 4-3 through 4-11. Each of those figures contains four plots which are the results for all 2025 runs broken up by each of the four market characteristics. Each plot has 3 lines, corresponding to low, medium and high levels of that characteristic–each line being the mean line for 1/3 of the data, 6075 runs. To spare the reader a lengthy, plot by plot analysis of these results, the general findings can be summarized by the effects of each characteristic:

Alumina Capacity Acquisition Delay:

The alumina Capacity Acquisition Delay has almost *no impact* on the cross-correlations of the metrics and the primary price.

Primary Capacity Acquisition Delay:

Unlike the Alumina Capacity Acquisition Delay, the Primary Capacity Acquisition Delay has a very pronounced effect on the cross correlations—particularly for the four metrics that showed strong correlations in the average results, Alumina Price, Primary Marginal Cost, Recycling Efficiency, and Static Depletion Index. In all of those cases, and with the addition of Normalized Mining Acceleration, which did not show a strong result in the plots of all results, l*ow Primary Capacity Acquisition Delays lead to much stronger cross-correlations between the metrics and primary price*, with more pronounced alternations between positive and negative correlations. Conversely, with longer delays, the cross-correlations became significantly weaker.

Goods Demand Elasticity:

For the metrics, Alumina Price, Primary Marginal Cost, and Static Depletion Index, *higher Goods Demand Elasticity increased the strength of the correlations*, but no effect was seen on Recycling Efficiency.

Secondary Substitutability:

For the metrics, Alumina Price, Primary Marginal Cost, and Recycling Efficiency, *higher Secondary Substitutability increased the strength of the correlations*, but no strong effect was seen on the Static Depletion Index. For the most part, it seems that Secondary Substitutability behaved like elasticity, in that it increased the strength of the correlations between the metrics and primary price.

4.5 Conclusions

There are two main conclusions that can be drawn from the experimental results: First, some of the modeled scarcity risk metrics do correlate weakly with primary price; namely, Recycling Efficiency, Static Depletion Index, Alumina Price, and Primary Marginal Cost; but the other metrics showed a very weak to no correlation. The second conclusion is that stronger correlations can be found for a specific set of market characteristics—low demand elasticity, high secondary substitutability, and low primary capacity acquisition delays. What that set of characteristics has in common is that they all tend to stabilize

Metric	Peak Cross-Corr. (All Runs)	Lag at Peak Correlation	Alumina Cap. Acquisition Delay	Primary Cap. Acquisition Delay	Goods Demand Elasticity	Secondary Substitutability
Primary Price	1	0		-		
Recycling Efficiency	0.37	0.25		-		+
Static Depletion Index	0.32	0.25		-	+	
Recycling Rate	-0.20	0				
Primary Marginal Cost	0.20	0		-	+	+
Alumina Price	-0.12	1.25		-	+	+
Norm. Min. Accel.	0.05	0		-		
Alumina Marginal Cost	0.04	0.75				
Herfindahl Index	0.04	0.25				

Table 4.2: Results Summary: This table displays the peak value of the cross-correlation coefficient for each metric, the lag time at which this peak occurred, and whether each of the four market characteristics changed the correlations. The '+' signs indicate that as the size of the market characteristic increased, the strength of the correlation increased; '-' signs indicate that as the size of the market characteristic increased, the strength of the correlation decreased.

price. They tend to stabilize price, because they increase the responsiveness of either supply or demand to price. So, it seems that in markets where supply or demand responds more quickly to price, the scarcity risk metrics correlate more strongly with primary price. Conversely, when supply and demand respond more slowly, the other metrics have much less predictive power—approaching none whatsoever. Or, another way of describing the result is that when the market characteristics are such that the price is very volatile, the other scarcity risk metrics do not provide much information on price.

This suggests that any attempts by firms to change the market characteristics towards those that provide a more stable price, would not just help stabilize price, but could open up opportunities to learn about price motions from scarcity risk metrics. The price elasticity of aluminum goods demand would be out of the control of firms that make aluminum products, and the delays in adding capacity are most likely due to the fundamental difficulty of the task; however, the secondary substitutability is at least partially within the control of the firm, as mentioned in section 3.8. It is possible that if firms increased the secondary substitutability of their products it would have the effect of increasing the amount of information available to them about price movements, as well as help stabilize price.

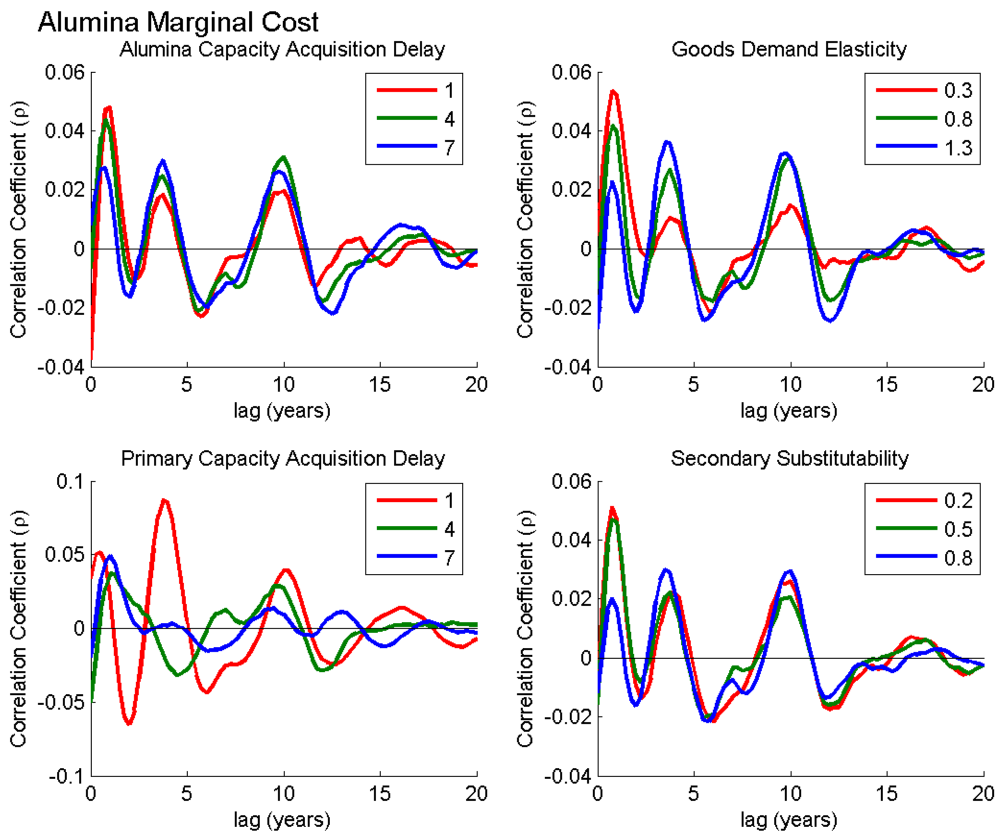

Figure 4-3: **Alumina Marginal Cost Correlogram, by Characteristic:** Cross-Correlations between alumina Marginal Cost and Primary Price are all very weak. Primary Capacity Acquisition Delay increases the strength and alternations of the cross-correlations, but it is hard to say if this is meaningful since they are all so low.

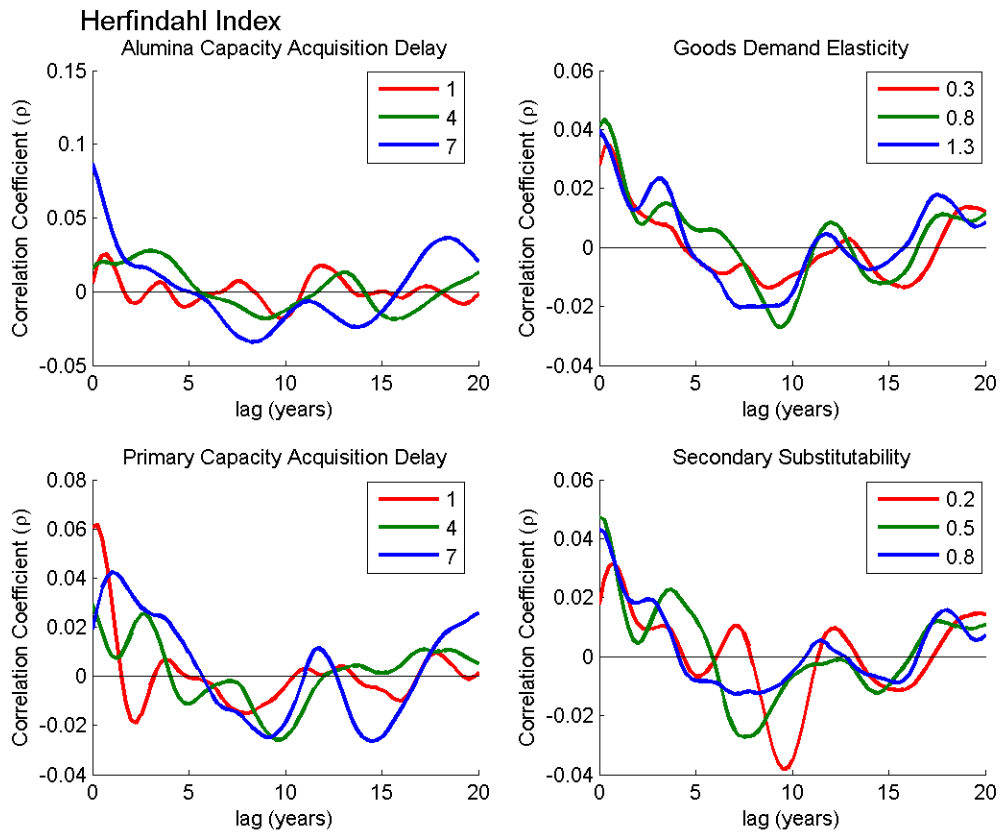

Figure 4-4: **Herfindahl Index Correlogram, by Characteristic:** As with the previous graph, the strength of the correlations is so low it is hard to say anything meaningful about the effects of the market characteristics

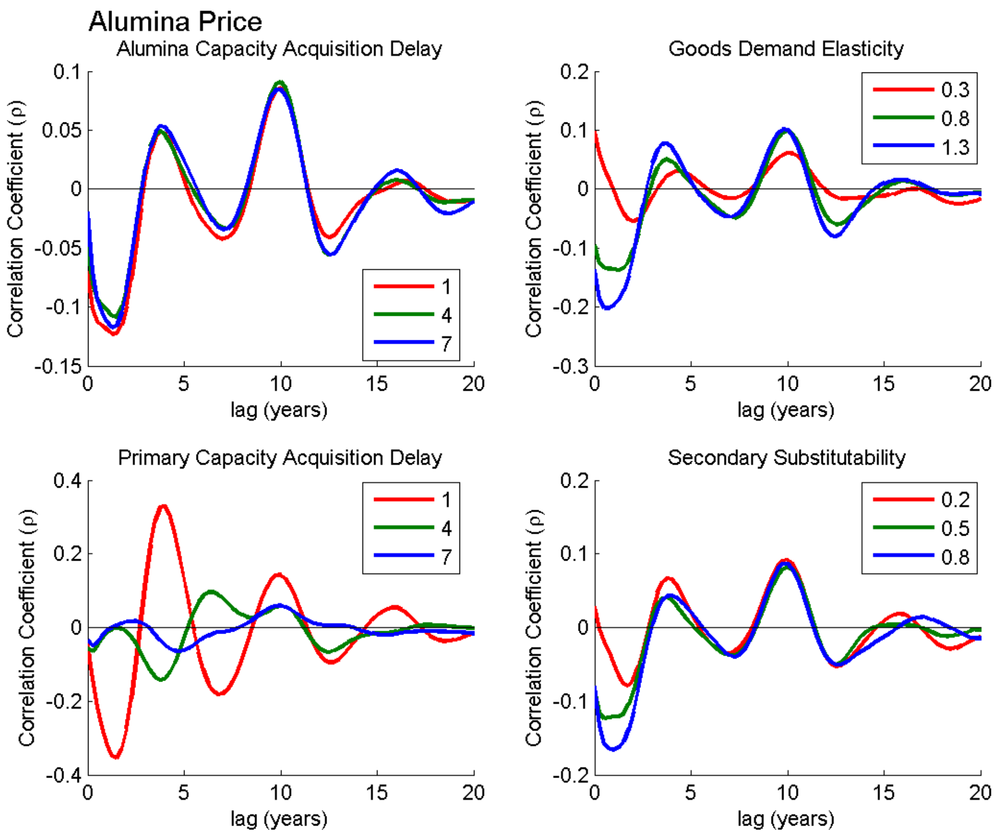

Figure 4-5: **Alumina Price Correlograms, by Characteristic:** Alumina Price moderately correlates with Primary Price. Increasing Primary Capacity Acquisition Delay decreases the correlation, goods demand elasticity increases it, and secondary substitutability increases it.

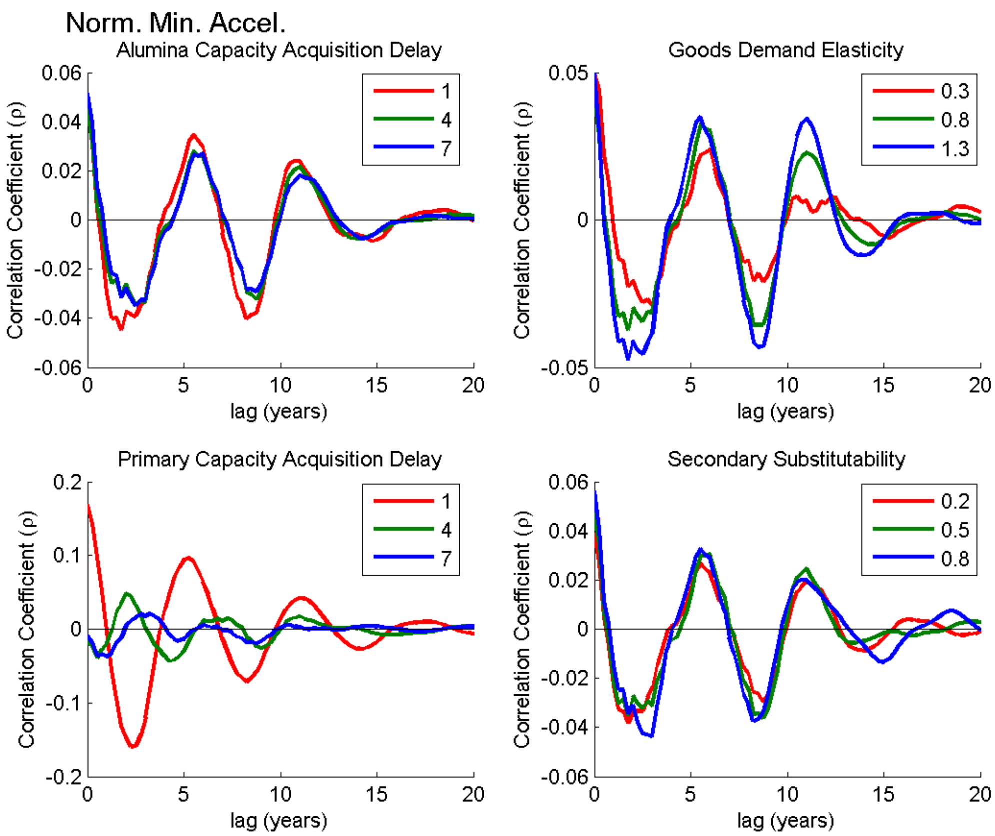

Figure 4-6: **Normalized Mining Acceleration Correlograms, by Characteristic:** Normalized Mining Acceleration has a very weak correlation with Primary Price, but it does become stronger with low Primary Capacity Acquisition Delay.

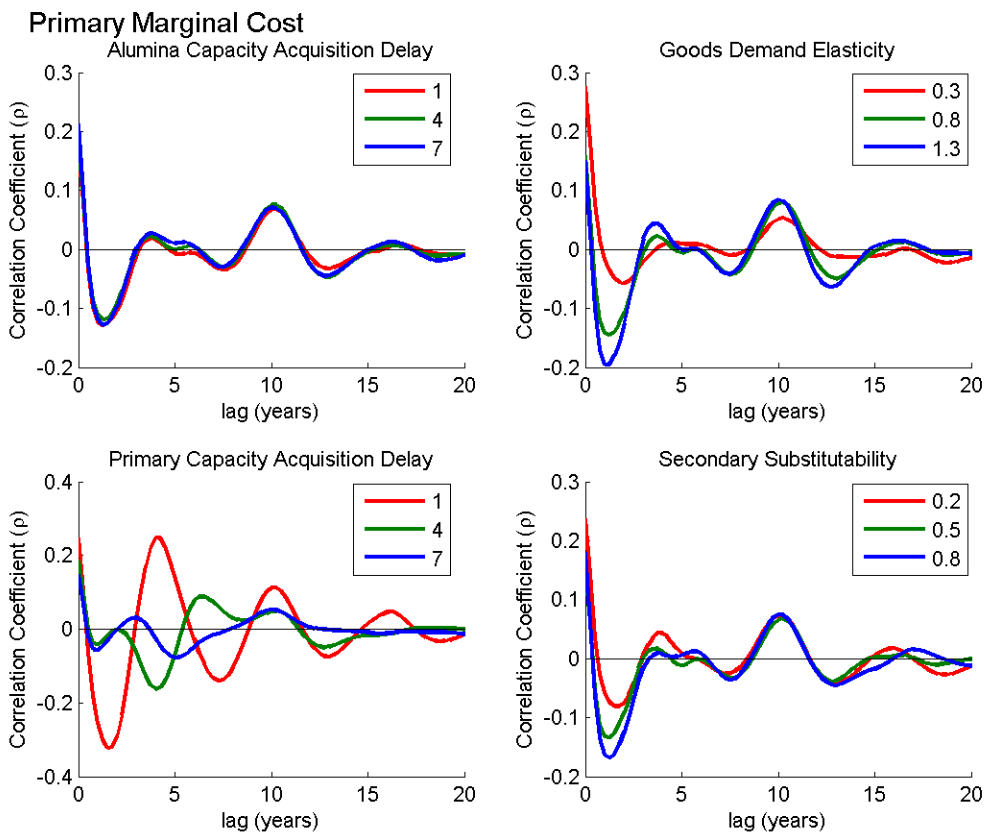

Figure 4-7: **Primary Marginal Cost Correlograms, by Characteristic:** Primary Capacity Acquisition Delay decreases correlation, goods demand elasticity increases it, and secondary substitutability increases it.

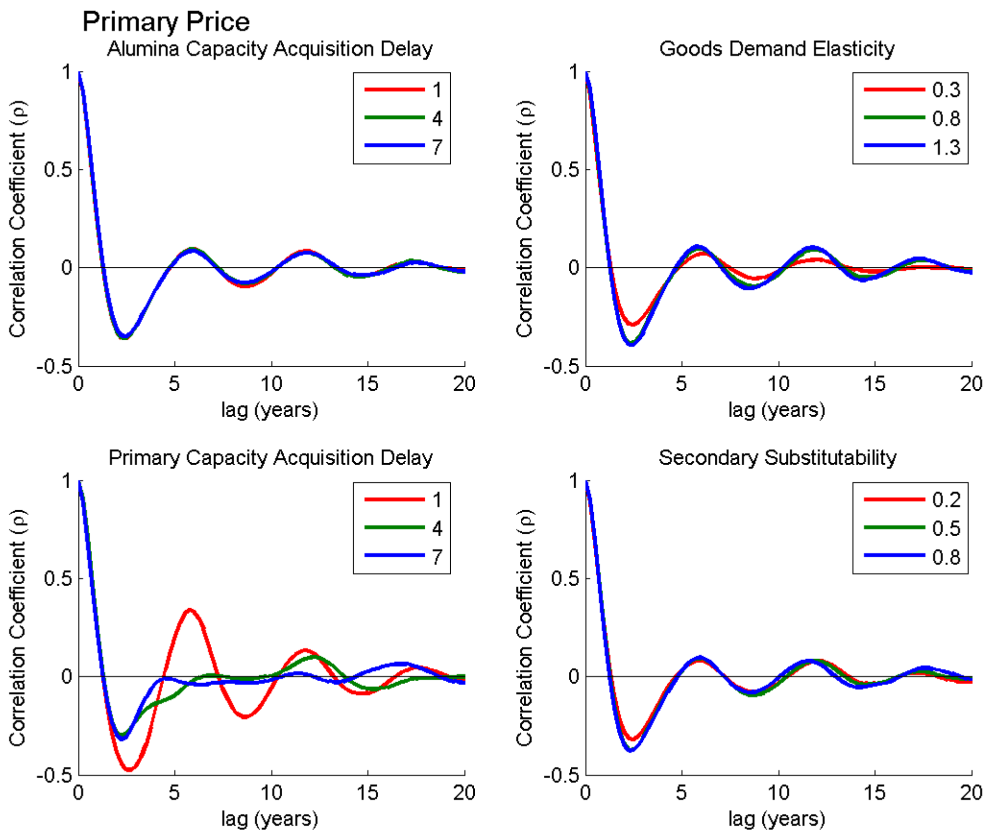

Figure 4-8: **Primary Price Auto-Correlograms, by Characteristic:** Primary Price is autocorrelated out to about 1 year, after which the correlation becomes negative, and then diminishes. Primary Capacity Acquisition Delay decreases the oscillation of the autocorrelation.

