Stochastic methods for improving secondary production decisions under compositional uncertainty

by

Gabrielle G. Gaustad

B.S. Ceramic Engineering
New York State College of Ceramics at Alfred University, 2004

Submitted to the School of Engineering
in partial fulfillment of the requirements for the degree of

Master of Science in Computation for Design and Optimization

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2009

Author...
School of Engineering
January 16, 2009

Certified by...
Olivier L. de Weck
Associate Professor of Aeronautics and Astronautics & Engineering Systems
Thesis Supervisor

Accepted by..
Jaime Peraire
Professor of Aeronautics and Astronautics
Director, Computation for Design and Optimization Program

Stochastic methods for improving secondary production decisions under compositional uncertainty

by

Gabrielle G. Gaustad

Submitted to the School of Engineering
on January 16, 2009, in partial fulfillment of the requirements for the degree of
Master of Science in Computation for Design and Optimization

Abstract
A key element for realizing long term sustainable use of any metal will be a robust secondary recovery industry. Secondary recovery forestalls depletion of non-renewable resources and avoids the deleterious effects of extraction and winning (albeit by substituting some effects of its own). For most metals, the latter provides strong motivation for recycling; for light metals, like aluminum, the motivation is compelling. Along aluminum's life-cycle there are a variety of leverage points for increasing the usage of secondary or recycled materials. This thesis aims to improve materials decision making in two of these key areas: 1) blending decisions in manufacturing, and 2) alloy design decisions in product development.

The usage of recycled aluminum in alloy blends is greatly hindered by variation in the raw material composition. Currently, to accommodate compositional variation, firms commonly set production targets well inside the window of compositional specification required for performance reasons. Window narrowing, while effective, does not make use of statistical sampling data, leading to sub-optimal usage of recycled materials. This work explores the use of stochastic programming techniques which allow explicit consideration of statistical information on composition. The computational complexity of several methods is quantified in order to select a single method for comparison to deterministic models, in this case, a chance-constrained model was optimal. The framework and a case study of cast and wrought production with available scrap materials are presented. Results show that it is possible to increase the use of recycled material without compromising the likelihood of batch errors, when using this method compared to conventional window narrowing.

The chance-constrained framework was then extended to improving the alloy design process. Currently, few systematic methods exist to measure and direct the metallurgical alloy design process to create alloys that are most able to be produced from scrap. This is due, in part, to the difficulty in evaluating such a context-dependent property as recyclability of an alloy, which will depend on the types of scraps available to producers, the compositional characteristics of those scraps, their yield, and the alloy itself. Results show that this method is effective in, a) characterizing the challenge of developing recycling-friendly alloys due to the contextual sensitivity of that property; b) demonstrating how such models can be used to evaluate the potential scrap usage of alloys; and (c) exploring the value of sensitivity analysis information to proactively identify effective alloy modifications that can drive increased potential scrap use.

Thesis Supervisor: Olivier L. de Weck
Title: Associate Professor of Aeronautics and Astronautics & Engineering Systems

Acknowledgements

I find that acknowledgements are often incredibly clichéd as well as too lengthy and containing far more personal information than any reader interested in (insert your thesis key words here) cares to know about. As such, I plan to keep these short with no syrupy platitudes to advisors, friends, and families. If you are included, you know what's up.

Prof. Randolph Kirchain and Prof. Olivier de Weck for taking the time out of their busy schedules to a) read this thing, and b) put up with me. Members of the Materials Systems Lab for creating a unique, professionally helpful, and hilarious working environment, most notably: Joel Clark, Frank Field, Rich Roth, Terra Cholfin, Tommy Rand-Nash, Sup, Anna Nicholson, Elsa Olivetti, Ashish Kar, Erica Fuchs, Preston Li and all other past and present MSL students who contributed to a great experience. Others toiling in the MIT trenches: Corinne Packard, Scott Bradley, and John Maloney. The ladies who helped me forget the MIT trenches: Becca Sorkin (bff), Carrie Sheffield, Livia Kelley, and Katie Casey. My husband, who is awesome: Jeff Povelaites. My fantastic family: Richard, Wendy, and Alexandra Gaustad.

Thanks.

Table of Contents

Table of Figures

Table of Tables

Chapter 1. Introduction

One of the key engineering challenges of the 21st century will be reducing the harmful effects associated with a growing population and the attendant flows of materials[1, 2]. The materials community is uniquely positioned to play a central role in addressing these problems by fundamentally changing the materials and processes used by society. For this to happen, materials experts must begin to consider the environmental impacts of their design choices and will require additional analytical tools to quantify those broader implications. This thesis begins to address the need for these analytical tools for at least one element of a material's environmental performance – the ability to be produced from secondary resources. Materials that perform well in this regard will be referred to herein as recycling-friendly. Within a materials' life-cycle or production chain there are a variety of leverage points for increasing environmental performance (Figure 1). This work aims to improve materials decision making in two of these key areas: 1) blending decisions in manufacturing, and 2) alloy design decisions in product development.

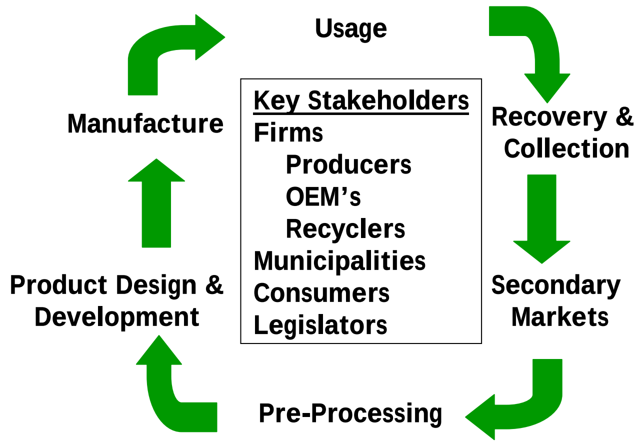

Figure 1. Product life-cycle showing key leverage points and major stakeholders

1.1 Production uncertainties effecting blending decisions

Uncertainty is a reality that confronts all businesses; materials producers are no exception. When business plans do not accommodate actual operating conditions, businesses are left with the negative economic impact of inefficient use of capital, materials, or potential market consumption. A significant set of economic disincentives emerge due to the various types of operational uncertainty that confront secondary metal processors [3-5]. In particular, depending on where one is in the production chain, business-critical sources of uncertainty include capricious demand, unstable availability of raw materials (particularly scrap materials), the precise composition of those raw materials, and the cost of factor inputs. An appreciation of these uncertainties can be gained by examining Figure 2 through Figure 5. Figure 2 shows the year to year change in aluminum apparent consumption in the United States over several decades, an illustration of the volatility of alloy demand. Scrap availability shows similar volatility (Figure 3), especially over the past few decades where much scrap has begun to be exported to rapidly industrializing countries such as China and Brazil.

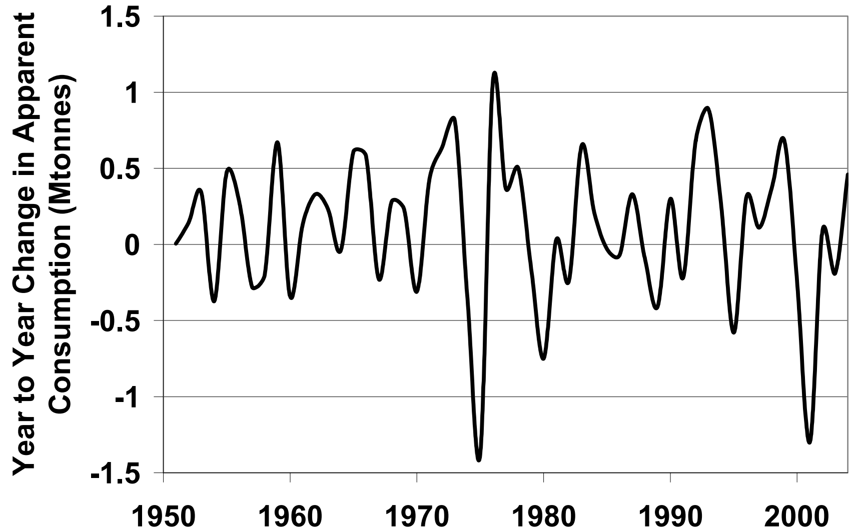

Figure 2. Year to year change in apparent aluminum consumption in the US[6]

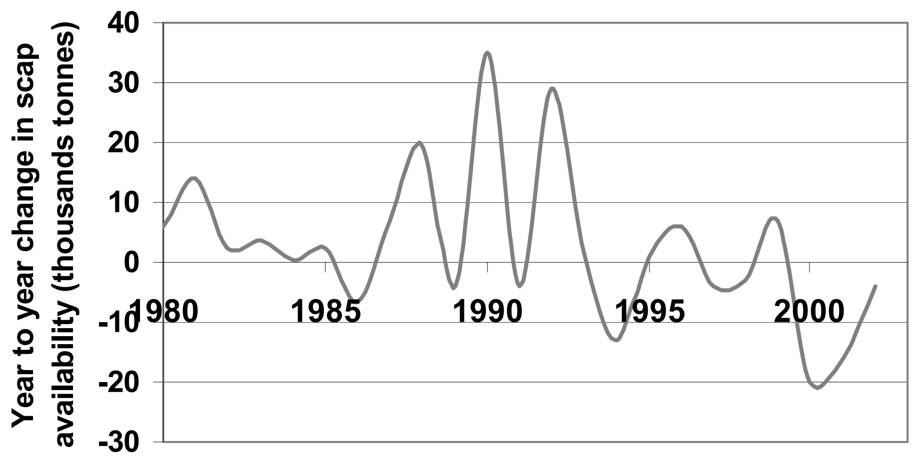

Figure 3. Year to year change in scrap generated in the US over the last two decades in thousands of metric tonnes[7]

Figure 4 shows the normalized London Metals Exchange price for primary aluminum over several months. Although the overall price trend of primary over a longer period (ie. the last four decades) is one which is clearly favorable to all aluminum producers, the significant variance of price represents not only a direct form of operating uncertainty, but also belies the underlying swings in demand which confront operational decision-makers. Scrap shows even larger volatility in price, due in part to geographic and regional price differences for different types of scrap materials. For example, scrap dealers near cities have larger supply of UBC's (used beverage cans) and therefore can offer lower prices; large scrap dealers in the Midwest have access to large amounts of automotive-heavy mixed scraps and therefore lower prices on those types. Such large price differences (e.g. approximately 47% difference between the maximum and minimum prices for auto wheels) over such a short period of time (less than two years) can lead to significantly different blending decisions. All of these uncertainties have the largest adverse effect on those furthest from the customer, e.g. materials producers, due to the feedback mechanisms inherent to typical market-based supply-chains [8]. Nevertheless, despite real uncertainties, definite business-critical decisions must be made on a daily basis.

Figure 4. Recent normalized London Metals Exchange daily cash settlement prices, Jan – Sept 2005 (Jan. 4, 2005 = 1)

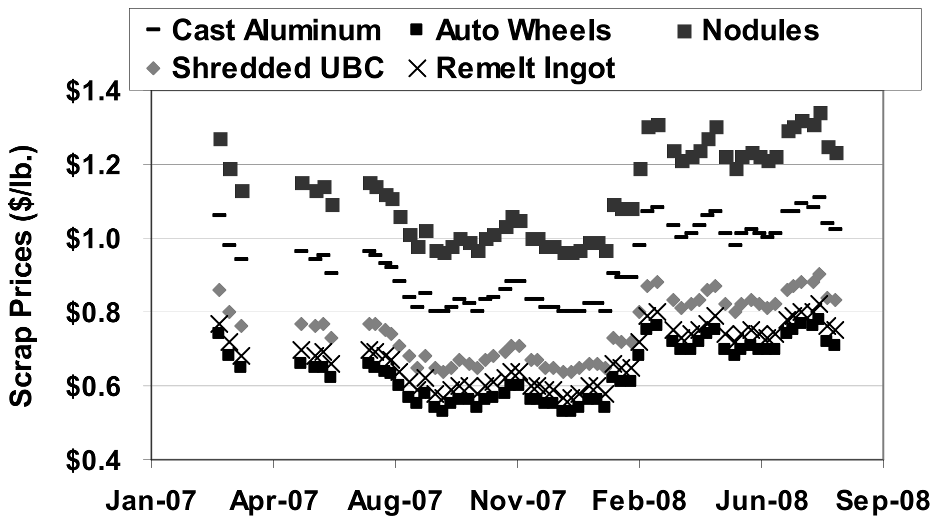

Figure 5. Scrap prices over two year period from multiple dealers[9]

Previous work has shown that explicit consideration of operational uncertainties (i.e., through the use of stochastic programming), in production planning can improve batch operator decisions both in terms of reduced operating costs and increased scrap consumption [10-15]. Such results are consistent with improvements observed in other contexts [16-21]. This thesis introduces the use of a related framework – referred to as stochastic programming – to explicitly address a previously unexplored source of uncertainty that confounds secondary batch planning – compositional variation of secondary raw materials. Elemental considerations for scrap have been identified as the most significant source of uncertainty in the production process [5, 22]. To provide an indication of the scope of this form of uncertainty, Figure 6 shows composition and standard deviation of recycled aluminum siding sampled over a period of a year; one can see the wide range in both mean and variance. Considering the many types of recycled materials secondary producers utilize multiplied by the dozens of relevant compositional elements, it is clear that compositional uncertainty makes it difficult to meet quality specifications and, thereby, creates a strong disincentive to use.

The stochastic optimization methods that will be presented attempt to address this problem by modifying conventional methods of dealing with compositional uncertainty. These conventional methods are deterministic or static in nature. A comparison of stochastic and conventional methods will be made in regards to both their scrap use and operational economics; these are examined through case studies and targeted simulations.

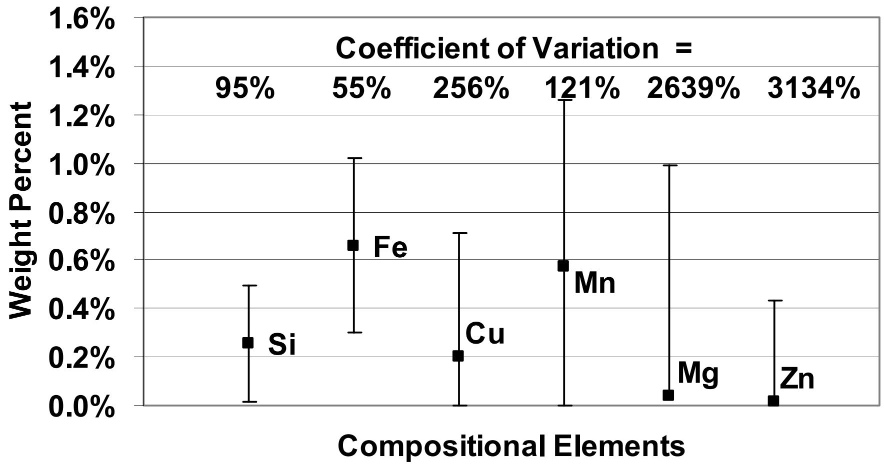

Figure 6. Compositional uncertainty (mean and standard deviation of various elements) in scrap aluminum siding sampled over the course of a year [4]

1.2 Recycling friendly alloy design

How to design alloys that are more recycling friendly or, in other words, more able to accommodate scrap materials in their production portfolios, is a challenging question. Industry experts and literature have provided a variety of suggestions. A much discussed strategy involves the development of single alloys that could meet the performance requirements currently filled by multiple alloys. For example, an alloy that could replace both 5XXX and 6XXX materials in transportation applications[23, 24]. Some even propose legislation or regulations to limit the number of alloys that can be used in certain products such as cars or aircraft[25]. Other suggestions involve modifying the forming and joining of aluminum, for example, replacing conventional welding with mechanical joining, laser welding, or friction stir welding[26]. Specific suggestions concerning the modifications of alloys include higher maximum compositional specifications for certain elements that will not adversely affect alloy properties, wider specification targets (i.e. higher maximums and lower minimums), or translating compositional constraints to specifications based on performance[27, 28]. However, no quantitative assessments of the efficacy of these suggestions on the ability of a recycler or recycling system to use more secondary raw materials have been reported in the literature. Furthermore, no methodology has been discussed that would quantitatively assess in what context these strategies should be applied.

Despite these specific issues, a range of literature has examined the use of decision-analysis models to improve the economic and resource use performance of recycling operations. The most pertinent include those that apply a range of mathematical programming techniques to improve decisions about raw materials purchasing strategy, technology selection[19, 29, 30], and the application of upgrading and sorting for secondary raw materials[31, 32]. Although notionally these models can be used iteratively to evaluate how some change would affect the ability to use secondary raw materials, none are applicable to evaluating the design of multi-specification alloys.

The primary challenge of evaluating the recycling-friendliness of multi-specification alloys is that it is a context dependent property; how much scrap an alloy can accommodate will be based on not only the compositional characteristics of the alloy itself, but also the types of scraps available to producers, the compositional characteristics of those scraps, and their metallic yield. As a result, a method to evaluate recycling-friendliness must be able to account for the confluence of these detailed effects.

Two sets of previous work on decision-analysis models have been specifically applied to recycling performance of secondary aluminum production and form the basis for addressing this need. The first set of studies by Reuter, van Schaik, and others[19, 29-31, 33-35] utilized optimization methods and dynamic modeling to optimize the recycling system for end-of-life vehicles, including the light-metals within them. This work and the models it presents can be used to guide operational and technological decisions by recyclers and to provide reasonable recovery expectations for recyclers, and more broadly, policy-makers. The second set is previous work by the author [36] and Rong[5] which describes schematic, mathematical programming models that identify the optimal raw materials mix to blend to produce a given multi-specification alloy or alloy production portfolio. These models are extensions of batch planning models that have been developed for decades and are available and used within the secondary metals industry today. To date, such models have been used to evaluate the effectiveness of improved operational practices and alloy substitution to increase potential scrap use. In previous work, it was hypothesized that model-derived sensitivity analysis information could be used to direct alloy design and demonstrated that, for a stylized case, such sensitivity analysis information varied significantly across specification and alloy.

This thesis extends this previous work by a) characterizing the challenge of developing recycling-friendly alloys due to the contextual sensitivity of that property; b) demonstrating how such decision-support models can be used to evaluate post-facto the potential scrap usage of alloys across a range of raw materials contexts; and (c) exploring the value of sensitivity analysis information to proactively identify the most effective alloy modification strategies that can drive increased potential scrap use. With regard to the latter, this thesis extends previous discussions by exploring in detail how sensitivity analysis information actually correlates with potential scrap use performance and how both the sensitivity analysis information and the associated potential scrap use effect changes with individual and coordinated specification modifications. In exploring the latter for two distinct cases, this thesis suggests that real potential exists for increasing potential scrap use through alloy redesign while remaining within established compositional specifications. Finally, this thesis extends previous work in this space by presenting a schematic algorithm for explicitly incorporating uncertain metal yield into the analyses of alloy design specifically, and recycler operational decisions more broadly.

Both the model and cases discussed herein are intended to be schematic in nature. Much work still remains to capture the metallurgical complexity of the recycling process; nevertheless, the results presented show that this framework holds promise to be a valuable tool in the metallurgists toolkit. Furthermore, a recycling evaluation tool, irrespective of scope or fidelity, would always be but a single element in the overall alloy design process. Traditional and emerging metallurgical methods will be required to identify alloys capable of meeting demanding physical performance requirements. Nevertheless, efficient design of resource-conscious materials depends upon analytical tools capable of projecting the impact of design choices on recycling performance.

Chapter 2. Batch planning models

2.1 Linear optimization

A large variety of modeling tools are available to help support the decisions of batch planners; many producers make use of linear optimization techniques [29]. Blending problems have been addressed with linear programming models for decades[37]. These models can improve decisions about raw materials purchasing and mixing as well as the upgrading and sorting of secondary materials [11, 19, 30]. Additionally, linear optimization techniques have been implemented to address larger scale aluminum recycling questions as well. Studies by [33-35, 39] have utilized dynamic modeling and large datasets to calculate optimized recovery rates for end-of-life vehicles in order to guide operational and technological decisions by recyclers and to provide reasonable recovery expectations for recyclers, and more broadly, legislators.

Analytical approaches may be used within such optimization tools to embed consideration of uncertainty in the decision-making, but generally this occurs through the use of statistical analyses that are used to forecast expected outcomes. Although this combination of statistical analysis and modeling can be powerful, it suffers from two fundamental limitations. First of all, implicit assessments based on mean expected conditions assume that deviation from that value has symmetric consequences. For many production related decisions within the cast-house, the repercussion of missing a forecast are inherently non-symmetrical. Additionally, deterministic approaches generally do not provide proactive mechanisms to modify production strategies as prevailing conditions evolve.

2.2 Modifying deterministic batch planning models

2.2.1 Mean based conventional method

Before exploring more advanced blending models, it is first helpful to understand how raw material compositional variation is handled presently within many batch planning tools. Although it is well known that most scraps will have some sort of compositional distribution, producers tend to model raw materials using their mean or expected value (A_i) or some other measure derived from these (e.g., mean + 10%). The goal of batch operators is then to combine these raw materials in such a way that the composition of finished goods falls within desired maximum and minimum targets. This can be stated mathematically by the following:

$$A_f^{\min} B_f \leq \sum_i A_i x_i \leq A_f^{\max} B_f \qquad (1)$$

where the symbol A represents the amount, expressed as a mass fraction, of a certain element such as silicon, magnesium, zinc, etc. that is either contained within a raw material, i, or cannot be exceeded / fallen below within a finished alloy, f. It is important to note that the minimum and maximum A_f are specifications while A_i is the actual composition. B is the amount of finished product f produced, while x is the amount of raw material used in the production of f. For clarity, the yield of A from raw material into finished alloy has been assumed to be one but could be changed to represent actual yields. Producing within specification based on the initial furnace charge is a key business objective of cast-shop operators [5]. Missed specifications require rework in the form of compositional additions or, even worse, dilution. Such rework is costly because it increases consumption of primary raw materials, energy, and time.

This method of representing recycled materials solely by their mean has one major shortcoming: it cannot differentiate between two scraps with the same mean but different variances. To illustrate this shortcoming, imagine a finished alloy is going to be made of only one raw material and can select between recycled materials: Scrap A or Scrap B. Considering only one compositional alloying element, (silicon, for example), assume both scraps have the same mean but scrap B has a larger standard deviation (variability). If these scraps have high silicon levels, they will be combined with primary raw materials until their mean composition is diluted inside the finished good's target minimum (Figure 7). Three pricing scenarios can be imagined. If scrap A were less expensive than scrap B, the model would choose all scrap A; this is highly unlikely. If scrap B is less expensive than scrap A as would be expected due to its higher level of compositional uncertainty, the mean based model would choose all scrap B. This would result in a finished good with a high likelihood of being out of specification. If the two scraps were equal in price, the mean based model would choose between the two indifferently which is also not optimal. Additionally, as will be shown subsequently, it is difficult to guarantee reasonable error rates at mid to high levels of compositional variation for mean-based methods.

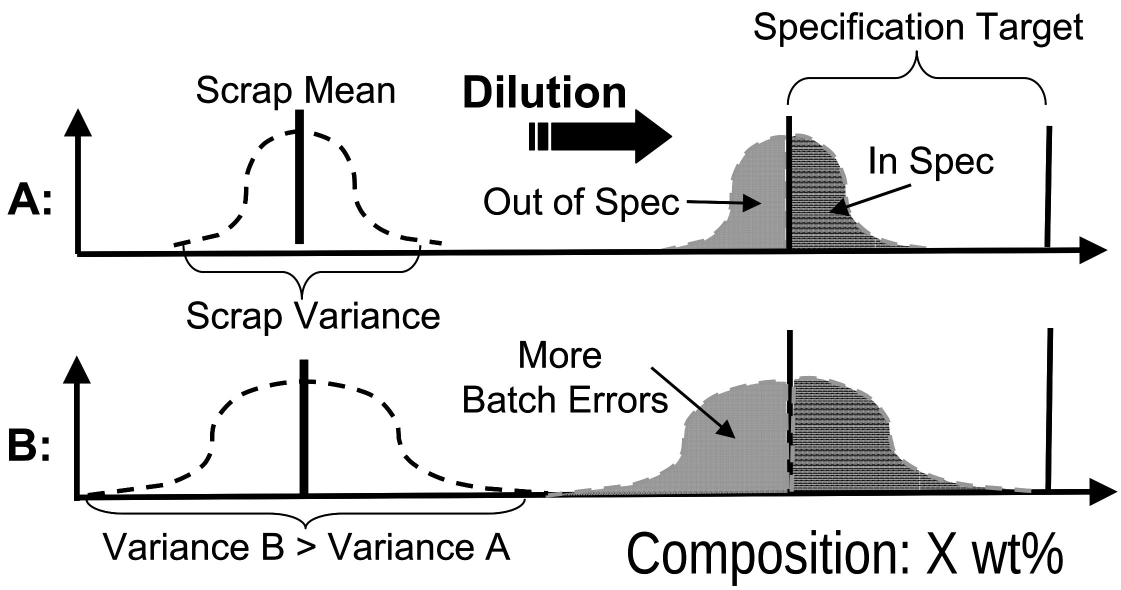

Figure 7. Schematic of scrap mixing illustrating shortcomings of mean based conventional method. Typical mean-based optimizations cannot differentiate between scrap A (small variance) and scrap B (large variance).

2.2.2 Compositional window conventional method

To address this shortcoming, some producers model scrap composition as a range or window. In this method, the width of the window (i.e., the modeled maximum and minimum compositional limits for the scrap) is set based on the scrap's mean and variance. These maximum and minimum limits are then compared to the maximum and minimum specifications of the finished goods to establish batch plans. This method for representing scrap composition is compared schematically to the mean based method in Figure 8**.** Using this scrap compositional window method, batch planning models seek out combinations of raw materials such that recycled materials are combined with primary metal until the minimum composition limit of the mixture is inside the finished good's target minimum and / or the maximum limit is inside of the maximum specification. The scrap compositional window method has advantages over the mean based method in that it can differentiate between scraps with the same mean but different variances and will provide a much lower error rate for being within specification, especially as scrap becomes more variable. However, shortcomings still exist; Figure 9 shows that when the compositional window is set at one standard deviation (σ) from the mean, a significant number of batch errors can still occur.

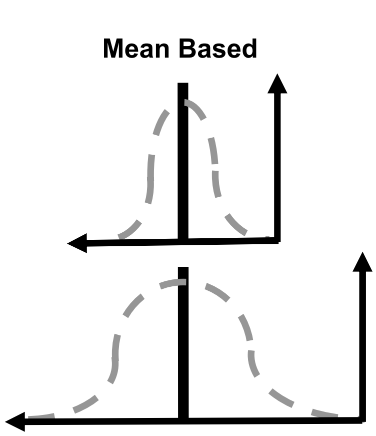

Figure 8. Schematic comparing the conventional methods for representing scrap composition as a mean or as a compositional window.

