Consumption Inequality in Canada, 1997 to 2009*

Sam Norris† and Krishna Pendakur‡

September 11, 2014

Abstract

We assess the evolution of consumption inequality in Canada over the years 1997 to 2009. We correct the imputation of shelter consumption for homeowners to allow for unobserved differences in housing quality correlated with selection into rental tenure, and we account for measurement error in this imputation. Using the Surveys of Household Spending 1997-2009, we find that household-level consumption inequality measured by the Gini coefficient increased from 0.251 to 0.275 over 1997 to 2006, and then declined to 0.264 by 2009. The Gini coefficient for individual level inequality similarly followed a hump-shaped pattern: it increased from 0.199 in 1997 to 0.216 in 2006, and then fell to 0.207 in 2009. In contrast, the Gini coefficient for household-level income inequality followed a similar hump-shaped pattern, but the post-2006 decline was large enough to entirely wipe out pre-2006 increase. We also explore a possible correction for tail non-response bias in inequality measurement, and find that the increase in measured consumption inequality is robust to this correction.

I. Introduction

Measures of inequality that are based on household consumption data hold a number of advantages over the traditional choice of income. Most strikingly, consumption is chosen by the household with full knowledge of past income and some knowledge of future income, so it better reflects permanent (or,

*We would like to thank Arthur Lewbel and two anonymous referees for their thoughts and comments on this work. Pendakur acknowledges the financial support of the Social Sciences and Humanities Research Council of Canada through its Standard Research Grants program. Norris acknowledges the financial support of SSHRC through its Doctoral Fellowship Awards.

†Northwestern University
‡Simon Fraser University

lifetime) income. Furthermore, consumption is what ultimately generates well-being — consumption, not income, is typically the argument of utility functions in economic models. Finally, recent work suggests that household survey data may better measure consumption than income, particularly in the tails of the distribution (Brzozowski and Crossley, 2011).

Measuring consumption inequality using household-level microdata presents a number of technical problems. Measurement typically proceeds by adding up expenditures on different commodities at the household level. Thus, an important problem is how to account for owner-occupiers, who do not report shelter expenditures that equal their shelter consumption flows. About half of households live in owner-occupied housing, but their reported mortgage and upkeep costs vary widely, and do not necessarily equal the consumption flow value of the house (in rental-tenure households paying market rent, this is simply the rent paid). The importance of this correction naturally depends on the share of the population that lives in owner-occupied housing. In particular, it can help explain differences in inequality between provinces with different population share of groups with varying home ownership rates. In developed countries, the groups most likely to have significantly different ownership rates than the norm are the very young and the elderly (Fisher et al., 2007). Following Norris and Pendakur (2013), we impute rent using a new estimator that accounts for potentially unobserved characteristics correlated with selection into home ownership.

We implement these methods with the 1997-2009 Surveys of Household Spending to update the literature on consumption inequality in Canada. We find that household consumption inequality rose over the period 1997 to 2006, before falling somewhat between 2007 and 2009. The Gini coefficient for household-level real consumption, accounting for both rent imputation and for price differences over time and across provinces, increased from 0.251 to 0.275 over 1997 to 2006, and then declined to 0.264 by 2009. The overall increase of 0.013 Gini points is a medium-size increase compared to historical variation since the 1970s (see, e.g., Pendakur 2001).

Our results are driven primarily by changes in expenditure rather than in prices, and are robust to the choice of price deflator. We also find that despite the run-up in housing prices over the past 15 years, non-housing inequality increased at a faster rate than overall inequality. Finally, we show that income inequality followed a similar hump-shaped pattern. However, the increase in income inequality was shallower over the period 1997-2006, and the decrease faster afterwards, meaning that income inequality was essentially flat over the entire study period. This is particularly interesting given that income inequality in the United States increased over the entire study period (see, e.g., Meyer & Sullivan 2013).

In a concluding subsection, we examine the problem of survey non-response. The SHS has relatively low non-response for a developed-country consumption survey; however 'low' still means 30% in a typical year. Since non-response in expenditure surveys is likely concentrated in the upper and lower tails of the distribution, naive inequality measures will be biased downwards (towards showing less inequality). We explore a potential remedy to this problem.

The idea is simple. Empirically, consumption distributions 'look' roughly lognormal. Gibrat's Law states that lognormality is a long-run equilibrium shape for the distribution of consumption when income is the product of a series of multiplicative, uncorrelated shocks (Hall, 1979). Battistin et al. (2009) note that consumption is indeed more log-normal than income, and show that Gibrat's Law applies to permanent income rather than within-period income (Brzozowski et. al (2010) provide welcome evidence that this pattern holds true in Canada). This suggests that consumption, which depends crucially on household-level permanent income, should be be close to log-normal.

Given a known functional form for consumption, we show that a sufficient statistic for the Gini coefficient can be calculated from a truncated sample near the middle of the distribution, where non-response may be less prevalent. We find that our main conclusion that consumption inequality rose in the late 1990s and early 2000s is robust to this method of correcting for tail non-response. But, we also find evidence of non-normality even in the truncated distribution, suggesting that our correction strategy may be insufficiently flexible for the real data we observe.

II. Methodology

A. Estimation and Imputation of Household Consumption

Much of the following material is presented in Norris and Pendakur (2013, hereafter NP), where we examined a similar problem of rent imputation, but in the context of estimating consumption poverty rates. There we found that correcting housing flow imputations for selection into home ownership had a large effect on poverty rates and trend: in 1997, the first year of our sample, this change decreased the number of people observed in poverty by about 40%. By 2009, however, the difference between the poverty rates calculated using our method and the standard approach was statistically and meaningfully zero. This means that our imputation dramatically flattened the observed reduction in poverty rates during the 2000s as compared to a naive estimator. We also explored whether accounting for the reduced variation in imputed housing expenditure would change observed poverty rates, and found that it had the opposite effect of our selection correction, increasing 1997 poverty rates by about 20% but having almost no effect in 2009. The following discussion of the rent imputation is therefore somewhat abridged; more detail can be found in NP.