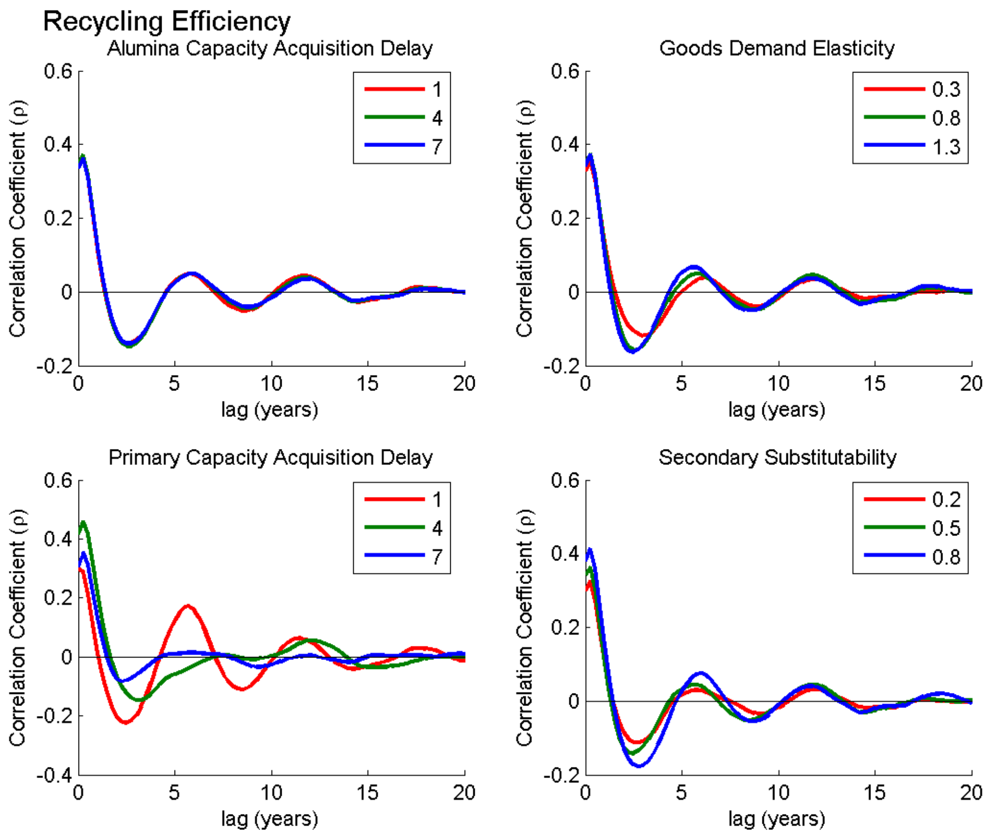

Figure 4-9: **Recycling Efficiency Correlograms, by Characteristic:** Primary Capacity Acquisition Delay decreases correlation, while secondary substitutability slightly increases it.

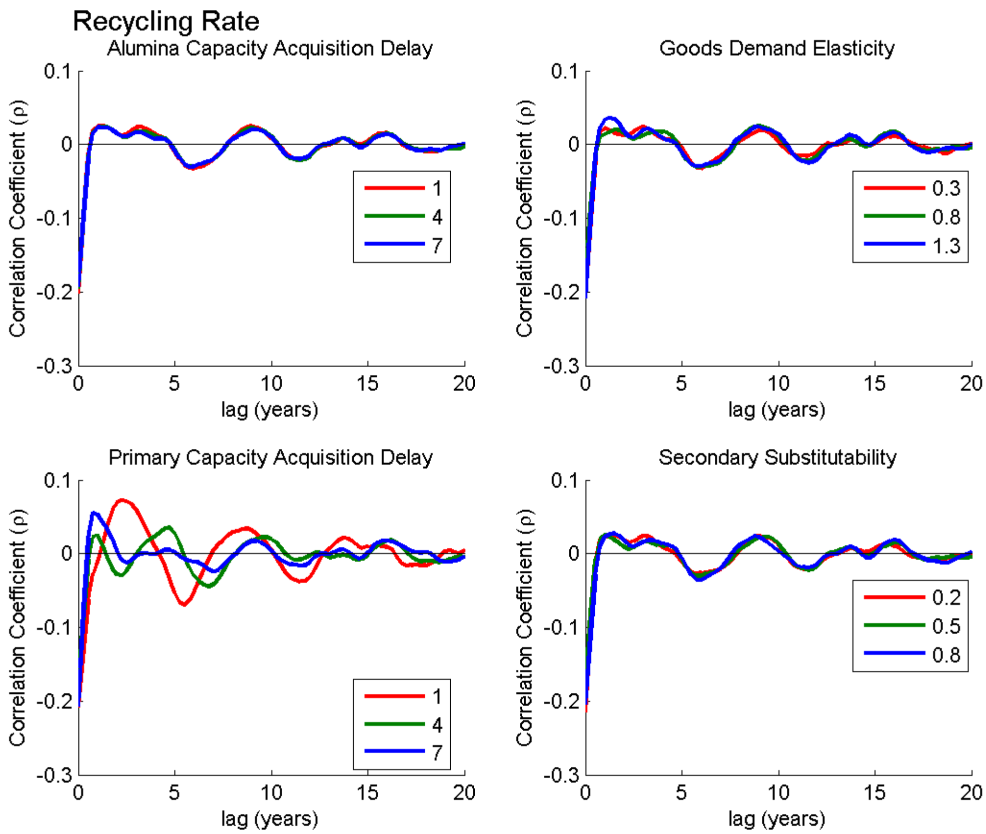

Figure 4-10: **Recycling Rate Correlograms, by Characteristic:** The correlation between Recycling Rate and Primary price is too small to say much about the effect of market characteristics.

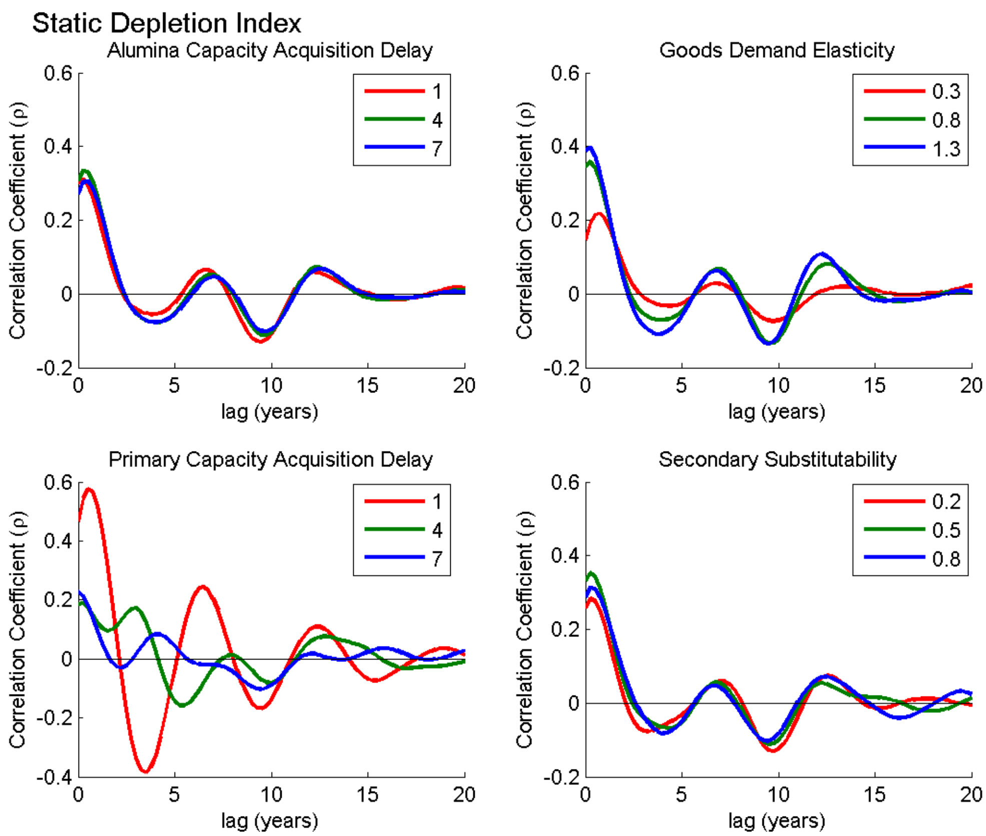

Figure 4-11: **Static Depletion Index Correlograms, by Characteristic:** Primary Capacity Acquisition Delay greatly decreases the correlation between Static Depletion Index and Primary Price, and goods demand elasticity increases it.

Chapter 5

Conclusions

5.1 Conclusions and Recommendations for Firms

Metal price volatility is a concern for firms that use metals as raw materials, because volatility in prices can translate into volatility of costs for those firms. Unexpectedly high costs may destroy profitability. Even unexpectedly low prices can be troublesome, because they may lure firms into selecting the material for use in a product, only to rise later and make the product unprofitable. Undesirable volatility in possible substitute materials may prevent firms from selecting them for their products, thus limiting their material selection decisions. Due to these concerns, firms have an incentive in trying to gain advanced information regarding changes in prices—to understand when they are at risk for a price change, and to find ways to manage price volatility, or at least decrease their exposure to it.

The first major conclusion of this study is that, in the simulation results, firms can influence the primary metal price by changing the degree to which secondary metal can be substituted for primary (secondary substitutability.) The larger the limit for secondary metal in a firm's product, the more elastic that firm's secondary use will be with respect to price. This increased elasticity allows firms to use less of the more expensive primary metal when its price increases, and thereby lowers the demand for the primary metal—decreasing its price.

It is important to make sure that this result is not confused with another more obvious benefit that the higher technical limits provide, which is greater flexibility to use secondary metal. Of course, if firms have the option to use 20% recycled metal as opposed to 10% they will have more flexibility as to how much they use. The results of the analysis show much more than that: consider a firm that uses secondary metal with a technical limit of 10%; they will not use all of that 10% due to variability in scrap quality. They will use a fraction of that 10%, say 40% of that limit, or 4% secondary metal. That percentage is set because it optimizes the savings from using scrap against the risks of exceeding batch targets due to variability in scrap quality. If the primary price increases, the balance shifts and the optimum percentage of the technical limit changes, say to 50% of the limit or 5% secondary metal.

Now, suppose that this firm changes the technical limit to 10% from 20%. The results predict that this change alone increases the fraction of the technical limit that it is optimal for them to use, from 40% to say 50% of that limit, which means they will use 10% secondary metal. By doubling their limit, they have more than doubled the actual (optimal) secondary metal use, going from 4% to 10% secondary metal. Moreover, for any change in price, the optimal fraction of the technical limit will change more under the higher technical limit. A change in primary price that produced a change from say 40% to 50% of the technical limit under the lower technical limit might produce a change from 50% to 70% under the higher limit; in secondary metal terms this would be the difference between going from 4% to 5% recycled metal under the low limit, but 10% to 14% under the higher limit. In other words, for a given price change, the change in the optimal recycling rate will be higher when the technical limit is higher; or, in economic terms, the price elasticity of substitution between primary and secondary will be higher. The benefits of the greater elasticity of substitution are both on the firm level and industry-wide: When primary price changes, the higher elasticity firm makes bigger adjustments away from the more expensive metal—larger both proportionally and in absolute terms. On the industry level, this can lower demand for the more expensive metal, thus preventing further increases in (or lowering) its price.

To reiterate, the model results suggest that *by raising their technical limits for secondary metal,*

1. Firms can increase their optimal secondary use by a greater percentage than the percentage by which they changed the limit.
2. They increase their price elasticity of substitution between primary and secondary metal—decreasing their exposure to price increases in primary metal.
3. Greater industry-wide price elasticity of substitution between secondary and primary

metal can have the effect of lowering primary price volatility.

So, according to the model, increasing secondary substitutability can help manage price volatility, it also increased the correlation between scarcity risk metrics and primary price.

The next part of the study evaluated whether other metrics of scarcity risk besides primary price could be used to provide information on changes in the price of primary metal in the simulation. The results show, that although many scarcity risk metrics examined in the model did not correlate well with future primary price, some metrics do weakly correlate with the future primary price. Those model scarcity risk metrics are Static Depletion Index, Alumina Price, Primary Marginal Cost, and Recycling Efficiency. The latter would not be especially useful to a firm, because it itself is a function of primary price (it increases and decreases with the primary price, all else being equal,) and likely only correlates with future primary price because primary price correlates with past values of itself (autocorrelation.) Furthermore, the Recycling Efficiency is not easy to calculate: The amount of metal recycled may be calculated but the amount disposed off is not tracked. Primary Marginal Cost is also not especially useful, because it is not a publicly available number.

Of the four parameters, Static Depletion Index and Alumina Price, afford useful insight into potential scarcity pricing risks for a firm. These have the advantage that they can be calculated without great difficulty. Alumina is publicly traded, so there is a market price. To the extent that the model results hold in the real world, the Alumina price should be a bellwether for primary price. The Static Depletion Index can also be calculated with publicly available information, though perhaps not with great frequency. However, the model shows that the metrics only correlate with future primary price when the market has certain characteristics: In the case of the Static Depletion Index, it correlates best when the primary elasticity is high, or the delays in adding primary capacity are low. Likewise for Alumina Price, with the addition that high secondary substitutability also increases its informative power. In general, the scarcity risk metrics correlated better with the primary price when the primary capacity acquisition delay was low, the goods demand elasticity was high, and the secondary substitutability was high.

These results can be summarized as:

1. In the aluminum simulation, some scarcity metrics weakly correlate with future primary price; namely, Static Depletion Index, Alumina Price, Primary Marginal Cost,

and Recycling Efficiency.

2. Correlations were stronger when the delay in adding primary capacity was low, when the finished good's price elasticity of demand was high, or the secondary substitutability was high.

To the extent that the model results hold in the real world, firms may be able to use the aforementioned (and potentially other) scarcity risk metrics to learn of possible changes in primary price. The strength of the correlations depends on the characteristics of the market; one of which is the degree to which secondary metal can be substituted for primary. This substitutability may be in the power of the firm to change. The results from the studies in this work suggest that increasing the secondary substitutability of products may lower a firm's exposure to primary price volatility; furthermore, market-wide increases in secondary substitutability may decrease the price volatility and increase the ability of firms to gather advanced information about changes in price from scarcity risk metrics.

5.2 Future Work

This research only examined the effect of a few market characteristics on the ability of the use of scarcity metrics to provide advanced information about changes in primary price. Future work might explore others. Furthermore, it would be interesting to look at the interaction effects of different market characteristics. For instance, high elasticity and or short primary capacity acquisition delays increased the predictive power of some of the metrics; well, what happens when the two happen simultaneously?

In this model, there was only one grade of scrap; what is the effect of having multiple grades? Future work might include the modeling of multiple grades of scrap. It would not be difficult to include material flows for multiple grades in the model. In fact the model was built to do so; however, the challenge is determining the optimal recycling rate. Doing so might require some built in linear programming in the system dynamics model, which has been done by others successfully (Martinez et al., 1999).

Bibliography

E. Alonso, J. Gregory, F. Field, and R. Kirchain. Material availability and the supply chain: Risks, effects, and responses. *Environmental Science & Technology*, 41(19):66496656, 2007a.

Elisa Alonso. *Material scarcity from the perspective of manufacturing firms : case studies of platinum and cobalt.* PhD thesis, Massachusetts Institute of Technology, 2010.

Elisa Alonso, Frank Field, Jeremy Gregory, and Randolph E Kirchain. *Materials Availability and the Supply Chain: Risks, Effects, and Responses.* MIT - Materials Systems Laboratory, January 2007b.

Robert U Ayres, Leslie Ayres, and Ingrid Rde. *The life cycle of copper, its co-products and byproducts.* Eco-efficiency in industry and science ; v. 13. Kluwer Academic, Dordrecht ; Boston, 2003.

Harold J Barnett and Chandler Morse. *Scarcity and growth; the economics of natural resource availability.* Resources for the Future, Johns Hopkins Press, Washington, 1963.

E. Lee Bray. Bauxite and alumina [advance release]. In *2009 Minerals Yearbook*, volume I. Interior Dept., Geological Survey, 2009.

Michael J Brennan. The supply of storage. *The American Economic Review*, 48(1):5072, 1958.

David Brooks. *Minerals: an Expanding Or a Dwindling Resource?* 1973.

Celso Brunetti and Christopher L Gilbert. Metals price volatility, 1972-1995. *Resources Policy*, 21(4):237254, 1995.

D.A. Buckingham. *Aluminum Statistics.* U.S. Geological Survey, October 2010. URL `http://minerals.usgs.gov/ds/2005/140/aluminum.pdf`.

Alfred Cavallo. Hubbert's petroleum production model: an evaluation and implications for world oil production forecasts. *Natural Resources Research*, 13(4):211221, December 2004. ISSN 1520-7439.

P. F Chapman and F. Roberts. *Metal resources and energy.* Butterworths monographs in materials. Butterworths, UK, 1983.

Christopher Chatfield. *The analysis of time series: an introduction.* CRC Press, 2004. ISBN 9781584883173.

William T. Choate and John A. S. Green. *U.S. Energy Requirements for Aluminum Production.* U. S. Department of Energy, 2003.

C. J Cleveland. An exploration of alternative measures of natural-resource scarcity - the case of petroleum resources in the United-States. *Ecological Economics*, 7(2):123157, 1993.

Nurhan Davutyan and Mark C. Roberts. Cyclicality in metal prices. *Resources Policy*, 20 (1):4957, March 1994. ISSN 0301-4207. doi: 10.1016/0301-4207(94)90040-X.

S. Devarajan and A. C Fisher. Hotelling economics of exhaustible resources - 50 years later. *Journal of Economic Literature*, 19(1):6573, 1981.

Steven Durlauf. *The new Palgrave dictionary of economics.* Palgrave Macmillan, Basingstoke Hampshire ;;New York, 2nd ed. edition, 2008. ISBN 9780333786765.

EIA. *EIA annual energy review 2009.* U.S. Energy Information Administration, 2010.

Eugene Fama. The behavior of Stock-Market prices. *Journal of Business*, 38(1):34, 1965. ISSN 0021-9398.

Ronald Fisher. *The Design of Experiments, By Sir Ronald A. Fisher.* New York Hafner, (8Th ed.) edition, 1971.

Jay Forrester. *Industrial dynamics.* M.I.T. Press, [Cambridge Mass.], 1961.

R. Fowler. The substitution of scrap for Pig-Iron in the manufacture of steel. *The Quarterly Journal of Economics*, 52(1):129, 1937. ISSN 0033-5533.

Milton Friedman. *Price theory.* Aldine Publishing, New York, 1976. ISBN 9780202060743.

Gabriel Gaustad. *Towards sustainable material usage : time-dependent evaluation of upgrading technologies for recycling.* PhD thesis, Massachusetts Institute of Technology, 2009.

R. B Gordon, M. Bertram, and T. E Graedel. Metal stocks and sustainability. *Proceedings of the National Academy of Sciences of the United States of America*, 103(5):12091214, 2006.

