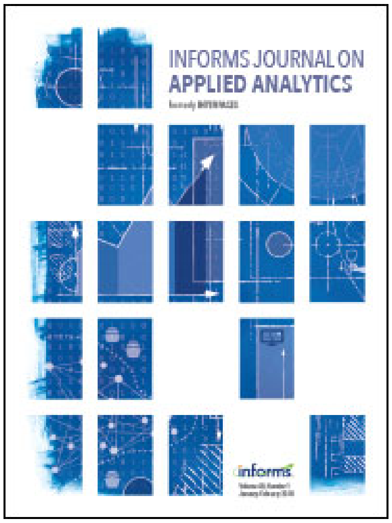

INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Inventory Management Using a Weekly Review (s, S) Policy at the Bank of Canada

Antoine Legrain, Johnathan Patrick

informs
http://pubsonline.informs.org/journal/inte
INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 52, No. 2, March–April 2022, pp. 210–225
ISSN 2644-0865 (print), ISSN 2644-0873 (online)

Inventory Management Using a Weekly Review (*s*, *S*) Policy at the Bank of Canada

Antoine Legrain,[a,*] Johnathan Patrick[b]

[a] Polytechnique Montreal, Centre-ville, Montréal, Québec H3C 3A7, Canada; [b] Telfer School of Management, Ottawa, Ontario K1N 6N5, Canada
*Corresponding author

Contact: antoine.legrain@polymtl.ca, https://orcid.org/0000-0003-2903-9593 (AL); patrick@telfer.uottawa.ca, https://orcid.org/0000-0003-1364-1803 (JP)

Received: February 28, 2020
Revised: October 13, 2020; November 24, 2020
Accepted: December 4, 2020
Published Online in Articles in Advance: November 17, 2021

https://doi.org/10.1287/inte.2021.1089

Abstract. One of the roles of the Bank of Canada (BoC) is to ensure that there is sufficient inventory of each denomination of currency all across the country. To accomplish this, the BoC has established 43 regional distribution centers (RDCs) to store inventory. Each RDC has its own limit in terms of the total value of currency that can be stored without penalty. Local banks can request withdrawals from the RDCs or else can choose to deposit money into the RDCs. In addition, the BoC often has to recall currency that has become unfit for service. Thus, the BoC is faced with a two-way inventory management problem wherein demand can be both positive (withdrawals) or negative (deposits), requiring the BoC to guard against both too little inventory (leading to shortages) and too much inventory (leading to penalties for exceeding value limit). We developed an adapted weekly review (*s*, *S*) policy for the management of the inventory at each RDC that determines the weekly decision as to how much inventory ships both to and from each RDC. The BoC implemented the adapted policy through the first half of 2018. Comparisons with the previous two years demonstrated a drop in transportation costs of around 15% in the 12 months following implementation compared with the average of the two years prior—despite the fact that costs had been increasing prior to implementation from one year to the next.

History: This paper was refereed.

Keywords: banks • financial institutions • inventory • production • heuristics • networks • graphs • simulation • transportation

Introduction

Sometimes the greatest compliment is to go unnoticed. If one were to take a moment to think about it, one would recognize that there are a multitude of services provided to us on a day-to-day basis that we largely take for granted unless they malfunction. One such service is the provision of cash when we want it and in the note denominations we prefer. How that is accomplished remains a matter of little concern to the user, and yet the management of such a network of note distribution across an entire country is by no means without its challenges.

Central banks across the world face the daunting task of trying to maintain sufficient inventory of bank notes of each denomination in order to meet a highly variable and geographically dispersed demand. This is accomplished in part by creating depositories (usually in partnership with local banks) across the country in order to maintain readily available cash in each region. The unpredictability of demand is exacerbated by the fact that demand from the central bank depositories comes not directly from consumers but from private banks who carry their own inventory and tend to deposit into or withdraw from the central bank in large lump sums. Large withdrawals need to be satisfied from existing inventory while still maintaining sufficient capacity in a given depository in order to accommodate a large deposit.

Thus, central banks face a two-way inventory management problem with highly unpredictable, large-sum, positive (withdrawal) and negative (deposit) demand. They must ship notes from central production centers to capacitated regional distribution points (RDCs). They must also ship notes back to the central production centers, either because of insufficient capacity in a region's depository or because the notes themselves have become unfit for circulation. Unlike other perishable goods, central banks cannot simply dispose of unfit notes but must collect, count, and verify them in order to protect against counterfeit production.

Finally, transportation of bank notes between the central production center and the regional distribution points is costly because of the significant precautions required that are due to the value of the shipments. This leads to quite complicated costing arrangements with armored car carrier companies that often depend on the volume, weight, and cash value of the shipment.

It also means that suboptimal inventory management policies can quite easily incur significant costs.

This paper discusses the optimization of shipment decisions by a central bank in charge of managing bank note distribution across a wide geographic region with multiple regional distribution points. We discuss the unique challenges of this inventory management problem and propose a modified (s, S) policy as a means of managing bank note distribution efficiently. We demonstrate the superiority of this policy to the current practice of the Bank of Canada (BoC) through a simulation.

Literature Review

Inventory management with returns has been studied extensively in the operations research literature (Toktay et al. 1997, Dowlatshahi 2000, Fleischmann et al. 2003, DeCroix and Zipkin 2005, Mitra 2012). Within this larger literature is a smaller subset that focuses on the specific inventory management challenges faced by central banks such as the BoC. For a good overview of the challenges of managing a currency supply chain from the perspective of multiple stakeholders, we refer the reader to the work of Geismar et al. (2017).

Massoud (2005) provides a detailed dynamic model for a central bank's production and inventory management problem. The model provides a production schedule based on demand and the rate at which notes become unfit for circulation. However, it considers the inventory management problem only at the level of the country rather than the region and thus does not consider the optimization of the transportation of notes from the central bank's main holding centers to the regional distribution points. These transportation costs are assumed to be fixed and given.

Geismar et al. (2007) deal with the same inventory management problem as our paper but from the perspective of the banks thus determining a policy to help them better manage their deposits and withdrawals from the central bank (Federal Reserve in the United States). Their model deals specifically with new guidelines that were recently introduced by the Federal Reserve in the United States at the writing of the paper. In a similar vein, Mehrotra et al. (2010) also examine the same set of guidelines from the perspective of the individual banks. Their aim is to provide banks with a decision-making tool to best determine their operating policies (from a fixed set of potential policies) in light of the new regulations.

Chen and Simchi-Levi (2009) provide a sophisticated theoretical model for the stochastic cash balance problem. This is an inventory management problem similar to that encountered by central banks in that demand can be both positive and negative and stock can be either shipped out or ordered in. Although clearly related, the model does not deal with the type of multisite inventory management problem faced by central banks with multiple regional distribution points.

Zhu et al. (2015) model a transportation problem for a head office shipping cash to multiple regional branches. Their model determines the optimal number of trucks as well as the optimal transportation schedule for each truck, including decoy trucks (used for security purposes). Unlike a typical central bank challenge, their problem deals with a relatively local problem where shipments are made only by truck and on a daily basis. The cost structure is thus much simpler. The authors derive a mixed-integer programming (MIP) model and a subsequent heuristic policy that can be easily implemented.

Finally, Huang et al. (2017) provide a MIP model for a country's currency supply network. They break the problem down into two subproblems. The first subproblem optimizes the capacity of branches within the network. The second subproblem seeks to optimize currency flow. This is undoubtedly the closest model to the setting we were presented with through our partnership with the BoC. However, their model optimizes the flow (which regional vault is served by which central vault) rather than how to determine the frequency and size of shipments.

This paper describes a research partnership with the Bank of Canada where we look to improve on their current inventory management policy without creating significant disruptions to current practice. We first provide an overview of the BoC operations culminating in a deterministic model of the BoC inventory management problem. We next introduce the impact of demand uncertainty and describe how the BoC currently tackles this issue. We propose an adapted weekly review (s, S) policy that we fine-tune based on a simulation using historical demand. A more analytical approach is simply not available given the complexity of the costing function. We demonstrate the significant savings the weekly review (s, S) policy provides compared with current practice. Finally, we present the results and implementation challenges following 15 months of use by the BoC.

BoC Network and Operations

The Bank of Canada manages the bank note holdings of 43 regional district centers (RDCs) gathered within 10 RDPs distributed across Canada. The RDCs are operated by nine different financial institutions (FIs) that participate in the Bank Network Distribution System (BNDS). The rationale for the BoC network distribution (illustrated in Figure 1) is to have the right quantity of bank notes at the right time and in the right place. This is achieved through the shipment of bank notes from an agency operation center (AOC) to each RDC (referred to as *replenishments*), the return of unfit

Figure 1. (Color online) The Bank Note Distribution Network

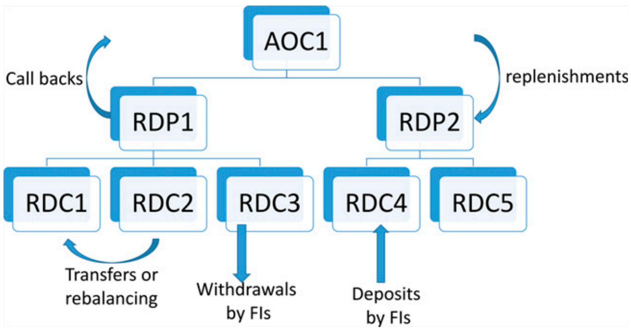

or excess bank notes from the RDCs to the AOCs (referred to as *callbacks*), or through transfers of bank notes between RDCs within the same RDP (referred to as *rebalancing and transfers*).