Let 'housing flow' refer to the money value of the consumption flow from shelter. Let x_i denote the total nominal consumption of households, $i = 1, ..., N$, and break total nominal consumption flow into two parts:

$$x_i = n_i + r_i \tag{1}$$

where r_i is the housing flow and n_i is the non-shelter consumption flow. We assume that n_i is observed and simply equal to the total expenditure on non-shelter goods and services. Let t_i be an indicator of rental tenure equal to 1 for households who are renters, defined as those who pay at least \$100 in rent for the year, do not pay reduced or subsidized rent, and do not pay any of their rent in kind. This strict definition of rental tenure is intended to capture only those households for whom the rent paid may be presumed to equal the consumption flow derived from their shelter.

We wish to estimate the housing flow of households for whom rent is not observed. Our problem is the familiar selection-correction model of Heckman (1979), modified to allow for heteroskedasticity in both equations. Let v_{1i} be a vector of variables that influence housing flow, including commodity prices, household demographics and dwelling characteristics. Let v_{2i} be a vector of variables that influence the tenure decision but do not influence rental expenditures (aka instruments). We assume that the real housing flow (housing flow divided by the housing price, p_{rent}) is linear in v_{1i}, that tenure is given by a probit model in v_{1i} and v_{2i}, and that the errors in these equations are heteroskedastic but jointly normally distributed. Our model is thus:

$$r_i/p_{rent} = v_{1i}\beta + u_{1i} \qquad \text{if } t_i = 1, \tag{2}$$

$$t_i = 1 \qquad \text{if } v_{1i}\Gamma_1 + v_{2i}\Gamma_2 + u_{2i} > 0, \tag{3}$$

$$t_i = 0 \qquad\qquad \text{otherwise,}$$

where

$$\begin{matrix} u_{1i} \\ u_{2i} \end{matrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_{1i}^2 & \rho\sigma_{1i}\sigma_{2i} \\ \rho_i\sigma_{1i}\sigma_{2i} & \sigma_{2i}^2 \end{pmatrix} \right), \tag{4}$$

$$\sigma_{1i} = v_{1i}\Delta_1, \tag{5}$$

$$\sigma_{2i} = exp(v_{1i}\Delta_2), \tag{6}$$

and

$$\rho_i = \rho_0 + v_{1i}C\rho_v, \tag{7}$$

where $v_{1i}C$ is a vector of the first two principal components of v_1 and ρ_v is vector of coefficients on each of these principal components. The last three equations allow the error terms to be heteroskedastic, which is important in a context where inequality is the object of interest. The standard deviations σ_{1i} and σ_{2i} are linear and exponential in v_{1i}, respectively, with the latter choice made for convenience.[1] The last equation could have been written as an unrestricted index in v_{1i}, but we use just the first two principal components of v_1 to reduce the dimension of the estimation problem.

This model could be estimated by full-information maximum-likelihood, but this turns out to be hard to implement due to the heteroskedastic components. We instead follow NP and estimate the model via a two-step limited-information maximum-likelihood procedure.

For households with $t_i = 1$, we have their total housing expenditures in the data. For households with $t_i = 0$, we must impute their real housing flow. Under our model, the real housing flow is normally distributed, so the predicted housing expenditure for a household with $t_i = 0$ is not a level, but rather a distribution of real housing flows. For each household with $t_i = 0$, the real housing flow is normally distributed as follows:

$$r_i/p_{rent} \sim N(\mu_i, s_i^2) \tag{8}$$

where

$$\mu_i = v_{1i}\beta - \rho_i\sigma_{1i}\lambda_i,$$

$$s_i^2 = \frac{\rho_i^2\sigma_{1i}^2}{\sigma_{2i}^2}\sigma_{1i}^2\left(1 + z_i\lambda_i - \lambda_i^2\right) + \sigma_{1i}^2\left(1 - \rho_i^2\right)$$

$$z_i = \frac{v_{1i}\Gamma_1 + v_{2i}\Gamma_2}{\sigma_{2i}},$$

[1] We tested alternate specifications for rent imputation, such as including v_2 in equation 6. This makes some difference in terms of imputed rent: with v_2 included, mean imputed rent is \$372 smaller. However, the spread is sufficiently small that it makes no practical difference in terms of Gini coefficients. Over the years of our sample period, for example, the largest difference between the Gini coefficient for nominal expenditure is 0.0004. Since the focus of this paper is inequality, we do not include results for different rent imputation specifications.

and

$$\lambda_i = \frac{\phi(-z_i)}{\Phi(-z_i)}. \tag{9}$$

The derivation of these formulae are given in NP. We pause only to note that the inverse Mills ratio, λ_i, responds to the heteroskedasticity in the probit, σ_{2i}. Further, because both λ_i and σ_{1i} are strictly positive and the estimated value of ρ_i is negative for all observations, the selection correction term $\rho_i \sigma_{1i} \lambda_i$ acts to push the mean of the real rental flow upwards for non-rental households.

Following NP, we compute *level-based* and *probability-based* Gini coefficients. For our level statistics, we assume that each household can be assigned an exact real housing flow. For all households, this is equal to μ_i, the expected value of the real housing flow. We estimate μ_i using estimated values of the parameters $\beta, \Gamma_1, \Gamma_2, \Delta_1, \Delta_2, \rho_0, \rho_v$. We note that although we know the exact value of expenditure for rental-tenure households, this observed rent includes the error term u_{1i}. Thus, to maintain comparability of rents for renters and imputed rents for non-renters, when we estimate our level-based Gini coefficients (which average out error terms), we impute for both non-renters and renters.

In contrast, for our probability-based Gini coefficients, each non-rental household is assigned a real housing flow which is normally distributed with mean μ_i and variance s_i^2. This real housing flow has variance, which accounts for the error term u_{1i}. For renters, the observed rent includes the error term u_{1i}, and so is comparable to the imputed consumption distribution of non-renters.

Since some of our household expenditure data are observed (that for renters) and some is imputed (that for non-renters), we have some observations where equivalent consumption is degenerate (renters) and others where it is a household-specific distribution (non-renters). To facilitate inequality measurement, we proceed by replacing each non-renter observation with 100 observations each with 1/100 the weight of the original observation. Then, we assign each of these a simulated real rental flow drawn from their household-specific normal distribution. Finally, inequality measures are calculated from the simulated distribution of real consumption.