R. L. Gordon and J. E. Tilton. Mineral economics: Overview of a discipline. *Resources Policy*, 2008.

Sanford Grossman. On the impossibility of informationally efficient markets. *The American Economic Review*, 70(3):393, 1980. ISSN 0002-8282.

Oliver Hart and David Kreps. Price destabilizing speculation. *The Journal of Political Economy*, 94(5):927, 1986. ISSN 0022-3808.

Albert Hirschman. The paternity of an index. *The American Economic Review*, 54(5):761, 1964. ISSN 0002-8282.

Harold Hotelling. The economics of exhaustible resources. *The Journal of Political Economy*, 39(2):137175, 1931.

M. King Hubbert. *Energy resources.* A report to the Committee on Natural Resources of the National Academy of Sciences-National Research Council. National Academy of Sciences-NRC, Washington, 1962.

Thomas Kelly. *Historical statistics for mineral and material commodities in the United States.* U.S. Geological Survey,, [Reston, Va.] :, 2005.

Jeffrey Krautkraemer. Nonrenewable resource scarcity. *Journal of Economic Literature*, 36 (4):2065, 1998. ISSN 0022-0515.

T. Malthus. *An essay on the principle of population or, a view of its past and present effects on human happiness, with an inquiry into our prospects respecting the future removal or mitigation of the evils.* R.D. Irwin, Homewood Ill., 1963.

Thomas Robert Malthus. *An Essay on the principle of population, as it affects the future improvement of society with remarks on the speculations of Mr. Godwin, M. Condorcet and other writers.* J. Johnson, May 2006.

Silvio Martinez, Francisco Gordillo, Elena Lopez, and Ismael Alcala. A case of interaction between systems dynamics and linear programming: The Rapim-Pirenaica model. In *The 17th International Conference of The System Dynamics Society and the 5th Australian & New Zealand Systems Conference*, Wellington, New Zealand, 1999.

Grecia Matos and Lorie Wagner. Consumption of materials in the US, 1900-1995. *Annual Review of Energy and the Environment*, 23(1):107122, 1998.

James McClements and Russ Cranswick. World mining overview. *Engineering & Mining Journal*, 202(11):2837, 2001.

Dennis L Meadows. *Dynamics of commodity production cycles.* Wright-Allen Press, Cambridge, Mass., 1970.

Donella H Meadows, Dennis Lynn Meadows, and Potomac Associates. *Limits to growth; a report for the Club of Rome's project on the predicament of mankind.* Universe Books, New York, 1972.

MMSD. *Breaking new ground: the report of the Mining, Minerals and Sustainable Development project.* Earthscan Publications Ltd., January 2006. ISBN 1-85383-907-8.

David Newbery. *The manipulation of futures markets by a dominant producer.* New York City, 1982.

William D Nordhaus, Hendrik Houthakker, and Robert Solow. The allocation of energy resources. *Brookings Papers on Economic Activity*, 1973(3):529576, 1973.

Sam H Patterson. *World bauxite resources : an updating of information on the geologic and geographic occurrences of world bauxite resources and discussions of uses, exploration and mining methods, and production.* Geology and resources of aluminum. U.S. G.P.O., Washington, 1986.

R Peterson. Scrap variability and its effects on producing alloys to specification. *Light metals.*, page 855, 1999. ISSN 0147-0809.

Robert S Pindyck. Inventories and the Short-Run dynamics of commodity prices. *The Rand Journal of Economics*, 25(1):141159, 1994.

Patricia Plunkert. *Metal prices in the United States through 1998.* U.S. Dept. of the Interior U.S. Geological Survey ;;For sale by the U.S. G.P.O., [Reston Va.] ;Washington DC, 1999.

Marian Radetzki and Carl Van Duyne. The demand for scrap and primary metal ores after a decline in secular growth. *Canadian Journal of Economics*, 18(2):435449, 1985.

Judith Rees. *Natural resources : allocation, economics, and policy.* Methuen, London ;;New York, 1985. ISBN 9780416319903.

David Ricardo. *On the principles of political economy and taxation.* John Murray, May 2006.

Martin Ruhrberg. Assessing the recycling efficiency of copper from end-of-life products in western europe. *Resources, Conservation and Recycling*, 48(2):141165, 2006.

Donald Rutherford. *Routledge dictionary of economics.* Routledge, London ;;New York, 2nd ed. edition, 2002. ISBN 9780415250917.

Paul A. Samuelson. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review*, 6(2), 1965.

Christoph Schmitz. *Handbook of aluminium recycling.* Vulkan-Verlag GmbH, 2006. ISBN 9783802729362.

Robert Shumway. *Time series analysis and its applications with R examples.* Springer,, New York :, 2006. ISBN 9780387293172.

H. A. Simon. Rational choice and the structure of the environment. *Psychological review*, 63(2):12938, March 1956. ISSN 0033-295X.

Margaret E Slade. Trends in natural-resource commodity prices: An analysis of the time domain. *Journal of Environmental Economics and Management*, 9(2):122137, 1982.

Margaret E Slade. *Pricing of metals.* Centre for Resource Studies, Queen's University, Kingston, Ont., 1988.

Margaret E Slade. *Do markets underprice natural - resource commodities?*, volume 962 of *Policy Research Working Paper Series.* The World Bank, Washington, DC, USA, 1992.

Robert M Solow. The economics of resources or the resources of economics. *The American Economic Review*, 64(2):114, 1974.

John Sterman. *Business dynamics : systems thinking and modeling for a complex world.* Irwin/McGraw-Hill, Boston, 2000.

Daniel Sullivan. *20th century U.S. mineral prices decline in constant dollars.* U.S. Dept. of the Interior, U.S. Geological Survey,, [Reston, Va.], 2000.

Swiss Centre for Life-Cycle Inventories. Ecoinvent 2.2 database: Accessed via SimaPro, May 2010.

John Sznopek. *Drivers of U.S. mineral demand.* U.S. Geological Survey,, Reston, Va. :, 2006.

Thomas H Tietenberg. *Environmental and natural resource economics.* HarperCollins Publishers, New York, NY, 2009.

J. E Tilton. The future of recycling. *Resources Policy*, 25(3):197204, 1999.

United Nations. *World population to 2300.* United Nations, New York, 2004. ISBN 9781417577156.

Randall Urbance. MAGNESIUM: PRODUCTION, PROPERTIES, AND MARKETS - market model simulation: The impact of increased automotive interest in magnesium. *JOM : the journal of the Minerals, Metals & Materials Society.*, 54(8):25, 2002. ISSN 1047-4838.

USGS. *Principles of a Resource/Reserve Classification for Minerals*, volume Circular 831. 1980. ISBN (200) Ci no.831.

U.S.G.S. *Minerals information.* Number April 2009. U.S. Department of the Interior, 2009.

Lorie Wagner. *Materials in the economy, material flows, scarcity, and the environment.* U.S. Dept. of the Interior U.S. Geological Survey, Denver CO, 2002. ISBN 9780607984125.

Lorie Wagner. *Economic drivers of mineral supply.* U.S. Dept. of the Interior, U.S. Geological Survey,, [Reston, Va.] :, 2003.

Michael Wetzstein. *Microeconomic theory : concepts and connections.* Thomson/South-Western, Mason OH, 2005. ISBN 9780324260298.

F. Helmut Weymar. *The dynamics of the world cocoa market.* M.I.T. monographs in economics, 9. M.I.T. Press, Cambridge, 1968.

World Bank. *Global economic prospects 2006: economic implications of remittances and migration.* Number November 2005. The World Bank Group, March 2006. ISBN 0-8213-6344-1.

Irene Xiarchos. *Three essays in environmental markets dynamic behavior, market interactions, policy implications.* [West Virginia University Libraries], Morgantown, W. Va. :, 2006.

Irene Xiarchos. Price and volatility transmission between primary and scrap metal markets. *Resources, conservation, and recycling.*, 53(12):664, 2009. ISSN 0921-3449.

Appendix A

Model Diagrams

The model description in Chapter 2 contains a general description of the system dynamics model. The following pages contains diagrams of the model structure taken directly from the Vensim® software used to build the model. Rectangles symbolize stocks, double lines represent flows; any text not inside a box is a model variable (or constant.) The blue arrows between variables, stocks and flows indicate that the item at the end of the arrow is a function of the item at the beginning of the arrow. Grayed out variables enclosed in angle brackets are variables from another diagram in the model. Some pipes go to a nebulous shape, called a cloud. Clouds are infinite sources or sinks for flows to pull from or dump into. For example, orders can be considered a flow into a backlog of orders; when orders are fulfilled they flow out of the backlog. But, they don't flow to any destination where they need to be tracked, so they flow to a cloud—and out of the model.

Not all of the structure built was utilized in the analyses. For instance, the model was built to accommodate many grades of scrap. This feature was not used in the analysis, but the structure remained–it was simply redundant, but left in place for future work.

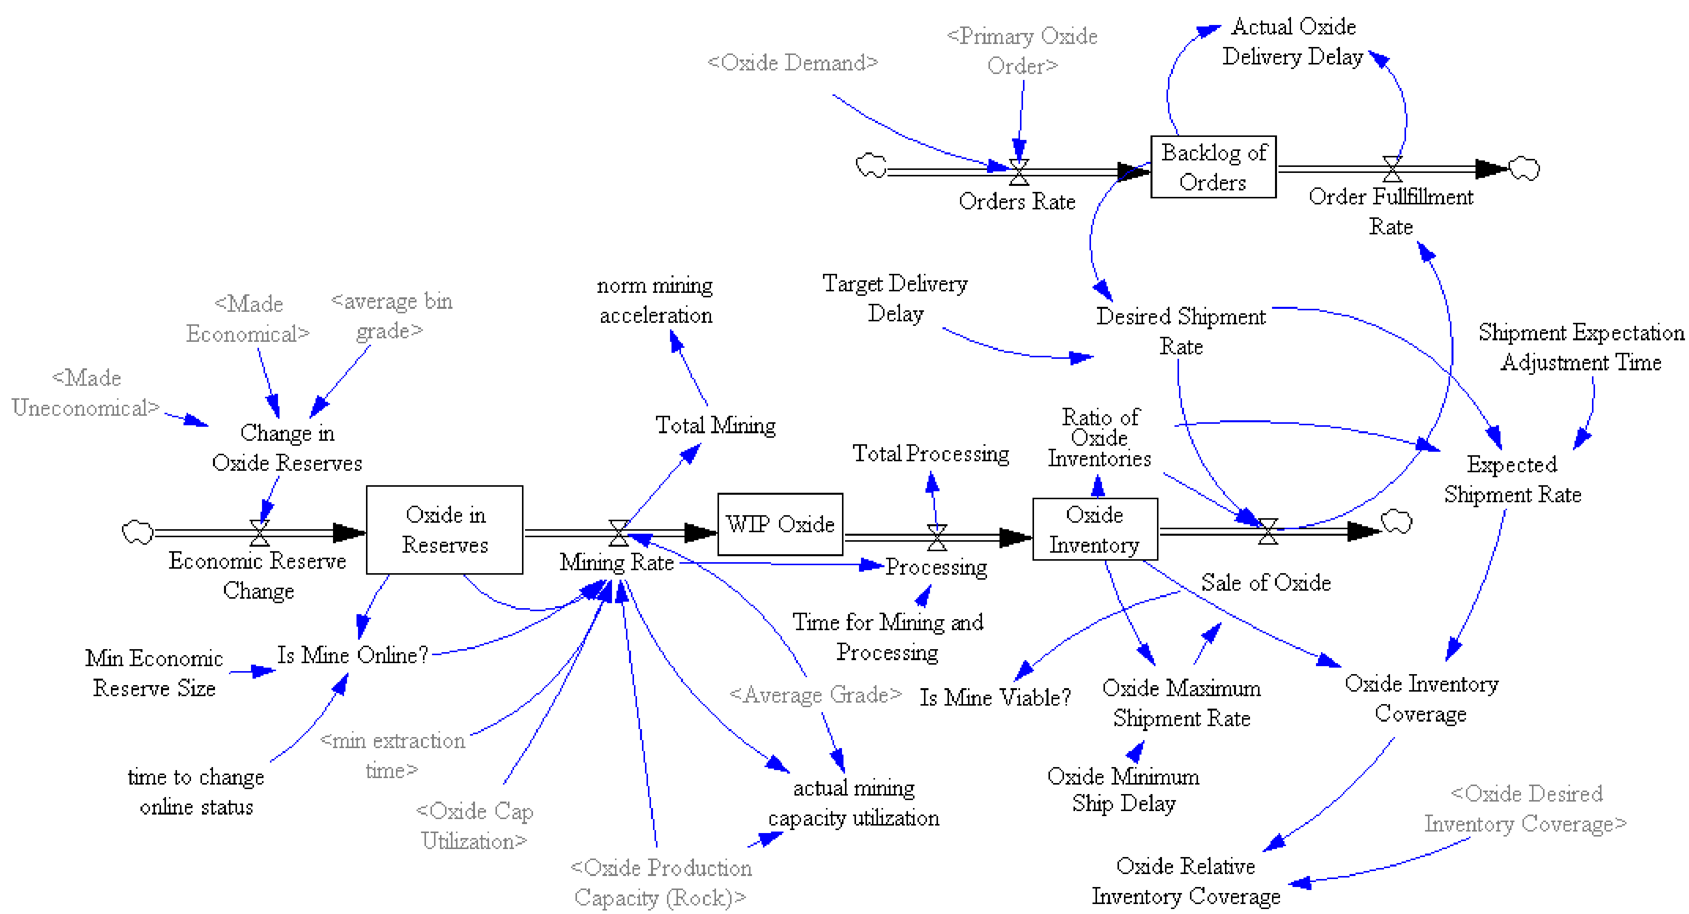

Figure A-1: Oxide Flow: This diagram shows the flow of oxide from being added to the reserve (discovery) to its shipment to customers.

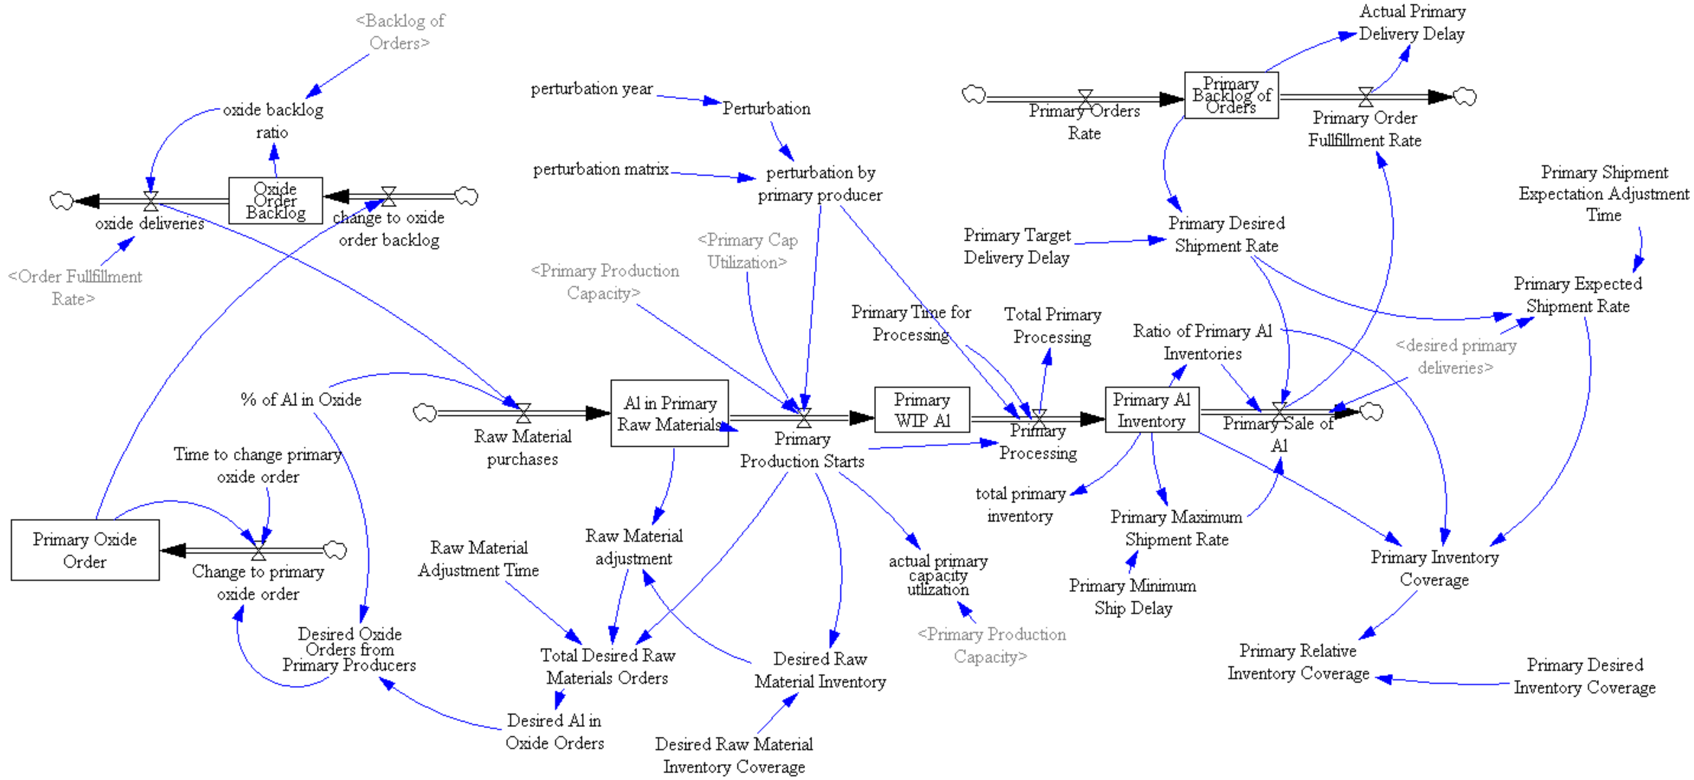

Figure A-2: Primary Flow: This diagram shows the flow of Primary Metal from being purchased in the form of oxide to its shipment out to customers.