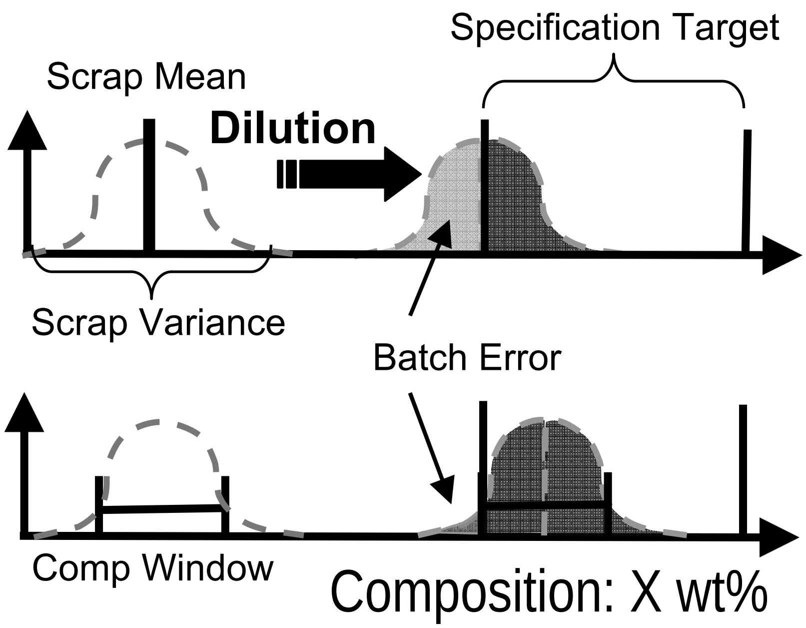

Figure 9. Schematic of scrap mixing illustrating the conventional method of representing scrap as a compositional window; when the window is set to be one standard deviation from the mean, a significant amount of batch errors will still occur

2.2.3 Window narrowing of finished good specifications

To further limit the incidence of rework, operators often modify the mean based or compositional window methods by generating batch plans based on more narrow finished goods specification targets[40]. This narrowing of the target window creates a margin of safety around compositional specifications such that a high level of likelihood is maintained for the finished goods compositions to fall within their actual specifications, shown schematically in Figure 10. This percentage of the original finished good specification is then used as the sole adjustment for dealing with uncertainty.

This practice modifies the constraint expressed in Equation (1) as shown in Equations (2) and (3) in which F is defined to be an adjustment factor of the specification range. The combination of representing scrap as a composition window and using a narrowed finished good specification window is the most robust conventional method for modeling scrap mixing decision and will be used as a conventional base case and referred to as the window narrowing (WN) method throughout the balance of this document.

$$\left.\begin{aligned}\sum_i A_i^{\max} x_i &\leq \frac{(1-F)}{2} A_f^{\max} B_f \\ \sum_i A_i^{\min} x_i &\geq \frac{(1+F)}{2} A_f^{\min} B_f\end{aligned}\right\} \quad , A_f^{\min} > 0 \qquad (2)$$

$$\left.\begin{aligned}\sum_i A_i^{\max} x_i &\leq (1-F) A_f^{\max} B_f \\ \sum_i A_i^{\min} x_i &\geq A_f^{\min} B_f\end{aligned}\right\} \quad , A_f^{\min} = 0 \qquad (3)$$

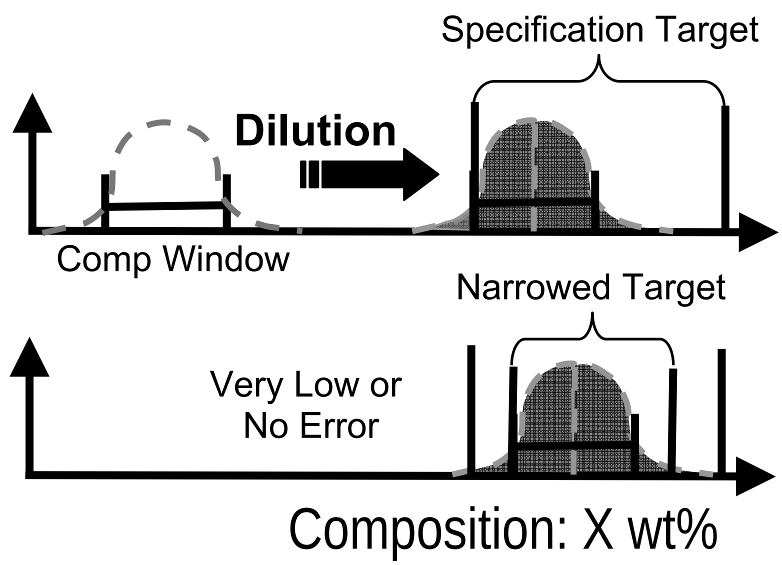

Figure 10. Scrap represented as compositional window combined with narrowed finished good specification targets: the window narrowing conventional method

2.3 Shortcomings of conventional methods

Although window narrowing does offer a mechanism to control the risk of errant production, various issues arise with this practice. Generally, the window narrowing method is the most conservative method for a given error rate, or number of batches that fall outside of the finished goods specifications. Other prevalent methods presented in the literature also aim to maintain a linear performance constraint while considering variability in scrap composition[41, 42]. Work done by Debeau[43] in the steel industry cautions that simply using linear programming for batch planning problems is not sufficient and suggests additional constraints to deal with variable raw materials. While in the same way as window narrowing these methods help to decrease the probability of batches being out of specification, they all share the trait of insufficient accounting for the impact on variance when combining scraps.

To illustrate this shortcoming, imagine a finished alloy is going to be made of only two raw materials: Scrap A and B. Considering only one compositional alloying element, (silicon, for example), Equation (4) shows the expected outcome of combining *a* kilograms of Scrap A and *b* kilograms of Scrap B from the perspective of the window in window method. The mean and variance of the combination is assumed to be simply the weighted sum of the individual pile's mean and variance. For example, combining 2 kilograms of Scrap A (μ=10% silicon σ=5%) and 2 kilograms of Scrap B (μ=10% silicon σ=5%) will result in a combination with μ= 2 kg*10% + 2 kg*10% = 40%/ 4 kg = 10% silicon and σ= 2kg*5% + 2kg*5% = 20%/4kg = 5%. In summary, the variance of the composition of a finished alloy will increase linearly as scrap piles are combined.

$$\begin{aligned} \mu_{A+B} &= a\mu_A + b\mu_B \\ \sigma_{A+B} &= a\sigma_A + b\sigma_B \end{aligned} \quad (4)$$

However, fundamental statistics shows that the combined scrap will have actual mean and variance according to Equation (5)[44]. So taking the example above (2 kilograms from each scrap pile) will result in the same mean as calculated above but the variance will be much different. In a case where the scrap pile compositions are completely uncorrelated, the variance will be σ= $((2\text{kg})^2 *(5\%)^2 + (2\text{kg})^2 *(5\%)^2)^{1/2}$= 14%/4kg = 3.5%. In reality, some correlations between scrap pile compositions would be expected and for these cases, the variance would be slightly higher than the example given, however, never as high as the weighted sum. So for all realistic cases, the variance of the composition of a finished alloy will increase sub-linearly as scrap piles are combined. The conservative nature of the window narrowing method is a direct consequence of its overestimation of compositional variance.

$$\begin{aligned} \mu_{A+B} &= a\mu_A + b\mu_B \\ \sigma_{A+B} &= \sqrt{a^2\sigma_A^2 + b^2\sigma_B^2 + 2ab\,\mathrm{cov}(A,B)} \end{aligned} \quad (5)$$

Another shortcoming of conventional window narrowing is that it may be difficult to determine the necessary amount of narrowing and therefore over-compensation can often occur. Tightening the window too much could cause the system to under utilize available secondary materials. Not meeting the finished good compositional specifications is an even larger problem that results in rework requiring time, energy, and money. An implied margin of error will exist depending on the size of the window; however, the static nature of the method does not provide an intrinsic method to tailor practices to these different error rates and therefore may penalize the system's ability to utilize scrap efficiently. For example, referring back to Figure 7, implementing the same window narrowing strategy for both scraps would result in under utilizing Scrap A and having high error rate when using Scrap B. For the operator tracking the compositional distribution of incoming raw materials, a mean based method would be more sensitive to these changes. Relating the margin of error to the underlying uncertainty in the chemical compositions of the raw materials may be a better practice and thereby provide both increased scrap consumption and associated economic benefits. The following section describes a method which provides just this capability.

Chapter 3. Methods

3.1 Stochastic programming

Stochastic programming techniques encompass a large set of problems that deal with uncertainty of one form or another in the formulation. Often, this is accomplished by considering some form of probability and/or statistics within the objective function or constraints. Pioneers in the field have done work in an extremely wide range of applications including: agricultural applications including farm management[45, 46] and irrigation, economics[47, 48], finance[49, 50], assignment [51], facility location[52], inventory, water storage and reservoir management, energy, and production. Stochastic programs have been in use for solving product mix and blending problems for decades as well. Particular applications include nutrient blending for crops[53], humans[54], and livestock[55, 56], product mix[57], coal blending[58], asphalt mixing[59], fertilizer mixing at a chemical plant[60], and most prevalent, petrochemicals including gasoline[61]. Methods in stochastic programming are generally divided into two broad categories: single stage and multi-stage stochastic programs[62]. This work will explore the use of a multi-stage recourse problem for dealing with compositional uncertainty as well as two single stage methods: the use of a penalty within the objective function and the use of chance-constraints within the compositional constraints. The computational complexity and scaling of these three stochastic methods compared to deterministic linear programming methods will be the metrics used to select an appropriate method for managing raw material uncertainty in metals blending problems (Section 4).

3.1.1 Multi-stage recourse stochastic programs

A recourse model consists of multi-stage decisions that must be made prior to an uncertain event. The associated consequence or recourse is a second stage reaction that depends on the decision made at the first stage and the outcome of the uncertain event, shown schematically in Figure 3. Recourse models have a wide range of applications including product planning[63], transportation networks[64], energy and electricity networks[42, 65], supply chain, inventory, and manufacturing networks[66-68], finance[69], and human resource planning[70].

The objective function for a recourse problem, which is to be either minimized or maximized, can be mathematically stated as follows:

$$f(C, X^1) + g(C, p, X^2) \quad (6)$$

In Eq (6), the contribution from stage one to the objective function is given by the function *f*(.). X^1 is the vector of decision variables from stage one. The contribution from stage two to the objective function is given by the function *g*(.). X^2 is the vector of stage-two recourse variables over all possible uncertain outcomes and p is the vector of the probabilities of those outcomes. The overall cost of the recourse decisions to the overall objective is weighted by those probabilities. Therefore, the objective is an expected average objective rather than a deterministic objective. C is the cost vector whose aggregate contribution to the objective function is being maximized or minimized in this optimization problem.

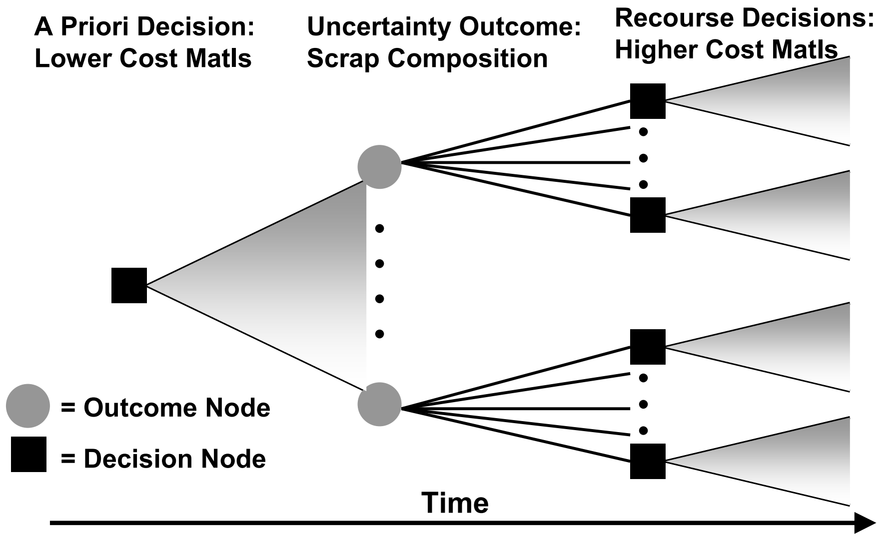

Figure 3. Schematic representation of recourse model

The variables to solve for are X^1_i, X^1_{ijz} and X^2_{ijz} which are translated into the recourse batch planning optimization model together with other notations used in the problem formulation (Eq (7) to (13)).

$$Min: \sum_i C_i X_i^1 + \sum_{r,j,z} C_{rework} P_z X_{ijz}^2 - \sum_{i,z} SC_i P_z R_{iz} \quad (7)$$

$$s.t.: \quad \forall_i \quad X_i^1 \le A_i \quad (8)$$

The amount of residual raw materials for each scenario is calculated as:

$$\forall_{i,z} \quad R_{iz} = X_i^1 - \sum_j X_{ijz}^1 \quad (9)$$

For each compositional scenario z there are supply constraints as determined by the amount of raw materials pre-purchased,

$$\forall_{i,z} \quad \sum_j X_{ijz}^1 \le X_i^1 \quad (10)$$

Eq (10) enforces the aforementioned condition that raw materials at the reduced cost (C_i) must be ordered before final production. As such, at production time, a cost of rework (C_{rework}, where $C_{rework} >> C_i$) must be assessed for all raw materials purchased in the second stage decision. Similarly, a production constraint exists for each scenario, quantifying how much of each alloy must be produced:

$$\forall_j \quad \sum_i X_{ijz}^1 + \sum_i X_{ijz}^2 = B_{jz} \ge M_{jz} \quad (11)$$

For each alloying element *k,* the composition of each alloy produced must meet production specifications [28]:

$$\forall_{j,k} \quad \sum_i x_{ij}^1 \overline{\varepsilon}_{ik} + \sum_i x_{ij}^2 \overline{\varepsilon}_{ik} \le B_j \varepsilon_{jk}^{\max} \qquad (12)$$

$$\forall_{j,k} \quad \sum_i x_{ij}^1 \overline{\varepsilon}_{ik} + \sum_i x_{ij}^2 \overline{\varepsilon}_{ik} \ge B_j \varepsilon_{jk}^{\min} \qquad (13)$$

All other variables are defined below:
C_i = unit cost (\$/T) of raw material i
A_i = mass of raw material i available for purchasing
$\overline{\varepsilon}_{ik}$ = average mass element k in raw material i
B_j = mass of finished good j produced M_j = mass of finished good j demanded
ε^{max}_{jk} = max. mass element k in finished good j
ε^{min}_{jk} = min. mass element k in finished good j
R_{sz} = Residual amount of raw material i unused in scenario z
D = 1 – discount on the value of unused scrap materials
P_z = probability of occurrence for compositional scenario z
x_{ijz} = amount of raw material i to be acquired on demand for the production of finished good j under demand scenario z
A_i = amount of raw material i available for pre-purchasing

In the notation of Eq (6), the objective function can be decomposed into two parts:

$$f(C, X^1) = \sum_i C_i \, X_i^1 \qquad (14)$$

$$g(C, p, X^2) = \sum_{i,j,z} C_{rework} P_z X_{ijz}^2 - \sum_{i,z} SC_i \, P_z R_{iz} \qquad (15)$$

The stage one component, Eq (14), consists of only cost contributions from the raw materials purchased initially at a lower cost. The stage two cost components, Eq (15), consist of the effects of additional higher priced raw material purchased, the additional cost of having to fix batches that are out of specification, and the salvage value of unused raw materials.

It is prohibitively computationally expensive to include a continuous distribution for the compositional scenarios. Therefore, each scenario is represented discretely with some associated probability (Figure 11). the problem will then scale with the number of discrete compositions used to represent some variance around the mean composition, this will be explored more fully in Chapter 4.

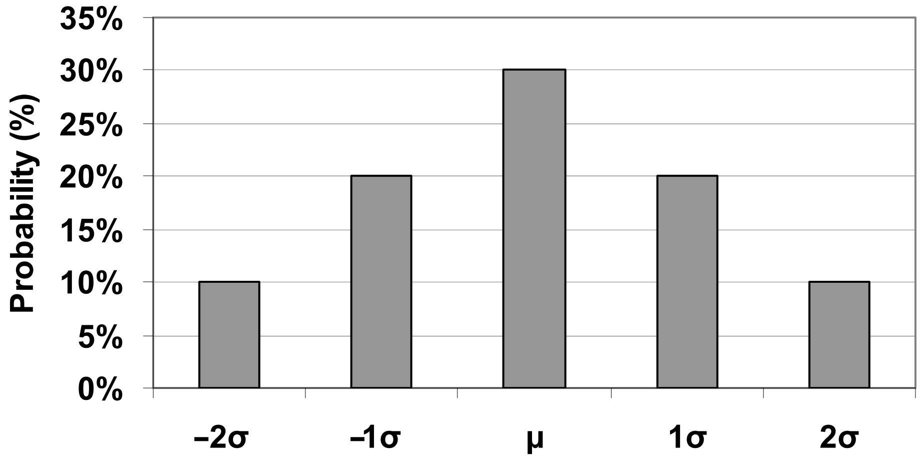

Figure 11. Discretely modeled probability distribution for compositional scenarios

3.1.2 Chance-constrained stochastic programming

Stochastic programming methods including chance-constrained variants were first formulated by Charnes and Cooper [71, 72] as mechanisms to embed a more rich set of statistical information into optimization based decision models. This technique is often used when the consequences for not meeting certain constraints are unknown. Typical applications include feed mixing[55, 56], reservoir management[73], nutritional planning, inventory control, metals blending[32, 74], scheduling problems[75], and water quality. Miller and Wagner [76] were the first to apply joint probabilistic constraints, however, could only consider independent random variables on one side of the equation, not simultaneously. These mathematical techniques became more useful for real world sized problems in the 1970's, when computational power was beginning to grow rapidly. Prekopa [77] introduced stochastic constraints as they are used currently: stochastically dependent joint probabilistic constraints. An excellent review of stochastic programming including the use of joint chance-constraints can be found in [78] while a variety of recent numerical applications are available [19, 73, 79].

In the context of cast-shop batch planning, the chance-constrained method allows the compositional constraints to be modified such that 1) the model embeds knowledge of both the mean and variance of available raw materials and 2) the user can query the model for optimal solutions which provide a specified level of confidence for meeting the compositional specifications. This method relates the desired level of confidence to the underlying standard deviations observed in the sampled raw materials. With the understanding that the compositional constraints will not always be satisfied due to inherent uncertainty, they can be rewritten as probabilistic expressions and transformed into their deterministic equivalents. For the problem at hand, the constraints expressed in Equations (2) and (3) are transformed into Equations (16) and (17):

$$\Pr\left\{\sum_i A_i x_i \leq F_f^{\max} A_f^{\max} B_f\right\} \geq \alpha \qquad (16)$$

$$\rightarrow \sum_i A_i x_i + X(\alpha)\Big(\sum_i \sum_j \rho_{ij}\sigma_i\sigma_j x_i x_j\Big)^{1/2} \leq F_f^{\max} A_f^{\max} B_f$$

$$\Pr\left\{\sum_i A_i x_i \geq F_f^{\min} A_f^{\min} B_f\right\} \geq \beta \qquad (17)$$

$$\rightarrow \sum_i A_i x_i + X(1-\beta)(\sum_i \sum_j \rho_{ij}\sigma_i\sigma_j x_i x_j)^{1/2} \geq F_f^{\min} A_f^{\min} B_f$$

The statements Pr(.) state that those constraints which were originally required to be strictly satisfied are now only satisfied α and β percent of the time. Thus α and β are desired levels of confidence factors which the operator can use to adjust his or her sense of importance for that particular elemental composition to be within specifications. The function *X(.)* is the inverse of a normalized cumulative Gaussian distribution function[1] which relates the underlying raw material composition standard deviations to the desired level of confidence. Work by Peterson[4] on scrap variability has shown this estimation to be applicable in many cases. The symbol ρ_{ij} represents the correlation between the fluctuations in composition of raw material *i* and *j*. By definition $\rho_{ij} = 1$ when $i = j$.

The chance-constrained problem can be formulated as follows in Equations (18) through (22) as a linear optimization model. The goal of this model is to identify the production plan, referred to subsequently as a batch plan, that will minimize the overall expected production costs (*C(x)*) of meeting finished good compositional specifications through optimal and efficient use of primary and secondary raw materials (Eq.(18)). To more accurately capture the behavior of a recycling operation, this simple objective is subject to a number of specific constraints. Firstly, raw materials cannot be prescribed in the batch plan in excess of the quantity available (Eq.(19)). Secondly, the batch plan must lead to production quantities that meet or exceed the established target for each alloy (Eq.(20)). Finally, the likelihood that a batch plan produces alloys within compositional specifications must exceed a specified probability (Eqs. (21) and (22)). This is accomplished by relating the likelihood of achieving a certain finished alloy composition to the underlying statistical characteristics of the raw material compositions (i.e., ε_{ik},$\sigma_{(\varepsilon)ik}$, and $\rho_{(\varepsilon)ilk}$) and the desired confidence limits, as established by the parameters α and β. With the understanding that the compositional constraints will not always be satisfied due to inherent uncertainty, they can be rewritten as probabilistic expressions and transformed into their deterministic equivalents.

Min: $$\sum_i C_i X_i \qquad (18)$$

Subject to: $$\forall_i \quad \sum_j x_{ij} = X_i \leq A_i \qquad (19)$$

$$\forall_j \quad \sum_i x_{ij} = B_j \geq M_j \qquad (20)$$

$$\forall_{j,k} \quad \sum_i x_{ij}\bar{\varepsilon}_{ik} + X(\alpha)(\sum_i \sum_l \rho_{(\varepsilon)ilk}\sigma_{(\varepsilon)ik}\sigma_{(\varepsilon)lk}x_{ijk}x_{ljk})^{1/2} \leq B_j \varepsilon_{jk}^{\max} \qquad (21)$$

$$\forall_{j,k} \quad \sum_i x_{ij}\bar{\varepsilon}_{ik} + X(1-\beta)(\sum_i \sum_l \rho_{(\varepsilon)ilk}\sigma_{(\varepsilon)ik}\sigma_{(\varepsilon)lk}x_{ijk}x_{ljk})^{1/2} \geq B_j \varepsilon_{jk}^{\min} \qquad (22)$$

All other variables are defined below:
C_i = unit cost ($/T) of raw material *i*
x_{ij} = mass of raw material *i* used in making finished good *j*
X_i = mass (kt) purchased raw material *i (both primary and scrap)*
A_i = mass of raw material *i* available for purchasing

[1] Other statistical distributions can also be assumed.

$\overline{\varepsilon}_{ik}$ = average mass element k in raw material i
$\sigma_{(\varepsilon)ik}$ = standard deviation of the composition (ε) of element k in raw material i
$\rho_{(\varepsilon)ilk}$ = correlation coefficient between composition (ε) of element k in raw materials i and l (ρ_{il} = 1 when $i=l$)
B_j = mass of finished good j produced
M_j = mass of finished good j demanded
ε^{max}_{jk} = max. mass element k in finished good j
ε^{min}_{jk} = min. mass element k in finished good j
$X(\_)$ = inverse of a normalized cumulative Gaussian distribution function
α = likelihood that the actual composition will fall below the upper limit of final alloy composition
β = likelihood that the actual composition will fall above the lower limit of final alloy composition
Of these, the ones that may require further clarification are α, β, and $X()$. Individually, α and β represent the likelihood that the batch plan identified by the model will result in a composition that is lower than the upper compositional limit and greater than the lower compositional limit, respectively. $X()$, the inverse normalized cumulative Gaussian distribution function, characterizes the relative distance from the mean that corresponds to the designated level of likelihood.

3.1.3 Penalty functions

Another stochastic programming method described in the literature [80-82] is the use of a penalty function added to the least cost objective function. While the majority of the applications are for blending within the chemicals industry, their extension to metals batch mixing is quite straightforward as formulated in Eq.(23) through (27). Typically these models multiply some cost of rework (C_{rework}) with the probability of being out of specification. The objective function then becomes:

$$\textit{Min.:}\quad \sum_i C_i X_i + C_{rework} * \left[\forall_{j,k} \Pr\left\{ \sum_i x_{ij}\overline{\varepsilon}_{ik} \le B_j \varepsilon^{\max}_{jk} \right\} \forall_{j,k} + \Pr\left\{ \sum_i x_{ij}\overline{\varepsilon}_{ik} \ge B_j \varepsilon^{\min}_{jk} \right\} \right] \tag{23}$$

This causes the problem to have nonlinearities in the objective function; each of the constraints will be unchanged from the deterministic formulations and therefore will remain linear as shown below.

Subject to:

$$\forall_i \quad \sum_j x_{ij} = X_i \le A_i \tag{24}$$

$$\forall_j \quad \sum_i x_{ij} = B_j \ge M_j \tag{25}$$

$$\forall_{j,k} \quad \sum_i x_{ij}\overline{\varepsilon}_{ik} \le B_j \varepsilon^{\max}_{jk} \tag{26}$$

$$\forall_{j,k} \quad \sum_i x_{ij}\overline{\varepsilon}_{ik} \ge B_j \varepsilon^{\min}_{jk} \tag{27}$$

3.2 Addition of stochastic yield to chance-constrained formulation

To make the chance-constrained formulation more flexible to capture the metallurgical complexity of modern recycling processes it can be amended to comprehend the effects of material yield. To increase the precision of the model and to accommodate current data collection practices within industry, yield is represented by two parameters: elemental yield and gross yield. The elemental yield, Y_{ik} , represents the fractional increase or decrease of the mass of alloying element *k* in the melt derived from incoming material inflow *i*. The mechanisms that drive such mass change vary widely by element. For example, iron is expected to increase due to pick-up from processing equipment and silicon will increase due to pick-up from the furnace refractories. Other elements may decrease due to oxide formation (zinc and magnesium may form spinel), volatization, or sinking/settling of the melt. The second parameter used to capture the effects of yield is the gross melt yield, G_i, which represents total metal loss (i.e. aluminum and other elements) as a fraction of incoming raw materials *i*. Such loss occurs due to dross formation, spills, etc. Notably, both forms of yield are expected to exhibit some variation from batch to batch. As such, they will be incorporated into the model as stochastic parameters.

It is important to point out that while this treatment of yield does significantly increase the flexibility and the potential fidelity of the model framework, it does not capture all possible effects. In contexts where elemental interaction is larger than the natural stochastic variation of the yield, additional second order terms would need to be incorporated. Similarly, works by van Schaik and Reuter[83] [84, 85] have shown that particle size distribution and scrap conformation are key factors in metal yield and, therefore, operational decisions. The formulation presented herein does not address such effects directly. For business critical decisions where the operational decisions will affect the yield variation associated with particle size, analysts should carefully consider the explicit treatment of these issues for their specific context. It is important to note, however, that irrespective of the form of the performance function (i.e. the left-hand sides of Equations (21) and (22) which relate the composition of the incoming raw materials to the composition of the outgoing finished alloy) most modern optimization packages can generate shadow price information on the constraint. As such, the types of work demonstrated herein should be extensible even to other blending models.