III. Data

We use expenditure data drawn from the public use files of the annual Surveys of Household Spending (SHS) 1997-2009, and price data from Statistics Canada. We use expenditure and price data for all years, and for all provinces except Prince Edward Island (dropped due to data masking). Our household consumption expenditure variable is the sum of expenditure in 10 categories: food purchased at home; food purchased in restaurants; housing; fuel for principal accommodation; electricity; clothing and footwear; health and personal care; recreation, education and reading; alcohol and tobacco; and transportation (excluding car and RV purchase). The time span for expenditure recall is one year. In our sample, these 10 categories account for 84% of total current consumption.[2]

The SHS contains a rich set of demographic data. Included in the vector of control variables, v_{1i}, are the following: the age of the household reference person less 42 and its square; a dummy indicating that the household's reference person is a female; year of the survey minus 2002 and its square; Environment Canada's heating and cooling degree-days for each year/province less the overall average of these quantities; a car non-ownership dummy; household type dummies for couple only, couple with children, couple with children and others, single parent, other with relatives only, and other; dummies for households of size 2, 3, 4, 5, 6 and larger; a indicator that the household receives more than 10% of its income as government transfer payments; a dummy for living in a small urban area (less than 100,000 residents), and a dummy for living in a rural area. The car non-ownership dummy assigns 1 to households who spend less than $50 on gasoline and 0 to all others.

Note that not all of these variables are strictly exogeneous. In particular, car ownership is chosen as an expenditure category. Our results should therefore be understood as conditional on the distribution of car ownership. Since car ownership changes very little over our sample, we consider this conditionality worthwhile in exchange for more accurate estimates of housing consumption flow, which feeds directly into the Gini coefficient, our main object of interest.

We include the logged prices of our 10 consumption goods (where the shelter price is that of rental shelter), varying by province and year, in v_{1i}. We also include a set of dwelling characteristics in v_{1i}: the number of bedrooms and bathrooms, the repair condition, date of construction and detached dwelling status.

For our purposes, a household is considered a renter if and only if they

[2]Left-out categories are household operation, car and RV purchases, charitable contributions, and the SHS's other expenses (financial services, union dues, gambling, taxes, retirement fund savings, monetary gifts and donations). These were excluded either because we could not find a satisfactory source of price data, or we considered the category too durable.

report: that they rent their accommodation; they spent more than $100 on rent in the year; they do not pay reduced rent; and they do not pay any of their rent in-kind. To ensure that reported rents reflect market rents, we exclude subsidized renters and others whose reported rent is not informative. We note that 22 per cent of rental-tenure households report either reduced/subsidized rent or payment via in-kind.

The instruments for the rent imputation, v_{2i}, are: the log of the price of owned accommodation, the square and cube of the difference between the logged rent and owned accommodation prices, logged mean national mortgage rates for the previous three years, provincial unemployment rates, and dummies for married and single (excluded is separated/divorced). The idea here is that these variables affect the rent versus own decision, but do not directly affect rents or housing flows.[3]

Table 1 gives descriptive statistics for rent, imputed rent and instruments v_2, by rental status. Unsurprisingly, owners are richer and more likely to be married. The other important thing to note is that the standard deviation of the housing flow is approximately twice as large for both renters and owners in the probability measure as it is in the level measure. This is because the latter imputation has a variance of zero for each household. Summary statistics for v_1 can be found in Online Appendix Tables 1-3.

In our data, we observe commodity prices for each province and year (comparable across provinces and year, see NP). We consider variations on our model that account for with price variation in one of three ways. First, we ignore price variation and estimate Gini coefficients for nominal distributions. Since the Gini coefficient is homogeneous of degree zero in its arguments, if prices affect all households equally, then nominal inequality equals real inequality. Second, we divide each household's nominal consumption by a household-specific Stone price index. This price index is equal to the weighted geometric mean of the prices of each good, where weights are that household's budget share for each good. Since the weights are specific to each household, this strategy allows differential effects of price changes across households; in particular, it allows price changes to affect poor households differently from rich households. Third, we estimate the EASI demand system of Lewbel and Pendakur (2009) to more fully account for possible substitution effects across goods in the price index. This model-based price index is a simple generalisation of the Stone index. In our main results, we consider all three ways of dealing with price variation.

[3]The F-statistic on excluding these instruments from the tenure probit regression is 30, with a p-value of 0.0001. Thus, we most likely do not face a weak instrument problem.

IV. Results

Table 2 and Figure 1 shows Gini coefficients for income, for non-housing expenditure, and for total consumption corrected for differential housing quality between renters and owners. We provide estimates using a number of different price indices, and using our level and probability methodology. We focus on household-level results, and do not adjust household expenditure with an equivalence scale. This has two main advantages. First, this allows us to sidestep the choice of equivalence scale, although column 8 shows that this choice does little to the trend results. Second, our decision to focus on inequality at the household level also allows us to consider a particular strategy for calculating Gini coefficients when there is survey non-response, which we address more fully in the final section. That said, we include estimates of individual-level with inequality, which we show follow the same pattern as household-level inequality.

In each column, asymptotic standard errors are obtained from 50 bootstrap repetitions. We note that these standard errors are slightly smaller than they should be because they do not account for the first-stage selection process of the SHS. This is unavoidable, however, given that the SHS does not release the sampling information in the public-use files.

Starting on the left hand side of Table 2, income inequality is significantly higher than consumption inequality. This is true for two reasons, neither of which are unique to this time period or our methodologies: income has a higher variance than expenditure due to overtime, transitory job loss, and profits for business owners; and households tend to smooth consumption over their lifetimes, so seniors will have high consumption relative to income, and working-aged households the opposite. The trend, however, may be meaningful. Income inequality increases from 1999 to 2006, then falls in 2007, 2008, and 2009. However, it rises more slowly than the consumption measures over the first part of the study period, and falls at a faster rate after 2006. Over the entire period, household income inequality rises 0.001 Gini points, a statistically negligible increase.