Typically, for each denomination, an RDC will have a "withdrawer" or "depositor" profile, depending on whether FIs generally deposit or withdraw bank notes of that denomination from the RDC.

Objectives

There are two main performance objectives that the BoC seeks to achieve while managing the holdings of each RDC. The main objective is to avoid shortages within each RDP, whereas the secondary goal seeks to minimize the number of overcaps at any RDC. A shortage means that there are insufficient bank notes for a 24-hour period to meet the demand of at least one denomination. An overcap occurs if the total dollar value of holdings in an RDC exceeds a predefined threshold based on insurance contracts. If an overcap occurs, then the FI is not allowed to deposit into the BoC reserve, and, therefore, the BoC is forced to pay a fee to compensate for its interest revenue loss.

Given these performance objectives, the goal of the BoC is to determine an inventory management policy that achieves satisfactory performance at minimal cost. Costs are primarily the result of shipment fees. Transportation is done by air (at a higher price) or by ground based on the location of the RDP being served. The total dollar value of a shipment is the main pricing component, with different charges in effect for various ranges of the dollar value of the shipment. However, the price is also dependent on the size of the shipment—either the total weight and number of bags for air transportation or the number of containers/bags required for ground transportation. Finally, if a replenishment and callback is performed in the same week for the same RDP, then there is a discount applied to the cost.

Actions

The replenishments and callbacks are planned weekly by the BoC to ensure that there is sufficient supply for all RDCs while not exceeding a maximum dollar value or space constraints. This inventory management policy can be summarized as a weekly review system (Hopp and Spearman 2011). Supply is divided by denominations ($100, $50, $20, $10, and $5 bills) and type (new, fit, and unfit).

Typically, each RDC will manage its own requests for withdrawals and store any deposits. However, despite the best efforts of the BoC, the stochastic nature of deposits and withdrawals by the FIs may lead to unexpected shortfalls or surpluses at a given RDC. A number of options are available in this case:

1. Fit demand: If there are not enough fit notes to meet the demand, then the BoC will use new notes from the same RDC to compensate for the shortfall. (Note that the reverse—using fit notes to meet demand for new notes—is not possible.)
2. Rebalancing: Once a week, surplus notes at one RDC may be spread to other RDCs within the same RDP that are in need of these same notes.
3. Transfers: If, between the weekly rebalancing, there is a risk of an overcap, then the surplus notes can be moved to another RDC where there is either a need for these notes or simply no risk of an overcap. This is essentially an unscheduled rebalancing that has cost implications.
4. Withdrawal fulfillment from another RDC: If an RDC does not have enough notes (fit or new) to meet the demand, then the request will be fulfilled, if possible, through holdings at another RDC within the same RDP.
5. Night courier: If all the previous mechanisms do not prevent a shortage, then the BoC may use a night courier to replenish the RDC at significantly raised rates.

The costs of most of these options are covered by the BoC (except for the rebalancing and the withdrawal fulfillment) and thus should be avoided when possible. The rebalancing and transfer movements are planned one day before, whereas the other three options are enacted as needed.

Mathematical Formulation

Formulations (1)–(8) summarize the BoC inventory management problem for a given RDP over a fixed horizon. This deterministic formulation does not take into account the dynamics of the problem but serves to illustrate the nature of the optimization challenge. It is worth noting that Equation (2) ensures that demand is met. See the appendix for a more detailed mathematical formulation.

$$\begin{aligned}\min\ &\text{Sum of transportation costs for replenishments/}\\&\text{callbacks, transfers,}\\&\text{and night couriers} + \text{new notes costs} + \text{overcap fees}\end{aligned} \tag{1}$$

subject to:

Conservation of each RDC inventory from day to day (2)

Tracking of any overcap (3)

Shortage interdiction (4)

Meeting the demand for fit notes with new notes (5)

Interdiction of shipping unfit notes to the RDCs (6)

Authorizing replenishments/callbacks only on the scheduled days (7)

Ensuring the positivity of the variables (8)

Demand Uncertainty

The aforementioned model provides a deterministic representation of the inventory management decisions of the BoC but ignores the significant impact of demand variability. To plan their decisions in the face of this uncertainty, the BoC uses a demand forecast black-box software that implements a time-series model using moving weighted averages of the same day (i.e., the second Monday of July) from the past four years. Despite its simplicity, this in-house forecasting software is a key element of the current decision process; its code has been developed by contractors, and, unfortunately, it does not belong to the BoC. Although it is possible that it could be improved with a deeper study of past data, a lengthy analysis (that included Arima models, dynamic harmonic regression, hierarchical forecasting, and multilayer perceptron models) failed to produce significant improvement, primarily because past data on demand fit no discernible distribution, consisting, as it does, of multiple weeks with zero activity interspersed with high spikes of withdrawals/deposits. This is likely the result of banks themselves having some holdings leading to an accumulation of demand prior to accessing the repositories of the BoC. This lack of a readily available distribution for demand makes most of the models in the literature difficult to apply.

Planning in the Face of Uncertainty at the RDP Level

The BoC and the FIs have agreed on volume inventory targets (S) for each RDC and for each bank note to ease the management of the inventory. These targets are used for the weekly review system and are revised twice a year to take into account modifications in demand patterns and trends but are otherwise untouched. The multiday lag time between when a shipment is planned and when it is received creates significant challenges to planning shipments.

Based on current practice, the replenishments seek to ensure that each RDP reaches the expected target inventory for each denomination for both fit and new notes *on the delivery date.* The anticipated inventory without replenishment on the day of delivery is estimated by the forecasting software. Thus the replenishment planning decision process can be summarized in three parts. First, the cumulative demand of both fit and new bank notes between the planning day and the delivery date is estimated for each denomination at each RDC. Next, sufficient bank notes are shipped so that the expected holdings on the delivery date are equal to the RDP targets (equal to the sum of the RDC targets within that RDP). Finally, upon arrival of the shipment, the bank notes are spread to each RDC, allocating any shortfall or excess evenly between all RDCs (a shortfall or an excess occurs if the actual withdrawals or deposits within the RDP do not follow the forecast for a given denomination).

The callbacks remove any unfit bank notes on the planning day as well as any expected surpluses within the RDP. A surplus occurs if the expected total value of the fit and new notes in inventory on the shipping day is sufficiently close to the maximum value limit (typically within 10%) to pose a risk of an overcap. Note that, when making this decision, the replenishment is added to the inventory, as usually the shipment is en route when planning the callback.

Table 1 illustrates the rules that the BoC uses for planning replenishments and callbacks for a fictitious RDP that has two RDCs and two denominations ($5 and $10). It provides the number of notes (in thousands) in the RDC *on the planning date* along with the *forecasted* number of notes in the RDC *on the delivery date* before and after the delivery, and the targets. The bold numbers show the holdings that have been impacted by the planning decisions. For the replenishment and callback decisions, all that matters is the total forecasted holdings for each bank note type and denomination on the day of delivery. Following their rules, the BoC would plan the following replenishment and callback:

- There are no $5 notes sent, as the *RDP* inventory levels are anticipated to be at or above the targets on the delivery date.
- There are 5,000 fit and 10,000 new $10 bank notes sent.
- There are 40,000 unfit bank notes sent back to the AOC (the total value of unfit notes).
- There are no other notes called back as the small surplus of $5 notes in RDC 2 is likely insufficient to pose a risk of an overcap: $570,000 < 585,000 = 0.9 * 650,000$.

Finally, the surplus of $5 notes in RDC 2 (if it still exists once demand is realized) will be transferred to RDC 1 during the rebalancing, provided that RDC 1 has a "withdrawer" profile for $5 notes.

Table 1. Actual Holdings on the Planning Day, Forecasted Holdings on the Delivery Day Before and After the Delivery, and the Target Inventories for Each RDC and for the Whole RDP

	RDC 1 holdings				RDC 2 holdings				RDP holdings			
		Forecast on delivery day				Forecast on delivery day				Forecast on delivery day		
	Actual	Before	After	Target or max value	Actual	Before	After	Target or max value	Actual	Before	After	Target or max value
$5 fit	5,000	3,000	3,000	10,000	25,000	30,000	30,000	20,000	30,000	33,000	33,000	30,000
$5 new	2,000	2,000	2,000	1,000	4,000	4,000	4,000	3,000	6,000	6,000	6,000	4,000
$10 fit	15,000	10,000	10,000	10,000	20,000	25,000	**30,000**	30,000	35,000	35,000	**40,000**	40,000
$10 new	1,000	2,000	**5,000**	5,000	3,000	3,000	**10,000**	10,000	4,000	5,000	**15,000**	15,000
Unfit $	$10,000	$20,000	**$0**	$0	$20,000	$20,000	**$0**	$0	$30,000	$40,000	**$0**	$0
Total $	$205,000	$165,000	**$175,000**	$220,000	$395,000	$470,000	**$570,000**	$650,000	$600,000	$635,000	**$745,000**	$870,000

The rules utilized by the BoC successfully avoid undue shortages and overcaps while respecting operational constraints. However, the weekly replenishments and callbacks are very expensive, especially by air. We thus sought to find a means of modifying the frequency of these shipments so as to more efficiently meet demand while providing the same level of service. The proposed modification is explained in the following section.