If the metallurgical yield is in fact stochastic in nature, its effect cannot be accounted for with a scalar transformation of the above formulation. Instead, the constraints on production quantity (Eq.(20)) and composition (Eqs.(21) and (22)) must be redeveloped. The probabilistic constraints incorporating yield can be expressed as:

$$\forall_j \quad \Pr\left\{ B_j^{'} = \sum_i x_{ij} G_i \le M_j \right\} \ge \gamma \tag{28}$$

$$\forall_{j,k} \quad \Pr\left\{ a_{jk} = \sum_i Y_{ik} \varepsilon_{ik} x_{ij} \le B_j^{'} \varepsilon_{jk}^{\max} \right\} \ge \alpha^{'} \tag{29}$$

$$\forall_{j,k} \quad \Pr\left\{ a_{jk} = \sum_i Y_{ik} \varepsilon_{ik} x_{ij} \ge B_j^{'} \varepsilon_{jk}^{\min} \right\} \ge \beta^{'} \tag{30}$$

where Y_{ik} and G_i are random variables as defined above, ε_{ik} is a random variable representing the actual amount of k in raw material *i*, $B_j^{'}$ is the actual quantity of batch *j* produced, a_{jk} is the actual quantity of element *k* in finished batch *j*, and γ, α' and β' are the minimum acceptable likelihood of the respective conditions being true.

The former (Eq.(28)) can be transformed directly into its deterministic equivalent as follows:

$$\forall_j \quad \sum_i x_{ij}\bar{G}_i + \mathrm{X}(\gamma)\sqrt{\sum_i\sum_l \rho_{(G)il}\sigma_{(G)i}\sigma_{(G)l}x_{ij}x_{lj}} \le M_j \tag{31}$$

Realizing the same for the new compositional constraint requires additional manipulation. First, it is helpful to replace the product of the two random variables $Y_{ik}\varepsilon_{ik}$ with Ψ_{ik} which is also a random variable. Formally, the distribution of Ψ_{ik} is a modified Bessel function of the second kind [27]. However, for typical characteristics of elemental yield and elemental content, Ψ_{ik} can be well approximated with a normal distribution such that:

$$\Psi_{ik} \cong Y_{ik}\varepsilon_{ik} \text{ and } \Psi_{ik} \sim N(\bar{\Psi}_{ik}, \sigma^2_{(\Psi)ik}) \tag{32}$$

$$\text{where } \bar{\Psi}_{ik} = \bar{Y}_{ik}\bar{\varepsilon}_{ik} \text{ and } \sigma^2_{(\Psi)ik} = \bar{\varepsilon}^2_{ik}\sigma^2_{(Y)i} + \bar{Y}^2_{ik}\sigma^2_{(\varepsilon)ik} + \sigma^2_{(Y)i}\sigma^2_{(\varepsilon)ik} \quad [84, 85] \tag{33}$$

Incorporating these relationships into and rearranging Eq (29) yields the following probabilistic constraint:

$$\forall_{j,k} \quad \Pr\left\{\sum_i\left(\Psi_{ik} - G_i\varepsilon^{\max}_{jk}\right)x_{ij} \le 0\right\} \ge \alpha' \tag{34}$$

Again, to simplify, it is helpful to define a new variable, $\phi^{\max}_{\mathrm{ijk}}$, such that:

$$\forall_{i,j,k} \quad \phi^{\max}_{\mathrm{ijk}} \equiv \Psi_{ik} - G_i\varepsilon^{\max}_{jk} \tag{35}$$

$$\text{where } \bar{\phi}^{\max}_{\mathrm{ijk}} = \bar{\Psi}_{ik} - \bar{G}_i\varepsilon^{\max}_{jk} \text{ and } \sigma_{(\phi)ijk} = \sqrt{\sigma^2_{(\Psi)ik} + (\varepsilon^{\max}_{jk})^2\sigma^2_{(G)i}} \tag{36}$$

Using these definitions it is possible to convert the probabilistic constraint Eq.(34) into its deterministic equivalent of the form:

$$\forall_{j,k} \quad \sum_i \bar{\phi}^{\max}_{\mathrm{ijk}}x_{ij} + X(\alpha')\sqrt{\sum_i\sum_l \rho_{(\phi)ilk}\sigma_{(\phi)ijk}\sigma_{(\phi)ljk}x_{ij}x_{lj}} \le 0 \tag{37}$$

Using an identical approach it is possible to develop Eq. (37) into an analogous constraint on minimum compositional specifications of the form:

$$\forall_{j,k} \quad \sum_i \bar{\phi}^{\min}_{\mathrm{ijk}}x_{ij} + X(1-\beta')\sqrt{\sum_i\sum_l \rho_{(\phi)ilk}\sigma_{(\phi)ijk}\sigma_{(\phi)ljk}x_{ij}x_{lj}} \le 0 \tag{38}$$

$$\text{where } \phi^{\min}_{\mathrm{ijk}} \equiv \Psi_{ik} - G_i\varepsilon^{\min}_{jk} \tag{39}$$

Replacing equations (20), (21), and (22) with their analogs (31), (37), and (38), respectively, creates a schematic model that explicitly addresses uncertainty in both raw materials composition and yield.

3.3 Monte Carlo simulations

To test the optimal batch solutions given by both the deterministic and stochastic methods, Monte Carlo simulations were run to evaluate the robustness of batch plans to variation in the composition of incoming scrap materials. The Monte Carlo method uses pseudo-random numbers (i.e. not truly random in the sense that they are generated by numerical algorithms) to statistically simulate random variables, in this case, scrap composition.

Chapter 4. Computational complexity and scaling

4.1 Projected scaling

The formulations of the stochastic methods (recourse, penalty function, and chance-constrained) as coded in Lingo by Lindo Systems are listed in Appendix A. The deterministic formulation has been omitted due to its simplicity.

The three main inputs that scale the size of the batch mixing problem are the number of products *j* (alloys in this case study), number of raw materials *i* (this includes scrap, primary aluminum, and alloying elements), and the number of compositions *k* being tracked (elements such as Si, Mg, etc.). For this scaling analysis, the chance constrained formulation without yield as an addition (see Eq. through) was compared to a deterministic linear programming batch plan. The addition of yield to the model does not significantly change these results as the yield terms are absorbed into the stochastic composition number. The recourse and penalty function stochastic methods are also presented. The effect of these input parameters is presented for the number of iterations required for convergence (either to global or local optimum) as well as storage, which would be of interest to users of these optimization techniques. For the deterministic method the total number of decision variables will be i x j (*O(mn)*) and the total number of constraints will be *i* (availability of raw materials) + *j* (batch size constraints) + 2 x *j* x *k* (compositional constraints); this will be dominated by the compositional terms and therefore will be *O(2mn)*.

For the chance constrained method, the number of decision variables will remain unchanged but the number of compositional constraints will now double due to the addition of the variance parameters described in section 3.1.2 and will therefore be $O(mn^2)$. More significant to the scaling however is that this includes the introduction of both non-linear decision variables and constraints which will greatly affect the computational intensity as explored further in Section 4.1.2. The penalty function has similar scaling to the chance-constrained formulation with the exception that the variance terms are added to the objective function instead of the compositional constraints. This will increase the number of decision variables because instead of an confidence interval (α or β) being specified and fed into the model, it now becomes a free variable. The compositional constraints for the penalty function formulation are the same as for the deterministic linear programming formulation. The recourse model has significantly different scaling than the other stochastic methods due to having an additional scaling parameter of z or the number of probability levels to represent the compositional scenarios. This reflects the granularity of the discrete representation of a continuous function, therefore, higher z reflects better accuracy. The recourse formulation will therefore be $O(2mnz^2)$.

Table I. Projected scaling results for three stochastic methods (CC-chance constrained, Pen-penalty function, and Rec- multi-stage recourse) compared to deterministic (Det) linear program for small scale case and typical production size case

	Det	CC	Pen	Rec	Det	CC	Pen	Rec	Rec
Products *j*	2	2	2	2	15	15	15	15	15
Raw Materials *i*	2	2	2	2	25	25	25	25	25
Compositions *k*	2	2	2	2	15	15	15	15	15
Prob. Levels *z*	NA	NA	NA	3	NA	NA	NA	3	6
Decision Variables	4	4	20	24	375	375	1275	2250	4500

Availability	2	2	2	2	25	25	25	25	25
Batch Sizes	2	2	2	2	15	15	15	15	15
Composition	8	16	8	24	450	900	450	1350	2700
Total Constraints	12	20	12	28	490	940	490	1390	2740

4.2 Actual scaling in Lingo

Run time is a very important parameter in determining which methods may be of use for actual batch planning in a metals production facility. Because of the recourse models poor scaling, even at high granularity (low accuracy) for representing the probability distribution of compositional scenarios, it was not included for further analysis. To run such a model would be prohibitively computationally expensive for day-to-day operational decision-making. The two remaining stochastic methods (chance-constraints and penalty function) were compared to a deterministic linear program. Source code for these runs in Lingo by Lindo Systems is available in the Appendix. Table III through Table VII list the number of variables, constrains, and non-zeros (both linear and non-linear – NL) with varying number of raw materials, products, and compositions. The time for local (L) or global (G) convergence are reported in the following section (4.3). Because of slightly different implementations in Lingo, one will notice that the numbers of these parameters do not exactly match those listed in Table I for projected scaling. Of note are the number of non-linear decision variables and constraints for the stochastic methods as well as the strong scaling dependence of the compositional constraints(Table VI and Table VII).

Table II. Run parameters with varying number of raw materials (products=15 and compositions=10)

Variables	733	415	665	351	563	255	478	175	398	95
Variables-NL	570	0	510	0	420	0	345	0	345	0
Constraints	674	374	666	366	654	354	644	344	639	339
Constraints-NL	600	0	600	0	600	0	600	0	600	0
Nonzeros	7244	6612	6205	5590	4207	3282	2542	1617	1880	1280
Nonzeros-NL	6000	0	4800	0	3000	0	1500	0	1500	0
Products	15	15	15	15	15	15	15	15	15	15
Raw Materials	25	25	21	21	15	15	10	10	5	5
Compositions	10	10	10	10	10	10	10	10	10	10
Det/CC	CC	Det	CC	Det	CC	Det	CC	Det	CC	Det
Local/Global	L	G	L	G	L	G	L	G	L	G

Table III. Run parameters with varying number of raw materials (products=15 and compositions=10) for penalty function

Variables	1015	951	855	775	695
Variables-NL	675	615	525	450	375
Constraints	674	666	654	644	638
Constraints-NL	600	600	600	600	600
Nonzeros	7843	6511	4513	2848	2267
Nonzeros-NL	5910	4710	2910	1410	990
Products	15	15	15	15	15
Raw Materials	25	21	15	10	5

Compositions	10	10	10	10	10
Local/Global	L	L	L	L	L

Table IV. Run parameters with varying number of products (raw materials=5 and compositions=10)

Variables	787	185	527	125	398	95	164	41	86	23
Variables-NL	660	0	440	0	345	0	138	0	69	0
Constraints	1268	669	848	449	639	339	261	141	135	75
Constraints-NL	1200	0	800	0	600	0	240	0	120	0
Nonzeros	3265	2548	2181	1703	1880	1281	760	520	387	267
Nonzeros-NL	2400	0	1600	0	1500	0	600	0	300	0
Products	30	30	20	20	15	15	6	6	3	3
Raw Materials	5	5	5	5	5	5	5	5	5	5
Compositions	10	10	10	10	10	10	10	10	10	10
Det/CC	CC	Det	CC	Det	CC	Det	CC	Det	CC	Det
Local/Global	L	G	L	G	L	G	G	G	G	G

Table V. Run parameters with varying number of products (raw materials=5 and compositions=10) for penalty function

Variables	1385	925	695	281	143
Variables-NL	750	500	375	150	75
Constraints	1268	848	638	260	134
Constraints-NL	1200	800	600	240	120
Nonzeros	4519	3018	2267	914	463
Nonzeros-NL	1980	1320	990	396	198
Products	30	20	15	6	3
Raw Materials	5	5	5	5	5
Compositions	10	10	10	10	10
Local/Global	L	L	L	L	L

Table VI. Run parameters with varying number of compositions (raw materials=15 and products=15)

Variables	863	255	713	255	563	255	443	255	383	255
Variables-NL	720	0	570	0	420	0	300	0	240	0
Constraints	1254	654	954	504	654	354	414	234	294	174
Constraints-NL	1200	0	900	0	600	0	360	0	240	0
Nonzeros	7124	5564	5519	4379	4207	3194	2714	2264	2011	1681
Nonzeros-NL	6000	0	4500	0	3000	0	1800	0	1200	0
Products	15	15	15	15	15	15	15	15	15	15
Raw Materials	15	15	15	15	15	15	15	15	15	15
Compositions	20	20	15	15	10	10	6	6	4	4
Det/CC	CC	Det	CC	Det	CC	Det	CC	Det	CC	Det
Local/Global	L	G	L	G	L	G	L	G	G	G

Table VII. Run parameters with varying number of compositions (raw materials=15 and products=15) for penalty function

Variables	1455	1155	855	615	495
Variables-NL	825	675	525	390	300
Constraints	1254	954	654	414	294

Constraints-NL	1200	900	600	360	240
Nonzeros	8318	6413	4508	2978	2156
Nonzeros-NL	5670	4290	2910	1800	1200
Products	15	15	15	15	15
Raw Materials	15	15	15	15	15
Compositions	20	15	10	6	4
Local/Global	L	L	L	L	L

4.3 Iterations to convergence

In Figure 12 through Figure 14, one can see the effect that the number of raw materials, products, and compositions has on the number of iterations required for convergence of the three methods, either to global optimum for the deterministic case or a best local optimum for the case of the stochastic methods. Lingo, an optimization solver package was used to perform these numerical experiments. A multi-start barrier method is used for solving the non-linear stochastic problems while a typical LP solver is used for the deterministic case. The case sizes presented (<25 raw materials, <30 products, <20 compositions) reflect those that are able to converge in a reasonable amount of time to be useful in batch planning (less than an hour). One can see that the stochastic methods require substantially more iterations to converge when compared to the deterministic method. This is due to the non-linearities present in either the compositional constraint (CC) or the objective function (Penalty). Most of the graphs would suggest that the scaling, while certainly not linear, is not quite exponential or power law. The number of compositions considered causes the most rapid scaling of the stochastic methods as would be expected due to the additional computations dealing with uncertainty directly relate to the number of elements being considered. In particular, the penalty function formulation scales very rapidly with number of compositions considered, hence only including data points in the lower numbers (Figure 14).

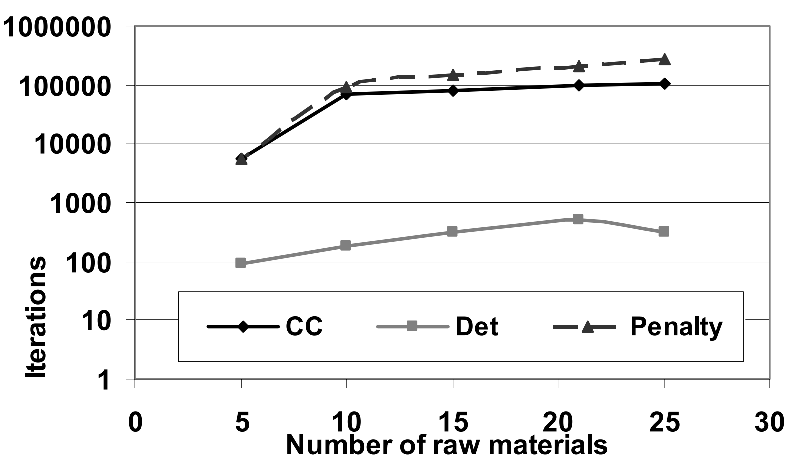

Figure 12. Number of iterations to global or local optimum convergence for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of raw materials

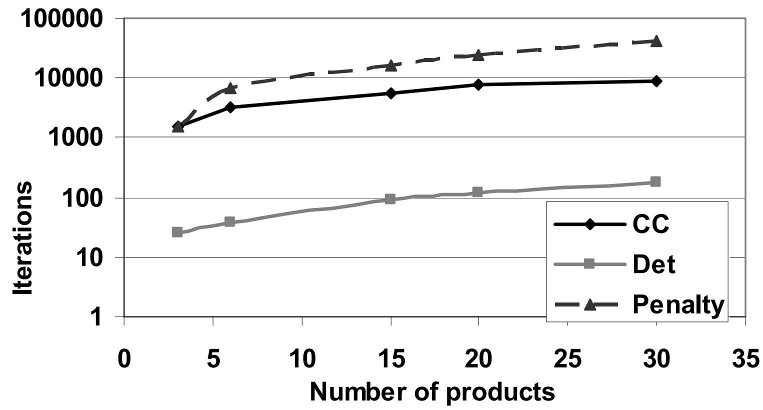

Figure 13. Number of iterations to global or local optimum convergence for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of products

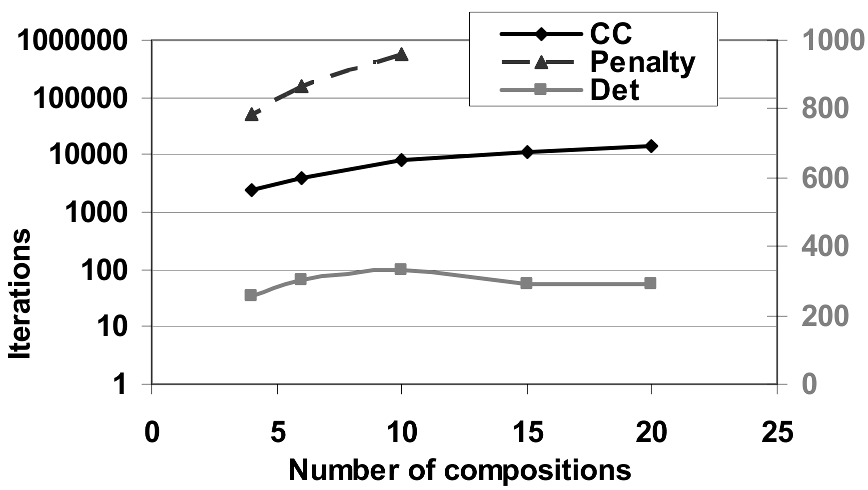

Figure 14. Number of iterations to global or local optimum convergence for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of compositions

4.4 Storage and memory

In Figure 15 through Figure 17 one can see the effect on the memory required for each of the three methods as it scales with products, raw materials, and composition. One will notice these are now plotted on a linear scale for both the x and y-axis. Generally, the chance constrained method requires the most storage and each scale linearly with the parameters shown. Often the amount of storage required can be an indicator for poor performance of some methods due to the number of cache misses caused by moving parts of the problem into smaller computer storage sections. However, these results would indicate that the poor performance of the penalty function method in terms of number of iterations is more likely due to its actual formulation.

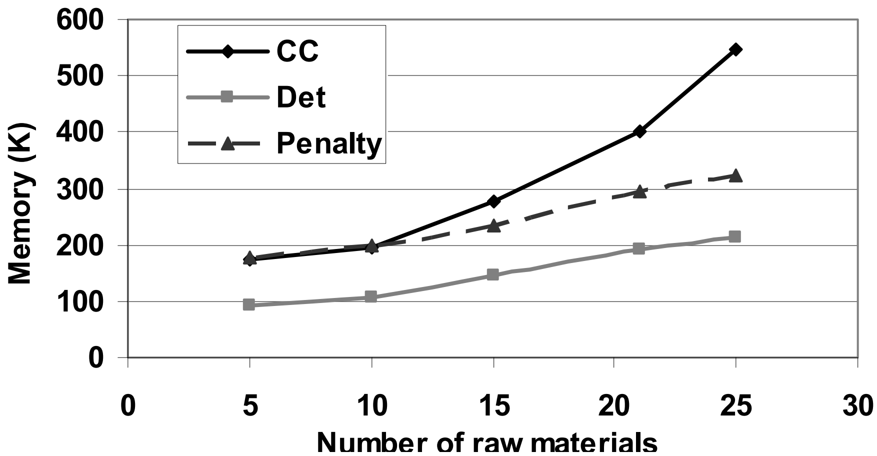

Figure 15. Memory for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of raw materials

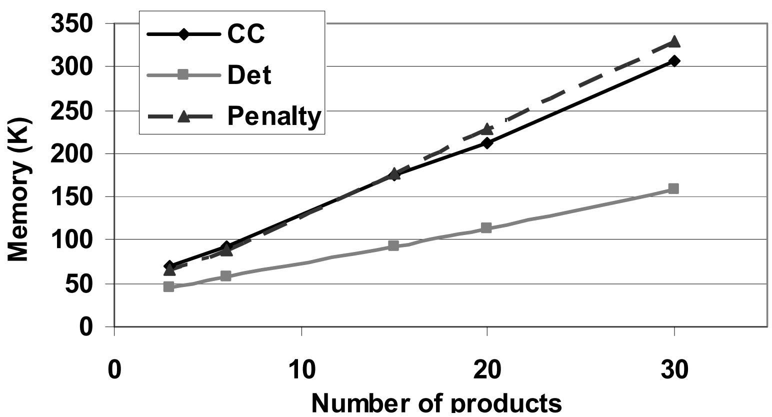

Figure 16. Memory for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of products

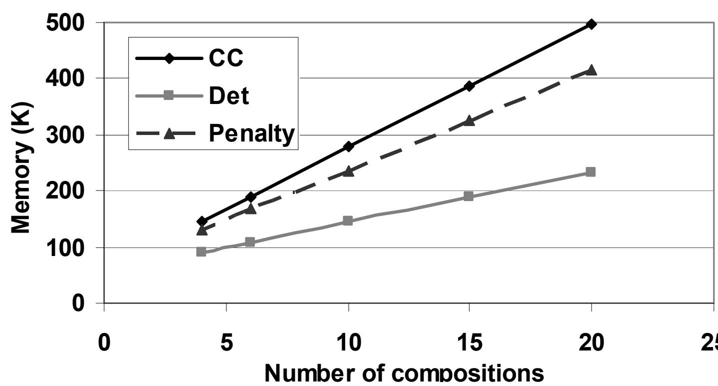

Figure 17. Memory for the chance-constrained formulation (CC) and the penalty function formulation (Penalty) vs. the deterministic batch planning model (Det) with number of compositions

While these scaling results are informative as to the robustness and usefulness of the methods explored, one of the most important features that will be differences in the optimal batch plans themselves. Metrics of interest to metals batch planners will not be performance of the algorithm but the error rate for the batch plan, ie. the probability of having one of the mixture's elements be outside of the alloy specifications. Another highly important metric will be the production cost of the batch; many cast-shop operators also want to increase the utilization of scrap or secondary materials in the batch. These goals often go hand in hand as scrap materials have a lower cost when compared to virgin metal or alloying elements, however, the compositional constraints will usually limit the amount of secondary that can be utilized in the batch plan.

The scaling analysis points strongly to the chance-constrained model as being the stochastic programming method of choice for managing compositional uncertainty within linear programming batch plans. Therefore, it will be the sole stochastic method explored further for its implications on the metrics of production cost, scrap utilization, and batch error rate compared to the prevalent deterministic linear programming used today.

Chapter 5. Managing compositional uncertainty in blending decisions

5.1 Case study set-up

To evaluate the impact of the chance-constrained formulation on operational decisions as well as both scrap use and economic outcomes, a hypothetical case study involving the production of both cast and wrought alloys was developed. Specifically, the case involved the production of equal amounts of 380, 390, 3003 and 6061 finished goods from seven secondary raw materials and appropriate primary raw materials. Prices and compositions used within the model for both input materials and the finished alloy products are summarized in Table VIII and Table X, respectively. Average prices on primaries and scrap materials as well as recent prices on alloying elements were taken from the London Metals Exchange. The particular scraps and scrap compositions used in this case are based on studies by Gorban[86] reflecting the scrap materials which would be expected to derive from post-consumer automotive scrap. It is important to note that these compositions assume some imperfect separation and uptake of non-aluminum materials particularly iron.

Finished good compositional specifications are based on international industry specifications and do not reflect production targets of any specific firm. In order to ensure that results are not biased towards any particular product type, all products were modeled using the same average demand. Furthermore, all raw materials were assumed to be unlimited in availability in order to avoid the potential effects of limited raw materials supplies. The model framework presented herein can be used for cases of constrained scrap supply with no modification. Notably, van Schaik, Reuter, and others [34, 35] have shown that the physical form of incoming raw materials, in particular particle size, can have real impact on the metal yield from a given feedstock. Also, some uncertainty also exists with regards to this metal yield for specific elements within incoming scrap materials. Due to the reactivity of the aluminum and the volatility of many of the alloying elements, melt loss can occur in a manner that does influence the composition and final metal yield of the finished goods[87-90]. Although treatment of these effects is important, the presented model omits these in the interest of simplicity and to focus on uncertainty. Model formulations intended to address business critical decisions should carefully consider the extension of the framework presented herein to comprehend particle size effects and specific element yield uncertainty.

For the base case, within the chance-constrained formulation, the scrap raw materials were modeled with a coefficient of variation (COV) of 50% on composition for all elements. Sensitivities around this number were also explored. For the deterministic window narrowing runs, the maximum and minimum scrap compositions were set to be two standard deviations (SD) from the mean values reported in; COVs and SDs are noted for each scenario. For chance-constrained runs, compositions were also assumed to be perfectly uncorrelated.

For this case, a normal distribution around the compositional mean of each of the scraps elemental considerations (Si, Mg, Fe, Cu, Mn, Zn) was assumed with mean and variance according to the assumptions used in the optimization model. The optimal solution (either chance- constrained or window narrowing) was tested against samples from these distributions 10,000 times. The number of batches that would have incurred any errors (i.e., the final composition of finished alloys would have been expected to fall out of specification) was recorded. This is reported subsequently as the Error Rate with one minus this amount being the batch Success Rate. Figure 18A shows 10,000 simulated magnesium compositions for alloy 390 while Figure 18B shows the same for copper compositions in alloy 380. Target alloy specifications are also noted on the graph; due to the high level of success rate specified by the α and β parameters, one can barely see the infrequent cases where the alloy is outside of specification.

