Bitcoin Adoption and Beliefs in Canada

A Comment on Daniela Balutel, Christopher Henry, Jorge Vásquez, and Marcel Voia (*Canadian Journal of Economics*, 2022)

Leo Van Hove *

Journal of Comments and Replications in Economics, Volume 4, 2025-6, DOI: 10.18718/81781.45

JEL: D83, O33
Keywords: Bitcoin, Adoption, Network effects

Data Availability: The Excel file to reproduce the results of this comment can be downloaded at JCRE's data archive (DOI: 10.15456/jl.2025169.1141885422).

Please Cite As: Van Hove, L (2025). Bitcoin Adoption and Beliefs in Canada. A Comment on Daniela Balutel, Christopher Henry, Jorge Vásquez and Marcel Voia (*Canadian Journal of Economics* , 2022). *Journal of Comments and Replications in Economics*, Vol.4 (2025-6). DOI: 10.18718/81781.45

Abstract
This comment shows that the empirical part of Balutel et al.'s (2022a) article suffers from an internal validity problem. Specifically, their network size variable is problematic. Hence, Balutel et al. fail to provide reliable evidence that network effects impact Bitcoin adoption.

*FIRE Research Group and Department of Applied Economics (APEC), Vrije Universiteit Brussel, Leo.Van.Hove@vub.be

Received October 10, 2024; Revised March 20, 2025; Accepted June 16, 2025; Published December 05, 2025.

1 Introduction

In a recent article, Balutel et al. (2022a) show, for the case of Canada, that individuals' probability of Bitcoin adoption depends on network and learning effects. This comment argues that Balutel et al.'s empirical analysis has several issues. In particular, their network size variable is problematic. Hence, while network effects may well play a role in Bitcoin adoption, Balutel et al.'s results do not demonstrate this.

This comment unfolds as follows. In Section 3, I proffer a number of conceptual remarks about the way in which Balutel et al. estimate Bitcoin network effects. Subsequently, in Section 4, I explain why the operationalisation of network size is critically compromised. But first I summarise (the relevant parts of) Balutel et al.'s method and results.

2 Balutel et al.'s approach

In the first part of their article, Balutel et al. develop a novel theoretical model of Bitcoin adoption with network effects and social learning. In the second part, they then test this model with unique, detailed data from the Bank of Canada's annual Bitcoin Omnibus Survey (BTCOS) for the years 2017 and 2018.

The BTCOS collects information on Canadians' awareness and ownership of Bitcoin; in addition, the survey also asks respondents about their reasons for owning or not owning Bitcoin, their knowledge of Bitcoin features, their beliefs about its survival, etc. In 2017, sample size was 2,623, of which 117 individuals (or 4.7%) reported that they owned Bitcoin. In 2018, respondents numbered 1,987 and Bitcoin owners 99 (5.0%) (o.c., p. 1740).

The theoretical model in the first part of Balutel et al.'s article suggests a specification of the following form for the empirical estimates (o.c., p. 1742):

$$a_{it} \ = \ F\Big(\beta_0 + \beta_1 A_{it} + \beta_2 \xi_{it} + \beta_3 A_{it}\xi_{it} + \beta_c X_i\Big) + \varepsilon_{it}, \tag{1}$$

where a_{it} is a dummy for Bitcoin adoption of individual i on their survey day, F denotes a logistic cumulative distribution function, and A_{it} and ξ_{it} measure, respectively, network size (i.e., adoption by peers) and Bitcoin beliefs. β_1 and β_2 supposedly capture, respectively, direct network effects and individual learning effects; the coefficient on the interaction term $A_{it}\xi_{it}$ captures social learning effects. The control variables X_i consist of multiple demographic characteristics, as well as regional dummies.

