Caving 2026 – A van As, D Cumming-Potvin & J Wesseloo (eds)

doi:10.36487/ACG_repo/2645_29

Increasing net present value in sublevel caving: incorporating discount rate as an essential input for mineable inventory and scheduling

Scott Loewen [a,*]

[a] *SRK Consulting, Canada*

Abstract

Net present value (NPV) is a fundamental metric in project evaluation, providing a basis for comparing the economic strengths of various projects and mine plans. Central to an NPV calculation is the consideration of the time value of money through the application of a discount rate.

This paper investigates the integration of the discount rate as a primary parameter within mineable inventory selection and production scheduling. The objective is to optimise the footprint and shut-off strategy to increase discounted operating cash flow (DOC) and hence the NPV of sublevel caving (SLC) mine plans relative to those generated by conventional methods using a static cut-off value. Employing a simplified SLC model, the paper details a mathematical approach for determining a dynamic cut-off value that prevents reduction in DOC due to mining a too-low grade of ore early within the mine schedule. Additionally, this paper presents a computer program created by the author which applies this methodology to generate SLC inventories and production schedules via material flow simulations. Comparative examples are provided to illustrate the impact on mine plans using this methodology and those developed without considering the time value of money.

Keywords: *net present value, discount rate, operating cash flow, optimisation, sublevel caving, cut-off, shut-off, hill of value, simulation*

1 Introduction

This paper presents an algorithmic approach for generating sublevel caving (SLC) footprints and draw strategies that incorporate the time value of money directly into mineable inventory selection and scheduling. Unlike conventional approaches that apply a fixed cut-off grade or rely on heuristic hill of value (HoV) analyses, the method derives a time-variable profit threshold that reflects the impact of discounting on future cash flows. The resulting production strategies are typically comparable to near-optimal HoV outcomes while requiring significantly less trial-and-error analysis.

The method does not explicitly solve for a globally optimal net present value (NPV). Instead, it defines a time-dependent minimum profit curve (MPC) representing the minimum profit per tonne required at each point in the schedule to avoid reducing discounted operating cash flow (DOC). Mining material below this threshold reduces DOC, while mining material above it can increase DOC. In practice, solutions generated using the MPC tend to lie close to the maximum of an HoV curve for a given set of assumptions.

The MPC changes over the life of the mine, generally declining towards a break-even profit margin at the end of the schedule as future cash flows diminish. This behaviour encourages the prioritisation of higher-margin material earlier in the mine life, which increases the present value of cashflows without relying on changes to production rate, capital expenditure or scheduling assumptions.

[*] Corresponding author. Email address: sloewen@srk.com

By advancing higher-margin material in the production schedule, the method reduces exposure to uncertainty in future cash flows captured within the discount rate. Unlike approaches that rely on increasing production rates or accelerating development, this method improves value primarily through sequencing decisions rather than changes to operating or capital assumptions.

The primary trade-off associated with this approach arises from the use of a higher-margin cut off that prioritises higher-value material, resulting in reduced resource utilisation and a shorter mine life compared to strategies using lower-margin cut offs. However, unlike a static elevated cut off or shut off, the dynamic MPC provides a compromise by allowing lower-margin material to be mined later in the schedule, thereby balancing investment return and resource extraction.

The algorithm is not intended to replace detailed case-matrix analyses across a range of economic and operating assumptions. Rather, it provides a systematic means of generating well-formulated production strategies for a specified discount rate, which can then be evaluated alongside alternative scenarios to identify the most appropriate outcome for a given project.

The algorithm presented in this paper does not consider capital costs; to make this distinction from NPV the author has defined the term Discounted Operating Cashflow (DOC). While DOC is used throughout this paper to evaluate the impact of scheduling decisions, changes in DOC are directly comparable to changes in NPV, where capital costs are assumed to remain constant.

2 Conceptual model

Generally, the impacts of discounting are considered only for larger timescales and cashflows. Examples include time to first production, delaying or minimising capital, or accelerated cashflows (i.e. increased production rates). But it still applies on smaller timescales, such as whether the profit on a tonne of material covers the reduction in the present value of future cashflows for the time it takes to mine it.

As an example, if a project has a DOC of 1 billion at a discount rate of 10% and a production rate of 10 ktpd, the reduction in DOC would be 26.09 for each additional tonne mined at the start of the project, according to the following equation:

$$\Delta DOC = DOC\left(1 - \frac{1}{(1 + r)^{\tau/P}}\right)$$

where:

ΔDOC	=	change in discounted operating cash flow
DOC	=	discounted operating cash flow
r	=	discount rate
τ	=	duration of delay, expressed in the same time units as the discounting period, P
P	=	discounting period (e.g. 1 year if annual discounting).

Therefore, any material mined at the start of the project must generate a minimum profit of 26.09/t to avoid reducing the project's DOC. This minimum required profit will vary throughout the lifetime of the project as the DOC of the future cash flow changes, reducing to zero at the end of the mine life. The variation of this minimum profit over time is referred to as the MPC.

A simplified SLC model presented in Figure 1 illustrates how this calculation is used to derive the time-variable MPC for a given deposit and discount rate.

Mining Sequence →

	Block		1	2	3	4	5	6	7	8	9	10	**Totals**
(a)	**Profit**		**10.00**	**5.00**	**2.00**	**5.00**	**10.00**	**15.00**	**25.00**	**20.00**	**20.00**	**10.00**	**122.00**
(b)	**DOC @10%**		**9.53**	**4.33**	**1.58**	**3.58**	**6.51**	**8.88**	**13.46**	**9.79**	**8.90**	**4.04**	**70.60**
(c)	**Future DOC**	70.60	67.17	68.64	73.41	75.51	72.57	64.10	44.28	27.74	9.53	0.00	
(d)	**Minimum Profit**		6.40	6.54	7.00	7.20	6.92	6.11	4.22	2.64	0.91	0.00	
(e)	**Updated DOC**		~~**9.53**~~	~~**0.00**~~	~~**0.00**~~	~~**0.00**~~	~~**8.67**~~	~~**11.82**~~	~~**17.53**~~	~~**13.02**~~	~~**11.84**~~	~~**5.38**~~	**77.80**

Figure 1 Simplified sublevel caving model showing steps to deriving time-variable cut-off strategy. (a) Sequence of blocks with associated profit per block; (b) Profit of block discounted to the start of mining at 10% per period, based on mining per period; (c) The future discounted operating cash flow (DOC) at the start and end of each period; (d) Minimum profit margin based on the cost of delaying future DOC by one period; (e) The updated mine plan that increases DOC by skipping blocks below the calculated minimum profit calculated for that period

The model consists of a stack of 10 blocks (a). Each block has an associated value that represents the operating profit to be made by mining the block. It should be emphasised that the blocks do not represent levels within an SLC but a sequential stream of material that must be mined according to the following set of simple rules:

- One block can be mined per period.
- The blocks must be mined in sequence.
- Any block can be skipped, bringing forward all subsequent blocks by one period.
- All profit from any skipped block is lost and cannot be recovered.