Table VIII. Prices of Raw Materials

Primary Al & Elements	Cost / Metric Ton	Scrap Materials	Cost / Metric Ton
P1020 (Al)	$1,360	Brake	$1,000
Silicon (Si)	$1,880	Transmission	$1,000
Magnesium (Mg)	$2,270	Media Scrap	$1,000
Iron (Fe)	$320	Heat Exchange	$1,000
Copper (Cu)	$2,660	Bumper	$1,000
Manganese (Mn)	$2,020	Body Sheet	$1,000
Zinc (Zn)	$980	All Al Eng. & Trans.	$1,000

Table IX. Mean Compositions of Scrap Materials

Raw Materials	Mean Compositions (wt. %)					
	Si	Mg	Fe	Cu	Mn	Zn
Brake	1.54	1.23	0.40	0.62	0.14	0.12
Transmission	10.30	0.21	0.90	3.79	0.28	2.17
Media Scrap	4.88	0.64	0.53	1.00	0.11	1.00
Heat Exchange	2.88	0.21	0.44	0.68	0.59	0.20
Bumper	0.39	0.78	0.38	0.32	0.09	0.75
Body Sheet	0.47	1.34	0.21	0.57	0.19	0.07
All Al Engines	8.61	0.30	0.68	2.69	0.27	1.26

Table X. Finished Goods Chemical Specifications

Finished Alloys		Specifications (wt. %)					
		Si	Mg	Fe	Cu	Mn	Zn
380	Min	7.50	0	0	3.00	0	0
	Max	9.50	0.20	2.00	4.00	0.50	3.00
390	Min	16.0	0.45	0	4.00	0	0
	Max	18.0	1.30	1.30	5.00	0.10	0.10
3003	Min	0	0	0	0.05	1.00	0
	Max	0.60	0.05	0.70	0.20	1.50	0.10
6061	Min	0.40	0.80	0	0.15	0	0
	Max	0.80	2.40	0.70	0.40	0.15	0.25

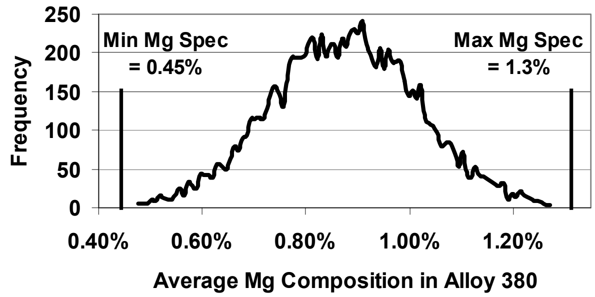

A)

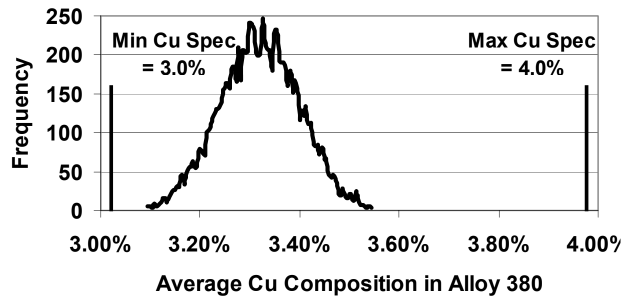

B)

Figure 18. Example Monte Carlo simulations of A) magnesium composition in alloy 390 and B) copper composition in alloy 380 for the base case chance constrained batch plan. Target alloy specifications are indicated by the bars; due to the high success rate specified by the α and β parameters, cases where the alloy is out of specification are infrequent. For all chance constrained cases, α was set to be 99.99% and β was set to be (1-α) or 0.01% unless otherwise indicated.

5.2 Base case results

Base case parameters for the mean-based method (MB), chance-constrained (CC) method, and the window narrowing (WN) method are summarized in Table XI. Results (Table XII) show that though the mean-based method gives an optimal solution that uses more scrap and thereby costs less than the other two methods, it has an extremely high expected error rate of 98.6%. This indicates that using this optimal solution for batch production for scraps with distributions in compositions as assumed here would be expected to result in at least one finished good composition being out of specification 98.6% of the time. Figure 19 shows that getting an error rate for the mean-based method to be equivalent to that for the chance-constrained method is not possible with any degree of window narrowing for the Base Case conditions. Even with decreased variation in scrap composition (COV=20%), the error rates for the mean based method reduce only to 43% with maximum window narrowing, a value still not comparable to the chance-constrained method (Figure 20). Because of this high error rate, only the compositional scrap window combined with specification narrowing (window narrowing) method will be considered in further analysis.

When comparing the CC method versus the WN method, both exhibit comparable and high expected success rates of 97.7% and 98.0% respectively (cf. Table XII). However, results show a marked improvement in scrap consumption of 29.7% with an associated decrease in primary purchased for the CC method versus the WN method. The chance-constrained method shows a slight improvement in production cost of 2.0%. Although these financial benefits are small, they come with no required investment or change in technology; they emerge strictly through improved planning tools. Table XIII shows the detailed batch plan results for the three methods. As will be explored in detail later, one of the most striking features of the detailed batch plans is the variation in diversity of scrap use between the deterministic and CC methods. Universally, the CC method uses a considerably more diverse set of scrap materials.

Table XI. Parameters for the base case analysis

		Confidence Interval (α)	Window Size	Std. Dev.	Variability (COV)
Mean Based	MB	—	100%	0	50%
Window Narrowing	WN	—	70%	2	50%
Chance Constrained	CC	99.99%	—	—	50%

Table XII. Base case results showing mean based method and comparison of chance-constrained method versus 30% window narrowing (70% of original composition window) case at 2 standard deviations and COV=50% (kt = kilotonnes)

	Mean Based	Window narrowing	Chance-Constrained	Improvement of CC over WN method
Scrap Usage	48.96 kt	15.75 kt	20.44 kt	↑29.7%
Production Cost	$1.17M/kg	$1.34M/kg	$1.31M/kg	↓2.0%
Success Rate	1.4%	98.0%	97.7%	
Error Rate	98.6%	2.0%	2.3%	

Table XIII. Base case results by alloy showing chance-constrained method versus 30% window narrowing case at 2 standard deviations and COV=50%. (MB = Mean Based, WN= Window narrowing, CC = Chance-Constrained)

	Alloy 380			Alloy 390			Alloy 3003			Alloy 6016		
	MB	WN	CC	MB	WN	CC	MB	WN	CC	MB	WN	CC
Brake	-	-	0.12	11.91	4.72	3.51	-	-	-	-	-	1.03
Transmission	16.19	-	0.41	-	-	-	-	-	0.03	-	-	0.03
Media Scrap	-	-	0.67	-	-	0.18	-	-	0.08	-	-	0.16
Heat Exchanger	-	4.79	3.78	-	0.04	0.05	3.23	1.45	1.33	-	-	0.44
Bumper	-	0.50	0.77	0.36	0.17	0.43	0.22	0.06	0.19	4.73	2.16	2.12
Body Sheet	-	-	0.06	-	-	2.20	-	-	-	10.07	1.86	2.22
All Al Engine	-	-	0.53	-	-	0.02	-	-	0.06	-	-	0.05
Total	16.19	5.30	6.33	12.27	4.93	6.38	3.45	1.51	1.69	14.79	4.02	6.05

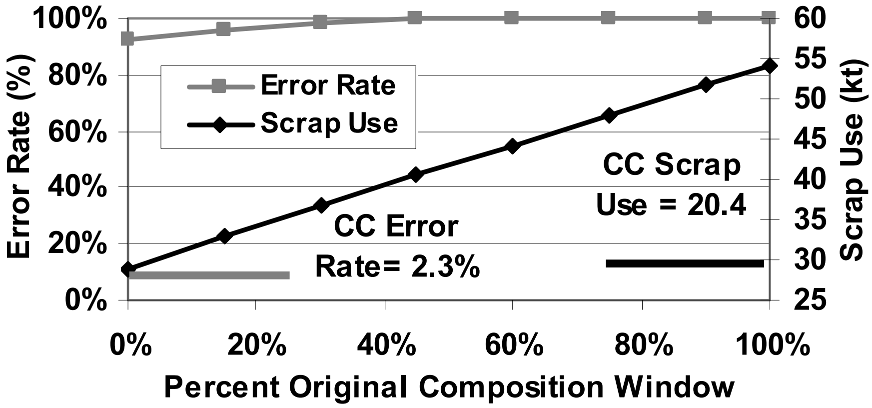

Figure 19. Error rates and scrap use for the mean based method at various degrees of window narrowing. Over all conditions, the error rate is much higher than the chance-constrained method when the COV=50%

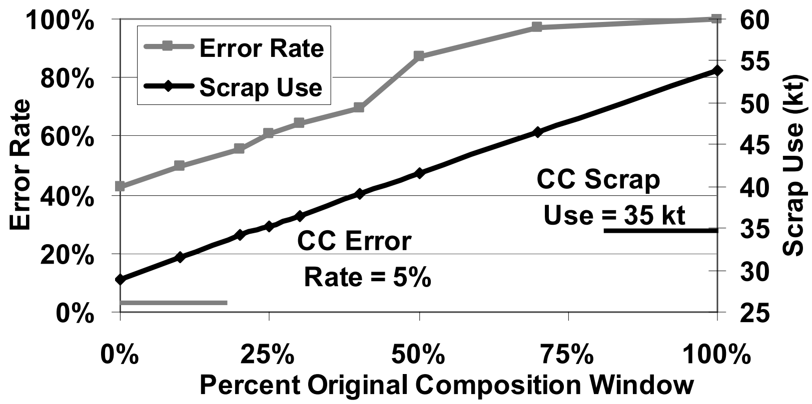

Figure 20. Error rates and scrap use for the mean based method at various degrees of window narrowing. At 75% window narrowing (25% original compositional window), the mean based method has equivalent scrap use to chance-constrained method but still a much higher error rate even with reduced variation (COV=20%)

Interestingly, although the CC method yields improvements in scrap consumption and associated economic benefits across a broad range of conditions, the degree of improvement in scrap consumption varies considerably across both alloys and scrap types. When comparing the individual production of each alloy, 6061 shows the largest improvement in scrap usage while 3003 shows the smallest at 51% and 12%, respectively (Figure 21). Comparing the impact on the various scraps, "Body Sheet" scrap showed an improvement in consumption of 16% using the chance-constrained method while "Bumper" scrap showed a much smaller increase of 1.2% (Figure 22). More interestingly, some of the scraps (Transmission, Media Scrap, and All Al Engines) which were utilized only a little or not at all when using the window narrowing method are now being utilized using the chance-constrained method. This suggests that this stochastic method is enabling more scrap opportunities; scrap consumption is not just being increased but opportunities for using lower quality scraps are being identified. This chance-constrained result of "diversifying" the scrap portfolio utilized is very noticeable in Table XIII, the CC method clearly uses more scrap types than either the mean based or window narrowed methods. This mechanism of diversification will be explored further in section 5.3. The increased opportunities for lower quality scraps also indicates that the assumption of all scraps being equal in price

(Table VIII) is most probably muting this effect in regards to cost savings. Scraps that have higher variability in composition would be considered "dirtier" and would therefore have lower prices. Identifying uses for these types of scrap would cause the chance-constrained method to be even lower in cost than indicated in the base case results.

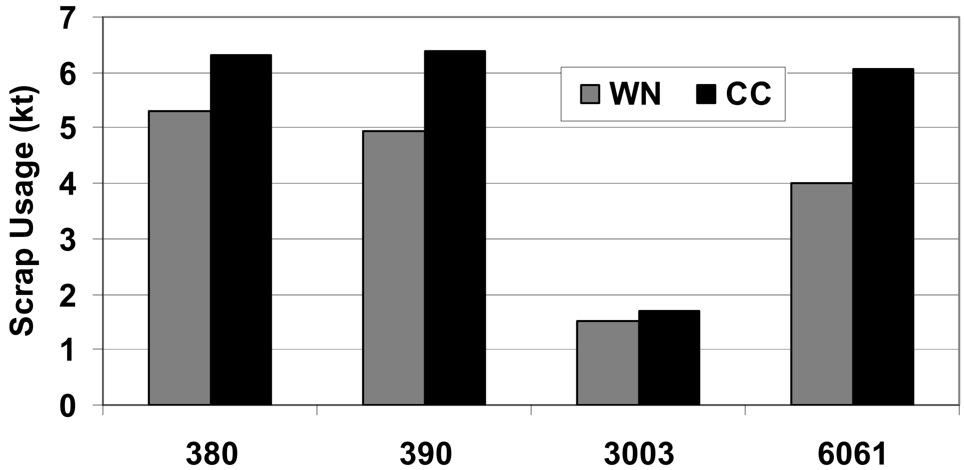

Figure 21. Comparison between base case results for chance-constrained and 30% window narrowing showing varying degrees of improved scrap consumption for the production of individual alloys (50% COV, 2 SD)

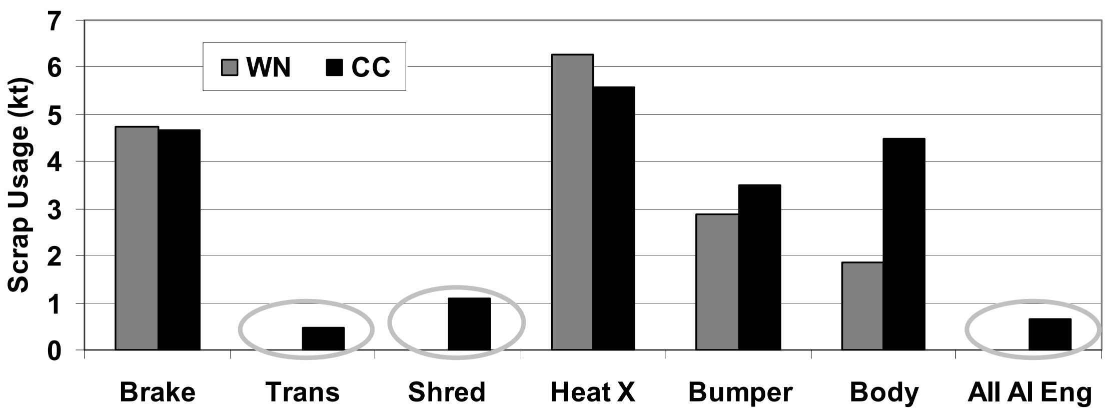

Figure 22. Comparison between base case results for chance-constrained and 30% window narrowing showing varying degrees of improved scrap consumption across scrap types. Scraps where the CC method finds usage opportunities are circled (50% COV, 2 SD)

5.3 Exploring model details: sensitivity analysis

A key question with regard to any form of modeling, and particularly with optimization, is the impact of changes in model assumptions. For the problem at hand, there are many assumptions about operating conditions and raw material characteristics. The following sections will examine the impact of 1) window sizes, 2) compositional limits, 3) confidence intervals, 4) uncertainty, and 5) shadow prices with regard to the ability of the CC method to identify batch mixing solutions that deliver improved economics and scrap use.

5.3.1 Window narrowing

One assumption that may be difficult to determine *a priori* (for reasons outlined in Section 2.3) is what size window to use for window narrowing. Figure 23 shows how scrap usage drops for the Base Case as the finished goods window target gets smaller[2]. The window narrowing solution shows scrap usage equivalent to the CC solution at a window that is approximately 87% of the original specification window. However, as Figure 23 shows, at an 87% window, the WN derived solution would be expected to have a much higher error rate than the CC solution. Conversely, with a 70% window, the WN method shows an error rate equivalent to the CC method (but with an inferior scrap usage). These results were used to identify the Base Case comparison (CC vs. 30%WN) described above. These figures also reveal that the WN method can only approach or exceed the scrap usage levels of the CC method by sacrificing reliability in the form of higher error rates.

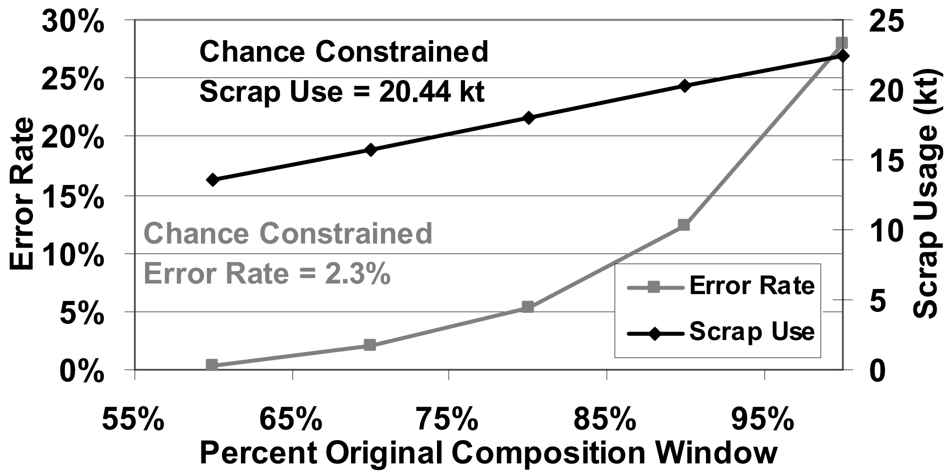

Figure 23. Scrap usage and error rate of the chance-constrained method compared to varying window sizes for the window narrowing method. A window narrowing size of 30% (70% original composition window) was chosen for the base case due to nearly equivalent error rate with the chance constrained method

5.3.2 Controlling error rate in both methods

In the window narrowing method, the size of the compositional window used to represent the scraps composition (refer back to Figure 8 & Figure 9) is another model variable. As Figure 24 shows, large compositional windows ensure low error rates but penalize the systems ability to utilize scrap effectively, while small compositional windows allow increased scrap use but with higher error rates.

[2] Because decreases in scrap usage cause monotonic increases in production cost, only scrap usage and error rate will be reported for this sensitivity analysis section.

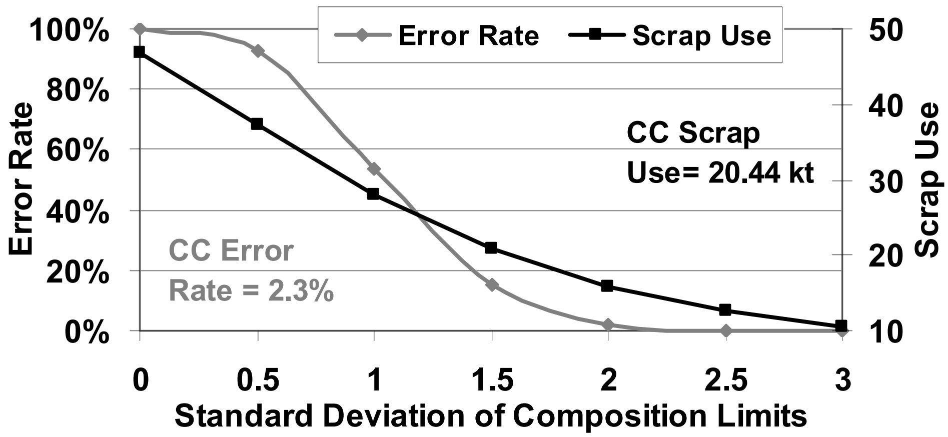

Figure 24. Scrap use and error rates for varying scrap compositional window representations for the window narrowing method (COV=50%, 30% WN). For a given variation in scrap composition, the finished good specification window size controls the error rate for this method. To maintain constant error rate, window size must be adjusted whenever scrap characteristics or availability change.

In the chance-constrained method, the error rate can be tailored by adjusting the confidence intervals on the compositional specifications, specifically the α and β values in Equations (16) and (17). Similarly to the window narrowing method, as the confidence intervals are tightened the scrap use decreases rapidly with a related decrease in the number of expected missed batches (Figure 25). The major difference between the methods in this regard though is that the chance-constrained method uses the confidence interval as a direct link to controlling the error rate while the window narrowing method requires trial and error experimentation and does not dynamically adjust to variability of the scrap portfolio. This concept will be explored further in the next section.

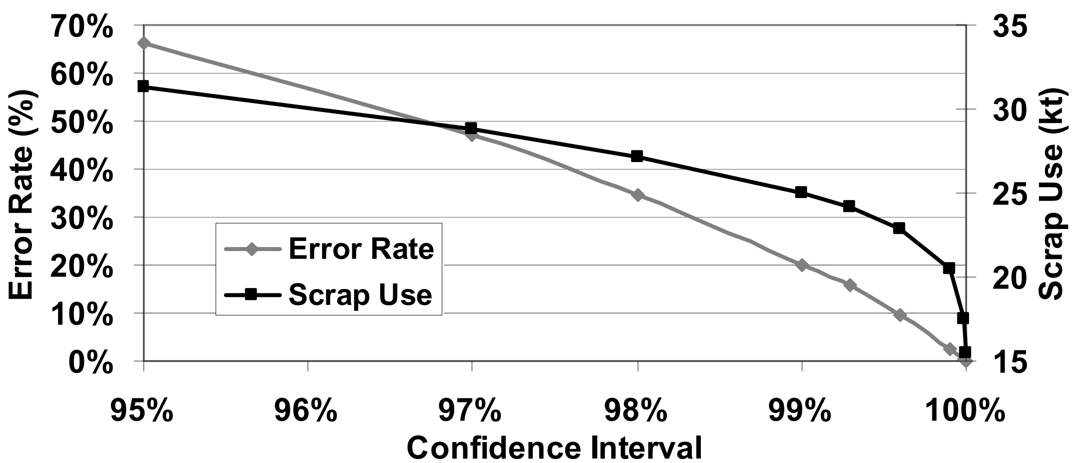

Figure 25. Scrap use and error rates for varying confidence intervals for the chance-constrained method. Error rate for the CC method is easily controlled through the alpha value specification and does not change significantly with scrap characteristics or availability. The base case alpha value of 99.90% gives a suitably low batch error rate.

5.3.3 Increasing scrap variability

One of the novel pieces of information that is being considered in this paper is the magnitude of uncertainty in the composition of the scrap material. Figure 26 shows the impact of that uncertainty on the advantage afforded by the chance-constrained solution, here expressed as the improvement over the WN approach at three different window sizes: 70%, 80%, and 90% of the original composition window. The CC solution outperforms the WN-derived solutions for 70% and 80% windows across the range of examined variation, with benefits as high as 36% additional scrap use. Compared to the 90% window, the CC-derived solution uses an equivalent amount of scrap at 10% COV and additional scrap for COVs equal to or above 50%. Although the CC solution suggests using less scrap than the 10% WN solution for COVs between 10% and 50%, the 10% WN solutions lead to dramatically higher expected error rates (cf. Figure 28). As shown in Figure 28, variability causes less error rate fluctuations for the chance-constrained method which maintains an error rate of approximately 2% while the window narrowing method increases from 0% to 2% for a 70% window, from 1% to 7% for an 80% window, and from 2% to 14% for a 90% window (Figure 28). Increasing scrap variability also drives up production cost; Figure 27 shows this increase to be approximately the same for all methods, but with the CC solution providing equal or better costs than each for the three WN solutions across the range of variability.

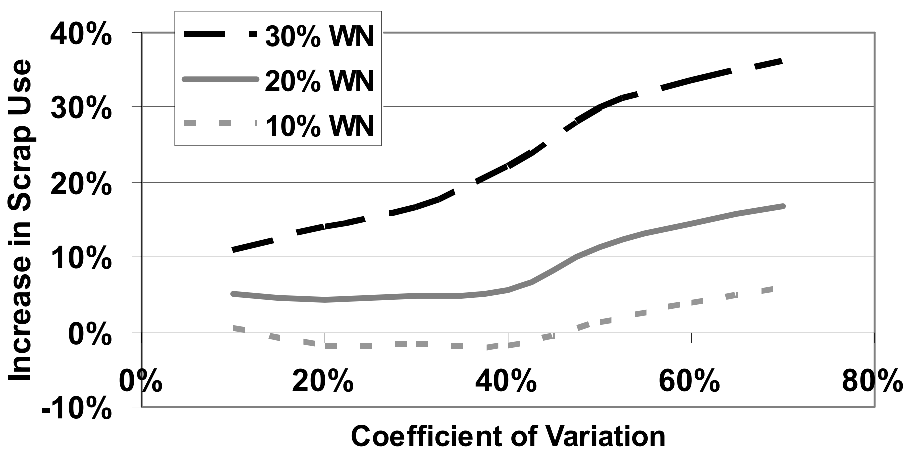

Figure 26. Improvement in scrap utilization of chance-constrained method over 30%, 20%, and 10% window narrowing methods (70%, 80%, and 90% windows) for a range of uncertainty conditions (SD=2) Base Case Comparison: COV=50%, 30% WN

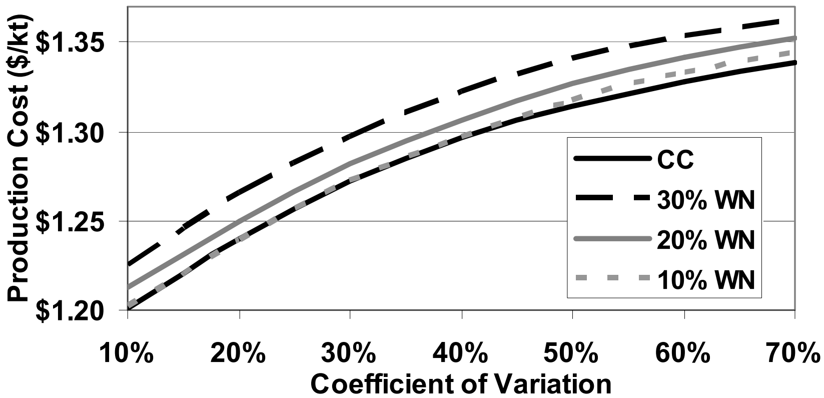

Figure 27. Improvement in production cost of chance-constrained method over 30%, 20%, and 10% window narrowing methods (70%, 80%, and 90% windows) for a range of uncertainty conditions (SD=2) Base Case Comparison: COV=50%, 30% WN

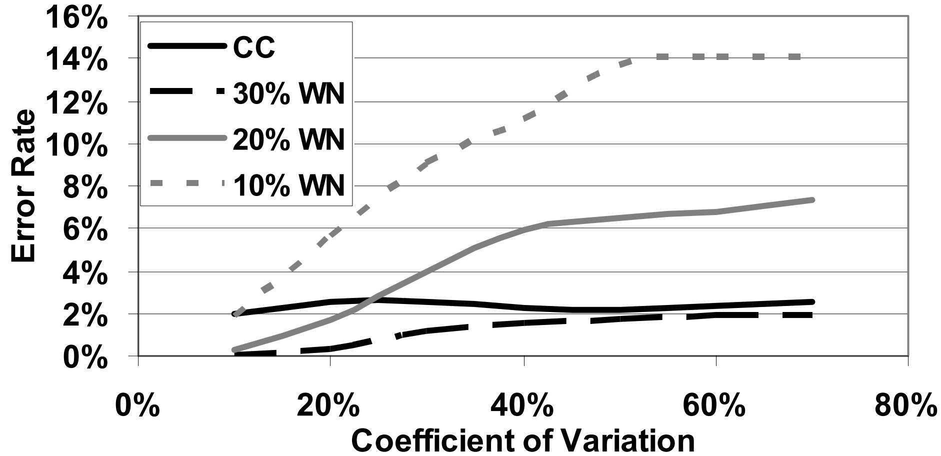

Figure 28. Fluctuating error rates with increasing scrap variability at various window sizes for all variations of window narrowing method. CC method maintains consistent error rates. Base Case Comparison: COV=50%, 30% WN, SD=2

5.3.4 Shadow prices

The results of a linear optimization problem are a set of decision variables that give the optimal objective function. In the case of secondary alloy production planning, these decision variables are the amounts of scrap and primary raw materials to be purchased. However, linear optimization solutions also provide a powerful set of information that quantifies the sensitivity of these results to changes in assumptions. These sensitivity parameters are known as "shadow prices". Specifically, a shadow price is the change in the objective function at the optimum when a specific constraint is changed by one unit[36] as expressed in Equation (40). Each shadow price has a range of validity associated with it.

$$SP_{\text{Constraint}} = \frac{\delta(\text{Production Cost})}{\delta(\text{Constraint})} \quad (40)$$

For the problem at hand, there are three potential classes of shadow prices that are reported, one for each of the types of constraints. These are shadow prices on 1) availability (derived from Eq.