However, Balutel et al. point out, with reason, that there is a potential simultaneity problem with specification (1): individuals with high beliefs are more likely to adopt, but, at the same time, adopters are likely to develop higher beliefs. To address this, Balutel et al. employ an identification strategy with a control function. In the first stage of this approach, they model Bitcoin beliefs as a function of demographic characteristics together with an exclusion restriction that comes from the supply side, namely the regional growth in Bitcoin automated teller machines (ATMs):

$$\xi_{it} \ = \ \alpha_0 + \alpha_1 \Delta ATM_{jt} + \alpha_2 \text{Age}_i^2 + \alpha_c X_i + u_{it}, \tag{2}$$

with ΔATM_{jt} the growth in Bitcoin ATMs in region j at time t. In the second stage, the residual from the first stage is used as a control function. That is, (1) is augmented with CF_{it} as follows (o.c., p. 1747):

$$a_{it} \ = \ F\Big(\beta_0 + \beta_1 A_{it} + \beta_2 \xi_{it} + \beta_3 A_{it}\xi_{it} + \beta_c X_i + \beta_{CF} CF_{it}\Big) + \varepsilon_{it}, \tag{3}$$

Equation (3) is then estimated with a logit-based likelihood.

Balutel et al. find that both network size and beliefs (as a proxy for individual learning effects) have a positive direct impact on Bitcoin adoption. The coefficient on beliefs is significant at the 1% level in all models. For network size, the relationship is significant at the 10% level in 2017, and at 1% in 2018. Finally, social learning attenuates the effect of network size on Bitcoin adoption; that is, a marginal increase in network size raises the likelihood of adoption less when beliefs are high. For a discussion of the results for the demographic characteristics — results that are not commented upon below — I refer to Balutel et al. (2022a, pp. 1750–1752).

3 Conceptual remarks

A first, high-level conceptual question that can be raised concerning Balutel et al.'s research is whether it can be justified to model, as Balutel et al. do, Bitcoin adoption primarily from a payment-instrument point of view. Tellingly, they define Bitcoin as "a form of decentralized electronic fiat money" and emphasise that it "allows agents to make peer-to-peer payments and transactions without needing a trusted third party" (o.c., p. 1730; my emphasis).

However, Henry et al. (2019, p. 26) report that in the 2017 BTCOS the primary motivation for Bitcoin owners was to hold it 'as an investment', with 58 per cent choosing this response. Further, only about half of Bitcoin owners could be considered 'transactors'; that is, individuals who used Bitcoin to pay for goods and services and/or send money to other people at least once per year.[1]

Network size obviously matters for financial assets too — as it increases liquidity — but the shape and strength of any network effects may well differ. The same is true for the (geographical) scope of the relevant market (Van Hove, 2016). Individuals who see Bitcoin as an investment vehicle are unlikely to trade it only with 'peers' or 'friends' (see below).

To their credit, in a robustness check Balutel et al. do consider the role of financial investment in Bitcoin adoption. They do so by measuring the financial return that respondents expect from holding Bitcoin and adding this as an explanatory variable to their 2018 regression for Bitcoin beliefs (o.c., p. 1748, Table 5). However, this only relates to the first stage of the estimation approach; the second stage does not take into account how network size might matter in the adoption of Bitcoin for investment reasons.

[1] See also Balutel et al. (2022b, p. 2): "..., the stunning increase in the price of Bitcoin, which rose from $1,000 USD in late 2016 to a peak of almost $20,000 USD in late 2017, has led many to view Bitcoin as something more akin to a 'cryptoasset' than a cryptocurrency".

A second point is that Balutel et al. consider only direct network effects, whereby one user's adoption has a direct positive impact on the utility of other users (Van Hove, 1999, p. 139). . Indeed, Balutel et al. (2022a, p. 1736; my emphasis) "assume that ... the benefit of using Bitcoin is increasing in how many other individuals use Bitcoin". Such direct network effects matter for peer-to-peer payments, but, in addition, there can — at least in principle — also be indirect network effects. The latter imply that increased usage of Bitcoin to pay for goods or services will incite more merchants to accept it; increased acceptance by merchants, in turn, makes Bitcoin more useful for consumers (Halaburda et al., 2022, p, 996-997).. Here one could argue that Bitcoin acceptance by merchants was (and is) low (o.c., p. 999), but this does not justify the neglect of indirect network effects. The lack of such effects could be a reason for individuals not to adopt Bitcoin.

Third, one could question whether the scope of Balutel et al.'s network size variable is appropriate. As is explained in more detail in Section 4, network size is proxied by respondents' probability of having friends who own Bitcoin. The implicit assumption is thus that the utility which individuals derive from using Bitcoin comes solely, or at least primarily, from exchanging coins with friends. In other words, Balutel et al. assume that the relevant network is not only local but even social.