We next compute inequality for non-housing expenditure. This specification is as close to the data as possible, and therefore provides a useful baseline. Non-housing consumption inequality is significantly higher than overall consumption inequality, which is unsurprising given that housing has high fixed costs and absorbs a large share of expenditure. We calculate inequality using both nominal consumption and and consumption deflated by a household-specific Stone price index.

We find that Stone-deflated non-shelter consumption inequality grows faster than nominal non-shelter inequality over the study period; the Stone-deflated Gini grows by 0.021 Gini points as opposed to 0.013 for the nominal measure. This means that variation in household-level Stone indices contributed increasingly to inequality over time. In effect, rich people's prices rose slower than did poor people's prices.

Moving to our preferred, imputed rent specifications, there is an overall increase of a little over 0.010 Gini points over the period. In all four specifications, inequality rises about 2 percentage point from 1999 to 2006, then declines about 1 percentage point until 2009. As with non-housing inequality, the rise in inequality is larger for Stone-deflated than nominal consumption, although the difference is smaller when housing is accounted for. The sixth column uses the EASI price index, which modifies the Stone index to more flexibly capture substitution effects (Lewbel and Pendakur, 2009). But this added flexibility in household-specific price indices adds approximately nothing to the inequality measurement; in most years the difference between the Stone and EASI columns is in the fifth decimal place. The point here is that the Stone index captures the major portion of how price changes hurt poor households more than rich households. Adding complexity to the price index does not change this picture.

The probability measure shows slightly higher inequality in all but one year. This is consistent with our expectations given that the imputing a distribution adds variance to the imputed rents of non-rental households. But, this increase in measured inequality is small. The differences are swamped by the over-time variation—in most years the difference is statatistically insignificant and no more than 0.003 Gini points. This is in contrast to the relatively large differences between level and probability *poverty* measures documented in NP. The reason is clear: for a poverty line in a steep part of the consumption distribution, the number of households in the neighborhood to the left of poverty line is much smaller than the number of households to the right, so allowing for a probability of poverty can significantly increase observed poverty. For inequality, however, no such story holds, and the effect of using the probability measure is to slightly flatten the consumption distribution, marginally increasing inequality. We therefore concentrate our attention on the fifth column (imputing a level of rent for every household and using the Stone index) as the most practical, easy-to-implement inequality measurement strategy. There, we see an increase of 2.4 percentage points in the Gini coefficient over 1997 to 2006, and a decrease of 1.1 percentage points over 2007 to 2009.

The last column of Table 2 contains Gini coefficients for individual-level consumption. We use the venerable square root equivalence scale, and weight observations by the SHS-provided weight multiplied by household size. Unsurprisingly, the level of inequality is lower — households may be large because the wage-earners can support more children, or simply because there are more wage-earners. The pattern is much the same as in our other specifications, although inequality peaks one year earlier, in 2005. Over the entire period, inequality increases by 0.008 Gini points, which is significant at the 5% level.

Table 3 presents household-level nominal consumption inequality estimates at the province level. There are several striking patterns. First, inequality in Newfoundland increases drastically between 1997 and 2000, rising from 0.202 to 0.258, after which it remains roughly constant for the remainder of the study period. Second, most provinces follow the same hump-shaped pattern observed at the national level, with declines in inequality starting in about 2005. The main exceptions to this trend are Ontario and Nova Scotia, where inequality briefly dipped in 2006 and 2007 before climbing again in the final two years of our sample (by 0.010 and 0.019, respectively). Conversely, the largest post-2006 declines were found in British Columbia, Saskatchewan and New Brunswick.

The patterns of inequality in Canada over our study period differed from the American experience in some important ways. First, the path of consumption inequality was surprisingly similar, with a peak in the middle of the 2000s, followed by a relatively rapid decline. Meyer and Sullivan (2013) report that the ratio of expenditure at the 90^{th} and 10^{th} percentiles peaked in 2005 at 4.5, then fell to the 2000 level of 4.1 by the end of 2011. Note that although the authors use a different inequality measure, because different inequality measures tend to folow the same trends it almost certainly makes no difference: Attanasio et. al (2012) track a third measure, the coefficient of variation, which itself follows the trend documented by Meyer and Sullivan.

Despite the similar trend in consumption inequality, income inequality in the United States and Canada follows dramatically different paths. In the US, income income inequality dipped slightly after 2005, but then continued growing from 2008 to 2011, with the 90/10 ratio increasing from about 5.3 to 6.3 over the 11 years since 2000. We believe that there are two possible explanations for why Canadian income inequality continued on a downward trend after 2006. First, Canadian unemployment increased much less than American unemployment during the Great Recession: from 6% to 8%, compared to an increase from 4.5% to 10%. In particular, of the four provinces with the lowest unemployment rates during the last half of the decade (British Columbia, Alberta, Saskatchewan and Manitoba), all except BC saw large decreases in income inequality in the years since 2006 (and as we note earlier, there was a significant decline in consumption inequality in BC over the time period). The opposite pattern was observed Nova Scotia, Ontario and Quebec, all of which have had relatively high unemployment rates since 2006.

Second, and in contrast to the first half of the decade, after 2005 wages increased more at the bottom of the wage distribution than at the top (Fortin et al., 2012). In particular, wages for women at the bottom of the distribution grew. Since women tend to make less money than men, wage growth at the bottom for women decreases inequality even more than wage growth for low-income men. Fortin et. al conjecture that much of this is due to minimum wage increases since the midpoint of the decade, which took place in all Canadian provinces except British Columbia.[4]

V. A Strategy to Correct for Tail Non-Response

Response rates for the SHS and other comparable national surveys are low; overall, the response rate from the SHS's sample frame was about 70% throughout the period. Certain populations are known to be less likely to respond: aboriginals; poor households; young people, particularly men; and the wealthy, among other groups. These groups are disproportionately found in the tails of the consumption distribution. While it is not immediately obvious in which direction their underrepresentation will bias mean income and consumption measurements, it will likely reduce the observed variance of the distribution, reducing observed inequality.