The Weekly Review (*s*, *S*) Policy

Because a major component of the inventory management expenses depends on the number of replenishments and callbacks performed, it is important to examine any possible means of decreasing these expenses. The current rules act like a weekly review base stock policy. Such a policy is rarely optimal and is most often outperformed by an (*s*, *S*) policy (Johnson 1968, Zheng and Federgruen 1991, Feng and Xiao 2000).

The Definition of a Weekly Review (*s*, *S*) Policy

The (*s*, *S*) policy is extensively used in inventory management and has been proved to be optimal for a wide range of situations, especially when items are often withdrawn (see the aforementioned references). The main idea is to control the inventory level of each item with two triggers: a minimum level, *s*, and a maximum level, *S*. When the inventory level falls below *s*, it triggers a replenishment action to restock this inventory to the level *S*. A weekly review (*s*, *S*) policy controls the inventory only once a week instead of at any time. Figure 2 presents a weekly review (*s*, *S*) policy for the replenishment of new $20 bank notes. Two triggers are defined, with *s* equal to 30,000 notes and *S* equal to 70,000 notes. The first time that a replenishment is planned occurs on April 13th, as the inventory level is below *s*. This replenishment arrives on April 14th and restocks the holdings to a level of 70,000 new $20 notes. However, on April 20th, no replenishment is planned as, despite some withdrawals, the current holdings are still over 30,000 notes.

In contrast, the current BoC practice works also with a weekly review system but uses a single trigger so that replenishments are made when the inventory falls below the trigger and inventory is restocked up to that same trigger value. In practice, this results in a weekly replenishment and callback, as there is invariably at least one RDC and denomination with inventory below the target within the RDP. There is also invariably some unfit notes within the RDP. The weekly review (*s*, *S*) policy has the potential to reduce the number of replenishments and callbacks compared with the current BoC policy as it does in the aforementioned example.

Figure 2. (Color online) A Weekly Review (s, S) Policy for the Replenishment of the New $20 Bank Notes

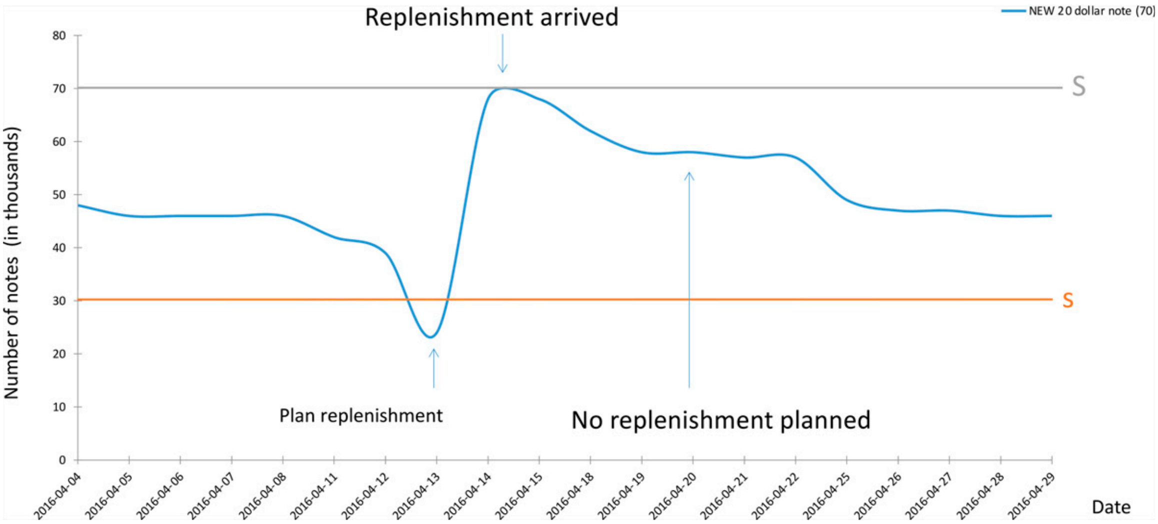

A Weekly Review (*s*, *S*) Policy for the BoC

The BoC inventory management problem has a number of characteristics that set it apart from other settings and that have mostly been mentioned earlier. We summarize these differences below:

- Inventory can increase as well as decrease stochastically as FIs deposit and withdraw bank notes.
- Inventory is shipped both ways through replenishments and callbacks.
- Some inventory items are interchangeable, as an RDC can use new notes to meet demand for fit notes.
- Not only are shortages a problem, but overcaps also incur additional costs.
- There are price discounts for combining a callback and a replenishment shipment.
- Callbacks can be triggered either by a large amount of unfit notes in an RDP or when the total holdings in an RDP approaches the overcap.

There is not, to the best of our knowledge, a method within the literature to optimally choose s and S in such a setting. Moreover the targets (S) cannot be changed without discussion with the FIs. From the perspective of the BoC, it might very well be advantageous to set higher targets, but this may not suit the FIs, which may balk at holding larger inventories. Thus, for the initial project, it was determined that it was best to keep the targets at their current levels. The values of s, however, are entirely within the control of the BoC. For both these reasons, the BoC needs an (s, S) policy adapted to its setting that incorporates both replenishments and callbacks. Mathematical details are in the appendix.

A Weekly Review (*s*, *S*) Policy for Replenishments

The (s, S) policy for replenishments needs to control the level of both fit and new notes (i.e., the notes that can be replenished). As a request for fit notes can be filled with new notes, the BoC needs to track the sum of the inventory level for fit and new notes combined, instead of concentrating on fit notes by themselves; but the BoC also needs to track new notes on their own, as demand for new notes cannot be met with fit notes. Furthermore, as notes can be transferred within the RDP at minimal cost, the (s, S) policy needs only to control the inventory levels for the whole RDP. Finally, the BoC is currently using the forecast to determine the expected holdings on the replenishment delivery date. The (s, S) policy can make use of these forecasts in much the same way by triggering a replenishment only if the *expected* holdings on the delivery date are below the trigger value, s. We thus proposed an (s, S) policy for the whole RDP for each denomination, for the combination of fit and new notes, and for new notes on their own. This process is defined as follows:

- If the expected holdings of fit plus new notes for any denomination are below a trigger value s *or* if the expected holdings of new notes for any denomination are below a different trigger value s, then a replenishment is planned so that the expected inventory within the RDP on the delivery date is equal to the target for any notes.
- Otherwise, no shipment is made.

Note that a deficit in any denomination triggers a shipment that includes all denominations below their respective targets.

A Weekly Review (*s*, *S*) Policy for Callbacks

The callbacks are used to control the total value of all holdings in the RDP (region) so as to avoid overcaps. In this case, the (s, S) policy depends on the total dollar value at each RDC and the dollar value of unfit notes within an RDP:

Table 2. Forecasted Holdings on the Delivery Day

	RDC 1		RDC 2		RDP	
	Notes	Target	Notes	Target	Notes	Target
$5 fit note	3,000	10,000	30,000	20,000	33,000	30,000
$5 new note	2,000	1,000	4,000	3,000	6,000	4,000
$10 fit note	10,000	10,000	25,000	30,000	35,000	40,000
$10 new note	2,000	5,000	3,000	10,000	5,000	15,000
Unfit note	$20,000	$0	$20,000	$0	$40,000	$0
Total value	$165,000	$205,000	$470,000	$515,000	$635,000	$720,000

1. If the expected dollar value of holdings of unfit notes in an RDP is over a trigger value s, a callback is planned, and all unfit notes in the entire RDP are shipped back to the AOC. Otherwise, nothing is done.

2. If the expected total dollar value in a RDC is over a certain minimum value (i.e., a fraction of the maximum inventory value), a callback is planned as follows: all unfit notes within the RDP are shipped back as well as all fit notes over the targets for which the RDC in question has a depositor status. Otherwise, nothing is done.

Finally, because of the cost discount when a callback is coupled with a replenishment, it is worthwhile to have different (and lower) trigger values in cases in which a replenishment has occurred than for cases in which the callback would be done without a replenishment.

Proposed Process

The current targets used by the BoC determine the size of replenishment and are based on an agreement between the BoC and the FIs. Although there is no guarantee that these targets are optimal, we were encouraged (as mentioned earlier) to derive an (s, S) policy using the same targets. This was deemed to greatly ease the challenge of implementation. Thus the major difference between current practice and the (s, S) policy is the result of choosing the triggers that determine whether a replenishment or a callback is indeed needed. For the replenishments, the triggers s are set as a percentage of the inventory targets S. Once the targets are decided, the trigger values are thus easy to compute. For the callbacks, the triggers s are set as a percentage of the total value limit. Computation of the optimal values of the targets would ideally be based on past data and current trends in demand. However, this is not easily done, as the demand does not follow any well-known distribution. However, a more evidence-based approach to determining the targets could potentially yield additional benefits. This would need to be done in partnership with the FIs, as they are averse to carrying larger inventories and would also need to provide some insight into what drives their withdrawal and deposit decisions so that demand can be more easily predicted.