This is at odds with the intrinsic global nature of Bitcoin. Owners of Bitcoin will obviously not interact with all other users in the country, let alone around the world. But, for a payment instrument, acceptance beyond one's social network — and even beyond the circle of merchants and websites one patronises regularly — does matter. Consider the case of PayPal. Knowing that if one spots something interesting on a P2P site such as eBay or Etsy, there is a decent chance that the seller (with whom one typically will not have dealt previously) will accept PayPal, adds to the value of PayPal for consumers. If one takes this reasoning to its extreme, one could even question whether the network size variable should carry a subscript i — were it not for the role played by expectations.

Indeed, the theory on network externalities holds that economic agents base their decision to adopt a network good — a decision that Balutel et al. (2022a, p. 1731) term "costly" — not only on the current network size but also on the future (read: expected) network size (Van Hove, 1999, p. 144). Clearly, such expectations can differ substantially across potential users, as descriptive statistics for Balutel et al.'s 'Bitcoin beliefs' variable, ξ_{it}, would most probably show.

From this perspective, it could be argued — and this is my fourth remark — that ξ_{it} to some extent captures network effects (besides, or rather than, individual learning effects). The explanation is as follows. The beliefs variable is based on the following BTCOS question: "How likely do you think it is that the Bitcoin system will fail or survive in the next 15 years? Please use the sliding scale where 0 means that the system will certainly fail and 100 means the system will certainly survive" (Henry et al., 2019, p. 28). Respondents who believe that Bitcoin has a low (high) probability of survival likely have low (high) expectations of network size. In other words, it could be argued that the beliefs variable tracks expected network size. As explained, the latter drives adoption. Hence, ξ_{it} might also capture network effects, rather than (only) individual learning effects.

A final conceptual problem stems from a discrepancy between the theoretical and empirical parts of the article. The theoretical model is clearly about adoption: "In our model, there is a continuum of risk-neutral agents who, at every period, choose whether to make the costly adoption of Bitcoin" (Balutel et al., 2022a, p. 1731; my emphasis); see also o.c., p. 1736. The empirical part of the paper also predominantly uses the term 'adoption', but whereas it is claimed that the BTCOS yields information on "Bitcoin adoption decisions" (o.c., p. 1731)[2], the empirics are, in fact, about Bitcoin ownership. Indeed, Bitcoin 'adopters' are identified based on the question "Do you currently have or own any Bitcoin?" (o.c., p. 1741). Clearly, a respondent who reported owning Bitcoins in 2018 may have adopted already in 2016 (or earlier)[3].

This is not merely a linguistic squabble. In my example, the respondent based their decision to adopt on Bitcoin penetration in 2016 and on their 2016 expectations of future network size. Any later changes in network size should not affect the respondent's ownership of Bitcoin, provided that the (expected) network size does not fall below the minimum required by the respondent for Bitcoin to be sufficiently useful. As explained earlier, in Balutel et al.'s approach 'network size' should be read as respondents' probability of having friends who own Bitcoin, but that does not detract from the argument that their regressions do assume that ownership depends on current network size.

As for the extent of the problem, according to the BTCOS, ownership of Bitcoin increased from 4.3% to 5.2% between 2017 and 2018 (o.c., p. 1741). Hence, there clearly must have been new adoption. However, by the same token, these numbers do suggest that quite some of the 2018 respondents already owned Bitcoin in 2017, although it is impossible to tell precisely how many. For this, one would need panel data. Interestingly, "one of the key innovations" in the 2021 BTCOS was the addition of a question asking Bitcoin owners when they first obtained it (Balutel et al., 2024, p. 780).

4 Technical issues

The previous section has set out a number of conceptual problems with the empirical part of Balutel et al.'s research. However, the most important issue is of a more technical nature and relates to the construction of the network size variable A_{it}.