Statistics Canada accounts for this problem in several ways. High-income neighborhoods are oversampled in the first stage to account for lower second-stage uptake. There is also stratification by a large list of demographic variables. After the survey has been undertaken, sample weights are given by the inverse probablity of selection, and are adjusted for non-response along the same demographic characteristics. Finally, interview weights are adjusted so that the sample statistics for a number of household characteristics are the same as for the population at large. These include census statistics like number of persons in each age group and number of households for a given size, as well as income data from the Canada Revenue Agency (CRA). In each province, population counts are created for income groups defined by the percentiles 0-25, 25-50, 50-65, 65-75, 75-95, and 95-100. Then, for each of a number of calibration groups formed by demographic characteristics, the weights are simultaneously adjusted so that the weighted number of

[4]BC increased its minimum wage in 2011.

households matches the CRA population in each income group. Under this method, the observed and weighted consumption distribution will deviate from the true population distribution if either consumption is insufficiently correlated with income, or if there is sufficient consumption inequality within the income groups.

A potential solution for this problem lies in the observation that non-response isn't the problem *per se*, but rather disproportionate non-response in a particular part of the consumption distribution. Intuition suggests that this is likely to be the top and bottom, which implies that some portion of the distribution near the middle may be unaffected by disproportionate non-response. If we know the parametric distribution function $f(y;\theta)$ for consumption, we can estimate it near the middle and calculate the Gini directly from the parameters of the distribution. This is the strategy that we now pursue.

Implementation of this method requires prior knowledge of the functional form of the true distribution, and it is helpful if that true distribution has an analytical form for its truncated analogue. For consumption distributions, there exist two bodies of theory that suggest that the log-normal distribution should provide a good approximation to the true distribution and thus a good choice for implementing our methodology. First, an old literature starting with McAlister (1879) and Pareto (1895) (see Dagum 1999 for an interesting review of this literature) suggests that the log-normal distribution well approximates actual observed distributions. Second, and more recently, work stemming from Zipf's Law and Gibrat's Law (see especially Batistin et al. 2009) suggests a theoretical rationale for the proposition that permanent income and consumption distributions, respectively, should be lognormal. Gibrat's argument is essentially that if income shocks are cumulative multiplicative shocks, then log-normality arises a consequence of the law of large numbers. Battistin et al. suggest that in fact Gibrat's law should more reasonably apply to marginal utilities, and therefore to the distribution of both lifetime wealth and consumption. They indeed find that although income is roughly log-normal, consumption distributions are much closer to exactly log-normal than are income distributions. Thus, there exists some evidence that the true distribution of consumption—the distribution uncorrupted by tail nonresponse—should be close to log-normal.

In this section, we assume that the true distribution of household consumption y_i is log-normal, and that a range near the median of the true distribution is 'safe' in the sense that the rate of non-response is constant within that range. We then use data within that range to compute the parameters of the true distribution.

The log-normal distribution has the feature that the mode, median and mean of the logs are equal to each other. Since we have unknown tail nonresponse, we cannot estimate the median or mean straightforwardly from the observed data. However, we can identify the mode of the logged distribution, $\ln y_{mode}$, and use this to identify the median and mean of the true distribution. We use a nonparametric density estimator to recover an estimator $\ln \hat{y}_{mode}$ of the mode (, median, and mean) of $\ln y_i$. This is done simply by nonparametrically estimating $f(y;\theta)$ and choosing the value with the largest density.

Given an estimate, $\ln \hat{y}_{mode}$, we choose a neighbourhood around $\ln \hat{y}_{mode}$, and estimate the standard deviation of $\ln y_i$, $\widehat{\sigma}_y$, using the analytical form of the truncated log-normal. This is easily implemented, e.g., in Stata via `truncreg`. The standard error of $\widehat{\sigma}_y$ is recovered via the bootstrap, with 50 repetitions. Once we have $\widehat{\sigma}_y$ in hand, we can calculate the Gini coefficient for a log-normal variate as

$$Gini = 2\Phi\left(\frac{\widehat{\sigma}_y}{\sqrt{2}}\right) - 1 \tag{10}$$

where Φ is the cumulative distribution function of the standard normal. The appropriate standard errors are calculated by the delta method.

Our theory does not predict how small our neighbourhood around the median should be. Since we are operating under the assumption of log-normality, it seems sensible to choose a neighbourhood that does not fail a test for truncated log-normality. To operationalize this intuition, for each year we estimate neighbourhoods covering the mode that have the largest number of observations and pass the Kolmogorov-Smirnov test for the truncated normal distribution at the 1 and 5 per cent levels. To facilitate comparability between years, we use the median left- and right-endpoints, taken over the years, for each year. The result of this procedure is the interval $[\hat{y}_{mode} - 1.375, \hat{y}_{mode} + 0.75]$ and $[\hat{y}_{mode} - 1.05, \hat{y}_{mode} + 0.975]$ for the 5 and 1 per cent levels, respectively.

Table 4 presents our main results. The leftmost column reproduces the fifth column from Table 2. The rightmost columns give our estimated Gini coefficients using log-normality and the truncated distributions. Here, we see that our main finding that consumption inequality rose over 1997 to 2006 and then declined subsequently is robust to this correction for tail nonresponse. However, our correction shows a much larger increase in inequality than the uncorrected results. Considering the results for the larger neighbourhood (1% level), we see that the Gini coefficient rose by 0.029 points over 1997 to 2006. This is a relatively large increase, comparable to the increase in consumption inequality in Canada over 1978-1992 (Pendakur, 1998). Then, after 2006, we see a medium-sized decrease of 0.011 in the Gini coefficient.

Although it is comforting that the naive results from Table 2 are in-line with the non-response corrected results in Table 3, it would be useful to assess whether or not the log-normality necessary for this approach to work is actually obtained in this empirical example. Our framework gives three predictions: first, that the inequality from the model will be larger than the empirical measure; second, that the standard deviation calculated from the truncated normal will be larger than that from the observed data; and third, that the pdf of the observed data will not lie above that of the predicted lognormal distribution (since we are allowing for the possibility of missing data and not extra data). Table 3 corroborates the first prediction, but not second or third. In particular, comparing the first and seventh columns, the Gini coefficient is higher in some years for the model-based approach. However, the observed SD (column 2) is higher than the 5% SD in about half of years (column 4), and higher than the 1% SD (column 6) for all years.