(19), 2) demand (derived from Eq. (20), and 3) composition from the chance-constraints (Eq. (21) and (22)). Because none of the cases examined in this paper placed any constraint on the quantity of scrap available, no shadow price on availability was reported. Shadow prices on demand have the potential to guide material producers in choosing raw materials that provide both economic and environmental benefit; shadow prices on composition can help product designers to choose appropriate alloys for both use and end-of-life. Interested readers should consult [11, 91] for a lengthier discussion on the value of this information for decision-makers.

Shadow prices on demand are positive as seen in Table XIV; if demand were to increase then the total production cost would increase. Studies[11] have shown shadow prices on demand to be useful for directing producers with regard to alloy choices. Differences in shadow prices can provide targets for improving production costs; for example, taking a higher shadow price alloy and substituting part of its production with a lower shadow price alloy would be economically beneficial. Technological issues may prevent this specific substitution of alloys but the example still illustrates how shadow prices can be utilized to leverage alloy choices. For Base Case conditions, the CC model generates shadow prices on demand that are universally lower than the WN method. Although this sample of alloys is too small for generalized observation, for Base Case conditions there is no reordering of the shadow prices compared to the WN result. More extensive investigation would need to be carried out before reaching any conclusion on the implications of the CC method on past recommendations of alloy substitution.

Table XIV. Shadow Prices on Alloy Demand for Base Case comparing CC vs. WN

	WN	CC
6061	$1.300	$ 1.256
380	$1.323	$ 1.295
3003	$1.336	$ 1.330
390	$1.407	$ 1.375

Table XV shows the shadow prices and "modified" shadow prices[3] on binding compositional constraints for all of the finished goods and compositional alloying elements as reported by the both the CC and WN models for Base Case conditions. The table is ordered according to the absolute magnitude of the WN-derived modified shadow price. As has been shown previously in this paper, compositional constraints have one of the largest effects on the optimized scrap use and production cost. For example, in the window narrowing method, if the constraint on the minimum allowable level of copper in alloy 6061 were tightened (increased by 1%), the production cost would increase by $6.23M/80 kilotonne total production. For the chance-constrained method, if the maximum allowable level of manganese in alloy 390 were loosened (raised by one unit), then the production cost would decrease by $17.89M/ 80 kilotonne total production.

[3] The modified shadow price on composition is derived in (91.Cosquer, A., *Optimizing the Reuse of Light Metals from End-of-Life Vehicles*, in *Engineering Systems Division*. 2003, Massachusetts Institute of Technology: Cambridge, MA. p. 98.) as: $\frac{\partial(C(x))}{\partial(U_{fc})} \simeq B_f * \frac{\partial(C(x))}{\partial b} * \frac{1}{100}$ in order to provide more direct guidance on the significance of specific compositional constraints. For this study, since all finished good batch sizes are equal both as-reported and modified shadow prices follow the same rank order.

It is also no surprise that more of the binding constraints are maximums for the chance constrained method (12 of the top 20); the amount of contaminants in a scrap usually determines how much dilution with primary aluminum is required and is therefore the major limiting factor. One will notice that shadow prices on the maximum compositional constraints are negative and the shadow prices on the minimum compositional constraints are positive. The shadow prices on magnesium, copper, and manganese are typically higher because these three alloying elements are the most expensive (>$2,000/ metric ton) and therefore have the highest impact on the production cost. It should be noted that shadow prices have a range of validity around the optimum associated with them and therefore cannot be applied to all cases of changed production.

More interesting than the actual shadow prices themselves though is their relative ranking to each other and across the two optimization methods. Notably, the CC method not only leads to a reordering of WN-derived shadow prices, but completely eliminates several of the more binding constraints derived from the WN result. Specifically, the chance-constrained method eliminates three of the top ten shadow prices (i.e., makes the constraint no longer binding) associated with the window narrowed results; most of these are minimums, namely minimum magnesium in alloy 380, minimum magnesium in alloy 3003, and minimum zinc in alloy 390. The results in Table XV suggest that the CC method may alter the most attractive candidates for alloy modification suggested by Cosquer. Nevertheless, the large range in these shadow prices, even for the CC method, indicates that a tool is necessary in order to systematically and efficiently target for development alloys that would provide the most significant improvements in scrap reuse and production cost. The sensitivity analysis results that emerge from this optimization approach provide just this tool.

Table XV. Shadow prices on binding compositional constraints for CC base case and WN base case (COV=50%, SD=2, window=70%). Values in parentheses are negative.

Constraint	Alloy	Element	WN - SP	Modified SP	Rank	CC	Modified SP	CC Rank
Min	6061	Cu	$31.16	$6.23	1	$23.94	$4.79	5
Max	6061	Cu	($29.86)	($5.97)	2	($22.64)	($4.53)	6
Min	390	Mn	$23.30	$4.66	3	$41.62	$8.32	2
Max	390	Mn	($22.64)	($4.53)	4	($89.47)	($17.89)	1
Min	380	Mg	$20.79	$4.16	5	--	--	30
Min	3003	Mg	$20.79	$4.16	5	--	--	30
Max	3003	Mg	($19.88)	($3.98)	7	($25.86)	($5.17)	3
Max	380	Mg	($19.88)	($3.98)	7	($24.90)	($4.98)	4
Min	390	Mg	$11.69	$2.34	9	$0.91	$0.18	22
Max	390	Mg	($10.78)	($2.16)	10	--	--	30
Max	6061	Zn	($10.72)	($2.14)	11	($10.45)	($2.09)	8
Min	6061	Zn	$10.34	$2.07	12	$10.07	$2.01	9
Max	390	Zn	($9.53)	($1.91)	13	($22.51)	($4.50)	7
Min	390	Zn	$9.15	$1.83	14	--	--	30
Min	380	Si	$5.14	$1.03	15	$3.59	$0.72	12
Min	3003	Si	$5.14	$1.03	15	$0.16	$0.03	29
Max	380	Si	($4.61)	($0.92)	17	($3.06)	($0.61)	13

Max	3003	Si	($4.61)	($0.92)	17	($2.92)	($0.58)	14
Min	380	Cu	$1.30	$0.26	19	$1.30	$0.26	15
Min	390	Cu	$1.30	$0.26	19	$1.30	$0.26	16
Min	3003	Cu	$1.30	$0.26	19	$1.30	$0.26	17
Max	3003	Fe	($1.04)	($0.21)	22	($1.04)	($0.21)	18
Max	380	Fe	($1.04)	($0.21)	22	($1.04)	($0.21)	19
Max	390	Fe	($1.04)	($0.21)	22	($1.04)	($0.21)	19
Max	6061	Fe	($1.04)	($0.21)	22	($1.04)	($0.21)	19
Min	6061	Mg	$0.91	$0.18	26	$0.91	$0.18	22
Min	380	Mn	$0.66	$0.13	27	$0.66	$0.13	24
Min	3003	Mn	$0.66	$0.13	27	$0.66	$0.13	25
Min	6061	Mn	$0.66	$0.13	27	--	--	30
Min	6061	Si	$0.53	$0.11	30	$6.90	$1.38	10
Min	390	Si	$0.53	$0.11	30	$0.53	$0.11	26
Max	3003	Zn	($0.38)	($0.08)	32	($0.38)	($0.08)	27
Max	380	Zn	($0.38)	($0.08)	32	($0.38)	($0.08)	28
Max	6061	Si	$0.00	$0.00	34	($6.38)	($1.28)	11

5.4 Mechanism of benefit: scrap portfolio diversification

It was hypothesized that the chance-constrained method would provide benefit over the current practice of window narrowing because it accurately accounts for the sub-linear increase in variance when combining scraps. Since this effect grows as the number of sampled distributions increases, this would suggest that scrap portfolio diversification would drive the difference in scrap consumption between the two methods. The Base Case results, as shown in Figure 22, seem to support this hypothesis. Clearly, the most striking difference between the chance-constrained and window-narrowing result is the dramatic increase in the number of scraps used. While the window-narrowed model suggests using four scraps, the chance-constrained model suggests using all seven scraps. To explore this hypothesis more directly, two systematic permutations of the Base Case were explored that directly affect the possibility of diversification – one that limits the opportunities for diversification and one that expands those opportunities.

In the first modification of the Base Case, models were executed against scenarios with progressively fewer available types of scrap (ie. 7 types, 6 types, 5 types…). Scraps were eliminated in reverse order of prevalence in use. As such, the first scrap eliminated was "Transmission", the secondary raw material used least by the CC model and not at all by the WN model in the Base Case. Figure 29 shows that for the permuted Base Case with COV=20% and WN=30%, the chance-constrained method begins to provide cost savings as soon as diversification among scrap types is possible, i.e. two scraps and higher. This benefit grows as the number of available scrap types increases. At the Base Case level of compositional variation (COV=50%), Figure 29 shows that the CC method begins to provide benefit at three available scrap types and above. This lag for higher variation conditions is an artifact of the higher error rate of the WN method at one and two scrap types; although the WN method uses more scrap, it does so at a considerably higher error rate (cf. Figure 30). Figure 31 shows the cost savings of the CC method compared to the WN method for both low and high scrap variation. As uncertainty and opportunity for diversification grows, so does the economic benefit of the CC method.

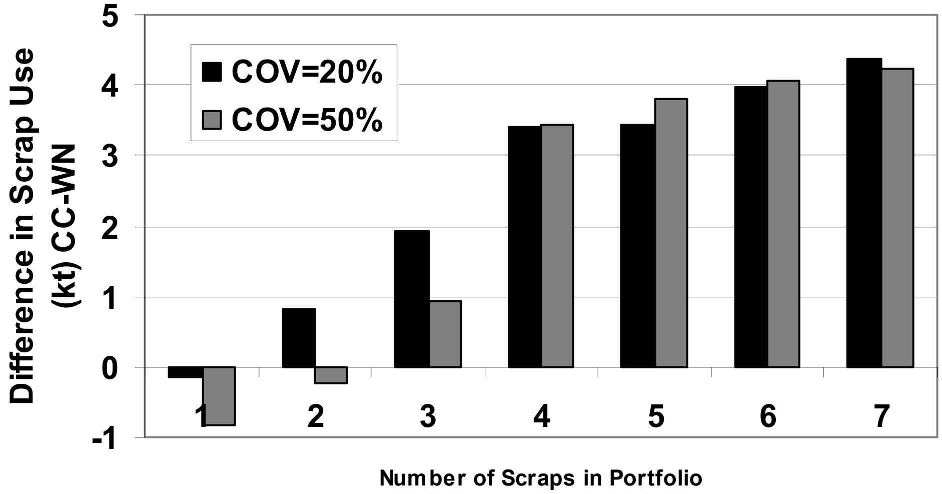

Figure 29. Scrap usage comparison between chance-constrained method and conventional window narrowing with increasing scrap portfolio diversification, taking away in order of least used (30% WN, COV=20%, 50%, SD=2)

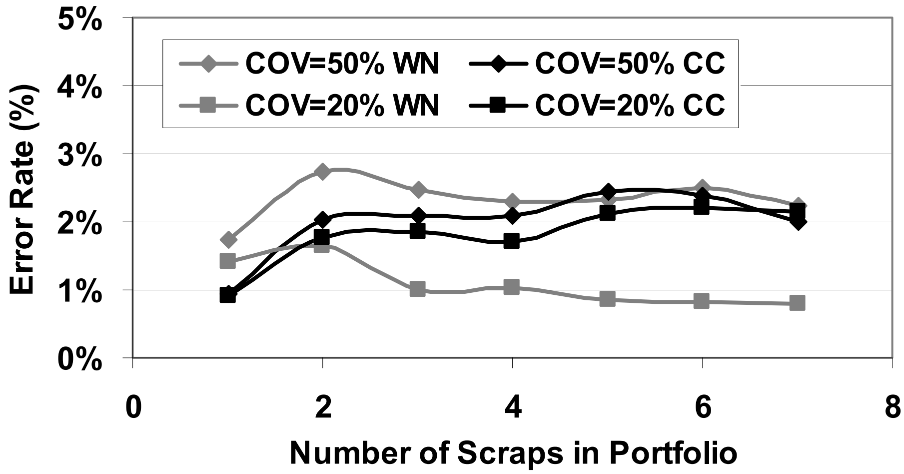

Figure 30. Error rate with increasing scrap diversification (COV=20% or 50%, SD=2, WN=70%)

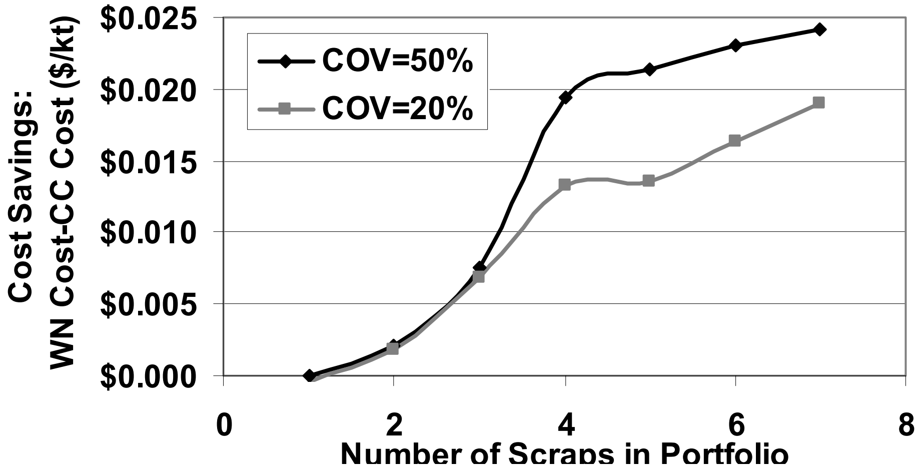

Figure 31. Cost savings of chance-constrained method over conventional window narrowing method with increasing scrap portfolio diversification (COV=20% or 50%, SD=2, WN=70%)

Figure 32 explores the diversification of the chance-constrained method in detail for the production of only alloy 6061; as more scrap become available the window narrowing case continues to utilize only Bumper and Body Sheet scrap while the chance-constrained method samples from all available piles. By doing so, the CC model reduces the consumption of the two most highly utilized scraps, Bumper and Body Sheet, thereby, buffering the production from adverse variation in the composition of those scraps (e.g., both being high in a given element). It is important to note that both methods produce low and equivalent expected error rate for these conditions (Figure 33).. However, though the error rates are nearly equivalent, the chance-constrained method selects a portfolio of materials that has a slightly reduced standard deviation in finished alloy composition as can be seen in Figure 34. In doing so, the CC method is able to utilize more scrap (30.23 kilotonne as compared to 22.96 kilotonne), without compromising the likelihood that the finished good composition of any element is out of specification.

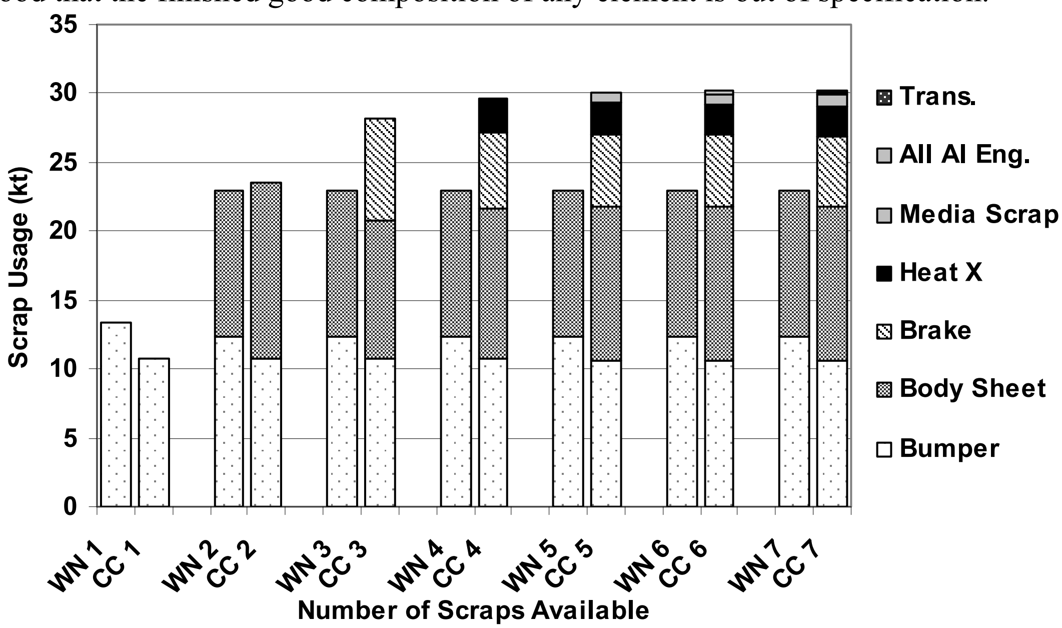

Figure 32. Comparison of detailed scrap consumption of chance-constrained (CC) and window narrowing (WN) methods with increasing scrap type availability. While the number of scraps available for consumption

increases from one to seven, the WN method continues to select only two scrap types. (COV=50%, SD=2, WN=80%)

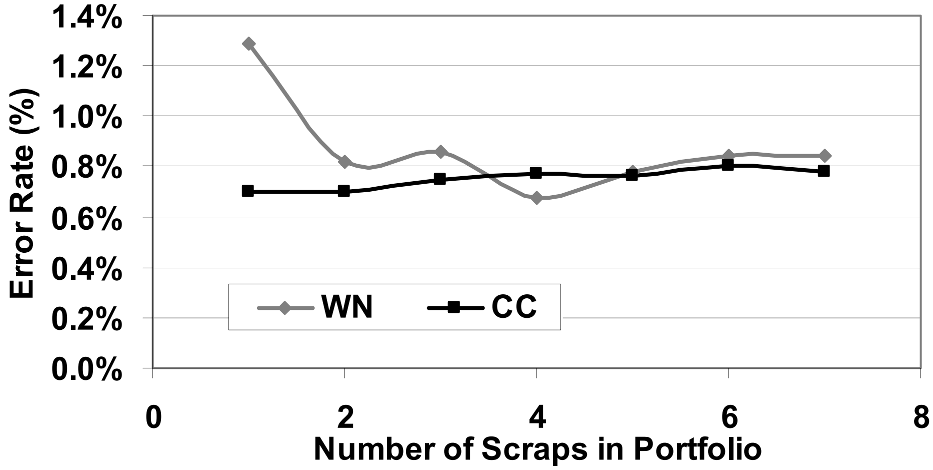

Figure 33. Comparing error rate for chance-constrained and window narrowing method for production of 100 kilotonne of 6061 alloy (COV=50%, 2 SD, WN=80%)

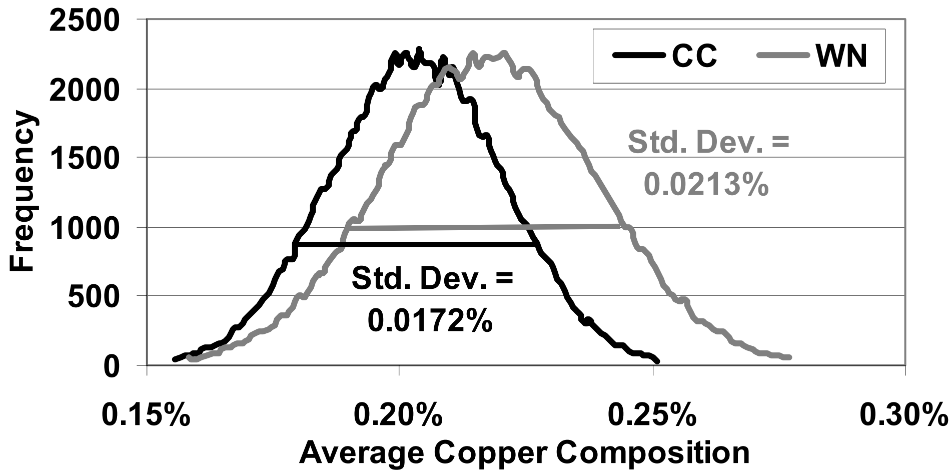

Figure 34. Comparison of stochastic Monte Carlo simulation of average copper composition for alloy 6061 produced according to production plans derived from either chance-constrained (CC) or window narrowing (WN) methods.

The second modification of the Base Case explores the impact of expanding the opportunity for diversification. Because the source for scrap compositional information was limited to only seven types a hypothetical case was constructed where a twin source was "created" with the same mean and standard deviation of the compositional specifications for each of the original seven available scrap types. In net, this leads to a scenario with fourteen available scrap types and, therefore, dramatically increased potential for diversification. Figure 35A compares individual scrap consumption when sampling from both the Base Case and the Diversified Base Case with equivalent scraps summed in the latter. Net scrap consumption increases in the Diversified for all scraps except the Media Scrap. The aggregate effect of this sampling diversification is an increase in consumption of 27.6% shown in Figure 35B.

In the end, both sets of analyses whether limiting or expanding opportunities support the hypothesis that the chance-constrained method leverages diversification to identify batch plans that provide both scrap utilization and economic benefit.

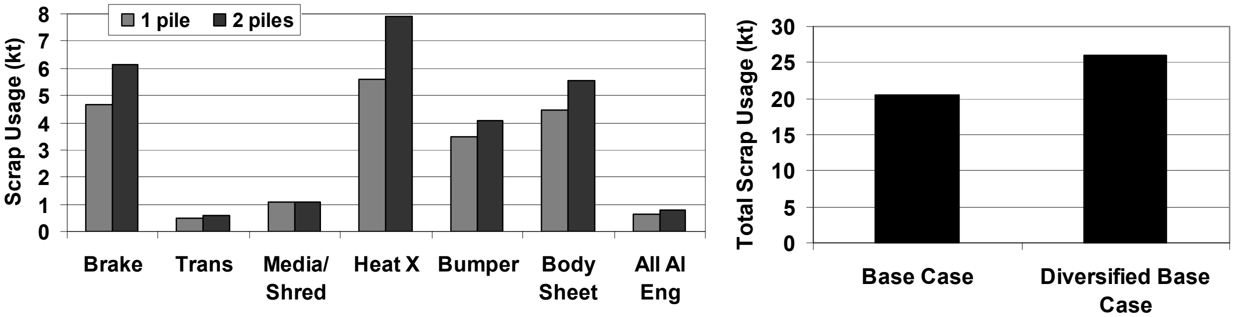

A) B)
Figure 35. A) Comparison of individual scrap consumption when sampling from one or two piles of each scrap using the chance-constrained method B) Total effect of sampling diversification within the chance-constrained method (COV=50%)

Chapter 6. Recycling friendly alloy design

The chance-constrained model presented in chapter 3 was utilized in two hypothetical case studies in order to: 1) test its capability in evaluating the expected recycling performance of specific proposed "recycling friendly" alloys from literature[6] and 2) demonstrate a framework that utilizes model results to *a priori* guide the design of alloys to facilitate increased recycling.

6.1 Case study set-up: common data and assumptions

For both studies, scrap types were selected to be representative of end use shipments of aluminum products in the United States and Canada. The largest categories of applications for aluminum are 1) transportation (which includes both automotive and aerospace), 2) containers and packaging, and 3) construction and building materials (Figure 36). These three categories were used to define the scrap sets used in both of the case studies and detailed in Table XVI, with specific scraps being selected from publicly available compositional data. The transportation scrap set is biased toward automotive related streams due to minimal scope of current aerospace recycling and includes mixed automotive castings, high copper car radiators, segregated alloy 6061 extrusions, and automotive shredder residue (ASR), termed Zorba on the scrap market and often having large amounts of impurities. The container and packaging set includes used beverage cans (UBC), thick foil scrap (Foil), thin foil scrap (Alumifoil), and a segregated mix of alloy 1100 and 3003. Building and construction scraps include mixed aluminum wires and cables, segregated alloy 5052 clippings, clean end-of-life building siding, and segregated alloy 6063 architectural extrusions.

Finally, in addition to the three industry-specific scrap sets, a General Set of scrap was defined to be representative of the overall flow of scrap in North America. Of the old scrap consumed in 2005, UBC, castings, shredded automotives, mixed wrought scraps, and extrusions made up the majority (v. Figure 36). The General scrap set, therefore, was based on a selection of scraps from the three industry specific scrap sets that closely matched this portfolio. These include: UBC, mixed automotive castings, radiators, wire & cable scrap, and mixed turnings.

Compositional data for all scraps were estimated from EU standards[9, 92, 93] which are listed in the Appendix. The EU standards list maximum compositional specifications under which certain scrap types must fall; mean scrap compositions were estimated to be approximately 75% of these values and are listed in Table XVI along with their corresponding EU standard number. Non-proprietary information on the characteristics of available scrap streams is scarce, partly because of confidentiality and partly because of variability across the industry. The authors worked with several industry experts to devise the dataset used herein without compromising firm specific information. While the compositional specifics of any of the scrap sets modeled herein may not match with those used by or available to a specific firm, they do represent the diversity of scrap sources available. The author believes that this diversity drives the challenge of identifying broadly scrap-friendly alloy improvements.

The elemental and gross yields used subsequently are given in Table XVII and Table XVIII; these were estimated from EU standards as well as input from industry experts. Elements such as silicon and iron have yields higher than one as they will generally increase due to melt contact with processing equipment and refractory. Other elements, including aluminum, will have melt loss due to dross/oxide formation, spills, etc. Notably, only very limited data is collected currently on the specifics of yield loss. Although the figures used herein represent the collective input of three independent industry sources, detailed inquiry should be undertaken to more accurately quantify these values. Nevertheless, incorporating these values in the case analyses makes it possible to characterize the magnitude of effects attributable to yield related change.

Prices used in both case studies for primary aluminum and alloying elements were taken from USGS 2005 averages[4, 22] as shown in the Appendix. Scrap prices were estimated from various on-line scrap dealers[94] by averaging costs of similarly named and described scrap categories. Prices vary greatly by day and location, therefore cost results should be used for relative comparison only. All raw materials were assumed to be unlimited in availability in order to avoid the potential effects of limited raw materials supplies. The model framework presented herein can be used for cases of constrained scrap supply with no modification.

Within the chance-constrained formulation, the scrap raw materials were modeled with a coefficient of variation (standard deviation divided by the mean) of 50% on composition for all elements for the base case. Literature on the variability of aluminum secondary materials[27] cites coefficients of variation on elemental means ranging from 55% to as high as 3100%. Sensitivities around this number were also explored. Compositions were assumed to be perfectly uncorrelated. As additional data becomes available, the implications of this assumption should be explored in future work. Collectively this set of data will be referred to as the Baseline in subsequent discussions.

