Staff Working Paper/Document de travail du personnel — 2021-64
Last updated: December 17, 2021

Quantifying the Economic Benefits of Payments Modernization: the Case of the Large-Value Payment System

by Neville Arjani,[1] Fuchun Li[2] and Zhentong Lu[3]

[1] Canada Deposit Insurance Corporation

[2] Payments Canada

[3] Financial Stability Department
Bank of Canada, Ottawa, Ontario, Canada K1A 0G9

narjani@cdic.ca, fli@payments.ca, zlu@bankofcanada.ca

ISSN 1701-9397

Acknowledgements
The authors thank Segun Bewaji, Anne Butler, Cyrielle Chiron, Jonathan Chiu, Walter Engert, Charles Kahn, Andrew McFarlane, and Andrew Usher for their helpful suggestions, as well as participants of the 19th Simulator Seminar (2021) at the Bank of Finland, the 2020 Bank of Canada–PayCan Quarterly Research Workshop, and the Bank of Canada BBL seminar for their useful comments. The views expressed are ours and not the opinion of any organizations. All errors are ours.

Abstract

In this paper, we develop a discrete choice framework to quantify the economic benefits of payments modernization in Canada. Focusing on Canada's large-value transfer system (LVTS), we first estimate participants' preferences for liquidity cost, payment safety and the network effect by exploiting intraday variations in the relative choice probabilities of the two substitutable sub-systems in the LVTS (i.e., Tranches 1 and 2). Then, with the estimated model, we conduct counterfactual simulations to calculate the changes in participants' welfare when the LVTS is replaced by a real-time gross settlement system (RTGS), like Lynx (as an important part of the payments modernization initiative). The results show that, first, compared to the old system, Lynx has higher liquidity costs but is more secure, while the former is considered a more important factor by system participants. Second, when over 90% of current LVTS payments migrate to Lynx, there is an overall welfare gain; however, it maybe difficult to achieve such a high migration ratio in the new market equilibrium. Third, accounting for equilibrium adjustment, about a 75% service level improvement is needed to generate overall net economic benefits to participants. Among other things, adopting a liquidity savings mechanism and reducing risks in the new system could help achieve this improvement. Finally, the welfare changes are quite heterogeneous, especially between large and small participants.

Topics: Financial institutions; Financial system regulation and policies; Payment clearing and settlement systems
JEL codes: C3, E4, E42, G1, G2, G28

1 Introduction

Payment systems play a crucial role in an economy by providing the mechanisms through which consumers, financial institutions, and governments can purchase goods and services, make financial investments, and transfer funds. Well-functioning payments systems can enhance the stability of the financial system, lower transaction costs, promote the efficient use of financial resources, and facilitate the conduct of monetary policy. Therefore, countries around the world have devoted much effort to monitoring, regulating, and upgrading their payments systems with the latest technological developments, international messaging standards, and modernizing their regulatory and risk control frameworks.

In Canada, there are currently two core payments systems: the large-value transfer system (LVTS) and the automated clearing settlement system (ACSS). Although these two payments systems are still functioning well, they were designed more than 20 years ago. As a result, their limited functionality and outdated technologies and risk control measures indicate that these systems are not suitable as future foundational platforms for payments clearing and settlement. Enhancements are required to establish a truly modern payments ecosystem that is fast, flexible, secure, promotes innovation, and strengthens Canada's competitive position internationally. To achieve the necessary enhancements, Canada is undertaking a large initiative to modernize its payments ecosystem. In the modernized world, there will be three new core payments systems: a real-time gross settlement system (RTGS) for large-value payments (named Lynx), a deferred net settlement system (DNS) for clearing lower valued and less urgent payments (a new batch retail system that has not yet been named), and a real-time payment system for processing small-value payments called RTR (Real-Time Rail). These three payments systems will coexist and complement each other to serve their intended purpose, which is to provide a richer set of viable payments options to meet Canadians' needs.

Although this modernized ecosystem is expected to bring large benefits to Canadian financial markets and the overall economy, limited work has been done to provide an economic model-based quantitative assessment of these benefits.[1] Also, payments modernization involves substantial investment costs and these new systems might generate new types of risk. Therefore, it is crucial to provide a quantitative assessment of the potential benefits the new payments system could bring. Such an assessment would provide useful information for the ongoing payments modernization initiative.

Evaluating the overall benefits of payments modernization is an ambitious task. Given the available data, in this paper we take a first step and focus on Canada's large-value payments system modernization; i.e., the replacement of the LVTS with Lynx. To do this, we build an empirical model based on the discrete choice approach for consumer welfare evaluation developed by McFadden (1981), Small and Rosen (1981), Trajtenberg (1989), and Petrin (2002), among others. Exploiting the intraday variations in participants' system choice behaviors, which are recorded in the historical LVTS and ACSS data, we estimate each participant's payoff function (preference) from sending an inter-bank payment. Then, using the estimated model, we conduct counterfactual experiments to calculate the welfare changes that occur after the LVTS is replaced by Lynx.

A participant's (typically a financial institution) payoff from sending a payment via a large-value payment system depends on many factors. For starters, we consider the two most prominent ones that govern participants' key trade-offs; i.e., the liquidity cost and the risk of payment failure (or delay). To explicitly measure these risks, we construct two indicators from the characteristics of both the payments system and the payment in question. These indicators capture the payment-by-payment variations in the financial institution's (FI) incentives when sending payments. Besides the two indicators, the payoff of sending a payment also depends on how likely other

[1] One exception is Arjani (2015), who applies discounted cash flow (DCF) analysis to study the potential economic benefits of adopting the ISO 20022 payments messaging standard for payments modernization. However, given the limitations of the DCF approach in estimating future cash flows, Arjani (2015) suggests the payments research community use an economic model-based approach to quantifying the economic benefits of payments modernization.

participants also use the same system (known as the "network effect"), as the payments game exhibits clear strategic complementarity (see, among others, Bech and Garratt (2003)). Finally, our model includes an unobserved, system specific "service quality" level in the payoff function to capture any residual factors beyond the ones mentioned above.

Exploiting the LVTS's special feature in that it effectively constitutes two systems, Tranche 1 (T1) and Tranche 2 (T2), we estimate the participants' payoff function based on their realized system choices when sending different payments. The key to our estimation strategy is that the two constructed indicators and the network effect term sufficiently vary by payment to allow us to identify their coefficients in the payoff function. With the estimated payoff function, we calculate a participant's welfare from sending each payment and then aggregate this over all payments to determine the total welfare from using the LVTS.

To evaluate the potential welfare gain; i.e., the economic benefit of replacing the LVTS with Lynx, we run a counterfactual simulation by letting all of the LVTS payments run through a baseline RTGS system (mimicking Lynx) and then record the output data, especially the two indicators we constructed. These indicators summarize the key differences between Lynx and the LVTS along the two dimensions: the liquidity cost and the payment risk. Also, since the network effect in the payoff function depends on the aggregate outcome (the choice probabilities of the two systems), we re-calculate the new equilibrium choice probability of Lynx. Finally, we calculate the total welfare of Lynx and compare it with that of the LVTS under alternative assumptions on equilibrium adjustment and service quality improvement, among others.

Our results show that, first, when comparing the LVTS to Lynx, the liquidity cost is higher and the payment risk is lower, with the system participants considering the former a more important factor. Second, the choice of a large-value payments system exhibits rather strong network effects, which means it is important to take the equilibrium adjustments into account when the LVTS is replaced by Lynx. Third, the net economic benefits of Lynx over the LVTS depends crucially on the extent of the payments migration from the old system to the new one and the resulting changes in the unobserved service quality level. In particular, we find that there is an overall welfare gain to participants if over 90% of current LVTS payments migrate to Lynx, which seems unlikely given our equilibrium adjustment calculation, or Lynx results in a 75% improvement over the LVTS in terms of service quality. For example, adopting a liquidity-saving mechanism (LSM) and/or reducing the credit risk (because Lynx is moving towards an RTGS system) in the system can help achieve this improvement. Fourth, the welfare changes are heterogeneous between large and small participants: smaller participants tend to have lower welfare gains (or higher welfare losses) than large ones when the LVTS is replaced by Lynx.

The rest of this paper is organized as follows. Section 2 provides background information on the LVTS and its modernized version, Lynx. Section 3 describes the historical payments data used in this paper. Section 4 proposes two indicators to quantify the liquidity cost and the perceived risk of delay a participant faces when sending a payment. In Section 5 we present an empirical model of the participants' choices of payments system. The estimation results of the model are presented in Section 6. The economic benefits from Lynx are computed and analyzed in Section 7. Section 8 concludes.

2 Payments Modernization: Large-Value Payments Systems

A pivotal part of the current payments modernization initiative in Canada is the replacement of the current large-value payments system with its modernized version, Lynx. In this section, we briefly summarize background information on each system.[2] This will help us identify the key changes participants experience when moving from the LVTS to Lynx, the findings of which will be used to set up our model for quantifying the economic benefits of replacing the old payments system with the new one.

2.1 The Legacy System: LVTS

The LVTS is Canada's core electronic payments system for processing inter-bank large-value payments. It is the only "systemically important" payments system in Canada that is operated by Payments Canada and overseen by the Bank of Canada (BoC). This study used data for 2019, a time during which Canada's LVTS had 17 direct participants, including the BoC.[3] The LVTS consists of two sub-systems, Tranche 1 (T1) or Tranche 2 (T2). A participant can choose either one when sending a payment. T1 and T2 differ, especially in their distinct collateral requirements and risk control measures.