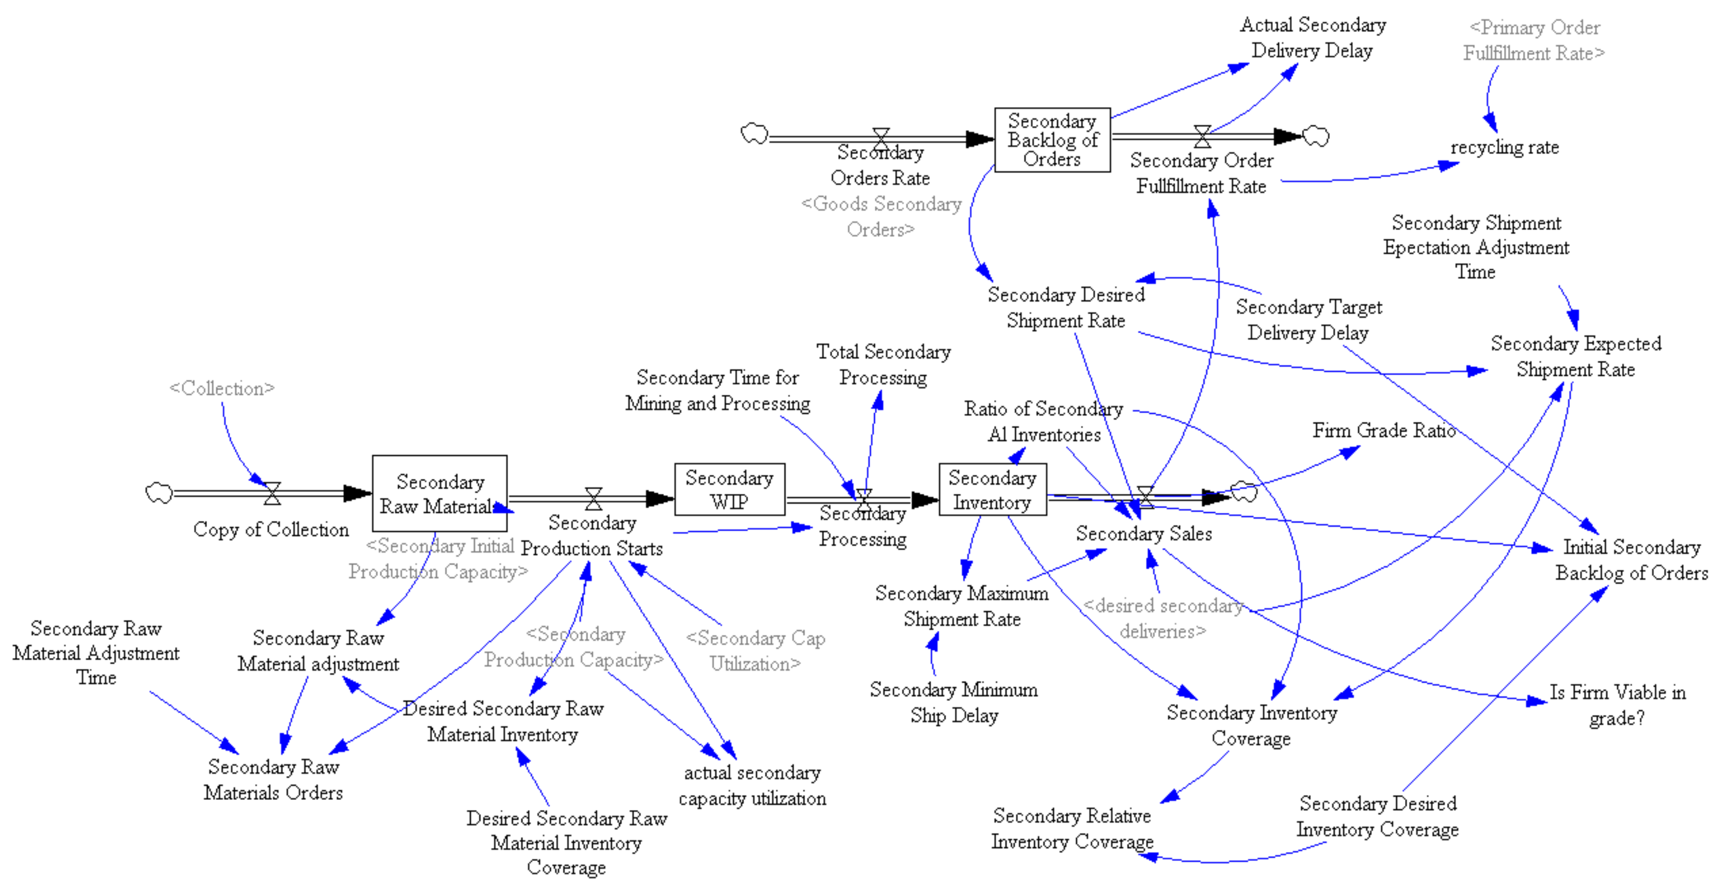

Figure A-3: page 3

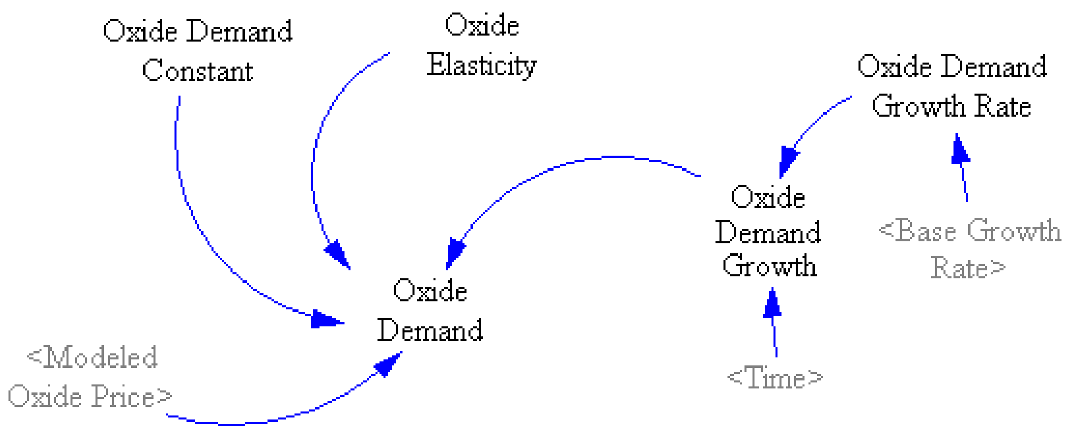

Figure A-4: Oxide Demand as represented by a constant elasticity demand curve.

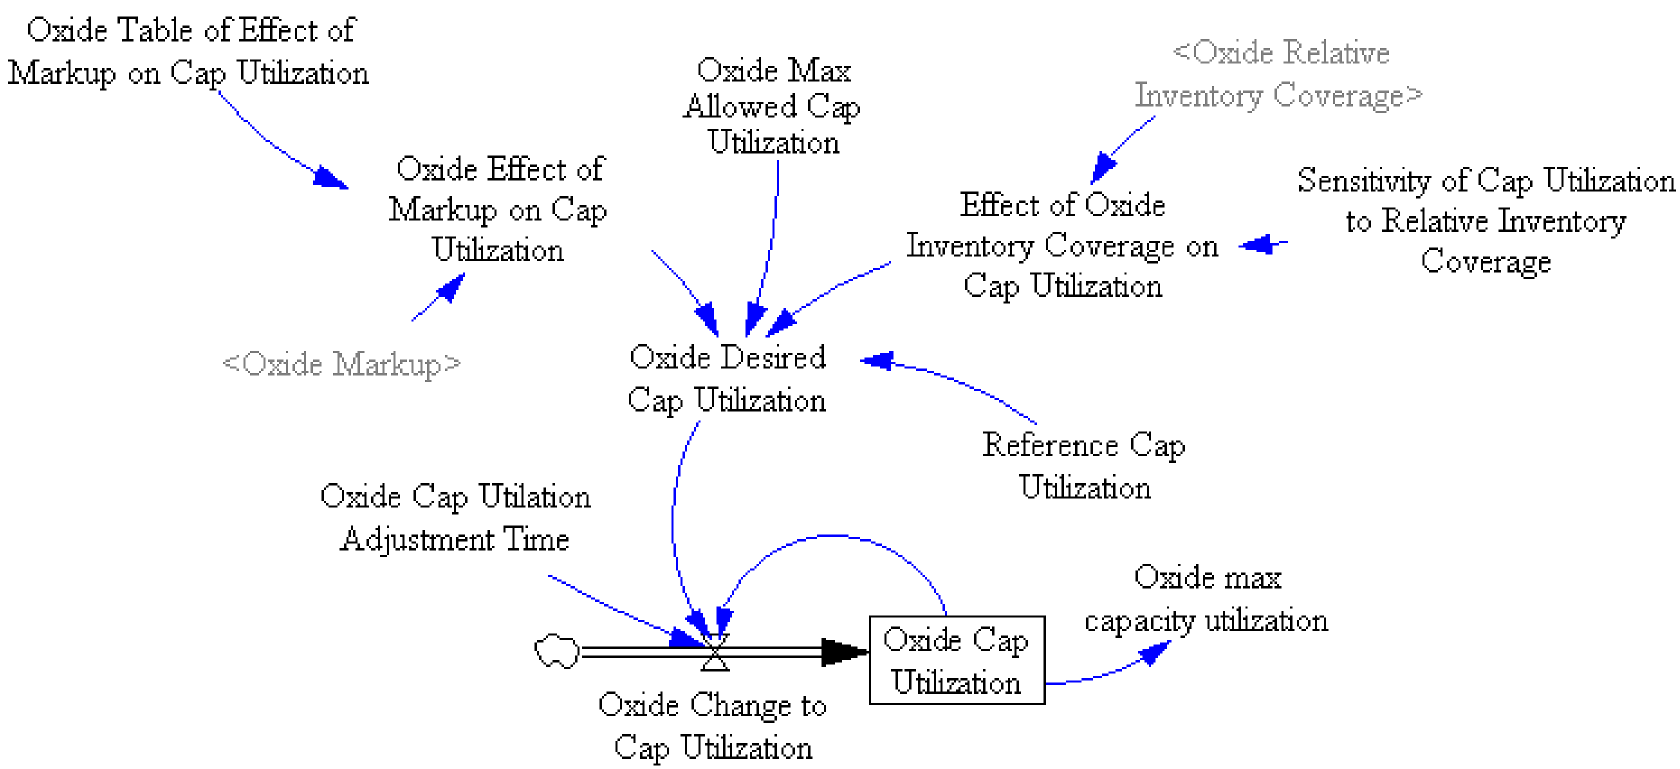

Figure A-5: Oxide Capacity Utilization

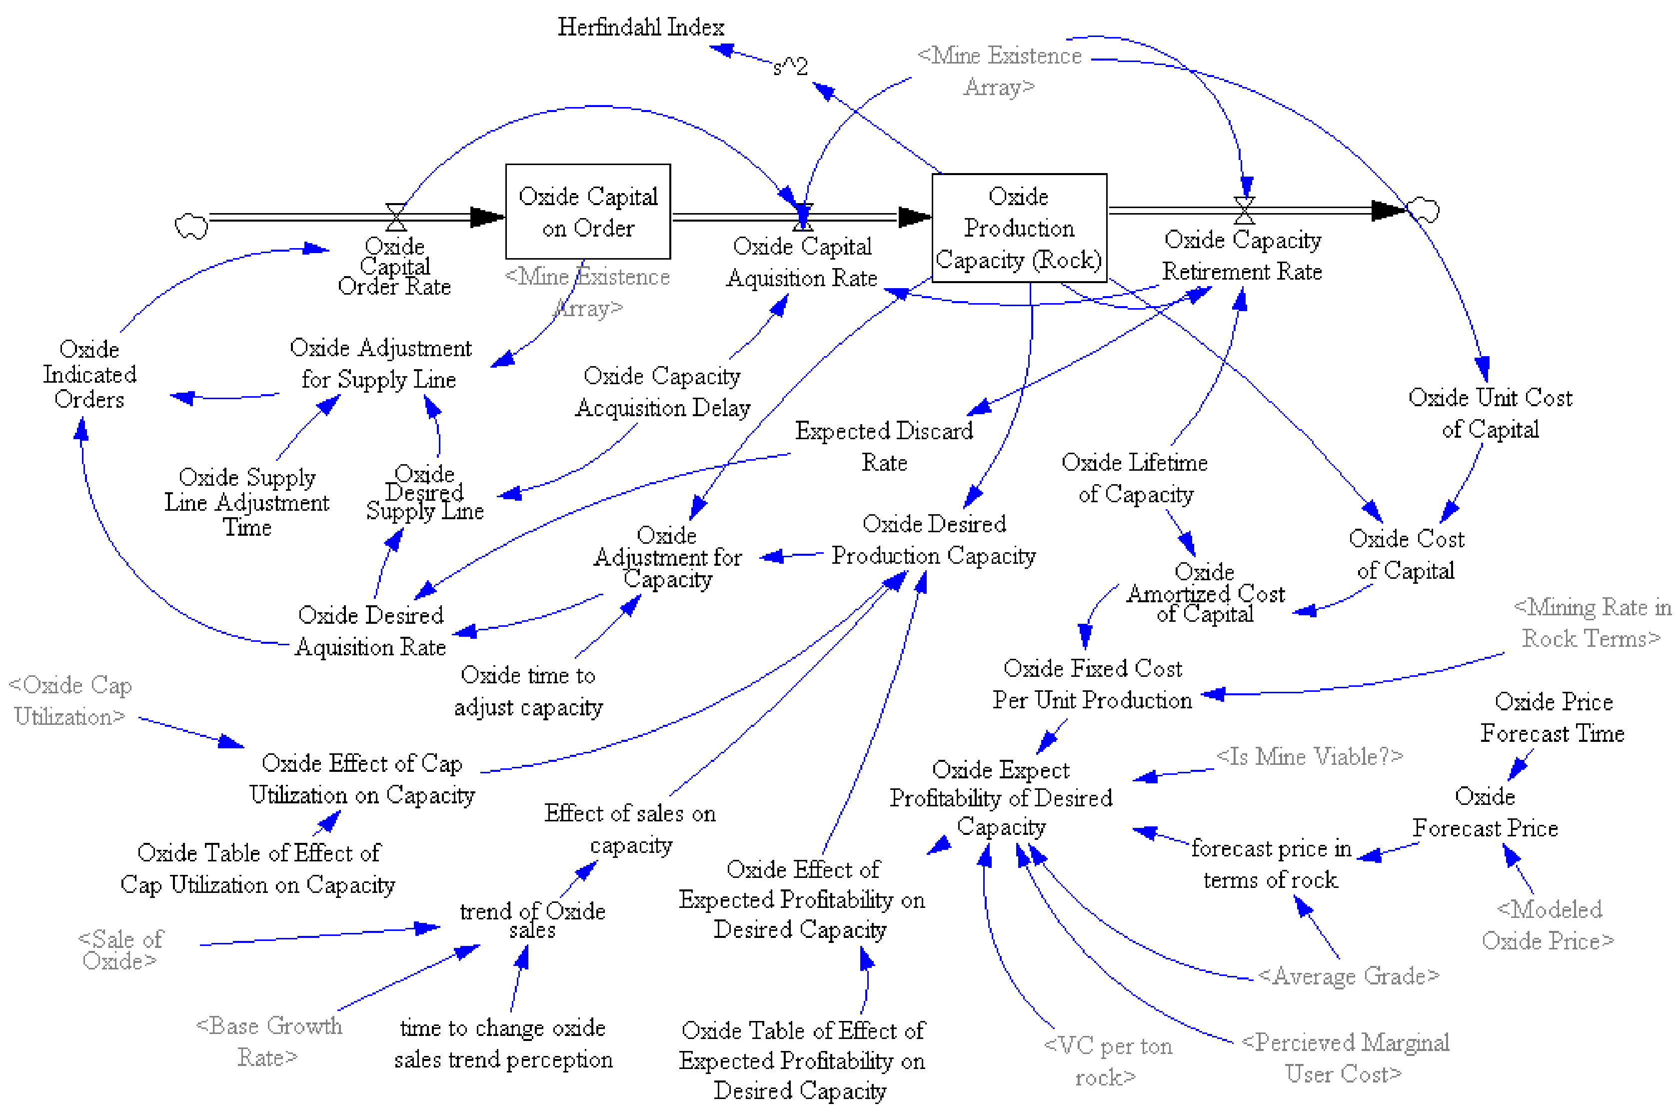

Figure A-6: Oxide Capacity

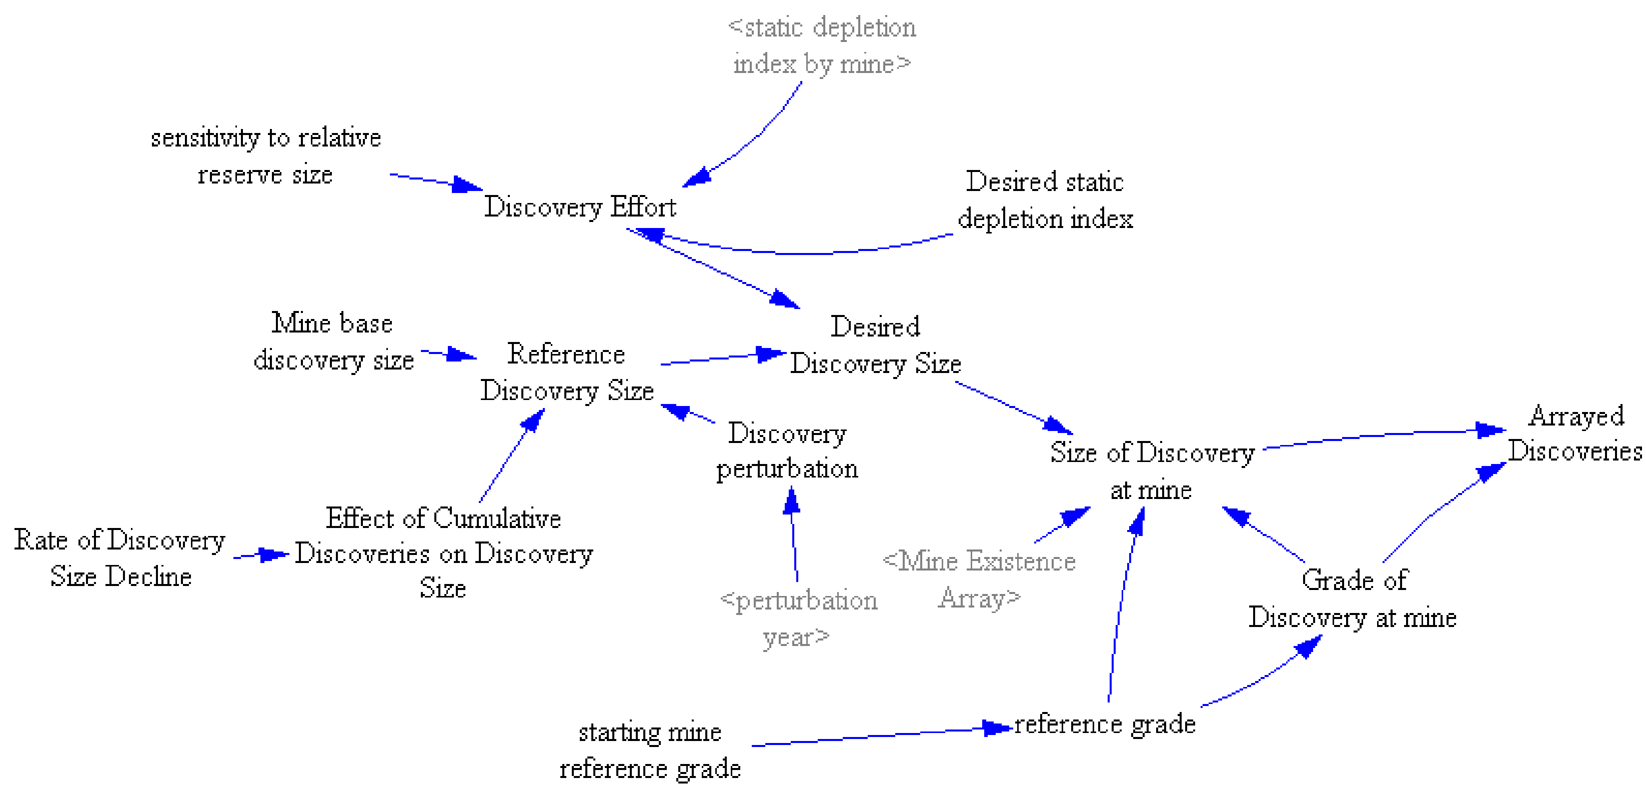

Figure A-7: Discovery: The structure governing the discovery of new ore deposits. It is driven by the static depletion indices of the individual mines–lower indices, more effort into discovery.