In the example shown, all blocks have a positive profit value; thus the maximum undiscounted value is achieved by mining all blocks in the sequence.

Figure 1b shows the profit of each block discounted at 10% per period to the start of mining, assuming all blocks are mined. The sum of these values represents the DOC of mining all ten blocks.

Figure 1c shows the future DOC at the start and end of each block/period. The DOC at the start of the first period is equivalent to the total DOC shown in (b). Future DOC reduces to zero at the end of the last period as there is no future revenue at this point.

A minimum profit value is calculated for each period (d) based on the reduction in that period's future DOC if it was delayed by one period. Any block or period with a profit below this value would result in lower DOC.

Any block below the minimum profit is then skipped (e), zeroing out its profit and bringing the future cash flow forward by one period. Although profitable material is removed from the mine plan, the DOC is increased due to bringing forward the revenue from future periods. In this example, the DOC increases by more than 10% compared to mining all 10 blocks, despite mining 30% less material.

3 Implementation

The methodology presented in the previous section depends on knowing the future cash flow. In our simple example this was not a problem as our actions in any period did not impact the future, but this is not the case in a real SLC or with detailed simulations that account for mixing, material flow and ore loss.

An iterative approach is therefore required, where subsequent simulations are run with the cash flow of the preceding schedule used to derive the MPC used for the ring selection and draw strategy in the next iteration. This process is repeated until the MPC and resulting mine plan do not significantly change between iterations. Figure 2 shows a schematic of the overall process.

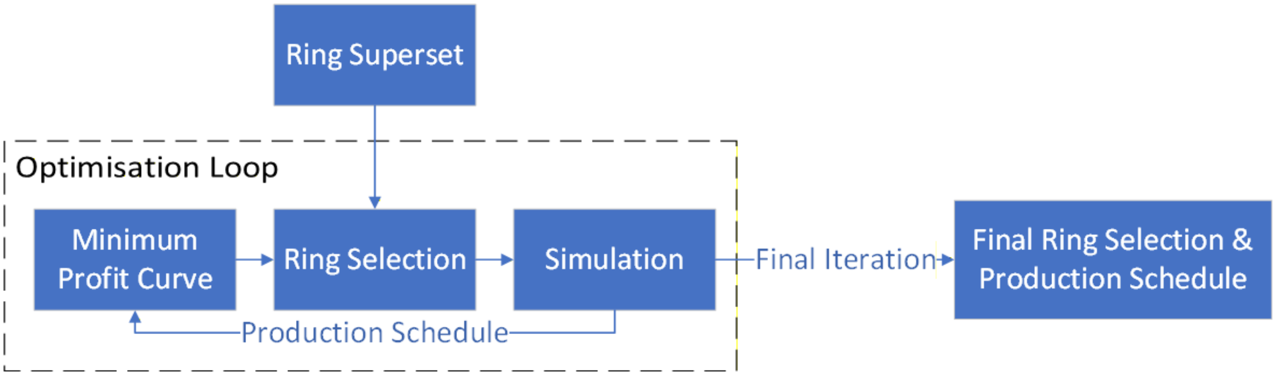

Figure 2 Schematic of the iterative optimisation process

To implement this process a proprietary SLC simulation software was developed that automates each step and will run the process for a user-defined number of iterations. The individual steps in the process are discussed in the following.

3.1 Ring superset

The process begins with the creation of a ring superset. The superset represents the largest possible extents of the SLC. First a cut off is defined; this is usually at or below the lowest expected operating cost for the SLC. Then a contiguous set of rings is created to cover all blocks within the block model above the cut off. Any intervening blocks that are below the cut off are also included to avoid gaps within the superset. For smaller deposits the superset may cover all mineralisation in the block model. Figure 3 shows an example superset created to cover all blocks above the cut off.

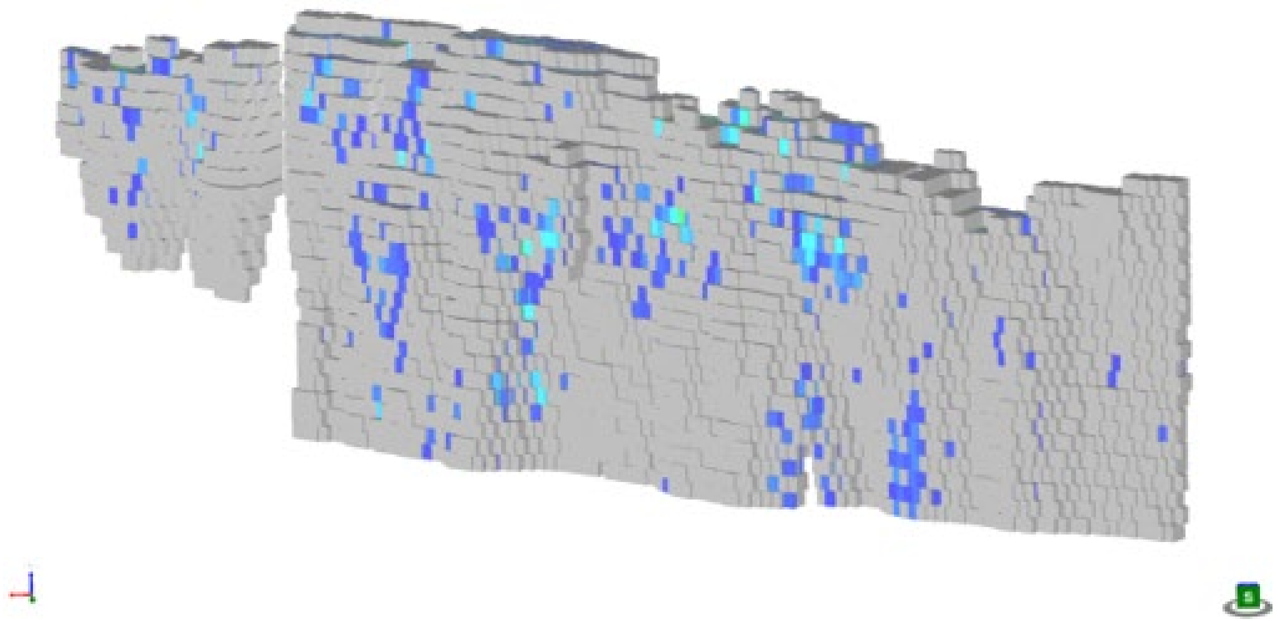

Figure 3 Example of a ring superset covering all blocks above the cut off

In addition to the initial superset a dilution model is created which includes any overlying material that could flow into the SLC as dilution†. This includes mineralised material from the block model below the cut off which is modelled as additional rings outside the superset. Overburden and backfill material from outside the block model are modelled as external dilution sources that are typically connected to the boundary of the superset and dilution rings.

3.2 Minimum profit curve

The MPC defines a minimum profit per tonne for each period. This is used both in the ring selection and to define the shut off used in the draw strategy.

For the initial iteration the MPC is held constant at zero throughout the mine life. The MPC is updated after each simulation based on the net operating cash flow (NOC) of the resulting production schedule. The updated MPC is then used in the ring selection and draw strategy of the next iteration. Figure 4 shows a typical progression of the MPC produced from each iteration.