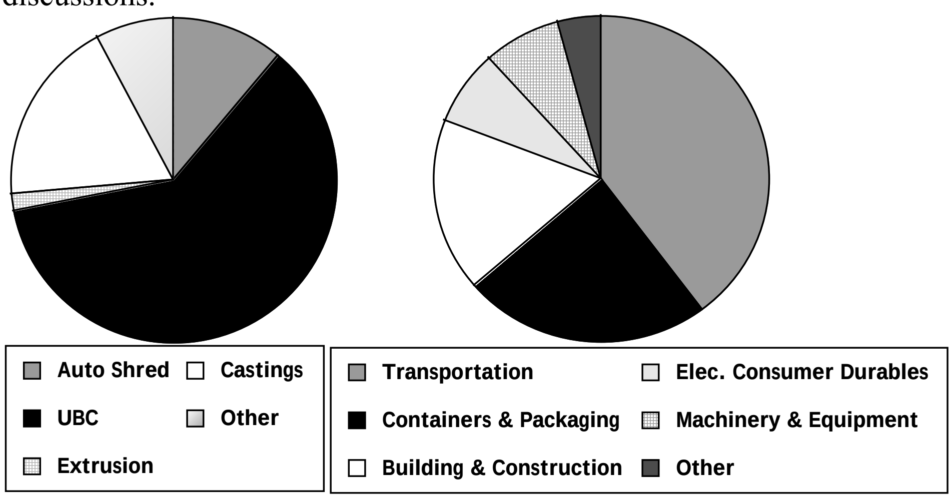

Figure 36. Percentages of old scrap consumed (total 1,154,000 metric tones) and distribution of end use shipments (total 9,699,000 metric tones) by category in the United States and Canada in 2005[95]

Table XVI. Average compositions for scrap sets and prices

Scrap Descriptions		Price	Si	Mg	Fe	Cu	Mn	Zn	EN #
General	Mixed Turnings	$0.74	0.0675	0.0023	0.0075	0.0263	0.0038	0.0113	13
Construction	Wire & Cable	$1.15	0.0019	0.0045	0.0030	0.0004	0.0004	0.0005	3-1
	5052 Clippings	$1.04	0.0023	0.0188	0.0038	0.0008	0.0045	0.0019	5-5

	Clean Siding	$1.02	0.0045	0.0098	0.0045	0.0015	0.0098	0.0015	5-3
	6063 Arch Ext	$1.00	0.0045	0.0038	0.0038	0.0015	0.0011	0	5-6
Automotive	Auto Castings	$0.77	0.1013	0.0023	0.0083	0.0263	0.0038	0.0090	7
	Cu-Al Radiator	$1.60	0	0	0.0053	0.3000	0	0	11
	Zorba	$0.82	0.0675	0.0038	0.0083	0.0263	0.0038	0.0090	8
	6061 Alum Ext	$0.98	0.0045	0.0038	0.0038	0.0015	0.0011	0.0019	5-6
Packaging	UBC	$0.88	0.0023	0.0098	0.0038	0.0015	0.0083	0.0004	10
	Foil	$0.95	0.0075	0.0045	0.0060	0.0060	0.0038	0.0038	6-2
	Alumifoil	$0.75	0.0075	0.0015	0.0075	0.0188	0.0030	0.0060	15
	Seg 1100/3003	$1.02	0.0071	0.0008	0.0015	0.0004	0.0004	0	4

Table XVII. Gross melt yield for scraps and primary

Raw Material Descriptions		Yield	Std.Dev.
General Scrap	Mixed Turnings	0.900	±0.00900
Construction Scraps	Wire & Cable	0.975	±0.00975
	5052 Clippings	0.975	±0.00975
	Clean Siding	0.975	±0.00975
	6063 Arch Ext	0.975	±0.00975
Automotive Scraps	Auto Castings	0.940	±0.00940
	Cu-Al Radiator	0.975	±0.00975
	Zorba	0.940	±0.00940
	6061 Alum Ext	0.975	±0.00975
Packaging Scraps	UBC	0.975	±0.00975
	Foil	0.975	±0.00975
	Alumifoil	0.900	±0.00900
	Seg 1100/3003	0.975	±0.00975
Primary & Alloying Elements	P1020	0.990	±0.00495
	Silicon	1	±0.00500
	Manganese	0.990	±0.00495
	Iron	1	±0.00500
	Copper	0.990	±0.00495
	Zinc	0.990	±0.00495
	Magnesium	0.990	±0.00495

Table XVIII. Elemental yield

	Yield	Std. Dev.
Silicon	1.01	±0.0101
Magnesium	0.98	±0.0098
Iron	1.01	±0.0101
Copper	0.99	±0.0099
Manganese	0.99	±0.0099
Zinc	0.985	±0.0098

6.1.1 Case one specific data and assumptions

The first case study involved the evaluation of three different alloy sets (R, M1, and M2), each comprising six predominant end-market aluminum alloys; one selected from each major alloy series. Set R are "recycling friendly" alloys suggested by Das[27] while Set M1 and Set M2 are the currently available alloys that most closely match the compositions of Set R. Each set was evaluated for potential scrap use within the model under conditions that would reflect production of 100 kilotonnes of each alloy for a total production of 600 kilotonnes. Because the evaluation was carried out under conditions with no limitation on the availability of raw materials, the subsequent results are independent of production scale. The scale of 100 kilotons was selected simply for statistical convenience. Maximum and minimum compositional constraints for Sets M1 & M2 are based on international industry specifications and do not reflect production targets of any specific firm; they are based on guidelines set by the Aluminum Association. These compositions are listed in Table XIX.

6.1.2 Case two specific data and assumptions

For the second case study, two alloys, 6063 and 3004, were evaluated to identify *a priori* specification modifications that would improve potential scrap consumption. Compositional shadow prices were used to target which specifications should be modified (i.e., loosened or tightened); the subsequent impact on potential total scrap consumption was evaluated *post facto* in the model.

Table XIX. Maximum and minimum compositional specifications for finished alloys in weight fraction[96]

	Si		Mg		Fe		Cu		Mn		Zn	
Set R	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min
A(2XXX)	0.007	0	0.006	0	0.07	0.055	0.004	0.002	0.007	0	0.005	0
B(3XXX)	0.007	0	0.006	0	0.004	0	0.015	0.01	0.015	0.008	0.005	0
C(4XXX)	0.14	0.1	0.01	0	0.015	0.005	0.003	0	0.015	0.008	0.005	0
D(5XXX)	0.007	0	0.006	0	0.003	0	0.0035	0.0005	0.03	0.02	0.005	0
E(6XXX)	0.01	0.003	0.006	0	0.003	0	0.003	0	0.01	0.004	0.005	0
F(7XXX)	0.005	0	0.006	0	0.012	0.005	0.003	0	0.028	0.02	0.06	0.04
Set M1												
2014	0.012	0.005	0.008	0.002	0.007	0	0.05	0.039	0.012	0.004	0.0025	0
3005	0.006	0	0.006	0.002	0.007	0	0.003	0	0.015	0.01	0.0025	0
4045	0.11	0.09	0.0005	0	0.008	0	0.003	0	0.0005	0	0.001	0
5454	0.002	0	0.03	0.024	0.002	0	0.001	0	0.01	0.005	0.0025	0
6063	0.006	0.002	0.009	0.0045	0.0035	0	0.001	0	0.001	0	0.001	0
7005	0.0035	0	0.018	0.01	0.004	0	0.001	0	0.007	0.002	0.05	0.04

Set M2											
2219	0.002	0	0.0002	0	0.003	0	0.068	0.058	0.004	0.002	0.001	0
3004	0.003	0	0.013	0.008	0.007	0	0.0025	0	0.015	0.01	0.0025	0
4032	0.135	0.11	0.013	0.008	0.01	0	0.013	0.005	0.005	0	0.0025	0
5052	0.0025	0	0.028	0.022	0.004	0	0.001	0	0.001	0	0.001	0
6061	0.008	0.004	0.012	0.008	0.007	0	0.004	0.0015	0.0015	0	0.0025	0
7075	0.004	0	0.029	0.021	0.005	0	0.02	0.012	0.003	0	0.061	0.051

6.2 Case one – evaluating alloy ability to utilize scrap: comparison of "recycling friendly" alloys to representative market alloys

The first case study involved the evaluation of the potential for scrap use (i.e., recycling-friendliness) for three different alloy sets (R, M1, and M2), each comprised of six predominant end-market aluminum alloys; one selected from each major alloy series. Set R are "recycling friendly" alloys suggested by Das[97] while Set M1 and Set M2 are the currently available alloys that most closely match the compositions of Set R.

Table XX compares the results for the Baseline conditions described above for each of three alloys sets. Results show a total improvement in potential scrap consumption of 67.9% and 65.6% respectively, with an associated decrease in primary purchased, for the scrap friendly alloy set (R) over the currently used market alloy Sets M1 & M2. These base cases were evaluated using Monte Carlo simulations to have comparably low expected error rates of 0.18%, 0.21%, and 0.25% respectively. The scrap friendly alloy set shows an associated decrease in modeled production cost of 13.4% and 13.6% over the other alloy sets.

Figure 37 , however, provides some indication of the challenge of creating a broadly recycling-friendly alloy. Specifically, Figure 37 shows that the batch plans for the recycling friendly alloys outperform their market counterparts for *most*, but not all, of the alloy series investigated in terms of potential scrap use. Most notably, Alloy C(4XXX) outperforms 4032 by 10X (917%) and 4045 by almost 9X (775%); Alloy B(3XXX) outperforms 3005 by 47% and 3004 by 26%. However, the recycling friendly alloys do not outperform all of the market alloys; Alloy D(5XXX) has about the same potential scrap consumption as alloy 5052 while alloy F(7XXX) is outperformed by both its comparative market counterparts, 7005 and 7075, respectively. Later sections will explore in detail the source of this underperformance.

Table XX. Baseline results showing comparison of scrap friendly alloys (Set R) with current alloys (Sets M1&M2) (CC Method, 50% COV, Total Production = 600 kt)

	Alloy Set R	Alloy Set M1	% Δ R-M1	Alloy Set M2	% Δ R-M2
Scrap Use (kt)	272.5	162.3	+67.9%	164.5	+65.6%
Production Cost (M$/kt)	$1.783	$2.060	-13.4%	$2.064	-13.6%
Error Rate	0.18%	0.21%		0.25%	

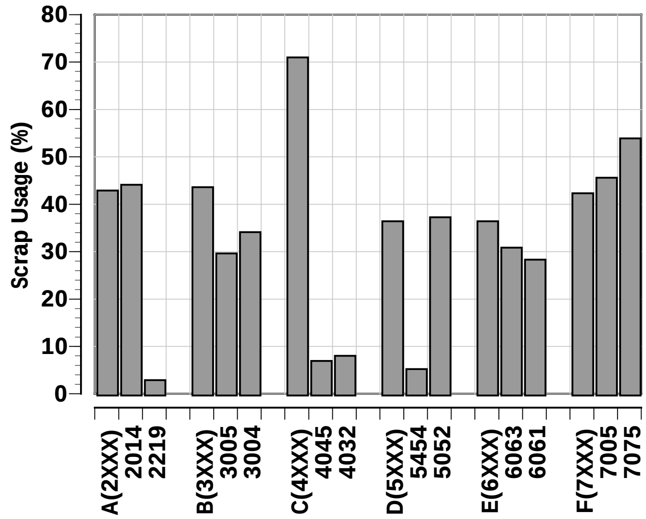

Figure 37. General scrap set consumption comparison for individual alloy sets, organized by series (CC Method, COV=50%, α=99.99%).

6.2.1 Evaluating recycling performance: sensitivity analysis

The recycling-friendliness of an alloy is not only dependent on the compositional variability of the available scraps, but also the make-up of the scrap portfolio itself. Many producers, depending on the size and location of their facility, have access to scrap portfolios that will be heavy in scraps from one industry over another. To systematically explore these implications, the three candidate alloy sets were tested against three scrap portfolios, created based on the three major aluminum markets shown in Figure 36: 1) transportation, 2) containers and packaging, and 3) building and construction. The scrap portfolios, their average compositions, and modeled prices are listed in Table XVI. In this section, the original scrap portfolio will be referred to as General, reflecting its composition from all of the other scrap classes.

As Table XXI shows, for each alloy set, the amount of scrap used is highly dependent on the available scrap "portfolio". For example, many automotive scraps go through a shredding process and therefore have a much higher accumulation of iron than other scraps[31]; one would expect the maximum potential usage of these scraps per application to be lower than other types which is confirmed by the results. Both packaging and construction scraps have extremely low accumulation of undesirable elements and can therefore be highly utilized by all the alloy sets.

Looking at the scrap consumption broken down by specific alloy (Figure 38A), one can see an even greater range of usage differences. Alloy 7005 and alloy 2014 accommodate large amounts of the construction heavy scrap in their production portfolios due to being compositionally close to clean unpainted siding. Packaging scraps, most notably used aluminum foil, have extremely low magnesium and manganese content and therefore can be utilized by alloy 4045 while that alloy can normally accommodate little to no secondary materials of other scrap types (Figure 38B).

Table XXI. Total scrap material usage for the varying scrap portfolios by alloy set.

Scrap Sets	Alloy Sets		
	Alloy Set R	Alloy Set M1	Alloy Set M2
Construction	348.04	214.90	201.83
Packaging	333.61	180.55	115.54
General	272.52	162.33	164.52
Automotive	200.91	115.31	87.73

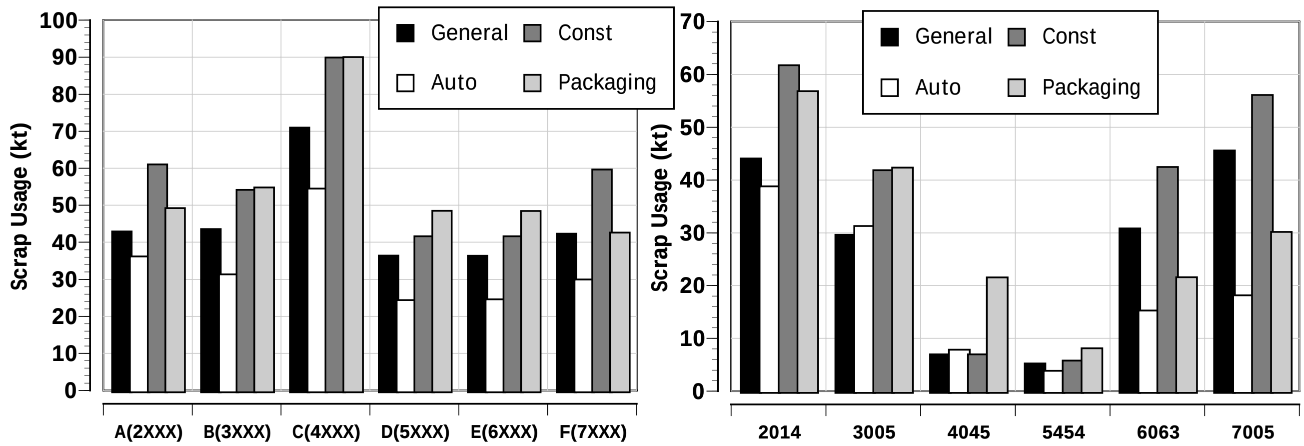

A) B)

Figure 38. Recycling friendly alloy set R (A) and market alloy (B) set M1 scrap use comparison for each different scrap "scenario"

In the end, the most important observation from this analysis is that the available scrap portfolio differentially impacts the quantity of scrap use even for analogous alloys (in the same series). As a result, the available scrap portfolio can change the relative performance of any given alloy, including ones from the recycling-friendly set, compared to its analogs in terms of potential scrap use.

As an example, for both the Construction and General scrap portfolios, market alloy 7075 has the potential to use the most scrap of the 7-series alloys considered, while for the Automotive and Packaging scenarios, the "recycling friendly" alloy F(7XXX) can use the most scrap (Figure 39.). Recycling friendly alloy 7XXX is more compositionally restrictive for Mg compared to market alloys 7005 and 7075, while market alloy 7075 is compositionally more restrictive for Fe and Mn compared to the recycling friendly alloy (see Table XIX). For the Automotive-heavy scrap portfolio case, one can see from Table XXII that the recycling friendly alloy is able to utilize more Zorba and 6061 aluminum extrusion than market alloy 7075; these scraps have fairly high iron and magnesium content. For the Construction-heavy scrap portfolio case, one can see that market alloy 7075 utilizes more wire and cable scrap compared to the other two alloys. Wire and cable scrap are desirable due to their low silicon content (cf. Table XVI), however, due to their high usage the magnesium becomes the constraining element. This allows alloy 7075 to accommodate more secondary materials in its portfolio overall.

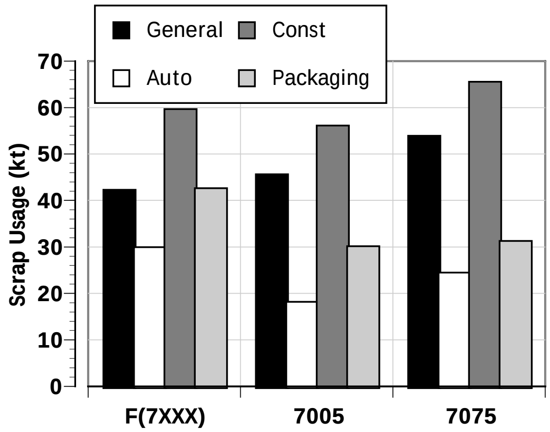

Figure 39. Comparison of scrap used in the production portfolio of the 7XXX series alloys for each of the scrap scenarios

Table XXII. Scrap usage (in kilotons) for 7XXX series alloys

Automotive Set	7XXX	7005	7075	Construction Set	7XXX	7005	7075
Auto Castings	0.1	0.1	0.1	Wire & Cable	25.2	26.8	37.8
Cu-Al Radiator	0.2	0	0.7	5052 Clippings	1.7	8.5	4.7
Zorba	0.1	0.1	0.1	Clean Siding	5.8	7.3	7.7
6061 Alum Ext	29.4	18.0	23.5	6063 Arch Ext	26.6	13.5	15.3

Similar results can be shown for the 5XXX alloys. D(5XXX) outperforms its market counterparts for the Automotive- and Packaging-heavy scrap portfolios, while for the General- and Construction-heavy portfolios, alloy 5052 can accommodate the most secondary materials in its production portfolio (Figure 40). In this case, for the Packaging-heavy portfolio, Seg 1100/3003 is desirable because it has a low iron content compared to the other scraps (cf. Table XVI), however, silicon becomes a constraining element in the production portfolio and therefore the recycling friendly alloy can consume more due to its less restrictive silicon specification. On the other hand, for the General portfolio, alloy 5052 has the potential to use more Wire and Cable scrap (Table XXIII) in its production portfolio because it is less compositionally restrictive in magnesium and manganese when compared to D(5XXX).

Ultimately, it is apparent that the recycling performance of a specific alloy can be strongly dependent on the operational context in which it is applied. This clearly confounds the process of designing alloys to improve the ability to accommodate scrap. Mathematical programming, like the chance-constrained model used herein, provides rapid, quantitative insight without the need for expensive and time-consuming experimentation. Nevertheless, given the range of potential operational settings (i.e. available scrap types, compositional variability, raw material prices, etc.) and the continuum of possible compositional modifications, *post facto* evaluation, even through a model, may not make effective alloy design tractable. Instead, the alloy designer needs insight into what modifications can provide the most benefit. The next section, through the use of a case study, explores the use and value of that information, referred to as a shadow price, to direct the development of more recyclable alloys.

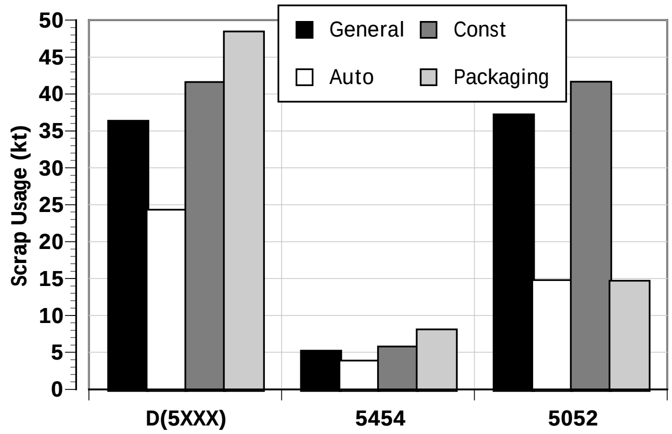

Figure 40. Comparison of scrap used in the production portfolio of the 5XXX series alloys for each of the scrap scenarios

Table XXIII. Scrap usage (in kilotons) for 5XXX series alloys

Packaging Set	5XXX	5454	5052	General Set	5XXX	5454	5052
UBC	16.6	2.3	2.9	UBC	14.5	2.3	2.9
Foil	4.4	0.2	3.1	Auto castings	1.2	0	0.0
Alumifoil	2.3	0	1.0	Cu-Al Radiator	0.2	0	0.1
Seg 1100/3003	25.0	5.7	7.7	Wire and Cable	18.6	2.9	34.0
Mixed turnings	0.1	0	0	Mixed turnings	1.9	0	0.3

6.3 Designing for scrap consumption: using compositional shadow prices

While it is clear that modifying alloy specifications could facilitate recycling, identifying the most effective modifications is not obvious. Unfortunately, the trivial solution of broadening all compositional specifications is sure to alter alloy materials properties. Instead, the alloy designer must selectively alter specifications, relaxing some, while tightening others, all without compromising performance specifications. Case Two explores the use of model outputs to guide the *a priori* modification of alloy specifications to improve potential for scrap use. Simple examples from Case One can serve to illustrate the challenge of realizing an effective design in the absence of such guidance.

Consider the relative performance of the 4XXX and 5XXX alloys as described in the previous section (i.e., Alloy C had the highest potential for scrap use of the 4XXX alloys while 5052 had the highest potential for scrap use of the 5XXX alloys (v. Figure 37) and their specifications as shown in Figure 41. Examining Figure 41, Alloy C, compared to its market equivalent 4XXX alloys, has a broader specification range for all elements with the exception of Cu and a higher maximum specification range for all elements except for Cu and Mg. This is in stark contrast to alloy 5052 which, compared to the other 5XXX alloys considered, has the highest potential for scrap use but neither the broadest nor highest maximum specification with the exception of Fe. These observations reinforce the key challenge of identifying the most effective modifications.

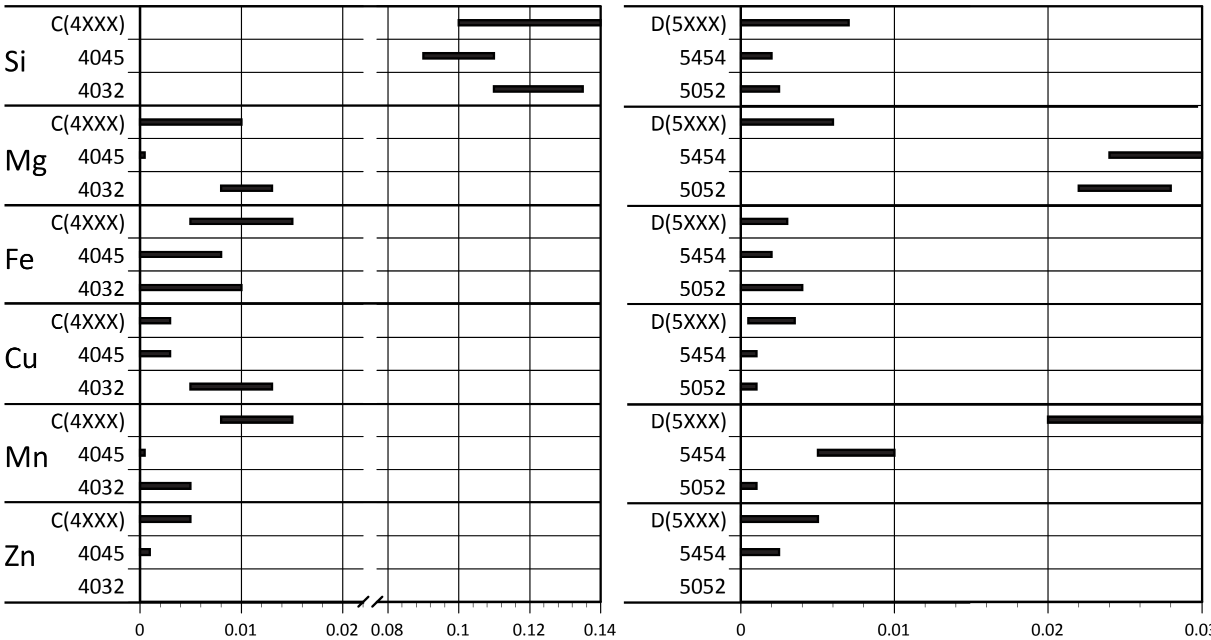

Figure 41. Schematic representation of alloy compositional specification windows for the 4XXX and 5 XXX series alloys of all three sets

With current specifications comprehending nearly two dozen elements, it is not practical for even fast models to iterate through every possible combination of specification modifications. Fortunately, the type of model presented in this paper as well as those used broadly in industry for daily batch planning and many other models discussed in the literature when executed in modern solution engines also generate a set of information that can direct the design process. As described above (see section 2.1), shadow prices identify those specifications with both the most and least impact on potential scrap use. Specifically, the magnitude and sign of the shadow prices on composition indicates how the production cost would change if the compositional specifications were tightened or loosened. Using this information, a prospective framework can be implemented to systematically and efficiently target alloy specifications for development that provide the most significant improvements in potential for increased scrap reuse and reduced production cost. Armed with such a framework, the design process should become more efficient.

To explore the value of shadow prices within the framework, modifications to two specific alloys were explored. These alloys meet the compositional specification of 6063 and 3004 as shown in Table XIX. To allow for alloy modifications that involve both tightening and loosening specifications, but still remain within the AA specification, the base case assessment will be carried out with the most constraining specification (by coincidence Mg in both cases) set to 10% of the width listed in Table XIX. This width reduction provides opportunity to explore constraint loosening and was carried out by changing only the maximum specification (the more binding specification[4]). As an example, for the analysis of 6063 the base line maximum

[4] It is no surprise that, for a model either minimizing raw material cost or maximizing scrap use, more of the binding compositional constraints are maximums (e.g., 30 of 52 for the recycling friendly alloy set); the amount of contaminants in a scrap usually determines how much dilution with primary aluminum is required and is therefore the major limiting factor. The shadow prices

specifications were set to the values as listed in Table XIX except for Mg which was set to ε^{min} + 10%(ε^{max}-ε^{min})=0.49 wt%. These alloys will be referred to in the subsequent discussions as 6063^{sr} for specification restricted. By establishing these alloys as a baseline for comparison, all of the modifications explored subsequently can be made without causing the resultant alloy to fall outside of AA specifications.