Example

The example used earlier to illustrate the BoC rules is reused here to demonstrate how a weekly review (s, S) policy will affect the current process.

Table 2 recaps the data from the previous example. Table 3 proposes different versions of the (s, S) policy that lead to different planning decisions. In particular, we consider two options for the trigger values for replenishments: one in which they are 30% of the value of the inventory targets S and another in which they are 50% of these targets.

For scenario 1, all the holdings are over the targets, so no replenishment is planned. However, for scenario 2, the number of $10 new notes is below $s = 7{,}500$, so a replenishment is planned, as illustrated in Table 1. Note that even though it was only the inventory level of new notes that triggered the decision to initiate a replenishment, once a replenishment is initiated, the shipment is such that *all* denominations and bank note types are brought up to the targets just as would have been done under the rules described earlier.

Table 3. Description of the Two (s, S) Policy Scenarios for the Replenishment

	RDP		Scenario 1		Scenario 2	
	Notes	Target	%	s	%	s
$5 fit + new note	39,000	34,000	30%	10,000	50%	17,000
$5 new note	6,000	4,000	30%	1,000	50%	2,000
$10 fit + new note	40,000	55,000	30%	16,500	50%	27,500
$10 new note	5,000	15,000	30%	4,500	50%	7,500

Table 4. Description of the Two (s, S) Policy Scenarios for the Callbacks

	Holdings		Scenario 1		Scenario 2	
	Unfit value	Value limit	%	s	%	s
RDP	$40,000	$780,000	5%	39,000	10%	78,000
	Total value	Value limit	%	s	%	s
RDC 1	$165,000	$230,000	80%	184,000	90%	207,000
RDC 2	$470,000	$550,000	80%	440,000	90%	495,000

Thus a key advantage of the proposed modifed (s, S) policy is that it provides a means of reducing the number of replenishments *without modifying the way replenishments are planned*.

Table 4 illustrates the planning of callbacks. Recall that these triggers are computed as a percentage of the total value limit. Note that, for callbacks, a large total dollar value of unfit notes or a high total dollar value of holdings is detrimental, and thus values greater than the trigger values result in a callback.

For scenario 1, the unfit value in the RDP and the total value in RDC 2 are over the trigger value s, so a callback is planned, whereas for scenario 2, all the values are below the triggers, so no callback is planned. For scenario 1, the planned callback will send $30,000 unfit notes (i.e., the *current* value of unfit notes in the RDP: 10,000 from RDC 1 and 20,000 from RDC 2). With this callback, the *expected* total value at RDC 2 drops to $450,000, which remains over the trigger value (i.e., $440,000). Thus, the callback will also include the *current* holdings at RDC 2 that are above the target and for which RDC 2 has a depositor profile (in this fictitious example). In this case, this means that 5,000 $5 fit notes are shipped back in order to reduce total holdings in RDC 2. The $5 fit notes are chosen because their current holdings are over the target (being, respectively, 25,000 and 20,000). Note that, even though the $5 new notes are also over the target, these are not sent back, as RDCs have a withdrawer profile for all new notes.

The aforementioned example demonstrates how the weekly review (s, S) policy works but does not describe how, in practice, one chooses the trigger values. This we do in the next section.

Simulation and Experiments

A discrete event simulation has been developed to measure the potential cost savings of the adapted weekly review (s, S) policy. The Java programming language has been used to build the simulation package described in Algorithm 1 (see the appendix). All the operational processes described previously have been implemented in the simulation. Furthermore, the real transportation contracts, distribution networks, and targets have been used in order to provide as much realism as possible.

The confidential data provided by the BoC contains the description of the network (AOCs, RDPs, RDCs, and their hierarchical structure), the inventory levels, the demand, the pricing agreements, as well as the targets S used for both weekly review policies P (BoC current practice or (s, S)). For confidentiality reasons, the RDPs will be described by letters (i.e., A, B, etc.). The pricing agreements provide a thorough description of the different transportation costs. Based on the RDP, the replenishments and callbacks are sent from a given AOC and follow a piecewise cost structure that depends on the total value of the shipment, its weight, and its volume. The transfer price is composed of only a fixed cost. Finally, the night courier has a high cost per shipped bag of notes.

Algorithm 1 is used to analyze and compare both policies on real data sets. The first set of tests measures the benefits of a weekly review (s, S) policy as a function of the percentage chosen for the trigger value in an RDP. The second set of tests demonstrates the benefit of the best weekly review (s, S) policy (taken from the first set of tests) over 10 months of historical data.

Determining the Lower-Bound Trigger Value for Replenishments

In these tests, we vary the trigger values for replenishments from 10% to 90% of the target in order to analyze the performance of the corresponding weekly review (s, S) policy. For each scenario, the software runs 30 replications of the BoC operations over 60 days to ensure the robustness of the results to fluctuations in future demand, which is generated based on the historical data. As no distribution has been found for the demand, the average value for the same day of the two previous years has been uniformly perturbed between [−50%, 50%]. This distribution, although arbitrary, draws its mean from the forecasting software and adds some noise to ensure a certain robustness of the results.

For the callback policy, as there is a cost discount when planning a replenishment and a callback in the

Figure 3. (Color online) Expected Improvements in the Total Cost as a Function of the s/S Ratio

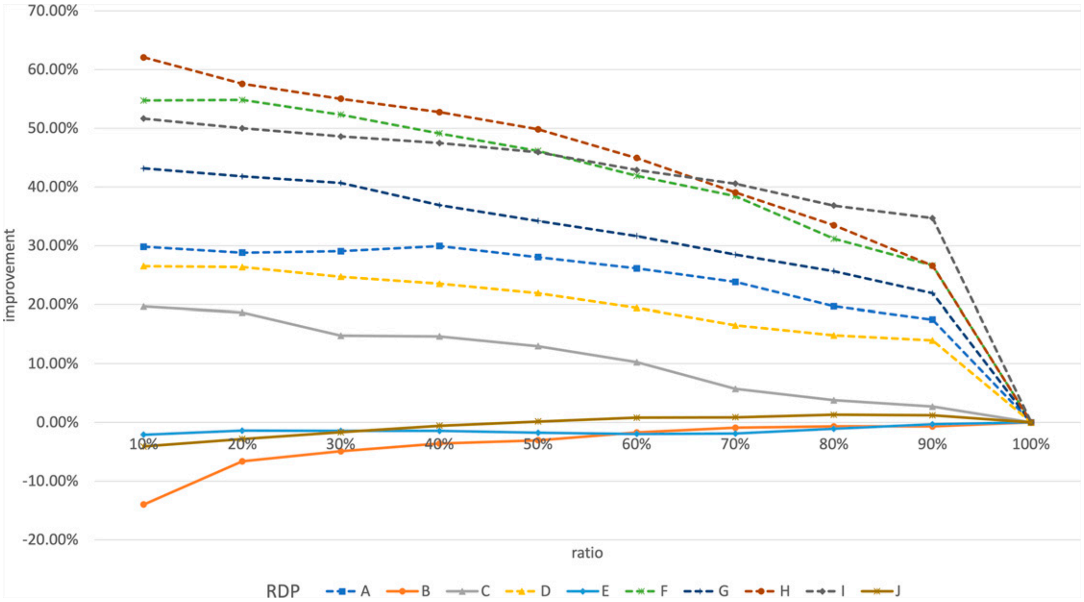

same week, the trigger s for the unfit notes and the minimum inventory value used in the simulation depend on whether there is a planned replenishment.

• If no replenishment is planned, then a callback is triggered if the total dollar value of unfit notes exceeds 10% of the total value limit.

• If a replenishment is planned, then a callback is triggered if the total dollar value of unfit notes exceeds 5% of the total value limit (in order to take advantage of the two-way shipment discount).

• A callback is also triggered if the total dollar value of all holdings exceeds 90% of the inventory value limit.

Later, we perform some sensitivity analyses around these two trigger values for the callback policy.

Figure 3 presents the simulated percentage of improvement in the total cost (averaged over the 30 replications) when using the weekly review (s, S) policy as opposed to the current BoC rules: RDPs served by air (A, D, F, G, H, and I) are represented by dashed lines and the ones served by air by plain lines. Note that the current BoC policy corresponds to a weekly review (s, S) policy with a ratio of 100%, namely, $s = S$. The total cost includes the shipment costs (i.e., the contract and size costs), the transfer costs, new note fees, and the overcap fees. As the trigger value s is reduced, there is a substantial decrease (up to 60%) in the total costs for those RDPs served by air (A, D, F, G, H, and I), whereas the RDPs served by ground transportation (the three bottom rows) appear to demonstrate little to no improvement. This is intuitively reasonable because of the very high fixed cost of air transportation that makes it cheaper to send larger replenishments and callbacks less frequently.