For T1, a participant can send a payment as long as its net debit position (when the current payment is made), calculated as the difference between all of the T1 payments it sent (including the current payment) and those it received, is no greater than the collateral the participant has pledged to the BoC to back up all of its T1 payments. If the participant defaults on its LVTS settlement obligations, the collateral it pledged will be used to cover any net negative position in its T1 account. For this reason, T1 payments are known as "defaulter pays."

[2]In the Canadian payments modernization program, the current two core payments systems (ACSS and LVTS) and the three core payments systems to be used in the future (RTR, the new batch retail system, and the LVTS) differ from each other in a variety of ways. See Kosse et al. (2020) for a detailed summary of these systems' main attributes.

[3]These participants include the Big Six Banks, BoC, Laurentian Bank, Manulife Bank, Caise Desjardins (the largest co-operative movement in Canada), Alberta Treasury Branches (a provincially owned deposit-taking institution), Credit Union Central of Canada (a credit union consortium) as well as foreign banks with branches in Canada (State Street Bank, Bank of American, BNP Paribas, HSBC, ICIC). Any deposit-taking institution and member of Payments Canada can be a member of the LVTS as long as they maintain an account with the BoC and can pledge the necessary collateral in LVTS. Deposit-taking institutions that are not members of the LVTS must send (or receive) their payments through one of its direct participants.

In T2, at the beginning of each day, each participant grants bilateral credit limits (BCLs) to every other participant in the system; this represents the largest bilateral net exposure it is willing to accept with respect to the other participant. In addition, each participant is subject to a multilateral net debit cap, calculated as the sum of all of the BCLs extended to it and then multiplied by a specified system-wide percentage (SWP, currently set at 30%) set by the BoC. The multilateral net debit cap (T2NDC) represents the maximum multilateral net debit position the participant can incur against all other participants during the trading day. Each participant pledges, to the BoC, collateral that is equal to the largest BCL it has extended to any other participants, multiplied by the SWP. If a participant defaults on its final settlement obligation, the collateral pool is used to cover the defaulter's remaining amounts owing.[4] For this reason, T2 payments are referred as "survivor pays."

From the above description, we can see the two key differences the participants perceive as existing between T1 and T2: T1 is more resilient to credit risk but is more costly in terms of liquidity because every payment needs to be fully collateralized in order to be processed (Kosse et al. (2021)). On the other hand, T2 uses liquidity more efficiently but it also has higher credit risk because it has uncollateralized BCLs. As a result, participants must take into account this key trade-off between the two subsystems when they are sending payments.

2.2 The Modernized System: Lynx

Lynx is Canada's new high-value payments system for processing large-value, time critical payments. One of the significant changes, when moving from the LVTS to Lynx, is the change in the financial risk model. The financial risk model in Lynx is intended to mitigate credit risks.[5] To achieve this objective, Lynx will be an RTGS

[4]In the event of a participant default, the surviving participants' losses are determined based on the BCL that had been granted to the defaulter (Arjani and McVanel (2006)).

[5]Note that the mitigation of the credit risk in Lynx does not mean that Lynx completely eliminates systemic risk. In fact, an important concern in Lynx is the risk of a liquidity shortage, which may trigger a systemic risk; e.g., the risk of gridlocking the whole system.

system and, as such, its participants will fully cover its credit risk exposure, which means it will no longer need to rely on either the "survivors-pay" collateral pool or the residual guarantee from the BoC.[6]

However, the reduction of credit risk in Lynx is traded off against a substantial increase in intraday liquidity requirements. As a result, to manage these liquidity requirements, Lynx will offer two distinct mechanisms: a liquidity savings mechanism (LSM) and an urgent payment mechanism (UPM). Lynx's LSM will enable participants to delay a payment and reduce the amount of liquidity required to settle a payment because it uses a combination of queuing, intraday liquidity recycling and payment offsetting. For payments that must be settled without delay, participants may use Lynx's UPM.

Lynx is very similar to T1 in the LVTS since every payment sent through T1 must be fully collateralized. Moreover, since most LVTS payments are sent through T2 (see Kosse et al. (2021)), the shift from the LVTS to Lynx implies an overall increase in liquidity costs and a decrease in credit risk. In this study, we use an economic model-based approach to conduct a quantitative analysis of how replacing the LVTS with Lynx will affect the overall welfare of the participants in Canada's large-value payments system.

3 Data

3.1 Data Source

Our main data set consists of information on the payments participants processed through the LVTS in 2019. For each payment, we observe its value, exact timing (date, hour and second), sender, receiver and the tranche through which it was sent:

[6]In the exceptional event of multiple participant defaults, if the collateral provided by the surviving participants, in addition to the collateral apportioned by the defaulters, is still not sufficient to cover the final net debit positions of the defaulting participants, the BoC will provide a guarantee of settlement for the remaining amount.

T1 or T2. We also observe each participant's intraday credit limits in both T1 and T2, which are determined by the amount of their pledged collateral and the rules of the LVTS, at any time during each trading day. Based on the payment level transaction data and participants' credit limits in T1 and T2, we will construct intraday time-varying indicators that can capture the key incentives participants face when sending their payments: the liquidity cost and the risk of payment failure.

Payments made to the BoC are mostly sent through T1 because of the small BCLs the bank grants. As this limits participants' freedom of choice (and thus cannot reveal the FIs's true preferences), we exclude from the data and analyses payments sent to the BoC. We also exclude payments of less than CAD 10 because these are mostly test payments.

We supplement the LVTS data with ACSS data, which contains daily aggregated bilateral payment values and volumes sent through ACSS in 2019, broken down by different pairs of participants.

The ACSS data will be used to construct a proxy for the market share of the "outside option" (besides options T1 and T2) in our estimation of the discrete choice model.

3.2 Data Description

When participants send payments through the LVTS, they must make decisions regarding where to send their payments: to T1 or T2. These realized decisions provide us information with which to build a measure of the relative usage of T1 versus T2; i.e., the share of the transactions in T1 relative to all of the transactions in the LVTS. Specifically, this share can be calculated as the ratio of the total volume (value) of payments sent through T1 to the total volume (value) sent through the LVTS.

If the share is close to zero, meaning most payments are processed through T2, this means that the LVTS operates at a level of bilateral trust that is high enough to allow it to overcome the need for costly collateral. As the share increases, it signals that the bilateral trust decreases and the counter-parties require more collateral to buffer the credit risk. In particular, an unexpectedly fast transition from a low to a high share of T1 transactions provides a strong risk signal that participants might not have enough collateral to make their payments.

In Figure 1, we plot the volume share of T1 transactions in each of 100 groups of payments, where the groups are defined by the percentiles of the value distribution of all LVTS payments in 2019. Two clear facts stand out from this figure. First, the overall usage of T1 is much less frequent than that of T2. This indicates that the credit-based transactions (i.e., T2 transactions) relative to all transactions are quite high, implying that on average, for the given sample period, the LVTS maintains an extremely high level of efficiency in its liquidity usage across different percentiles of the value distribution. Second, T1 is used mostly for payments that are very high-value. In particular, this figure shows that there is a relatively low usage of T1 for payments that are below the 71% value percentile. However, the pattern of payments above the 96% percentile shows a fairly high level of T1 usage, suggesting that participants are more concerned with safety when processing very high-value payments. These observations show that, given the trade-off between the efficient use of liquidity and safely processing payments, participants face very different incentives when sending payments that differ in value. [7]

Similarly, the value share of payments per hour through T1 varies over the course of a day in a fairly predictable pattern. The average hour-by-hour pattern of the value share of T1 payments processed in a day is reported in Figure 2, which shows the timing pattern of the value share of T1 payments made in a day. This figure indicates that, in general, T1 is used more heavily in the afternoon, especially during the system's last two hours before closing. The pattern of the afternoon peak in the value share of T1 payments suggests that participants are more concerned with

[7] This pattern of payments is also documented in Kosse et al. (2021). Using the transaction data of the LVTS for the period 2004 to 2018, Kosse et al. (2021) find that the large majority of payments in the small- and medium-value percentile bins are sent through T2. The share of T1 payments, however, starts to increase as the transaction values become larger.

payment safety and thus switch from T2 to T1 closer to the end of day.

Besides the LVTS data, we also need information about the payment activities that take place outside the LVTS. This helps us measure how substitutable the LVTS is to alternative payments systems or other payments methods more generally. Since it is virtually impossible to obtain a comprehensive data set on all payments systems/methods in Canada, we focus on the most direct substitute for the LVTS, the ACSS system, for which we have rather good data. Of course, focusing on the ACSS can potentially understate the "competition" the LVTS is facing from other systems. We shall discuss the implications for our model and the results later.

Now we turn to the ACSS data. ACSS is designed to process payments of smaller value and less urgency than those processed through the LVTS, which makes it more substitutable for smaller-value payments. To capture this substitution pattern, we divide the total value of the payments transacted through the ACSS in 2019 by the mean value of the LVTS payments in each percentile bin (as shown in Figure 1). This gives us a "normalized volume" of ACSS payments that is comparable to the LVTS payments in a given percentile bin. Figure 3 shows the share of the normalized volume of ACSS (relative to the LVTS and ACSS combined) for each value percentile bin. Note that this construction of the outside option captures the substitution pattern that ACSS is more substitutable for smaller payments in LVTS.