The failure of our second prediction suggests that the underlying distribution may not be sufficiently lognormal to justify the use of this method. This is confirmed by Figure 3. Figure 3 shows the observed consumption distribution for 2003, as well as the distribution from the 5% level prediction scaled up by a factor of 1/0.723, the reciprocal of the non-response rate in 2003. For our estimated distribution to be the true distribution, we would expect that the observed mass falls strictly under the scaled distribution. However, inspection of the Figure shows that this does not hold. Further, the estimated pdf at $\ln y = 8.8$ and $\ln y = 8.2$ is statistically significantly higher than the normal distribution (taken as a fixed threshold).

Dealing with non-response is increasingly important in the Canadian context. The loss of the Census long-form after 2006 means that we will have decreasingly useful reweighting strategies to correct for non-response. Thus, it is imperative that assessments of inequality make some attempt to deal with this problem. Appealing to log-normality of the consumption distribution is one possibility, and that this approach has testable restrictions is clearly a desirable feature. But, in this case, the data do not meet those restrictions, so we are left wanting an even better strategy.

VI. Conclusions

We show that household income inequality has been essentially flat over the period 1997-2009, while consumption inequality has grown moderately, with the Gini coefficient increasing from 0.251 to 0.264 in our preferred specification. Both have declined from 2006 peaks, which we speculate may be driven by two factors. First, due to minimum wage increases in most provinces and a relatively strong labour market, wages increased faster in the bottom of the distribution than at the top. Second, historically poor stock market performance depressed incomes for high-asset households, which tend to have higher permanent income than non-asset holders. The patterns that we uncover are somewhat different from those in the United States, where over the same period consumption inequality also followed a hump-shaped pattern, but there was no post-2006 decline in income inequality. Most provinces followed the same pattern in both income and consumption inequality; the largest declines in consumption inequality after 2006 were in British Columbia, Saskatchewan and Nova Scotia, while the declines in income inequality were greatest in Alberta, Saskatchewan and Manitoba. We also uncover two notable provincial exceptions to the hump-shaped pattern: Newfoundland, where inequality rose very sharply from 1997 to 2000, then levelled off for the rest of the period; and Ontario, which did not see significant declines in either type of inequality after 2006 (in fact, the highest measure was observed in 2008).

We find that individual-level consumption inequality follows a similar pattern as the household-level measure, increasing from 0.199 to 0.216 over 1999-2006, then falling back to 0.207 by 2009.

Despite the surprising decline in consumption inequality over the latter half study period, we find that prices (and particularly non-housing prices) for poor households grew at a faster rate than prices for rich households. Over the medium term, this trend could potentially reverse the drop in consumption inequality observed since 2006, especially if wage growth slows at the bottom of the distribution.

We investigate a potential method for addressing survey non-response. For lognormal consumption, we outline a simple procedure that can address survey non-response in the tails of the distribution. We implement our procedure and show that our main estimates are robust to this kind of non-response. However, we also show that the Canadian data is likely not sufficiently lognormal for our procedure to be warranted.

VII. Figures

Figure 1: Gini coefficient over time, all Canada

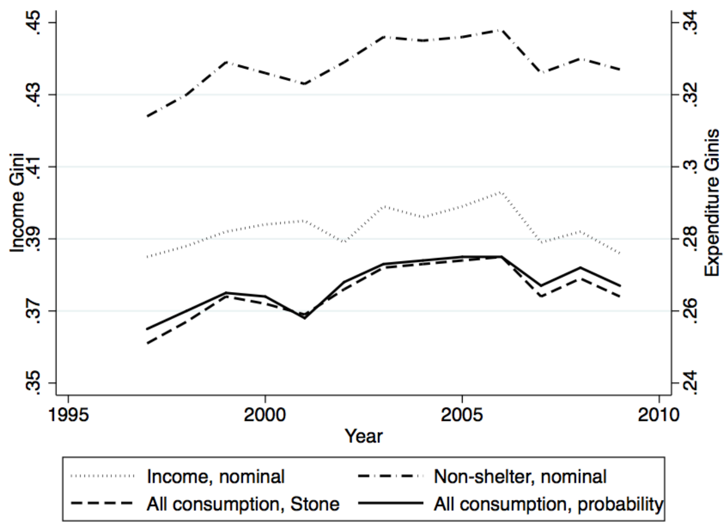

Empirical Gini coefficients for Canada, 1997-2009. Where housing expenditure is included in consumption, we impute it from observed rents, correcting for selection into homeownership. Note different axis labels for income and consumption.

Figure 2: Gini coefficient over time, for different truncation rules

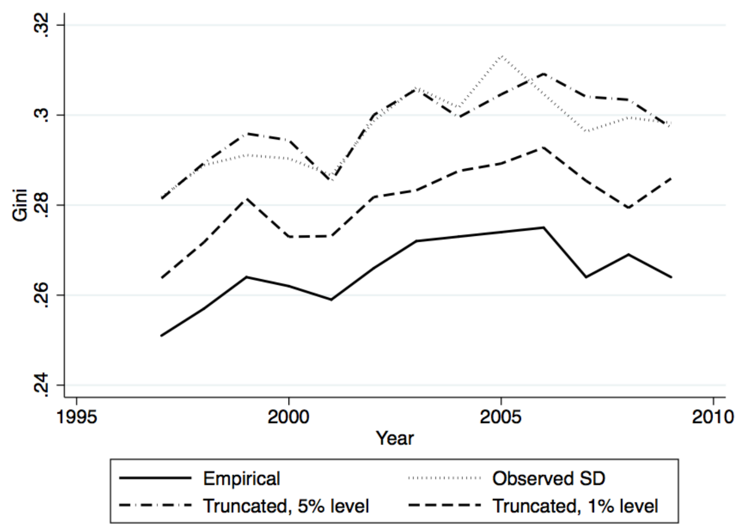

Model-based Gini coefficients for Canada, 1997-2009. Empirical consumption is deflated with a Stone price index and includes imputed rents, corrected for selection into homeownership. The observed SD Gini uses the Gini calculated from the standard deviation of Stone-deflated log expenditure under the (false) assumption that the observed data is lognormal. The final two entries calculate the Gini using the standard deviation of the truncated normal distribution of Stone-deflated log expenditure, where the truncation is chosen to maximize the number of observations while passing a K-S test at the indicated level.