To construct this variable, Balutel et al. exploit the BTCOS question "Please tell us your main reason for owning Bitcoin". This is a multiple-choice question with 12 answer options, one of which is "My friends own Bitcoin" (Henry et al., 2019, p. 28). Balutel et al. build on the latter answer. In a first step, they regress "the probability of having friends owning Bitcoin" on the observed demographic characteristics of the Bitcoin owners (Balutel et al., 2022a, p. 1742). To avoid contemporaneous collinearity between individual Bitcoin adoption and A_{it}, in this step they use data from the BTCOS survey of the year $t-1$; i.e., they use 2016 and 2017 BTCOS data for, respectively, the 2017 and 2018 analysis. In the second step, Balutel et al. then use the estimated betas from the previous step to impute the network size variable for all respondents in year t — including non-owners. This imputation is needed because the question that is exploited is asked only to owners. For robustness, Balutel et al. also use an alternative, non-parametric method to construct

[2] The full sentence reads: "Key to this paper is BTCOS information on both Bitcoin adoption decisions and individual beliefs about Bitcoin" (o.c., p. 1731).

[3] At a number of points in their paper, the authors themselves use the term ownership. Cf. "The BTCOS was commissioned in late 2016 by the Bank of Canada to gather information on the awareness and ownership of Bitcoin among Canadians" (o.c., p. 1731; my emphasis) and "The core components of the survey are as follows: awareness of Bitcoin; ownership/past ownership of Bitcoin ..." (o.c., p. 1739; my emphasis).

A_{it} (Balutel et al., 2021, p. 33-34).

Unfortunately, the above approach contains a critical logical error: the question that is exploited does not, in fact, allow one to correctly measure "the probability of having friends owning Bitcoin" amongst Bitcoin owners (and, by extension, in the broader sample). This is because the question asks about the *main reason* for owning Bitcoin; it is not a yes/no question on whether owners have friends who also own Bitcoin. A respondent who answers, say, that they own Bitcoin because "It is an investment" may well have Bitcoin-owning friends. Yet they are assigned a value of 0 in the first step of the two-step procedure. In short, Balutel et al. mismeasure the Bitcoin network of many of the owners. This problem obviously percolates through all subsequent steps of the construction of A_{it}. In other words, the imputed values for A_{it} will also be incorrect.

In order to gauge the extent of the problem, I have conducted a simulation. A detailed explanation can be found in the Appendix. The simulation shows that, if one assumes that respondents have 10 friends, up to 32.6% of the owners in the 2017 BTCOS may have erroneously been assigned a 0 for network size in the first step of the procedure. This would imply an underestimation of up to 273%. How this affects the subsequent steps depends on precisely which owners have Bitcoin-owning friends — because their demographic profile matters (see above). The number of possible scenarios for the distribution of Bitcoin ownership among the friends of owners is such that it is impossible to reliably compute the eventual magnitude of the error in A_{it}. However, given the potential mismeasurement in the first step, it is probably substantial.

Note that the problem also affects the non-parametric alternative, as this uses the same BTCOS question as its starting point (o.c., p. 33). As an aside, while the authors concede that "there are marginal differences" (o.c., p. 34) between the values of A_{it} obtained by means of the parametric and the non-parametric approach, on closer scrutiny these differences are quite sizeable, in particular for the year 2018. The mean values of A_{it} are, respectively, 0.0168 (parametric approach) and 0.0424 (non-parametric approach) — a difference of a factor of 2.5. Similarly, albeit in the other direction, the maximum values are 0.375 and 0.198.[4]

A second remark is that A_{it} should not be read as a measure of network size — in contrast to the impression that is created by the authors' discussion of their results (Balutel et al., 2022a, p. 1750; my emphasis):

"..., a one percentage point increase in size of the network raises the likelihood of Bitcoin adoption by 0.45 to 0.32 percentage points in 2017 and 2018, respectively. ... Thus, we find that a large network size directly increases the probability of adopting Bitcoin, indicating that high Bitcoin adoption among peers is associated with a high propensity to adopt."

Indeed, a value of 0.0182 for A_{it} — the average for the year 2017 (Balutel et al., 2021, p. 34). — does not imply that, on average, 1.82% of respondents' friends own Bitcoin. Rather it should be read as a 1.82% probability that respondents have at least one friend who owns Bitcoin — without it being possible to determine just how many.

[4]The authors do not report correlations between the values of Ait obtained by means of the parametric and the non-parametric approach.