4 Liquidity Efficiency and Payment Safety

The credit risk mitigation in T1 comes at the cost of high liquidity usage, while the high efficiency of liquidity usage in T2 comes at the expense of a high credit risk. Moreover, the high cost of accessing liquidity in T1 can increase the LVTS's exposure to liquidity risk.[8] Hence, the liquidity cost and risk (liquidity and/or credit) are the

[8]Since access to liquidity is costly for settling payments in T1, there is the potential (driven by the participant's incentive) for there to be insufficient intraday liquidity and, as result, delays or rejections in the settlement of transactions (Byck and Heijmans (2021)).

two key factors a decision maker must take into account when sending a payment. In this section, we propose two indicators that will be used as proxy variables in our empirical model: the liquidity cost and the risk associated with processing a payment.

4.1 Liquidity Cost Indicator

Given that the payment i can pass the risk control tests in the payment system j, where $j \in \{T1, T2\}$, the liquidity cost indicator, which measures the liquidity cost of settling payment i in terms of the amount of collateral, is defined as

$$LCI_{i,j} = \varphi_{i,j} \cdot \max\left\{V_i - NI_{i,j}, 0\right\}, \tag{1}$$

where V_i is the value of payment i, $NI_{i,j}$ is the cumulative net payment income up to payment i in the current payments cycle in system j, and $\varphi_{i,j}$ is a cost factor that measures the liquidity cost in terms of collateral spending. Given the design of the LVTS, if payment i is processed in T1, then $\varphi_{i,T1}$ equals 1; i.e., \$1 of collateral is required for spending \$1 of credit (Arjani and McVanel (2006)). If payment i is processed in T2, then $\varphi_{i,T2}$ is defined as $\frac{MaxASO_{i,T2}}{T2NDC_{i,T2}}$; i.e., the amount of collateral, on average, that is required for spending \$1 of the line of credit.[9]

The intuition behind the liquidity cost indicator in (1) is straightforward. For any payment i, if $NI_{i,j}$ is greater than V_i, then sending this payment does not cost any of the pledged collateral; i.e., the cost is 0. When $NI_{i,j}$ is less than V_i, the balance $V_i - NI_{i,j}$ needs to be paid through collateral, with different cost factors $\varphi_{i,j}$ for T1 and T2, respectively. Similar indicators are proposed and used in previous literature; see, among others, the recent report of CPMI (2015).

[9] $MaxASO_{i,T2}$ is the largest BCL participant i chooses to grant to any other participant, multiplied by the SWP, which is currently set at 30%. $T2NDC_{i,T2}$ represents the maximum multilateral T2 net debit position participant i can incur in relation to all other participants in T2 (Arjani and McVanel (2006)).

4.2 Safety Indicator

From the sender's point of view, the main risk is that their payment may be rejected or delayed due to a lack of liquidity, so they closely monitor their intraday liquidity position to ensure there is sufficient payment capacity (depends on available liquidity) to remain "safe," and this usually means not violating certain risk control criteria. In the following, we construct a payment-specific indicator that measures the sender's perceived safety regarding their payment capacity for the remainder of the day in which the payment is made. For T1, we build on the liquidity risk indicator proposed in Arjani et al. (2020) and define the safety indicator as

$$SI_{i,T1} = \frac{NI_{i,T1} + CL_{i,T1} + RPI_{i,T1}}{RPD_{i,T1} + V_i}, \tag{2}$$

where $CL_{i,T1}$ denotes the sender's intraday credit limit in the day of payment i, $RPI_{i,T1}$ is the sender's payment income to be received from other participants in the remainder of the day after payment i, and $RPD_{i,T1}$ is the intraday liquidity demand during the remainder of the day.

Assuming the sender can perfectly predict the payment income and demand for the remainder of the day,[10] a greater value of the safety indicator means that the sender is less likely to encounter a liquidity shortage after sending payment i. Hence, the indicator measures the sender's expected payment capability for the remainder of the day. Note that the safety indicator resembles the notion of the "clearing capacity" proposed by Lopez et al. (2021), which measures the value of the payments a participant can send under the design of a particular payments system.

In T2, besides having a multilateral credit limit similar to that in T1, senders are subject to bilateral credit limits. So for any payment i, the sender's perceived safety level depends on their multilateral and bilateral liquidity positions. Hence, we extend the safety indicator for T1 to account for both credit limits; i.e., the safety indicator

[10] The extension that would allow forecasting errors in the indicator is left for future research.

of payment i in T2 is defined as

$$SI_{i,T2} = \min\left\{\frac{NI_{i,T2} + CL_{i,T2} + RPI_{i,T2}}{RPD_{i,T2} + V_i}, \frac{BNI_{i,T2} + BCL_{i,T2} + BRPI_{i,T2}}{BRPD_{i,T2} + V_i}\right\}, \quad (3)$$

where the first expression in the brackets is the multilateral safety indicator that is analogous to $SI_{i,T1}$ and the second expression is the bilateral version with all of the variables defined by the sender's liquidity position with respect to the receiver of the payment i; i.e., $BNI_{i,T2}$, $BCL_{i,T2}$ $BRPI_{i,T2}$ and $BRPD_{i,T2}$ are the sender's bilateral net payment income from the receiver, the bilateral credit limit, the bilateral payment income to be received from the receiver, and the liquidity demand from the receiver for the remainder of the day, respectively.

Compared with $SI_{i,T1}$, $SI_{i,T2}$ is smaller, given the same multilateral liquidity safety indicator in T1 and T2, because of the additional bilateral constraints. Also, there are richer variations in $SI_{i,T2}$ because it depends on the receiver of the payment in question. Moreover, given that their designs differ (i.e., in risk control), T1 and T2 target liquidity and credit risks, respectively, so $SI_{i,T1}$ mostly captures the liquidity risk while $SI_{i,T2}$ largely reflects the credit risk that is associated with the payments.

5 Model

For a given intended payment i between two participants, deciding which payment system to use to settle the payment is a discrete choice problem. Based on the standard demand estimation framework developed by Berry et al. (1995), we propose a discrete choice demand model for the participants' decisions on which payment system to use. The estimated model will be informative about the participants' preferences and, thus, can be used to assess the economic benefits from replacing the LVTS with Lynx.

To begin with, let $\mathcal{J} = \{T1, T2, 0\}$ denote the set of payments systems from which each decision maker can choose to make payment i [11], where the alternative 0 represents the “outside option”; i.e., the option of not choosing either T1 or T2 but using an alternative system; e.g., ACSS for the settlement, or delaying or canceling the current payment for the next cycle, etc.

To map the discrete choice framework to the LVTS payments data, we group similar payments together by defining a series of markets, where a market m is defined as a combination of an “hour-sender-receiver-value percentile.” Given the definition of a market, we can aggregate all the payments in a market (these payments are “similar” in terms of timing/sender/receiver/size) to obtain the total volume of T1 and T2 in the market. The volume of the outside option is constructed as in Section 3.2 and is based on ACSS data. The volume shares of T1, T2 and the outside option correspond to the market shares implied by our discrete choice model, which measures the choice probabilities of the payments systems in the market, m, and are the main outcome variables of our discrete choice model.

Note that with the above specification, we abstract away from the timing decision for sending a payment, which rules out the possibility of a strategic delay. The disadvantage of this approach is that in a counterfactual scenario where the environment changes, our model cannot predict how timing decisions change. Incorporating these more complicated decisions into the model seems challenging and is beyond the focus of this paper; thus, we leave this for future investigations.

Given the choice set of payments systems, $\mathcal{J} = \{T1, T2, 0\}$, a decision maker's optimal choice is determined by their preference, which is represented by a random payoff function. Specifically, for each payment i in market m, the random payoff to the decision maker (i.e., referring to the sender and receiver collectively) of sending a

[11] In fact, the decision about which system to use is jointly made by the sender and the receiver of the payment. For simplicity, we do not model the details of the joint decision process but regard the sender and receiver collectively as one single entity and focus on their final joint decision on which payment system to use. This single entity is defined as the decision maker. For payments that are processed in the LVTS, the decision is mostly dominated by the sender, but we do not want to restrict the interpretation of our model to be a description of the sender's choice problem only.

payment through system $j \in \{T1, T2, 0\}$ is specified to follow a nested logit structure:

$$\pi_{i,j,m} = \alpha LCI_{j,m} + \beta SI_{j,m} + \gamma \bar{s}_{j,m} + X_m \rho + \xi_{j,m} + \zeta_{i,g,m} + (1-\lambda)\varepsilon_{i,j,m}, \qquad (4)$$

where

- $LCI_{j,m}$ is the log of the value-weighted average of the liquidity costs of all of the payments in m and $SI_{j,m}$ is the log of the value-weighted average of the safety indicator in m.

- $\bar{s}_{j,m}$ is the total market share of system j in the "neighboring markets of m," defined as markets that have the same sender as m but a different "receiver-hour-value percentile." This variable captures the sender's overall preference for a particular payments system. The preference could be driven by the economy of scale of using one payment system, or the expected benefits from payments coordination in one payment system; i.e., other participants also use the same payment system. The latter is due to the well-known strategic complementarity in payments games; see, among others, Bech and Garratt (2003)). And this is called the "network effect" or the "social interaction effect" in the network and social interaction literature (see, among others, Brock and Durlauf (2001)).

- X_m is a vector of the observed market specific characteristics, including the dummy variables for each sending participant, receiving participant, value percentile, and hour (in a day).

- $\xi_{j,m}$ represents an unobserved characteristic of payments system j in market m. It is important to include this system/market level demand shock in the model because it is impossible to include all of the factors that affect the demand for option j in market m; see, among others, Berry et al. (1995) for detailed justifications for the importance of this term. In our context, ξ may include factors that are not captured by $LCI_{j,m}$, $SI_{j,m}$,$\bar{s}_{j,m}$, and X_m; e.g., the service level of a payment system, and the operational and legal risks that participants face in a payments system, etc.