Table XXIV shows the shadow prices for the maximum specification constraints for original alloys 6063 and 3004 and the specification restricted 6063 and 3004. The largest shadow prices indicate those with the largest potential effect on production cost and, therefore, potential scrap use. For alloy 6063^{sr}, the largest compositional shadow price is that associated with the maximum specification for magnesium at 807.8. Based on this information, one would therefore propose a "recycling friendly" alloy with a relaxed constraint on magnesium. Such an alloy would be the same as market alloy 6063^{sr} but with the compositional constraints on maximum magnesium composition modified slightly. Specifically, Figure 42A shows the implications of relaxing this specification from a value of 0.49 wt% (i.e., 10% of the full listed AA specification window) to 0.90 wt.% (95% of full window). Expanding this specification over this range results in a 16-fold increase in potential scrap utilization. Increasing the specification beyond the AA specification could result in even higher scrap utilization; however, these changes may be less feasible due to processing and property restrictions.

It is intuitive that relaxing compositional constraints make it possible for an alloy to accommodate more scrap in its optimal production portfolio. However, to maintain properties, it is likely that this relaxation will need to be accommodated through the tightening of other specifications. Compositional shadow prices also indicate which of the constraints could be tightened with the least effect of potential scrap utilization. In this case, the lowest shadow prices indicate the constraints with the least effect on cost and, therefore, scrap usage. Figure 42B is used to explore the implications of complementary loosening of constraints. Within Figure 42B, each curve represents the impact of constraint tightening overlaid on the impact of Mg constraint loosening from the plot above. As with Figure 42A, the constraint tightening here is carried out by only modifying the maximum specification for the given alloying element, with a percent tightening of X% translating to a specification value of ε^{max} - X%(ε^{max}-ε^{min}). As an example, the height of the solid black curve at 60% on the x-axis represents the change in potential scrap usage associated with a 60% tightening of the Si and Zn constraints on top of the 60% relaxation of the Mg constraint captured in Figure 42A above. Similarly, the height of the gray solid line at 40% represents the change in potential scrap use associated with a 40% tightening of the Fe constraint in addition to a 40% relaxation of the Mg constraint. Obviously, specific alloy modifications do not have to mirror one another and could take on any given value. This presentation strategy was chosen to densely represent the range of possible alloy modifications that could improve potential scrap usage when such modifications are chosen carefully.

Looking back at the compositional shadow prices in Table XXIV, one would expect that the specifications on Si, Zn, and Fe for 6063^{sr} could be tightened with little negative impact on scrap

---

on magnesium, copper, and manganese are typically higher because these three alloying elements are the most expensive (>$2,000/tonne within the cases studied herein) and therefore have the highest impact on the production cost.

usage. Figure 42B shows that this hypothesis is correct for specification tightening of 0 to 40% for Fe and 0 to 60% for both Si and Zn. Tightening of the Fe specification beyond 40% begins to compromise the ability to use scrap. Similarly, tightening beyond 60% for Si and Zn begins to have a significant impact on scrap utilization.

Table XXIV. Compositional shadow prices for maximum specifications for 6XXX and 3XXX series market alloys

	Example 1		Example 2	
Element	Alloy 6063	Alloy 6063^{sr}	Alloy 3004	Alloy 3004^{sr}
Mg	74.3	807.8	83.8	222.7
Si	0.9	0.96	7.7	7.3
Fe	2.0	2.0	2.0	2.0
Cu	20.5	6.4	0.8	0.6
Mn	29.9	0	0.0	0.0
Zn	1.2	1.3	1.2	1.2

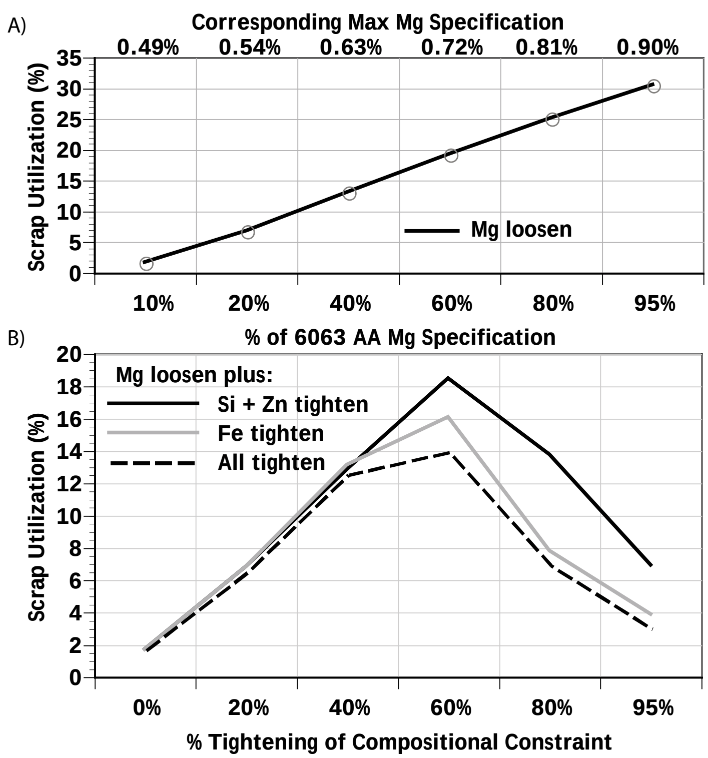

Figure 42. A) Improvement in potential scrap consumption from loosening Mg constraints on alloy 6063 from 10% to 95% of as-listed specification width (a given percentage on the x-axis translates into a Mg specification through the relationship $\varepsilon^{max} = \varepsilon^{min} + x\%(\varepsilon^{max}-\varepsilon^{min})$, B) Effect on scrap consumption of overlaying a tightening of other constraints in conjunction with loosening the Mg constraint corresponding to the same value on the x-axis of the graph above (Figure 42A). The maximum specification of a given element on a specific plot follows, $\varepsilon^{max} = \varepsilon^{max} - x\%(\varepsilon^{max}-\varepsilon^{min})$. The line marked A represents the effect of tightening all other specifications (Si, Fe, Cu, Mn, and Zn) in conjunction with a loosening of the Mg constraint of an equal percentage amount.

This analysis suggests that it should be possible to use shadow price information to identify effective alloy modification targets. The existence of the inflection points in each of the constraint tightening curves (Figure 42B) reflects the fact that all specifications, even ones with very low shadow prices, can eventually become a limitation to scrap use. To avoid exploring alloy modifications that involve constraints tightened to the extent that scrap use is compromised, it will be necessary to employ an iterative procedure of using shadow prices and model evaluation (as well as technical performance evaluation). As diagrammed in Figure 43, this procedure would involve executing the model described here or a similar batch planning model (c) for an alloy of interest (a) and a broad set of available secondary materials (b). This model execution will generate both an optimal production plan and shadow prices for the compositional specifications. Those shadow prices can be used to identify specifications that are candidates to modify to increase scrap use (d) and to compensate for those modifications through tightening without compromising scrap use (e). Promising specification modifications (f) can be tested within the model to understand the extent of scrap use improvement (g). This step is particularly important for any proposed compensatory constraint tightening to ensure that the amount of tightening does not reach a point that undermines other gains (h). New alloy specifications would eventually need to be tested for technical performance (h), a procedure that would lie outside of this method. Ideally new alloys with satisfactory performance could iterate again through this process to see if further scrap usage improvements were possible (j-yes). This iterative process could identify a set of promising alloy candidates, the alloy producer would select from the set that met performance requirements (j-no).

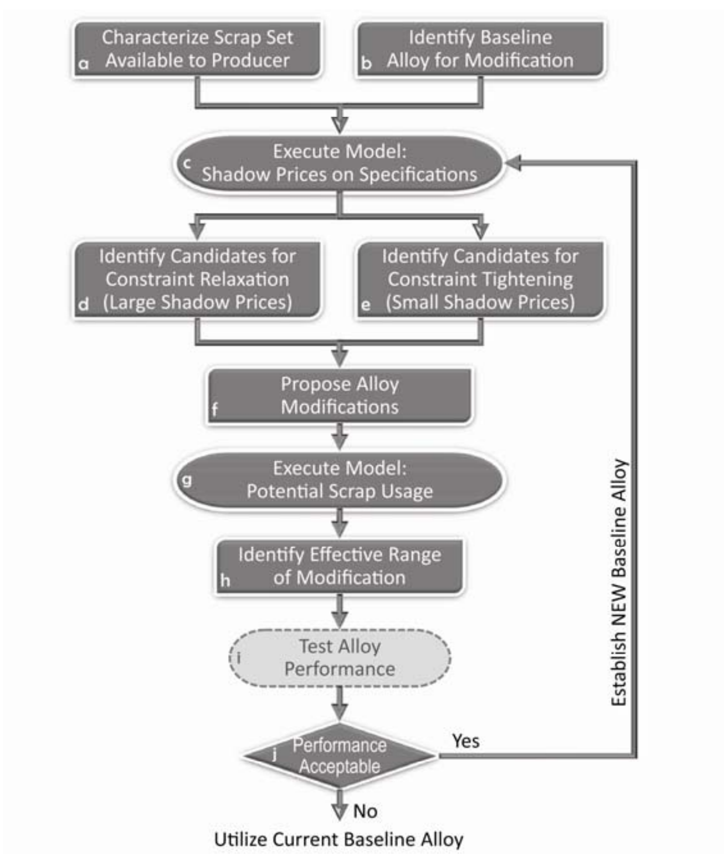

Figure 43. Framework to identify alloy modification targets and possible recycling-friendly alloys

A second example will be used to demonstrate such an approach more directly. One of the proposed "recycling friendly" alloys that did not increase scrap usage for every scrap type was the 3XXX series alloys; this is a case where the market alloy could accommodate more scrap than the proposed recycling friendly alloy in some cases. By looking at the compositional shadow prices on the maximum specifications for the 3XXX series market alloys (cf. Table XXIV), one can see that magnesium again has the highest shadow prices for alloy 3004^{sr} with a value of 222.7, 30 times larger than the next largest value for Si. One would therefore propose a "recycling friendly" alloy with relaxed constraints on magnesium.

Figure 44A shows the results of executing the model for modifications of the alloy 3004^{sr} that involve relaxing the maximum specification on Mg from the baseline value of 0.85 wt% to 1.3 wt% (a expansion / relaxation of the constraint window by 50%). Given the large shadow price, it is not surprising to find that expanding the magnesium specification across this range results in a nearly 40-fold increase in potential scrap utilization. For illustration purposes, we will select an alloy modification half way through this range at 25% relaxation. As a crude way of tracking the metallurgical implications of this and subsequent proposed modifications, it is useful to look at the minimum expected aluminum content within the alloy. By increasing the width of the magnesium specification window by 25%, the minimum aluminum content would drop from 96.15% to 95.95%. It is likely that this change will need to be compensated for by tightening other specifications.

Examining the shadow prices for the modified alloy, labeled $3004^{sr'}$ in Figure 44B, one hypothesizes that tightening the constraint on Mn would have the smallest effect on scrap usage. The plot in Figure 44B shows the results of testing modifications to the Mn specifications for $3004^{sr'}$ and reveals a broad range over which the constraint on Mn can be tightened without compromising scrap usage in the alloy. In fact, the width of the constraint can be reduced to a value only 2% of the original width (98% tightening on the plot) before scrap usage falls below the level of the baseline 3004^{sr} specification. For illustrative purposes, a value of 50% tightening will be selected, resulting in a Mn specification of 1.25 wt% for proposed alloy $3004^{sr''}$. Such a change would lead to a minimum aluminum content of 96.2 wt% in $3004^{sr''}$.

As a final example of iteration, looking at the shadow prices for $3004^{sr''}$ reveals that of the other specifications, Cu has the lowest shadow price and therefore represents a promising candidate for modification. As for manganese, executing the model for Cu constraint modifications on $3004^{sr''}$ confirms this expectation with only small reductions of potential scrap use for window tightening up to nearly 90%. Using this information and following the pattern above, one could propose a final alloy modification 3004sr''' at 50% window tightening for Cu, a point that translates into a maximum specification of 0.13 wt% Cu. At this point, minimum aluminum content would sit at 96.33% and potential scrap usage would have climbed to 18% (from a value of 1% for 30004^{sr}).

These examples are provided to illustrate how a properly constructed batch planning model and the shadow price information that it could generate can be used to inform the alloy design process. These examples also demonstrate that for most production constructs there are compositional constraints with both large (Mg for 6063 and 3004) and small (e.g., Si, Zn, and Fe for 6063 and Mn and Cu for 3004) shadow prices. Any actual alloy design process should also be informed by metallurgical considerations concerning the complementary (or conflicting) effects of various alloying elements.

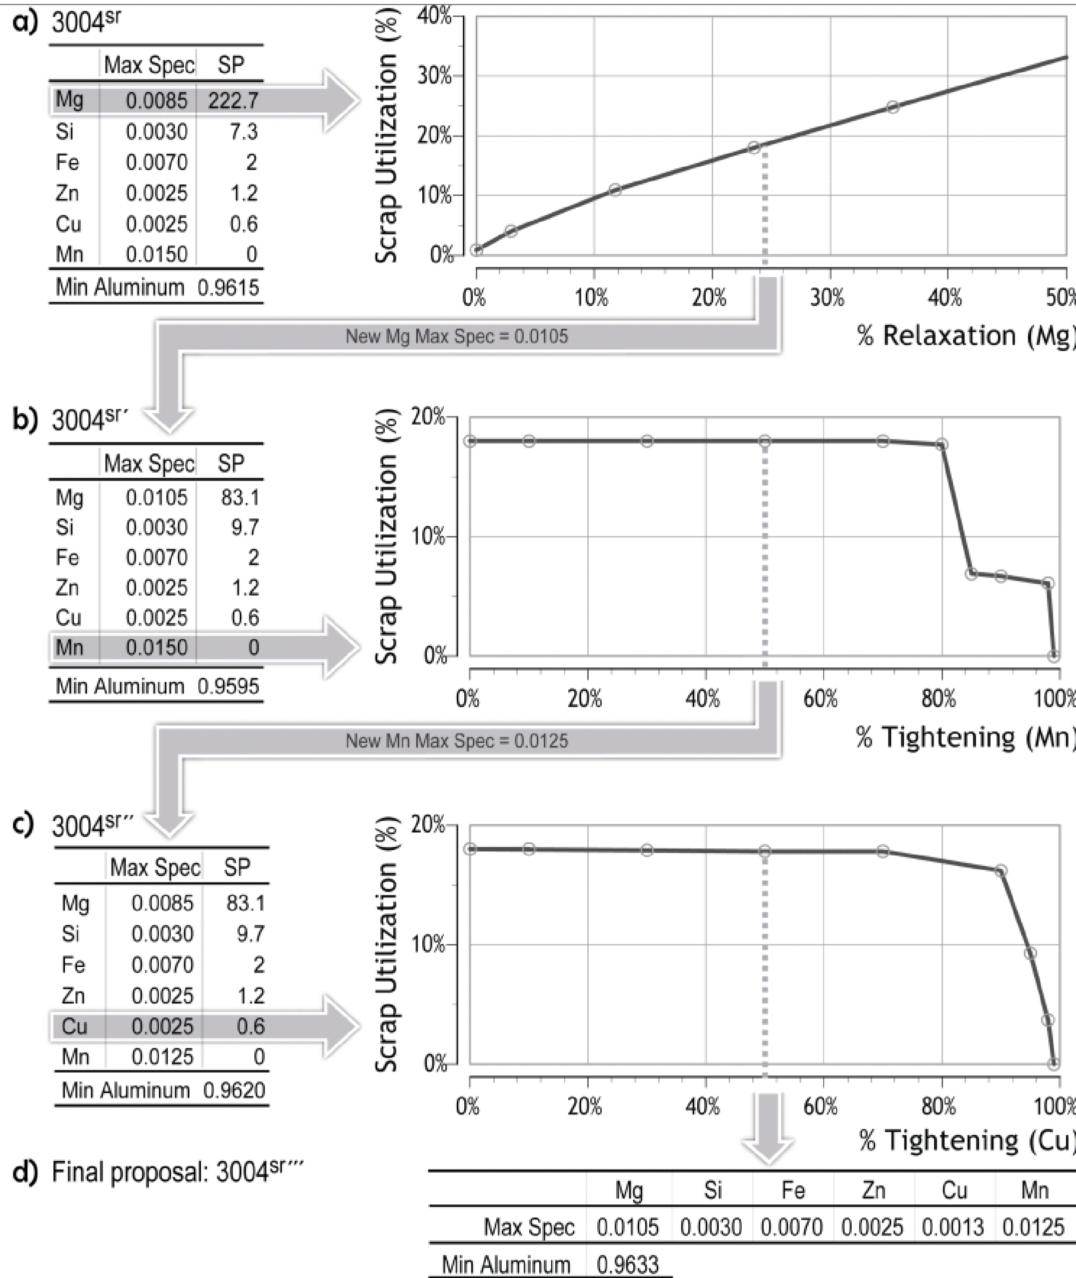

a) 3004^{sr}

	Max Spec	SP
Mg	0.0085	222.7
Si	0.0030	7.3
Fe	0.0070	2
Zn	0.0025	1.2
Cu	0.0025	0.6
Mn	0.0150	0
Min Aluminum	0.9615	

b) $3004^{sr'}$

	Max Spec	SP
Mg	0.0105	83.1
Si	0.0030	9.7
Fe	0.0070	2
Zn	0.0025	1.2
Cu	0.0025	0.6
Mn	0.0150	0
Min Aluminum	0.9595	

c) $3004^{sr''}$

	Max Spec	SP
Mg	0.0085	83.1
Si	0.0030	9.7
Fe	0.0070	2
Zn	0.0025	1.2
Cu	0.0025	0.6
Mn	0.0125	0
Min Aluminum	0.9620	

d) Final proposal: $3004^{sr'''}$

	Mg	Si	Fe	Zn	Cu	Mn
Max Spec	0.0105	0.0030	0.0070	0.0025	0.0013	0.0125
Min Aluminum	0.9633					

Figure 44. A) Improvement in scrap consumption of loosening Si constraints on alloy 3004, B) Effect on scrap consumption of now tightening other constraints in conjunction with loosening the Si constraint

In the end, shadow prices provide valuable information on which elemental specification modifications may provide the largest impacts on scrap use (both positive and negative). The large number of elemental considerations of today's aluminum alloys coupled with the often volatile state of the secondary materials market make alloy design for recycling an increasingly complex and challenging task. The model/tool presented in this paper aids the alloy designer by considering both compositional uncertainty of secondary materials and the impact their use has on an alloy's ability to be produced from scrap. The sensitivity analysis that accompanies the optimal production portfolio solutions can guide designers in terms of targeting which specification constraints could be relaxed or tightened.

Chapter 7. Conclusions

Growing industrial awareness of resource scarcity and environmental impact has highlighted the steadily increasing consumption of metals and materials in production. A key strategy for enabling a shift to more sustainable use of materials will be increased recycling. Although reaching full recycling potential will likely require changes throughout any given materials system, two strategies that may play an important role are 1) efficient blending plans and 2) the redesign of alloys to accommodate more scrap. Realizing such recycling improvements will require effective tools to evaluate an alloy's potential recycling performance; without such tools, both blending and redesign can only proceed by expensive trial-and-error.

This thesis has explored the use of several stochastic programming techniques including recourse, penalty functions, and chance-constrained programming to improve on deterministic linear programming batch planning methods in regards to managing compositional uncertainty inherent in recycled materials. To accommodate this uncertainty currently, most remelters modify their batch planning tools either by a hedge around scrap composition or through modification of target specifications. Both approaches require iterative analysis to select an appropriate hedge and do not dynamically accommodate changes in the portfolio of incoming scrap materials. Both of these shortcomings lead to batch plans with lower potential for scrap use and, therefore, reduced economic and environmental benefit. The chance-constrained stochastic programming method was found to overcome both of these shortcomings and has been demonstrated to provide both economic and scrap use benefit over a wide range of relevant conditions. Compared to the other stochastic methods, it also exhibited optimal performance in terms of lowest computational complexity (for both storage and scaling).

For the first set of case studies exploring this methods performance managing compositional uncertainty (cf. Chapter 5), the CC method suggested a batch plan that would increase scrap use by nearly 30% over the batch plan derived from the conventional WN method without compromising the expected number of batch errors (cf Table XII). Furthermore, the CC method demonstrated increased scrap use (or lower expected error rates) across the entire range of scrap compositional variation that was explored. This analysis also clearly demonstrated how the CC method provides a direct mechanism to control the expected batch error rate (i.e., through the confidence interval parameters α and β), while the traditional, deterministic methods require trial-and-error testing and do not adjust as raw material characteristics vary.

Based on fundamental understanding of the mathematical structure of the chance-constrained method and supported by the Base Case results, the authors hypothesized that increases in scrap use projected by the CC method derive primarily from raw material diversification; by making use of a broader array of scrap materials, the CC method is able to increase scrap use while controlling variation in the finished good. Model based experiments that both limited and expanded the opportunity for diversification supported this hypothesis. An exciting consequence of this outcome is that the CC method can identify economically-beneficial opportunities to make use of lower grades of scrap material (cf Table XIII).

From this limited case analysis, the chance-constrained method shows great promise for allowing secondary producers to mitigate the risk of quality variation in incoming raw materials.

Ultimately, this method provides a platform for firms to explore the use of a broader set of raw materials, to make use of available compositional data, and to identify production strategies that provide competitive advantage.

For the case studies exploring alloy design (cf. Chapter 6), the chance-constrained optimization framework, was demonstrated to be a systematic method to (1) evaluate an alloy's ability to accommodate recycled materials (scrap) in its production portfolio and (2) proactively identify the most effective alloy modification strategies that can drive increased potential scrap use. Additionally, this thesis extended previous work in this space by presenting a schematic algorithm for explicitly incorporating uncertain metal yield into the analyses of alloy design specifically, and recycler operational decisions more broadly.

In the cases presented, the model was shown to be effective at differentiating the potential scrap usage of a set of suggested "recycling friendly" alloys particularly as production context shifted across different scraps sets. The specifications for these alloys as suggested in the literature were generally shown to enable increased scrap usage, although this improvement was not uniform across all alloy series. This improvement in secondary material utilization was shown to be context dependent on both the product the alloy is replacing as well as the scraps that are available for its production.

The optimization framework was also shown to be useful in guiding the initial alloy design process in regards to which compositional specifications should be targeted for modification in order to increase recyclability. With regard to guiding such a design process, this paper extends previous discussions [6] by exploring in detail how sensitivity analysis information actually correlates with potential scrap use performance and how both the sensitivity analysis information and the associated potential scrap use effect changes with individual and coordinated specification modifications. The case analyses suggest that real potential exists for increasing potential scrap use through alloy redesign while remaining within established compositional specifications.

Both the model and cases discussed herein are intended to be schematic in nature. Much work still remains to capture the metallurgical complexity of the recycling process; nevertheless, the results presented show that such a framework holds promise to be a valuable member of the metallurgists toolkit. This is especially important because evaluating all possible specification combinations is not practical for efficient decision making. In addition to being highly efficient, this method makes use of a type of model currently in place for aluminum batch planning and therefore could be implemented without additional capital investment. In the end, that design process is still dependent on traditional and emerging methods to identify alloys that will satisfy demanding physical performance requirements. Nevertheless, efficient design of resource-conscious materials will require analytical tools capable of projecting the impact of design choices on recycling performance.

Chapter 8. Future Work

Revisiting Figure 1 (shown again here as Figure 45) illustrates a variety of other leverage points and stakeholders that one could target in a product's life-cycle in order to increase the recycling potential of the system as a whole. Key issues within recovery & collection, secondary markets, and pre-processing are outlined below.

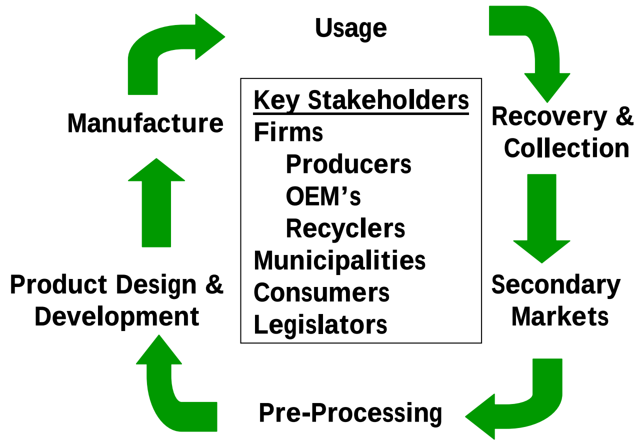

Figure 45. Product life-cycle showing key leverage points and major stakeholders

8.1 Recovery and collection: the challenge of accumulation

The chance-constrained method provides a tool to manage compositional uncertainty which can thereby increase the usage of recycled materials. However, when using scrap aluminum, it is often necessary to "dilute" the melt with primary aluminum in order to ensure the finished product meets compositional specifications. This is because recycled materials often include high levels of unwanted, or "tramp" elements. Many of these elements, for example, iron and silicon in the case of aluminum, are accumulated due to end-of-life processing such as shredding. A growing number of studies and literature would suggest that accumulation of unwanted elements is a growing problem, in all recycled material streams. Table XXV shows a brief literature review of various recycled material streams and their problematic accumulated elements. One can see that aluminum has a significant list. Though previous use of the chance-constrained model has used aluminum for specific case studies, this list illustrates that a generic model could be applied to a wide range of recycled material streams.