A third technical comment is that the first stage of the described approach hinges on a limited number of observations. In the 2017 BTCOS — which is used to estimate the 2018 values for A_{it}; see above — a mere 14 respondents out of 117 owners indicated owning Bitcoin because their friends owned some. The demographic characteristics of these 14 individuals eventually drive the imputed 'network sizes' for all 1,987 observations in the 2018 dataset. This implies that if the demographic characteristics of the 14 individuals do not reflect those of the Bitcoin owners with Bitcoin-owning friends, then the imputation in the second stage will be off target. Clearly, the lower the number of exemplars, the higher the probability of measurement errors. In this respect, note also that only one of the 14 exemplars was in the 35–54 age category; the 13 others were all 18 to 34 years of age (Balutel et al., 2022a, p. 1756).

A final remark is that the first step of the imputation procedure selects individuals who indicate that they bought Bitcoin because their friends own the cryptoasset; that is, individuals for whom (the size of) their network had an impact on adoption. In the second step, the profile of these respondents (in year $t-1$) is then used to assign values for A_{it} to all respondents (in year t) — implying that respondents with a similar profile will be assigned high values for 'network size'. This raises the question whether the significant result that is obtained for β_1 — the coefficient on network size in (1) — is not embedded in the approach.

5 *Conclusion*

This comment has shown that, albeit econometrically strong, Balutel et al.'s empirical analysis has several problems. In particular, their network size variable is critically compromised. I show that the question used to construct the variable does not allow one to accurately capture the probability of having Bitcoin-owning friends among Bitcoin owners (and, by extension, in the broader sample). In addition, the resulting variable does not measure the size of respondents' social Bitcoin network (that is, the number of Bitcoin-owning friends). I also argue that Balutel et al.'s network 'size' variable is at the wrong level, in that adoption beyond one's social network matters.

Given these problems, Balutel et al. cannot, based on their network size variable, reliably demonstrate that "network effects [. . .] were important drivers of Bitcoin adoption in 2017 and 2018 in Canada" (o.c., Abstract, p. 1729).

My suggestions for improvement are natural corollaries of the critiques just outlined. Arranged from better to best, one could ask Bitcoin owners whether they have friends who own Bitcoin (so that the problem that I document in my simulation is avoided); one could ask all respondents whether they have friends who own Bitcoin (so that there is no need for imputation); one could ask all respondents how many Bitcoin-owning friends they have, or what percentage of their friends own Bitcoin (so that A_{it} becomes a measure of social network size); and, finally, one could ask respondents about their perception of global Bitcoin penetration (so that A_{it} no longer measures just social network size).

More broadly, an important cautionary insight from Balutel et al.'s efforts is that network effects and learning effects are difficult to disentangle. To quote the authors "as the size of the network (or the number of adopters) increases, presumably the incentives to adopt Bitcoin increase [= network effects], as does the ability to learn about the unobserved technological quality of this innovation" (o.c., p. 1731; my emphasis).[5] Given that Balutel et al.'s network size variable has a crucial social dimension — in that it measures adoption by friends/peers — what they find for A_{it} may, if anything, well be social learning effects rather than network effects. Conversely, Balutel et al.'s beliefs variable, ξ_{it}, may well capture network effects rather than (just) individual learning effects.

[5] Cf. also: "To capture social learning effects, we assume that the arrival rate of news depends not only on the quality of the technology, but also on the level of adoption: the speed of learning rises as more people adopt Bitcoin" (o.c., p. 1732).

Bibliography

Balutel, D., Engert, W., Henry, C. S., Huynh, K. P., & Voia, M. (2024). "Explaining Bitcoin Ownership in Canada: Trends from 2016 to 2021." *Canadian Journal of Economics*, 57, 777–798. DOI: 10.1111/caje.12724.

Balutel, D., Henry, C. S., Vásquez, J., & Voia, M. (2021). "Bitcoin Adoption and Beliefs in Canada." *Staff Working Paper 2021-60*, Bank of Canada.

Balutel, D., Henry, C. S., Vásquez, J., & Voia, M. (2022a). "Bitcoin Adoption and Beliefs in Canada." *Canadian Journal of Economics*, 55, 1729–1761. DOI: 10.1111/caje.12620.