† Dilution is defined as any material drawn from an SLC draw point which has originated outside the blasted sublevel caving rings (Power 2004). Dilution may or may not be comprised of mineralised material.

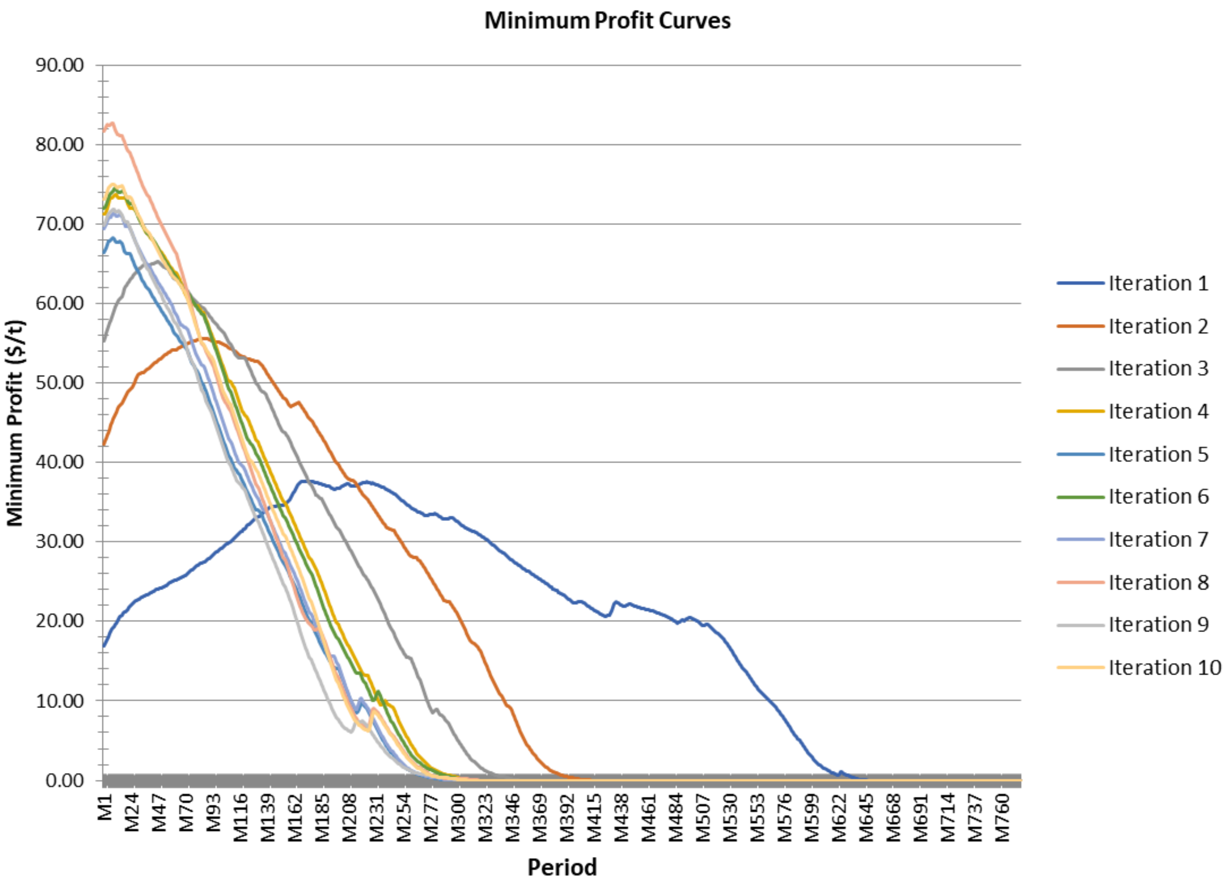

Figure 4 Example of a minimum profit curve per iteration

Iteration 1 MPC is used in iteration 2. Iteration 2 is used in iteration 3 and so on. Iteration 1 typically uses a flat zero profit MPC which is not plotted on the graph. The slope of MPC typically steepens with each subsequent iteration until there is little to no difference between successive iterations. In all cases the MPC value decreases to zero at the end of the mine life since there is no future cash flow.

3.3 Ring selection

For each iteration a subset of rings is selected from within the superset based on the MPC of the previous iteration. All rings within the ring selection will be drilled, blasted and mined in the simulation. To speed up and facilitate the automation of the process, the ring selection is produced algorithmically.

Net cash flow per ring is estimated based on the previous production schedule. The algorithm will then remove any ring with a NOC below the MPC for the period in which it is blasted if it can do so without violating a set of rules intended to maintain a practical SLC mine shape. The selection process will also remove any tunnels or levels that do not cover their associated development cost, provided their removal does not violate the selection rules.

Any rings removed from the ring selection are moved to the dilution model so that the in situ material, with its associated value and grades, can be recovered as dilution in the simulation. Figure 5 shows the progression from the ring superset to the ring selection of the first iteration and through to the ring selection of the last. Note that the final ring selection contains many sub-economic rings (grey) that could not be removed without violating the ring selection constraints.

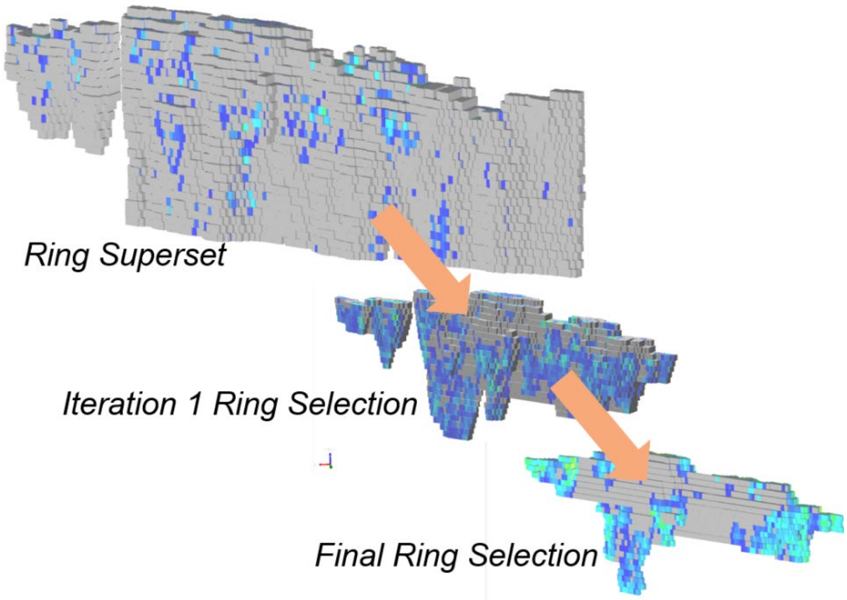

Figure 5 From superset to final ring selection

3.4 Simulation

The ring selection is then used in an SLC simulation to generate a production schedule and extract the mixed material from each ring.