- $\zeta_{i,g,m}+(1-\lambda)\varepsilon_{i,j,m}$ is the preference shock that follows the nested logit structure, with $\lambda \in [0,1)$ being the nesting parameter to be estimated. In particular $\zeta_{i,g,m}$ is an extreme value random variable that captures the interaction between the decision maker who is sending payment i and the nested group of payment systems, $\{0\}$ and $\{T1, T2\}$, which are labeled $g = 0$ and $g = 1$, respectively. And $\varepsilon_{i,j,m}$ is an extreme value variable and is identically and independently distributed.[12]

The nested logit structure is important for modeling the substitution patterns between the three options in the choice set. The nesting parameter λ in this model can be interpreted as a measure of the substitutability of the alternative payments systems across groups.[13] As the parameter λ approaches one, the within-group correlation of payoff levels goes to one. On the other hand, as λ approaches zero, the within-group correlation goes zero and, thereby, the nested logit model is reduced to the logit model. In our case, the substitution between T1 and T2 is naturally stronger than that between T1 (or T2) and the outside option. Therefore, the above specified nested logit model allows us to model a more realistic substitution pattern among the three options in the choice set, compared to the simple logit model, which suffers from the well-known drawback of the "independence of irrelevant alternatives" property.

Given the random payoff in (4), we denote the mean utility level of payment system j in market m as

$$\delta_{j,m} = \alpha LCI_{j,m} + \beta SI_{j,m} + \gamma \bar{s}_{j,m} + X_m \rho + \xi_{j,m}, \tag{5}$$

[12] In the nested logit specification, the variable $\zeta_{i,g,m}$ is a common variable to all systems in g. Cardell (1997) shows that the distribution of $\zeta_{i,g,m}$ is the unique distribution with the property where if $\varepsilon_{i,j,m}$ is an extreme value random variable, then $\zeta_{i,g,m} + (1-\lambda)\varepsilon_{i,j,m}$ is also an extreme value random variable.

[13] See, among others, Trajtenberg (1989) for details.

which will play an important role in deriving the market shares of payments system j. Note that the equation is also in regression form and will be convenient for estimation. Aggregating individual choices for each system j in group g, we obtain the within-group share of g in market m,

$$s_{j|g,m} = \frac{e^{\delta_{j,m}/(1-\lambda)}}{D_g}, \tag{6}$$

where $D_g = \sum_{j \in G_g} e^{\delta_{j,m}/(1-\lambda)}$ and G_g includes options in group g.

Similarly, we can obtain the market share of each payment system j in the choice set $\{T1, T2, 0\}$:

$$s_{j,m} = \frac{e^{\delta_{j,m}/(1-\lambda)}}{D_g^{\lambda}[\sum_g D_g^{(1-\lambda)}]}. \tag{7}$$

Based on (5), (6), and (7), with $\delta_{0,m}$ being normalized to zero, we apply the well-known choice probability inversion formula of the nested logit model (see Berry (1994)) to obtain the following regression equation:

$$\log\left(\frac{\hat{s}_{j,m}}{\hat{s}_{0,m}}\right) = \alpha LCI_{j,m} + \beta SI_{j,m} + \gamma \bar{s}_{j,m} + \lambda \log\left(\hat{s}_{j|g,m}\right) + X_m \rho + \xi_{j,m}, \tag{8}$$

where $\hat{s}_{j,m}$ is the observed market share of j in market m and $\hat{s}_{j|g,m}$ is the observed within-group market share of j in market m. Equation (8) forms the basis of estimating the parameters in the model.

To estimate (8), we need to impose statistical assumptions on the demand shock, $\xi_{j,m}$. Following Berry (1994) and Berry et al. (1995), we assume that the following mean independence condition holds:

$$E\left[\xi_{j,m}|Z_{j,m}\right] = 0, \tag{9}$$

where $Z_{j,m}$ is a set of exogenous variables that do not depend on $\xi_{j,m}$.

The two market share variables on the RHS, $\bar{s}_{j,m}$ and $\log\left(\hat{s}_{j|g,m}\right)$, in (8) are clearly endogenous because they depend on the market shares on the LHS. Thus we construct instrumental variables for them below. Regarding the other RHS variables, X_m is determined by the exogenous payment demand from outside the system and thus is unlikely to be correlated with $\xi_{j,t}$; both $LCI_{j,m}$ and $SI_{j,m}$ can depend on the market shares in market m as well as on their neighbors (and thus ξ's market share); however, we think such dependence is weak after directly controlling for $\bar{s}_{j,m}$ (assuming the dependence of $LCI_{j,m}$ and $SI_{j,m}$ on market shares mostly goes through $\bar{s}_{j,m}$). In this sense, from a purely econometric point of view, the inclusion of $\bar{s}_{j,m}$ in our model can be regarded as a way to control for the endogeneity of $LCI_{j,m}$ and $SI_{j,m}$. Hence, we treat $LCI_{j,m}$ and $SI_{j,m}$ as exogenous variables and only handle the endogeneity problem in the market share variables.

Specifically, for the two endogenous variables $\bar{s}_{j,B_m}, \log\left(\hat{s}_{j|g,m}\right)$, we construct the following instrumental variables:

$$\left[\bar{s}_{j,B_m'\setminus B_m}, \; \frac{1}{|\mathcal{M}_m|}\sum_{l\in\mathcal{M}_m}\log\left(\hat{s}_{j|g,l}\right)\right], \tag{10}$$

where $B_m^{'}$ is a superset of B_m (a bigger neighbor of market m), $\mathcal{M}_m$ is a set of markets "adjacent" to market m (excluding m itself); i.e., those having the same sender, submission hour, and value percentile, but different receivers. The construction of these instrumental variables is based on the following two ideas. First, the market shares of T1 and T2 in neighboring markets are informative about those in the market being considered (i.e., for their relevance). Second, the demand shocks in neighboring markets have limited correlation with those in the market in question (i.e., exogeneity). With the constructed instrument variables in (10) (along with other exogenous variables), we can estimate (8) using the standard two-stage least-squares method.

6 Estimation Results

Table 1 shows the estimation results of the discrete choice model introduced in the previous section. We consider four alternative specifications: a simple logit specification, a simple logit specification with fixed effects, a nested logit specification, and a nested logit specification with instrumental variables. For convenience, we refer to the first as “Logit,” the second as “FE Logit,” the third as “Nested Logit,” and the fourth as “IV Nested logit.” The advantage of presenting the logit results is that we can explore what occurs when we control for the fixed effects. We present the nested logit results to examine the effectiveness of the instrumental variables for the log of the within-group share $\left(s_{j|g,m}\right)$ in the nested logit specification.

In the first column of Table 1, we report the estimation results of the simple logit model. Although the coefficients of the safety indicator and the network effect are of the expected sign, the positive coefficient on the liquidity cost is an anomaly because we would expect the liquidity cost to yield a negative marginal payoff. Additionally, the logit model gives us an adjusted R^2 of 0.712, which implies that about 29% of the variance in the mean utility level is due to the unobserved characteristics. In the second column of Table 1, we report the estimation results of the FE logit model. All of the coefficients are significantly different from zero and have the expected sign. Also R^2 from the FE Logit model is fairly high at 0.903, a significant improvement from the logit model. All of these suggest that controlling for fixed effects is critical to obtaining reasonable estimations of the coefficients.

In Table 1, the third column reports the estimation results of the nested logit specificiation using ordinary least squares (OLS), while the fourth column reports the results of IV Nested Logit, using the instrument variables in (10). Comparing the estimation results in the third column with those in the fourth column, we find that the coefficients between the two columns differ noticeably, especially the coefficient on the network effects in the fourth column, which decreases substantially. This indicates the importance of correcting the endogeneity problem.

Focusing on the IV Nested Logit model, our favored specification, we can see that the coefficients on the liquidity cost, safety indicator, network effect, and log of the within-group market share all have the expected signs and are statistically significant. Also, the nesting parameter; i.e., the coefficient of the log of the within-group market share, is greater than 0.7, indicating a strong within-group correlation between the preferences for *T1* and *T2*, relative to the cross-group correlation between *T1* (or *T2*) and outside option 0. Also, the statistically significant coefficient of the log of the within-group market share suggests that the extension from the FE Logit to the IV Nested Logit seems necessary. Finally, the large value of the first-stage F test statistic supports the relevance of the constructed instruments.

7 Quantifying the Economic Benefits of Lynx

In this section, we use the estimated discrete choice model from the previous section to calculate the economic benefits of Lynx, which is replacing the LVTS as part of the payments modernization initiative.

7.1 Welfare Calculation

The discrete choice model allows us to calculate the welfare, or economic benefits, to participants from sending payments. Given our IV Nested Logit specification, using a similar approach as in Trajtenberg (1989), who suggests the use of discrete choice models to measure the benefits of product innovations, the expected maximum payoff of sending a payment in market m under the current LVTS regime is calculated as follows:

$$EV_{\text{LVTS,m}} = \log\left[1 + \left(\exp\left(\frac{\hat{\delta}_{\text{T1},m}}{1-\hat{\lambda}}\right) + \exp\left(\frac{\hat{\delta}_{\text{T2},m}}{1-\hat{\lambda}}\right)\right)^{1-\hat{\lambda}}\right], \qquad (11)$$

where $\hat{\delta}_{\mathrm{T1},m}$ and $\hat{\delta}_{\mathrm{T2},m}$ are the fitted values of the mean utilities defined as

$$\hat{\delta}_{j,m} = \hat{\alpha} LCI_{j,m} + \hat{\beta} SI_{j,m} + \hat{\gamma}\bar{s}_{j,B_m} + X_m\hat{\rho} + \hat{\xi}_{j,m}, j \in \{T1, T2\}. \tag{12}$$

Note that the estimated parameters in (11) and (12), $\hat{\lambda}$, $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$, $\hat{\rho}$, and $\hat{\xi}_{j,m}$, are obtained from the estimation results in Section 6.