Balutel, D., Henry, C. S., Huynh, K. P., & Voia, M. (2022b). "Cash in the Pocket, Cash in the Cloud: Cash Holdings of Bitcoin Owners." *Staff Working Paper 2022-26*, Bank of Canada.

Gandal, N., & Halaburda, H. (2016). "Can We Predict the Winner in a Market With Network Effects? Competition in Cryptocurrency Market." *Games*, 7, 16. DOI: 10.3390/g7030016.

Halaburda, H., Haeringer, G., Gans, J. S., & Gandal, N. (2022). "The Microeconomics of Cryptocurrencies." *Journal of Economic Literature*, 60, 971–1013. DOI: 10.1257/jel.20201593.

Henry, C. S., Huynh, K. P., & Nicholls, G. (2017). "Bitcoin Awareness and Usage in Canada." *Working Paper No. 2017-56*, Bank of Canada.

Henry, C. S., Huynh, K. P., & Nicholls, G. (2019). "Bitcoin Awareness and Usage in Canada: An Update." *The Journal of Investing*, 28, 21–31. DOI: 10.3905/joi.2019.28.3.021.

Van Hove, L. (1999). "Electronic Money and the Network Externalities Theory: Lessons for Real Life." *Netnomics*, 1, 137–171. DOI: 10.1023/A:1019105906374.

Van Hove, L. (2016). "Testing Metcalfe's Law: Pitfalls and Possibilities." *Information Economics and Policy*, 37, 67–76. DOI: 10.1016/j.infoecopol.2016.09.001.

Appendix – Simulation

As mentioned in Section 4 of the main text, to gauge the potential error in the construction of the network size variable A_{it}, I conducted a simulation for the first step in the procedure, where Balutel et al. (try to) identify the BTCOS Bitcoin owners who have friends who also own Bitcoin. As explained, the authors exploit the multiple-choice BTCOS question "Please tell us your main reason for owning Bitcoin" and assume that respondents who answered "My friends own Bitcoin" have Bitcoin-owning friends. This is obviously correct, but the authors apparently overlook that respondents who gave other reasons may well have Bitcoin-owning friends too.

The goal of the simulation was to determine, under certain assumptions, how many of the Bitcoin owners may, in this way, have erroneously been assigned a 0 for network size. The assumptions are as follows. First, every Bitcoin owner has the same number of friends, f. This gives a total number of friends amounting to f multiplied by the number of Bitcoin owners ($f \times 117$ in 2017, $f \times 99$ in 2018). Second, overall Bitcoin penetration among the friends mirrors penetration among the BTCOS respondents and, given that the BTCOS sample is representative[6], also among the broader adult Canadian population. Specifically, I set the penetration rate at 4.46% in 2017 and 4.98% in 2018[7].

I then considered two extreme scenarios for the distribution of Bitcoin ownership amongst the friends: highly concentrated vs. highly dispersed. The first scenario creates BTCOS Bitcoin owners whose f friends all have Bitcoin; the second creates Bitcoin owners who each have only one Bitcoin-owning friend. Both simulations stop once all Bitcoin-owning friends have been assigned to an owner. The first scenario allows determination of the minimum number of BTCOS Bitcoin owners with Bitcoin-owning friends; the second gives the maximum. In both simulations, a key constraint is that the respondents who indicated that the main reason why they adopted Bitcoin was that their friends own Bitcoin definitely need to be assigned Bitcoin-owning friends — f or 1, depending on the scenario.

[6] Cf. Balutel et al. (2022a, p. 1740): "Sampling for the survey is conducted to meet quota targets based on age, gender and region".

[7] These percentages differ from the 4.3% and 5.2% mentioned by Balutel et al. in their Table 2 (2022a, p. 1741), because in Table 2 data have been weighted to reflect the Canadian population. I prefer to use the raw numbers, mentioned on p. 1740, for reasons of consistency with the numbers that I mention elsewhere in the text (and that I work with in other parts of my calculations).