The impact of the weekly review (s, S) policy is to dramatically decrease the number of shipments and thus reduce the overall replenishment and callback costs. As these are the most significant costs for the BoC, their reduction while using the weekly review (s, S) policy account for the greater part of the improvement observed in the total costs shown in Figure 3. These improvements in the total costs, however, do not come without some potential sacrifice in terms of an increase in withdrawal fulfillments from another RDC within the same RDP and/or an increase in the risk of a shortage. Figure 4 demonstrates that, as the trigger value is reduced, the number of withdrawal fulfillments from another RDC increases. This is, again, intuitively reasonable because, as the number of shipments decreases, inventory levels at each RDC are allowed to decrease, and thus it may be more likely that a withdrawal request exceeds the inventory of a particular RDC, thus necessitating that the inventory at another RDC make up the shortfall. Although this undoubtedly places some additional burdens on the FIs, the increased cost of additional withdrawal fulfillments from another RDC is minimal in comparison with the cost reduction associated with fewer replenishments and callbacks. Figure 5 demonstrates the impact of reducing the trigger value on the number of potential shortages. Although there is an undeniable increase, it is also clear that a weekly review (s, S) policy can manage the increase in the number of potential shortages by keeping the ratio above a certain threshold. However, these last results need to be considered with caution, as this was the one aspect of the analysis where our

Figure 4. (Color online) Expected Number of Withdrawal Fulfillments from Another RDC as a Function of the s/S Ratio

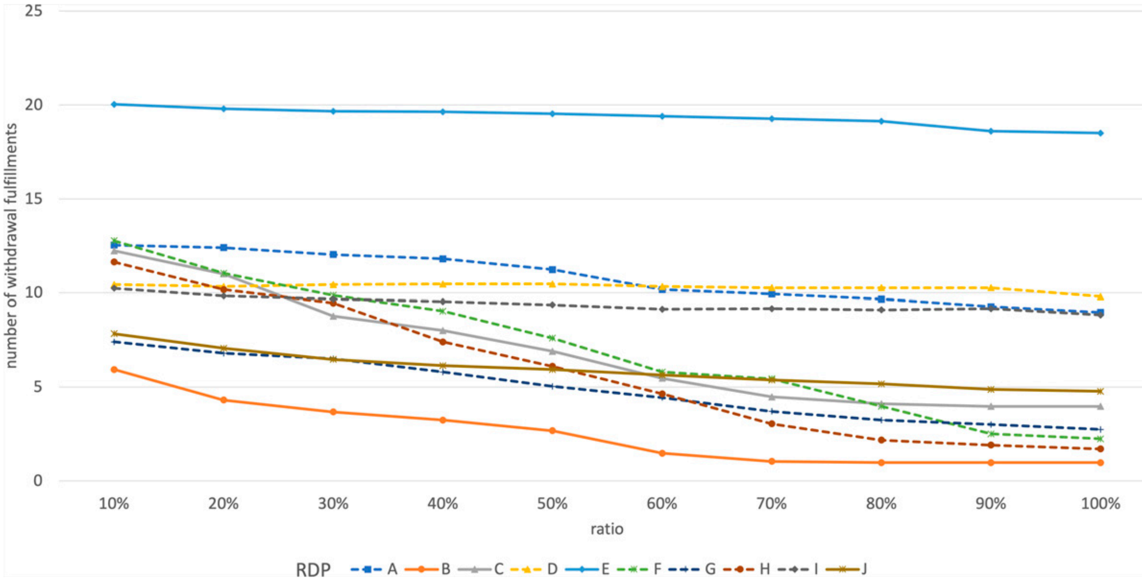

simulation of the BoC rules did not appear to match reality. One can see from Figure 5 that even in the case where the ratio is set to 100% (i.e., current practice) the expected number of shortages is nonzero, whereas, in actual fact, shortages are nonexistent under current practice. This is most likely explained by the current practice of rush shipments that can and do occur on occasion but that are not easily recognized in the data as such. These are used to avoid shortages and could continue to be used in the same fashion under the weekly review (s, S) policy. Thus, the number of shortages can be more helpfully interpreted as the number of rush (or overnight) shipments required, as these are the actions available to the BoC to prevent a shortage. The introduction of the weekly review (s, S) policy is likely to increase the frequency of rush shipments but not by an unmanageable amount.

Finally, our analysis showed no observable global trend for the overcap fees as a function of the ratio.

Table 5 shows the final ratios recommended to the BoC for use in the weekly review (s, S) policy for replenishment. These ratios provide a good balance between achieving the cost savings and maintaining the quality of service. Because the smaller the ratio, the higher the risk, the chosen ratios are small

Figure 5. (Color online) Expected Number of Potential Shortages as a Function of the s/S Ratio

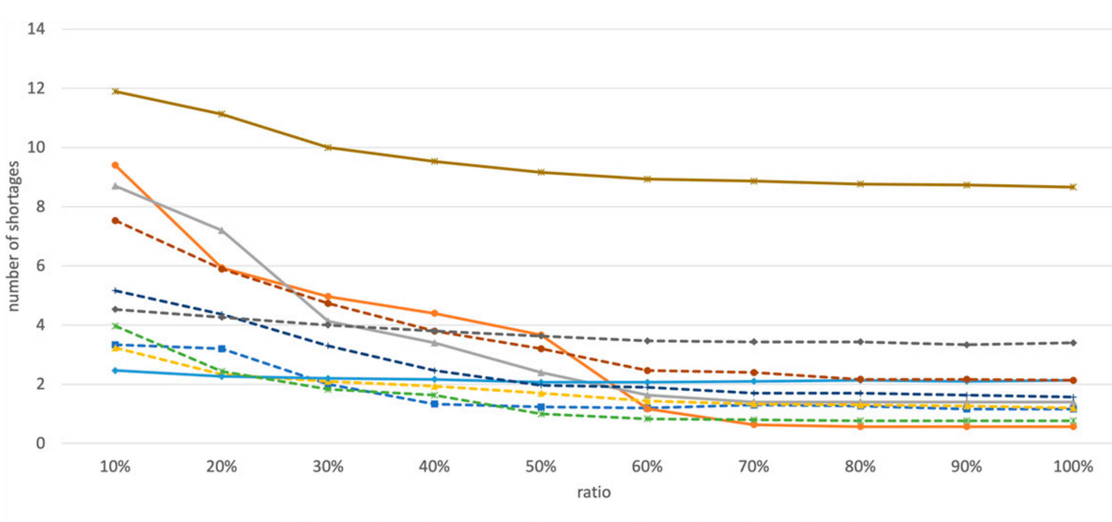

Table 5. Trigger Value, s, as a Percentage of the Target, S, for the Replenishment Weekly Review (s, S) Policy

	RDPs served by air					
	A	D	F	G	H	I
Percent of target	60%	20%	30%	40%	60%	30%
	RDPs served by ground					
	B	C	E	J		
Percent of target	100%	60%	100%	100%		

enough to provide important cost savings but high enough to avoid unnecessary risk (i.e., shortages). The appropriateness of these ratios should be monitored regularly and can be changed by the BoC whenever they are deemed no longer suitable to their operations. Note that the recommended triggers leave in place the current BoC rules (ratio of 100%) for three out of the four RDPs served by ground transportation.

Determining the Lower-Bound Trigger Value for Callbacks

In these tests, we vary the trigger values for the callback policy. These triggers are set according to two ratios: the unfit percentage and the value percentage.

- If no replenishment is planned, then a callback is triggered if the total dollar value of unfit notes exceeds the unfit percentage of the total value limit.
- If a replenishment is planned, then a callback is triggered if the total dollar value of unfit notes exceeds half of the unfit percentage of the total value limit (thus increasing the chance to benefit from the two-way shipment discount).
- A callback is also triggered if the total dollar value of all holdings is over the value percentage of the inventory value limit.

Figure 6 presents the percentage of expected improvements (compared with a policy that calls back every week) in the total cost for different callback policies when using the replenishment policy presented in Table 5. As the trigger value for unfit notes (unfit percentage) increases, there is a substantial improvement (up to 60%) in the total costs for those RDPs served by air. However, these improvements also result in a significant increase in the number of overcaps for certain RDPs, as shown in Figure 7. Note that the x axis of Figure 6 is sorted by the unfit percentage and then by the value percentage, whereas it is the opposite for Figure 7. The main reason is that the improvement mainly increases with the unfit percentage,

Figure 6. (Color online) Expected Improvements in the Total Costs as a Function of the Ratios (Unfit Percentage, Value Percentage) for the Callbacks