The opening date of each ring in the simulation is determined by the interaction of the desired mining sequence, development rates, undercutting constraints, drawpoint productivity and overall production target.

The resulting draw strategy is a combination of a minimum extraction, the MPC and extraction limits. The extraction limits define a maximum extraction for each ring based on the number of levels overlying the ring. This is done to ensure the formation of an ore blanket and prevent the creation of void space. A ring may be shut prior to reaching its limit if it has exceeded the minimum extraction and the grade of the rings falls below the shut off for the current period. The shut off for the periods is equal to the expected cost per milled tonne plus the MPC value for the current period. The shut off applied to a ring is often lower than the cut off applied to it since shut off does not necessarily include drill and blast costs, the shut off is applied later in the schedule and the MPC generally reduces with time.

From the simulation a production schedule is produced and the operating cash flow for each month is calculated from the value of the extracted material minus the estimated operating costs. This cash flow is then used to generate the MPC for the next iteration.

3.5 Final ring selection and production schedule

Iterations are repeated until there is little difference between the MPC, ring selection and production schedule created with each subsequent iteration. Figure 6 shows the resulting cumulative DOC for sequential iterations. After iteration 3 the results are nearly identical, with similar ring selections and production schedules.

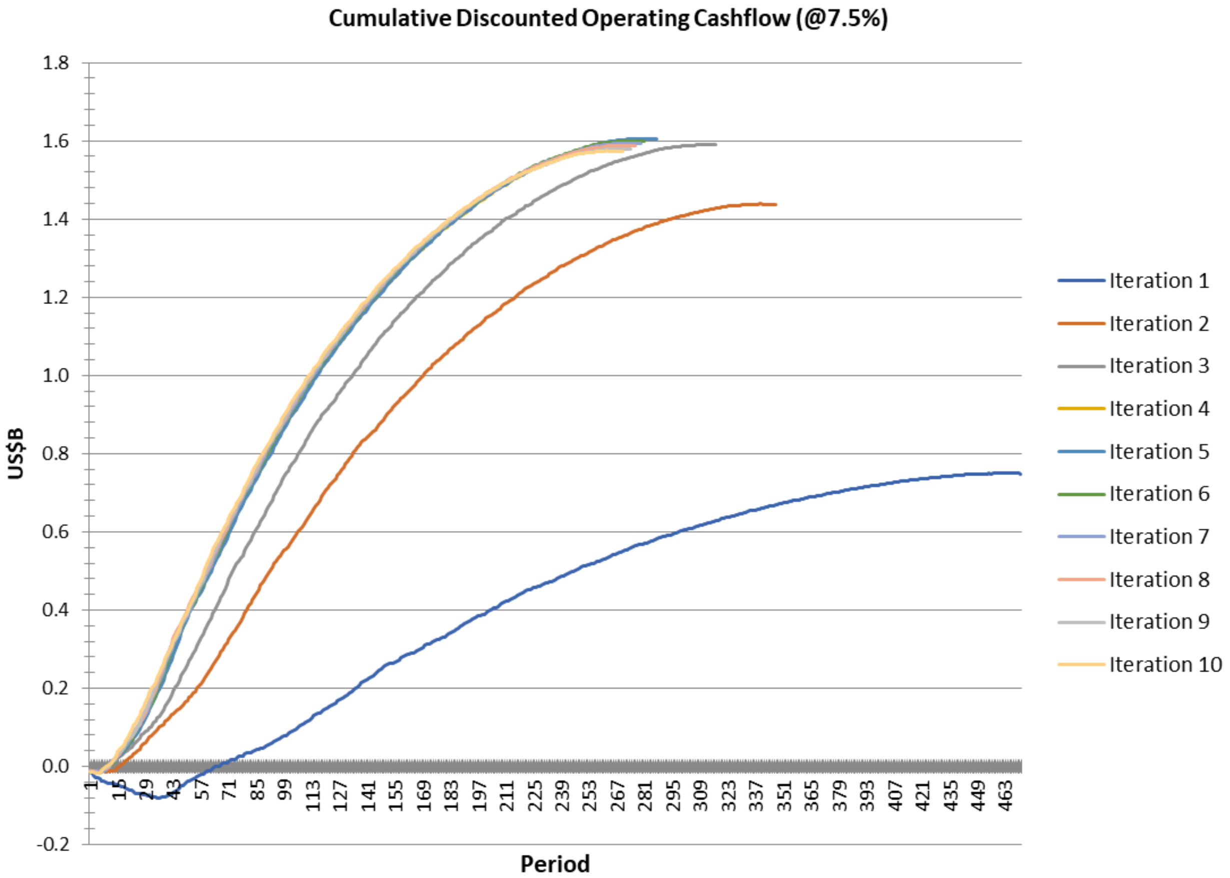

Figure 6 Cumulative discounted operating cash flow per iteration

4 Example analysis

To help illustrate the difference in outcomes between the typical fixed cut-off and shut-off approach used to generate an SLC footprint and draw strategy versus the algorithmic process presented in this paper, an example analysis is provided.

4.1 Example orebody

To create a realistic grade distribution the example presented is based on a portion of a real-world deposit, but the absolute revenue value of each ring has been scaled and the underlying commodity grades removed. Figure 7 shows the ring superset for the example deposit.