Table XXV. Possible tramp elements that increase with recycling

Material	Tramp Elements				
Steel	Ni, Cr, Sn[98, 99]	Cu[98-101]	Zn[98]		
Plastics	Cd				
Aluminum	Mg, Ni, Zn, Pb, Cr[102]	Fe[27, 102, 103]	Cu[102, 103]	V, Mn[103]	Si[27, 103]
Brass	Pb				
Copper	Fe, Pb, Ni, Cr, Sb, Bi, Se, Te[102]				
Glass	Al, SiC, C, Chromite, Carborundum[102]				
Cast Iron	Mn, Ni, Mo, Zn, Co[104]				

Metals recycling is a metallurgical process and is therefore governed by the laws of thermodynamics. The removal of unwanted elements in the scrap stream is dictated by the energy considerations of the melt process. In the case of aluminum, the thermodynamic barrier to the removal of most elements is quite large. Figure 46 shows an Ellingham diagram for alumina reduction illustrating the Gibbs free energy change as a function of temperature for various reactions. The main reaction of note, reduction of alumina to aluminum metal as expressed in Eq. (0.41) is the thick black line in the middle of Figure 46. One can see that the majority of equilibrium lines are at a higher free energy position than the aluminum line, so no partial pressure of oxygen would allow them to be oxidized into the slag. Only magnesium and calcium can be effectively removed from the melt by simple oxidation. Compared to many metals, aluminum has a high degree of difficulty in the removal of tramp elements. In the case of steel for Figure 46, only copper has a higher free energy than iron oxide reduction and therefore all other elements listed can be removed from the melt. These thermodynamics prevent a simple solution for accumulated tramp element removal.

$$\frac{4}{3}\mathrm{Al} + \mathrm{O}_2 \rightleftarrows \frac{2}{3}\mathrm{Al}_2\mathrm{O}_3 \quad (0.41)$$

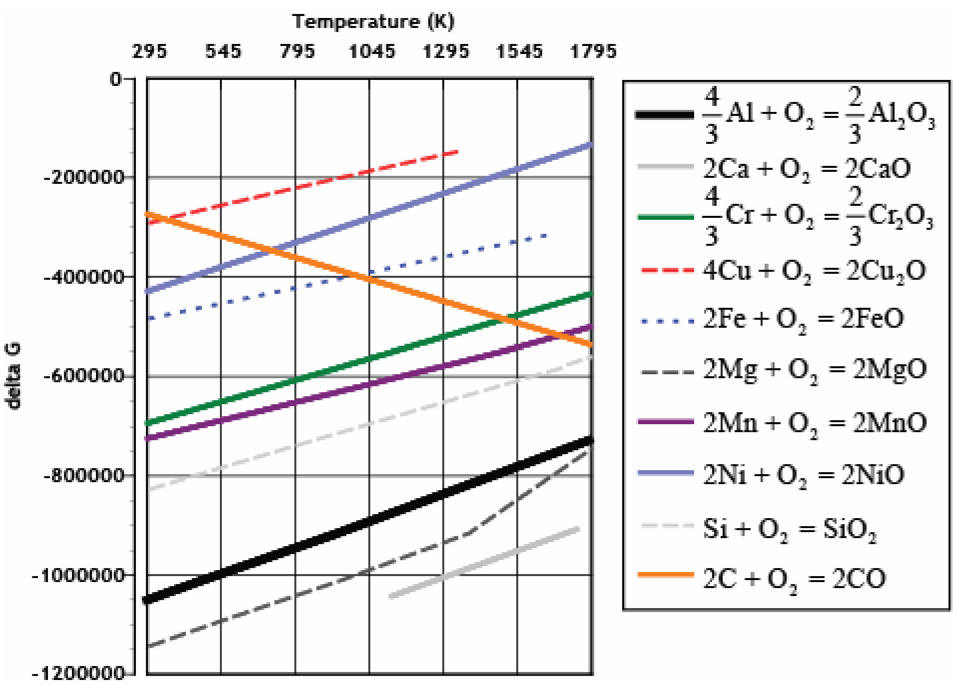

Figure 46. Ellingham diagram for various reactions[105, 106]

8.2 Secondary markets: material flow analysis

Accumulation of tramp elements is problematic and thermodynamics prevents an obvious solution. Devising strategies to mitigate this problem depends first upon understanding the mechanisms of accumulation in a multi-generation closed loop recycling system. Therefore, the first portion of future work will need to develop an understanding of the materials flow of secondary aluminum. Traditional analyses of scrap flows have relied upon market-wide statistical metrics that tend to mask fine technical structures that might offer desired insights into the management of compositional drift.

Figure 47 shows how the aluminum market is divided between various end-use products with transportation, containers and packaging, and construction making up the top three categories (75% of total products in 2003)[7]. The wide range in lifetimes for these products creates complexities in determining the availability and composition of the returning aluminum scrap stream. The vast majority of scrap in the containers and packaging category are used beverage cans which can have lifetimes as short as 60 days[107]. In contrast, the lifetime for aluminum used in construction applications (siding for houses, roofing, etc.) can be decades. Automotives which make up the majority of the transportation use sector, fall in the middle of these ranges with average lifetimes exceeding 15 years[108]. The wide range in lifetimes, varying quantities of recycled products, and variety in collection and recycling processes make modeling compositional accumulation highly complex.

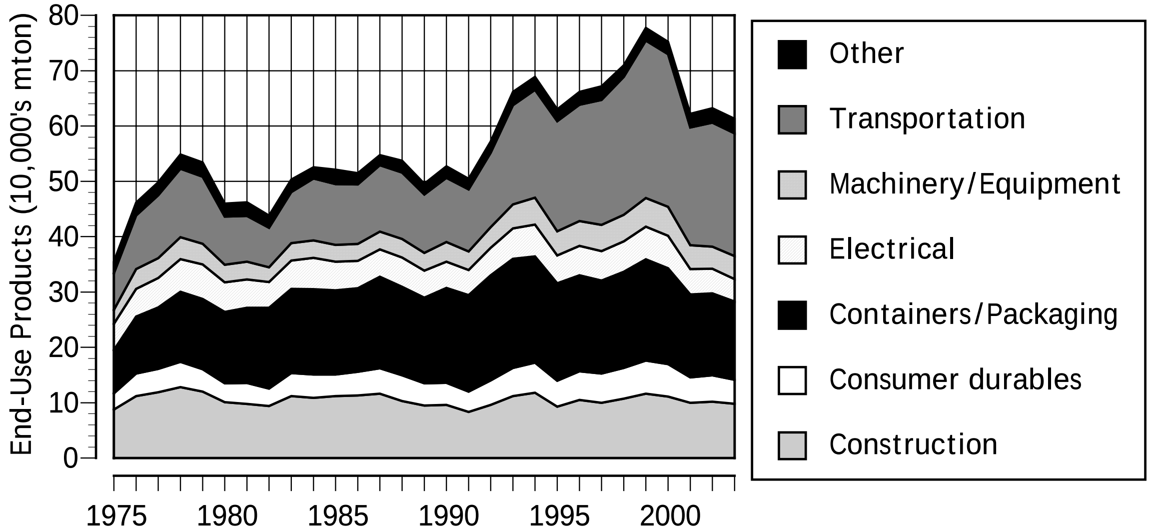

Figure 47. Aluminum products in the United States by end use sector[7]

8.3 Pre-processing: upgrading technologies evaluation

There are a variety of solutions to deal with the negative impact on recycling due to accumulation of undesired elements; each will have a trade-off between cost and scrap utilization (or recycling) as estimated by Figure 48. Dilution with primary is the most common; this has a negative impact on recycling as the required dilution results in a compositionally determined cap to recycling rates. "Down-cycling", where materials are recycled into lower value products, is another common method of dealing with highly contaminated secondary materials; this enables higher usage of recycled materials but negatively effects recycling economics. A specific example of down-cycling is when wrought scrap is used in cast products due to their ability to accommodate higher silicon contamination.

Other current processing solutions to the accumulation problem are dismantling of end-of-life products, spectrographic or magnetic sorting of shredded scrap, and "filtration" technologies that attempt to remove elements in the melt such as fractional crystallization and vacuum distillation. A variety of other technologies may exist that are still in the early stages of research and development. Although, qualitatively it is clear that such technologies could be useful, it is not clear that they would be economic and/or efficient. Understanding the cost and scrap utilization tradeoffs between these various strategies, as estimated schematically in Figure 48, is critical to determining their value. Therefore, future work will require a set of analytical tools to quantify the potential value of scrap "upgrading" technologies, including possibly "filtration", segregation, and sorting technologies. An important part of this analysis will be the extension of the single period batch mixing model to a multiple generation mixing model capable of characterizing a closed loop recycling system.

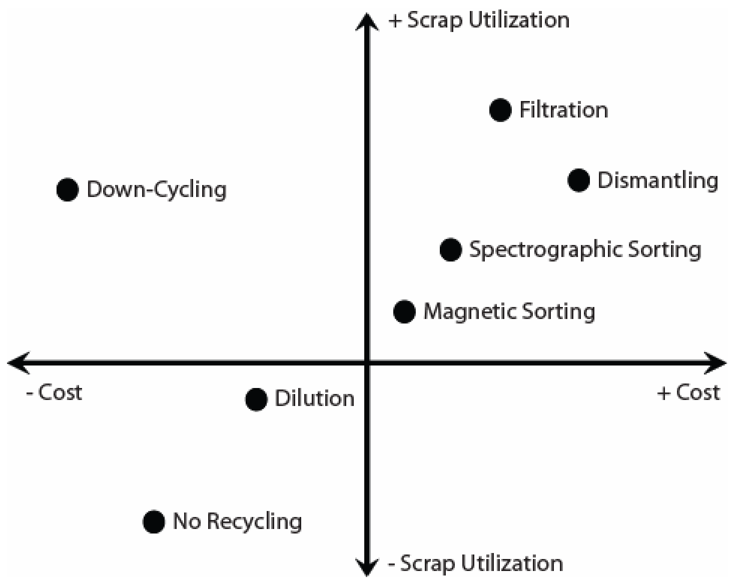

Figure 48. The cost and scrap utilization trade-offs of various strategies for dealing with compositional accumulation

Chapter 9. References

1. Graedel, T.E. and B.R. Allenby, *Industrial ecology*. 2nd edition Prentice-Hall international series in industrial and systems engineering. Vol. xix. 2003, Upper Saddle River, NJ: Prentice Hall. 363.
2. Frosch, R.A. and N.E. Gallopoulos, *Strategies for Manufacturing.* Scientific American, 1989. **261**(3): p. 144-152.
3. Khoei, A.R., I. Masters, and D.T. Gethin, *Design optimisation of aluminium recycling processes using Taguchi technique.* Journal of Materials Processing Technology, 2002. **127**: p. 96-106.
4. Peterson, R.D. *Scrap Variability and its Effects on Producing Alloys to Specification*. in *TMS: The Metals, Minerals, and Materials Society*. 1999.
5. Rong, A. and R. Lahdelma, *Fuzzy Chance Constrained Linear Programming Based Scrap Charge Optimization in Steel Production*. 2006, University of Turku, Department of Information Technology: Turku, Finland. p. 1-19.
6. Kelly, T., et al., *Historical Statistics for Mineral and Material Commodities in the United States*, U.G. Survey, Editor. 2004, USGS. p. 1-5.
7. Kelly, T.D. and G.R. Matos, *Historical statistics for mineral and material commodities in the United States*, in *Aluminum*. 2006, US Geological Survey.
8. Lee, H.L., V. Padmanabhan, and S. Whang, *The Bullwhip Effect in Supply Chains.* Sloan Management Review, 1997. **38**(3): p. 93-102.
9. *Global Scrap*. 2007, www.globalscrap.com.
10. Gaustad, G., P. Li, and R. Kirchain. *Modeling Methods for Managing Raw Material Compositional Uncertainty in Alloy Production*. in *TMS: The Minerals, Metals, and Materials Society*. 2006. San Antonio, Texas.
11. Cosquer, A. and R. Kirchain, *Optimizing The Reuse of Light Metals from End-of-Life Vehicles: Lessons from Closed Loop Simulations.* TMS: The Minerals, Metals, and Materials Society, 2003. **AA3**.
12. Li, P., *Strategies For Aluminum Recycling: Insights From Material System Optimization*, in *Department of Materials Science and Engineering (DMSE)*. 2005, Massachusetts Institute of Technology: Cambridge, MA. p. 183.
13. Li, P. and R. Kirchain. *Modeling Scrap Sourcing Decisions Given Uncertain Demand*. in *TMS: The Minerals, Metals, and Materials Society*. 2005.
14. Li, P., et al., *Identifying Economic and Scrap Reuse Benefits of Light Metals Sorting Technologies.* TMS: The Minerals, Metals, and Materials Society, 2005. **EPD Congress 2005**.
15. Gaustad, G., P. Li, and R. Kirchain. *Raw Material Usage Strategies Under Conditions of Uncertain Alloy Demand*. in *TMS: The Minerals, Metals, and Materials Society*. 2006. San Antonio, Texas.
16. Cheng, L., E. Subrahmanian, and A.W. Westerberg, *Design and planning under uncertainty: issues on problem formulation and solution.* Computers & Chemical Engineering, 2003. **27**: p. 781-801.
17. Harrison, J.M. and J.A. Van Mieghem, *Multi-resource investment strategies: Operational hedging under demand uncertainty.* European Journal of Operational Research, 1999. **113**: p. 17-29.
18. Sahinidis, N., *Optimization under uncertainty: state-of-the-art and opportunities.* Computers & Chemical Engineering, 2004. **28**: p. 971-983.
19. Shih, J.S. and H.C. Frey, *Coal blending optimization under uncertainty.* European Journal of Operational Research, 1995. **83**: p. 452-465.
20. Tang, O. and R. Grubbstrom, *Planning and replanning the master production schedule under demand uncertainty.* International Journal of Production Economics, 2002. **78**: p. 323-334.
21. Vladimirou, H. and S. Zenios, *Stochastic linear programs with restricted recourse.* European Journal of Operational Research, 1997. **101**: p. 177-192.
22. Liu, Z.-K., *Effect of Impurities on Alloys*, I.T. Program, Editor. 2003, Energy Efficiency and Renewable Energy, US Department of Energy: Washington, DC.
23. Miller, W.S., et al., *Recent development in aluminium alloys for the automotive industry.* Materials Science and Engineering A, 2000. **280**(1): p. 37-49.
24. Das, S.K., J.A.S. Green, and J.G. Kaufman. *Aluminum recycling - an integrated, industry-wide approach*. in *Global Symposium on Recycling, Waste Treatment, and Clean Technology (REWAS)*. 2008. Cancun, Mexico: TMS.
25. Woodward, R., *Where next wrought aluminium alloys?*, in *Aluminium Today*. 1997.

26. Sutherland, J.W., et al., *A global perspective on the environmental challenges facing the automotive industry: state-of-the-art and directions for the future.* International Journal of Vehicle Design, 2004. **35**(1/2).
27. Das, S.K. *Emerging Trends in Aluminum Recycling: Reasons and Responses*. in *TMS (The Minerals, Metals & Materials Society*. 2006. San Antonio, Texas.
28. Das, S.K., J.C.A. Green, and J.G. Kaufman, *The development of recycle-friendly automotive aluminum alloys.* Journal of Materials, 2007. **59**(11): p. 47-51.
29. Lund, J.R., et al., *Linear Programming for Analysis of Material Recovery Facilities.* Journal of Environmental Engineering, 1994. **120**(5): p. 1082-1094.
30. Stuart, J.A. and Q. Lu, *A Model for Discrete Processing Decisions for Bulk Recycling of Electronics Equipment.* IEEE Transactions on Electronics Packaging Manufacturing, 2000. **23**(4): p. 314-320.
31. Kirchain, R. and A. Cosquer, *Strategies for Maintaining Light Metal Reuse: Lessons from Closed Loop Modeling of End-of-Life Vehicles.* Resources Conservation & Recycling, 2007. **51**(2).
32. Gaustad, G., P. Li, and R. Kirchain, *Modeling Methods for Managing Raw Material Compositional Uncertainty in Alloy Production.* Resources Conservation & Recycling, 2007. **52**(2): p. 180-207.
33. van Schaik, A., et al., *Dynamic modelling and optimisation of the resource cycle of passenger vehicles.* Minerals Engineering, 2002. **15**: p. 1001-1016.
34. Reuter, M.A., et al., *Fundamental limits for the recycling of end-of-life vehicles.* Minerals Engineering, 2006. **19**: p. 433-449.
35. van Schaik, A. and M.A. Reuter, *The time-varying factors influencing the recycling rate of products.* Resources Conservation & Recycling, 2004. **40**: p. 301-328.
36. Neufville, R.d., *Applied Systems Analysis: Engineering Planning and Technology Management*. 1990, New York: McGraw-Hill Publishing Company. 470.
37. Metzger, R.W. and R. Schwarzbek, *A Linear Programming Application to Cupola Charging.* The Journal of Industrial Engineering, 1961. **12**(2): p. 87-93.
38. Kattentidt, H.U.R., T.P.R. de Jong, and W.L. Dalmijn, *Multi-sensor identification and sorting of bulk solids.* Control Engineering Practice, 2003. **11**: p. 41-47.
39. Schaik, A.v., et al., *Dynamic modelling and optimisation of the resource cycle of passenger vehicles.* Minerals Engineering, 2002. **15**: p. 1001-1016.
40. Bliss, N.G. *Scrap management and processing, cradle to ladle*. in *57th Electric Furnace Conference*. 1999. Pittsburgh, PA.
41. Tintner, G., *A note on Stochastic Linear Programming.* Econometrica, 1960. **28**(2): p. 490-496.
42. Dupacova, J., *Applications of stochastic programming: Achievements and questions.* European Journal of Operational Research, 2002. **140**: p. 281-290.
43. Debeau, D., *Linear Programming isn't Always the Answer.* Operations Research, 1957. **5**(3): p. 429-433.
44. Larsen, R.J. and M.L. Marx, *An Introduction to Mathematical Statistics and Its Applications*. 1981, Englewood Cliffs, NJ: Prentice-Hall, Inc. 536.
45. Johnson, S., K. Tefertiller, and D. Moore, *Stochastic Linear Programming and Feasibility Problems in Farm Growth Analysis.* Journal of Farm Economics, 1967. **49**: p. 908-919.
46. Moruyama, Y., *A Truncated Maximin Approach to Farm Planning under Uncertainty with Discrete Probability Distributions.* American Journal of Agricultural Economics, 1972. **54**: p. 192-250.
47. Sengupta, J.K., G. Tintner, and B. Morrison, *Stochastic Linear Programming with Applications to Economic Models.* Economica, 1963. **30**: p. 262-276.
48. Tintner, G., *Some Aspects of Stochastic Economics.* Stochastics, 1973. **1**: p. 71-86.
49. Ziemba, W.T., C. Parkan, and R. Brooks-Hill, *Caclulation of Investment Portfolios with Risk Free Borrowing and Lending.* Management Science, 1974. **21**: p. 209-222.
50. Bradley, S. and D. Crane, *A Dynamic Model for Bond Portfolio Management.* Management Science, 1972. **19**: p. 139-151.
51. King, W.R., *A Stochastic Personnel Assignment Model.* Operations Research, 1965. **13**: p. 67-81.
52. Seppala, Y., *On a Stochastic Multifacility Location Problem.* AAIE Transactions, 1975. **7**: p. 56-62.
53. Glen, J.J., *A mixed integer programming model for fertilizer policy evaluation.* European Journal of Operational Research, 1988. **35**: p. 165-171.
54. Balintfy, J.L. and A. Prekopa, *Nature of Random Variation in the Nutrient Composition of Meals.* Health Services Research, 1966. **1**: p. 148-169.
55. van de Panne, C. and W. Popp, *Minimum Cost Cattle Feed under Probabalistic Protein Constraints.* Management Science, 1963. **9**(3): p. 405-430.

56. Rahman, S. and F.E. Bender, *Linear Programming Approximation of Least Cost Feed Mixes with Probability Restrictions.* American Journal of Agricultural Economics, 1971. **53**: p. 612-618.
57. Hodges, S.D. and P.G. Moore, *The Product-Mix problem under Stochastic Seasonal Demand.* Management Science, 1970. **17**: p. 107-114.
58. Candler, W., *Coal blending with acceptance sampling.* Computers and Operations Research, 1991. **18**: p. 591-596.
59. Martin, C.H. and S. Lubin, *Optimization modeling for business planning at Trumbull Asphalt.* Interfaces, 1985. **15**: p. 66-72.
60. Ashayeri, J., A.G.M. van Eijs, and P. Nederstigt, *Blending modeling in process manufacturing: A case study.* European Journal of Operational Research, 1994. **72**: p. 460-468.
61. Rigby, B., L.S. Lasdon, and A.D. Warren, *The evolution of Texaco's blending systems: From OMEGA to StarBlend.* Interfaces, 1995. **25**: p. 64-83.
62. Kall, P. and J. Mayer, *Stochastic Linear Programming: Models, Theory, and Computation*. International Series in Operations Research & Management Science, ed. F.S. Hillier. 2005, New York, NY: Springer.
63. Cattani, K., G. Ferrer, and W. Gilland, *Simultaneous Production of Market-Specific and Global Products: A Two-Stage Stochastic Program with Additional Demand After Recourse.* Naval Research Logistics, 2003. **50**: p. 438-461.
64. Liu, C., Y. Fan, and F. Ordonez, *A two-stage stochastic programming model for transportation network protection.* Computers & Operations Research, 2009. **36**: p. 1582-1590.
65. Lin, Q.G., et al., *IFTEM: An interval-fuzzy two-stage stochastic optimization model for regional energy systems planning under uncertainty.* Energy Policy, 2008. **10**.
66. Francas, D. and S. Minner, *Manufacturing network configuration in supply chains with product recovery.* Omega, 2009. **37**(4): p. 757-769.
67. Petruzzi, N.C. and M. Dada, *Information and Inventory Recourse for a Two-Market, Price-Setting Retailer.* Manufacturing & Service Operations Management, 2001. **3**: p. 242-263.
68. Kira, D., M. Kusy, and I. Rakita, *A stochastic linear programming approach to hierarchical production planning.* Journal of the Operational Research Society, 1997. **48**: p. 207-211.
69. Castro, J., *A stochastic programming approach to cash management in banking.* European Journal of Operational Research, 2009. **192**: p. 963-974.
70. Martel, A. and W. Price, *Stochastic programming applied to human resource planning.* Journal of Operations Research Society, 1982. **32**(3).
71. Charnes, A. and W.W. Cooper, *Chance constrained programming.* Management Science, 1959. **6**: p. 73-89.
72. Charnes, A. and W.W. Cooper, *Deterministic Equivalents for Optimizing and Satisficing under Chance Constraints.* Operations Research, 1963. **11**: p. 18-39.
73. Azaiez, M.N., M. Hariga, and I. Al-Harkan, *A chance-constrained multi-period model for a special multi-resevoir system.* Computers & Operations Research, 2005. **32**: p. 1337-1351.
74. Rong, A. and R. Lahdelma, *Fuzzy chance constrained linear programming model for optimizing the scrap charge in steel production.* European Journal of Operational Research, 2008. **186**: p. 953-964.
75. Cao, C., X. Gu, and Z. Xin, *Chance constrained programming models for refinery short-term crude oil scheduling problem.* Applied Mathematical Modelling, 2009. **33**: p. 1696-1707.
76. Miller, B.L. and H.M. Wagner, *Chance Constrained Programming with Joint Constaints.* Operations Research, 1965. **13**(6): p. 930-945.
77. Prekopa, A., *A class of stochastic programming decision problems.* Matematische Operations forschung und Statistik, 1972: p. 349-354.
78. Ruszczynski, A. and A. Shapiro, *Stochastic Programming Models*, in *Handbooks in Operations Research and Management Science*. 2003. p. 1-64.
79. Growe, N., *Estimated stochastic programs with chance constraints.* European Journal of Operational Research, 1997. **101**: p. 285-305.
80. Mihailidis, D. and K.R. Chelst, *Optimal blending of out-of-specifications substances.* Journal of the Operational Research Society, 1998. **49**: p. 458-466.
81. Evers, W.H., *A new model for stochastic linear programming.* Management Science, 1967. **13**(9): p. 680-693.
82. Karmarkar, U.S. and K. Rajaram, *Grade Selection and Blending to Optimize Cost and Quality.* Operations Research, 2001. **49**(2): p. 271-280.
83. Springer, M.D., *The algebra of random variables*. 1979, New York: Wiley. 470.

84. Cornwell, L.W., L.A. Aroian, and V.S. Taneja, *Numerical Evaluation of the Distribution of the Product of Two Normal Variables.* Journal of statistical computation and simulation, 1978. **7**: p. 123-131.
85. Aroian, L.A., V.S. Taneja, and L.W. Cornwell, *Mathematical Forms of the Distribution of the Product of Two Normal Varaibles.* Communications in statistics, theory, and methods, 1978. **7**(2): p. 164-172.
86. Gorban, L.T., G.K. Ng, and M.B. Tessieri. *An In-Depth Analysis of Automotive Aluminum Recycling in the Year 2010*. in *SAE International Congress*. 1994. Detroit, MI.
87. Xiao, Y. and M.A. Reuter, *Recycling of distributed aluminium turning scrap.* Minerals Engineering, 2002. **15**(11, Supplement 1): p. 963-970.
88. Das, S.K., *Final Technical Report: Reduction of Oxidative Melt Loss of Aluminum and Its Alloys*. 2006, Secat, Inc.: Lexington, KY.
89. Amini, S.H., et al., *Quantifying the quality loss and resource efficiency of recycling by means of exergy analysis.* Journal of Cleaner Production, 2007. **15**(10): p. 907-913.
90. Rossel, H. *Fundamental Investigations About Metal Loss During Remelting of Extrusion and Rolling Fabrication Scrap*. in *The Minerals, Metals, and Materials Society (TMS) Annual Meeting*. 1990.
91. Cosquer, A., *Optimizing the Reuse of Light Metals from End-of-Life Vehicles*, in *Engineering Systems Division*. 2003, Massachusetts Institute of Technology: Cambridge, MA. p. 98.
92. Market, A.M., *Non-ferrous scrap prices*. 2006.
93. *Recycling Xchange*. 2006, www.recyclingxchange.com.
94. Plunkert, P., *Aluminum*, in *US Geological Survey Minerals Yearbook*. 2005, USGS. p. 20.
95. Gesing, A., *Recycling light metals from end-of-life vehicles.* Journal of Materials, 2001. **53**(11).
96. Gesing, A., *Assuring the Continued Recycling of Light Metals in End-of-Life Vehicles: A Global Perspective.* Journal of Materials, 2004. **August 2004**: p. 18-27.
97. Gaustad, G., P. Li, and R. Kirchain. *Managing Raw Materials Compositional Uncertainty: Chance Constrained Modeling for Alloy Production*. in *Materials Science and Technology (MS&T) 2006*. 2006. Cincinnati, OH.
98. Noro, K., M. Takeuchi, and Y. Mizukami, *Necessity of Scrap Reclamation Technologies and Present Conditions of Technical Development.* Iron and Steel Institute of Japan (ISIJ) International, 1997. **37**(3): p. 198-206.
99. Menad, N., *Modelling of environmental, technical, and economical aspects of impurity accumulation in the recycling of steel*. 1999.
100. Cho, W., et al., *A New Approach for the Removal of Copper from Solid Ferrous Scrap.* AISTech 2004: Iron and Steel Conference Proceedings, 2004. **I&II**: p. 861-867.
101. Daigo, I., et al., *Development of a Dynamic Model for Assessing Environmenal Impact Associated with Cyclic Use of Steel.* Journal of the Iron and Steel Institute of Japan (Tetsu-to-Hagane), 2004.
102. Lundqvist, U., et al., *Design for Recycling in the Transport Sector- Future Scenarios and Challenges*. 2004, Goteborg, Sweden: Chalmers University of Technology, Goteberg University. 335.
103. Kim, J.-Y., et al. *Aging Characteristics of recycled ACSR wires for distribution lines*. in *Electrical Insulation Conference and Electrical Manufacturing & Coil Winding Conference*. 1997. Rosemont, IL.
104. Anigstein, R., et al., *Potential Recycling of Scrap Metal from Nuclear Facilities*. 2001, S. Cohen & Associates.
105. Kubaschewski, O., E.L.L. Evans, and C.B. Alcock, *Metallurgical Thermochemistry*. 4th ed. ed. 1979, Oxford: Pergamon. 329.
106. Ragone, D.V., *Thermodynamics of Materials*. MIT Series in Materials Science and Engineering. Vol. Vol. I. 1995, New York: John Wiley & Sons, Inc. 311.
107. Kelly, P., *How are aluminum cans recycled?* 2008, The Aluminum Association.
108. Davis, S. and S. Diegel, *Transportation Energy Data Book*. Vol. ORNL-6967. 2002, Oak Ridge, TN: Oak Ridge National Laboratory.