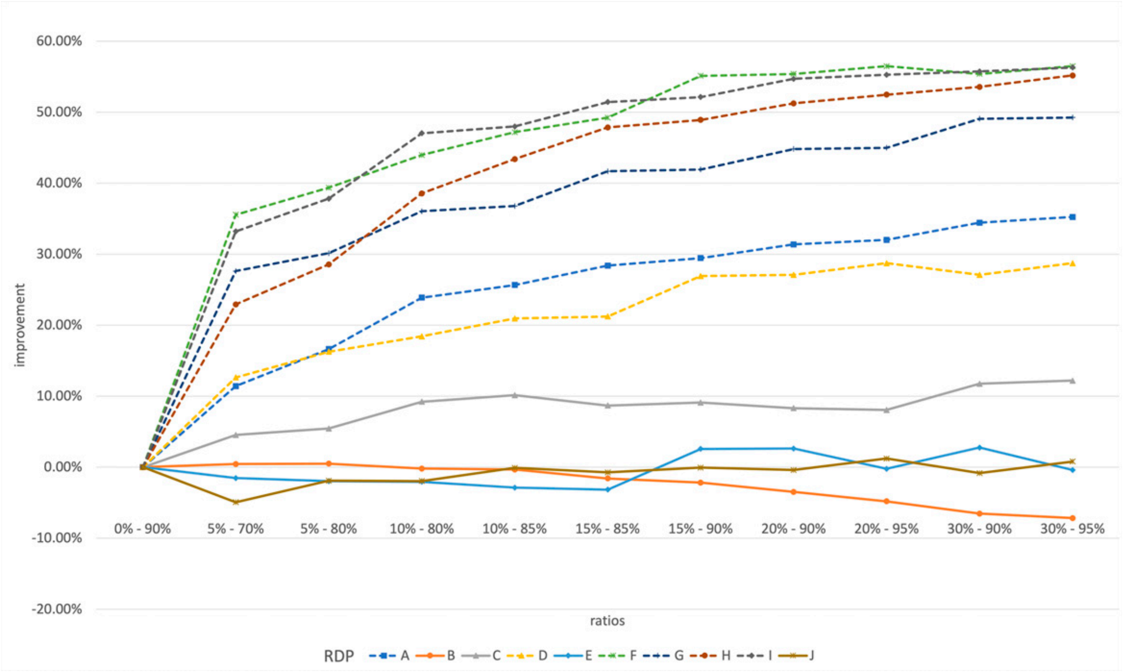

Figure 7. (Color online) Expected Number of Overcaps as a Function of the Ratios (Unfit Percentage, Value Percentage) for the Callbacks

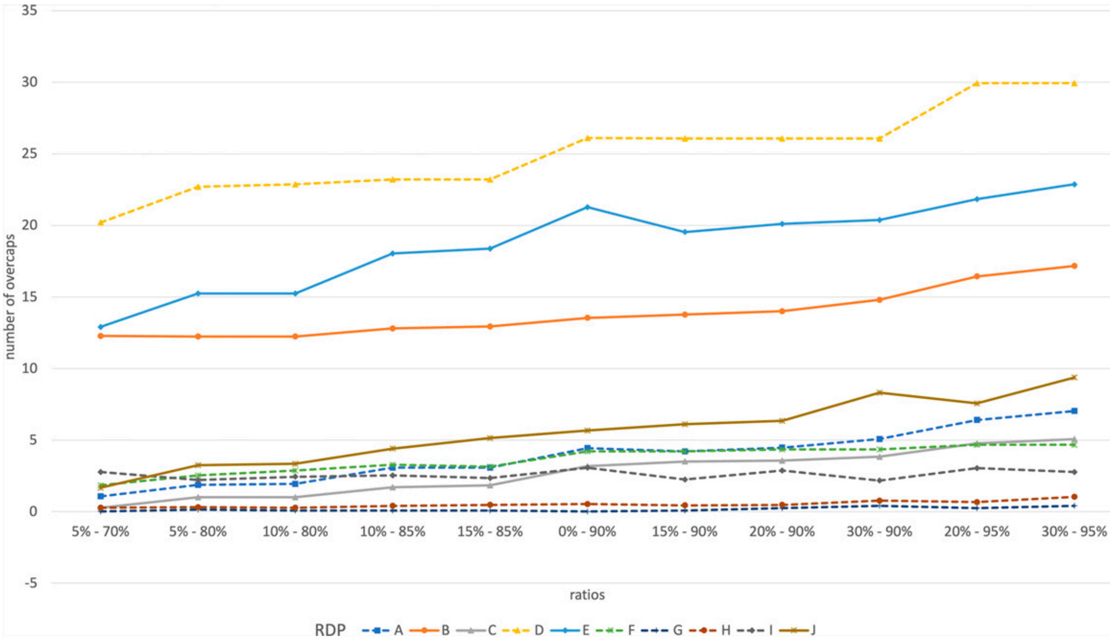

whereas the number of overcaps increases slowly with the value percentage.

To avoid increasing the number of overcaps, we propose that the value percentage that the BoC is currently using (i.e., 90%) be kept, as, combined with the weekly review (s, S) policy for replenishments, it already provides a substantial improvement without increasing the risk of overcaps. We also propose that a low unfit percentage of 10% be taken for the unfit notes, as higher ratios do not seem to produce much improvement while substantially increasing the number of overcaps. The main reason for setting this ratio lower than in the tests (i.e., 15%) is so that the BoC can have regular callbacks and thus detect counterfeit notes in a reasonable time frame.

Determining the Impact of the Proposed Weekly Review (*s*, *S*) Policy

Part of the difficulty in attempting to predict the potential impact of the BoC's switching to the weekly review (s, S) policy is that the BoC already will, on occasion, deviate from their own rules when the situation seems to make it clear that it would be of benefit. Thus, comparing these rules against the weekly review (s, S) policy in a simulation may in fact overestimate the potential impact. To prevent this, we compared a strict adherence to the current rules against both the weekly review (s, S) policy and the actual historical costs over 10 months of real demand data (February 2, 2015, to November 27, 2015). The Christmas period was not included because of the significant difference in demand patterns during that time. To manage the very high demand over Christmas, the BoC policy ships as many notes as possible in December and callbacks as much as possible in January. Table 6 provides the cost improvements associated with using the weekly review (s, S) policy with the ratios presented in Table 5. These results suggest that the proposed weekly review (s, S) policy could have saved approximately 25% on the total costs over those 10 months compared with the BoC rules (if strictly followed) while maintaining the same service quality (i.e., approximately the same number of potential shortages or overnight shipments), though with a small transfer of costs (handling and transportation) to the FIs. In addition, although the historical performance of the BoC was generally an improvement over a strict following of their rules for those RDPs served by air, it in fact led to higher costs for those RDPs served by ground transport. By contrast, for each RDP, the weekly review (s, S) policy did at least as well as those rules (and significantly better for those RDPs served by air) and uniformly outperformed historical practice.

Benefits

Our analysis demonstrates the potential benefit of the weekly review (s, S) policy compared with the current

Table 6. Performance of the Weekly Review (s, S) Policy Compared with Current Practice

	RDPs served by air						
	A	D	F	G	H	I	Total
Historical performances	5.94%	17.88%	19.34%	9.25%	26.12%	22.47%	16.20%
Performances using weekly review (s, S) policy	14.64%	23.00%	40.09%	23.68%	37.72%	31.63%	26.87%
Increase in potential shortages	1	1	0	0	0	1	3
Increase in withdrawal fulfillments from another RDC	0	4	9	7	24	6	50

	RDPs served by ground				
	B	C	E	J	Total
Historical performances	−2.64%	−69.21%	10.85%	−12.48%	−8.53%
Performances using weekly review (s, S) policy	0.00%	13.13%	0.00%	0.00%	1.34%
Increase in potential shortages	0	0	0	0	0
Increase in withdrawal fulfilments from another RDC	0	2	0	0	2

BoC operations for the inventory management of the RDCs. The main impacts are fourfold:

- The policy decreases the transportation activity at a national level by scheduling fewer replenishments and callbacks, thus generating significant savings on the transportation costs (especially by air).
- By scheduling fewer shipments, the policy also increases the size of the shipments. Thus, there is an opportunity to further decrease operational costs by renegotiating the contracts with the carriers to decrease the cost of larger shipments.
- As there are fewer replenishments, the transportation activity increases marginally at the regional level.
- The lack of any benefit to the utilization of the weekly review (s, S) policy for those regions that are served by ground is largely due to the lack of any significant discount for larger shipments. A renegotiation of the contract could provide an incentive to switch to a weekly review (s, S) policy, even for those regions served by ground.

Too aggressive a choice of s in the weekly review (s, S) policy for replenishments may increase the number of potential shortages and thus decrease the quality of service. Consequently, the ratios of the policy need to be set carefully to prevent such negative impacts. Furthermore, the weekly review (s, S) policy depends on the chosen targets, so the cost savings will thus depend on how they are set. An in-depth analysis of past data could lead to better targets and thus provide even further reduction to the costs. The relatively low cost of overcaps suggests that higher targets could be set that would result in even fewer shipments.

Implementation and Challenges

A pilot project of only one RDP served by air transportation was conducted by the BoC during the month of November 2017 to test the theoretical results of this project without requiring significant disruption to current practice. With ratios of 50% for the replenishments and 10% for the callbacks, the BoC reduced transportation by 50% in comparison with November 2016 and did not encounter any shortage during the pilot. Moreover, there were no complaints by the FIs with regards to increased transfers. In fact, they were not even aware of a change in practice by the BoC.