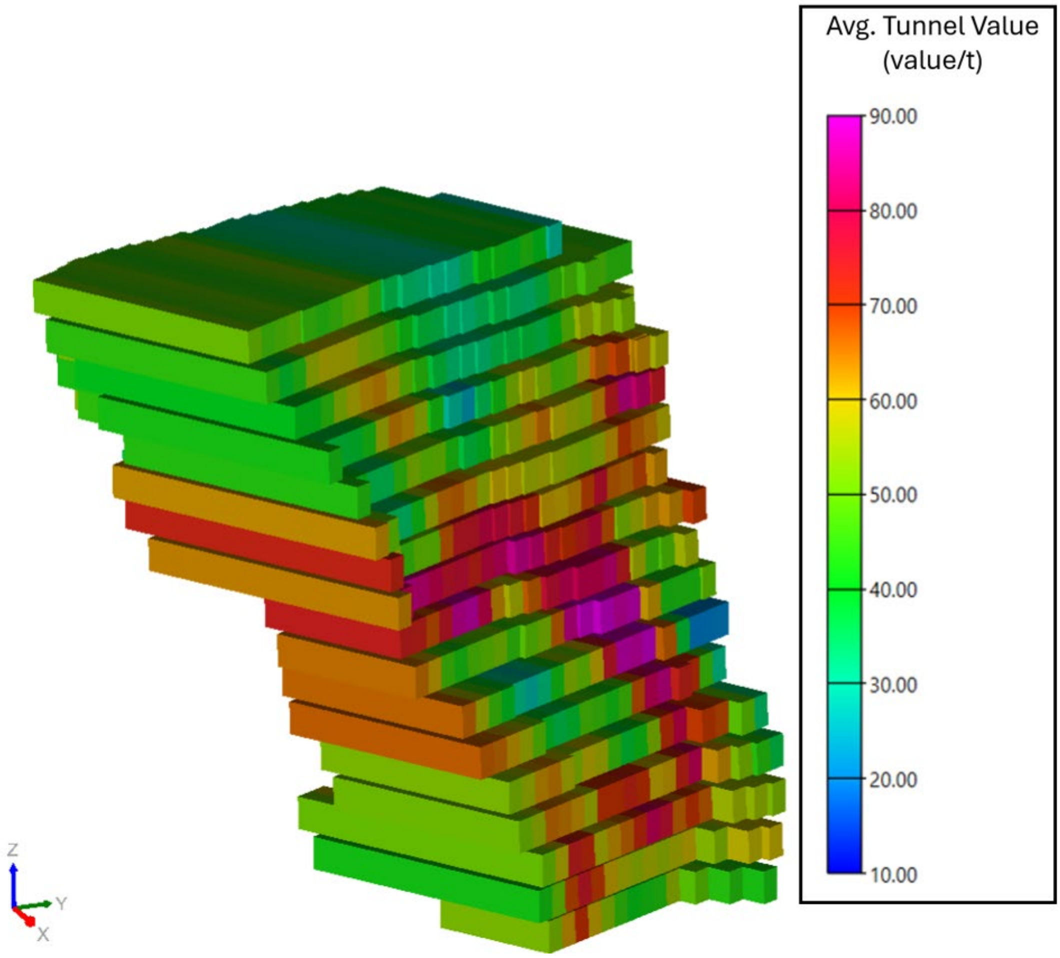

Figure 7 Ring superset of an example deposit

The superset is comprised of just over 25 thousand rings covering all material with a value greater than 15/t within the portion of the block model. To show the relative grade distribution within the deposit the image is showing the average value of each tunnel which ranges from 0/t to plus 90/t.

Figure 8 shows the dilution model containing all mineralised material overlying the superset. Barren hanging wall material and overburden are modelled using external dilution sources that are connected to the top of the dilution rings.

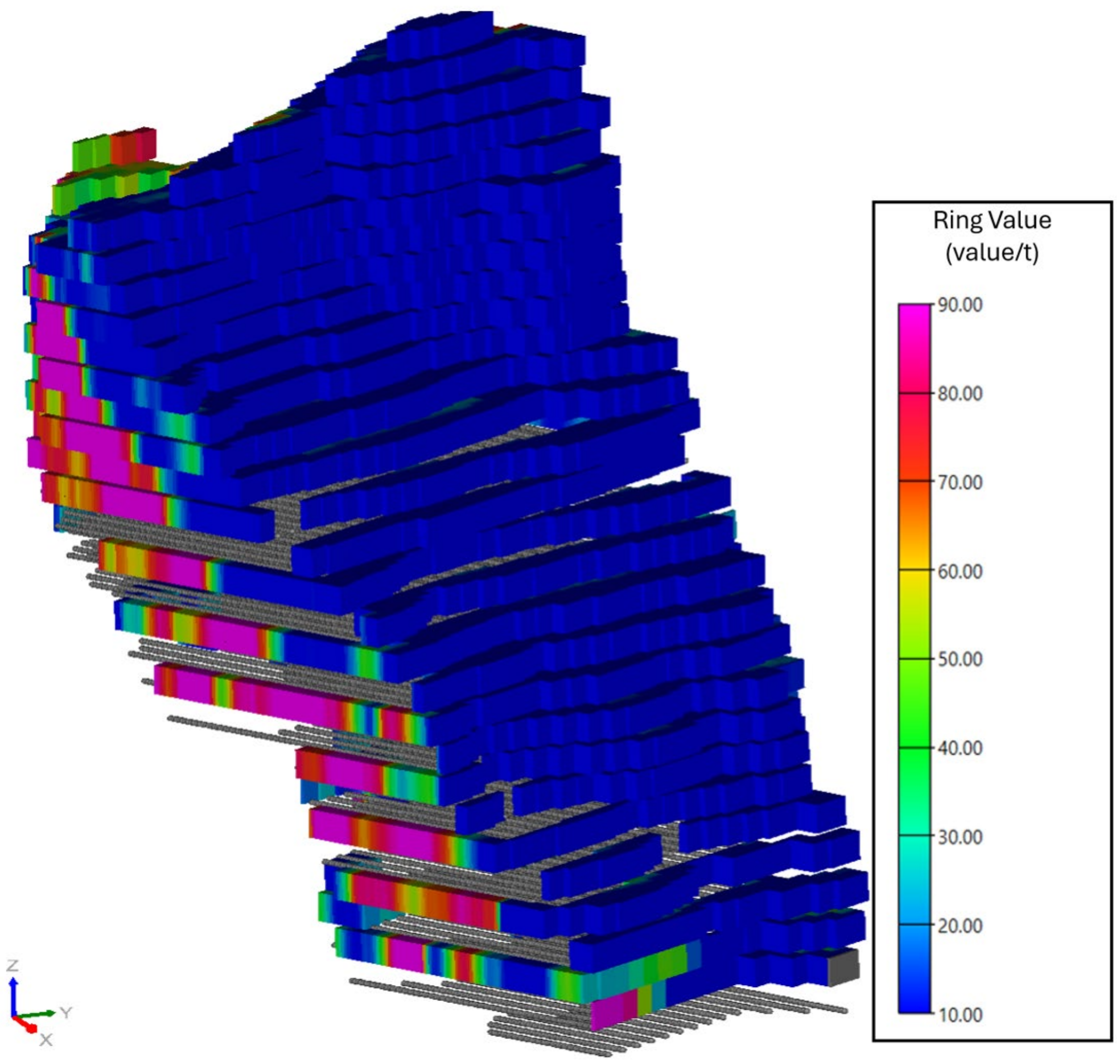

Figure 8 Superset (shown as grey spheres) with an overlying dilution model

4.2 Scenarios

From the superset, 5 production schedules were produced for various combinations of operating costs and discount rates as outlined in Table 1.

Table 1 List of scenarios and process used to generate sublevel caving ring selection and draw strategies

Scenario no.	Process	Discount rate (%)
1	Algorithmic	5
2	Algorithmic	10
Scenario no.	**Process**	**Cut off/shut off (value/t)**
3	Static cut off/shut off	35
4	static cut off/shut off	45
5	static cut off/shut off	55
6	static cut off/shut off	65

A single operating cost of 25/t was used for all scenarios. For the algorithmic process two discount rates (5 and 10%) were defined. For the static cut-off/shut-off scenarios, values from 35/t to 65/t were used, which represent a 10/t to 40/t minimum operating margin.

The production rate for all scenarios was 3 Mtpa. The analysis did not evaluate for different production rates as any comparison without including the increased capital costs associated with higher production rates would render the comparison meaningless, higher production rates invariably produce higher DOC values.

4.3 Example algorithmic process

Five iterations were run for each of the 3 algorithmic scenarios. For each scenario the iteration with maximum DOC was chosen as the final production schedule. Scenario 2, using a discount rate of 10%, is presented as an example of this iterative process.