After the LVTS is replaced by Lynx, the choice set of the payments systems the participants are facing becomes $\{0, \mathrm{Lynx}\}$. And the expected maximum payoff of sending a payment in market m boils down to

$$EV_{\mathrm{Lynx},m} = \log\left[1 + \exp\left(\hat{\delta}_{\mathrm{Lynx},m}\right)\right], \tag{13}$$

where $\hat{\delta}_{\mathrm{Lynx},m}$ is the mean payoff of sending a payment through Lynx and can be decomposed as

$$\hat{\delta}_{\mathrm{Lynx},m} = \hat{\alpha} LCI_{\mathrm{Lynx},m} + \hat{\beta} SI_{\mathrm{Lynx},m} + \hat{\gamma}\bar{s}_{\mathrm{Lynx},B_m} + X_m\hat{\rho} + \hat{\xi}_{\mathrm{Lynx},m}. \tag{14}$$

Note that $\hat{\alpha}$, $\hat{\beta}$, $\hat{\gamma}$ and $X_m\hat{\rho}$ in (14) are the same as in (12); however, other variables in (14) are unknown and we will discuss how we assign their values in the next subsection.

Given the expected maximum payoffs of both the LVTS and Lynx in market m, we can use the difference in the required amount of collateral (i.e., in Canadian dollars) to measure the welfare change from replacing the LVTS with Lynx; i.e.,

$$\Delta EB = \sum_m \mathrm{W}_m \cdot \mathrm{Vol}_m \cdot \left(EV_{\mathrm{Lynx},m} - EV_{\mathrm{LVTS},m}\right), \tag{15}$$

where Vol_m is the total payments volume in market m, and W_m is a factor that translates the payoff level of each payment made in a payments system into the "willingness-to-pay," which is measured in terms of collateral (in Canadian dollars).

Since the liquidity cost enters the payoff function in log form, we can define the factor as

$$\mathrm{W}_m = \frac{\exp\left(LCI_{T1,m}\right) + \exp\left(LCI_{T2,m}\right)}{|\hat{\alpha}|}. \tag{16}$$

It is possible to further translate the collateral amount into the actual financial cost (e.g., using the appropriate interest rates); however, this is not straightforward because it can be tricky to measure collateral costs that vary by participant and timing, etc. As this is less relevant to our current analysis, we leave it for future investigation.[14]

7.2 Payoff Function of Lynx

Since Lynx had not been launched by the time of this analysis, we do not have realized data on the new system. Hence, we have to use simulation to construct synthetic data for Lynx. Simulating a full-scale Lynx with all of its design features is challenging, so we focus on a simplified version of Lynx with only its core characteristics; i.e., this is a real-time settlement system where each payment must be fully collateralized in order to be processed.

In particular, we let all of the payment instructions from the 2019 LVTS data that includes both T1 and T2 run through the simplified Lynx and we record the outcome; e.g., we evaluate the cost and safety indicators in this new system. In this process, we maintain the following assumptions: (1) all of the LVTS payments migrate to the Lynx system; (2) the timing and order of all of the payments for each day remain unchanged; (3) the participants pledge sufficient collateral to Lynx so they can make all of their payments in a fully collateralized manner; and (4) the risk control test for each payment in Lynx is the same as in T1.

Figure 4 shows a comparison between the LVTS and Lynx, in terms of the liquidity cost, for different hours in a day and different sending participants. As expected, the

[14] See McPhail and Vakos (2003) for some high-level estimates of the collateral costs associated with using the LVTS.

liquidity cost of Lynx is higher than that of the LVTS. Also, the difference is rather heterogeneous across participants and different hours. For example, the participants with higher overall liquidity costs, which are mostly driven by the larger overall volume of their payments, have to bear greater increases in their liquidity costs when Lynx replaces the LVTS.

Figure 5 plots the safety indicators for Lynx and the LVTS, respectively, across different participants and hours in a day. It is clear that the safety indicator for Lynx is higher than that for the LVTS, mostly due to the higher credit limits (i.e., combining T1 and T2) and the removal of the bilateral risk control test for T2 (reflecting the reduction in credit risk). Compared with the liquidity cost, the safety indicator is much smaller in scale and thus its difference between Lynx and the LVTS will have less effect on participants' welfare changes.

With the above simulation, we obtain values for the liquidity and safety indicators, $LCI_{\text{Lynx},m}$ and $SI_{\text{Lynx},m}$, in the mean payoff function of Lynx (14). Regarding the unobserved characteristics of Lynx, $\hat{\xi}_{\text{Lynx}}$, which captures the relative payoff (e.g., the service quality level) compared to the outside option, we assume that this is proportional to the average of its counterparts to T1 and T2; i.e.,

$$\hat{\xi}_{\text{Lynx}} = \frac{\theta_1}{2}\left(\hat{\xi}_{\text{T1}} + \hat{\xi}_{\text{T2}}\right),$$

where θ_1 is a tuning parameter that captures the unobserved service quality level of Lynx (compared to the LVTS).

Finally, the model includes a network effect that is captured by $\bar{s}_{j,m}$ in (4) that needs to be determined for Lynx. We consider two ways of assigning this value. First, we simply assume that it is proportional to the total market share of the LVTS; i.e.,

$$\bar{s}_{\text{Lynx},m} = \theta_2\left(\bar{s}_{\text{T1},m} + \bar{s}_{\text{T2},m}\right),$$

where θ_2 represents the fraction of the LVTS payments that migrate to Lynx. Second,

$\bar{s}_{\text{Lynx},m}$ can be endogenously determined in the equilibrium adjustment of market shares. Specifically, with the new system Lynx replacing the LVTS, the model implies that the market shares of Lynx differ from those of the LVTS. Then the network effect term, which depends on the market shares in neighboring markets, needs to be updated. This change in turn implies a vector of the new market shares of Lynx. This heuristic adjustment process can be formalized as an iterative approach to computing the new equilibrium market shares of Lynx. Formally, the equilibrium market shares of Lynx can be computed iteratively as follows:

$$s^{r+1}_{\text{Lynx},m} = \frac{\exp\left(\hat{\alpha}LCI_{\text{Lynx},m} + \hat{\beta}SI_{\text{Lynx},m} + \hat{\gamma}\bar{s}_{\text{Lynx},m}\left(\mathbf{s}^{r}_{\text{Lynx}}\right) + X_m\hat{\rho} + \xi_{\text{Lynx},m}\right)}{1 + \exp\left(\hat{\alpha}LCI_{\text{Lynx},m} + \hat{\beta}SI_{\text{Lynx},m} + \hat{\gamma}\bar{s}_{\text{Lynx},m}\left(\mathbf{s}^{r}_{\text{Lynx}}\right) + X_m\hat{\rho} + \xi_{\text{Lynx},m}\right)}, \tag{17}$$

where $\mathbf{s}^{r}_{\text{Lynx}}$ is the vector of market shares in the rth iteration (with some starting value $\mathbf{s}^{0}_{\text{Lynx}}$). This iteration process is, in fact, a contraction mapping, given the value of our estimated parameter $\hat{\gamma} = 1.549$; see Brock and Durlauf (2001) for details.

7.3 Results

We first consider the simple case where $\theta_1 = 1$ and there is no equilibrium adjustment; i.e., by treating the network effect term as exogenous and examining how it contributes to the welfare measure. Figure 6 shows the total welfare change against θ_2, the fraction of payments in the LVTS that migrate to Lynx. There are three lines in the graph: the "Baseline" refers to the simple Lynx described above; the "20% Cost Reduction" is the case where we assume that Lynx adopts a certain liquidity saving mechanism such that its liquidity cost is overall 20% lower than the baseline simple Lynx and the "100% Safety Increase" refers to the case where Lynx improves the safety indicator by 100%.

We can see that as θ_2 increases, the net economic benefits of Lynx also increase and exceed 0 when θ reaches around 0.9, which means that there is a welfare gain when over 90% of the current LVTS payments migrate to Lynx. Also, lowering the liquidity cost by 20% has a non-negligible effect on welfare, but increasing the payment safety indicator can do very little.

Next, we consider the case with an equilibrium adjustment. It turns out the new equilibrium implies a migration ratio that is lower than 90%, around 55%.[15] Under this lower migration ratio, the baseline Lynx is likely to cause a welfare loss to participants; however, a quality improvement (e.g., an increase in $\hat{\xi}_{\text{Lynx}}$) can mitigate the loss and even generate welfare gains. Then a natural question is, how much quality improvement is needed? Figure 7 illustrates the overall welfare changes against θ_1, the percentage of the quality improvement. We can see that Lynx needs an almost 75% improvement in its service level, over the LVTS, to achieve an overall welfare gain for the participants. Again, lowering the liquidity cost (e.g., through liquidity saving mechanisms) has much larger effects than improving the payment safety indicator.