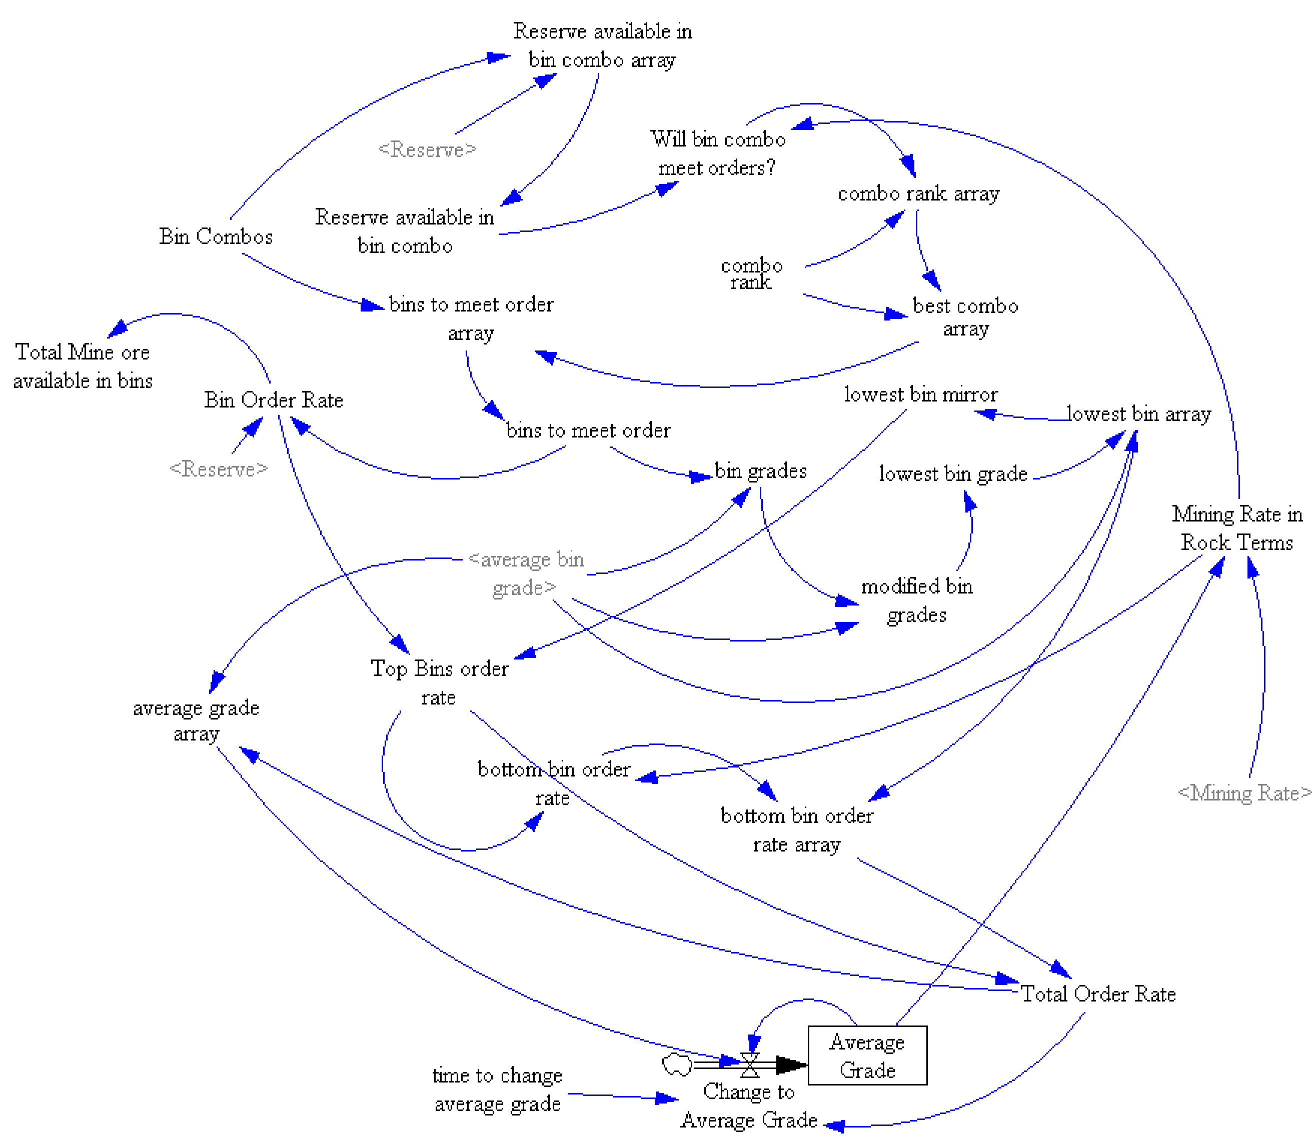

Figure A-8: Bin Control: Ore is binned by quality. Mines use the ore from the largest bins first, because it is the cheapest. This structure controls the pulling of ore from bins, and calculates the average grade of the ore being mined.

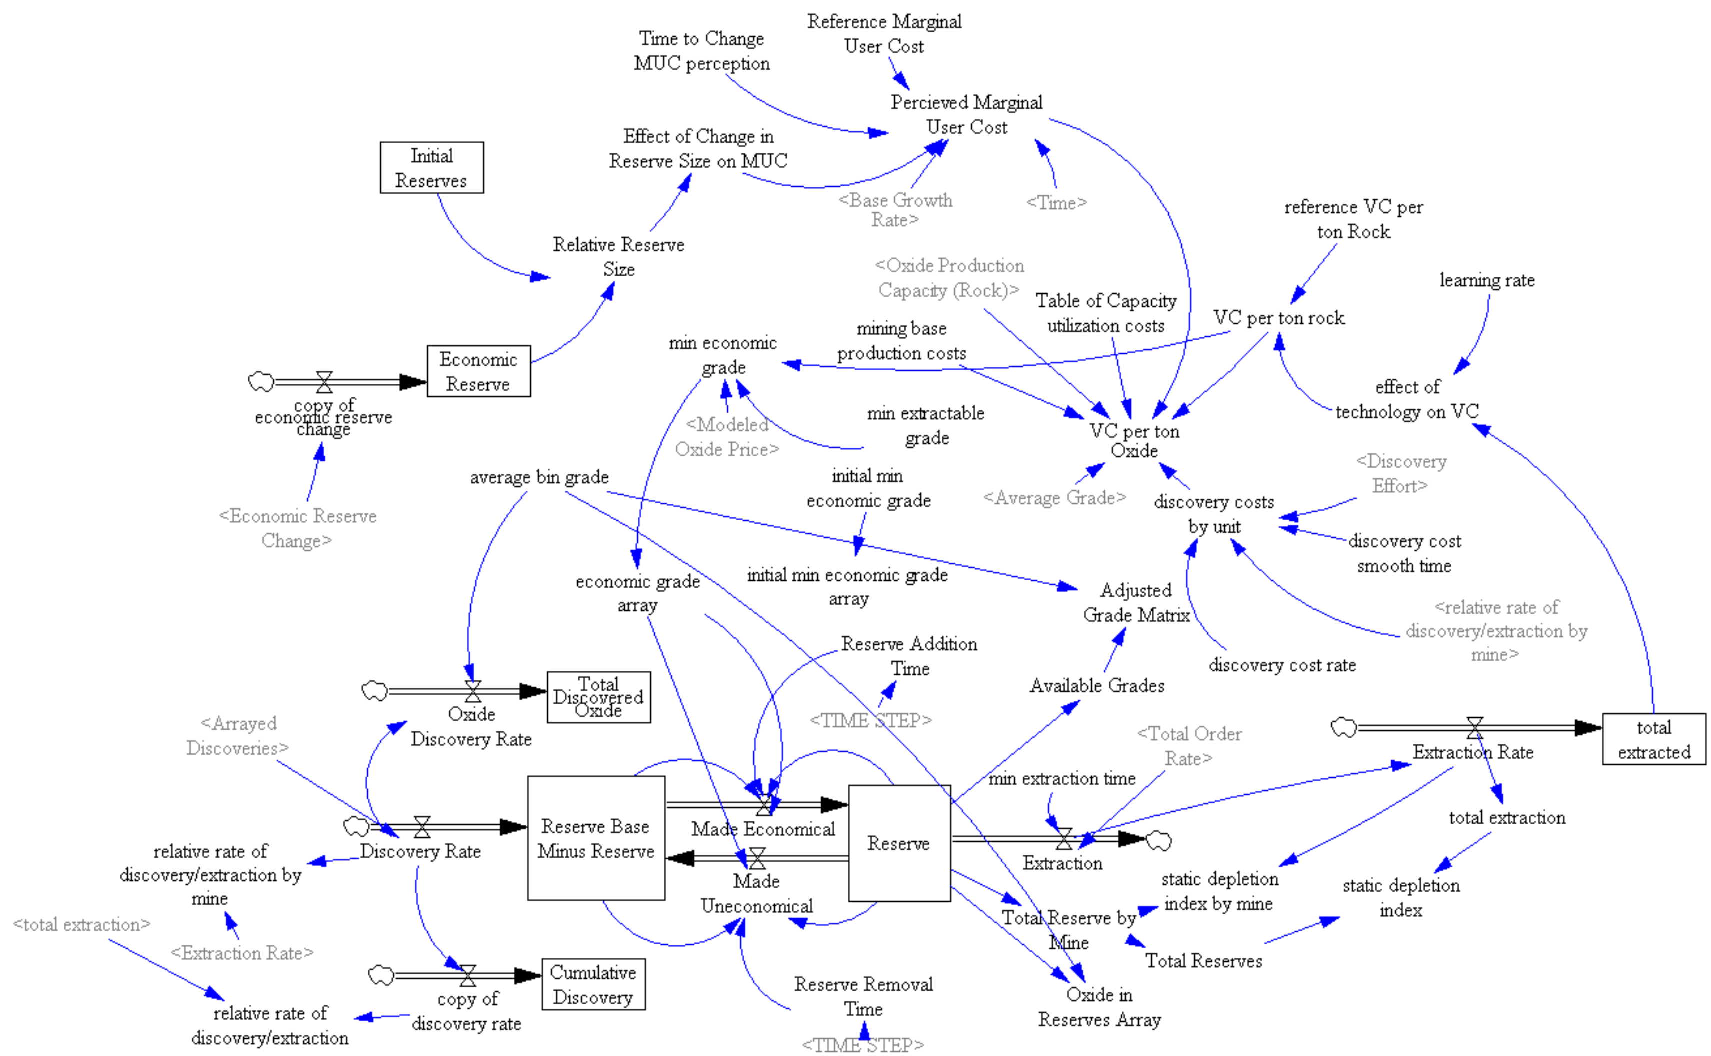

Figure A-9: Mining Reserves and Costs: This structure tracks the costs of mining and the movement of ore in and out of the reserve (based on the current market price and costs of mining)