Figure 9 shows the MPC calculated after each iteration for scenario 2. Each MPC is used in the subsequent iteration for ring selection and draw strategy. Since there is no simulation prior to iteration 1, a static zero profit margin is used, with both cut off and shut off equal to the break-even operating cost.

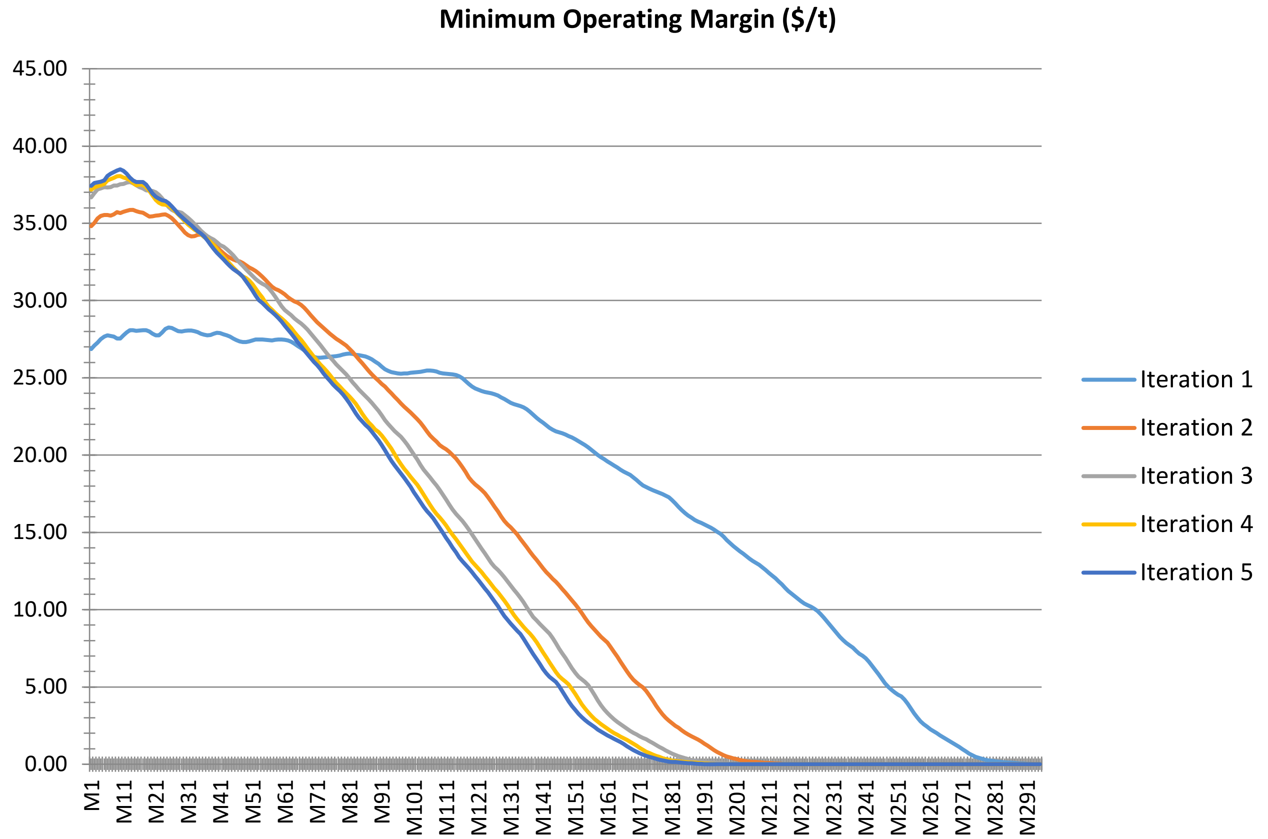

Figure 9 Minimum profit curve for scenario 2

Figure 9 shows typical MPC curves for each iteration with minimum profit margin starting high and dropping towards zero at the end of the mine life. For iteration 3 and above, the derived MPC curves are quite similar. The similarity of these curves produces similar results for iterations 3 to 5, as shown in Table 2. The differences between subsequent iterations become smaller with each iteration as the algorithm converges on a solution.

Table 2 Iteration results for scenario 2

Iteration	No. of rings	In situ mass (Mt)	Extraction (%)	Dilution (%)	Revenue grade (value/t)	Gross revenue (M)	Net op. revenue (M)	DOC@10% (M)
1	23,815	71.6	94.3	15.7	54.63	3,688	2,001	779
2	18,479	55.6	87.0	16.5	64.28	3,106	1,898	974
3	16,434	49.4	89.6	19.7	66.89	2,959	1,853	1,013
4	15,479	46.5	91.5	20.7	67.58	2,872	1,810	1,017
5	15,060	45.2	91.9	21.4	68.03	2,827	1,788	1,017

While iterations 3 through 5 are similar, iteration 4 was chosen as the final production schedule for the scenario as it had the highest DOC. Iteration 4 has a 30% higher DOC, 20% lower gross revenue and 7% less net revenue compared to iteration 1 which used a static break-even cut off and shut off of 25/t.

Figure 10 shows the ring selection created by the algorithm for iteration 4.