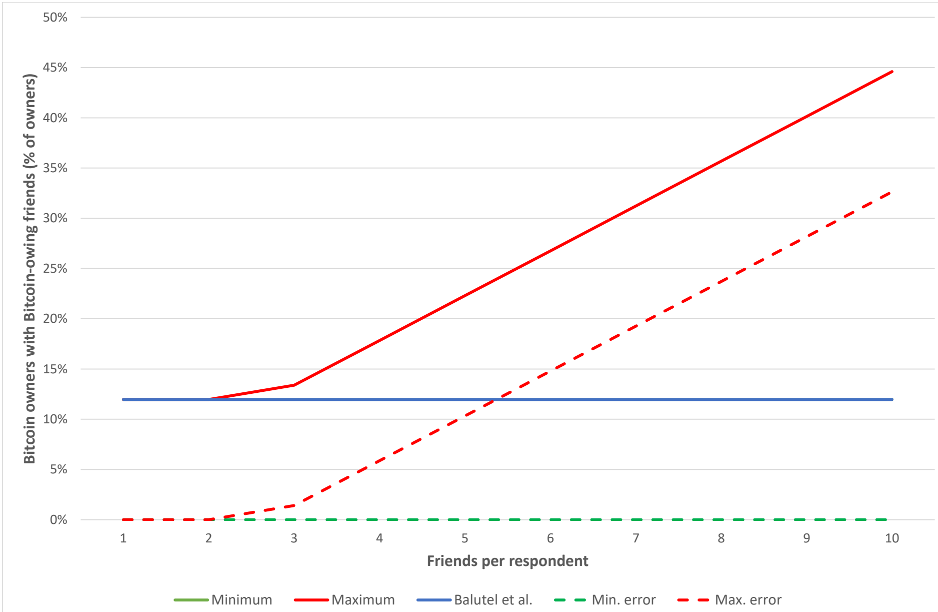

Figure A1: Owners with Bitcoin-owning friends, 2017 BTCOS, adoption rate = 4.46%.

A first remark relates to the observation that the full green line coincides with the blue line. This indicates that there is effectively a possibility that Balutel et al. have the number of Bitcoin owners with Bitcoin-owning friends correct, namely — hence the coincidence[8] — in the 'highly concentrated' scenario. However, one needs to realise that in this scenario literally all Bitcoin owners among the friends are assigned to the 14 Bitcoin owners in the 2017 BTCOS whose main reason for owning Bitcoin is that their friends own Bitcoin; none of the other 103 BTCOS Bitcoin owners in that year would have any friends who own Bitcoin. This does not seem realistic. Remember that it is, on purpose, an extreme scenario.

As the figure suggests, in the other scenario — the 'highly dispersed' scenario — Balutel et al.'s estimate is off the mark as soon as the BTCOS Bitcoin owners have 3 friends, and the error becomes very substantial for higher numbers. This raises the question of what value of f is realistic. Interestingly, detailed numbers can be found in Statistics Canada's General Social Survey on Social Identity (GSSSI); see Table A1. As can be seen, the survey makes a distinction between close friends and other friends. The first are defined as "people who are not your relatives, but who you feel at ease with, can talk to about what is on your mind, or call on for help?"[9] The BTCOS questionnaire, for its part, does not offer a definition of friends. I therefore consider both close and other friends.

[8]This coincidence is not a coincidence but stems from the set-up of the simulation. The intuition is as follows. There are f times as many friends as Bitcoin owners, but as long as the Bitcoin adoption rate (here: 4.46%) is lower than the relative importance, among the BTCOS Bitcoin owners, of respondents whose main reason for owning Bitcoin is that their friends own Bitcoin (here: 11.97%), the Bitcoin-owning friends can never outnumber the total number friends of the Bitcoin owners. Hence, in the 'highly concentrated' scenario no Bitcoin owning friends need to be assigned to non-owners. The upshot is that the number of BTCOS Bitcoin owners with Bitcoin-owning friends is limited to the 'known ones', so to speak.

[9]Source: https://www.statcan.gc.ca/en/statistical-programs/instrument/5024_Q1_V4, last visited on March 10, 2025.