Figure 3: Observed and predicted mass, 2003 at 5% level

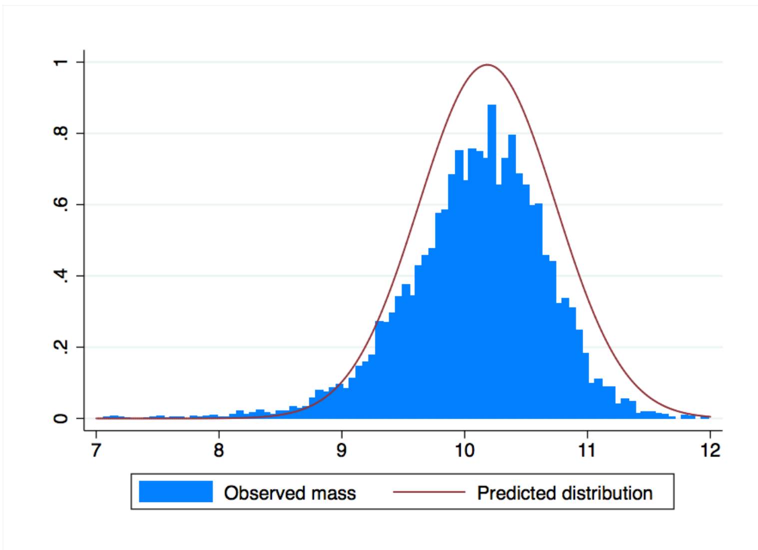

Observed distribution of Stone-deflated log consumption in 2003, with normal distribution. The normal distribution has parameters estimated from the 5% level prediction, and is scaled by 1/0.723, the non-response rate for that year. Were the true distribution lognormal, one would expect the entire mass to be under the normal distribution pdf.

VIII. Tables

Table 1: Summary statistics

	Owner	Renter
Panel A: Rent and expenditure summary statistics		
Expenditure	26,879	25,226
	(15,685)	(13,618)
Expenditure, non-housing	26,458	18,115
	(15,883)	(12,072)
Rent, real	5,577	8,237
	(3,400)	(3,730)
Imputed rent, real (level)	11,468	8,226
	(2,694)	(1,748)
Imputed rent, real (probability)	11,904	8,237
	(5,252)	(3,730)
Panel B: Instruments for rent imputation		
Married (=1)	0.689	0.394
	(0.463)	(0.489)
Single (=1)	0.0981	0.312
	(0.297)	(0.463)
Owned housing price	-0.169	-0.146
	(0.228)	(0.213)
Rented - owned housing price, squared	0.00484	0.00494
	(0.00718)	(0.00682)
Rented - owned housing price, cubed	0.000469	0.000441
	(0.00129)	(0.00120)
Unemployment rate	7.612	7.754
	(2.231)	(1.967)
Mortgage rates	1.907	1.904
	(0.115)	(0.115)
Observations	124,112	45,766

Summary statistics from rent imputation. Panel A contains summary statistics or expenditure and rent, Panel B the instrument vector v_2 for the first stage of the rent imputation. Summary statistics for v_1, the vector of second-stage RHS variables, can be found in Tables 1-3 of the Online Appendix.

Table 2: Household-level Gini coefficients

	Income	Nonshelter consumption		All consumption, imputed rent				
	Nominal	Nominal	Stone prices	Nominal	Stone prices	EASI prices	Probability	Indiv. level
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1997	0.385	0.314	0.312	0.260	0.251	0.251	0.255	0.199
	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)
1998	0.388	0.32	0.319	0.266	0.257	0.257	0.26	0.203
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)
1999	0.388	0.329	0.327	0.273	0.264	0.264	0.265	0.204
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)	(0.002)
2000	0.394	0.326	0.324	0.274	0.262	0.262	0.264	0.203
	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)
2001	0.395	0.323	0.324	0.273	0.259	0.259	0.258	0.206
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)
2002	0.389	0.329	0.329	0.278	0.266	0.266	0.268	0.207
	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
2003	0.399	0.336	0.336	0.283	0.272	0.272	0.273	0.215
	(0.004)	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.001)	(0.002)
2004	0.396	0.335	0.336	0.284	0.273	0.273	0.274	0.213
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
2005	0.399	0.336	0.338	0.284	0.274	0.274	0.275	0.216
	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.002)	(0.001)	(0.002)
2006	0.403	0.338	0.342	0.285	0.275	0.275	0.275	0.214
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.001)	(0.002)
2007	0.389	0.326	0.329	0.274	0.264	0.264	0.267	0.209
	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)
2008	0.392	0.33	0.335	0.278	0.269	0.269	0.272	0.210
	(0.004)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)
2009	0.386	0.327	0.333	0.271	0.264	0.264	0.267	0.207
	(0.004)	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.002)	(0.003)

Household-level Gini coefficients under different estimation methods. Column 1 calculates the income Gini using nominal prices. Column 2 displays the Gini for nonshelter nominal consumption; Column 3 deflates nominal consumption by the geometric mean of prices, weighted by consumption shares at the household level. For columns 4-8, rent is imputed, accounting for selection into homeownership with a heteroskedasticity-consistent Heckman two-step estimator. Column 6 deflates consumption with the EASI price index, which extends the Stone index to allow for second order price effects. Column 7 is calculated from a simulated consumption distribution, where imputed rent for each household is assumed to follow a normal distribution with the imputed mean and calculated standard deviation. Column 8 measures individual-level inequality: we use a square root equivalence scale and use the SHS-provided weights multiplied by family size.