Besides the overall welfare change, in Figure 8 we show the heterogeneous effects across participants for the special case where $\theta_1 = 1.75$ (with an equilibrium adjustment); i.e., the overall welfare change is 0. We can see that the welfare changes are rather heterogeneous across participants; e.g., Bank 1 (a big participant) enjoys quite a large welfare gain, while the small participants as a whole (labeled as "Others") suffer from a welfare loss. The differences are mainly driven by the heterogeneous effects of replacing the LVTS with Lynx on the random payoff functions across different participants. This heterogeneous effect raises some interesting policy questions; e.g., the central bank and the payments system operator may consider providing certain incentives for some participants that suffer from welfare losses as a result of participating in the new system, given that participation is vital for the new system to achieve its public objectives.

[15] This result is in the same ballpark as the payment modernization patterns predicted by Kosse et al. (2021). However, recall that this result is obtained under the assumption that Lynx is a pure RTGS system, which is not exactly the case because of the adoption of liquidity saving mechanisms in Lynx. Hence, we expect that this predicted migration ratio might be a reasonable "lower bound" for the actual migration ratio that will be realized in the near future.

8 Concluding Remarks

In this paper, we propose a discrete choice demand framework to model the participants' decisions on which payment system to use and we apply it to measure the benefits of payments modernization in Canada. Focusing on the modernization of the large-value payments system, we first use historical data of the LVTS to estimate participants' preferences regarding liquidity cost and payment safety, as well as the network effect; then with the estimated preferences, we use a counterfactual simulation to calculate participants' welfare change when the LVTS is replaced by Lynx. Our results suggest that a high migration ratio and/or improvements in service quality (e.g., new liquidity saving/safety features) are crucial to generating overall net economic benefits to participants.

Our study is the first to quantify the economic benefits of a payments system modernization. There are several caveats in our current results:

- We construct and include only two indicators that describe the incentives participants face when making a payment. Although these indicators capture two important factors, it is clear that this approach has not exhausted all of the important considerations participants encounter when making payments. Including more variables that capture the incentives participants face can improve our results, although this is not an easy task.

- Our measure of the "outside option" (relative to the LVTS) is based on ACSS data. This assumption may underestimate the share of the actual outside option, especially for small-value payments, because other payments systems/networks are available for financial institutions to process payments. Whether this would change our conclusion is not clear. For example, if Lynx has certain new features that can "steal" market shares from other payments systems, then we would underestimate the welfare gain of the new system.

- Our evaluation focuses on the modernization of the large-value payments system and, thus, only covers part of the payments ecosystem. Hence, the welfare results should be interpreted with caution. For example, even if Lynx cannot generate a welfare gain when replacing the LVTS, other modernized payments systems; e.g., the new retail payments system, may produce a sufficiently high surplus to make the overall benefits of payments modernization positive.

- We only measure the economic benefits to the participants in the system and do not consider the overall potential benefit or loss to all of society. Given that payments systems have clear externalities in terms of systemic risk, it is important to extend the current analysis along this line.

Despite the above mentioned caveats, which are mostly driven by data limitations, our proposed framework for evaluating the economic benefits of payments systems is rather general and flexible. With more and richer data in the future, especially data generated by Lynx, the framework could be extended and enriched to provide more-comprehensive assessments of the welfare implications of payments modernization.

References

Arjani, Neville, "The economic benefit of adopting the ISO 20022 payment message standard in Canada," Technical Report, Canadian Payments Association 2015.

_ **and Darcey McVanel**, "A primer on Canada's large value transfer system," Technical Report, Bank of Canada 2006.

_ **, Fuchun Li, and Leonard Sabetti**, "Monitoring Intraday Liquidity Risks in a Real Time Gross Settlement System," Technical Report, Payments Canada Discussion Paper 2020.

Bech, Morten L. and Rod Garratt, "The intraday liquidity management game," *Journal of Economic Theory*, April 2003, *109* (2), 198–219.

Berry, S., J. Levinsohn, and A. Pakes, "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 1995, pp. 841–890.

Berry, S.T., "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, 1994, pp. 242–262.

Brock, William A and Steven N Durlauf, "Discrete choice with social interactions," *The Review of Economic Studies*, 2001, *68* (2), 235–260.

Byck, Shaun and Ronald Heijmans, "How much liquidity would a liquidity-saving mechanism save if a liquidity-saving mechanism could save liquidity? A simulation approach for Canada's large-value payment system," *Journal of Financial Market Infrastructures*, 2021, *9* (3).

Cardell, N. Scott, "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity," *Econometric Theory*, 1997, *13* (2), 185–213.

CPMI, "Liquidity efficiency of the large value payment systems around the world," Technical Report 2015.

Kosse, Anneke, Zhentong Lu, and Gabriel Xerri, "An economic perspective on payments migration," Technical Report, Bank of Canada 2020.

_ , _ , and _ , "Predicting Payment Migration in Canada," *Journal of Financial Market Infrastructure*, 2021.

Lopez, Jorge Cruz, Charles M. Kahn, and Gabriel Rodriguez Rondon, "Joint Determination of Counterparty and Liquidity Risk in Payment Systems," Technical Report, The University of Western Ontario 2021.

McFadden, D., "Econometric Models of Probabilistic Choice," in C. Manski and D. McFadden, eds., *Structural Analysis of Discrete Data with Econometric Applications*, Cambridge: MIT Press, 1981, chapter 5, pp. 198–272.

McPhail, Kim and Anastasia Vakos, "Excess Collateral in the LVTS: How Much is Too Much?," Technical Report, Bank of Canada 2003.

Petrin, Amil, "Quantifying the benefits of new products: The case of the minivan," *Journal of political Economy*, 2002, *110* (4), 705–729.

Small, Kenneth A and Harvey S Rosen, "Applied welfare economics with discrete choice models," *Econometrica: Journal of the Econometric Society*, 1981, pp. 105–130.

Trajtenberg, Manuel, "The welfare analysis of product innovations, with an application to computed tomography scanners," *Journal of Political Economy*, 1989, *97* (2), 444–479.

Tables

Table 1: Demand Estimation Results

	Simple Logit		Nested Logit	
	Without FE	With FE	Without IV	With IV
Liquidity Cost	0.564	-0.0443	-0.0220	-0.0299
	(0.00250)	(0.00467)	(0.00440)	(0.00438)
Safety Indicator	0.0154	0.0246	0.0264	0.0202
	(0.00248)	(0.00187)	(0.00181)	(0.00180)
Network Effect	6.191	9.788	6.001	1.549
	(0.0175)	(0.260)	(0.223)	(0.117)
Nesting Parameter			0.515	0.724
			(0.00775)	(0.0218)
Constant	-8.140	-7.082	-5.262	-4.522
	(0.0335)	(0.130)	(0.123)	(0.157)
Sender FE		✓	✓	✓
Receiver FE		✓	✓	✓
Hour FE		✓	✓	✓
Value Pctile FE		✓	✓	✓
Cragg-Donald Wald F				7869.96
# Obs.	104,707	104,707	104,707	100,350
Adj. R^2	0.712	0.903	0.909	0.913

Note: Table 1 shows the estimation results of the discrete choice model with four alternative specifications. These specifications are a simple logit specification without fixed effects, a simple logit specification with fixed effects, a nested logit specification without IV, and a nested logit specification with IVs. The standard errors of the estimated parameters are reported in parentheses. All of the estimates shown in the table are statistically significant at the 1% level.

Figures

Figure 1: Volume Share of T1 for Different Payments Sizes

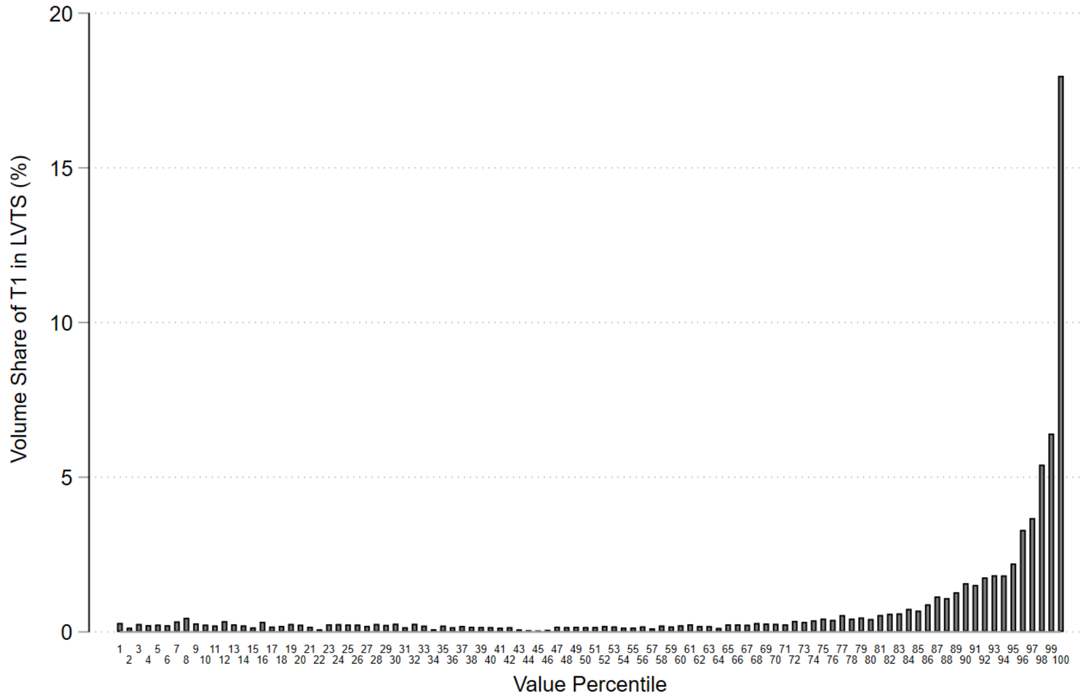

Note: Figure 1 plots the volume share of T1 in each of 100 groups of payments, where the groups are defined by the percentiles of the value distribution of all of the LVTS payments in 2019. In each group, the volume share is calculated as the ratio of the total volume of payments sent through T1 to the total volume sent through the LVTS.