Table A1: Social Connections with Friends: Canada, 2020

	Number	Percent
Number of close friends		
0	3,010	8.84
1	3,568	10.48
2	6,111	17.95
3	4,878	14.33
4	3,745	11.00
5	4,054	11.91
6	2,166	6.36
7	752	2.21
8	1,072	3.15
9	197	0.58
10 to 19	3,242	9.52
20 to 29	624	1.83
30 or more	256	0.75
Number of other friends		
0	2,883	8.54
1	933	2.76
2 to 19	18,785	55.63
20 to 49	6,800	20.14
50 to 79	2,544	7.53
80 or more	1,825	5.40

Notes: Numbers and percentages do not include 'don't know' and 'not stated' answers.
Source: Statistics Canada, General Social Survey on Social Identity, 2020.

To start with the first, in the 2020 GSSSI, the weighted average of the number of close friends is 4.9, and 62% of Canadians have three or more such friends — three being the point from which Balutel et al.'s estimate begins to diverge (cf. supra)[10]. With $f = 5$ (the rounded value of 4.9), 10.3% of the Bitcoin owners are erroneously assigned a 0 for network size in Figure A1 (for the year 2017). With BTCOS data for 2018, the percentage is 18.9%.

The weighted average of the number of 'other' friends reported in the GSSSI amounts to 22.0. Hence, if one uses a broader definition of friends than just close friends, even $f = 10$ is not unreal-

[10]Source: own calculations, based on the data in Table A1. For the calculation of the weighted average, I have, when needed, used the midpoints of the categories. For the upper category, I have used the lower bound.

istic. For $f = 10$, Balutel et al. erroneously assign a 0 for network size to, respectively, 32.6% and 43.8% of the Bitcoin owners in 2017 and 2018. Put differently, for 2017, at 44.60% the simulated share of Bitcoin owners with Bitcoin-owning friends is 3.7 times higher than Balutel et al.'s 11.97%. For 2018, the ratio is even 8.2 (49.8% vs. 6.06%).

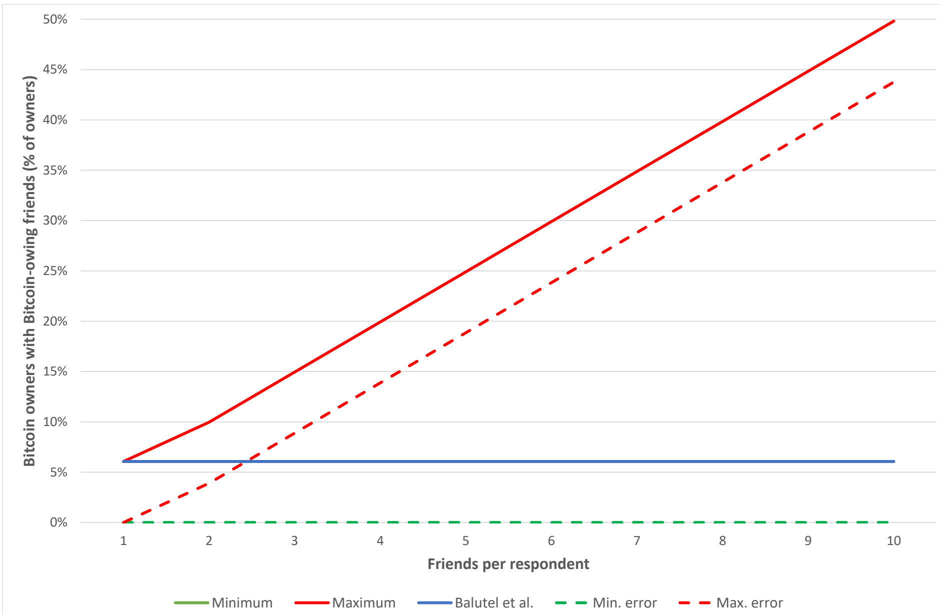

Figure A2: Owners with Bitcoin-owning friends, 2018 BTCOS, adoption rate = 4.98%.

How these potential errors affect the subsequent steps of Balutel et al.'s procedure to construct the network size variable A_{it} is difficult to determine. As explained in the main text, the next step consists in regressing the probability of having friends owning Bitcoin on the demographic characteristics of the Bitcoin owners. This implies that it matters precisely which owners have Bitcoin-owning friends. The number of possibilities for the distribution of Bitcoin ownership among the friends of owners is so high that it is impossible to reliably compute the eventual magnitude of the error in A_{it}. However, it is clear that the margins of error identified in Figures A1 and A2 do not bode well.