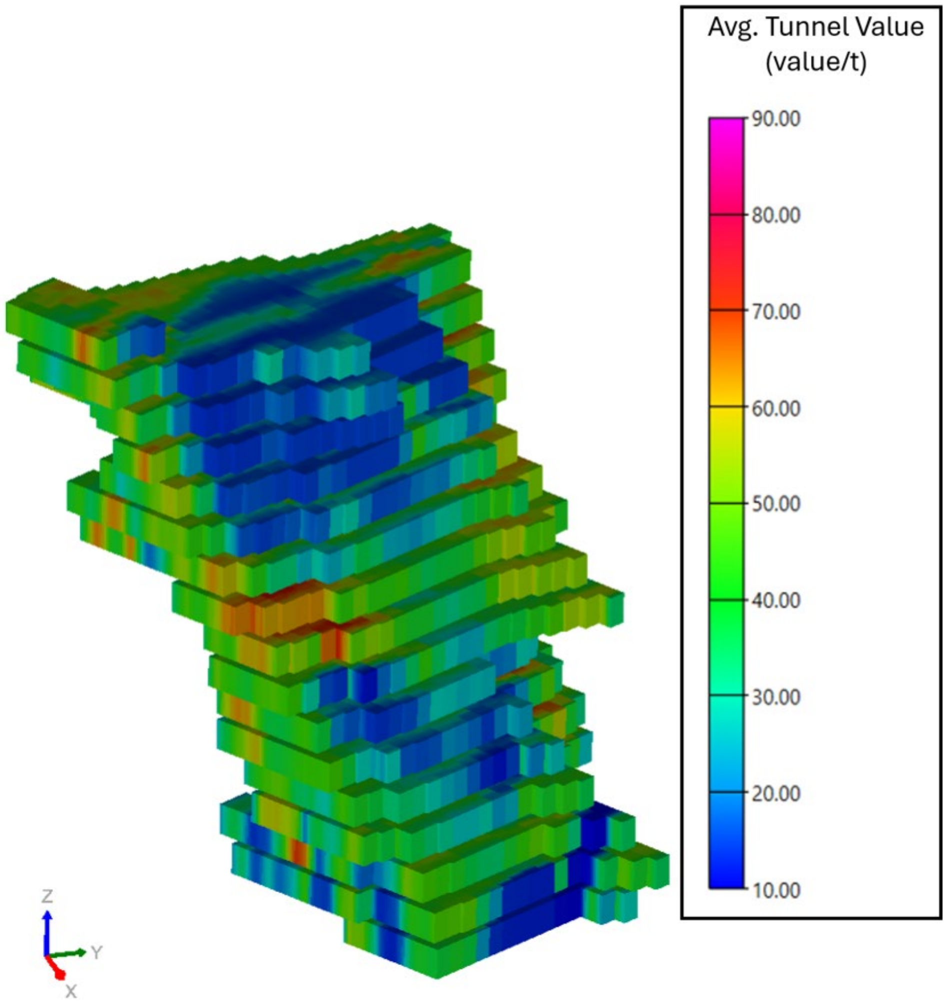

Figure 10 Ring selection of the scenario 2 iteration 4

Most of the rings removed in the process are from the upper levels of the deposit because these levels are mined early in the mine life when the MPC is highest. However, several rings remain in the hanging wall of the upper levels that are below the 25/t operating cost. This is due to a maximum step-out constraint that is configured in the software, among other constraints. The constraints are applied equally to all ring selections in the analysis. While the automated ring selection used by the algorithm does not produce perfect results, each is created using the same quantitative methodology and constraints, enabling better like-to-like comparisons versus ring selections produced manually. The results of the automated ring selection will vary by deposit but usually require some modification prior to being used in any final SLC design.

Figure 11 shows the variation in value per tonne of ore over the life of mine. Note the relatively high-value material at the end of iteration 1, which extends beyond the other iterations. It is by bringing this material forward through removing lower-grade early in the schedule that the process increases the DOC in subsequent iterations, despite a shorter mine life.

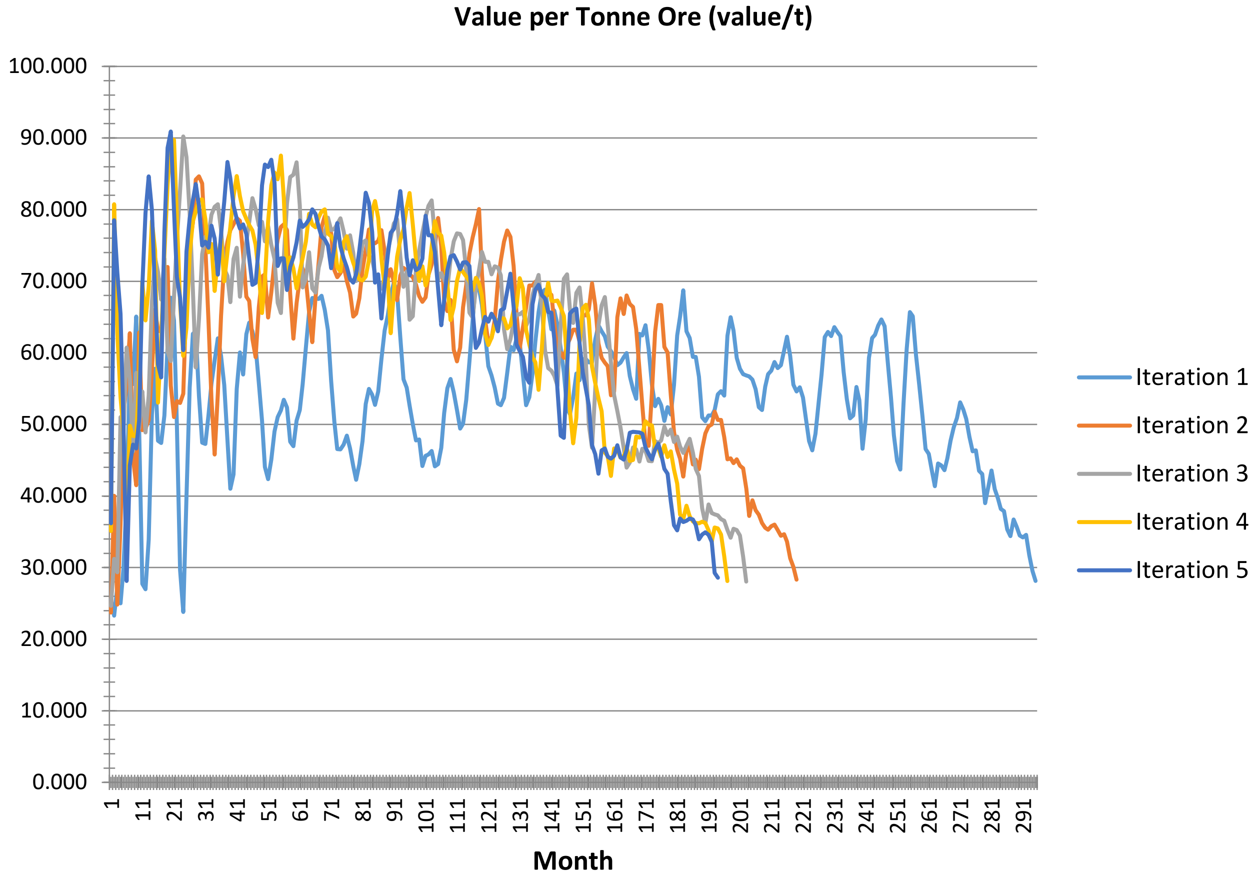

Figure 11 Value per tonne ore per iteration of scenario 2

Figure 12 shows the cumulative DOC for the 5 iterations of scenario 2.