Figure 2: Value Share of T1 for Different Hours

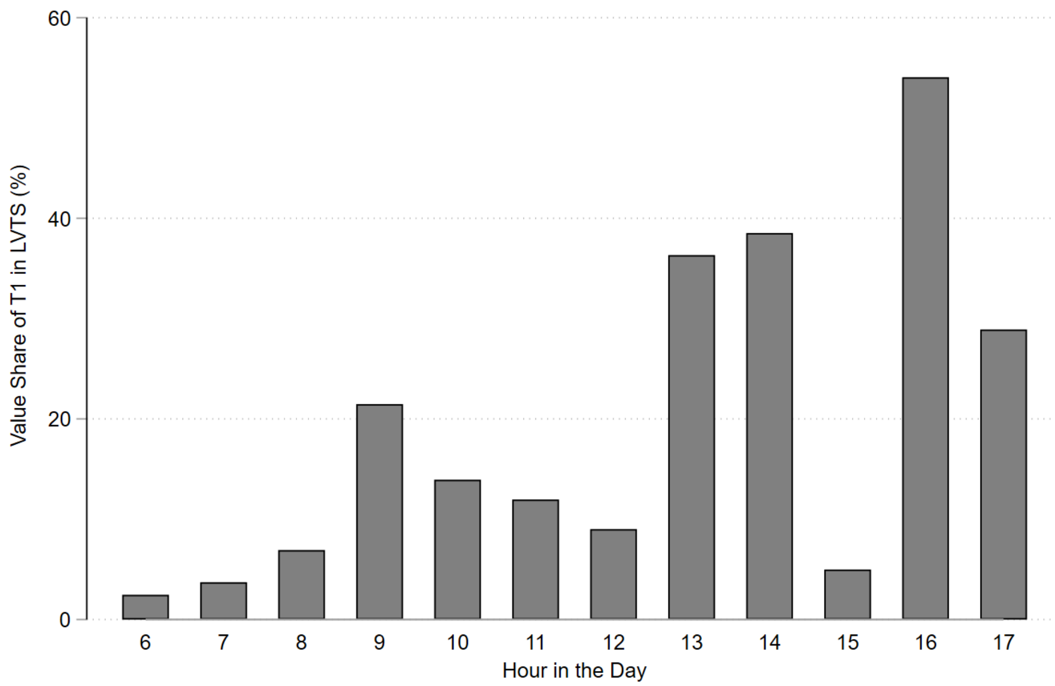

Note: Figure 2 plots the average hour-by-hour pattern of the value of the shares of payments made through T1 in a day. In each hour, the value of the shares is calculated as the ratio between the total value of payments sent through T1 and the total value of payments made through the LVTS.

Figure 3: Volume of the Share of the Outside Option

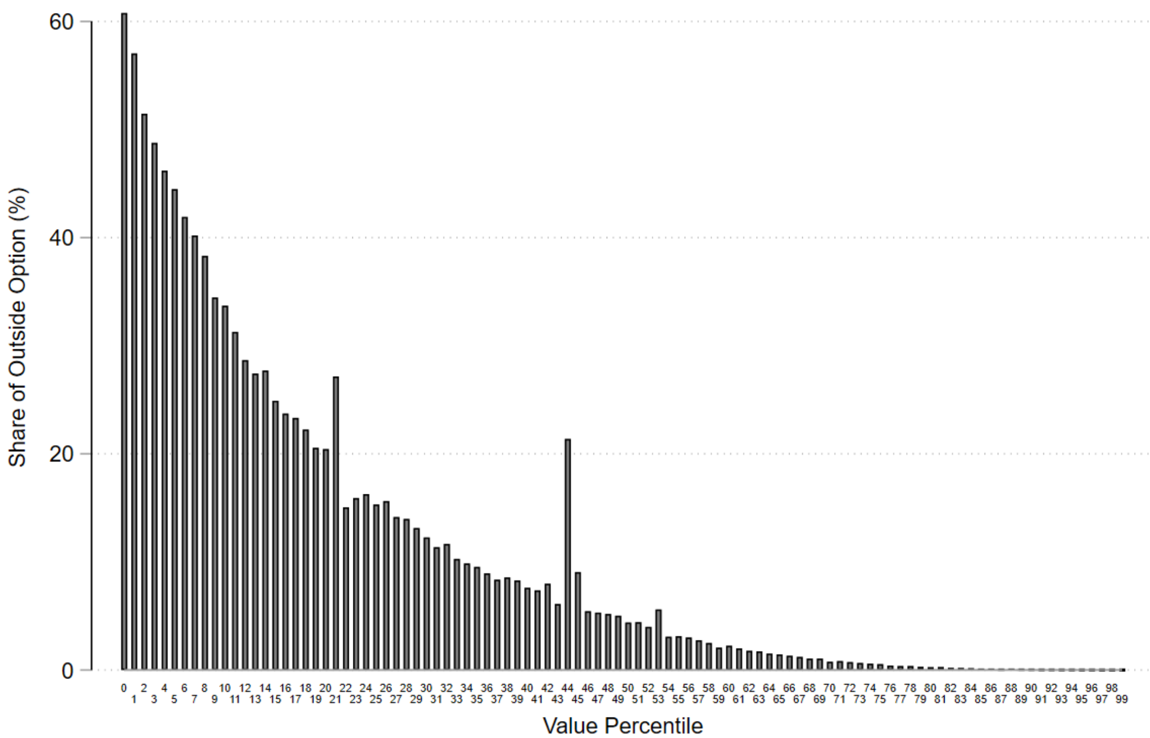

Note: Figure 3 shows the share of the normalized volume of ACSS in 2019 (relative to the LVTS and ACSS combined) for each value percentile bin.

Figure 4: Liquidity Cost: LVTS vs Lynx

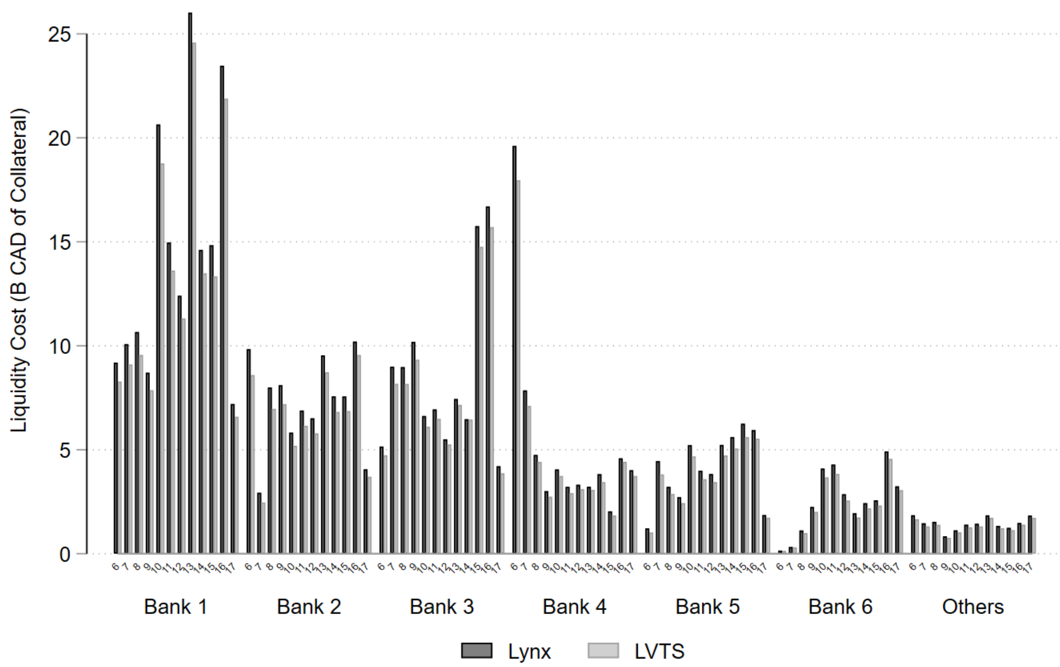

Note: Figure 4 shows the comparison of the liquidity cost between the LVTS and Lynx in 2019, for different hours in a day and different sending participants.