As a result of the success of the pilot project, the BoC began rolling out the weekly review (s, S) policy across all the regions served by air transportation in February and March of 2018. Regions H and I switched to the weekly review (s, S) policy in February, whereas regions A, D, F, and G transitioned in March. To provide a fair comparison, we compared the monthly average cost for a year's worth of data postimplementation (April 2018 through March 2019) to the two previous years (April 2016 through March 2018). Although we had data going all the way to the end of 2019, the BoC renegotiated their contract with the armored car carrier company in May 2019, so we have chosen not to include that data in order to avoid overstating the impact of the change in inventory policy (because the renegotiation of the contract also reduced the transportation costs). Table 7 provides the net and percent reduction in the number of trips to each region, as well as the corresponding reduction in cost for the six regions served by air. As can be seen, the percentage of cost reduction achieved over the 12 months is right in line with that anticipated by the simulation. Although there are some differences between regions, all regions demonstrated a substantial reduction in cost. The one exception is region D. This region is a net depositer, meaning that the region largely deposits currency rather than withdraws. It is also a high-traffic region, so that even with the weekly review (s, S) policy, shipments were almost weekly. This could perhaps be due to a target that is too low. Overall, average monthly savings were about \$45,000, yielding an anticipated yearly savings of around \$500,000. Savings are perhaps even higher than

Table 7. Comparison of 12 Months Post-implementation vs. the Average of the Two Years Prior to Implementation

	Trips saved per year		Cost savings per year	
Region	Number	Percent reduction	Total costs saved	Percent reduction
A	23	25.0%	$135,076	27.2%
D	1.5	1.8%	$33,615	2.9%
F	15	26.8%	$83,874	22.6%
G	23	35.9%	$63,737	18.5%
H	17	31.5%	$72,099	23.8%
I	20	34.5%	$148,353	30.6%
Total	99.5	24.5%	$536,754	15.3%

this, as a regression analysis of the 32 months prior to implementation led to an average monthly increase of $900 (albeit with a fair bit of variability from month to month). In other words, prior to implementation, transportation costs were experiencing some increase from month to month, so it is difficult to avoid the conclusion that the dramatic decrease in cost postimplementation was, in fact, the result of the change in policy.

Although the reduction in transportation costs is clear and substantial, there have been a number of challenges associated with implementation. First, although there has been no increase in the number of shortages (defined as a 24-hour period during which a region has insufficient notes of at least one denomination), there has been some increase in the number of overnight shipments used to prevent a shortage, as anticipated (though the additional cost of these clearly does not outweigh the cost savings resulting from fewer shipments, as they are included in the cost calculations in Table 7). This trade-off has been readily accepted by the BoC. Second, the switch from a prescheduled weekly shipment has led to more variability in the time required to process a shipment both at the AOCs and in the regions, leading to some challenges with respect to workload and staffing. In particular, the larger return shipments have, on occasion, caused challenges in receiving, opening, and processing the notes within an acceptable period of time. To address these issues, the BoC is considering the implementation of a fixed schedule for the callbacks. However, rather than reverting to a weekly callback, the BoC is experimenting with individualized schedules, where some regions will have weekly callbacks and others will space them out every two or three weeks. Third, fewer callbacks has led to some storage issues for some of the RDCs and some increase in the number of overages, though the minimal cost associated with these has, again, been accepted as a reasonable trade-off. Finally, although there have been few complaints from the FIs, the increased number of withdrawal fulfillments from another RDC within a single RDP has led to an exacerbation of a preexisting problem. It would appear that some FIs are less fastidious than others in terms of ensuring that so-called fit notes are, in fact, still fit for circulation. This has led to complaints by some FIs as to the quality of the notes received from a different RDC. Using the simulation, we have begun to explore the impact of imposing the further constraint that some RDCs are not allowed to satisfy demand outside of their own.

In short, the switch to the weekly review (s, S) policy has resulted in the type of cost savings anticipated by the simulation but with some predictable and manageable challenges. It has also led to a number of potential future projects involving the FIs more directly. First, future discussions with the FIs may lead to an opportunity to optimize the target values that may well lead to further reductions in the transportation costs, as it is fairly clear that current targets overestimate the importance of avoiding overcaps. Second, a better understanding of the inventory management policies used by the FIs could lead to a better forecast of demand from week to week, which would in turn allow for more aggressive choices of s and S. Some incentives could also be used by the BoC to push the FIs to deposit and withdraw on a daily or weekly basis in order to avoid any backlogging effects and thus build a better forecasting tool. Finally, more advanced models based on approximate dynamic programming could be developed to solve the mathematical formulation of the problem and lead to an optimization method to set the trigger values of the (s, S) policy.

Appendix
Mathematical Formulation

Formulation (A.1)–(A.9) represents the BoC inventory management problem for a given RDP over a fixed horizon. This deterministic formulation does not take into account the dynamics of the problem but serves to illustrate the nature of the optimization challenge. Let T and T_s be the set of all days within the planning horizon and the set of days on which a shipment (replenishment or/and callback) arrives/leaves an RDP, respectively. Each RDC, $i \in I$, has a maximum inventory value $V_i^{\max}$ and a fee $c^{\max}$ for each dollar that exceeds this maximum. Each note, $j \in B$, has a value v_j, where B is the set of all notes. The sets B^{fit}, B^{new}, and B^{unfit} contain, respectively, the fit, new, and unfit notes ($B^{\text{fit}} \cup B^{\text{new}} \cup B^{\text{unfit}} = B$). Finally, the demand D_{ijt} of RDC i for note j on day t is positive for a withdrawal and negative for a deposit.

The variable I_{ijt} represents the inventory level for note j at RDC i at the end of day t. The variables x^r_{ijt}, x^c_{ijt}, and x^n_{ijt} represent the number of notes j at RDC i on day t that are received from a replenishment, sent for a callback, and received from a night courier, respectively. Variable x^+_{ijt} counts the number of new notes j used to meet the demand in fit notes f_j of the same denomination at RDC i on day t. The count is positive for new notes, negative for fit notes, and null for unfit notes. Variable r_{iljt} gives the number of notes j that are transferred from RDC i to RDC l on day t. Variable z_{it} measures the value exceeding the maximum $V_i^{\max}$ for RDC i at the end of day t.

Finally, the functions $c^s(x^r_{ijt}, x^c_{ijt})$, $c^r_t(r_{iljt})$, and $c^n(x^n_{ijt})$ give the price on day t for a replenishment/callback, a transfer, and a night courier, respectively. It is worth noting that these costs are nonlinear, and, thus, if the formulation is solved with an integer-programming solver, then they would need to be linearized. Equation (A.1) provides the total cost, Equation (A.2) ensures that demand is met, and Equation (A.3) tracks any overcap. Equations (A.4)–(A.6) ensure that decisions regarding how demand is met are consistent with no demand unserved (inventory always positive), and Equation (A.7) ensures that decisions are made on the scheduled days.

$$\min \sum_{t \in T_s} c^s(x^r_{ijt}, x^c_{ijt}) + \sum_{t \in T} c^r_t(r_{iljt}) + \sum_{i \in I} \sum_{j \in B^{\text{new}}} \sum_{t \in T} c^+ x^+_{ijt} + \sum_{t \in T} c^n(x^n_{ijt}) + \sum_{i \in I} \sum_{t \in T} c^{\max} z_{it} \tag{A.1}$$

subject to

$$I_{ij(t-1)} + x^r_{ijt} + \sum_{l \in I} r_{lijt} + x^n_{ijt} = I_{ijt} + D_{ijt} + x^c_{ijt} + \sum_{l \in I} r_{iljt}, + x^+_{ijt} \qquad \forall i \in I, t \in T \tag{A.2}$$

$$\sum_{j \in B} v_j I_{ijt} \le V_i^{\max} + z_{it}, \qquad \forall i \in I, t \in T \tag{A.3}$$

$$I_{ijt} \ge 0, \qquad \forall i \in I, j \in B, t \in T \tag{A.4}$$

$$x^+_{if_jt} + x^+_{ijt} = 0, \qquad \forall i \in I, j \in B^{\text{new}}, t \in T \tag{A.5}$$

$$x^+_{ijt} = 0, \qquad \forall i \in I, j \in B^{\text{unfit}}, t \in T \tag{A.6}$$

$$x^r_{ijt} = x^c_{ijt} = 0, \qquad \forall i \in I, j \in B, t \notin T_s \tag{A.7}$$

$$x^r_{ijt}, x^c_{ijt}, r_{ii_1jt} \ge 0, \qquad \forall i, l \in I, j \in B, t \in T \tag{A.8}$$

$$x^+_{ijt} \ge 0, \qquad \forall i \in I, j \in B^{\text{new}}, t \in T \tag{A.9}$$

Mathematical Details for the Weekly Review (*s*, *S*) Policy

Let us define a few notations. The cumulative demand $\overline{D}_{ijt}$ of both fit and new bank notes between the planning day t_0 and the delivery date t is estimated for each denomination j at each RDC i with the demand forecast black-box software; S_{ij} is the volume inventory target for RDC i and for bank note j.