Table 3: Consumption inequality by province

Year	AB	BC	Man	NB	Nfl	NS	ON	Qc	Sask
1997	0.256	0.258	0.265	0.249	0.202	0.256	0.249	0.263	0.245
1998	0.250	0.258	0.255	0.246	0.219	0.251	0.26	0.268	0.260
1999	0.255	0.271	0.239	0.253	0.217	0.248	0.262	0.282	0.274
2000	0.251	0.261	0.270	0.249	0.258	0.253	0.268	0.279	0.266
2001	0.265	0.275	0.282	0.247	0.250	0.254	0.268	0.275	0.257
2002	0.256	0.269	0.279	0.259	0.254	0.254	0.280	0.284	0.259
2003	0.255	0.276	0.277	0.257	0.258	0.261	0.281	0.265	0.265
2004	0.264	0.278	0.284	0.259	0.262	0.269	0.278	0.272	0.267
2005	0.266	0.277	0.278	0.261	0.26	0.275	0.282	0.276	0.275
2006	0.254	0.269	0.278	0.243	0.254	0.263	0.265	0.266	0.264
2007	0.261	0.273	0.271	0.253	0.264	0.245	0.267	0.267	0.264
2008	0.257	0.262	0.257	0.255	0.240	0.262	0.287	0.271	0.251
2009	0.265	0.269	0.272	0.252	0.255	0.266	0.277	0.271	0.261

Household-level nominal consumption Gini coefficients by province. Standard errors available upon request but excluded for readability.

Table 4: Household-level Gini coefficients, by truncation rules

	Empirical	Observed		5% level		1% level	
	Gini	SD	Gini	SD	Gini	SD	Gini
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
1997	0.251	0.510	0.282	0.510	0.281	0.476	0.264
	(0.002)	(0.006)	(0.003)	(0.008)	(0.004)	(0.006)	(0.003)
1998	0.257	0.524	0.289	0.525	0.289	0.491	0.272
	(0.002)	(0.005)	(0.002)	(0.006)	(0.003)	(0.005)	(0.003)
1999	0.264	0.528	0.291	0.537	0.296	0.510	0.281
	(0.002)	(0.007)	(0.004)	(0.012)	(0.006)	(0.007)	(0.004)
2000	0.262	0.527	0.290	0.534	0.294	0.494	0.273
	(0.002)	(0.007)	(0.004)	(0.006)	(0.003)	(0.007)	(0.004)
2001	0.259	0.519	0.287	0.517	0.285	0.494	0.273
	(0.002)	(0.004)	(0.002)	(0.005)	(0.002)	(0.004)	(0.002)
2002	0.266	0.542	0.299	0.545	0.300	0.510	0.282
	(0.003)	(0.006)	(0.003)	(0.008)	(0.004)	(0.008)	(0.004)
2003	0.272	0.557	0.306	0.556	0.306	0.513	0.283
	(0.002)	(0.006)	(0.003)	(0.009)	(0.005)	(0.009)	(0.005)
2004	0.273	0.548	0.302	0.544	0.300	0.521	0.288
	(0.003)	(0.006)	(0.003)	(0.009)	(0.005)	(0.008)	(0.004)
2005	0.274	0.570	0.313	0.554	0.305	0.524	0.289
	(0.003)	(0.019)	(0.010)	(0.006)	(0.003)	(0.008)	(0.004)
2006	0.275	0.554	0.305	0.562	0.309	0.531	0.293
	(0.002)	(0.006)	(0.003)	(0.010)	(0.005)	(0.007)	(0.003)
2007	0.264	0.538	0.296	0.553	0.304	0.517	0.285
	(0.003)	(0.006)	(0.003)	(0.005)	(0.003)	(0.006)	(0.003)
2008	0.269	0.544	0.299	0.551	0.303	0.506	0.279
	(0.003)	(0.006)	(0.003)	(0.010)	(0.005)	(0.009)	(0.005)
2009	0.264	0.542	0.298	0.540	0.297	0.518	0.286
	(0.003)	(0.007)	(0.004)	(0.008)	(0.004)	(0.011)	(0.006)

Column 1 presents Gini coefficients for EASI-deflated consumption. Columns 3, 5, and 7 present our model-based Gini coefficients; columns 2, 4, and 6 the respective standard deviations. The 'observed' model uses the standard deviation of the distribution as a whole; for the other two models we calculate truncation points by the following procedure. In each year, truncation points are selected by gridsearch to maximize N conditional on not rejecting a hypothesis of truncated normality by K-S test at the indicated level. We then take the median truncation point over years to facilitate comparability.

References

Attanasio, O., Hurst, E., and Pistaferri, L. (2012). The evolution of income, consumption, and leisure inequality in the us, 1980-2010. Technical report, National Bureau of Economic Research.

Battistin, E., Blundell, R., and Lewbel, A. (2009). Why is consumption more log normal than income? gibrat's law revisited. *Journal of Political Economy*, 117(6):pp. 1140–1154.

Brzozowski, M. and Crossley, T. F. (2011). Viewpoint: Measuring the well-being of the poor with income or consumption: a canadian perspective. *Canadian Journal of Economics*, 44(1).

Brzozowski, M., Gervais, M., Klein, P., and Suzuki, M. (2010). Consumption, income, and wealth inequality in canada. *Review of economic dynamics*, 13(1):52–75.

Dagum, C. (1999). Linking the functional and personal distributions of income. In *Handbook of Income Inequality Measurement*, pages 101–132. Springer.

Fisher, J. D., Johnson, D. S., Marchand, J. T., Smeeding, T. M., and Torrey, B. B. (2007). No place like home: Older adults and their housing. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 62(2):S120–S128.

Fortin, N., Green, D., Lemieux, T., and Milligan, K. (2012). Canadian inequality: Recent developments and policy options. *Canadian Public Policy*, 38(2):121–145.

Hall, R. (1979). Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence. *NBER Working Paper No. R0015.*

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pages 153–161.

Lewbel, A. and Pendakur, K. (2009). Tricks with hicks: The easi demand system. *The American Economic Review*, pages 827–863.

McAlister, D. (1879). The law of the geometric mean. *Proceedings of the Royal Society of London*, 29(196-199):367–376.

Meyer, B. D. and Sullivan, J. X. (2013). Consumption and income inequality and the great recession. *The American Economic Review*, 103(3):178–183.

Norris, S. and Pendakur, K. (2013). Imputing rent in consumption measures, with an application to poverty in canada, 1997-2009. *Canadian Journal of Economics*.

Pareto, V. (1895). La legge della domanda. *Giornale degli economisti*, 10:59–68.

Pendakur, K. (1998). Changes in canadian family income and family consumption inequality between 1978 and 1992. *Review of Income and Wealth*, 44(2):259–282.