Figure 5: Safety Indicator: LVTS vs Lynx

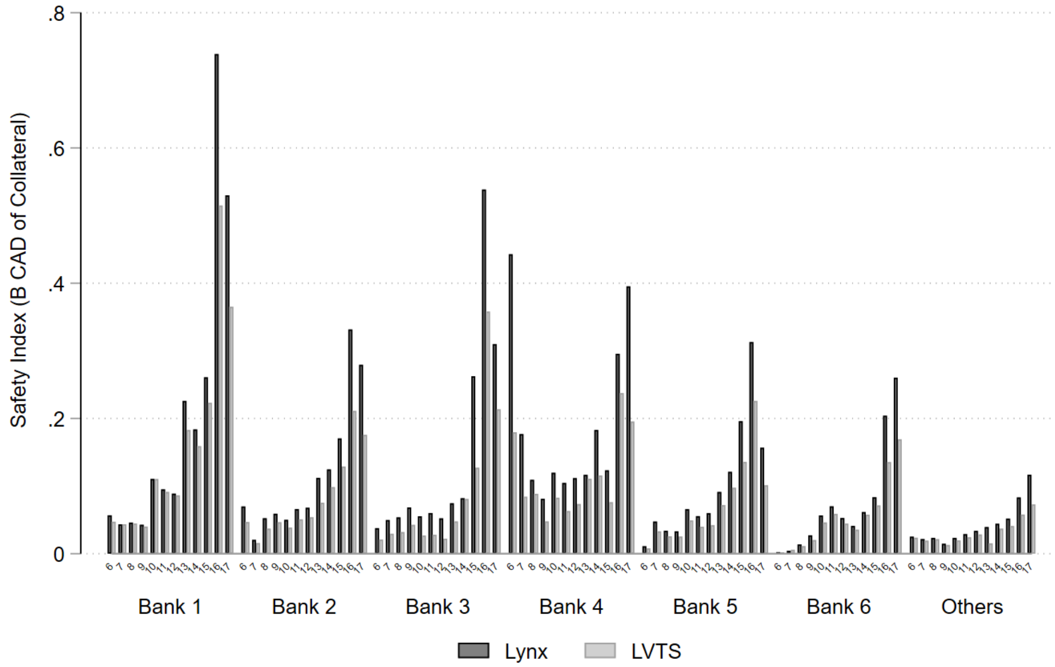

Note: Figure 5 plots the safety indicator for Lynx and the LVTS, respectively, across different participants and hours in a day.

Figure 6: Total Welfare Change: No Equilibrium Adjustment

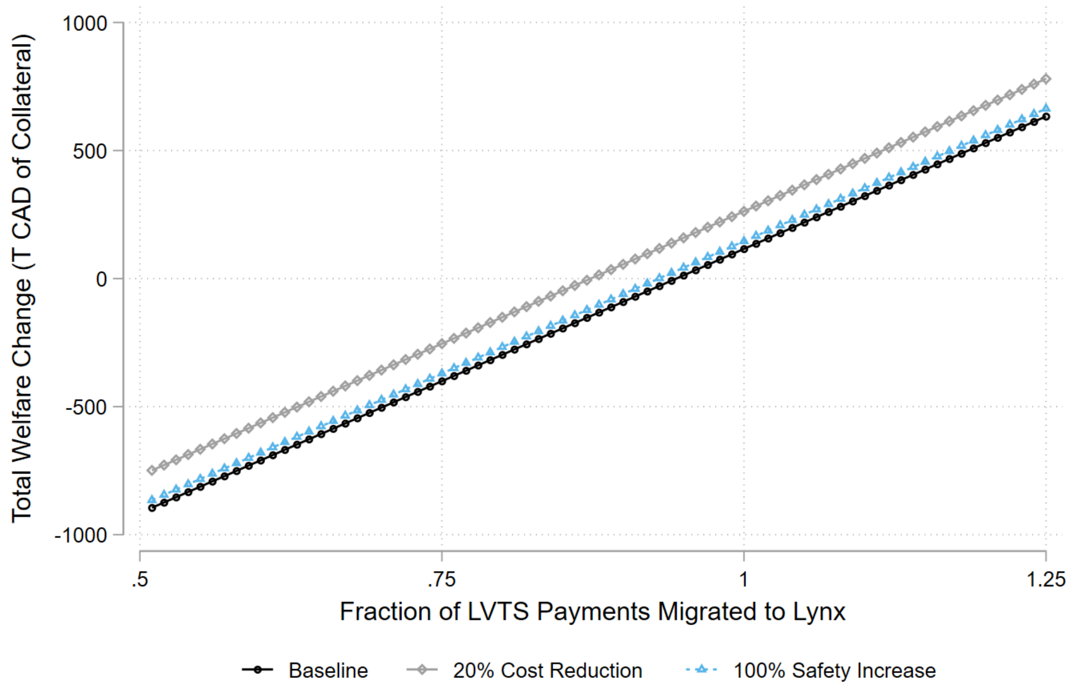

Note: Figure 6 shows the total welfare change against θ_2, the fraction of payments in the LVTS migrating to Lynx. There are three lines in the graph: "Baseline" refers to the simple Lynx described in the main text; "20% Cost Reduction" is the case where we assume that Lynx adopts a certain liquidity saving mechanism such that its liquidity cost is 20% lower overall than the baseline simple Lynx; "100% Safety Increase" refers to the case where Lynx improves the safety indicator by 100%.

Figure 7: Total Welfare Change: With Equilibrium Adjustment

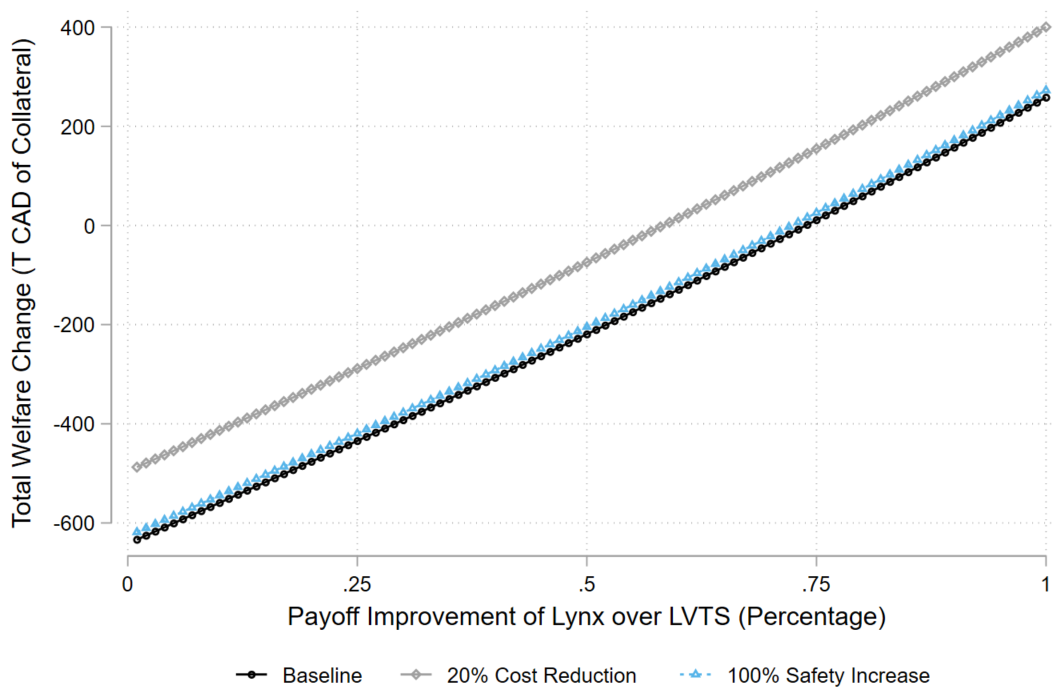

Note: Figure 7 illustrates the overall welfare changes against θ_1, which is a tuning parameter that captures the improvement in the service level of Lynx.

Figure 8: Welfare Change: Heterogeneity Across Banks

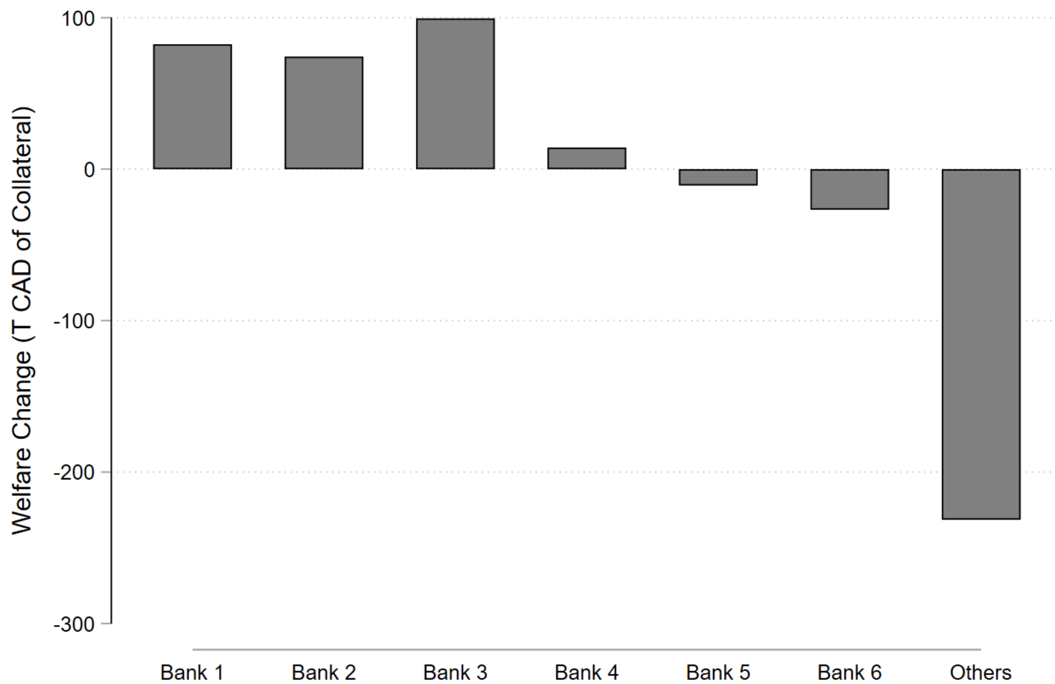

Note: Figure 8 shows the heterogeneous welfare changes across participants for a given level of the tuning parameter that captures the service quality of Lynx; i.e., $\theta_1 = 1.75$, for which the overall welfare change is 0.