A Weekly Review (*s*, *S*) Policy for Replenishments

The (s, S) policy for replenishments is mathematically defined as follows:

- If the expected holdings of fit notes j plus new notes n_j for any denomination are below a trigger value s_j ($\sum_{i \in I}(I_{ijt_0} + I_{in_jt_0} - \overline{D}_{ijt} - \overline{D}_{in_jt}) < s_j$) *or* if the expected holdings of new notes j for any denomination are below a different trigger value s_j ($\sum_{i \in I}(I_{ijt_0} - \overline{D}_{ijt}) < s_j$), then a replenishment is planned, so that the expected inventory within the RDP on the delivery date t is equal to S_j for any notes $j \in B^{\text{fit}} \cup B^{\text{new}}$. More precisely, for the new note j, $\sum_{k=j,f_j} \sum_{i \in I} x^r_{ikt} = \sum_{k=j,f_j} \sum_{i \in I}(S_{ik} + \bar{D}_{ikt} - I_{ikt_0})$ and $\sum_{i \in I} x^r_{ijt} = \sum_{i \in I}(S_{ij} + \bar{D}_{ijt} - I_{ijt_0})$, where note f_j corresponds to the fit note of the same denomination as new note j.
- Otherwise no shipment is made.

A Weekly Review (*s*, *S*) Policy for Callbacks

The (s, S) policy for callbacks is mathematically defined as follows:

- If the expected dollar value of holdings of unfit notes j in an RDP is over s_j ($\sum_{i \in I} I_{ijt_0} - \bar{D}_{ijt} > s_j$), then a callback is planned and all unfit notes in the entire RDP are shipped back to the AOC: $\sum_{i \in I} x^c_{ijt} = \sum_{i \in I}(I_{ijt_0} - \bar{D}_{ijt})$. Otherwise, nothing is done.
- If the expected total dollar value in a RDC i is over $V_i^{\max}$ (i.e., $\alpha V_i^{\max}$), then a callback is planned as follows: all unfit notes within the RDP are shipped back as well as all fit notes j over the targets for which the RDC i in question has a depositor status: $x^c_{ijt} = \max(0, I_{ijt_0} - \bar{D}_{ijt} - S_{ij})$. Otherwise, nothing is done.

Simulation Algorithm

Note that α corresponds to the value percentage.

Algorithm 1 (Simulation of the BoC Operations for a Given RDP and a Given Weekly Review Policy *P*)

1. Initialize the cost $c = 0$, the number of withdrawal fulfillments $N^W = 0$, and the number of shortages $N^- = 0$;
2. Initialize the inventory levels for the first day (I_{ij0});
3. **for** $t \in T \backslash \{0\}$ **do**
4. Initialize the inventory on day t: $I_{ijt} = I_{ij(t-1)}$;
5. Process, if needed, the shipments that arrive at the RDP or depart from the RDP: $I_{ijt} += x^r_{ijt} - x^c_{ijt} + \sum_{l \in I} r_{lijt} - \sum_{l \in I} r_{iljt}$;
6. Plan replenishement and callback with policy P if the right weekday, and update cost: $c += c^s(x^r_{ijt}, x^c_{ijt})$;
7. Plan rebalancing if the right weekday or transfer if there is a risk of overcap ($\sum_j v_j I_{ijt} \ge \alpha V_i^{\max}$) and update cost: $c += c^r_t(r_{iljt})$;
8. Update the inventory with the demand: $I_{ijt} -= D_{ijt}$;
   //Fix the inventory levels (can be negative or above the value limit) with the following recourses:
9. **if** *inventory level* I_{if_jt} is negative for fit notes f_j **then**
10. Try to use new notes j of the same denomination: $x^+_{ijt} = \min(-I_{if_jt}, I_{ijt})$, $I_{if_jt} += x^+_{ijt}$, and $I_{ijt} -= x^+_{ijt}$;
11. Update the cost: $c += c^+ x^+_{ijt}$;
12. **if** *there is a negative inventory* **then**
13. Try to transfer some notes based on the target values S_{ij} and "withdrawer" profiles: $I_{ijt} += \sum_{l \in I} r_{lijt} - \sum_{l \in I} r_{iljt}$;
14. Increment the number of withdrawal fulfillments $N^W += 1$;
15. **if** *there is an overcap at RDC i* ($\sum_j v_j I_{ijt} \ge V_i^{\max}$) **then**
16. Update the cost: $c += c^{\max}[V_i^{\max} - \sum_j v_j I_{ijt}]$
17. **if** *there is still a negative inventory at RDC i* **then**
18. Send notes through a night courier to cover the shortage: $I_{ijt} += x^n_{ijt}$;
19. Increment the number of shortages: $N^- += 1$;

References

Chen X, Simchi-Levi D (2009) A new approach to the stochastic cash balance problem with fixed costs. *Probab. Engrg. Inform. Sci.* 23(4):545–562.

DeCroix G, Zipkin P (2005) Inventory management for an assembly system with product or component returns. *Management Sci.* 51(8):1250–1265.

Dowlatshahi S (2000) Developing a theory of reverse logistics. *Interfaces* 30(3):143–155.

Feng Y, Xiao B (2000) A new algorithm for computing optimal (*s*,*S*) policies in a stochastic single item/location inventory system. *IIE Trans.* 32(11):1081–1090.

Fleischmann M, van Nunen J, Grave B (2003) Integrated closed-loop supply chains and spare parts management at IBM. *Interfaces* 33(6):44–56.

Geismar H, Sriskandarajah C, Zhu Y (2017) A review of operational issues in managing physical currency supply chains. *Production Oper. Management* 26(6):976–996.

Geismar H, Dawande M, Rajamani D, Sriskandarajah C (2007) Managing a bank's currency inventory under new federal reserve guidelines. *Manufacturing Service Oper. Management* 9(2):147–167.

Hopp WJ, Spearman ML (2011) *Factory Physics* (Waveland Press, Long Grove, IL).

Huang Y, Geismar HN, Rajamani D, Sethi S, Sriskandarajah C, Carlos M (2017) Optimizing logistics operations in a country's supply chain network. *IISE Trans.* 49(2):223–237.

Johnson E (1968) On (*s*,*S*) policies. *Management Sci.* 15(1):80–101.

Massoud N (2005) How should central banks determine and control their bank note inventory? *J. Banking Finance* 29(12):3099–3199.

Mehrotra M, Dawande M, Sriskandarajah S (2010) A depository institution's optimal currency supply network under the fed's new guidelines: Operating policies, logistics, and impact. *Production Oper. Management* 19(6):709–724.

Mitra S (2012) Inventory management in a two-echelon closed-loop supply chain with correlated demands and returns. *Comput. Indust. Engrg.* 62(4):870–879.

Toktay L, Wein L, Zenios S (1997) Inventory management of remanufacturable products. *Management Sci.* 46(11):1412–1426.

Zheng Y, Federgruen A (1991) Finding optimal (*s*, *S*) policies is about as simple as evaluating a single policy. *Oper. Res.* 39(4): 654–665.

Verification Letter

Pierre Roach, Senior Director, Quality & Operations Currency, Bank of Canada, 234 Wellington St., Ottawa, Ontario K1A 0G9, Canada, writes:

"The Bank of Canada (Boc) is responsible for ensuring that Canadians can access cash of all denominations from wherever they are in Canada. To that end, we have agreements with financial institutions across the country that allow us to store inventory in secure vaults in select cities. The BoC has long operated the management of these vaults by setting a target inventory for each denomination in each vault and shipping sufficient cash each week to ensure that the inventory in the vault stays close to the targets. However, the transportation costs associated with maintaining these inventories is substantial. We thus contacted Professor Patrick from the Telfer School of Management at the University of Ottawa in order to see if he might be willing to work with us to develop better inventory management practices. The BoC agreed to fund the postdoctoral position of Antoine Legrain, who then built a simulation model that demonstrated the potentially significant benefit to switching to an adapted (*s*, *S*) policy to manage the BoC holdings.

"In November 2017, following the submission of their final report, the BoC ran a pilot project using one of our regional distribution points to see if the expected reduction in transportation costs would in fact be realized in practice. The initial pilot project exceeded our expectations, as transportation costs were reduced by half. This was done without notifying the financial institutions, as we wanted to see if they would notice the change. We did not receive a single concern raised by the financial institution. In fact, the only negative impact of the change in policy that was noted was an increase in the processing time of returned unfit notes, as the new policy returned notes less frequently but in higher volume.

"Following the success of the pilot, we began rolling out the policy across all of our regional distribution points that are served by air transportation. In February, we transitioned two regions, and, in March, we transitioned three more. In July 2018, we completed the transition of the final regional distribution point serviced by air. At the end of October 2018, we compared the total transportation costs for the six regions served by air over the last three years from February through October of each year. There were 86 fewer shipments made in 2018 compared with the average of 2016 and 2017, with an average reduction in costs of $448,344. Thus, in 2018, our shipment costs to the six regional distribution points served by air was 80% of the costs incurred in 2017. This was achieved despite the fact that we did not transition to the new policy at all sites till July. There have been challenges associated with increased variability in the frequency and size of shipments, but we are working through how to address these without losing the substantial savings."

Antoine Legrain's research focuses on real-time operations management (online and stochastic optimization applied to dynamic problems) to improve access to healthcare systems and manage on-demand and multimodal transportation services. Professor Legrain seeks to propose new mathematical programming methodologies, as well as develop practical tools to tackle innovative industrial problems.

Johnathan Patrick's research applies the methods of operations research to improve the efficiency of healthcare management. His primary stream of research examines intelligent patient scheduling policies under setting with heterogenous patients and/or resources. Other interests research interests include capacity planning and staffing. On the theoretical side, Professor Patrick is interested in seeking to advance the field of approximate dynamic programming.