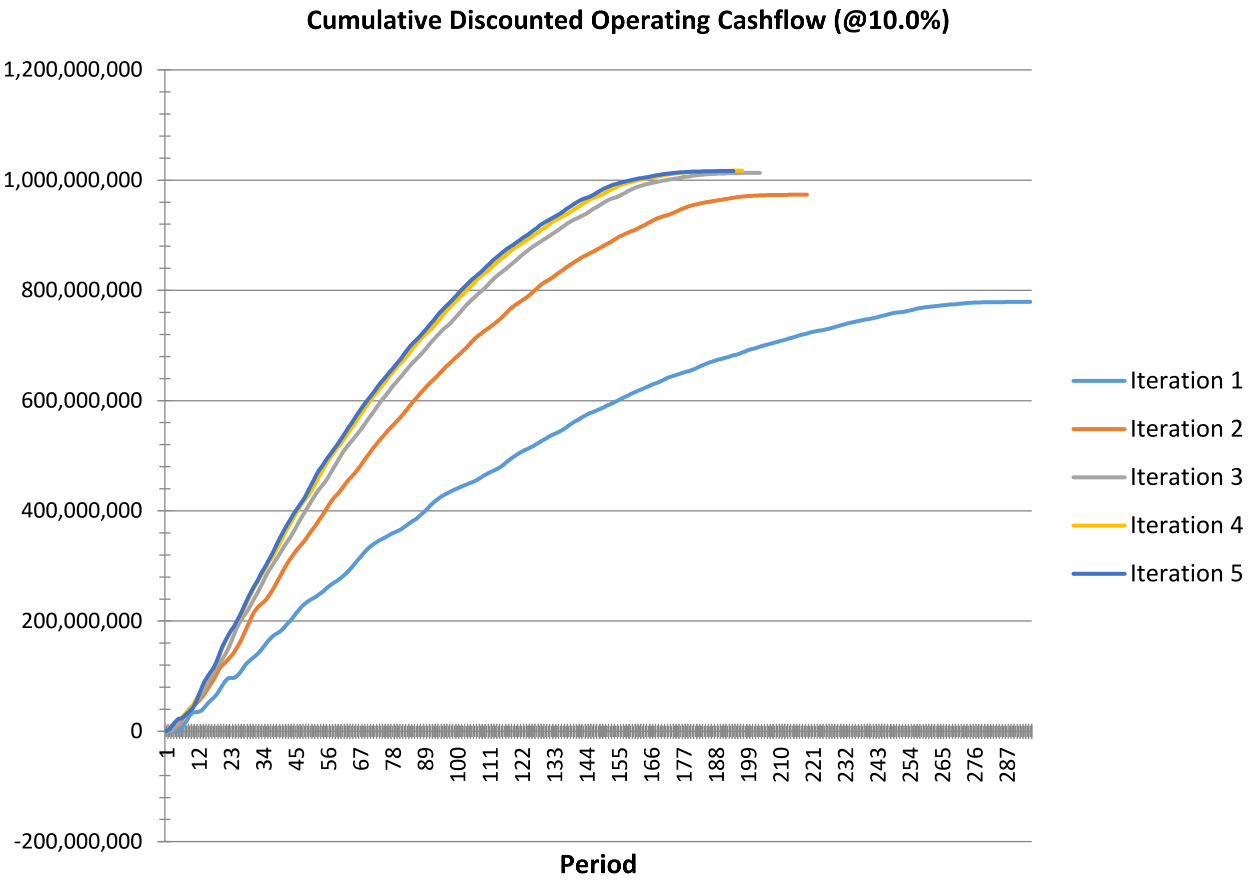

Figure 12 Cumulative discount sustaining operating cash flow per iteration for scenario 2

The total DOC increases significantly from iteration 1 to 2. The increase in DOC is noticeably smaller in going from iteration 2 to 3, but the slope of the cumulative DOC curve continues to increase, with iteration 3 reaching 95% of its peak value 15 months prior to iteration 2. As previously noted, iterations 3 and beyond are similar enough to be considered the same despite small fluctuations between iterations. The flat tail at the end of each curve is due to the MPC decreasing towards zero at the end of the mine life and the mining of low-grade marginal material. This contributes marginally to DOC but increases the overall utilisation of the resource.

4.4 Analysis results

Table 3 summarises the in situ and extracted tonnes and associated value for all 6 scenarios.

Table 3 Mined tonnes and value per scenario

Metric	Algorithmic		Static cut-off/shut-off			
	5%	10%	35	45	55	65
Number of rings	18,956	15,479	22,153	18,894	15,312	13,139
In situ mass (Mt)	56.9	46.5	66.6	56.9	46.1	39.6
In situ grade (value/t)	66.30	71.75	61.76	66.92	73.36	76.76
Extracted mass (Mt)	52.3	42.5	58.4	46.6	36.3	28.5
Extraction (%)	92	91	88	82	79	72
Extracted grade (value/t)	62.49	67.59	59.59	65.70	72.68	77.96
Gross revenue (B)	3.27	2.87	3.48	3.06	2.64	2.22
Net op. revenue (B)	1.96	1.81	2.02	1.90	1.73	1.51

Table 4 compares the DOC of the static cut-off/shut-off production schedules to the algorithmically produced schedule for both discount rates.

Table 4 Discounted operating cash flow per discount rate per scenario (values in billions)

Discount rate (%)	Algorithmic	Static cut off/shut off			
		35	45	55	65
5.0	1.33	1.26	1.29	1.27	1.18
10.0	1.02	0.86	0.93	0.98	0.94

For both discount rates, the algorithmically generated production schedule produced the highest DOC when evaluated at the corresponding discount rate.

However, the best static cut-off/shut-off scenarios produced DOC values within 2.9 and 4.0% of the algorithmic results for 5 and 10% discount rates, respectively. While these differences are small, the algorithmic scenarios also achieved 7 and 8% more net operating revenue on an undiscounted basis compared to the same static cut-off/shut-off scenarios.

For this example analysis, the algorithm has identified a solution only marginally better than the best-case static cut-off/shut-off scenario, but it has identified the solution in a deterministic manner, compared to the trial-and-error approach typically used in cut-off and shut-off optimisation.

5 Conclusion

This paper presents a methodology for generating SLC footprints and draw strategies that incorporate the time value of money directly into inventory selection and scheduling. While the methodology often produces discounted operating cash flow (DOC) values comparable to those obtained through detailed hill of value (HoV) analysis, it does so in a deterministic manner, enabling the efficient identification of high-quality production strategies without an extensive trial-and-error approach. In addition, compared to static elevated cut-off or shut-off strategies that achieve similar DOC or net present value outcomes, the dynamic minimum profit curve formulation can improve resource utilisation by allowing lower-margin material to be mined later in the schedule rather than being permanently excluded.

For a given set of economic and scheduling assumptions, the methodology produces a single production strategy corresponding to a specified discount rate. In this context, the discount rate can be viewed as an explicit parameter controlling the balance between competing investment priorities. Lower discount rates favour greater resource extraction, while higher discount rates prioritise higher-margin material earlier in the mine life.

The results obtained using this methodology depend on orebody geometry, grade distribution, and the mining sequence and rate. Where higher-value material is concentrated later in the sequence, the methodology is expected to increase DOC relative to low-margin static cut-off strategies. Conversely, where similar DOC outcomes can be achieved using a high-margin static cut off, the methodology instead improves resource utilisation by enabling lower-margin material to be mined later in the schedule. This effect is most pronounced where lower-grade material underlies a high-value core or where the orebody geometry supports efficient recovery of material left behind in earlier mining stages. For some orebodies, however, the methodology may produce ring selections and production schedules that do not differ materially from those generated using a well-defined static cut off approach.

This methodology is not intended to replace conventional HoV or sensitivity analyses. Rather, it provides a systematic means of generating alternative cut-off and shut-off strategies that can be incorporated into broader scenario evaluation. In this context, the methodology can help guide the range of economic parameters investigated in HoV analyses while reducing reliance on purely trial-and-error approaches.

Acknowledgement

The author would like to acknowledge SRK Consulting for providing support and encouragement to pursue this work, and Neil Winkelmann for his inspiration and valuable insights through numerous discussions.

References

Power, G 2004, 'Sub-level caving (SLC) draw control', *Proceedings of MassMin 2004,* pp. 235–240.