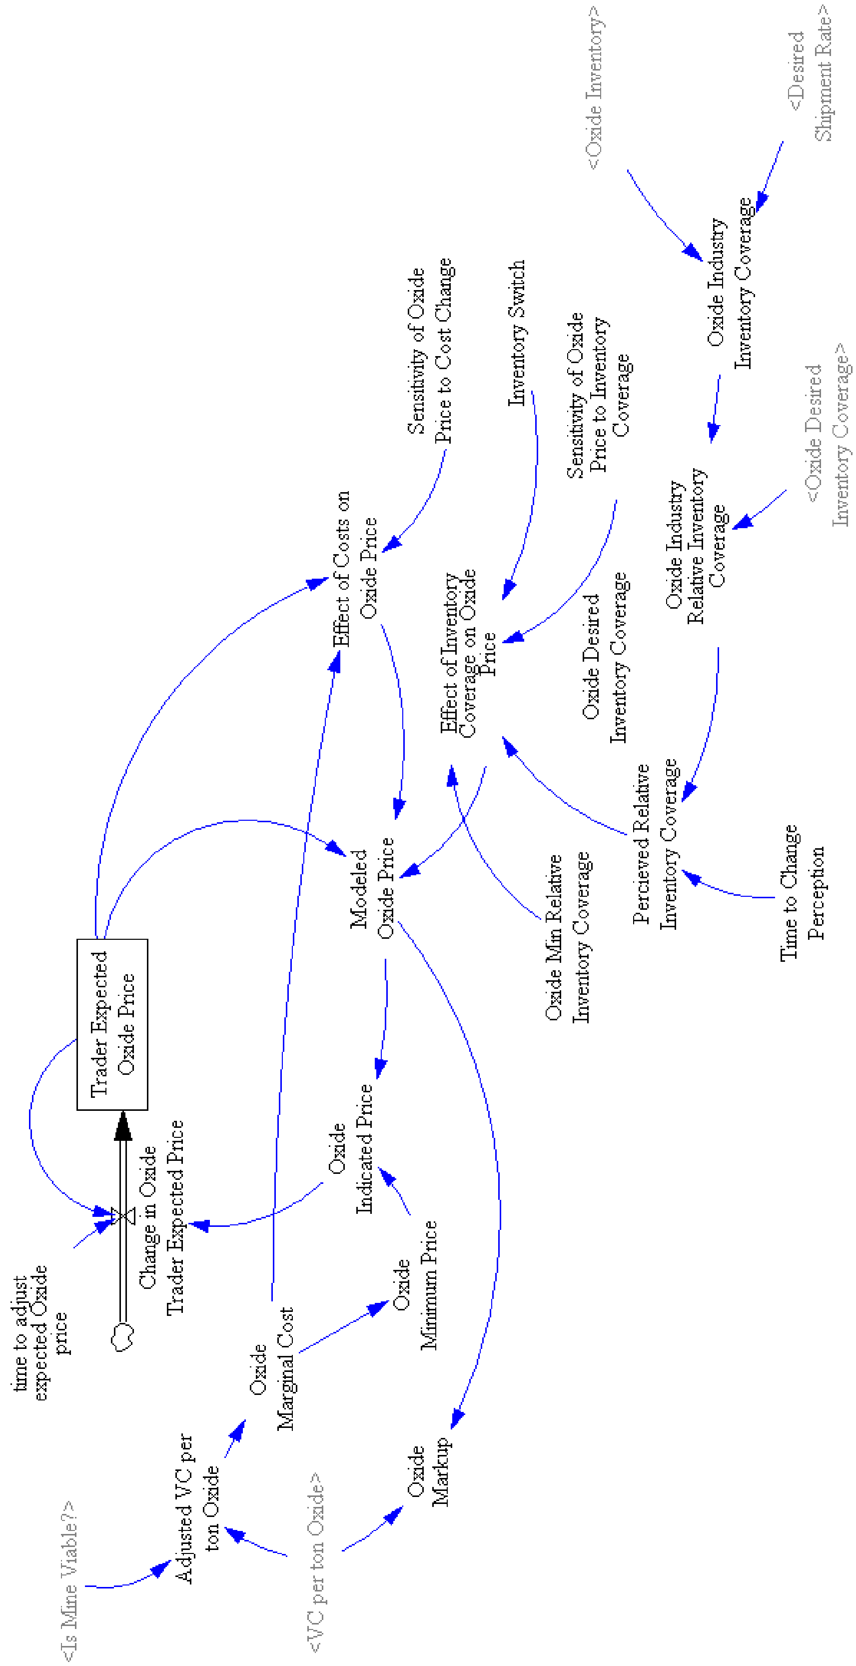

Figure A-10: Oxide Price

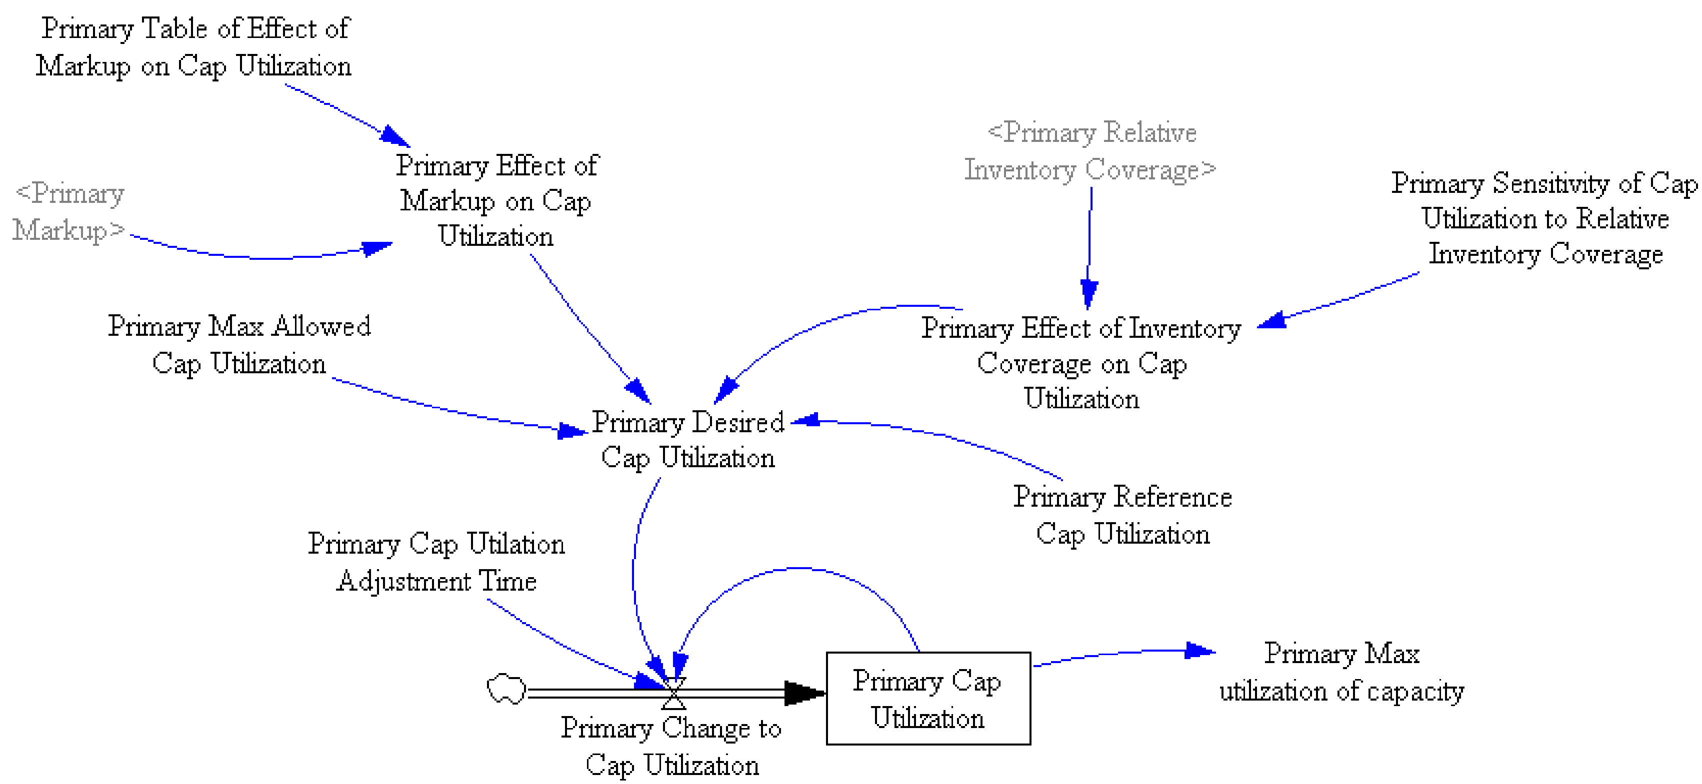

Figure A-11: Primary Capacity Utilization

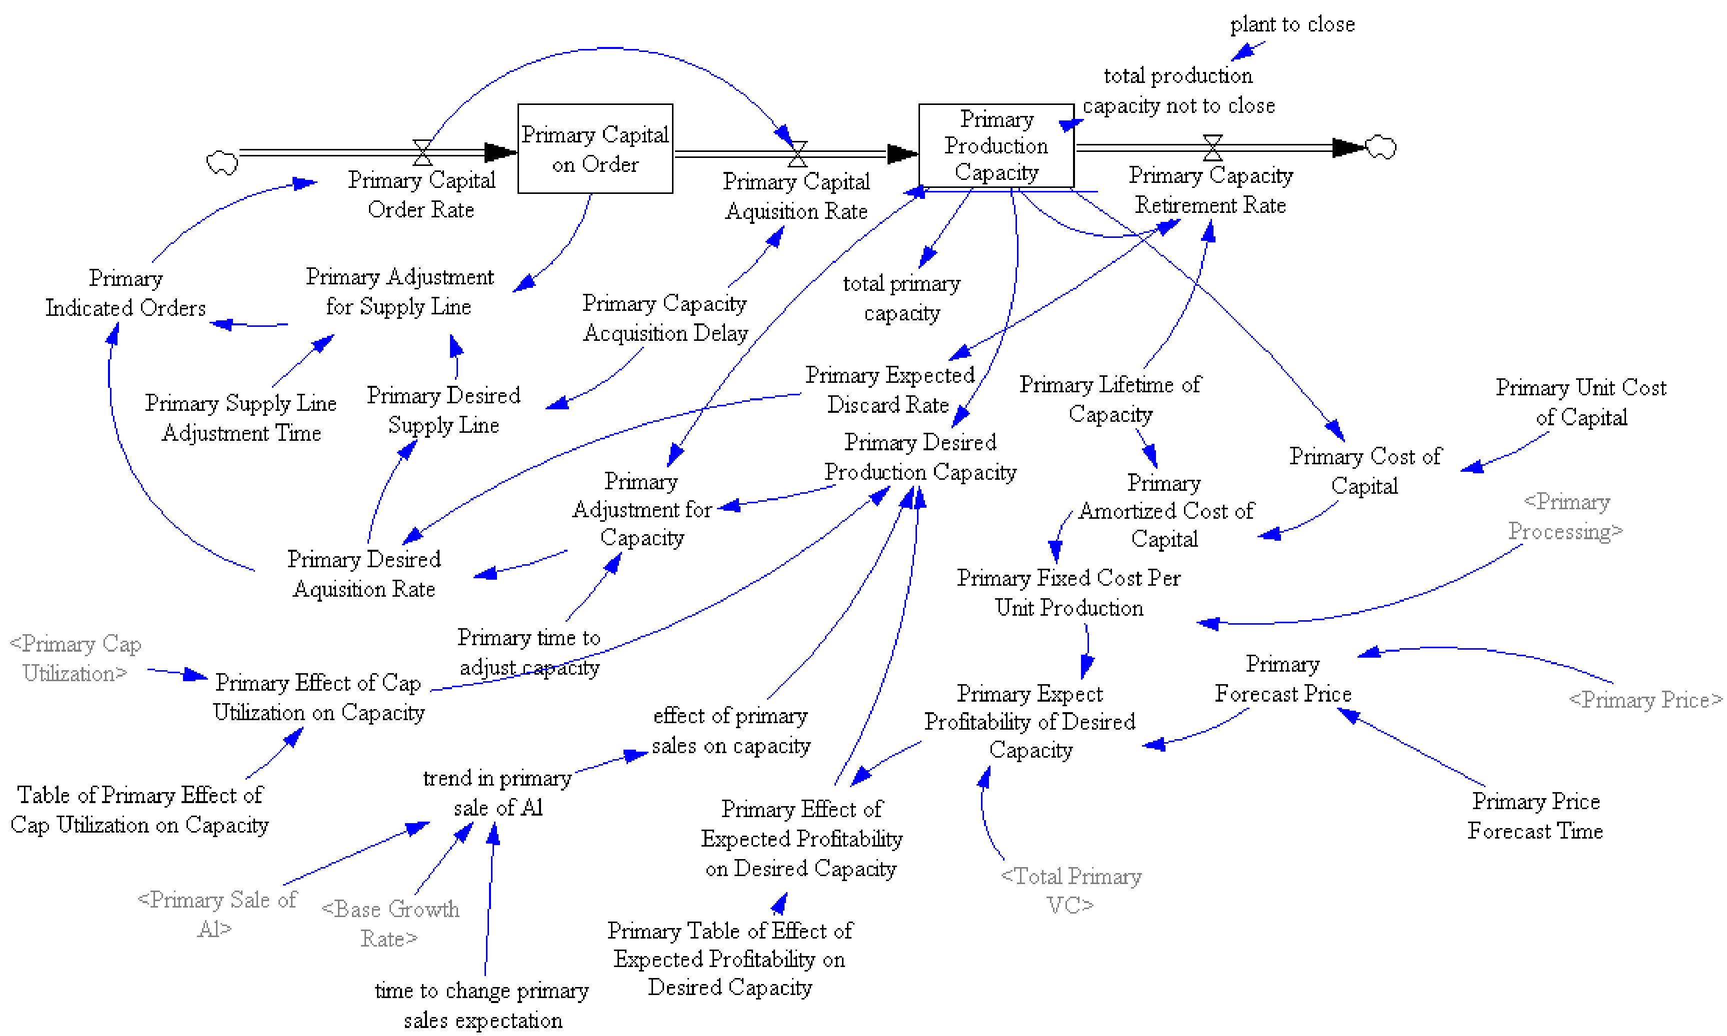

Figure A-12: Primary Capacity

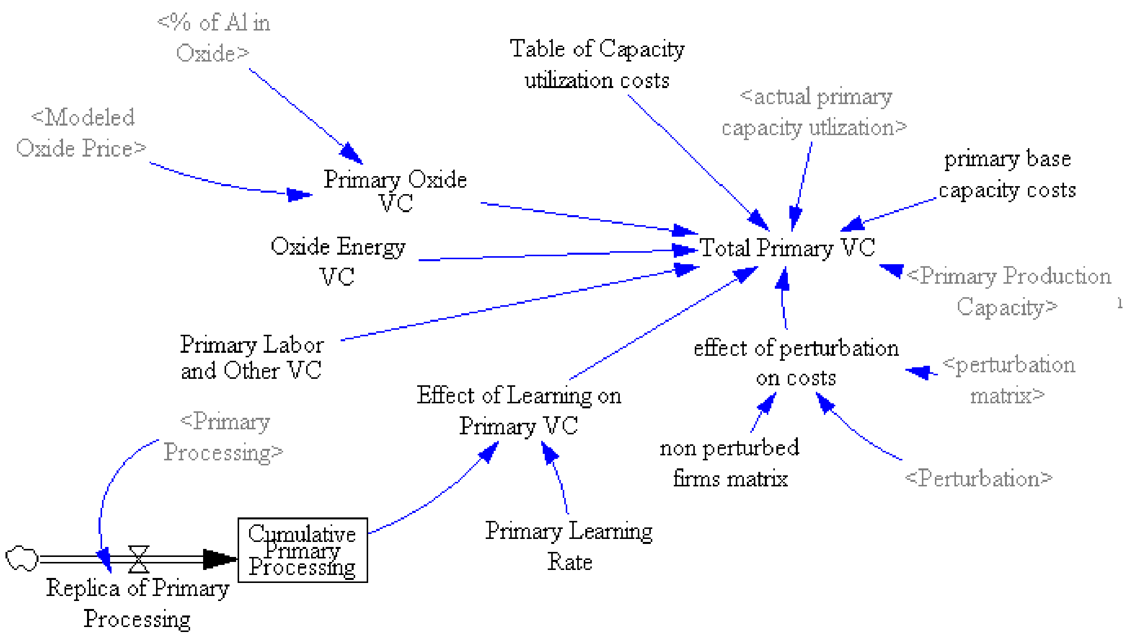

Figure A-13: Primary Variable Costs

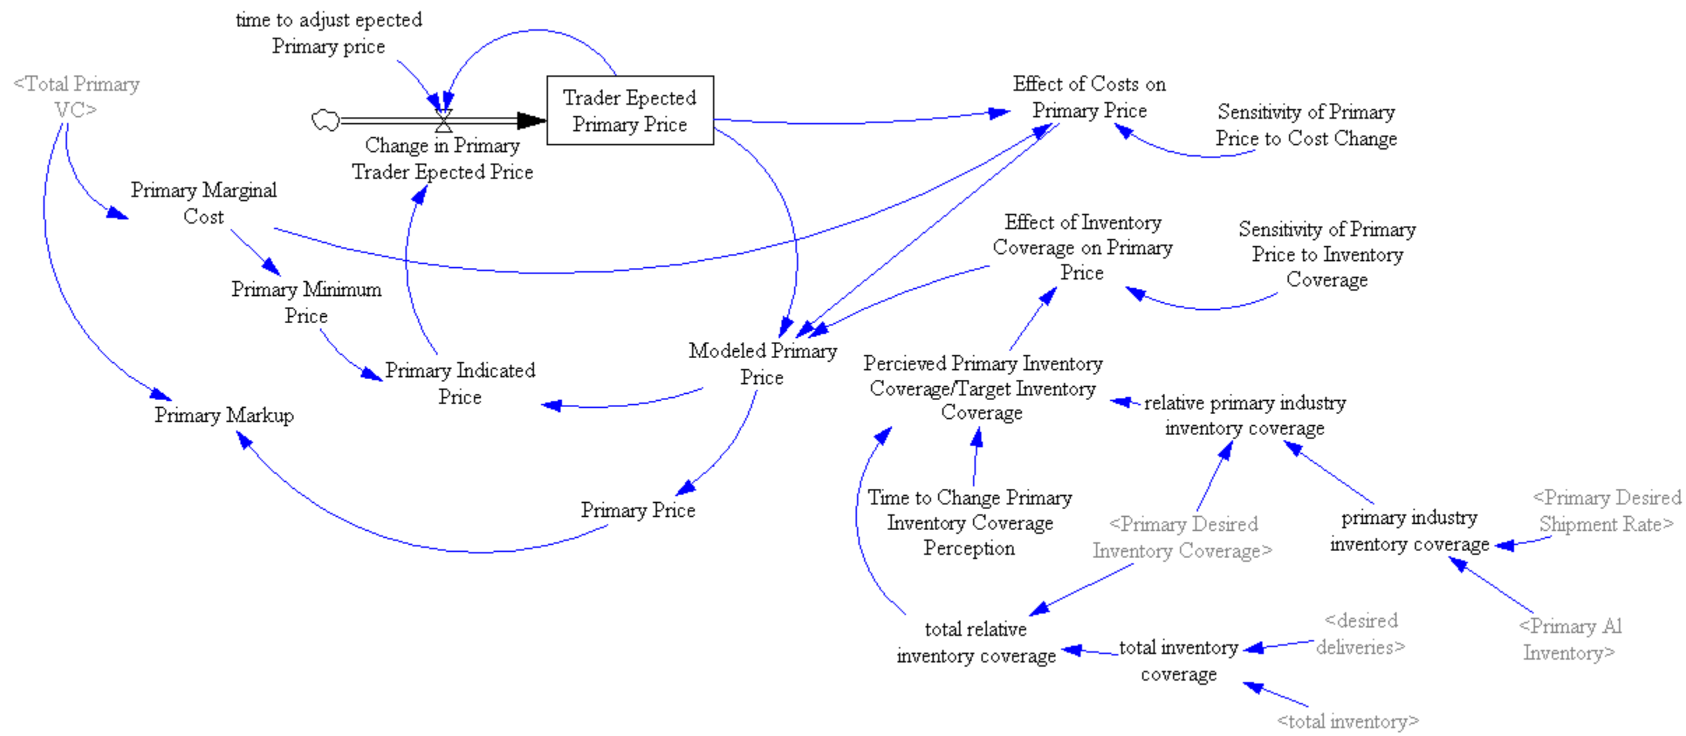

Figure A-14: Primary Price

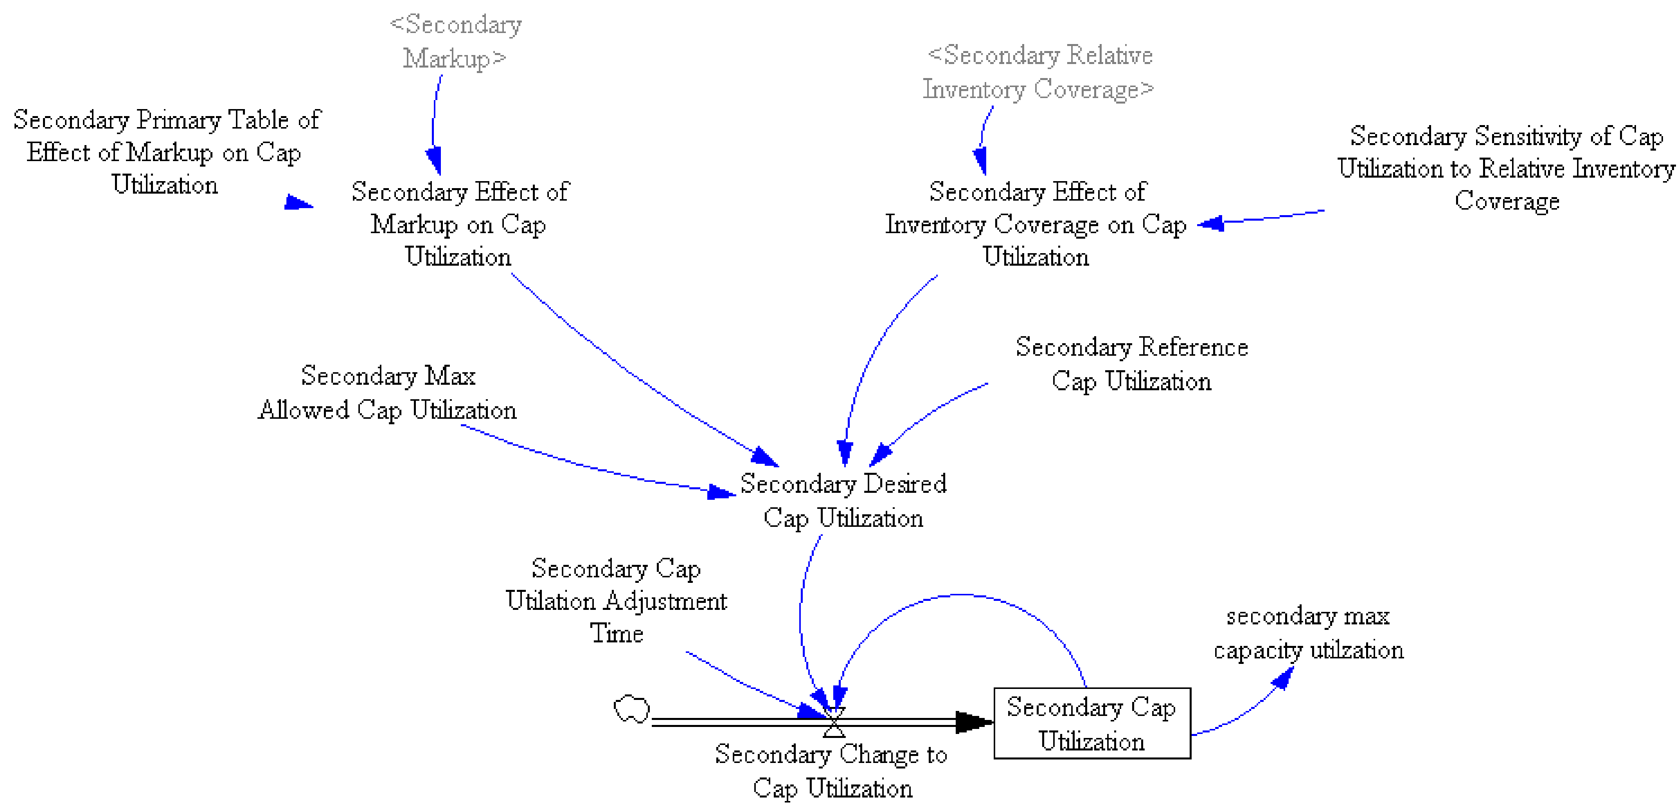

Figure A-15: Secondary Capacity Utilization

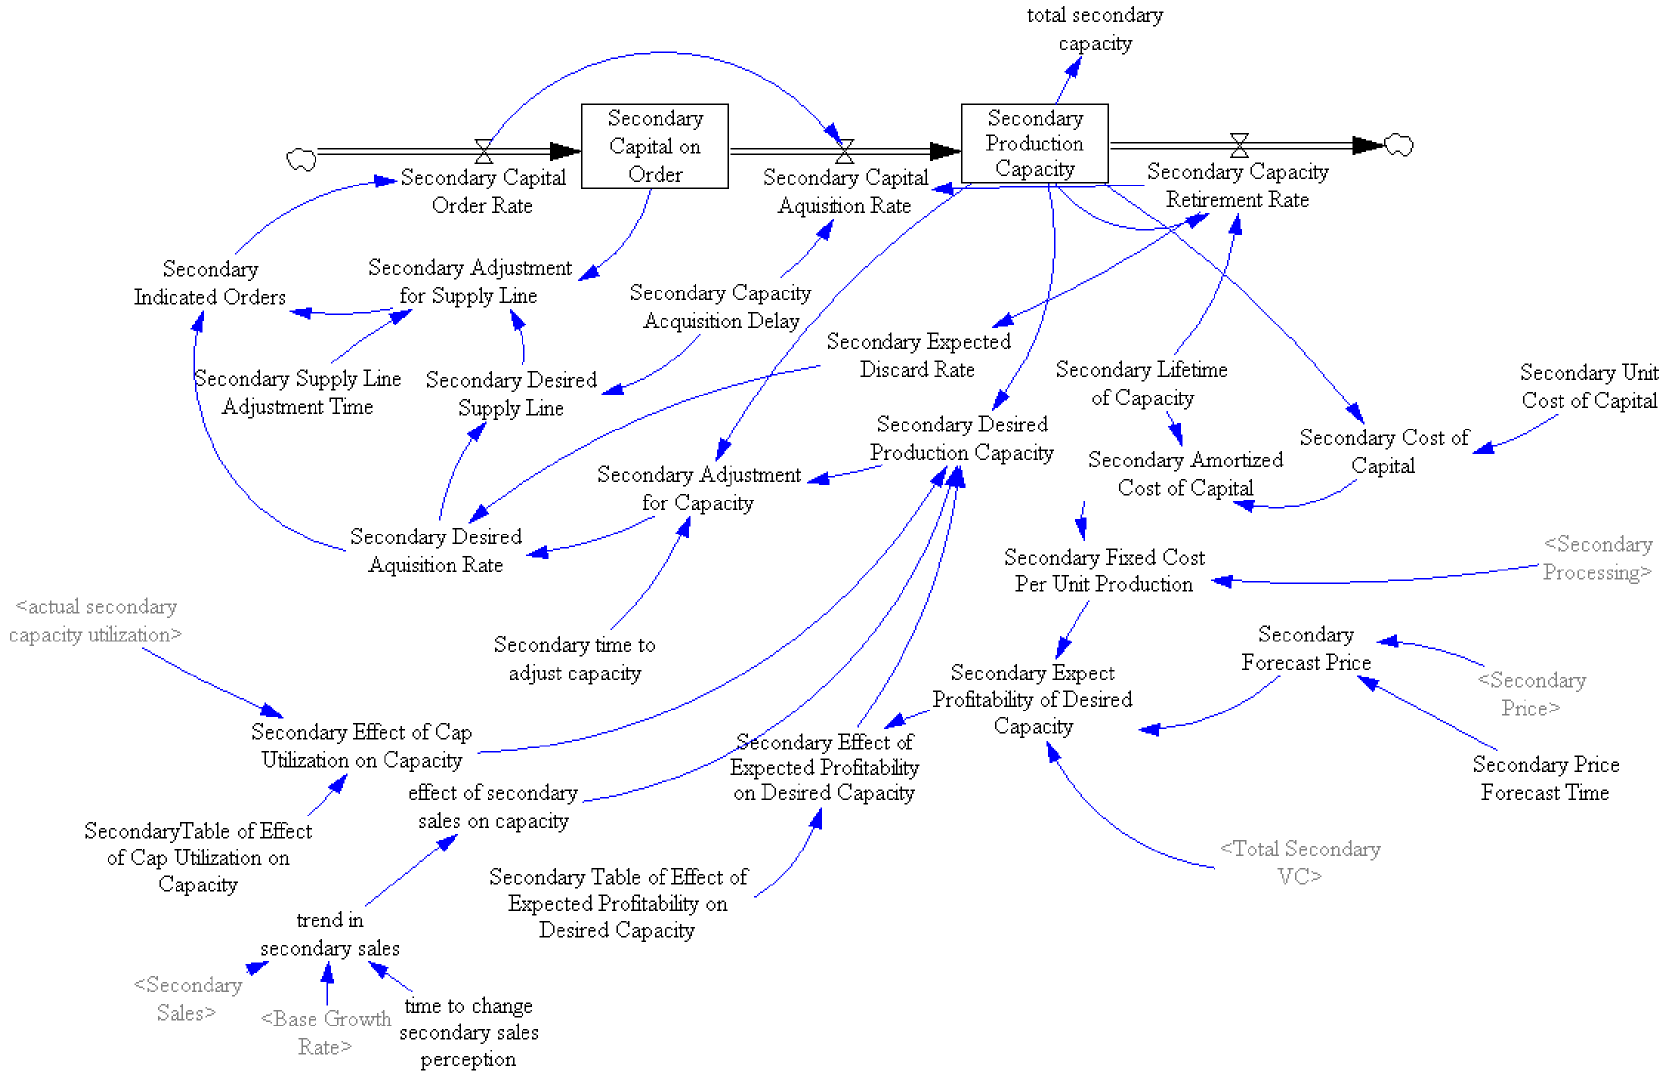

Figure A-16: Secondary Capacity

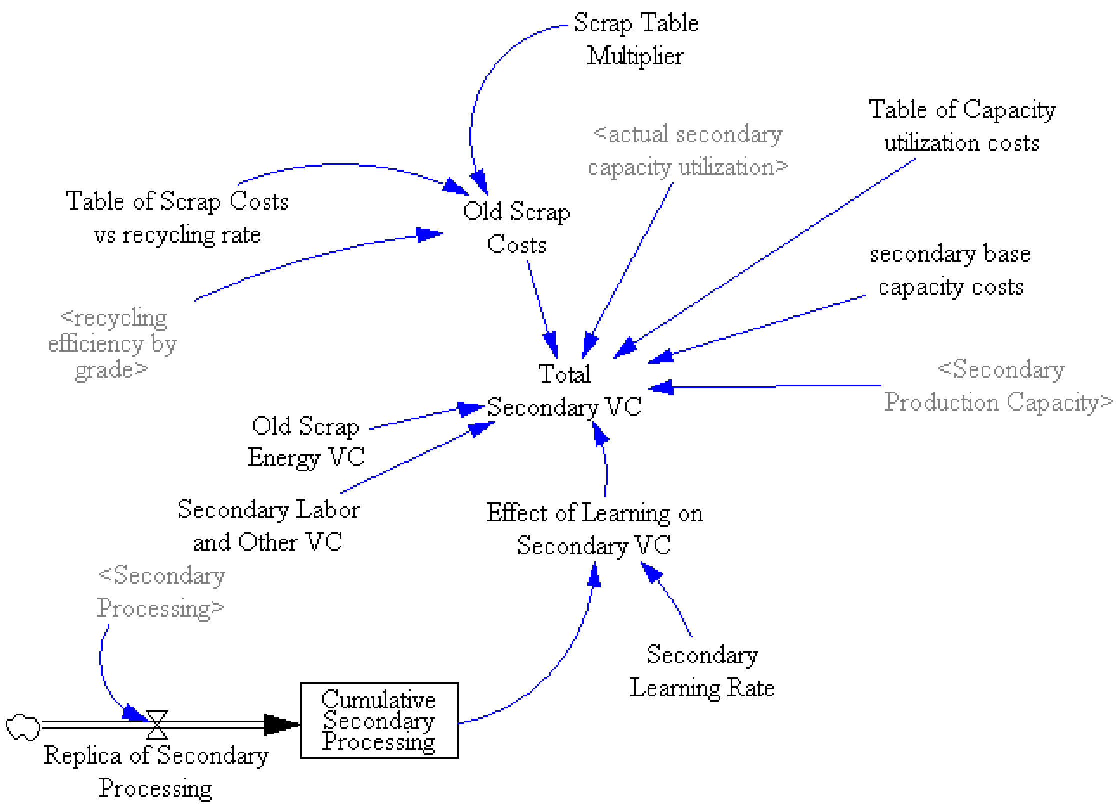

Figure A-17: Secondary Variable Costs

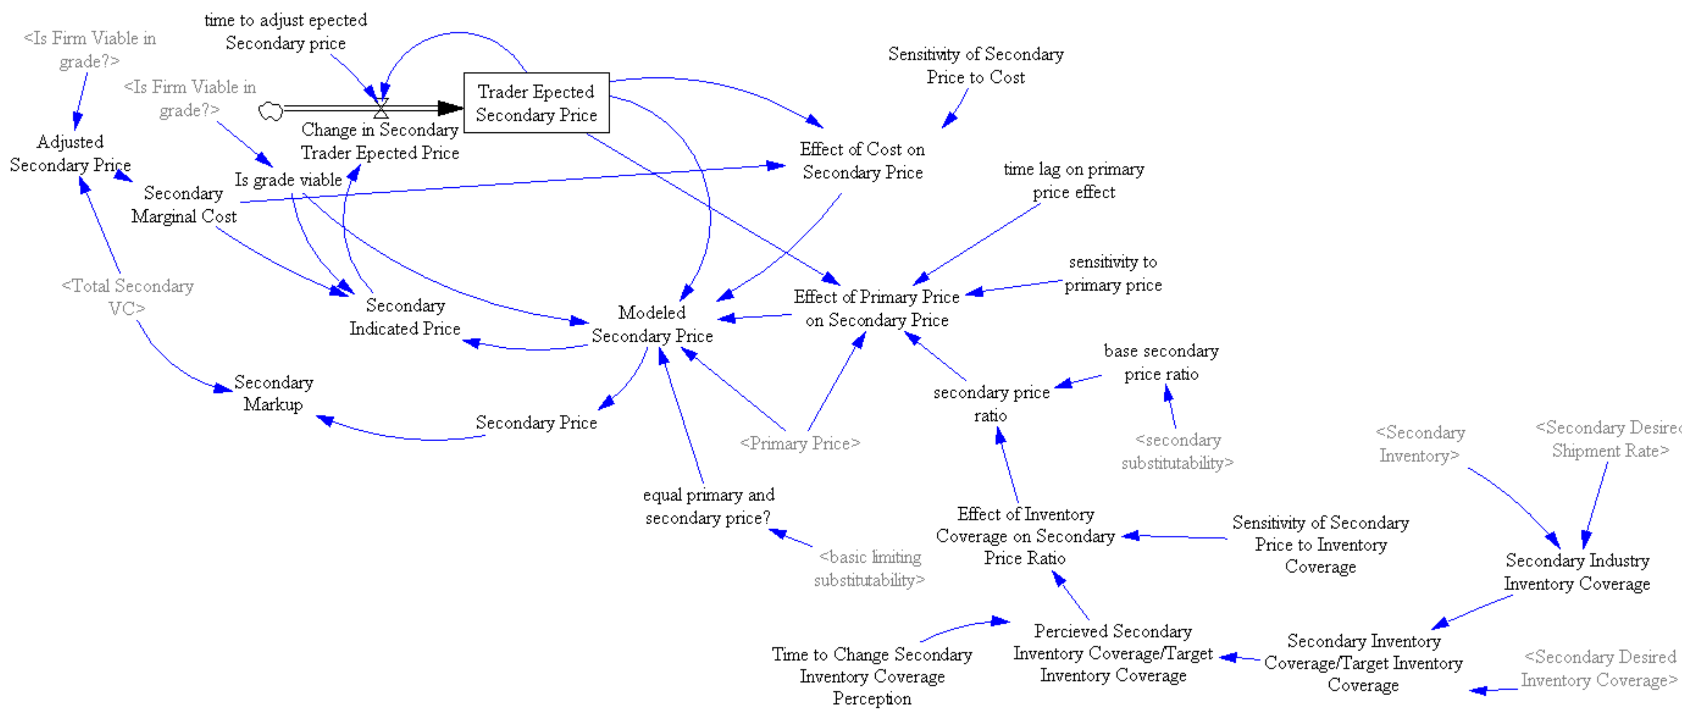

Figure A-18: Secondary Price

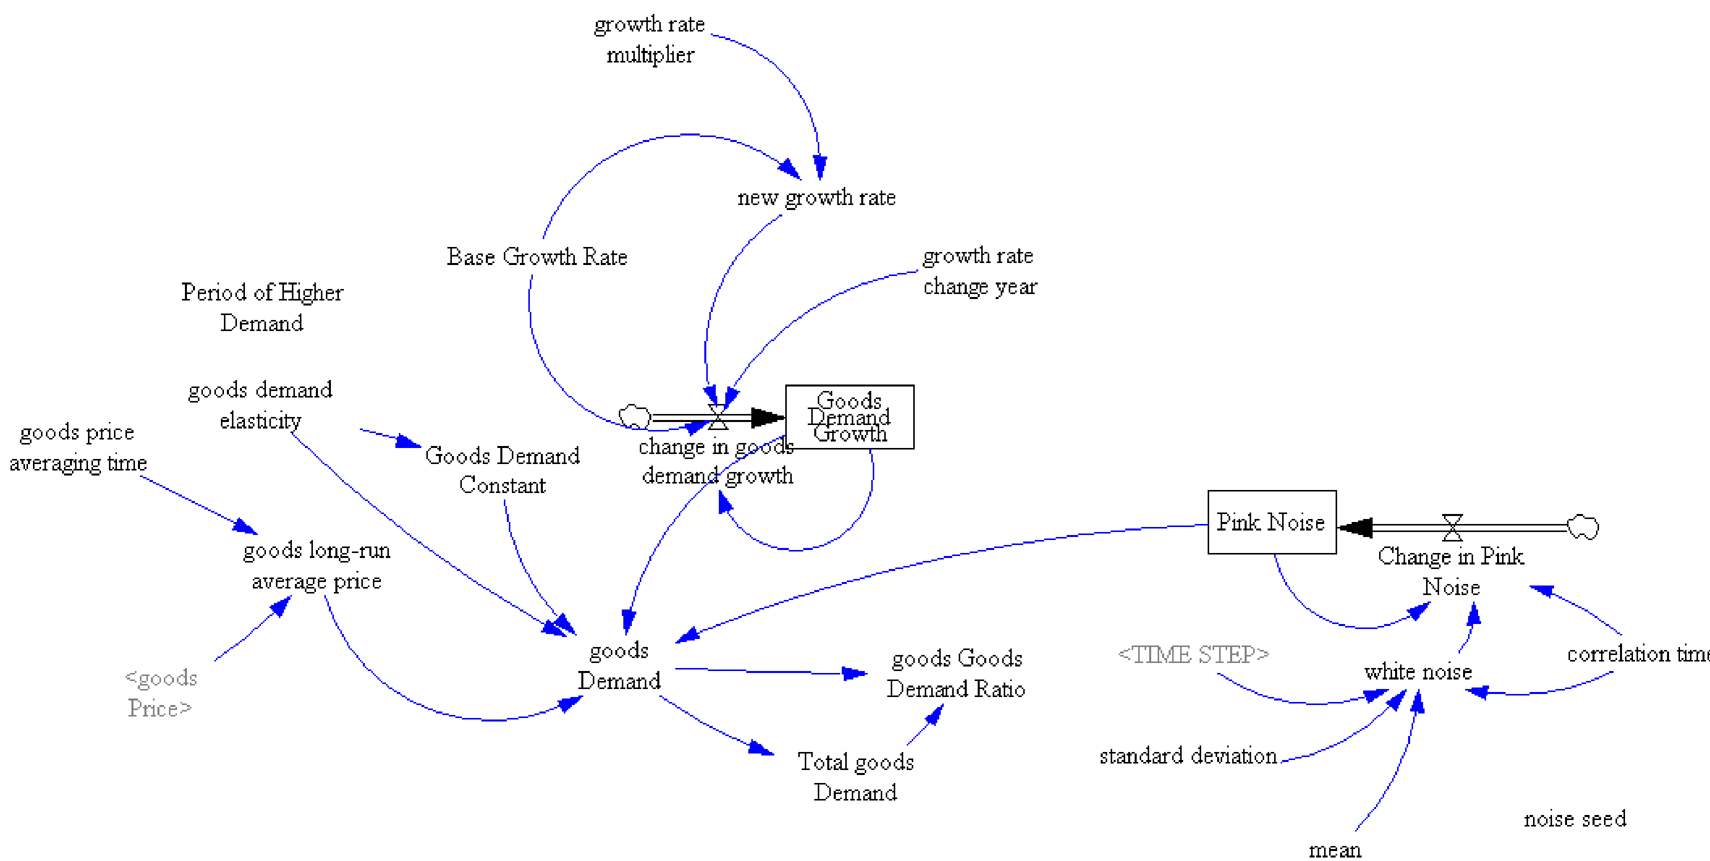

Figure A-19: Goods Demand

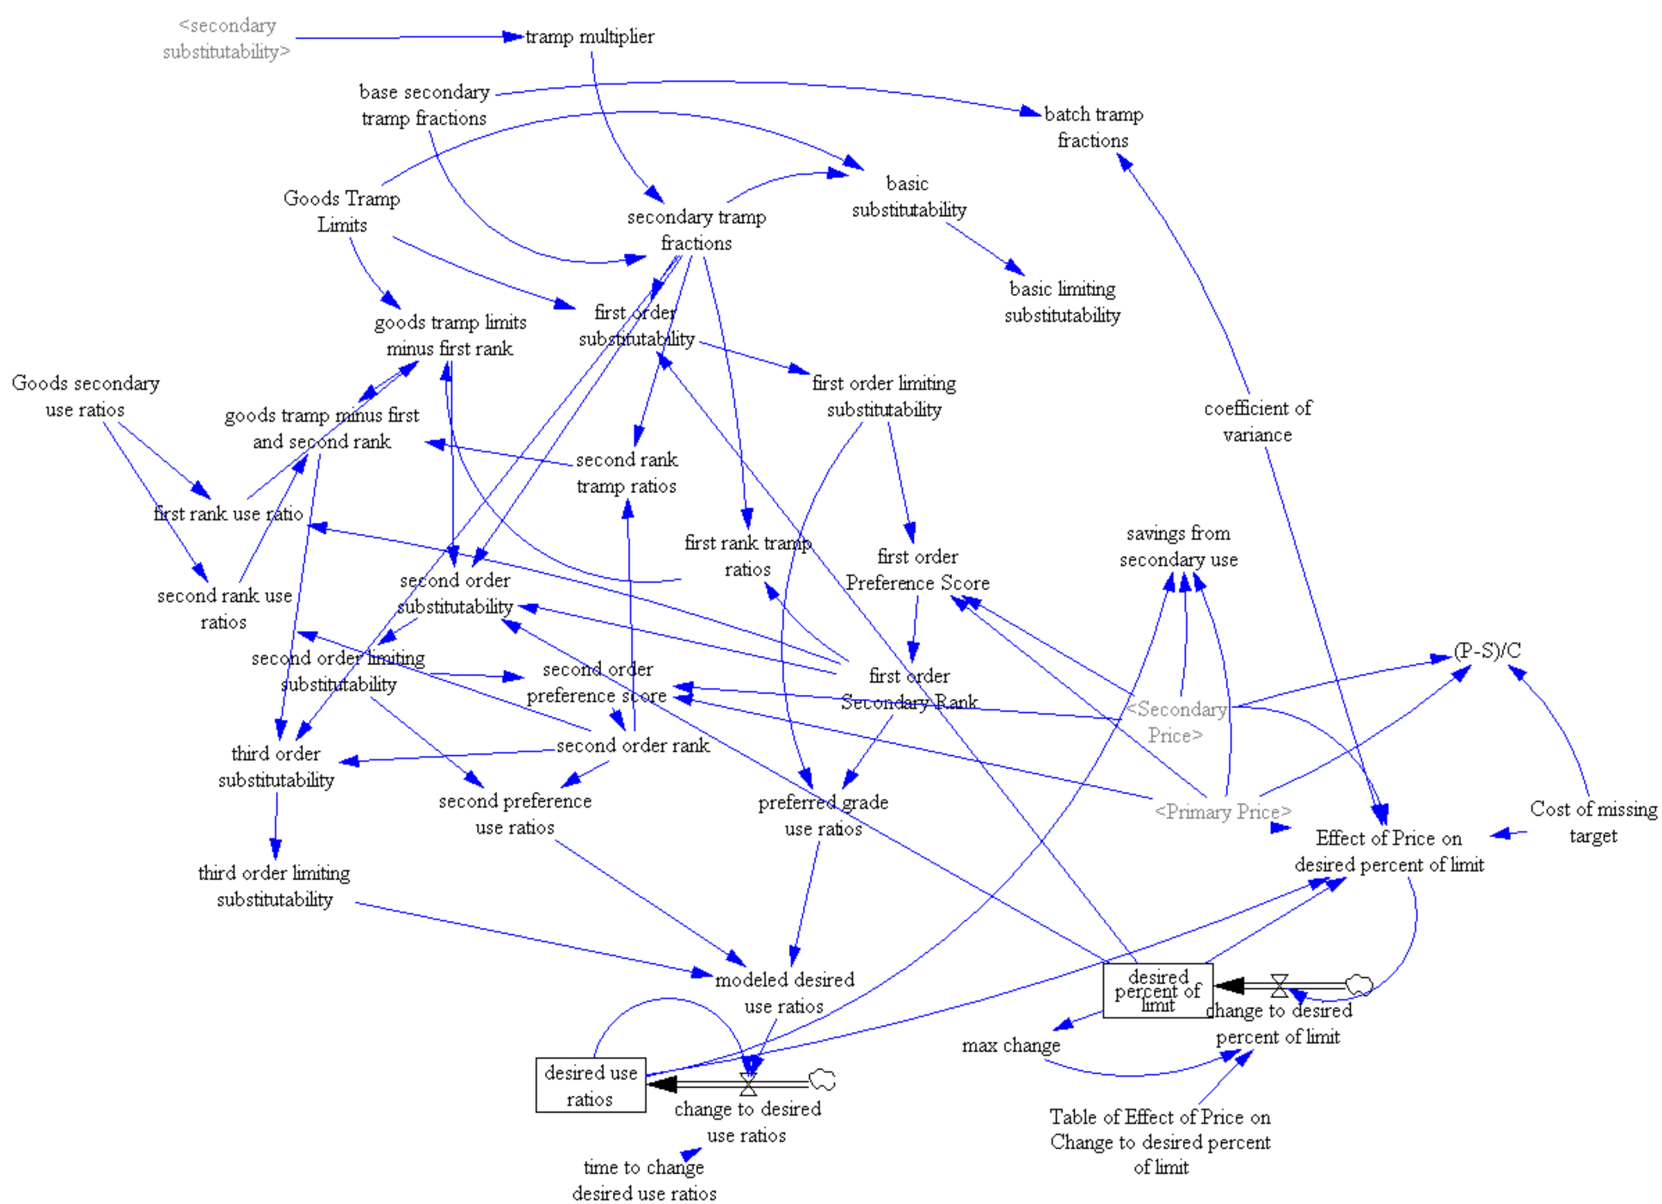

Figure A-20: Desired Secondary Use Ratios

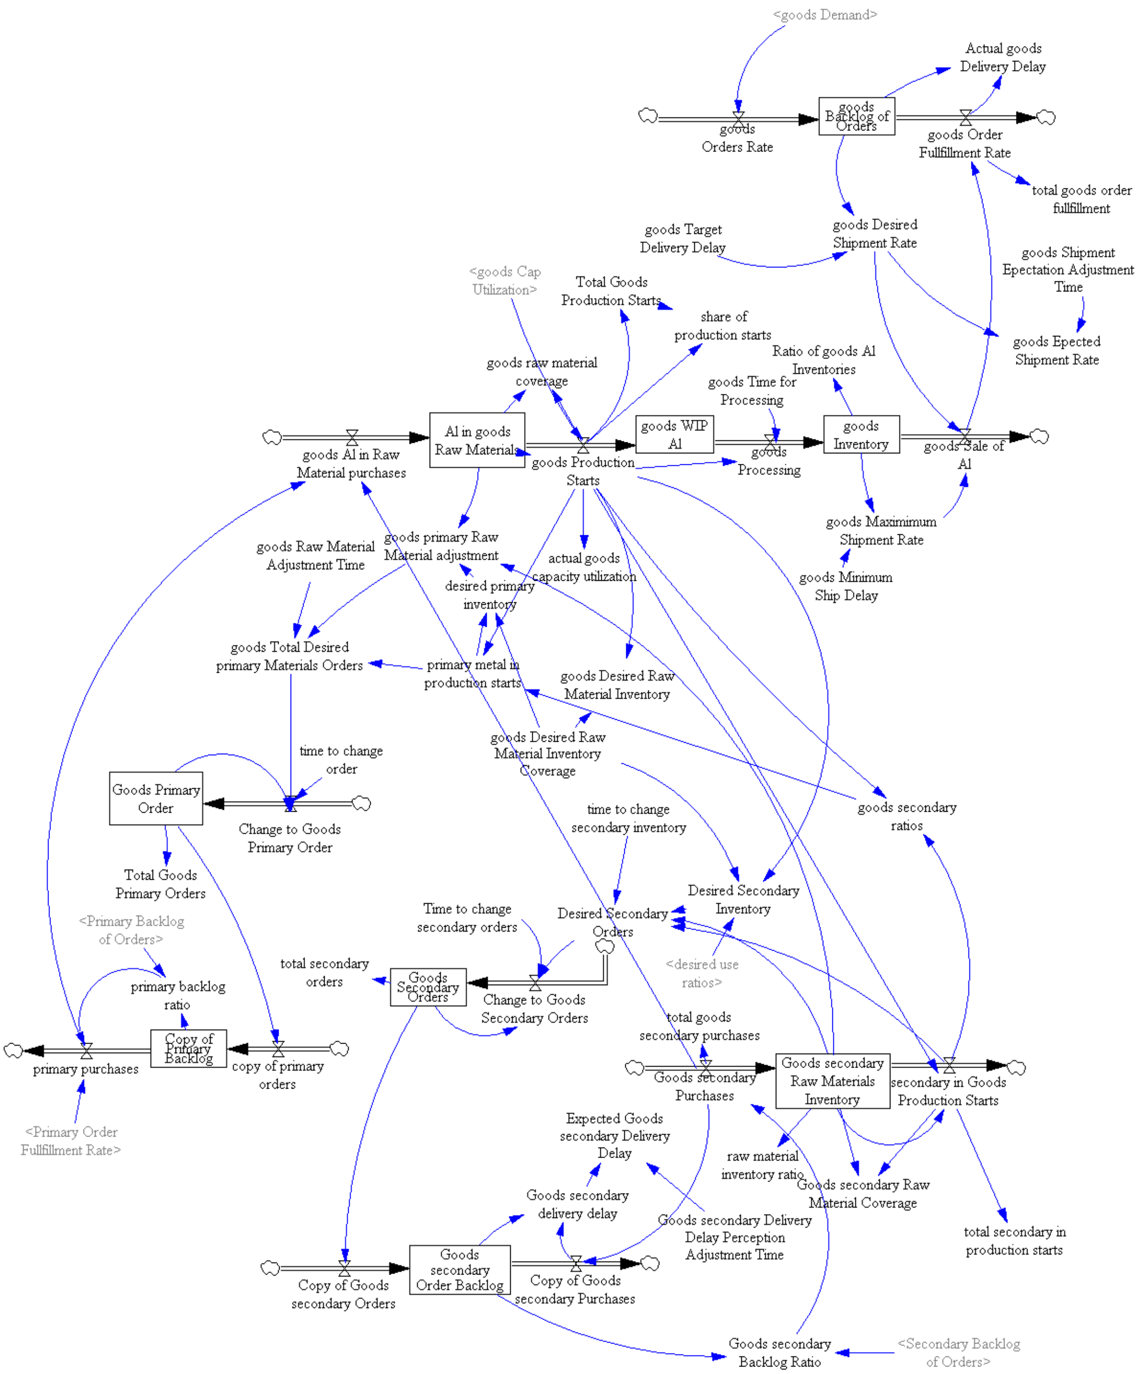

Figure A-21: Desired Secondary Use Ratios

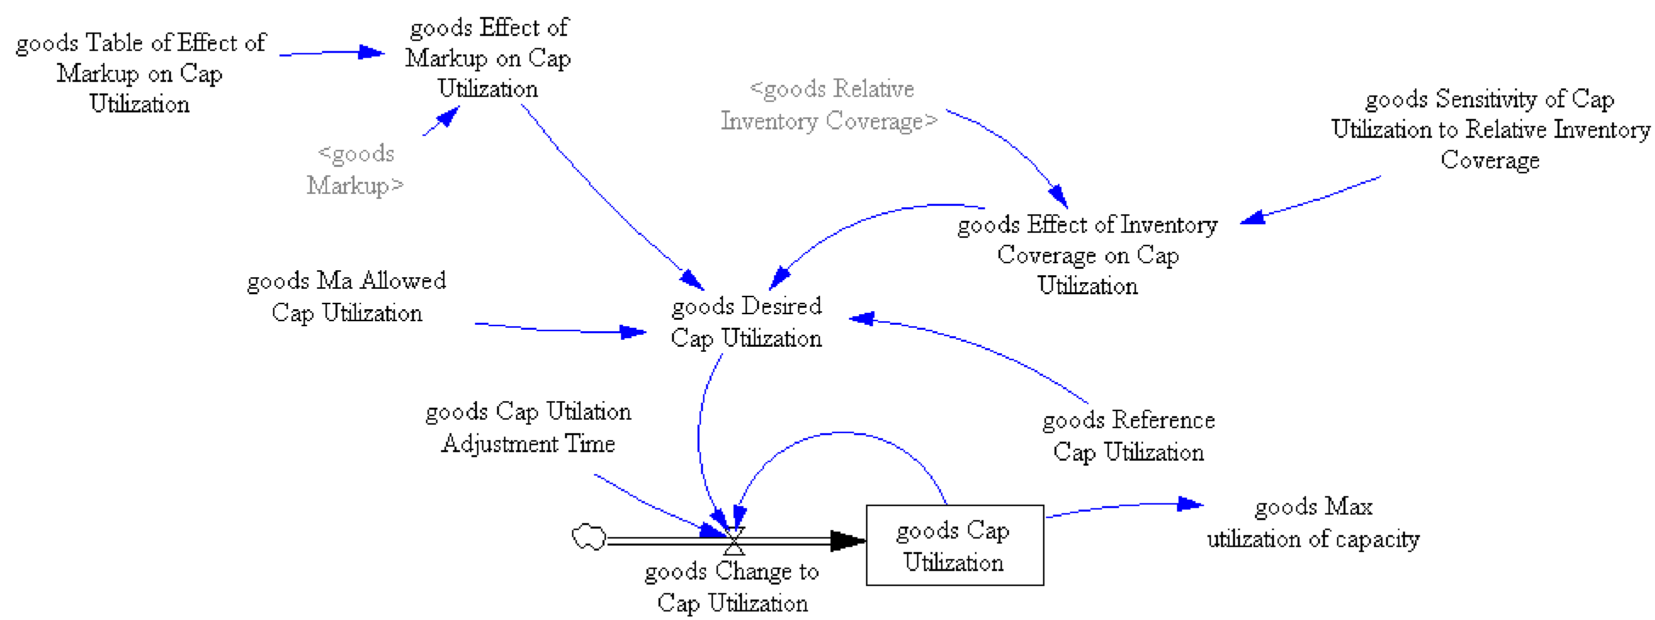

Figure A-22: Goods Capacity Utilization

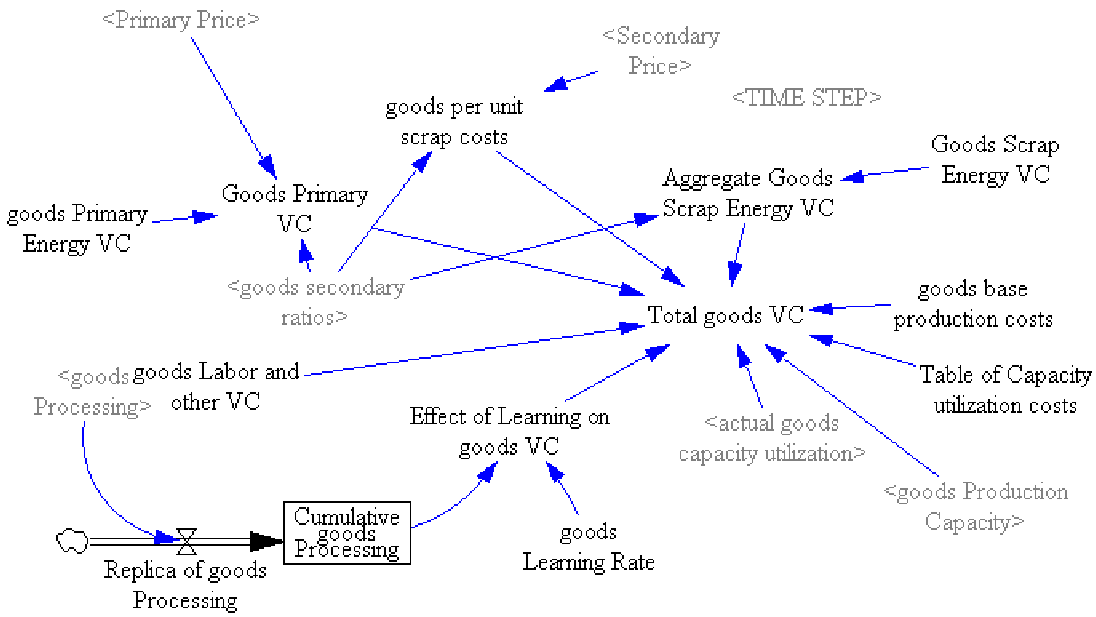

Figure A-23: Goods Variable Costs

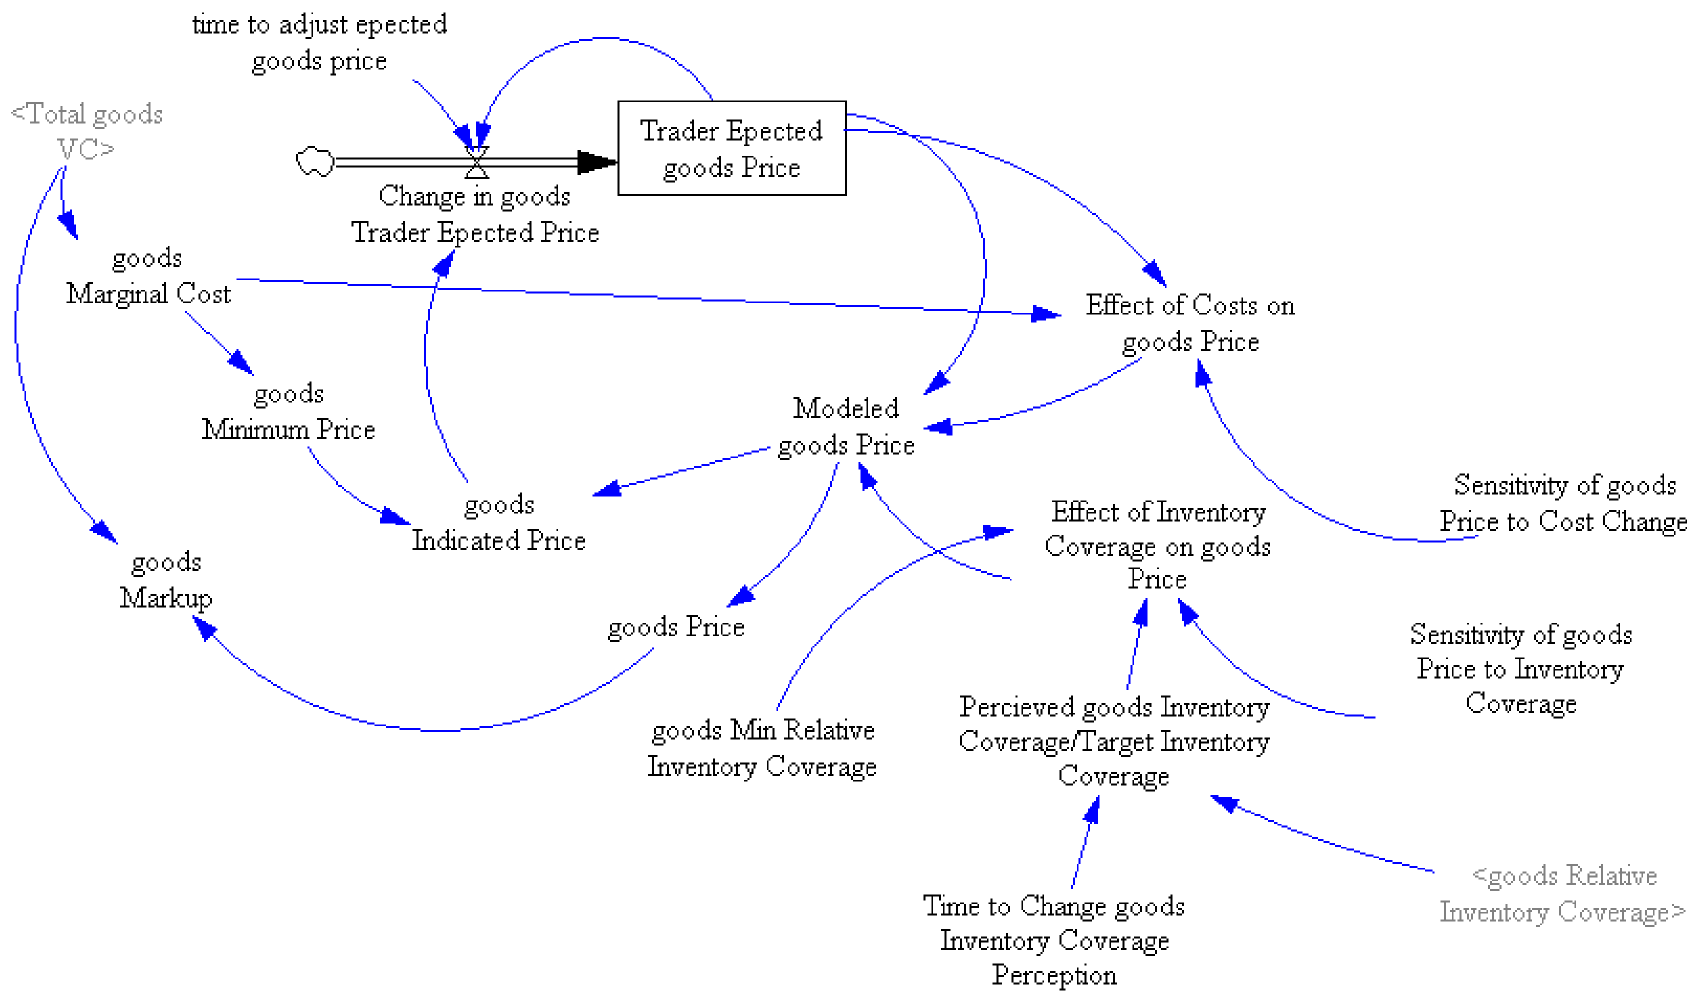

Figure A-24: Goods Price

Appendix B

MATLAB® Code

The following MATLAB® Code was used to load output data from Vensim® and produce the correlograms in Chapter 4.

B.1 Load Data

```

tic
%clear variables and screen
clear;
clc;

%import data
[data,varnames,casenames] = tblread('stochastic5.tab','\t');

% set variables
mtime = 201;
stime = 1985;
endtime = 2035;
dt=(endtime-starttime)/(mtime-1);
t=starttime:dt:endtime;
charrows=5;
replicates=25;
```

```
varnames=cellstr(varnames);

for i=1:charrows
    charnames(i)=cell(varnames(1));
    characteristic(:,i)=data(:,1);
    varnames(1)=[];
    data(:,1)=[];
end

charnames(3)=[];
characteristic(:,3)=[];
factlength=length(casenames)/replicates;
characteristic=characteristic(1:factlength,:);

[dummy,n]=size(data);
var=cell(n/mtime,1);
for j=1:length(var)
  var(j)=varnames(j*mtime);
end
[token,var] = strtok(var);
var=strtrim(var);

%move data into 3D matrix
mdata=zeros(factlength,mtime,length(var),replicates);
for n = 1:replicates
    for m = 1:length(var)
        for j = 1:factlength
            for k = 1:mtime
                mdata(j,k,m,n)=data(j+(n-1)...
                    *factlength,(m-1)*mtime+k);
            end
        end
```

```
    end
end

%truncate
mdata(:,1:41,:,:)=[];
t(1:41)=[];

toc
```

B.2 Produce Correlograms

```
tic
clc
close all
max_lag=80;

[m,n] = size(characteristic);
desmat = zeros(m,n);
for k=1:n
    for j=1:m
        if (characteristic(j,k) == max(characteristic(:,k)))
            desmat(j,k) = 1;
        elseif (characteristic(j,k) == min(characteristic(:,k)))
            desmat(j,k) = -1;
        else desmat(j,k) = 0;
        end
    end
end

C=zeros(factlength,2*max_lag+1,length(var),replicates);
L=zeros(factlength,2*max_lag+1,length(var),replicates);
```

```
difflev=zeros(factlength,length(var),replicates);

maxdiff=zeros(1,length(var));

%Difference data, calculate cross-covariance
maxdiff(:)=2;
for n=1:replicates
    for j=1:factlength
        for m=1:length(var)
            if maxdiff(m)==0;
                x=mdata(j,:,m,n)';
                y=mdata(j,:,6,n)';
            else
                x=diff(mdata(j,:,m,n),maxdiff(m))';
                y=diff(mdata(j,:,6,n),maxdiff(m))';
            end

            [C(j,:,m,n),L(j,:,m,n)]=xcov(x,y,max_lag,'coef');

        end
    end
end

AC=zeros(factlength,2*max_lag+1,length(var),replicates);
AL=zeros(factlength,2*max_lag+1,length(var),replicates);

%Difference data, calculate cross-covariance
for n=1:replicates
    for j=1:factlength
        for m=1:length(var)
            if maxdiff(m)==0;
                x=mdata(j,:,m,n)';
```

```
            else
                x=diff(mdata(j,:,m,n),maxdiff(m))';
            end

            [AC(j,:,m,n),AL(j,:,m,n)]=xcov(x,max_lag,'coef');

        end
    end
end

Cmean=squeeze(mean(mean(C,1),4));
Cmax=squeeze(max(max(C,[],1),[],4));
Cmin=squeeze(min(min(C,[],1),[],4));
Lmean=squeeze(mean(mean(L,1),4));
T=-Lmean*dt;

figure
set(gcf,'Position',[100,100,700,800])

%plot correlograms
for m=1:length(var)

    subplot(3,3,m)
    hold on
    plot(T(1:max_lag+1,m),Cmean(1:max_lag+1,m),...
        'LineWidth',1,'color','b')
    plot(T(1:max_lag+1,m),Cmax(1:max_lag+1,m),...
        ':','LineWidth',1,'color','b')
    plot(T(1:max_lag+1,m),Cmin(1:max_lag+1,m),...
        ':','LineWidth',1,'color','b')
    plot([T(1,m) T(max_lag+1,m)],[0 0],'k')
    axis([ T(max_lag+1,m) T(1,m) -1 1])
```

```
        title(var(m));
        xlabel('lag (years)');
        ylabel('\rho_{xy}');

end
saveas(gcf,'Corall', 'png')

%create factorial matrix for plotting grouped correlograms
[m,n] = size(characteristic);
desmat = zeros(m,n);
for k=1:n
        for j=1:m
                if (characteristic(j,k) == max(characteristic(:,k)))
                        desmat(j,k) = 1;
                elseif (characteristic(j,k) == min(characteristic(:,k)))
                        desmat(j,k) = -1;
                else desmat(j,k) = 0;
                end
        end
end

levels=[-1 0 1];
[dummy, cols]=size(characteristic);
Ccmean=zeros(2*max_lag+1,length(var),cols,3);
s=zeros(2*max_lag+1,length(var),cols,3);

for q=1:cols
        for r=1:3
                ind=(desmat(:,q)==levels(r));
                s(:,:,q,r)=squeeze(sum(sum(C(ind,:,:,:),1),4));
                N=sum(ind)*replicates;
                Ccmean(:,:,q,r)=s(:,:,q,r)./N;
```

```
        end
    end

    %plot grouped correlograms
    for m=1:length(var)%1:length(var)
        figure
        set(gcf,'Position',[100,100,800,600])
        w=1;
        for q=1:cols;%1:cols
            subplot(2,2,w)%subplot(2,2,q)
            w=w+1;
            hold on
            color=[1 0 0; 0 .5 0; 0 0 1];
            for r=1:3
                plot(T(1:max_lag+1,m),Ccmean(1:max_lag+1,m,q,r),...
                    'LineWidth', 2,'color',color(r,:))
            end
            %plot([T(1,m) T(max_lag+1,m)],[Rcrit Rcrit],'k:')
            plot([T(1,m) T(max_lag+1,m)],[0 0],'k')
            %plot([T(1,m) T(max_lag+1,m)],[-Rcrit -Rcrit],'k:')
            xlabel('lag (years)');
            ylabel('Correlation Coefficient (\rho)');
            title(charnames{q})
            legend(num2str(min(characteristic(:,q))),...
                num2str(mean(characteristic(:,q))),...
                num2str(max(characteristic(:,q))),'Location','Best')
        end
        mtit(var{m},'fontsize',14,'xoff',-.4,'yoff',.035);
        saveas(gcf,strcat('Cor4by', num2str(m)), 'png')
    end

    toc
```